# net>scaler.

## **NetScaler ingress controller**





## Contents

NetScaler cloud native solution for microservices based on Kubernetes	7
NetScaler Ingress Controller for Kubernetes	11
Supported platforms and deployments	11
Release notes	14
Deploy NetScaler Ingress Controller in Kubernetes cluster	15
Deploy CPX with NetScaler Ingress Controller in Kubernetes cluster	17
Deployment topologies	18
Unified ingress	25
Dual-tier ingress	30
Multi-cluster ingress	35
Deploy the NetScaler Ingress Controller on a Rancher managed Kubernetes cluster	48
Deploy the NetScaler Ingress Controller on a PKS managed Kubernetes cluster	50
Deploy NetScaler-Integrated Canary Deployment Solution	52
Deploy NetScaler IPAM controller	75
NetScaler API Gateway for Kubernetes	78
Deploying NetScaler API Gateway using Rancher	80
Deploy API Gateway with GitOps	82
GSLB overview and deployment topologies	91
Deploy NetScaler GSLB controller	99
NetScaler GSLB controller for single site	108
Service Mesh lite	121
NetScaler Operator	129
NetScaler Operator release notes	130

Upgrade NetScaler Operator	131
Deploy NetScaler Operator	134
Deploy NetScaler Ingress Controller in OpenShift using NetScaler Operator	138
Deploy NetScaler Ingress Controller (NSIC) with CPX in OpenShift using NetScaler Operato	or161
Deploy NetScaler Cloud Controller using NetScaler Operator	185
Deploy NetScaler GSLB Controller in OpenShift using NetScaler Operator	197
Deploy NetScaler Ingress Controller as a standalone pod by using NetScaler Operator	203
Deploy NetScaler Observability Exporter (NSOE) by using NetScaler Operator	209
Deploy the NetScaler Ingress Controller as an OpenShift router plug-in	216
Deploy the NetScaler Ingress Controller with OpenShift router sharding support	227
Deploy NetScaler CPX as an Ingress device in an Azure Kubernetes Service cluster	232
Deploy NetScaler Ingress Controller in an Azure Kubernetes Service cluster with NetScaler VPX	234
Deploy NetScaler CPX as an Ingress device in Google Cloud Platform	238
Deploy the NetScaler Ingress Controller in Anthos	240
Deploy NetScaler VPX in active-active high availability in EKS environment using Amazon ELB and NetScaler Ingress Controller	248
Deploy the NetScaler Ingress Controller for NetScaler with admin partitions	261
Deploy Citrix solution for service of type LoadBalancer in AWS	271
Multi-cloud and GSLB solution with Amazon EKS and Microsoft AKS clusters	275
Annotations	290
ConfigMap support for the NetScaler Ingress Controller	314
Ingress configurations	318
Ingress class support	323
Service class for services of type LoadBalancer	328

Configure HTTP, TCP, or SSL profiles on NetScaler	329
Log levels	346
SSL certificate for services of type LoadBalancer through the Kubernetes secret resource	347
BGP advertisement of external IP addresses for type LoadBalancer services and Ingresses using NetScaler CPX	351
NetScaler CPX integration with MetalLB in layer 2 mode for on-premises Kubernetes cluster	s363
Advanced content routing for Kubernetes Ingress using the HTTPRoute CRD	368
Profile support for the Listener CRD	371
IP address management using the IPAM controller	378
Apply CRDs through annotations	384
Listener CRD support for Ingress through annotation	386
Configuring consistent hashing algorithm using NetScaler Ingress Controller	391
Add DNS records using NetScaler Ingress Controller	392
Open policy agent support for Kubernetes with NetScaler	394
Exporting metrics directly to Prometheus	402
Configure static route on NetScaler VPX or MPX	404
Establish network between Kubernetes nodes and Ingress NetScaler using node controller	406
Expose Service of type NodePort using Ingress	408
Configure pod to pod communication using Calico	410
Enhancements for Kubernetes service of type LoadBalancer support in the NetScaler Ingress Controller	416
TLS certificates handling in NetScaler Ingress Controller	422
TLS client authentication support in NetScaler	430
TLS server authentication support in NetScaler using the NetScaler Ingress Controller	431
Install, link, and update certificates on a NetScaler using the NetScaler Ingress Controller	433

Configure SSL passthrough using Kubernetes Ingress	437
Automated certificate management with cert-manager	439
Deploy HTTPS web application on Kubernetes with the NetScaler Ingress Controller and Let's Encrypt using cert-manager	440
Deploy an HTTPS web application on Kubernetes with NetScaler Ingress Controller and HashiCorp Vault using cert-manager	459
Enable NetScaler certificate validation in the NetScaler Ingress Controller	476
Disable API server certificate verification	479
Create a self-signed certificate and linking into Kubernetes secret	479
Authentication and authorization policies for Kubernetes with NetScaler	481
Rate limiting in Kubernetes using NetScaler	501
Use Rewrite and Responder policies in Kubernetes	506
Remote content inspection or content transformation service using ICAP	526
VIP custom resource definition (CRD)	543
HTTPRoute	544
Advanced content routing for Kubernetes with NetScaler	565
Listener	571
Configure web application firewall policies with the NetScaler Ingress Controller	595
Configure bot management policies with the NetScaler Ingress Controller	610
Configure cross-origin resource sharing policies with NetScaler Ingress Controller	618
Enable request retry feature using AppQoE for NetScaler Ingress Controller	621
Configuring wildcard DNS domains through NetScaler Ingress Controller	623
View metrics of NetScalers using Prometheus and Grafana	625
Analytics and observability	634
Analytics configuration support using ConfigMap	637

IP address management using the IPAM controller	642
Securing Ingress	648
TCP use cases	653
HTTP use cases	670
HTTP callout with the rewrite and responder policy	680
Configure session affinity or persistence on the Ingress NetScaler	686
Allowlisting or blocklisting IP addresses	689
Interoperability with ExternalDNS	693
Using NetScaler credentials stored in a Vault server for the NetScaler Ingress Controller	694
How to use Kubernetes secrets for storing NetScaler credentials	701
How to load balance ingress traffic to UDP-based application	705
How to set up dual-tier deployment	708
Horizontal pod autoscaler for NetScaler CPX with custom metrics	715
Deploy Direct Server Return	719
Support for admission controller webhooks	723
Enable gRPC support using the NetScaler Ingress Controller	728
Policy-based routing for multiple Kubernetes clusters	734
Single tier NetScaler Ingress solution for MongoDB	741
Canary and blue-green deployment using NetScaler VPX and Azure pipelines for Kuber- netes based applications	750
Traffic management for external services	756
Support for external name service across namespaces	758
Troubleshooting NetScaler Ingress Controller	760
Upgrade NetScaler Ingress Controller	773

Entity name change	776
Licensing	779

## NetScaler cloud native solution for microservices based on Kubernetes

## August 6, 2024

As companies transform to innovate faster and get closer to customers, they are rearchitecting their internal process and breaking down boundaries within their organization. They are removing silos to pull together the right skill sets in the same team. One of the goals is to create and deliver software applications with speed, agility, and efficiency. In this regard, modern application architectures based on microservices are being adopted by a growing number of enterprises.

Using a microservices architecture, you can create applications as sets of loosely coupled services which can be deployed, updated, and scaled independently.

Cloud native is an approach that relies on the microservices architecture for building and deploying applications with the following key attributes:

- Deploys applications as loosely coupled microservices or containers
- Involves a very high degree of automation
- Implements agile DevOps processes and continuous delivery workflows
- Centers around APIs for interaction and collaboration

## How does Kubernetes help in the cloud native journey?

To provide the desired levels of agility and stability, cloud native applications require high levels of infrastructure automation, security, networking, and monitoring. You need a container orchestration system that can efficiently manage containers at a large scale. Kubernetes has emerged as the most popular platform for container deployment and orchestration. Kubernetes abstracts the complex task of running, deploying, and managing containers from developers and operators and automatically schedules containers among a cluster of nodes. Kubernetes and the cloud native computing foundation (CNCF) ecosystem helps you to build a platform for cloud native solutions.

Some of the key benefits of using Kubernetes:

- Simplifies application deployment be it on-premises, hybrid, or public cloud infrastructure
- Accelerates application development and deployment
- Increases agility, flexibility, and scalability of applications

## What is NetScaler cloud native solution?

To maximize the benefits of using Kubernetes in production, you need to integrate Kubernetes with several tools, vendor-sourced, and open-source components. Ensuring production grade reliability and security for their cloud native application is a challenge faced by many organizations. NetScaler

offers a NetScaler cloud native solution to address the challenges in a Kubernetes production environment.

NetScaler cloud native solution leverages the advanced traffic management, observability, and comprehensive security features of NetScaler to ensure enterprise grade reliability and security. It can provide complete visibility to application traffic in your Kubernetes environment, render immediate feedback, and help gain meaningful insights about the application performance.

The following table lists the key requirements of different stakeholders while implementing an Ingress solution.

Stakeholders	Job function	Needs
Platform administrators	Ensure availability of Kubernetes clusters	Simpler ways to manage applications deployed across multiple clusters, operation, and platform life cycle management
DevOps	Accelerate the deployment of applications to production	Integration with CI/CD pipeline, support for deployment techniques like Canary and blue-green for faster deployment
Developers	Develop and test microservices	Ways to bring traffic into the Kubernetes cluster, tracing and debugging, rate limiting for applications, and authentication for applications
SREs	Ensure availability of applications to meet service level agreements	Advanced telemetry for applications and infrastructure
SecOPs	Ensure security compliance	Secure Ingress traffic, API protection, service mesh for secure communication between microservices inside the Kubernetes cluster

The following diagram explains the NetScaler cloud native solution and how it addresses the various challenges faced by stakeholders in their cloud native journey.

	Å ↓		
Smarter CICD with L3-L7 metrics	Enterprise Ingress solution <b>net&gt;scaler</b> . Ingress	AAA, WAF, API GW	NetScaler ADM for better
+	NetScaler functions as Kubernetes APIs Applications, Databases, Message Qs	+	observability with L3-L7 app metrics
Spinnaker	Lightweight, High Performance, Enterprise Grade Service Mesh	spiffe	Prometheus Napility (Mon
CI/CD	Kubernetes Platform, CNI, CSI	Security	fluentd
Stakeholders:	DevOps Developers Platform Admins DevSecOps SRE		

NetScaler cloud native solution provides the following key benefits:

- Provides an advanced Kubernetes Ingress solution that caters to the needs of developers, SREs, devOps, and network or cluster administrators.
- Eliminates the need to rewrite legacy applications based on TCP or UDP traffic while moving them into a Kubernetes environment.
- Secures applications with NetScaler policies exposed as Kubernetes APIs.
- Helps to deploy high performing microservices for North-South traffic and East-West traffic.
- Provides an all-in-one view of all microservices using NetScaler ADM service graph.
- Enables faster troubleshooting of microservices across different kinds of traffic including TCP, UDP, HTTP, HTTPS, and SSL.
- Secures APIs.
- Automates CI/CD pipeline for Canary deployments.
- Provides out of the box integrations with CNCF open-source tools.

For more information on the various cloud native solutions offered by NetScaler, see the following links:

- Kubernetes Ingress solution
- Service mesh
- Solutions for observability
- API gateway for Kubernetes

## **Components of NetScaler cloud native solution**

The following table explains the major components of NetScaler cloud native solution:

Component	Description
NetScaler Ingress Controller	This container is an implementation of the Kubernetes Ingress Controller to manage and route traffic into your Kubernetes cluster using NetScaler ADCs (NetScaler CPX, BLX, VPX, or MPX). Using NetScaler Ingress Controller, you can configure NetScaler CPX, BLX, VPX, or MPX according to the Ingress rules and integrate your NetScaler ADCs with the Kubernetes
NetScaler Observability Exporter	NetScaler Observability Exporter is a container which collects metrics and transactions from NetScaler ADCs and transforms them to suitable formats (such as JSON, AVRO) for supported endpoints. You can export the data collected by NetScaler Observability Exporter to the desired endpoint. By analyzing the data exported to the endpoint, you can get valuable insights at a microservices level for applications proxied by
NetScaler xDS adapter	NetScaler ADCs. NetScaler xDS adapter is a container for integrating NetScaler with service mesh control plane implementations based on xDS APIs (Istio, Consul, and so on). It communicates with the service mesh control plane and listens for updates by acting as a gRPC client to the control plane API server. Based on the updates from the control plane, the NetScaler xDS-Adaptor generates the equivalent NetScaler configuration.

Component	Description
NetScaler CPX	NetScaler CPX is a container-based application
	delivery controller that can be provisioned on a
	Docker host. NetScaler CPX enables customers
	to leverage Docker engine capabilities and use
	NetScaler load balancing and traffic
	management features for container-based
	applications. You can deploy one or more
	NetScaler CPX instances as standalone instances
	on a Docker host.

## **NetScaler Ingress Controller for Kubernetes**

#### August 6, 2024

When you are running an application inside a Kubernetes cluster, you need to provide a way for external users to access the applications from outside the Kubernetes cluster. Kubernetes provides an object called Ingress which allows you to define the rules for accessing the services with in the Kubernetes cluster. It provides the most effective way to externally access multiple services running inside the cluster using a stable IP address.

An ingress controller is an application deployed inside the cluster that interprets rules defined in the Ingress. The Ingress controller converts the Ingress rules into configuration instructions for a load balancing application integrated with the cluster. The load balancer can be a software application running inside your Kubernetes cluster or a hardware appliance running outside the cluster.

NetScaler provides an implementation of the Kubernetes Ingress Controller to manage and route traffic into your Kubernetes cluster using NetScaler ADCs (NetScaler CPX, BLX, VPX, or MPX).

Using NetScaler Ingress Controller, you can configure NetScaler CPX, BLX, VPX, or MPX according to the ingress rules and integrate your NetScaler ADCs with the Kubernetes environment.

## Supported platforms and deployments

July 8, 2025

This topic provides details about supported Kubernetes platforms and compatibility of cloud native components with NetScaler.

## **Kubernetes platforms**

NetScaler Ingress Controller(NSIC) is supported on the following platforms:

- Kubernetes v1.21 (and later) on bare metal or self-hosted on public clouds such as AWS, GCP, or Azure.
- Google Kubernetes Engine (GKE)
- Elastic Kubernetes Service (EKS)
- Azure Kubernetes Service (AKS)
- Red Hat OpenShift version 3.11 and later
- Pivotal Container Service (PKS)
- Diamanti Enterprise Kubernetes Platform
- Mirantis Kubernetes Engine
- VMware Tanzu
- Rancher

## Compatibility with NetScaler

The following table lists the NetScaler versions that are compatible with any NetScaler Ingress Controller release:

NetScaler Ingress Controller version	NetScaler versions
3.1.34	VPX/MPX/CPX 14.1-43.50
3.0.5	VPX/MPX/CPX 14.1-38.53
2.3.15	VPX/MPX/CPX 14.1-38.53
2.2.10	VPX/MPX 14.1-25.56 CPX 14.1-25.111
2.1.4	VPX/MPX 14.1-25.56 CPX 14.1-25.111
2.0.6	VPX/MPX 14.1-25.56 CPX 14.1-25.111
1.43.7	VPX/MPX 14.1-25.56 and later CPX 14.1-25.109 and later

NetScaler Ingress Controller version	NetScaler versions
1.42.12	VPX/MPX 14.1-17.38 and later
	CPX 14.17.101 and later
1.41.5	VPX/MPX 14.1-17.38 and later
	CPX 14.17.101 and later
1.40.12	14.1-12.35 and later
1.39.6	14.1-12.35 and later
1.38.27	13.1-49.15 and later
1.37.5	13.1-49.15 and later
1.36.5	13.1-49.15 and later
1.35.6	13.1-49.15 and later

## NetScaler Operator compatibility with NetScaler cloud-native components

The table lists NetScaler cloud-native component releases that can be deployed using any NetScaler Operator release.

	NetScaler		GSLB		Cloud	NetScaler
OpenShift	Operator	NSIC	Controller	NSOE	Controller	IPAM
version	version	version	version	version	version	Controller
4.11 and	3.2.0	2.3.15	2.3.15	1.10.1	1.2.0	2.1.2
later						
4.11 and	3.1.2	2.2.10	2.2.10	1.10.1	1.2.0	_
later						

## Compatibility of cloud native components with NetScaler

The following cloud native components are compatible with NetScaler version 13.1 and later:

- NetScaler Node Controller
- NetScaler Observability Exporter
- NetScaler Console service
- NetScaler Console on-prem
- NetScaler xDS-adapter

## **Release notes**

July 8, 2025

## NetScaler Ingress Controller (NSIC)/NetScaler GSLB Controller

a. NSIC/GSLB releases

- 3.1.34
- 3.0.5
- 2.3.15
- 2.2.10
- 2.1.4
- 2.0.6
- 1.43.7

b. Helm chart releases

- 3.1.34
- 3.0.6
- 3.0.5
- 2.3.16
- 2.3.15
- 2.2.11
- 2.2.10
- 2.1.4
- 2.0.6

c. Operator releases

- 3.2.0
- 3.1.2

## **IPAM Controller**

- 2.1.2
- 2.0.2
- 2.0.1
- 1.2.0

## References

- For information on upgrading NSIC using Helm chart, see Upgrade NetScaler Ingress Controller.
- For information on upgrading NetScaler Operator, see Upgrade NetScaler Operator.

## **Deploy NetScaler Ingress Controller in Kubernetes cluster**

## October 16, 2024

This section helps network administrators who have access to NetScaler (MPX/SDX/VPX/BLX) and Kubernetes environment to deploy NetScaler Ingress Controller in Kubernetes.

## Prerequisites

- You have installed and set up a Kubernetes cluster.
- You have installed Helm version 3.x or later. To install Helm, see here.
- You have a system user account on NetScaler. NetScaler Ingress Controller uses this account to push configuration to NetScaler. See Create system user account for NetScaler Ingress Controller in NetScaler.
- You have a Kubernetes secret to store NetScaler user account credentials. See How to use Kubernetes secrets for storing NetScaler credentials.

## How to deploy NetScaler Ingress Controller in Kubernetes cluster

1. Add the NetScaler Helm chart repository to your local registry by using the following command.

```
1 helm repo add netscaler https://netscaler.github.io/netscaler-helm
        -charts/
```

If the NetScaler Helm chart repository is already added to your local registry, use the following command to update the repository:

1 helm repo update netscaler

2. Install NetScaler Ingress Controller by using the following command.

#### Note:

For information about all the parameters that can be configured during installation, see **Configuration**.

```
1 helm install netscaler-ingress-controller netscaler/netscaler-
ingress-controller --set license.accept=yes,adcCredentialSecret
=<Secret-for-NetScaler-credentials>,nsIP=<NetScalerIP>
```

#### 3. Verify the installation.

Note:

```
If any custom Helm release name is provided, then the the label is app=<releasename >-netscaler-ingress-controller.
```

```
1 kubectl get pods -l app=netscaler-ingress-controller
```

root@mc1-master1:~# kubectl get pods -l app=netscaler-ingress-controller NAME READY STATUS RESTARTS AGE netscaler-ingress-controller-56588cd548-5k7s7 1/1 Running 0 52s

4. (Optional) View the summary of NetScaler Ingress Controller pod.

```
1 kubectl describe pod <pod name>
```

Example:

```
1 kubectl describe pod netscaler-ingress-controller-8474d866d9b-
nqp6p
```

5. (Optional) Check the NetScaler Ingress Controller logs.

1 kubectl logs -f <pod name>

Example:

1 kubectl logs -f netscaler-ingress-controller-8474d866d9b-nqp6p

6. (Optional) Delete the installation.

```
1 helm uninstall netscaler-ingress-controller
```

## What's next

See the following topics for comprehensive information on deploying NetScaler Ingress Controller and to customize your installation accordingly.

- Deployment topologies
- Unified ingress

## Deploy CPX with NetScaler Ingress Controller in Kubernetes cluster

## August 5, 2024

This section helps platform administrators who have access to Kubernetes environment to deploy CPX with NetScaler Ingress Controller in Kubernetes.

## Prerequisites

- You have installed and set up a Kubernetes cluster.
- You have installed Helm version 3.x or later. To install Helm, see here.

## How to deploy NetScaler Ingress Controller in Kubernetes cluster

1. Add the NetScaler Helm chart repository to your local registry by using the following command.

1 helm repo add netscaler https://netscaler.github.io/netscaler-helm -charts/

If the NetScaler Helm chart repository is already added to your local registry, use the following command to update the repository:

```
1 helm repo update netscaler
```

2. Install NetScaler CPX with NSIC using the following command.

```
1 helm install netscaler-cpx-with-ingress-controller netscaler/
netscaler-cpx-with-ingress-controller --set license.accept=yes,
serviceType.nodePort.enabled=True
```

3. Verify the installation.

1 kubectl get pods -l app=netscaler-cpx-wi	ith-ing	ress-con	troller	
root@mc1-master1:~# kubectl get pods -l app=netscaler-cp	px-with-:	ingress-com	ntroller	
NAME netscaler-cpx-with-ingress-controller-78ffd4549d-mvxlm	READY	STATUS Running	RESTARTS 0	AGE 33m

There are 2 containers running in the same pod highlighted by 2/2 under READY column. One container is for NetScaler CPX proxy and another container is for NetScaler Ingress Controller.

4. View the details of both the containers.

```
1 kubectl describe pod $(kubectl get pods -l app=netscaler-cpx-with-
ingress-controller | awk '{
2 print $1 }
3 ' | grep netscaler-cpx-with-ingress-controller)
```

#### NetScaler ingress controller

Controlled By: ReplicaSet/msic-netscaler-ingress-controller-98657f7dc Containers: nsic: Container ID: containerd://259706cce8e3ca5f4c6c0c64ddef458021c607020e51e6171183f220c0049858 Image: quay.lo/netscaler/netscaler-k8a-ingress-controller:1.37.5 Image ID: quay.lo/netscaler/netscaler-k8a-ingress-controller8sha256:7d6a83ac72676e45e4d9812cfa0e2a4222c9ff82252b06d1b Fort: <forme> Nrgs: -contigmapdc=stalts:restcaler.ex8a-ingress-controller8sha256:7d6a83ac72676e45e4d9812cfa0e2a4222c9ff82252b06d1b Fort: <forme> Nrgs: -contigmapdc=stalts:restcaler.ex8a-ingress-controller8sha256:7d6a83ac72676e45e4d9812cfa0e2a4222c9ff82252b06d1b Fort: <forme> Nrgs: -contigmapdc=stalts:restcaler.ex8a-ingress-controller8sha256:7d6a83ac72676e45e4d9812cfa0e2a4222c9ff82252b06d1b Fort: <forme> Nrgs: -contigmapdc=stalts:restcaler.ex8a-ingress-controller8sha256:7d6a83ac72676e45e4d9812cfa0e2a4222c9ff82252b06d1b Fort: <forme> Nrgs: -contigmapdc=stalts:restcaler.ex8a-ingress-controller8sha256:7d6a83ac72676e45e4d9812cfa0e2a4222c9ff82252b06d1b Fort: <forme> Neguests: -cpu: 32m memory: 128Mi Environment: NS IP: 10.146.109.4 NS USER: <forme> NS FASSWORD: <forme> NS FASSWORD: <forme> NS AFPS_NMME_PREFIX: K8s Nounds: -fvar/run/secrets/kubernetes.io/serviceaccount from kube-api-access-bq255 (ro) Conditions: Type Stalts Type Stalts Type Stalts The ContinersReady True Projected (a volume that contains injected data from multiple sources) TokenExpirationSeconds: 3607 ConfidMapName: Con</forme></forme></forme></forme></forme></forme></forme></forme></forme>	1		
Container ID: containerd://259106ccce@3ca514c6c0c64ddef458021c607020e51e6171183f220c0049858 Image: quay.io/netscaler/netscaler/netscaler-k8s-ingress-controller@sha256:7d6a83ac72676e45e4d9812cfa0e2a4222c9ff82252b06d1b Port: <cone> Host Port: <cone> Host Por</cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone></cone>	Controlled By: Re	plicaSet/n	sic-netscaler-ingress-controller-98657f7dc
<pre>nsic: Container ID: containerd://259f06cce8e3ca5f4c6c0c64ddef458021c607020e51e6171183f220c0049858 Image: quay.io/netscaler/nstscaler=k8s=ingress=controller:1.37.5 Image ID: quay.io/netscaler/nstscaler=k8s=ingress=controller*8ha256;7d6a83ac72676e45e4d9812cfa0e2a4222c9ff82252b06d1b Port: <none> Mrgs: contemps default/nsic=nsic=configmap contemps default/nsic=nsic=configmap cathure=nod=watch true update=ingress=status yes State: Running Stated: Mon, 15 Jan 2024 00:11:06 +0000 Ready: True Restat: Count: 0 Reguests: cpu: 32m memory: 12Mi Environment: 10.146.109.4 NS JEF: 10.146.109.146.1</none></pre>	Containers:		
Container ID: containeri//259106ccceebac5f46c60c64ddef455021660720e051ed71183f220c0049858 Image : quay.io/netscaler/netscaler+&s-ingress-controller@sha256:7d6a83ac72676e45e4d9812cfa0e2a4222c9ff82252b06d1b Port: < none> Host Port: < none> Args: configmap default/nsic-nsic-configmap feature-node-watch true enable-onc-pbr false unable-onc-pbr false -	nsic:		
<pre>Image: quay.io/netscaler/netscaler=k8s=ingress=controller:1.37.5 Image 10: quay.io/netscaler/netscaler=k8s=ingress=controller@sha256:7d6a83ac72676e45e4d9812cfa0e2a4222c9ff82252b06d1b Port: <none> Args: configmap default/nsic=nsic=configmap feature=node=watch true enable=cnc=pbr false update=ingress=status yes State: Running Started: Mon, 15 Jan 2024 00:11:06 40000 Ready: True Reguests: cpu: 32m memory: 128Mi Environment: NS TF: 10.146.109.4 NS USER: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false NS PF: 10.146.109.4 NS USER: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false NS PF: vses NS APPS_NAME_PREFIX: k8s Mounts: /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access=bq255 (ro) Conditions: Type Status Initialized True Ready True ContainersReady True PodScheduled True Volumes: kube=api-access=bq255: Type: Projected (a volume that contains injected data from multiple sources) TokenExpirationSeconds: 3607 ConfidmapAmee: kube=root=ca.crt ConfidmapAmee: kube=root=ca.crt ConfidmapAmee: kube=root=ca.crt ConfidmapAmee: kube=root=ca.crt ConfidmapAmee: cone&gt; Node=kubernetes.io/not=ready:NoExecute op=Exists for 300s_ </set></set></none></pre>	Container ID:		d://259f06cce8e3ca5f4c6c0c64ddef458021c607020e51e6171183f220c0049858
<pre>Image 10: quay.io/netscaler/netscaler-k8s-ingress-controller@sha256:7d6a83ac72676e45e4d9812cfa0e2a4222c9ff82252b06d1b Fort: <code concept="" for="" of="" set="" set<="" th="" the=""><th>Image:</th><th>quay.io/n</th><th>etscaler/netscaler-k8s-ingress-controller:1.37.5</th></code></pre>	Image:	quay.io/n	etscaler/netscaler-k8s-ingress-controller:1.37.5
Port: <none>         Host Port:       <none>         Args:      </none></none>	Image ID:		etscaler/netscaler-k8s-ingress-controller@sha256:7d6a83ac72676e45e4d9812cfa0e2a4222c9ff82252b06d1b
Host Port: <none>         Args:      configmap default/nsic-nsic-configmap        centure=node=watch true        enable=conc=pbr false        update=ingress-status yes         State:       Running         Stated:       Non, 15 Jan 2024 00:11:06 +0000         Ready:       True         Restart Count:       0         Requests:      </none>	Port:		
<pre>Args: configmap default/nsic-nsic-configmap feature-node-watch true enable-cnc-pbr false update-ingress-status yes State: Running Started: Mon, 15 Jan 2024 00:11:06 +0000 Readerst: cpu: 32m nemory: 128M Environment: NS_IP: 10.146.109.4 NS_USER: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false NS_IP: 10.146.109.4 NS_USER: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSWORD: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSWORD: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSWORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSWORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSWORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false NS_PAPS_NAME_PREFIX: k8s Mounts: /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-bqz55 (ro) Conditions: Type Status Initialized True PodScheduled True PodScheduled True PodScheduled True Volumes: kube-api-access-bqz55: Type: Projected (a volume that contains injected data from multiple sources) TokenExpirationSeconds: 3607 ContigMapOptional: <nl>         Frojected (a volume that contains injected data from multiple sources) TokenExpirationSeconds: 3607 ContigMapOptional: <nl>         Kube-root-ca.crt ConfigMapOptional: <nl>         Kube-root-ca.crt ConfigMapOptional: <nl>         Kube-root-ca.crt ConfigMapOptional: <nl>         Kouternetes.io/not-ready:NoExecute op=Exists for 300s Tolerations: </nl></nl></nl></nl></nl></set></set></set></set></set></set></set></pre>	Host Port:		
<pre>configmap default/nsic=nsic=configmap feature=node=vatch true enable=cnc=pbr false update=ingress=status yes State: Running Started: Mon, 15 Jan 2024 00:11:06 +0000 Ready: True Restart Count: 0 Requests: cpu: 3.2m memory: 128Mi Bnvironment: NS_TP: 10.146.109.4 NS_USER: &lt; set to the key 'username' in secret 'nslogin'&gt; Optional: false NS_DSER: &lt; set to the key 'username' in secret 'nslogin'&gt; Optional: false SULA: yes NS_DFPS_NAME_PREFIX: k8s Mounts: -/var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-bq255 (ro) Conditions: Type Status Initialized True Ready True ContainersReady True PodScheduled True Welumes: kube-api-access=bq255: Type: Projected (a volume that contains injected data from multiple sources) Type: Projected (a volume that contains injected data from multiple sources) Type: Nube-rot-ca.crt ConfidMapOptional: &lt; inib- DownwardAPI: true CoS Class: Supration Node-Subernetes.io/not-ready:NoExecute op=Exists for 300s Node-Subertions: Tolerations: Tolerations: Node-Subernetes.io/not-ready:NoExecute op=Exists for 300s Node-Subernetes.io/not-ready:NoExecute op=Exists for 300s Node-Subernetes.io/not-ready:NoExecute op=Exists for 300s Node-Subernetes.io/not-ready:NoExecute op=Exists for 300s Node-Subernetes.io/not-ready:NoExecute op=Exists for 300s Node-Nubernetes.io/not-ready:NoExecute op=Exists f</pre>	Args:		
<pre>feature-node-watch true enable-non-pbr false update-ingress-status yes State: Running Started: Mon, 15 Jan 2024 00:11:06 +0000 Ready: True Restart Count: 0 Reguests: cpu: 32m memory: 120Mi Environment: NS_DSER: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSWORD: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSWORD: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSWORD: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSWORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSWORD: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false SULA: yes NS_APES_NAME_PREFIX: k8s Mounts: /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-bq255 (ro) Conditions: Type Status Initialized True Ready True ContainersReady True PodScheduled True Volumes: kube-api-access-bq255: Type: Projected (a volume that contains injected data from multiple sources) TokenExpirationSeconds: 3607 ConfigMapOptional: <ni> DownwardAPI: true OoS Class: Burstable Node-Selectors: <none> toke.kubernetes.io/not-ready:NoExecute op=Exists for 300s_ Colerations:</none></ni></set></set></set></set></set></set></pre>	configmap	default/ns	ic-nsic-configmap
<pre>enable-onc-pbr false update-ingress-status yes State: Running Statted: Mon, 15 Jan 2024 00:11:06 +0000 Ready: True Restat Count: 0 Requests: 32m memory: 120Mi Environment: 10.146.109.4 NS_USER: 4set to the key 'username' in secret 'nslogin'&gt; Optional: false NS_DRSWORD: 4set to the key 'username' in secret 'nslogin'&gt; Optional: false NS_DRSWORD: 4set to the key 'username' in secret 'nslogin'&gt; Optional: false NS_DRSWORD: 4set to the key 'username' in secret 'nslogin'&gt; Optional: false NS_DRSWORD: 4set to the key 'username' in secret 'nslogin'&gt; Optional: false NS_DRSWORD: 4set to the key 'username' in secret 'nslogin'&gt; Optional: false NS_DRSWORD: 4set to the key 'username' in secret 'nslogin'&gt; Optional: false NS_DRSWORD: 4set to the key 'username' in secret 'nslogin'&gt; Optional: false NS_DRSWORD: 4set to the key 'username' in secret 'nslogin'&gt; Optional: false NS_DRSWORD: 4set to the key 'username' in secret 'nslogin'&gt; Optional: false Nounts: /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-bqz55 (ro) Conditions: Type Status Initialized True PodScheduled True Volumes: kube-api-access-bqz55: Type: Projected (a volume that contains injected data from multiple sources) Type: Projected (a volume that contains injected data from multiple sources) Type: Projected (a volume that contains injected data from multiple sources) Type: Projected (a volume that contains injected data from multiple sources) TokenExpirationSecond: 3607 ConfigMapName: kube-root-ca.crt ConfigMapOptional: cnil&gt; DommardAPI: true QoS Class: Surstable Node-selectors: cnome&gt; Tolerations: node.kubernetes.io/not-ready:NoExecute op=Exists for 300s_ Nodes.kubernetes.io/not-ready:NoExecute op=Exists for 300s_ Nodes.kubernetes.io/not-ready:No</pre>	feature-no	de-watch t	
<pre>update-ingress-status yes State: Running Started: Mon, 15 Jan 2024 00:11:06 +0000 Ready: True Restart Count: 0 Requests: cpu: 32m memory: 128Mi Environment: NS_IP: 10.146.109.4 NS_USER: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSMORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSMORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSMORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSMORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSMORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSMORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSMORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false NS_PAPS_NAME_PREFIX: k8s Mounts: /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-bqz55 (ro) Conditions: Type Status Initialized True Ready True ContainersReady True PodScheduled True Volumes: kube-api-access-bqz55: Type: Projected (a volume that contains injected data from multiple sources) Type: Projected (a volume that contains injected data from multiple sources) Type: Projected (a volume that contains injected data from multiple sources) Type: Projected (a volume that contains injected data from multiple sources) Type: Status NS_PASSMORD: Signature ConfigMapMame: kube-root-ca.crt ConfigMapMame: Kube-root-ca.crt</set></set></set></set></set></set></set></set></pre>	enable-cnc	-pbr false	
State:       Running         Stated:       Mon, 15       Jan 2024 00:11:06 +0000         Ready:       True         Restart Count:       0         Requests:	update-ing		
<pre>Started: Mon, 15 Jan 2024 00:11:06 +0000 Ready: True Restart Count: 0 Reguests:     cpu: 32m     memory: 128Mi Environment:     NS IF: 10.146.109.4     NS USER: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false     NS PASSWORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false     NS PASSWORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false     NS PASSWORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false     NS PASSWORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false     NS PASSWORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false     NS PASSWORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false     NS APPS NAME PREFIX: k8s     Mounts:</set></set></set></set></set></set></set></pre>	State:	Running	
<pre>Ready: True Rest Count: 0 Requests:     cpu: 32m     memory: 122Mi Environment:     NS IP:</pre>	Started:	Mon, 15	Jan 2024 00:11:06 +0000
<pre>Restart Count: 0 Requests:     cpu: 32m     memory: 128Mi Environment:     ID.146.109.4     NS IP:</pre>	Ready:	True	
Requests:         cpu: 32m         memory: 128Mi         Environment:         NS_TP:       10.146.109.4         NS_USER: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false         EULA:       yes         NS_APPS_NAME_PREFIX:       k8s         Mounts:       /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-bqz55 (ro)         Conditions:       Type         Status       True         Initialized       True         PodScheduled       True         Volumes:       kube-api-access-bqz55         rype:       Projected (a volume that contains injected data from multiple sources)         TokenExpirationSeconds:       3607         ConfigMapName:       kube-root-ca.crt         ConfigMapOptional:       <ni><ni>         DownwardAPI:       true         Oode-Selectors:       <non>         Node-Selectors:       <non>         Toterations:       node.kubernetes.io/not-ready:NoExecute op=Exists for 300s</non></non></ni></ni></set>	Restart Count:		
<pre>cpu: 32m memory: 128Mi Environment: NS_IP: 10.146.109.4 NS_USER: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSWORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false NS_PASSWORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false ULA: yes NS_APPS_NAME_PREFIX: k8s Mounts: /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-bqz55 (ro) Conditions: Type Status Initialized True Ready True ContainersReady True PodScheduled True Volumes: kube-api-access-bqz55: Type: Projected (a volume that contains injected data from multiple sources) TokenExpirationSeconds: 3607 ConfigMapName: kube-root-ca.crt ConfigMapOptional: <nil> DownwardAPF: true QoS Class: Burstable Node-Selectors: <none> Toterations: node.kubernetes.io/not-ready:NoExecute op=Exists for 300s</none></nil></set></set></set></pre>	Requests:		
<pre>memory: 128Mi Environment: NS IP: 10.146.109.4 NS_USER: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false NS_USER: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false EULA: yes NS_APPS_NAME_PREFIX: k8s Mounts: /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-bqz55 (ro) Conditions: Type Status Initialized True Ready True PodScheduled True Volumes: kube-api-access-bqz55: Type: Projected (a volume that contains injected data from multiple sources) TokenExpirationSeconds: 3607 ConfigMapName: kube-root-ca.crt ConfigMapOptional: <ni></ni></set></set></pre>	cpu: 32m		
Environment:         NS IP:       10.146.109.4         NS_USER: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false         NS_DAPS_MAME_PREFIX:       yes         NS_APPS_NAME_PREFIX:       k8s         Mounts:       /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-bqz55 (ro)         Conditions:      </set>	memory: 128	Mi	
NS IP:       10.146.109.4         NS USER: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false         NS_PASSWORD:       <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false         EULA:       yes         NS_APPS_NAME_PREFIX:       K8s         Mounts:       /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-bqz55 (ro)         Conditions:       Type         Type       Status         Initialized       True         Ready       True         PodScheduled       True         Volumes:       kube-api-access-bqz55:         Type:       Projected (a volume that contains injected data from multiple sources)         TokenExpirationSeconds:       3607         ConfigMapName:       kube-root-ca.crt         ConfigMapOptional:       <ni>         DownwardAPI:       true         QoS Class:       Burstable         Node-Selectors:       <non>         Tolerations:       node.kubernetes.io/not-ready:NoExecute op=Exists for 300s</non></ni></set></set>	Environment:		
NS_USER: <set 'nslogin'="" 'username'="" in="" key="" secret="" the="" to=""> Optional: false         NS_PASSWORD:       <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false         EULA:       yes         NS_APPS_NAME_PREFIX:       k8s         Mounts:       /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-bqz55 (ro)         Conditions:       Type         Type       Status         Initialized       True         PodScheduled       True         Volumes:       kube-api-access-bqz55         Type:       Projected (a volume that contains injected data from multiple sources)         TokenExpirationSeconds:       3607         ConfigMapName:       kube-root-ca.crt         ConfigMapOptional:       <ni>         DownwardAPI:       true         Oode-Selectors:       <non>         Node-Selectors:       <non>         Tolerations:       node.kubernetes.io/not-ready:NoExecute op=Exists for 300s</non></non></ni></set></set>	NS_IP:		10.146.109.4
NS PASSWORD: <set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false         EULA:       yes         NS_APPS_NAME_PREFIX:       k8s         Mounts:       /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-bqz55 (ro)         Conditions:       Type         Type       Status         Initialized       True         Ready       True         PodScheduled       True         Volumes:       kube-api-access-bqz55         Kube-api-access-bqz55:       Projected (a volume that contains injected data from multiple sources)         TokenExpirationSeconds:       3607         ConfigMapName:       kube-root-ca.crt         ConfigMapOptional:       <nl>         DownwardAPI:       true         QoS Class:       Burstable         Node-Selectors:       <none>         Toterations:       node.kubernetes.io/not-ready:NoExecute op=Exists for 300s</none></nl></set>	NS_USER:		
EULA: yes NS APPS_NAME_PREFIX: k8s Mounts: /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-bqz55 (ro) Conditions: Type Status Initialized True Ready True PodScheduled True Volumes: kube-api-access-bqz55: Type: Projected (a volume that contains injected data from multiple sources) TokenExpirationSeconds: 3607 ConfigMapName: kube-root-ca.crt ConfigMapOptional: <ni>&gt; DownwardAPI: true QoS Class: Burstable Node-Selectors: <non>&gt; Toterations: node.kubernetes.io/not-ready:NoExecute op=Exists for 300s</non></ni>	NS_PASSWORD:		<set 'nslogin'="" 'password'="" in="" key="" secret="" the="" to=""> Optional: false</set>
NS_APPS_NAME_PREFIX: k8s Mounts: /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-bqz55 (ro) Conditions: Type Status Initialized True Ready True ContainersReady True PodScheduled True PodScheduled True Volumes: kube-api-access-bqz55: Type: Projected (a volume that contains injected data from multiple sources) TokenExpirationSeconds: 3607 ConfigMapOptional: sube-root-ca.crt ConfigMapOptional: kube-root-ca.crt ConfigMapOptional: sube-root-ca.crt ConfigMapOptional: sube-root-ca.crt Node-Selectors: sustable Node-Selectors: sustable Node-Selectors: node.kubernetes.io/not-ready:NoExecute op=Exists for 300s	EULA:		yes
Mounts:       /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-bqz55 (ro)         Conditions:       Type         Type       Status         Initialized       True         Ready       True         PodScheduled       True         volumes:       kube-api-access-bqz55:         Type:       Projected (a volume that contains injected data from multiple sources)         TokenExpirationSeconds:       3607         ConfigMapName:       kube-root-ca.crt         ConfigMapOptional: <nil>         DownwardAPI:       true         QoS Class:       Burstable         Node-Selectors:       <none>         Tolerations:       node.kubernetes.io/not-ready:NoExecute op=Exists for 300s</none></nil>	NS_APPS_NAME	PREFIX:	k8s
<pre>/var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-bqz55 (ro) Conditions: Type Status Initialized True Ready True PodScheduled True Volumes: kube-api-access-bqz55: Type: Projected (a volume that contains injected data from multiple sources) TokenExpirationSeconds: 3607 ConfigMapName: kube-root-ca.crt ConfigMapOptional: <ni>&gt; DownwardAPI: true OoS Class: Burstable Node-Selectors: <none> Tolerations: node.kubernetes.io/not-ready:NoExecute op=Exists for 300s</none></ni></pre>	Mounts:		
Conditions:         Type       Status         Initialized       True         Ready       True         ContainersReady       True         PodScheduled       True         Volumes:       ************************************	/var/run/sec	rets/kuber	netes.io/serviceaccount from kube-api-access-bqz55 (ro)
Type       Status         Initialized       True         Ready       True         PodScheduled       True         PodScheduled       True         Volumes:       ************************************	Conditions:		
Initialized       True         Ready       True         ContainersReady       True         PodScheduled       True         Volumes:       ************************************	Type		
Ready       True         ContainersReady       True         PodScheduled       True         Volumes:       true         kube-api-access-bqz55:       Projected (a volume that contains injected data from multiple sources)         TokenExpirationSeconds:       3607         ConfigMapName:       kube-root-ca.crt         ConfigMapOptional: <nil>         DownwardAPI:       true         QoS Class:       Burstable         Node-Selectors:       <none>         Tolerations:       node.kubernetes.io/not-ready:NoExecute op=Exists for 300s</none></nil>	Initialized		
ContainersReady True PodScheduled True Volumes: kube-api-access-bqz55: Type: Projected (a volume that contains injected data from multiple sources) TokenExpirationSeconds: 3607 ConfigMapAme: kube-root-ca.crt ConfigMapAme: kube-root-ca.crt ConfigMapOptional: <nl> TownwardAPI: true QoS Class: Burstable Node-Selectors: &lt;</nl>	Ready		
PodScheduled       True         Volumes:       kube-api-access-bqz55:         Type:       Projected (a volume that contains injected data from multiple sources)         TokenExpirationSeconds:       3607         ConfigMapName:       kube-root-ca.crt         ConfigMapOptional: <nl>         DownwardAPI:       true         QoS Class:       Burstable         Node-Selectors:       <none>         Tolerations:       node.kubernetes.io/not-ready:NoExecute op=Exists for 300s</none></nl>	ContainersReady		
Volumes:       kube-api-access-bqz55:         Type:       Projected (a volume that contains injected data from multiple sources)         TokenExpirationSeconds:       3607         ConfigMapName:       kube-root-ca.crt         ConfigMapOptional: <ni>&gt;         DownwardAPI:       true         QoS Class:       Burstable         Node-Selectors:       <non>         Tolerations:       node.kubernetes.jo/not-ready:NoExecute op=Exists for 300s</non></ni>	PodScheduled		
kube-api-access-bq255:       Projected (a volume that contains injected data from multiple sources)         TokenExpirationSeconds:       3607         ConfigMapName:       kube-root-ca.crt         ConfigMapOptional: <nl>         DownwardAPI:       true         QoS Class:       Burstable         Node-Selectors:       <none>         Tolerations:       node.kubernetes.io/not-ready:NoExecute op=Exists for 300s</none></nl>	Volumes:		
Type:     Projected (a volume that contains injected data from multiple sources)       TokenExpirationSeconds:     3607       ConfigMapName:     kube-root-ca.crt       ConfigMapOptional: <nl>       DownwardAPI:     true       QoS Class:     Burstable       Node-Selectors:     <none>       Tolerations:     node.kubernetes.io/not-ready:NoExecute op=Exists for 300s</none></nl>	kube-api-access-	bqz55:	
TokenExpirationSeconds:       3607         ConfigMapName:       kube-root-ca.crt         ConfigMapOptional: <nil>         DownwardAPI:       true         QoS Class:       Burstable         Node-Selectors:       <non>         Tolerations:       node.kubernetes.io/not-ready:NoExecute op=Exists for 300s</non></nil>	Type:		Projected (a volume that contains injected data from multiple sources)
ConfigMapName:     kuberoot-ca.crt       ConfigMapOptional: <nl>       DownwardAPI:     true       QoS Class:     Burstable       Node-Selectors:     <none>       Tolerations:     node.kubernetes.io/not-ready:NoExecute op=Exists for 300s</none></nl>	TokenExpirationSeconds:		3607
ConfigMapOptional: <ni>         DownwardAPI:       true         QoS Class:       Burstable         Node-Selectors:          Tolerations:       node.kubernetes.io/not-ready:NoExecute op=Exists for 300s</ni>	ConfigMapName:		
DownwardAPI:     true       QoS Class:     Burstable       Node-Selectors: <none>       Tolerations:     node.kubernetes.io/not-ready:NoExecute op=Exists for 300s</none>	ConfigMapOptional:		
QoS Class:     Burstable       Node-Selectors: <none>       Tolerations:     node.kubernetes.io/not-ready:NoExecute op=Exists for 300s</none>	DownwardAPI:		
Node-Selectors: <pre></pre> Concestions:  Concestion node.kubernetes.io/not-ready:NoExecute op=Exists for 300s	QoS Class:		
Tolerations: node.kubernetes.io/not-ready:NoExecute op=Exists for 300s	Node-Selectors:		
	Tolerations:		node.kubernetes.io/not-ready:NoExecute op=Exists for 300s

## What's next

You can see the following topic for comprehensive information on deploying NetScaler Ingress Controller and to customize your installation accordingly.

- Deployment topologies
- Dual-tier ingress

## **Deployment topologies**

#### May 9, 2024

NetScalers can be combined in powerful and flexible topologies that complement organizational boundaries. Dual-tier deployments employ high-capacity hardware or virtualized NetScalers (NetScaler MPX and VPX) in the first tier to offload security functions and implement relatively static organizational policies while segmenting control between network operators and Kubernetes operators.

In Dual-tier deployments, the second tier is within the Kubernetes Cluster (using the NetScaler CPX) and is under control of the service owners. This setup provides stability for network operators, while

allowing Kubernetes users to implement high-velocity changes. Single-tier topologies are suited to organizations that need to handle high rates of change.

## Single-Tier topology

In a Single-Tier topology, NetScaler MPX or VPX devices proxy the (North-South) traffic from the clients to microservices inside the cluster. The NetScaler Ingress Controller is deployed as a standalone pod in the Kubernetes cluster. The controller automates the configuration of NetScalers (MPX or VPX) based on the changes to the microservices or the Ingress resources.



## **Dual-Tier topology**

In Dual-Tier topology, NetScaler MPX or VPX devices in Tier-1 proxy the traffic (North-South) from the client to NetScaler CPXs in Tier-2. The Tier-2 NetScaler CPX then routes the traffic to the microservices in the Kubernetes cluster. The NetScaler Ingress Controller deployed as a standalone pod configures the Tier-1 devices. And, the sidecar controller in one or more NetScaler CPX pods configures the associated NetScaler CPX in the same pod.



## **Cloud topology**

Kubernetes clusters in public clouds such as Amazon Web Services (AWS), Google Cloud, and Microsoft Azure can use their native load balancing services such as, AWS Elastic Load Balancing, Google Cloud Load Balancing, and Microsoft Azure NLB as the first (relatively static) tier of load balancing to a second tier of NetScaler CPX. NetScaler CPX operates inside the Kubernetes cluster with the sidecar Ingress controller. The Kubernetes clusters can be self-hosted or managed by the cloud provider (for example, AWS EKS, Google GKE and Azure AKS) while using the NetScaler CPX as the Ingress. If the cloud-based Kubernetes cluster is self-hosted or self-managed, the NetScaler VPX can be used as the first tier in a Dual-tier topology.

Cloud deployment with NetScaler (VPX) in tier-1:



Cloud deployment with Cloud LB in tier-1:



## Service mesh lite

An Ingress solution typically performs layer 7 proxy functions for traffic from client to microservices inside the Kubernetes cluster (north-south traffic). The Service mesh lite architecture uses the same

Ingress solution to manage the traffic across services with in the Kubernetes cluster (east-west traffic) as well. Typically, a service mesh solution is used for managing east-west traffic, but it is heavier and complex to manage. Service mesh lite solution is a simplified version of the service mesh architecture and ideal when there is a need to manage both north-south and east-west traffic management. In a service mesh, there are as many sidecar proxies as the number of applications. But, in the service mesh lite architecture, a proxy is deployed as a standalone proxy managing multiple east-west connections. Hence, the Service mesh lite solution is lighter compared to a service mesh because the number of proxies required are less.

In a standard Kubernetes deployment, east-west traffic traverses the built-in kube-proxy deployed in each node. Kube-proxy being a L4 proxy can only do TCP/UDP based load balancing without the benefits of L7 proxy.

NetScaler (MPX, VPX, or CPX) can provide such benefits for east-west traffic such as:

- Mutual TLS or SSL offload
- Content based routing, allow, or block traffic based on HTTP or HTTPS header parameters
- Advanced load balancing algorithms (for example, least connections, least response time and so on.)
- Observability of east-west traffic through measuring golden signals (errors, latencies, saturation, or traffic volume).NetScaler ADM Service Graph is an observability solution to monitor and debug microservices.

## For more information, see Service mesh lite.

Following are some of the scenarios when service mesh lite topology is recommended:

- When you need both the north-south and the east-west traffic management for microservices.
- When you need the east-west traffic management through a proxy deployed as a standalone proxy and not as sidecar proxies to microservices.
- When you need the proxy inside the Kubernetes cluster to perform both north-south and eastwest traffic management.
- When you need the benefits of service mesh, but wants a lighter and simpler solution.

## Services of type LoadBalancer

Services of type LoadBalancer in Kubernetes enables you to directly expose services to the outside world without using an ingress resource. It is made available only by cloud providers, who spin up their own native cloud load balancers and assign an external IP address through which the service is accessed. This helps you to deploy microservices easily and expose them outside the Kubernetes cluster.

By default, in a bare metal Kubernetes cluster, service of type LoadBalancer simply exposes NodePorts for the service. And, it does not configure external load balancers.

The NetScaler Ingress Controller supports the services of type LoadBalancer. You can create a service of type LoadBalancer and expose it using the ingress NetScaler in Tier-1. The ingress NetScaler provisions a load balancer for the service and an external IP address is assigned to the service. The NetScaler Ingress Controller allocates the IP address using NetScaler IPAM controller.



For more information, see Expose services of type LoadBalancer.

## **Services of type NodePort**

By default, Kubernetes services are accessible using the cluster IP address. The cluster IP address is an internal IP address that can be accessed within the Kubernetes cluster. To make the service accessible from the outside of the Kubernetes cluster, you can create a service of the type NodePort.

The NetScaler Ingress Controller supports services of type NodePort. Using the Ingress NetScaler and NetScaler Ingress Controller, you can expose the service of type NodePort to the outside world.

For more information, see Expose services of type NodePort.



## Guidelines for choosing the topology

The following information helps you to choose the right deployment among the topologies Single-Tier and Dual-tier based on your needs.

## Single-Tier (Unified Ingress)

Following are some of the scenarios when the Single-Tier (unified ingress) topology is recommended and the benefits:

- Easy to start and adopt because you can use the existing NetScaler as the ingress proxy in front of the Kubernetes cluster.
- When the Network team manages both NetScaler and the Kubernetes deployment.
- Your workload running as microservices is less and a Kubernetes proxy inside the Kubernetes cluster is not required.
- More suitable for north-south traffic deployments.

## **Dual-Tier**

Following are some of the scenarios when the Dual-Tier ingress topology is preferred and the benefits:

- When you have significant workload running as microservices there is a need for a proxy inside the Kubernetes cluster.
- When the external proxy (managed by the network team) and Kubernetes proxies (managed by the platform team) are managed by two different teams.
- You need segregation of functions for proxies external to Kubernetes and for proxies inside Kubernetes. For example, WAF, and SSL offload on external NetScaler and policy enforcement and rate limiting on the Kubernetes proxy.
- The proxy inside the Kubernetes cluster performs north-south traffic management only.

## **Deployment using Helm charts**

For deploying NetScaler cloud native topologies, there are various options available using YAML and Helm charts. Helm charts are one of the easiest ways for deployment in a Kubernetes environment. When you deploy using the Helm charts, you can use a values.yaml file to specify the values of the configurable parameters instead of providing each parameter as an argument.

## **Unified ingress**

## June 26, 2025

The Unified ingress deployment is used to perform load balancing of North-South traffic. North-South traffic is the traffic that enters or exits Kubernetes cluster, that is, the traffic that flows between the client and the front-end microservices. In this deployment, NetScaler Ingress Controller automates the configuration of NetScaler, which load balances traffic to microservices in the Kubernetes environment. NetScaler Ingress Controller runs as a pod that monitors the Kubernetes API server and configures NetScaler MPX or NetScaler VPX.

## Deploy NetScaler Ingress Controller as a standalone pod in the Kubernetes cluster

Use the netscaler-ingress-controller Helm chart to deploy NetScaler Ingress Controller as a pod in your Kubernetes cluster.

The installation involves NetScaler Ingress Controller pod creation, cluster role creation, cluster role bindings, and service account creation.

When you deploy NSIC using the Helm chart, you can use a values.yaml file to specify the values of the configurable parameters instead of providing each parameter as an argument.

#### Note:

NSIC can also work with RBAC roles.

## Prerequisites

- Determine the NS\_IP (NetScaler IP) address using which NetScaler Ingress Controller communicates with NetScaler. The IP address might be any one of the following IP addresses depending on the type of NetScaler deployment:
  - NSIP (for standalone appliances) The management IP address of a standalone NetScaler appliance. For more information, see IP Addressing in NetScaler.
  - SNIP (for appliances in High Availability mode) The subnet IP address. For more information, see IP Addressing in NetScaler.
  - CLIP (for appliances in Cluster mode) The cluster management IP (CLIP) address for a cluster NetScaler deployment. For more information, see IP addressing for a cluster.
- The user name and password of a system user account of NetScaler MPX or NetScaler VPX used as the Ingress device. NetScaler needs to have a system user account (non-default) with certain privileges so that the NetScaler Ingress Controller can configure NetScaler MPX or NetScaler VPX. For instructions to create the system user account on NetScaler, see Create a system user account for NetScaler Ingress Controller in NetScaler.

You can directly pass the user name and password or use Kubernetes secrets. If you want to use Kubernetes secrets, create a secret for the user name and password using the following command:

```
1 kubectl create secret generic nslogin --from-literal=username=<
    username> --from-literal=password=<password>
```

- Version 1.21 and later in the Kubernetes environment.
- Version 4.8 and later in the OpenShift platform.
- Ensure that you have installed Helm version 3.x or later. To install Helm, see here.

**Create a system user account for NetScaler Ingress Controller in NetScaler** NetScaler Ingress Controller configures NetScaler using a system user account of NetScaler. The system user account should have certain privileges so that NetScaler Ingress Controller has permission to configure the following on NetScaler:

- Add, delete, or view content switching (CS) virtual server
- Configure CS policies and actions

- Configure Load Balancing (LB) virtual server
- Configure service groups
- Configure SSL certkeys
- Configure routes
- Configure user monitors
- Add system file (for uploading SSL certkeys from Kubernetes)
- Configure virtual IP address (VIP)
- Check the status of the NetScaler appliance

To create the system user account, perform the following:

- 1. Log on to NetScaler. Perform the following steps:
  - a) Use an SSH client, such as PuTTy, to open an SSH connection to NetScaler.
  - b) Log on to NetScaler using the administrator credentials.
- 2. Create the system user account using the following command:

```
1 add system user <username> <password>
```

For example:

```
1 add system user nsic mypassword
```

3. Create a policy to provide required permissions to the system user account. Use the following command:

```
add cmdpolicy nsic-policy ALLOW '^(\?!shell)(\?!sftp)(\?!scp)
(\?!batch)(\?!source)(\?!.*superuser)(\?!.*nsroot)(\?!install)
(\?!show\s+system\s+(user|cmdPolicy|file))(\?!(set|add|rm|create|
export|kill)\s+system)(\?!(unbind|bind)\s+system\s+(user|group))
(\?!diff\s+ns\s+config)(\?!(set|unset|add|rm|bind|unbind|switch)
\s+ns\s+partition).*|(^install\s*(wi|wf))|(^\S+\s+system\s+file)
^(\?!shell)(\?!sftp)(\?!scp)(\?!batch)(\?!source)(\?!.*superuser)
(\?!.*nsroot)(\?!install)(\?!show\s+system\s+(user|cmdPolicy|file
))(\?!(set|add|rm|create|export|kill)\s+system)(\?!(unbind|bind)\
s+system\s+(user|group))(\?!diff\s+ns\s+config)(\?!(set|unset|add
|rm|bind|unbind|switch)\s+ns\s+partition).*|(^install\s*(wi|wf))
|(^\S+\s+system\s+file)'
```

## Note:

The system user account has privileges based on the command policy that you define.

The command policy mentioned in *step 3* is similar to the built-in sysAdmin command policy with additional permission to upload files.

In the command policy specification provided, special characters which need to be escaped are already omitted to easily copy-paste into the NetScaler command line.

For configuring the command policy from NetScaler configuration wizard (GUI), use the following command policy specification.

```
^(?!shell)(?!sftp)(?!scp)(?!batch)(?!source)(?!.*superuser)(?!.*
nsroot)(?!install)(?!show\s+system\s+(user|cmdPolicy|file))(?!(
set|add|rm|create|export|kill)\s+system)(?!(unbind|bind)\s+system
\s+(user|group))(?!diff\s+ns\s+config)(?!(set|unset|add|rm|bind
|unbind|switch)\s+ns\s+partition).*|(^install\s*(wi|wf))|(^\S
+\s+system\s+file)^(?!shell)(?!sftp)(?!scp)(?!batch)(?!source)
(?!.*superuser)(?!.*nsroot)(?!install)(?!show\s+system\s+(user
|cmdPolicy|file))(?!(set|add|rm|create|export|kill)\s+system)
(?!(unbind|bind)\s+system\s+(user|group))(?!diff\s+ns\s+config)
(?!(set|unset|add|rm|bind|unbind|switch)\s+ns\s+partition).*|(^
install\s*(wi|wf))|(^\S+\s+system\s+file)
```

4. Bind the policy to the system user account using the following command:

```
1 bind system user nsic nsic-policy 0
```

## Deploy NetScaler Ingress Controller as a standalone pod in the Kubernetes cluster

To deploy NetScaler Ingress Controller as a standalone pod:

1. Add the NetScaler Helm chart repository to your local registry using the following command.

```
1 helm repo add netscaler https://netscaler.github.io/netscaler-helm
        -charts/
```

2. Install NetScaler Ingress Controller using the following command.

```
1 helm install nsic netscaler/netscaler-ingress-controller --set
nsIP=<NSIP>,license.accept=yes,adcCredentialSecret=<Secret-for-
NetScaler-credentials>
```

## Note:

You can also use a values.yaml file in the helm install command to specify the values of the configurable parameters instead of providing each parameter as a command-line argument. For example, helm install nsic netscaler/netscaler-ingress-controller -f values.yaml.

For information about the mandatory and optional parameters that you can configure during NSIC installation, see Configuration.

## Warning:

Do not create NSIC pod replicas in the Kubernetes cluster. When you create replicas of an NSIC pod by specifying replicatset>1, an unexpected behavior is observed with NSIC/VPX because the same APP\_PREFIX is shared among the NSIC PODs and each instance tries to configure the same resources.

## **Route addition in NetScaler MPX or NetScaler VPX**

For seamless functioning of services deployed in the Kubernetes cluster, NetScaler should be able to reach the underlying overlay network over which the pods are running. The feature-node-watch argument of NetScaler Ingress Controller can be used for automatic route configuration on NetScaler towards the pod network. Refer to Static Routing for more information. By default, feature-node-watch is set to false; set to true to enable automatic route configuration.

If your deployment uses a single NetScaler to load balance between multiple k8s clusters, there is a possibility of CNI subnets to overlap, causing the static routing to fail due to route conflicts. In such deployments, Policy Based Routing (PBR) can be used. PBR requires you to provide one or more subnet IP addresses unique for each kubernetes cluster either using an environment variable or ConfigMap, see PBR Support.

Use the following command to provide subnet IP addresses (SNIPs) to configure Policy Based Routes (PBR) on NetScaler.

```
1 helm install my-release netscaler/netscaler-ingress-controller --set
nsIP=<NSIP>,license.accept=yes,adcCredentialSecret=<Secret-for-
NetScaler-credentials>,nsSNIPS='[<NS_SNIP1>\, <NS_SNIP2>\, ...]'
```

## Note:

PBR can also be achieved by deploying NetScaler Node Controller. Netscaler Node Controller, by default, adds static routes while creating the VXLAN tunnel. To use Policy Based Routing (PBR) to avoid static route clash, both NetScaler Node Controller and NetScaler Ingress Controller have to work in conjunction and have to be started with specific arguments. For more information, refer Configure PBR using the NetScaler node controller.

Use the nsncPbr=<True/False> parameter in the helm install command of NSIC to inform NSIC that NetScaler Node Controller is configuring Policy Based Routes (PBR) on NetScaler.

```
1 helm install my-release netscaler/netscaler-ingress-controller --set
nsIP=<NSIP>,license.accept=yes,adcCredentialSecret=<Secret-for-
NetScaler-credentials>,clusterName=<unique-cluster-identifier>,
nsncPbr=<True/False>
```

## **Dual-tier ingress**

## August 23, 2024

In a dual-tier deployment, NetScaler VPX or NetScaler MPX is deployed outside the Kubernetes cluster (tier-1) and NetScaler CPXs are deployed inside the Kubernetes cluster (tier-2).

NetScaler MPX or NetScaler VPX devices in tier-1 proxy the traffic (North-South) from the client to NetScaler CPXs in tier-2. The tier-2 NetScaler CPX then routes the traffic to the microservices in the Kubernetes cluster. NetScaler Ingress Controller deployed as a standalone pod configures the tier-1 NetScaler. And, the sidecar NetScaler Ingress Controller in one or more NetScaler CPX pods configures the associated NetScaler CPX in the same pod.

The dual-tier deployment can be set up on Kubernetes in a bare metal environment or on public clouds such as AWS, GCP, or Azure.



## **Dual-tier ingress deployment**

## Prerequisites

• Create a Kubernetes cluster in cloud or on-premises. The Kubernetes cluster in the cloud can be a managed Kubernetes (for example: GKE, EKS, or AKS) or a custom created Kubernetes deployment.

## **Dual-tier Ingress deployment**

- 1. Deploy NetScaler MPX or NetScaler VPX on a multi-NIC deployment mode outside the Kubernetes cluster.
  - For instructions to deploy NetScaler MPX, see NetScaler documentation.
  - For instructions to deploy NetScaler VPX, see Deploy a NetScaler VPX instance.
- 2. Deploy a standalone NetScaler Ingress Controller to configure NetScaler VPX to direct the client requests to a CPX service in the Kubernetes cluster.

For information on how to deploy a standalone NetScaler Ingress Controller, see Unified Ingress. Set the ingressClass for this NSIC instance as nsic-vpx because the ingress resource deployed to configure NetScaler VPX in step 8 uses the same ingress class.

- 3. In cloud deployments, enable MAC-Based Forwarding mode on the tier-1 NetScaler VPX. As NetScaler VPX is deployed in multi-NIC mode, it does not have the return route to reach the POD CNI network or the client network. Hence, you must enable MAC-Based Forwarding mode on the tier-1 NetScaler VPX to handle this scenario.
- 4. Deploy a sample microservice by using the following command:

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: apps/v1
3 kind: Deployment
4 metadata:
5
   name: cnn-website
6
     labels:
7
      name: cnn-website
8
      app: cnn-website
9 spec:
10
   selector:
     matchLabels:
11
12
        app: cnn-website
  replicas: 2
13
   template:
14
15
     metadata:
16
        labels:
17
        name: cnn-website
```

```
app: cnn-website
18
19
      spec:
        containers:
         - name: cnn-website
21
22
          image: quay.io/sample-apps/cnn-website:v1.0.0
23
           ports:
24
           - name: http-80
            containerPort: 80
25
           - name: https-443
27
            containerPort: 443
28 ---
29 apiVersion: v1
30 kind: Service
31 metadata:
32
    name: cnn-website
    labels:
34
      app: cnn-website
35 spec:
    type: NodePort
     ports:
37
38
    - name: http-80
39
       port: 80
      targetPort: 80
40
41
     - name: https-443
     port: 443
42
43
      targetPort: 443
44 selector:
45
      name: cnn-website
46 EOF
```

- 5. Deploy NetScaler CPX as tier-2 ingress with NetScaler Ingress Controller as a sidecar.
  - a) Add the NetScaler Helm chart repository to your local registry by using the following command.

```
1 helm repo add netscaler https://netscaler.github.io/netscaler-
helm-charts/
```

b) Install NetScaler CPX with NetScaler Ingress Controller by using the following command.

```
1 helm install netscaler-cpx-with-ingress-controller netscaler/
netscaler-cpx-with-ingress-controller --set license.accept=
yes,serviceType.nodePort.enabled=True,ingressClass[0]=cpx
```

#### Note:

You can also use a values.yaml file in the helm install command to specify the values of the configurable parameters instead of providing each parameter as a command-line argument. For example,

```
helm install netscaler-cpx-with-ingress-controller netscaler/
netscaler-cpx-with-ingress-controller -f values.yaml
```

For information about the mandatory and optional parameters that you can configure during CPX with NSIC installation, see Configuration.

- 6. Create a Kubernetes secret using a certificate to copy the certificate to NetScaler when an ingress object is deployed. You can use the same secret in the ingress resource for both CPX and VPX. In this procedure, we'll use tls-secret. For information on how to generate Kubernetes secret, see Generate Kubernetes secret.
- 7. Create an ingress object to route traffic from NetScaler VPX or NetScaler MPX to NetScaler CPX in the Kubernetes cluster by using the following command:

```
kubectl apply -f - <<EOF</pre>
2 apiVersion: networking.k8s.io/v1
3 kind: Ingress
4 metadata:
5
     annotations:
6
       ingress.citrix.com/frontend-ip: <NSVIP> # vserver IP
7
       ingress.citrix.com/secure_backend: "True"
8
    name: vpxingress
9
    spec:
10
      tls:
         - secretName: tls-secret
12
      ingressClassName: nsic-vpx
13
      rules:
14
       - host: cnnwebsite.com
15
        http:
16
           paths:
17
           - backend:
18
               service:
                 name: netscaler-cpx-with-ingress-controller-cpx-
19
                     service
                 port:
21
                   number: 443
22
             path: /
23
             pathType: Prefix
24 ---
25 apiVersion: networking.k8s.io/v1
26 kind: IngressClass # If Helm chart is used to deploy NSIC, you
      need not create IngressClass
27 metadata:
28 name: nsic-vpx
29 spec:
    controller: citrix.com/ingress-controller
```

8. Create an ingress object for the tier-2 NetScaler CPX to route the traffic to the microservice application by using the following command:

```
1 kubectl apply -f - <<EOF
```

```
2 apiVersion: networking.k8s.io/v1
3 kind: Ingress
4 apiVersion: networking.k8s.io/v1
5 metadata:
6
    annotations:
7
    name: cpxingress
8 spec:
9 tls:
10
       - secretName: tls-secret
11
    ingressClassName: cpx
12
   rules:
       - host: cnnwebsite.com
13
14
        http:
15
           paths:
16
             - path: /
17
               pathType: Prefix
18
               backend:
19
                 service:
20
                   name: cnn-website
                   port:
21
22
                     number: 80
23 ---
24 apiVersion: networking.k8s.io/v1
25 kind: IngressClass # If Helm chart is used to deploy CPX with NSIC
      , you need not create IngressClass
26 metadata:
27
    name: cpx
28 spec:
     controller: citrix.com/ingress-controller
29
30 EOF
```

9. To send traffic to the cnn-website application, perform the following steps:

#### **For Windows:**

- a) Open NotePad as an administrator.
- b) Add VIP cnnwebsite.com entry in c:\Windows\System32\Drivers\etc\ hosts.

#### For MAC or Linux:

- a) Open the host files in a file editor using the following command: sudo nano /etc/ hosts.
- b) Add VIP cnnwebsite.com entry.
- 10. Send the traffic using the following command:

```
1 curl -kv https://cnn-website
```

## Warning:

Do not create NSIC pod replicas in the Kubernetes cluster. When you create replicas of an NSIC pod by specifying replicatset>1, an unexpected behavior is observed with NSIC/VPX because the same APP\_PREFIX is shared among the NSIC PODs and each instance tries to configure the same resources.

## Multi-cluster ingress

September 24, 2024

## Introduction

Multi-cluster Kubernetes solutions are ideal for distributing the workload across multiple clusters. NetScaler multi-cluster ingress solution enables NetScaler to load balance applications distributed across clusters using a single front-end IP address. The load-balanced applications can be either the same application, different applications of the same domain, or entirely different applications.

Earlier, to load balance applications in multiple clusters, a dedicated content switching virtual server was required on NetScaler for each instance of NetScaler Ingress Controller (NSIC) running in the clusters. With NetScaler multi-cluster ingress solution, multiple ingress controllers can share a content switching virtual server. Therefore, applications deployed across clusters can be load balanced using the same content switching virtual server IP (VIP) address.

To summarize, the multi-cluster ingress solution optimizes the use of load-balancing resources, thereby reducing operational costs.

Note:

The multi-cluster ingress solution is supported starting from NSIC version 2.0.6.

## **Deployment topology**

The following diagram describes the multi-cluster ingress deployment topology for two Kubernetes clusters in a data center/site. Here, NetScaler load balances applications distributed across clusters using a single front-end IP address.


- An NSIC is deployed in both the clusters. Both NSIC instances configure the same NetScaler.
- Depending on the nature of applications deployed in the clusters, we typically have the following use cases:
  - A different application of the same domain (company.website.com) deployed in each Kubernetes cluster: App-B in cluster1 and App-C in cluster2.

Here, a content switching virtual server on NetScaler load balances traffic across App–B and App–C. Separate content switching policies are configured for each application.

- The same application deployed across both Kubernetes clusters.

Here, a content switching virtual server on NetScaler load balances traffic for the same application App–A deployed in both the clusters. Only one content switching policy is created for App–A. which is used by both running instances (endpoints) of the application.

Let's understand how the same front-end IP address is used to load balance traffic to applications App-B and App-c deployed across cluster1 and cluster2 with the following diagram.



Configuring the following resources enables the multi-cluster ingress solution: Listener and Ingress.

## Listener

• The listener resource deployed in a cluster creates a content switching virtual server on NetScaler using the VIP specified in the listener YAML. In the multi-cluster ingress setup, the same listener resource is deployed in all the clusters.

Listener resource in the multi-cluster ingress solution handles the front-end traffic management. You must provide all secrets, ciphers, and front-end profile configuration required for the content switching virtual server. For more information on listener CRD, see Listener. Details on deploying listener resources for the multi-cluster ingress setup are available in the Steps to deploy a multi-cluster ingress solution section.

#### Ingress

• The ingress resources deployed in each cluster to route traffic to an application in the multicluster setup refer to the same listener resource with the shared VIP address. Therefore, NSIC refers to the same content virtual server that exists and creates only the content switching policies on NetScaler that are later bound to the content switching virtual server.

In this case, the content switching policy for each application, that is, CSPOL-App-B and CSPOL-App-C, are bound to the same content switching virtual server. When a request for example, company.

website/App-B is sent to the content switching virtual server IP (VIP) address, the content switching virtual server applies the CSPOL-App-B policy and routes the request to the LBvserver-App-B, which sends it to the App-B service.

## Multi-cluster ingress deployment

In this procedure, we deploy the same HTTPS application in two clusters, deploy the required listener and ingress resources in both the clusters, such that NetScaler VPX or NetScaler MPX load balances these applications using a single front-end IP address.

## Prerequisites

- Access to Kubernetes clusters hosted in cloud or on-premises. The Kubernetes cluster in the cloud can be a managed Kubernetes (for example: GKE, EKS, or AKS) or a custom created Kubernetes deployment. Kubernetes or OpenShift must be properly configured and operational. Ensure the network connectivity between NetScaler and the cluster's pod network.
- For Kubernetes deployment, you need access to Kubernetes clusters running version 1.21 or later.
- For OpenShift deployment, you need access to OpenShift clusters running version 4.11 or later.
- You have installed Helm version 3.x or later. To install Helm, see here.
- NetScaler MPX/VPX is deployed in a standalone, HA, or cluster setup depending on your specific needs.
  - For instructions to deploy NetScaler MPX, see NetScaler documentation.
  - For instructions to deploy NetScaler VPX, see Deploy a NetScaler VPX instance.
- Determine the NSIP (NetScaler IP) address using which NetScaler Ingress Controller communicates with NetScaler. The IP address might be any one of the following IP addresses depending on the type of NetScaler deployment:
  - NSIP (for standalone appliances) The management IP address of a standalone NetScaler appliance. For more information, see IP Addressing in NetScaler.
  - SNIP (for appliances in High Availability mode) The subnet IP address. For more information, see IP Addressing in NetScaler.
  - CLIP (for appliances in Cluster mode) The cluster management IP (CLIP) address for a cluster NetScaler deployment. For more information, see IP addressing for a cluster.
- A user account in NetScaler VPX or NetScaler MPX. NetScaler Ingress Controller uses a system user account in NetScaler to configure NetScaler MPX or NetScaler VPX. For instructions to create

the system user account on NetScaler, see Create System User Account for NetScaler Ingress Controller in NetScaler.

You can directly pass the user name and password or use Kubernetes secrets. If you want to use Kubernetes secrets, create a secret for the user name and password using the following command:

1 kubectl create secret generic nslogin --from-literal=username=< username> --from-literal=password=<password>

#### Steps to deploy a multi-cluster ingress solution

Follow these steps to deploy a multi-cluster ingress solution to expose the Cloud Native Networking (CNN) application deployed across two Kubernetes clusters using NetScaler and a shared front-end IP address.

Note:

Repeat the following steps in both the clusters to set up a multi-cluster ingress solution. There are exceptions added in some steps highlighting what needs to be done in a particular cluster; perform the steps accordingly.

#### 1. Deploy the NetScaler CNN application by using the following YAML.

The CNN application is an HTTP-based application that lists the solutions offered under the NetScaler cloud native portfolio.

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: apps/v1
3 kind: Deployment
4 metadata:
5
    name: cnn-website
6
    labels:
7
      name: cnn-website
8
      app: cnn-website
9 spec:
10 selector:
11
     matchLabels:
12
        app: cnn-website
13 replicas: 2
14 template:
15
     metadata:
16
        labels:
17
          name: cnn-website
18
           app: cnn-website
19
       spec:
20
        containers:
21
         - name: cnn-website
22
           image: quay.io/sample-apps/cnn-website:v1.0.0
```

```
ports:
24
          - name: http-80
            containerPort: 80
26
           - name: https-443
27
            containerPort: 443
28 ---
29 apiVersion: v1
30 kind: Service
31 metadata:
   name: cnn-website
   labels:
34
      app: cnn-website
35 spec:
    type: NodePort
36
    ports:
38
     - name: http-80
39
     port: 80
40
      targetPort: 80
    - name: https-443
41
42
      port: 443
    targetPort: 443
43
44
    selector:
45
    name: cnn-website
46 EOF
```

#### Note:

For an OpenShift deployment, run the following command oc adm policy addscc-to-user anyuid system:serviceaccount:<namespace>:default to grant a specific Security Context Constraint (SCC) to a service account. Replace < namespace> with the actual namespace where you have deployed the CNN application.

2. Add the NetScaler Helm chart repository to your local registry using the following command.

1 helm repo add netscaler https://netscaler.github.io/netscaler-helm -charts/

If the NetScaler Helm chart repository is already added to your local registry, use the following command to update the repository:

1 helm repo update netscaler

3. Update values.yaml to configure NetScaler Ingress Controller as described as following.

Example values.yaml for cluster1

```
1 license:
2 accept: yes
3 adcCredentialSecret: nslogin # K8s Secret Name
4 nsIP: <x.x.x> # CLIP (for appliances in Cluster mode), SNIP (for
appliances in High Availability mode), NSIP (for standalone
appliances)
```

```
openshift: false # set to true for OpenShift deployments
entityPrefix: cluster1 # unique for each NSIC instance.
clusterName: cluster1
ingressClass: ['nsic-vpx'] # ingress class used in the ingress
multiClusterPrefix: mc # Multi-cluster prefix for the NSIC
instance. Same value must be specified for a set of NSIC
instances configuring NetScaler in multi-cluster setup.
# serviceClass- To use service type LB, specify the service
class
```

Example values.yaml for cluster2

```
1
     license:
2
       accept: yes
     adcCredentialSecret: nslogin # K8s Secret Name
3
     nsIP: <x.x.x> # CLIP (for appliances in Cluster mode), SNIP (for
4
         appliances in High Availability mode), NSIP (for standalone
         appliances)
5
     openshift: false # set to true for OpenShift deployments
     entityPrefix: cluster2 # unique for each NSIC instance.
6
7
    clusterName: cluster2
     ingressClass: ['nsic-vpx'] # ingress class used in the ingress
8
        resources
    multiClusterPrefix: mc # Multi-cluster prefix for the NSIC
9
        instance. Same value must be specified for a set of NSIC
        instances configuring NetScaler in multi-cluster setup.
10
     # serviceClass- To use service type LB, specify the service
        class
```

#### Note:

For an OpenShift deployment, set the openshift parameter to **true** in values.yaml.

For information about the mandatory and optional parameters that you can configure during NSIC installation, see Configuration.

4. Deploy NetScaler Ingress Controller with the modified values.yaml.

1 helm install nsic netscaler/netscaler-ingress-controller -f values
 .yaml

#### Note:

If an earlier version of NSIC is already deployed in the cluster, deploy the listener CRD specification using the following command: kubectl apply -f https:// raw.githubusercontent.com/netscaler/netscaler-k8s-ingresscontroller/master/crd/contentrouting/Listener.yaml.

5. Deploy the following listener CRD resource.

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: citrix.com/v1
3 kind: Listener
4 metadata:
5
    name: mc-listener
6 spec:
7
   multicluster: True
8 ingressclass: nsic-vpx
9
   protocol: 'https'
10 vip: <Provide the shared front-end IP address>
11 certificates:
   # you need to specify either secret name or pre-configured cert
12
        -keyname
13
       - secret:
14
           name: <Provide k8s secret name> # provide k8s secret
              containing cert-key of the application
15
        default: true
       - preconfigured: <ADC-certkeyname> # provide pre-configured
16
          cert-key from ADC.
17 EOF
```

- Update the vip with the virtual IP address that is used to expose the application deployed in both the clusters.
- multicluster must be set as True.
- Either update a preconfigured secret with the ssl certificate-key name created on NetScaler or update the name of Kubernetes TLS Secret in the secret.name section. Refer Listener.certificates.

In this example, the listener resource mc-listener specifies the front-end configuration such as VIP and secrets. It creates a content switching virtual server in NetScaler on port 443 for HTTPS traffic.

6. Deploy the following ingress resource to expose the CNN application.

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: networking.k8s.io/v1
3 kind: Ingress
4 metadata:
5 name: frontend-ingress
   annotations:
6
    ingress.citrix.com/listener: mc-listener
7
8 spec:
    ingressClassName: nsic-vpx
9
10
   rules:
11
       - http:
12
           paths:
13
             - path: /
14
               pathType: Prefix
15
               backend:
```

```
service:
16
17
                   name: cnn-website
18
                   port:
                      number: 80
19
20 ---
21 apiVersion: networking.k8s.io/v1
22 kind: IngressClass
23 metadata:
24
   name: nsic-vpx
25 spec:
     controller: citrix.com/ingress-controller
27 EOF
```

• The listener resource created in the previous step must be passed using ingress. citrix.com/listener annotation.

In this example, the ingress resource frontend-ingress refers to the Listener resource mc-listener (created in step 5).

• The ingress class is mentioned as nsic-vpx.

Ingress creates a content switching policy, a load balancing virtual server, a service group and binds the application pod IP addresses as service group members.

#### Notes:

- When load balancing different applications in the multi-cluster ingress setup, separate content switching policies are created for each application. In such cases, if you require a particular sequence for policy binding, you must assign a priority number to the content switching policies by using the ingress.citrix.com/multicluster-policy-priority-order annotation. For more information, see Policy bindings.
- For information about listener and ingress resources when load balancing different applications in the multi-cluster ingress setup, see Multi-cluster ingress setup with different applications.
- 7. To validate the configuration, replace <VIP> with the virtual IP address provided in the listener resource in step 5.

```
1 curl -kv https://<VIP>
```

#### **Advanced use cases**

When the same application is deployed across multiple clusters, diverse deployment scenarios with varying traffic distribution requirements arise. The default behavior, known as Active-Active, involves all instances of the application receiving traffic based on the load balancing method. Let's

see the advanced use cases considering the same application deployed in two clusters as described here.

**Active-backup mode** In this mode, one application is always active, which receives traffic all the time. The other application receives traffic only when this cluster or application is down. You can define which application acts as a backup using ingress.citrix.com/multicluster-backup -order ingress annotation.

## Note:

The same listener resource must be deployed in both the clusters. For an example listener resource, see step 5 in the Steps to deploy a multi-cluster ingress solution section.

	Cluster 1 (Active)	Cluster2 (Backup)
Ingress annotation	ingress.citrix.com/multicluster- backup-order: "1"	ingress.citrix.com/multicluster- backup-order: "2"
Information	The default is active (1). Annotation can be skipped	Annotation is mandatory for the application to be considered as a backup.

• Application deployed in cluster1 actively receives traffic, whereas cluster2 acts as a backup. The application in cluster2 receives traffic only when cluster1 or application in cluster1 is down.

## Canary mode

Note:

For the following deployments, the same listener resource must be deployed in each cluster. For an example listener resource, see step 5 in the Steps to deploy a multi-cluster ingress solution section.

We support canary deployments with two distinct strategies: "Canary by weight" and "Canary by header".

**Canary deployment by weight** You can implement canary deployment by assigning different weights to instances of an application. For instance, if there are two clusters and a new version of the application is being rolled out, you can allocate a certain percentage of traffic to the new version (canary) and the remaining traffic to the existing version. This canary mode allows gradual rollout and monitoring of the new version's performance in a controlled manner.

	Cluster1 (Version1)	Cluster2 (Version2)
Ingress annotation	NA	ingress.citrix.com/multicluster- canary-weight: "10"
Information	Annotation is not required	Annotation is mandatory. This application receives 10% of client traffic.

**Canary deployment by header** You can implement a canary deployment based on the specific HTTP headers. For example, a specific header in the HTTP request such as "X-Canary-Version" can be used to control the routing of traffic. In this case, requests with the "X-Canary-Version" header are directed to the canary version of the app, while others are routed to the stable version. This canary mode provides a more granular control over canary deployments, allowing you to target specific subsets of users or requests for testing purposes.

	Cluster1 (Version1)	Cluster2 (Version2)
Annotation	NA	ingress.citrix.com/multicluster- canary-by-header: "version2"
Information	Annotation is not required	Annotation is mandatory for the canary application. All client traffic with the version2 header is sent to cluster2.

#### **Policy bindings**

If you require a particular sequence for policy binding, you must assign a priority number to the content switching policies by using the ingress.citrix.com/multicluster-policy-priority-order annotation. Lower the number, the higher the priority.

Let's consider the following ingress resource example to understand policy bindings.

```
1 kubectl apply -f - <<EOF
2 apiVersion: networking.k8s.io/v1
3 kind: Ingress
4 metadata:
5 name: mc-ing
6 annotations:
7 ingress.citrix.com/listener: mc-listener
8 ingress.citrix.com/multicluster-policy-priority-order: '{
9 "frontend": {
```

```
"80": "3", "9443": "1"
11
      "backend": "2" }
12
13 spec:
14
    ingressClassName: nsic-vpx
15
     rules:
16
       - http:
17
            paths:
18
              - path: /
19
                pathType: Prefix
20
                backend:
21
                  service:
22
                    name: frontend
23
                     port:
24
                      number: 80
25
              - path: /abc
                pathType: Prefix
26
27
                backend:
                  service:
28
29
                    name: frontend
                     port:
31
                       number: 9443
32
              - path: /xyz
33
                pathType: Prefix
34
                backend:
                  service:
                    name: backend
37
                     port:
38
                       number: 80
39
   EOF
```

This example resource includes ingress.citrix.com/multicluster-policy-priority -order annotation, which defines the priority order for content switching policies. Rather than assigning priorities randomly, NetScaler Ingress Controller uses the priority values you have provided to bind content switching policies to the content switching virtual server.

For the ingress.citrix.com/multicluster-policy-priority-order: '{"Front end": {"80": "3", "9443": "1"}, "back-end": "2"}'annotation, the priority is assigned as following:

- The content switching policy associated with the frontend:80 service is bound to the content switching virtual server with a priority of 3.
- The content switching policy associated with the frontend: 9443 service is bound to the content switching virtual server with a priority of 1.
- The content switching policy associated with the backend: 80 service is bound to the content switching virtual server with a priority of 2.

## Warning:

When a single service is mentioned in the ingress resource with different ports, you must explic-

itly specify a priority number for each port. Otherwise, random priorities are assigned.

#### Multi-cluster ingress setup with different applications

In the multi-cluster ingress setup for load balancing different applications, the listener resource deployed in each cluster must be the same. The ingress resource in each cluster must refer to the same listener. For the example described in the deployment topology section, mc-listener is the listener resource deployed in both cluster1 and cluster2. The following sample ingress resources deployed in cluster1 and cluster2 refer to mc-listener.

Ingress resource to expose App-B in cluster1:

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: networking.k8s.io/v1
3 kind: Ingress
4 metadata:
5 name: app-b
6 annotations:
       ingress.citrix.com/listener: mc-listener
7
8 spec:
9
   ingressClassName: cluster1
10
     rules:
    - http:
11
       paths:
12
13
         - path: /App-B
14
           pathType: Prefix
15
           backend:
16
            service:
17
              name: app-b-svc
18
               port:
19
                 number: 80
20 EOF
```

Ingress resource to expose App-C in cluster2:

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: networking.k8s.io/v1
3 kind: Ingress
4 metadata:
5 name: app-c
6 annotations:
7
      ingress.citrix.com/listener: mc-listener
8 spec:
   ingressClassName: cluster2
9
10 rules:
11
    - http:
         paths:
12
         - path: /App-C
13
14
           pathType: Prefix
15
           backend:
16
             service:
```

```
      17
      name: app-c-svc

      18
      port:

      19
      number: 80

      20
      EOF
```

# Deploy the NetScaler Ingress Controller on a Rancher managed Kubernetes cluster

#### December 31, 2023

Rancher is an open-source platform with an intuitive user interface that helps you to easily deploy and manage Kubernetes clusters. Rancher supports Kubernetes clusters on any infrastructure be on cloud or on-premises deployment. Rancher also allows you to centrally manage multiple clusters running across your organization.

The NetScaler Ingress Controller is built around the Kubernetes Ingress and it can automatically configure one or more NetScalers based on the Ingress resource configuration. You can deploy the NetScaler Ingress Controller in a Rancher managed Kubernetes cluster to extend the advanced load balancing and traffic management capabilities of NetScaler to your cluster.

#### Prerequisites

You must create a Kubernetes cluster and import the cluster on the Rancher platform.

## **Deployment options**

You can either deploy NetScaler CPXs as pods inside the cluster or deploy a NetScaler MPX or VPX appliance outside the Kubernetes cluster.

Based on how you want to use NetScaler, there are two ways to deploy the NetScaler Ingress Controller in a Kubernetes cluster on the Rancher platform:

- As a sidecar container alongside NetScaler CPX in the same pod: In this mode, NetScaler Ingress Controller configures the NetScaler CPX.
- As a standalone pod in the Kubernetes cluster: In this mode, you can control the NetScaler MPX or VPX appliance deployed outside the cluster.

## Deploy the NetScaler Ingress Controller as a sidecar with NetScaler CPX

In this deployment, you can use the NetScaler CPX instance for load balancing the North-South traffic to microservices in your Kubernetes cluster. NetScaler Ingress Controller is deployed as a sidecar alongside the NetScaler CPX container in the same pod using the citrix-k8s-cpx-ingress. yaml file.

Perform the following steps to deploy the NetScaler Ingress Controller as a sidecar with NetScaler CPX on the Rancher platform.

1. Download the citrix-k8s-cpx-ingress.yaml file using the following command.

```
1 wget https://raw.githubusercontent.com/citrix/citrix-k8s-ingress-
controller/master/deployment/baremetal/citrix-k8s-cpx-ingress.
yml
```

- 2. On the Rancher GUI cluster page, select Clusters from Global view.
- 3. From the Clusters page, open the cluster that you want to access.
- 4. Click Launch kubectl to open a terminal for interacting with your Kubernetes cluster.
- 5. Create a file named cpx.yaml in the launched terminal and then copy the contents of the modified citrix-k8s-cpx-ingress.yaml file to the cpx.yaml file.
- 6. Deploy the newly created YAML file using the following command.

1 kubectl create -f cpx.yaml

7. Verify if NetScaler Ingress Controller is deployed successfully using the following command.

```
1 kubectl get pods --all-namespaces
```

## Deploy the NetScaler Ingress Controller as a standalone pod

In this deployment, NetScaler Ingress Controller which runs as a stand-alone pod allows you to control the NetScaler MPX, or VPX appliance from the Kubernetes cluster. You can use the citrix-k8s-ingress-controller.yaml file for this deployment.

**Before you begin:** Ensure that you complete all the prerequisites required for deploying the NetScaler Ingress Controller.

To deploy the NetScaler Ingress Controller as a standalone pod on the Rancher platform:\*\*

1. Download the citrix-k8s-ingress-controller.yaml file using the following command: wget https://raw.githubusercontent.com/citrix/citrix-k8s-ingresscontroller/master/deployment/baremetal/citrix-k8s-ingresscontroller.yaml

2. Edit the citrix-k8s-ingress-controller.yaml file and enter the values of the environment variable using the information in Deploy NetScaler Ingress Controller as a pod.

**Note:** To update the Status.LoadBalancer.Ingress field of the Ingress resources managed by the NetScaler Ingress Controller with the allocated IP addresses, you must specify the command line argument --update-ingress-status yes when you start the NetScaler Ingress Controller. For more information, see Updating the Ingress status for the Ingress resources with the specified IP address.

- 3. On the Rancher GUI cluster page, select Clusters from Global view.
- 4. From the Clusters page, open the cluster that you want to access.
- 5. Click Launch kubectl to open a terminal for interacting with your Kubernetes cluster.
- 6. Create a file named cic.yaml in the launched terminal and then copy the content of the modified citrix-k8s-ingress-controller.yaml file to cic.yaml.
- 7. Deploy the cic.yaml file using the following command.

1 kubectl create -f cic.yaml

8. Verify if the NetScaler Ingress Controller is deployed successfully using the following command.

1 kubectl get pods --all-namespaces

# Deploy the NetScaler Ingress Controller on a PKS managed Kubernetes cluster

#### August 5, 2024

Pivotal Container Service (PKS) enables operators to provision, operate, and manage enterprisegrade Kubernetes clusters using BOSH and Pivotal Ops Manager.

The NetScaler Ingress Controller is built around the Kubernetes Ingress and it can automatically configure one or more NetScalers based on the Ingress resource configuration. You can deploy the NetScaler Ingress Controller in a PKS managed Kubernetes cluster to extend the advanced load balancing and traffic management capabilities of NetScaler to your cluster.

## Prerequisites

Before creating the Kubernetes cluster using PKS. Make sure that for all the plans available on the Pivotal Ops Manager, the following options are set:

- Enable Privileged Containers
- Disable DenyEscalatingExec

For detailed information on PKS Framework and other documentation, see Pivotal Container Service documentation.

After you have set the required options, create a Kubernetes cluster using the PKS CLI framework and set the context for the created cluster.

## **Deployment options**

You can either deploy NetScaler CPXs as pods inside the cluster or deploy a NetScaler MPX or VPX appliance outside the Kubernetes cluster.

Based on how you want to use NetScaler, there are two ways to deploy the NetScaler Ingress Controller in a Kubernetes cluster on the PKS:

- As a sidecar container alongside NetScaler CPX in the same pod: In this mode, NetScaler Ingress Controller configures the NetScaler CPX.
- As a standalone pod in the Kubernetes cluster: In this mode, you can control the NetScaler MPX or VPX appliance deployed outside the cluster.

## Deploy NetScaler Ingress Controller as a pod

Follow the instruction provided in topic: Deploy NetScaler Ingress Controller as a standalone pod in the Kubernetes cluster for NetScaler MPX or VPX appliances.

## Deploy NetScaler Ingress Controller as a sidecar with NetScaler CPX

Follow the instruction provided in topic: Deploy NetScaler Ingress Controller as a sidecar with NetScaler CPX.

## **Network Configuration**

For seamless functioning of the services deployed in the Kubernetes cluster, it is essential that Ingress NetScaler device should be able to reach the underlying overlay network over which Pods are running.

The NetScaler Ingress Controller allows you to configure network connectivity between the NetScaler device and service using Static Routing, node controller, services of type NodePort, or services of type LoadBalancer.

## **Deploy NetScaler-Integrated Canary Deployment Solution**

## June 26, 2025

Canary release is a technique to reduce the risk of introducing a new software version in production by first rolling out the change to a small subset of users. After the user validation, the application is rolled out to the larger set of users.

NetScaler provides the following options for canary deployment using the NetScaler Ingress Controller.

- Deploy canary using the Canary CRD
- Deploy canary using Ingress annotations

In a deployment using the Canary CRD, canary configuration is applied using a Kubernetes CRD. Citrix also supports a much simpler option for canary deployment using Ingress annotations.

## Deploy canary using the Canary CRD

This section provides information about how to perform Canary deployment using the Canary CRD.

NetScaler-Integrated Canary Deployment solution stitches together all components of continuous delivery (CD) and makes canary deployment easier for the application developers. This solution uses Spinnaker as the continuous delivery platform and Kayenta as the Spinnaker plug-in for canary analysis. Kayenta is an open-source canary analysis service that fetches user-configured metrics from their sources, runs statistical tests, and provides an aggregate score for the canary. The score from statistical tests and counters along with the success criteria is used to promote or fail the canary.

NetScaler comes with a rich application-centric configuration module and provides complete visibility to application traffic and health of application instances. The capabilities of NetScaler to generate accurate performance statistics can be leveraged for Canary analysis to take better decisions about the Canary deployment. In this solution, NetScaler is integrated with the Spinnaker platform and acts as a source for providing accurate metrics for analyzing Canary deployment using Kayenta.

NetScaler Metrics Exporter exports the application performance metrics to the open-source monitoring system Prometheus and you can configure Kayenta to fetch the metrics for canary deployment. Traffic distribution to the canary version can be regulated using the NetScaler policy infrastructure. If you want to divert a specific kind of traffic from production to baseline and canary, you can use match expressions to redirect traffic to baseline and canary leveraging the rich NetScaler policy infrastructure.

For example, you can divert traffic from production to canary and baseline using the match expression HTTP.REQ.URL.CONTAINS("citrix india"). The traffic which matches the expression is diverted to canary and baseline and the remaining traffic goes to production.

The components which are part of the Citrix-Integrated Canary Deployment Solution and their functionalities are explained as follows:

- GitHub: GitHub offers all the distributed version control and source code management functionalities provided by Git and has extra features.
   GitHub has many utilities available for integrating with other tools that form part of your CI/CD pipeline like Docker Hub and Spinnaker.
- Docker Hub: Docker Hub is a cloud-based repository service provided by Docker for sharing and finding Docker images. You can integrate GitHub with Docker Hub to automatically build images from the source code in GitHub and push the built image to Docker Hub.
- Spinnaker: Spinnaker is an open source, multi-cloud continuous delivery platform for releasing software changes with high velocity and reliance. You can use Spinnaker's application deployment features to construct and manage continuous delivery workflows. The key deployment management construct in Spinnaker is known as a pipeline. Pipelines in Spinnaker consist of a sequence of actions, known as stages. Spinnaker provides various stages for deploying an application, running a script, performing canary analysis, removing the deployment, and so on. You can integrate Spinnaker with many third-party tools to support many extra functionalities.
- Prometheus: Prometheus is an open-source systems monitoring and alerting toolkit. Prometheus is a monitoring system which can maintain a huge amount of data in a time series database. NetScaler Metrics exposes the performance metrics to Spinnaker through Prometheus.
- Jenkins: Jenkins is an open source automation server which helps to automate all sorts of tasks related to building, testing, and delivering or deploying software. Jenkins also supports running custom scripts as part of your deployment cycle.
- NetScaler Ingress Controller NetScaler provides an Ingress Controller for NetScaler MPX (hardware), NetScaler VPX (virtualized), and NetScaler CPX (containerized) for bare metal and cloud deployments. The NetScaler Ingress Controller is built around Kubernetes Ingress and automatically configures one or more NetScalers based on the Ingress resource configuration.

Following NetScaler software versions are required for Citrix-Integrated Canary Deployment Solution:

• NetScaler Ingress Controller build/version: quay.io/citrix/citrix-k8s-ingresscontroller:1.29.5.

- NetScaler CPX version: quay.io/citrix/citrix-k8s-cpx-ingress:13.0-83.27.
- NetScaler Metrics Exporter version: quay.io/citrix/netscaler-metrics-exporter :1.4.0.

#### Workflow of a Spinnaker pipeline for NetScaler-Integrated Canary Deployment Solution

The following diagram explains the workflow of a Spinnaker pipeline for NetScaler-Integrated Canary Deployment Solution.



The following steps explain the workflow specified in the diagram.

- 1. Developers maintain the source code in GitHub, make changes whenever required, and commit the changes to GitHub.
- 2. A webhook is configured in GitHub to listen for the source code changes. Whenever the source code is checked in to GitHub, the webhook is triggered and informs Docker Hub to build the image with the new source code. Once the docker image is created, a separate webhook configured in Docker Hub triggers a Spinnaker pipeline.
- 3. Once the Spinnaker pipeline is triggered, canary and baseline versions of the image are deployed.
- 4. Once the canary and baseline versions are deployed, some percentage of traffic from production is diverted to the canary and baseline versions. NetScaler collects the performance statistics and exports the statistics to Prometheus with the help of NetScaler Metrics Exporter. Prometheus feeds these statistics to Kayenta for canary analysis.
- 5. Kayenta performs a canary analysis based on the performance statistics and generates a score. Based on the score, the canary deployment is termed as success or failure and the image is

rolled out or rolled back.

## Deploy the NetScaler-Integrated Canary Deployment Solution in Google Cloud Platform

This section contains information on setting up Spinnaker, how to create a Spinnaker pipeline, and a sample canary deployment.

**Deploy Spinnaker in Google Cloud Platform** This topic contains information about deploying Spinnaker and how to integrate plug-ins with Spinnaker for canary deployment on Google Cloud Platform(GCP).

Perform the following steps to deploy Spinnaker and integrate plug-ins in GCP.

1. Set up the environment and create a GKE cluster using the following commands.

1 export GOOGLE\_CLOUD\_PROJECT=[PROJECT\_ID]
2 gcloud config set project \$GOOGLE\_CLOUD\_PROJECT
3 gcloud config set compute/zone us-central1-f
4 gcloud services enable container.googleapis.com
5 gcloud beta container clusters create kayenta-tutorial
6 --machine-type=n1-standard-2 --enable-stackdriver-kubernetes

2. Install the plug-in for integrating Prometheus with Stackdriver using the following command.



- 3. Deploy Spinnaker in the GKE cluster using the following steps.
  - a) Download the quick-install.yml file for Spinnaker from Spinnaker website.
  - b) Update the quick-install.yml file to integrate different components starting with Docker Hub. To integrate Spinnaker with Docker Hub, update the values of address, user name, password, email, and repository under ConfigMap in quick-install.yml file.

```
1 dockerRegistry:
2 enabled: true
3 accounts:
4 - name: my-docker-registry
```

5		<pre>requiredGroupMembership: []</pre>
6		providerVersion: V1
7		permissions: {
8	}	
9		
10		<pre>address: https://index.docker.io</pre>
11		username: <username></username>
12		password: <password></password>
13		email: <mail-id></mail-id>
14		cacheIntervalSeconds: 30
15		<pre>clientTimeoutMillis: 60000</pre>
16		cacheThreads: 1
17		paginateSize: 100
18		sortTagsByDate: <b>false</b>
19		trackDigests: <b>false</b>
20		insecureRegistry: <b>false</b>
21		repositories:- <repository-name></repository-name>
22		<pre>primaryAccount: my-docker-registry</pre>

c) (Optional) Perform the following steps to set up Jenkins.

#### Note:

If Jenkins is installed in one of the nodes of Kubernetes, you must update the firewall rules for that node for public access.

d) Update the following values in the quick-install.yml file for integrating Jenkins with Spinnaker.

```
1 data:igor.yml: |
2 enabled: true
3 skipLifeCycleManagement: false
4 ci:jenkins:
5 enabled: true
6 masters:
7 - name: master
8 address: <endpoint>
```

```
9 username: <username>
10 password: <password>
```

- e) To set up Prometheus and Grafana, see the Prometheus and Grafana Integration section in NetScaler Metrics Exporter and perform the steps.
- f) To integrate Prometheus with Spinnaker, update the following values in the quickinstall.yml file.

```
1
     data:
2
     config:
      deploymentConfigurations:
3
4
      canary:
5
      enabled: true
6 serviceIntegrations:
7
        - name: prometheus
8
   enabled: true
9 accounts:
10 - name: my-prometheus
11 endpoint:
        baseUrl: prometheus-endpoint
12
13
       supportedTypes:
14 - METRICS_STORE
15
   data:
16
    config:
      deploymentConfigurations:
17
      metricStores:
18
     prometheus:
19
20
        enabled: true
21
       add_source_metalabels: true
22
        stackdriver:
23
   enabled: true
24
      period: 30
       enabled: true
```

g) To integrate Slack for notification with Spinnaker, update the following values in the quick-install.yml file.

```
1 data:
2 config: |
3 deploymentConfigurations:
4 notifications:
5 slack:
6 enabled: true
7 botName: <BotName>
8 token: <token>
```

h) Once all the required components are integrated, deploy Spinnaker by performing the following step.

```
1 kubectl apply -f quick-install.yaml
```

i) Verify the progress of the deployment using the following command. Once the deployment is complete, this command outputs all the pods as Ready x/x.

```
1 watch kubectl -n spinnaker get pods
```

- 4. Once you deploy Spinnaker, you can test the deployment using the following steps:
  - a) Enable Spinnaker access by forwarding a local port to the deck component of Spinnaker using the following command:

```
1 DECK_POD=$(kubectl -n spinnaker get pods -l \
2 cluster=spin-deck,app=spin \
3 -o=jsonpath='{
4 .items[0].metadata.name }
5 ')
6 kubectl -n spinnaker port-forward $DECK_POD 8080:9000 >/dev/
null &
```

b) To access Spinnaker, in the Cloud Shell, click the **Web Preview icon** and select **Preview on port 8080**.

#### Note:

You can access Spinnaker securely or via HTTP. To expose Spinnaker securely, use the spin-ingress-ssl.yaml file to deploy the Ingress.

Once the Spinnaker application is publicly exposed, you can use the domain assigned for Spinnaker or the IP address of the Ingress to access it.

**Create a Spinnaker pipeline and configure automated canary deployment** Once you deploy Spinnaker, create a Spinnaker pipeline for an application and configure the automated canary deployment.

- 1. Create an application in Spinnaker.
- 2. Create a Spinnaker pipeline. You can edit the pipeline as a JSON file using the sample file provided in Sample JSON files.
- 3. Create an automated canary configuration in Spinnaker for automated canary analysis. You can use the configuration provided in the JSON file as a sample for automated canary configuration Sample JSON files.

**Deploy a sample application for canary** This example shows how to run the canary deployment of a sample application using NetScaler-Integrated Canary Deployment Solution. In this example, NetScaler CPX, MPX, or VPX is deployed as an Ingress device for a GKE cluster. NetScaler generates the performance metrics required for canary analysis.

As a prerequisite, you must complete the following step before deploying the sample application.

• Install Spinnaker and the required plug-ins in Google cloud platform using Deploy Spinnaker in Google Cloud Platform.

**Deploy the sample application** Perform the following steps to deploy a sample application as a canary release.

1. Create the necessary RBAC rules for NetScaler by deploying the rbac.yaml file.

```
1 kubectl apply -f rbac.yaml
```

2. You can either deploy the NetScaler Ingress Controller as a sidecar with NetScaler CPX or as a standalone pod which controls NetScaler VPX or MPX.

Use the cpx-with-cic-sidecar.yml file to deploy the NetScaler Ingress Controller as a sidecar with NetScaler CPX. It also deploys NetScaler Metrics Exporter on the same pod.

1 kubectl apply -f cpx-with-cic-sidecar.yml

To deploy the NetScaler Ingress Controller as a stand-alone pod for NetScaler VPX or MPX use the cic-vpx.yaml file. In this deployment, you should use the exporter.yaml file to deploy NetScaler Metrics Exporter.

```
1 kubectl apply -f cic-vpx.yaml
2 kubectl apply -f exporter.yaml
```

#### Note:

Depending on how you are deploying the NetScaler Ingress Controller, you must edit the YAML file for NetScaler Ingress Controller deployment and modify values for the environmental variables as provided in **deploying NetScaler Ingress Controller**.

3. Deploy the Ingress for securely exposing Spinnaker using the spin-ingress-ssl.yaml file.

1 kubectl apply -f spin-ingress-ssl.yaml

Note:

For more information on creating a TLS certificate for Ingress, see TLS certificates in NetScaler Ingress Controller.

- 4. Once Spinnaker is exposed using NetScaler, access Spinnaker and perform the steps in Create a Spinnaker pipeline and configure automated canary deployment if the steps are not already done.
- 5. Deploy the production version of the application using the production.yaml file.

1 kubectl apply -f production.yaml

6. Create the Ingress resource rule to expose traffic from outside the cluster to services inside the cluster using the ingress.yaml file.

1 kubectl apply -f ingress.yaml

7. Create a Kubernetes service for the application that needs canary deployment using the service.yaml file.

1 kubectl apply -f service.yaml

8. Deploy the canary CRD that defines the canary configuration using the canary-crd-class.yaml file.

```
1 kubectl apply -f canary-crd-class.yaml
```

Note:

Once you create the CRD, wait for 10 seconds before you apply the CRD object.

9. Create a CRD object canary-crd-object.yaml based on the canary CRD for customizing the canary configuration.

1 kubectl apply -f canary-crd-object.yaml

The following table explains the fields in the canary CRD object.

Field	Description
serviceNames	List of services on which this CRD has to be applied
deployment	Specifies the deployment strategy as Kayenta.
percentage	Specifies the percentage of traffic to be diverted from production to baseline and canary.
matchExpression (optional)	Any NetScaler supported policy that can be used to define the subset of users to be directed to canary and baseline versions. If x percentage of traffic is configured, then from within subset of users which matches the matchExpression only x percentage of users are diverted to baseline and canary. Remaining users are diverted to production.
Spinnaker	Specifies the Spinnaker pipeline configurations you want to apply for your services.

Field	Description
domain	IP address or domain name of the Spinnaker gate.
port	Port number of the Spinnaker gate.
applicationName	The name of the application in Spinnaker.
pipelineName	The name of the pipeline under the Spinnaker application.
serviceName	Specifies the name of the service to which you want to apply the Spinnaker configuration.

#### 10. Deploy canary and baseline versions of the application.

#### Note:

If you are fully automating the canary deployment, deploy canary and baseline versions using the **Deploy** (Manifest) stage in Spinnaker pipeline and there is no need to perform this step.

For manually deploying canary and baseline versions, use canary.yaml and baseline.yaml files.

```
1 kubectl apply -f canary.yaml
2 kubectl apply -f baseline.yaml
```

#### Troubleshooting

For troubleshooting the deployment, perform the following steps.

- 1. Check the pod logs for the respective components like Spinnaker, Prometheus, Kayenta, NetScaler CPX, NetScaler Metrics Exporter, NetScaler Ingress Controller.
- 2. Check the pod logs of the NetScaler Ingress Controller for any configuration-related errors while configuring the NetScaler proxy.
- 3. Search for the exception/Exception keyword in the NetScaler Ingress Controller pod logs to narrow down the issues.
- 4. Check for the logs preceding the search. Check for the configuration that failed and caused the issue.
- 5. Check for the reason of failures during configuration.
- 6. If the failure happened because of incorrect configuration, correct the configuration.

## Sample JSON files

1 {

This topic contains sample JSON files for Spinnaker pipeline configuration and automated canary configuration. These files can be used as a reference while creating Spinnaker pipeline and automated canary configuration.

A sample JSON file for Spinnaker pipeline configuration\*\*

```
2
     "appConfig": {
3
4
     }
5
     "description": "This pipeline deploys a canary version of the
6
         application, and a baseline (identical to production) version.\nIt
          compares them, and if the canary is OK, it triggers the
         production deployment pipeline.",
7
     "executionEngine": "v2",
8
     "expectedArtifacts": [
9
       {
10
11
         "defaultArtifact": {
12
13
           "kind": "custom"
14
          }
15
    ,
         "id": "ac842617-988f-48dc-a7a4-7f020d93cc42",
16
         "matchArtifact": {
17
18
19
           "kind": "docker",
           "name": "index.docker.io/sample/demo",
           "type": "docker/image"
21
          }
23
    ,
24
         "useDefaultArtifact": false,
25
         "usePriorExecution": false
26
        }
27
28
     ],
     "keepWaitingPipelines": false,
29
     "lastModifiedBy": "anonymous",
     "limitConcurrent": true,
31
32
     "parallel": true,
     "parameterConfig": [],
33
34
     "stages": [
       {
         "account": "my-kubernetes-account",
37
         "cloudProvider": "kubernetes",
38
         "kinds": [
39
40
           "Deployment",
           "ConfigMap"
41
         ],
42
```

```
"labelSelectors": {
43
44
            "selectors": [
45
46
              {
47
                "key": "version",
48
                "kind": "EQUALS",
49
                "values": [
50
51
                   "canary"
52
                ]
53
               }
54
55
            ]
           }
56
57
     ,
          "location": "default",
58
59
          "name": "Delete Canary",
          "options": {
60
61
62
           "cascading": true
63
           }
64
     ,
          "refId": "12",
65
66
          "requisiteStageRefIds": [
            "19",
67
            "26"
68
69
          ],
70
          "type": "deleteManifest"
71
         }
72
     ,
73
        {
74
75
          "account": "my-kubernetes-account",
          "cloudProvider": "kubernetes",
76
          "kinds": [
77
            "Deployment"
78
          ],
"labelSelectors": {
79
80
81
82
            "selectors": [
83
              {
84
85
                "key": "version",
                "kind": "EQUALS",
86
                "values": [
87
                  "baseline"
88
89
                ]
90
               }
91
92
            ]
93
           }
94
          "location": "default",
95
```

```
"name": "Delete Baseline",
96
97
           "options": {
98
99
            "cascading": true
100
            }
101
     ,
           "refId": "13",
102
103
           "requisiteStageRefIds": [
            "19",
104
            "26"
          ],
106
107
          "type": "deleteManifest"
108
         }
109
    ,
        {
110
111
           "name": "Successful deployment",
112
          "preconditions": [],
113
           "refId": "14",
114
           "requisiteStageRefIds": [
115
            "12",
116
            "13"
117
          ],
118
          "type": "checkPreconditions"
119
120
         }
121
     ,
122
        {
123
124
           "application": "sampleapplicaion",
           "expectedArtifacts": [
125
126
            {
127
               "defaultArtifact": {
128
129
130
                "kind": "custom"
131
                }
132
     ,
               "id": "9185c756-c6cd-49bc-beee-e3f7118f3412",
133
134
               "matchArtifact": {
135
                 "kind": "docker",
136
                 "name": "index.docker.io/sample/demo",
137
138
                 "type": "docker/image"
139
                }
140
     ,
               "useDefaultArtifact": false,
141
               "usePriorExecution": false
142
143
              }
144
145
          ],
           "failPipeline": true,
146
           "name": "Deploy to Production",
147
148
           "pipeline": "7048e5ac-2464-4557-a05a-bec8bdf868fc",
```

```
"refId": "19",
149
150
           "requisiteStageRefIds": [
             "25"
151
152
           ],
           "stageEnabled": {
153
154
             "expression": "\"${
155
156
      #stage('Canary Analysis')['status'].toString() == 'SUCCEEDED' }
157
     \"",
             "type": "expression"
158
159
            }
160
     ,
161
           "type": "pipeline",
           "waitForCompletion": true
          }
163
164
      ,
165
         {
166
167
           "account": "my-kubernetes-account",
           "cloudProvider": "kubernetes",
168
169
           "manifestArtifactAccount": "embedded-artifact",
           "manifests": [
170
171
             {
172
               "apiVersion": "apps/v1",
173
               "kind": "Deployment",
174
175
               "metadata": {
176
                 "labels": {
177
178
                    "name": "sampleapplicaion-prod",
179
180
                    "version": "baseline"
181
                   }
182
      ,
183
                 "name": "sampleapplicaion-baseline-deployment",
                 "namespace": "default"
184
                }
185
186
      ,
187
               "spec": {
188
                 "replicas": 4,
189
190
                 "strategy": {
191
                    "rollingUpdate": {
193
                      "maxSurge": 10,
194
                      "maxUnavailable": 10
195
196
                     }
197
      ,
198
                    "type": "RollingUpdate"
                   }
200
      ,
                  "template": {
201
```

```
202
                    "metadata": {
204
                      "labels": {
206
207
                        "name": "sampleapplicaion-prod"
                       }
208
209
210
                     }
      ,
                    "spec": {
213
214
                      "containers": [
215
                         {
                           "image": "index.docker.io/sample/demo:v1",
217
218
                           "imagePullPolicy": "Always",
                           "name": "sampleapplicaion-prod",
219
                           "ports": [
221
                             {
222
223
                               "containerPort": 8080,
                               "name": "port-8080"
224
225
                              }
226
227
                           ]
228
                          }
229
230
                      ]
                     }
232
233
                   }
234
                 }
              }
237
238
239
           ],
240
           "moniker": {
241
             "app": "sampleapplicaion"
242
243
            }
244
     ,
           "name": "Deploy Baseline",
245
           "refId": "20",
246
           "relationships": {
247
248
249
             "loadBalancers": [],
             "securityGroups": []
250
251
            }
252
      ,
253
           "requisiteStageRefIds": [],
           "source": "text",
254
```

```
NetScaler ingress controller
```

"type": "deployManifest" 255 } 257 , 258 { 259 260 "account": "my-kubernetes-account", "cloudProvider": "kubernetes", 261 "manifestArtifactAccount": "embedded-artifact", 262 "manifests": [ 263 264 { "apiVersion": "apps/v1", 266 267 "kind": "Deployment", "metadata": { 268 269 "labels": { 270 271 "name": "sampleapplicaion-prod", 272 "version": "canary" 273 274 } 275 , "name": "sampleapplicaion-canary-deployment", 276 "namespace": "default" 277 } 278 279 , "spec": { 281 "replicas": 4, 283 "strategy": { 284 "rollingUpdate": { "maxSurge": 10, "maxUnavailable": 10 288 289 } , "type": "RollingUpdate" 291 292 } 293 , 294 "template": { 296 "metadata": { 297 298 "labels": { 299 "name": "sampleapplicaion-prod" } 301 302 } 304 , "spec": { 306 307 "containers": [

```
ł
                          "image": "index.docker.io/sample/demo",
311
                          "imagePullPolicy": "Always",
                          "name": "sampleapplicaion-prod",
312
                          "ports": [
313
314
                            {
315
                               "containerPort": 8080,
                              "name": "port-8080"
317
318
                              }
319
                          ]
                         }
322
323
                      ]
324
                     }
                  }
327
328
                }
329
              }
331
332
           ],
           "moniker": {
334
             "app": "sampleapplicaion"
336
            }
337
     ,
338
           "name": "Deploy Canary",
339
           "refId": "21",
           "relationships": {
340
341
342
             "loadBalancers": [],
             "securityGroups": []
343
344
            }
345
     ,
346
           "requiredArtifactIds": [
             "ac842617-988f-48dc-a7a4-7f020d93cc42"
347
           ],
348
349
           "requisiteStageRefIds": [],
           "source": "text",
           "type": "deployManifest"
351
          }
353
     ,
354
         {
           "analysisType": "realTime",
357
           "canaryConfig": {
358
359
             "beginCanaryAnalysisAfterMins": "2",
             "canaryAnalysisIntervalMins": "",
```

361		"canaryConfigId": "7bdb4ab4-f933-4a41-865f-6d3e9c786351",
362		"combinedCanaryResultStrategy", "LOWEST"
262		"lifetimeDuration". IDT045M"
203		
364		"metricsAccountName": "my-prometheus",
365		"scopes": [
366		{
367		
368		"controllocation". "default"
200		"Sentrol Second", "Was complemented as default 20 kgs
309		controlscope . Kos-sample apprication derautt. 80. Kos-
		sampleapplication.default.8080.svc-baseline",
370		"experimentLocation": "default",
371		"experimentScope": "k8s-sampleapplicaion.default.80.k8s-
		sampleapplicaion.default.8080.svc-canary",
372		"extendedScopeParams"
373	l	
274	ſ	
374	,	
375		"scopeName": "default"
376		}
377		
378		1.
379		"scoreThresholds": {
380		
201		
381		"marginal": "0",
382		"pass" "70"
383		}
384	,	
385		"storageAccountName": "kayenta-minio"
386		}
387		
388	,	"name", "Canary Analysis"
200		
200		
390		
391		"20",
392		"21"
393		],
394		"type": "kayentaCanary"
395		}
396		
397	,	{
200		L L L L L L L L L L L L L L L L L L L
398		Here and the set of the Here Colling
399		"continueripeline": <b>Talse</b> ,
400		"failPipeline": true,
401		"job": "NJob",
402		"master": "master",
403		"name": "Auto Cleanup: GCR Image and code revert".
404		"parameters": {
405	ι	
405	ſ	
400	,	
407		"reila": "26",
408		"requisitestageketids": [
409		"25"
410		],
411		"stageEnabled": {

```
412
413
            "type": "expression"
414
           }
415
     ,
          "type": "jenkins"
416
417
          }
418
419
      ],
420
      "triggers": [
421
       {
422
423
           "account": "my-docker-registry",
424
           "enabled": true,
           "expectedArtifactIds": [
425
            "ac842617-988f-48dc-a7a4-7f020d93cc42"
426
427
          ],
428
           "organization": "sample",
           "payloadConstraints": {
429
430
    }
431
     ,
          "registry": "index.docker.io",
432
           "repository": "sample/demo",
433
          "source": "dockerhub",
434
435
          "type": "webhook"
436
          }
437
438
      ],
439
      "updateTs": "1553144362000"
440
     }
```

#### A sample JSON file for automated canary configuration

Following is a sample JSON file for automated canary configuration.

```
1 {
2
3
     "applications": [
       "sampleapplicaion"
4
5
     ],
     "classifier": {
6
7
8
        "groupWeights": {
9
          "Group 1": 70,
          "Group 2": 30
11
12
         }
13
     ,
        "scoreThresholds": {
14
15
          "marginal": 75,
16
          "pass": 95
17
18
         }
```

```
NetScaler ingress controller
```

```
19
20
      }
21
     "configVersion": "1",
22
23
     "createdTimestamp": 1552650414234,
     "createdTimestampIso": "2019-03-15T11:46:54.234Z",
24
25
     "description": "Canary Config",
     "judge": {
26
27
28
        "judgeConfigurations": {
29
     }
     ,
31
        "name": "NetflixACAJudge-v1.0"
32
      }
33
     ''metrics": [
34
        {
          "analysisConfigurations": {
37
38
39
            "canary": {
40
              "direction": "increase"
41
42
             }
43
44
           }
45
     ,
          "groups": [
46
           "Group 1"
47
48
          ],
          "name": "Server Response Errors - 5XX",
49
50
          "query": {
51
            "customFilterTemplate": "tot_requests",
52
            "metricName": "netscaler_lb_vserver_svr_busy_err_rate",
53
            "serviceType": "prometheus",
54
            "type": "prometheus"
55
56
           }
57
          "scopeName": "default"
58
59
         }
     ,
61
        {
62
63
          "analysisConfigurations": {
64
            "canary": {
              "direction": "either",
67
              "nanStrategy": "replace"
68
             }
70
71
           }
```
```
72
73
          "groups": [
             "Group 2"
74
          ],
          "name": "Server Response Latency - TTFB",
76
77
          "query": {
78
79
             "customFilterTemplate": "ttfb",
             "metricName": "netscaler_lb_vserver_hits_total",
81
            "serviceType": "prometheus",
            "type": "prometheus"
82
           }
84
          "scopeName": "default"
85
         }
86
87
      ],
      "name": "canary-config",
89
      "templates": {
90
91
92
        "tot_requests": "lb_vserver_name = \"${
93
     scope }
94
     \""
        ,
"ttfb": "lb_vserver_name = \"${
96
     scope }
     \""
97
98
       }
99
     'updatedTimestamp': 1553098513495,
100
      "updatedTimestampIso": "2019-03-20T16:15:13.495Z"
101
102
     }
```

## Simplified canary deployment using Ingress annotations

This topic provides information about the simplified Canary deployment using Ingress annotations. While NetScaler provides multiple options to support canary deployment, this is a simpler type of Canary deployment.

Canary using Ingress annotations is a rule based canary deployment. In this approach, you need to define an additional Ingress object with specific annotations to indicate that the application request needs to be served based on the rule based canary deployment strategy. In the Citrix solution, Canary based traffic routing at the Ingress level can be achieved by defining various sets of rules as follows:

- Applying the canary rules based on weight
- Applying the canary rules based on the HTTP request header
- Applying the canary rules based on the HTTP header value

The order of precedence of the canary rules is as follows:

Canary by HTTP request header value -> canary by HTTP request header -> canary by weight

## Canary deployment based on weight

Weight based canary deployment is a widely used canary deployment approach. In this approach, you can set the weight as a range from 0 to 100 which decides the percentage of traffic to be directed to the canary version and the production version of an application.

Following is the workflow for the weight based canary deployment:

- Initially the weight can be set to zero which indicates that the traffic is not forwarded to the canary version.
- Once you decide to start canary deployment, change the weight to the required percentage to make sure the traffic is directed to canary version as well.
- Finally, when you determine that the canary version is ready to be released, change the weight to 100 to ensure that all the traffic is being directed to the canary version.

For deploying weight based canary using the NetScaler Ingress Controller, create a new Ingress with a canary annotation ingress.citrix.com/canary-weight: and specify the percentage of traffic to be directed to the canary version.

## Canary deployment based on the HTTP request header

You can configure canary deployment based on the HTTP request header which is controlled by clients. The request header notifies the Ingress to route the request to the service specified in the canary Ingress. When the request header contains the value mentioned in the Ingress annotation ingress. citrix.com/canary-by-header:, the request is routed to the service specified in the canary Ingress.

## Canary deployment based on the HTTP request header value

You can also configure canary deployment based on values of the HTTP request header which is an extension of canary by header. In this deployment, along with the ingress.citrix.com/canary -by-header: annotation, you also specify the ingress.citrix.com/canary-by-header -value: annotation. When the request header value matches with the value specified in the Ingress annotation ingress.citrix.com/canary-by-header-value: the request is routed to the service specified in the canary Ingress. You can specify multiple header values as a list of strings.

Following is a sample annotation for canary deployment based on the HTTP request header values:

ingress.citrix.com/canary-by-header-value: '["value1","value2","value3","value4"]'

## Configure canary deployment using Ingress annotations

Perform the following steps to deploy a sample application as a canary release.

- 1. Deploy the NetScaler Ingress Controller using the steps in deploy the NetScaler Ingress Controller. You can either deploy the NetScaler Ingress Controller as a sidecar with NetScaler CPX or as a standalone pod which controls NetScaler VPX or MPX.
- 2. Deploy the Guestbook application using the guestbook-deploy.yaml file.

 $1 \quad {\tt kubectl apply -f guestbook-deploy.yaml} \\$ 

3. Deploy a service to expose the Guestbook application using the guestbook-service.yaml file.

1 kubectl apply -f guestbook-service.yaml

4. Deploy the Ingress object for the Guestbook application using the guestbook-ingress.yaml file.

1 kubectl apply -f guestbook-ingress.yaml

5. Deploy a canary version of the Guestbook application using the canary-deployment.yaml file.

1 kubectl apply – f canary-deployment.yaml

6. Deploy a service to expose the canary version of the Guestbook application using the canaryservice.yaml file.

1 kubectl apply - f canary-service.yaml

7. Deploy an Ingress object with annotations for the canary version of the Guestbook application using the canary-ingress.yaml file.

```
kubectl apply - f canary-ingress.yaml
1
2
3
4
5
          apiVersion: networking.k8s.io/v1
6
          kind: Ingress
7
          metadata:
8
            annotations:
9
                ingress.citrix.com/canary-weight: "10"
10
            name: canary-by-weight
11
          spec:
12
            ingressClassName: citrix
13
            rules:
14
            - host: webapp.com
15
              http:
                paths:
16
17
                - backend:
                     service:
18
```

19	name: guestbook-canary
20	port:
21	number: 80
22	path: /
23	pathType: Prefix

Here, the annotation ingress.citrix.com/canary-weight: "10" is the annotation for the weight based canary. This annotation specifies the NetScaler Ingress Controller to configure the NetScaler in such a way that 10 percent of the total requests destined to webapp. com is sent to the guestbook-canary service. This is the service for the canary version of the Guestbook application.

For deploying the HTTP header based canary using the NetScaler Ingress Controller, replace the canary annotation ingress.citrix.com/canary-weight: with the ingress.citrix.com /canary-by-header: annotation in the canary-ingress.yaml file.

For deploying the HTTP header value based canary using the NetScaler Ingress Controller, replace the ingress.citrix.com/canary-weight: annotation with the ingress.citrix.com/canary-by-header: and ingress.citrix.com/canary-by-header-value: annotations in the canary-ingress.yaml file.

Note:

You can see the Canary example YAMLs for achieving canary based on header and canary based on header value.

# **Deploy NetScaler IPAM controller**

July 1, 2024

NetScaler provides an IPAM controller for IP address management. NetScaler IPAM controller runs in parallel to NetScaler Ingress Controller in the Kubernetes cluster. NetScaler IPAM controller allocates IP addresses to services of type LoadBalancer and ingress resources from a specified IP address range.

NetScaler IPAM controller requires NetScaler's VIP custom resource definition (CRD). The VIP CRD is used for internal communication between NetScaler Ingress Controller and NetScaler IPAM controller.

## Prerequisites

• Kubernetes cluster and a kubectl command-line tool to communicate with the cluster.

• Create a namespace called netscaler to isolate resources. Run the following command to create a namespace:

```
kubectl create namespace netscaler
```

• Install NetScaler Ingress Controller for your NetScaler VPX or NetScaler MPX using the following Helm commands.

#### Note:

Ensure to create a secret using NetScaler VPX or NetScaler MPX credentials before running the following commands.

```
1 helm repo add netscaler https://netscaler.github.io/netscaler-helm-
charts/
2
3 helm install NetScaler-ingress-controller netscaler/NetScaler-ingress-
controller --set nsIP=<NSIP of MPX/VPX>,license.accept=yes,
adcCredentialSecret=<Secret-for-ADC-credentials>,ingressClass[0]=
netscaler,serviceClass[0]=netscaler,ipam=true,crds.install=true -n
netscaler
```

For detailed information about deploying and configuring NetScaler Ingress Controller using Helm charts, see the Helm chart repository.

## **Deploy IPAM controller**

1. Add the NetScaler Helm chart repository to your local registry using the following command.

```
1 helm repo add netscaler https://netscaler.github.io/netscaler-
helm-charts/
```

2. Install NetScaler IPAM controller using the following command.

```
1 helm install netscaler-ipam-controller netscaler/netscaler-ipam-
controller --set vipRange='[{
2 "<VIP-range-key>": ["<ip-range>"] }
3 ]' -n netscaler
```

For information about all the configurable parameters while installing the IPAM controller using Helm charts, see the Helm chart repository.

#### **IP address allocations**

• For services of type LoadBalancer, a unique IP address is allocated to each service from the VIP range.

- For an ingress resource, an IP address in the specified IP range is allocated. When more ingress resources refer to the same VIP range, the IP address allocated to the first ingress resource is allocated to all the other ingress resources.
- Both services of type LoadBalancer and ingress resources can use NetScaler IPAM controller for IP address allocations at the same time. If an IP address is allocated to any one resource type, it is not available for another resource type. But, the same IP address can be allocated to multiple ingress resources.

## **Environment variables in IPAM controller**

This section provides information about the environment variables in NetScaler IPAM controller.

**VIP\_RANGE** The VIP\_RANGE environment variable allows you to define the IP address range. You can either define an IP address range or an IP address range associated with a unique name.

**IP address range** You can define the IP address range from a subnet or multiple subnets. Also, you can use the – character to define the IP address range. The IPAM controller assigns the IP address from this IP address range to the service.

**IP address range associated with a unique name** You can assign a unique name to the IP address range and define the range in the VIP\_RANGE environment variable. This way of assigning the IP address range enables you to differentiate between the IP address ranges. When you create the services of type LoadBalancer, you can use the service.citrix.com/ipam-range annotation in the service definition to specify the IP address range to use for IP address allocation.

**reuseIngressVip** The reuseIngressVip variable enables you to configure the IPAM controller to assign either the same IP address or a different IP address to each ingress resource referring to the same VIP range.

reuseIngressVip value	Description
true (Default)	Assigns the same IP address to each ingress
	resource
false	Assigns a different IP address to each ingress
	resource

## Reference

• For information about exposing services of type LoadBalancer with IP addresses assigned by the IPAM controller, see this section.

# **NetScaler API Gateway for Kubernetes**

## November 4, 2024

An API gateway acts as the single entry point for your APIs and ensures secure and reliable access to multiple APIs and microservices in your system.

NetScaler provides an enterprise grade API gateway for North-South API traffic into the Kubernetes cluster. The API gateway integrates with Kubernetes through the NetScaler Ingress Controller and the NetScaler (NetScaler MPX, VPX, or CPX) deployed as the Ingress Gateway for on-premises or cloud deployments.



The following diagram shows a dual-tier topology for the API gateway.

Using the API gateway offered by Citrix, you can perform the following functionalities:

- Enforce authentication policies
- Rate limit access to services
- Advanced content routing

- Flexible and comprehensive transformation of HTTP transactions using the rewrite and responder policies
- Enforce web application firewall policies

## How does the API gateway work

NetScaler API Gateway is built on top of the NetScaler Ingress Gateway and leverages Kubernetes API extensions such as custom resource definitions (CRDs). Using CRDs, you can automatically configure a NetScaler and API gateway in the same instance.

NetScaler provides the following CRDs for the API gateway:

- Auth CRD
- Rate limit CRD
- Content routing CRD
- Rewrite and responder CRD
- WAF CRD

## Key benefits of using the API gateway

Following are the key benefits of the API gateway offered by Citrix:

- Leverages the advanced traffic management and comprehensive security features of NetScaler
- Optimizes your deployments by consolidating multiple network functions into a single component of the NetScaler Ingress Gateway.
- Reduces the operational complexity and cost involved in deploying multiple components
- Ensures better performance for your application traffic by reducing multiple hops of TCP or TLS decryption while using separate components
- Simplifies deploy and integrate in your Kubernetes environments either by directly using YAMLs or helm charts

## **Deploying NetScaler API Gateway**

For more information on how to configure NetScaler API Gateway features using CRDs, see the following:

- Authentication
- Rate limiting
- Advanced content routing
- Rewrite and responder policies
- Web application firewall policies

# **Deploying NetScaler API Gateway using Rancher**

## December 31, 2023

NetScaler API Gateway provides a single entry point for APIs by ensuring secure and reliable access to APIs and microservices on your system. NetScaler provides an enterprise-grade API gateway for North-South API traffic for Kubernetes clusters.

NetScaler API Gateway integrates with Kubernetes through the NetScaler Ingress Controller and the NetScaler (NetScaler MPX, VPX, or CPX) deployed as the Ingress Gateway for on-premises and cloud deployments.

You can use the Rancher platform to deploy NetScaler API Gateway. Rancher provides a catalog of application templates that help you to deploy NetScaler API Gateway.

## Prerequisites

You must import the cluster, in which you want to deploy the API gateway, to the Rancher platform.

## Import the cluster to the Rancher platform

Perform the following steps to import your cluster to the Rancher platform:

- 1. Log in to the Rancher platform.
- 2. In the Clusters page, click Add Cluster.

<b>17</b> 7	Global 🗸	Clusters	Apps	Users	Settings	Security 🗸	Tools 🗸						~
Clusters													dd Cluster
Delete 💼													
State													
Activ	е	cluster-can	al					Imported v1.15.1	3	<b>0.7/8 (</b> 9%	Cores	0.1/7.5 GiB 2%	:
Active								Imported v1.11.0+d4cacc0		<b>0.7/6 (</b> 12%	Cores	2.5/22.6 GiB 11%	:

- 3. In the Add Cluster Select Cluster Type page, choose the Import an existing cluster option.
- 4. Specify the **Cluster Name**.
- 5. Specify Member Roles, Labels, and Annotations.
- 6. Click Create.

## **Deploy NetScaler API Gateway using the Rancher platform**

Perform the following steps to deploy the API gateway on the cluster using the Rancher platform:

- 1. Log in to the Rancher platform.
- 2. From the **Global** drop-down list, select the cluster that you have imported.
- 3. Select the **Apps** tab and click **Launch**.

	cluster-canal Default	Resources 🗸	Namespaces	Members	Tools 🗸		<b></b> ~
Apps						🕸 Manage Catalogs	Launch

4. From the Catalog page, choose the citrix-api-gateway template.

Cluster-canal Default Resources ~	Apps Namespaces Members Tools 🗸		<b>T</b>
Catalog		• Refresh All Ca	itegories 🗸 🤇 citrix 🗙
library			
dmyx	cimp	cimp	Citege
PARTNER	PARTNER	PARTNER	PARTNER
citrix-adc-istio-ingress-gateway	citrix-api-gateway	citrix-cpx-istio-sidecar-injector	citrix-k8s-cpx-ingress-controller
dTRXX			
PARTNER			
citrix-k8s-ingress-controller			

5. Specify the mandatory and required fields under **Configuration Options** (includes deployment settings, ADC settings, the NetScaler Ingress Controller image settings, and exporter settings).

The mandatory fields include:

- **Namespace:** Specify the namespace where you want to create the NetScaler Ingress Controller. You can also use the **Edit as YAML** option to specify the same in the YAML file.
- Accept License: Select Yes to accept the terms and conditions of the NetScaler license.
- Login File Name: Specify the name of the Kubernetes secret. The secret file is used for the NetScaler login.
- **NetScaler IP:** It is the NSIP or SNIP of the NetScaler device. For high availability, specify the SNIP as the IP address.
- 6. Click **Preview** to verify the information and click **Launch**.

# **Deploy API Gateway with GitOps**

## May 13, 2025

Custom Resource Definitions (CRDs) are the primary way of configuring API gateway policies in cloud native deployments. Operations teams create the configuration policies (routing, authentication, rewrite, Web Application Firewall (WAF), and so on) and apply them in the form of CRDs. In an API Gateway context, these policies are applied on the specific APIs and upstream hosting these APIs.

API developers document the API details in an Open API specification format for the client software developers and peer service implementation teams for using the API details. API documents contain information such as base path, path, method, authentication, and authorization.

Operation teams can use the information in an API specification document to configure the API Gateway. Git, a source control solution, is used extensively by developers and operations teams. The GitOps solution makes the collaboration and communication that take place between development and operations teams easier. GitOps helps to create a faster, more streamlined, and continuous delivery for Kubernetes without losing stability.

The API Gateway deployment with the GitOps solution enables operations teams to use the API specification document created by software developers in the API gateway configuration. This solution automates the tasks and information exchange between API development and operations teams.

## About the GitOps solution for API Gateway

The GitOps solution is constituted mainly by three entities:

- Open API specification document
- Policy template CRDs
- API Gateway deployment CRD

## **Open API Specification document**

Created by API developers or API designers, the document provides an API information. The GitOps solution uses the following details from an Open API specification document:

- Base path
- Path
- Method
- Tags
- Authentication

Authorization

The following is a sample Open API specification file with the details (in red) that are used to automatically create policies.



## **Policy template CRDs**

CRDs are the primary way of configuring an API gateway instance. The operations team creates and manages the CRD implementations. In the traditional workflow, as part of creating the policies, the operations team manually fills the target details such as upstream and API path in the CRD instances. In the GitOps solution, the API path and upstream details are derived automatically. Operations team creates the CRDs without any target details and the solution refer to such CRD instances as policy templates.

The GitOps solution supports the following policy templates:

- Rewrite policy
- Rate limit policy
- Authentication policy
- WAF

The following is a sample rewrite policy template:

Note: For information on how to create a CRD instance, see the individual CRDs.



## API Gateway deployment CRD

API Gateway deployment CRD binds the API specification document with policy templates. This CRD enables mapping of API resources with upstream services and API gateway policies related to routing and security. The API Gateway deployment CRD is maintained by the operations team with the data received from the development team.

The API Gateway deployment CRD configures the following:

- Git repository details
- Endpoint listener
- API to upstream mapping
- API to policy mapping
- Open API authentication policy references to authentication policy template mapping

Alternatively, API Gateway CRD supports non-Git sources for fetching OpenAPI Specification (OAS) documents. Currently, both HTTP and HTTPS URL sources are supported. These URLs can be password protected and basic HTTP authentication is supported. Credentials can be configured using the same fields as that of Git based OAS file sources.

The following image shows the API Gateway deployment CRD binding the API specification with policy templates using the API selectors and policy mappings.



APIs that start with the */pet* regular expression is selected with the *path regexp* pattern and APIs with */play* is selected with the *play* tag. Security definitions in the API specification document are mapped with the available authentication, authorization, and auditing configurations in the authentication CRD template.

## **Configure API Gateway CRD**

The API Gateway CRD binds the API resources defined in the Swagger specification with policies defined in the other CRDs.

## Prerequisites

Apply CRD definitions for the following CRD objects:

- Listener
- HTTP route

- Rate limit
- Rewrite
- Authentication
- WAF

The following sections provide information about the various elements in the API Gateway CRD configuration file:

## **API definition**

It provides information about the Git repository in which the Git watcher monitors for the Open API specification files.

#### API defenition: Git repository access details

Field	Description
Repository	Specifies the Git repository URL.
Branch	Specifies the Git branch name (By default, master).
oas_secret_ref	Specifies the Git access secret reference as a Kubernetes secret object name. <b>Note:</b> When
	creating a secret, keep the username and
	password as the secret field names for Git access credentials.
Files	The credentials for these OAS URLs can be
	accessed from the <code>oas_secret_ref</code> field or
	user_name and password field combinations.

#### **API proxy**

It provides information about the endpoint (VIP) configuration that is used to expose the APIs on the API Gateway front end.

## api\_proxy: VIP details

Field	Description
ip_address	Specifies the IP address of the end point (VIP).
port	Specifies the endpoint port.
protocol	Specifies the protocol (HTTP/HTTPs).
secret	Specifies the SSL certificate secret for the endpoint configuration.

# **Policy mappings**

It maps the API resources with the upstream services and policy templates. Some information in this section is collected from the developers when the operations team creating the CRD.

Section	Sub section	Field	Sub field	Description
Section Policies	Subsection	Field Name API	Sub field	Description Specifies the policy and upstream mapping. Specifies the name of the policy. It is unique in a CRD instance. A list of filters for selecting the API resources. Specifies the <i>Regexp</i> pattern for the API selection. All the APIs that match
				with this pattern are selected for applying policies from this block.

# NetScaler ingress controller

Section	Sub section	Field	Sub field	Description
		method		A list of HTTP
		method		verbs if the API
				resource verb
				matches with ANV
				in the list it is
				in the list, it is
		Tags		A list of tags to
		Tags		A list of tags to
				API. These lags
				are matched with
				tags in the API
				specification
				document. You
				can use either
				ragexp based
				path patterns or
				tags to match a
				policy.
	Upstream.			Specifies the
				upstream for the
				selected policy.
		Service		Specifies the
				back-end service
				name.
		Port		Specifies the
				back-end service
				port.
	Policy-binding			Specifies the
				policy list to be
				applied on the
				selected API.
			Type of the policy	Specifies the
			template	exact type of
				policy.
				Supported types
				are WAF, rewrite
				policy, and rate
				limit.

Section	Sub section	Field	Sub field	Description
			Name	Specifies the name of the policy template.

## AAA mappings

It maps the authentication references in the API specification document with the available policy definition sections in the authentication CRD template.

Section	Sub section	Field	Sub field	Description
ааа				Authentication,
				authorization,
				and auditing
				policy section
				mappings.
		Crd_name		Specifies the
				name of the
				authentication
				CRD template.
	Mappings			Mapping API
				specification
				security policy
				references with
				the appropriate
				sections in the
				authentication
				CRD template.
				Note: If the API
				specification
				refer to string
				matches with the
				policy section
				name in the CRD
				template, explicit
				mapping is not
				required.

Perform the following steps to deploy the API Gateway CRD:

- 1. Download the API Gateway CRD.
- 2. Deploy the API Gateway CRD using the following command:

```
1 kubectl create -f apigateway-crd.yaml`
```

The following is an example API Gateway CRD configuration:

```
apiVersion: citrix.com/v1beta1
 1
2 kind: apigatewaypolicy
3 metadata:
     name: apigatewaypolicyinstance
4
5 spec:
6
       api_definition:
7
           repository: "<repository name>"
8
           branch: "modify-test-branch"
           oas_secret_ref: "mysecret"
9
10
           files:
                - "test_gitwatcher/petstore.yaml"
11
12
                - "test_gitwatcher/playstore.yaml"
13
       api_proxy:
14
           ipaddress: "10.106.172.83"
15
           port: 80
           protocol: "http"
16
17
           secret: "listner-secret"
18
       policies:
         - name: "p1"
19
20
           selector:
               - api: "/pet.*"
21
                 method: ["GET", "POST"]
22
23
           upstream:
24
               service: "pet-service"
25
                port: 80
26
           policy_bindings:
27
               ratelimit:
28
                    name: "ratelimit-gitops-slow"
          - name: "p2"
29
30
           selector:
                - api: "/user.*"
31
32
                 method: ["GET", "POST"]
           upstream:
                service: "user-service"
34
35
                port: 80
36
           policy_bindings:
37
                ratelimit:
38
                    name: "ratelimit-gitops-slow"
39
         - name: "p3"
40
           selector:
41
                - tags: ["play"]
42
           upstream:
                service: "play-service"
43
44
                port: 80
```

```
45
           policy_bindings:
46
               ratelimit:
                   name: "ratelimit-gitops"
47
48
                rewritepolicy:
49
                   name: "prefixurl"
50
                waf:
                    name: "buffoverflow"
51
52
       aaa:
53
         - crd_name: authgitops
54
           mappings:
              - petstore_auth: jwt-auth-provider
              - api_key: introspect-auth-provider
```

## Support for web insight based analytics

Web insight based analytics is now supported with the API gateway CRD. When you use GitOps, the following web insight parameters are enabled by default:

- httpurl
- httpuseragent
- httphost
- httpmethod
- httpcontenttype

## **GSLB** overview and deployment topologies

April 22, 2025

## **Overview**

For ensuring high availability, proximity-based load balancing, and scalability, you need to deploy an application in multiple distributed Kubernetes clusters. When an application is deployed in multiple Kubernetes clusters dispersed across geographically distributed locations, a load balancing decision has to be taken to distribute traffic among application instances.

NetScaler GSLB controller configures NetScaler (GSLB device) to load balance services among geographically distributed locations. GSLB solution ensures better performance and reliability for your Kubernetes services that are exposed using ingress or service type LoadBalancer. In the GSLB topology, a GSLB device is deployed in each region; one of the GSLB devices acts as the primary ADC and others act as the secondary ADCs. The GSLB primary ADC is configured by the GSLB controller deployed in each cluster deployed across sites. This GSLB device load balances services deployed in multiple clusters across sites.

For more information about GSLB, see Global Server Load Balancing.

Note:

The NetScaler GSLB controller image is the same as that of NetScaler Ingress Controller.

## **Deployment topologies**

The components of GSLB deployment topology are described here:

- **GSLB device**: NetScaler MPX or NetScaler VPX is used as a global server load balancing (GSLB) device. A GSLB device is configured for each data center. In each GSLB device, one site is configured as a local site representing the local data center. The other sites are configured as remote sites. NetScaler MPX or NetScaler VPX used as the GSLB device can also be used as the ingress device with NetScaler Ingress Controller.
- **Ingress load balancer**: NetScaler CPX or any third-party application is deployed as the ingress load balancer in each Kubernetes cluster.
- **Ingress controller**: The ingress controller can be a NetScaler Ingress Controller or any thirdparty ingress controller.
- **GSLB controller**: Each cluster in the deployment runs a GSLB controller instance. Each GSLB controller configures the GSLB primary ADC for the applications deployed in its respective cluster. The global server load balancing (GSLB) configuration synchronization option is used to copy the GSLB configuration on the primary site to all the GSLB sites in the GSLB setup. The NetScaler on which you configure GSLB synchronization is referred as the primary site and the sites to which the configuration is copied are referred as the secondary sites.

The following deployment diagrams show sample topologies for NetScaler GSLB controller. Each sample topology contains two data centers (sites) in different regions and each data center contains a Kubernetes cluster.

Let's consider two sample deployment topologies based on the type of ingress controller used in the GSLB sites:

- NetScaler Ingress Controller in both GSLB sites
- Any third-party ingress controller in both GSLB sites

Note:

GSLB controller deployment is also supported with NetScaler Ingress Controller in one GSLB site and a third-party ingress controller in another GSLB site.

## NetScaler Ingress Controller in both GSLB sites

## Note:

NetScaler MPX or NetScaler VPX used for GSLB and ingress or service type LoadBalancer can be the same or different. In the following example, the same is used for GSLB and ingress or service type LoadBalancer.

# The following diagram explains the deployment topology for NetScaler GSLB controller in presence of NetScaler Ingress Controller.



The numbers in the following steps map to the numbers in the earlier diagram.

- 1. In each cluster, NetScaler Ingress Controller configures NetScaler either using ingress or using the service of type LoadBalancer configuration.
- 2. In each cluster, NetScaler GSLB controller configures the GSLB device in the primary site with the GSLB configuration.
- 3. GSLB configuration synchronizes automatically between the GSLB devices in different GSLB sites.
- 4. A DNS query for application hostname or FQDN is sent to the GSLB virtual server configured on NetScaler. The DNS resolution on the GSLB virtual server resolves to an IP address on any one of the clusters based on the configured global traffic policy (GTP).
- 5. Based on the DNS resolution, data traffic lands on either the Ingress front-end IP address or service type LoadBalancer IP address of one of the clusters.
- 6. The required application is accessed through the GSLB device.

## Any third-party ingress controller in both GSLB sites

The following diagram explains the deployment topology for NetScaler GSLB controller in the presence of any third-party ingress controller.



The following items explain the previous diagram:

- 1. NetScaler GSLB controller configures the GSLB primary ADC in the primary site with the GSLB configuration.
- 2. GSLB configuration synchronizes automatically between the GSLB devices in different GSLB sites.
- 3. A DNS query for application hostname or FQDN is sent to the GSLB virtual server configured on NetScaler. The DNS resolution on the GSLB virtual server resolves to an IP address on any one of the clusters based on the configured global traffic policy (GTP).
- 4. Based on the DNS resolution, data traffic lands on either the ingress front-end IP address or service type LoadBalancer front-end IP address of one of the clusters.
- 5. The required application is accessed through proxy.

## **GSLB** methods and supported deployment types

The following global load balancing methods are supported:

• Round trip time (RTT)

- Static proximity
- Round robin (RR)

The following deployment types are supported:

- Local first: In a local first deployment, when an application wants to communicate with another application, it prefers a local application in the same cluster. When the application is not available locally, the request is directed to other clusters or regions.
- Canary: Canary release is a technique to reduce the risk of introducing a new software version in production by first rolling out the change to a small subset of users. In this solution, canary deployment can be used when you want to roll out new versions of the application to selected clusters before moving it to production.
- Failover: A failover deployment is used when you want to deploy applications in an active/passive configuration when they cannot be deployed in active/active mode.
- Round trip time (RTT): In an RTT deployment, the real-time status of the network is monitored and dynamically directs the client request to the data center with the lowest RTT value.
- Static proximity: In a static proximity deployment, an IP-address based static proximity database is used to determine the proximity between the client's local DNS server and the GSLB sites. The requests are sent to the site that best matches the proximity criteria.
- Round robin: In a round robin deployment, the GSLB device continuously rotates a list of the services that are bound to it. When it receives a request, it assigns the connection to the first service in the list, and then moves that service to the bottom of the list.

Note:

Currently, IPv6 is not supported.

# CRDs for configuring NetScaler GSLB controller for applications deployed in distributed Kubernetes clusters

The following CRDs are introduced to support NetScaler configuration for performing GSLB of Kubernetes applications.

- Global traffic policy (GTP)
- Global service entry (GSE)

## GTP CRD

GTP CRD accepts the parameters for configuring GSLB on NetScaler including deployment type (canary, failover), GSLB domain, health monitor for the ingress, and service type.

## The GTP CRD spec is available here.

## Note:

GTP CRD is the same across all the clusters for a given domain.

## The following table explains the GTP CRD attributes.

Field	Description
ірТуре	Specifies the DNS record type as A or AAAA.
	Currently, only A record type is supported
serviceType	Specifies the protocol to which GSLB support is
	applied.
host	Specifies the domain for which GSLB support is
	applied.
trafficPolicy	Specifies the traffic distribution policy
	supported in a GSLB deployment.
sourceIpPersistenceId	Specifies the unique source IP persistence ID.
	This attribute enables persistence based on the
	source IP address for the inbound packets. The
	sourceIpPersistenceId attribute should
	be a multiple of 100 and should be unique.
secLbMethod	Specifies the traffic distribution policy
	supported among clusters under a group in
	local-first, canary, or failover.
destination	Specifies the Ingress or LoadBalancer service
	endpoint in each cluster. The destination name
	should match with the name of GSE.
weight	Specifies the proportion of traffic to be
	distributed across clusters. For canary
	deployment, the proportion is specified as
	percentage.
CIDR	Specifies the CIDR to be used in local-first to
	determine the scope of the locality.
primary	Specifies whether the destination is a primary
	cluster or a backup cluster in the failover
	deployment.

Field	Description
monType	Specifies the type of probe to determine the
	health of the GSLB endpoint. When the monitor
	type is HTTPS, SNI is enabled by default during
	the TLS handshake.
uri	Specifies the path to be probed for the health of
	the GSLB endpoint for HTTP and HTTPS
	protocols.
respCode	Specifies the response code expected to mark
	the GSLB endpoint as healthy for HTTP and
	HTTPS protocols.
customHeader	Specifies the custom header that you want to
	add to the GSLB-endpoint monitoring traffic.

#### **GSE CRD**

GSE CRD dictates the endpoint information (any Kubernetes object which routes traffic into the cluster) in each cluster.

The GSE CRD spec is available in the NetScaler Ingress Controller GitHub repo at: gse-crd.yaml.

The following table explains the **GSE CRD** attributes.

Field	Description
ipv4address	Local cluster ingress ipv4 address or service of
	type LoadBalancer endpoint ipv4 address
domainName	Local cluster ingress domain name or service of
	type LoadBalancer domain name
monitorPort	Listening port of local cluster ingress or Listening
	port of service of type LoadBalancer

## Notes:

- GSE CRD is different for each cluster.
- For GSE CRD auto generation with ingress, the host name must match the host name specified in the GTP CRD instance. For both ingress and service of type LoadBalancer, the GSE CRD is generated only for the first port specified.
- For a service of type LoadBalancer, the GSE CRD is auto generated if the service is re-

ferred in the GTP CRD instance and the status-loadbalancer-ip/hostname field is already populated.

## ConfigMap for configuring local site preference

When configuring GSLB devices, it's important to select the virtual IP addresses and applications deployed in the Kubernetes or OpenShift Cluster that is geographically closest to the ADNS IP address for efficient traffic routing and minimized latency.

You can use the local site preference configuration for GSLB devices to resolve to the local virtual IP address if the services are available. You can achieve this scenario by setting the priority order for GSLB services using load balancing policy commands. For more information, see Configure priority order for GSLB services using load balancing policy commands.

The priority order for services or service groups within GSLB devices enables effective network traffic management and optimization. It is typically configured when binding a service group to a GSLB virtual server. Prioritizing specific services or groups ensures that critical applications are given preference during load balancing, especially when demand is high or low-latency connections are required.

The localSiteSelection parameter is added to the Netscaler-gslb-controller Helm chart to enable local site preference. Setting this parameter to true automatically adds the configuration to the GSLB device. The parameter creates a ConfigMap for the GSLB controller, which supports on-the-fly addition, modification, and deletion of the configuration.

Note:

- For any configuration (addition or modification or deletion) for local site preference through the ConfigMap, traffic gets affected as priority order for service groups must be configured or unconfigured.
- The ConfigMap must be configured in all clusters.

## Sample local site priority configuration for the GSLB devices:

1	add lb action	<site1-act></site1-act>	-type	SELECTIONORDER -	-value 1 2 3	
2	add lb action	<site2-act></site2-act>	-type	SELECTIONORDER -	-value 2 3 1	
3	add lb action	<site3-act></site3-act>	-type	SELECTIONORDER -	-value 3 1 2	
4	add lb policy	<site1-pol></site1-pol>	-rule	"CLIENT.IP.DST.E	EQ <site1_ip>'</site1_ip>	" -action <
	site1-act>					
5	add lb policy	<site2-pol></site2-pol>	-rule	"CLIENT.IP.DST.E	EQ <site2_ip></site2_ip>	" -action <
	site2-act>					
6	add lb policy	<site3-pol></site3-pol>	-rule	"CLIENT.IP.DST.E	EQ <site3_ip></site3_ip>	" -action <
	site3-act>					
7	bind gslb vser	<sup>-</sup> ver <gslb_na< td=""><td>me&gt; -:</td><td>serviceGroupName</td><td><sg_name_1></sg_name_1></td><td>-order 1</td></gslb_na<>	me> -:	serviceGroupName	<sg_name_1></sg_name_1>	-order 1
8	bind gslb vser	<sup>-</sup> ver <gslb_na< td=""><td>me&gt; -:</td><td>serviceGroupName</td><td><sg_name_2></sg_name_2></td><td>-order 2</td></gslb_na<>	me> -:	serviceGroupName	<sg_name_2></sg_name_2>	-order 2
9	bind gslb vser	<sup>-</sup> ver <gslb_na< td=""><td>me&gt; -:</td><td>serviceGroupName</td><td><sg_name_3></sg_name_3></td><td>-order 3</td></gslb_na<>	me> -:	serviceGroupName	<sg_name_3></sg_name_3>	-order 3

```
10 bind gslb vserver <gslb_name> -policyName <site1-pol> -priority 1
11 bind gslb vserver <gslb_name> -policyName <site2-pol> -priority 2
12 bind gslb vserver <gslb_name> -policyName <site3-pol> -priority 3
```

For configuring ConfigMap using the Helm chart, see https://github.com/netscaler/netscaler-helmcharts/tree/master/netscaler-gslb-controller.

For configuring ConfigMap using the OpenShift operator, see https://docs.netscaler.com/en-us/net scaler-k8s-ingress-controller/deploy/gslb-openshift-operator.

## **Deploy NetScaler GSLB controller**

#### February 5, 2025

The following steps describe how to deploy a GSLB controller in a cluster.

Note:

Repeat the steps 1 through 5 to deploy a GSLB controller in other clusters.

1. Create the secrets required for the GSLB controller to connect to GSLB devices and push the configuration from the GSLB controller.

```
1 kubectl create secret generic secret-1 --from-literal=username=<
    username for gslb device1> --from-literal=password=<password
    for gslb device1> --from-literal=sitesyncpassword=<password1
    for secure site-to-site communication>
```

```
1 kubectl create secret generic secret-2 --from-literal=username=<
    username for gslb device2> --from-literal=password=<password
    for gslb device2> --from-literal=sitesyncpassword=<password2
    for secure site-to-site communication>
```

#### Notes:

- These secrets are provided as parameters while installing GSLB controller using helm install command for the respective sites. The username and password in the command specifies the credentials of a NetScaler GSLB device user. For information about creating a system user account on NetScaler, see Create a system user account for NetScaler Ingress Controller in NetScaler.
- Starting from GSLB controller version 2.3.15, an additional key sitesyncpassword is supported when creating the above secrets to enhance the security of communication between the GSLB sites. If the sitesyncpassword key is not provided, the default password key is used for communication between the GSLB sites.

2. Add the NetScaler Helm chart repository to your local Helm registry using the following command:

1 helm repo add netscaler https://netscaler.github.io/netscaler-helm -charts/

If the NetScaler Helm chart repository is already added to your local registry, use the following command to update the repository:

helm repo update netscaler

3. Install the GSLB controller on a cluster using the Helm chart by running the following command.

#### Note:

For information about installing GSLB controller using NetScaler Operator, see Deploy NetScaler GSLB Controller in OpenShift using NetScaler Operator.

1 helm install gslb-release netscaler/netscaler-gslb-controller -f
values.yaml

#### Note:

The chart installs the recommended RBAC roles and role bindings by default.

Example values.yaml file:

```
1 license:
2
    accept: yes
3 localRegion: "east"
4 localCluster: "cluster1"
5 openshift: false # set to true for OpenShift deployments
6 entityPrefix: "k8s"
7
8 sitedata:
9 - siteName: "site1"
10 siteIp: "x.x.x.x"
   siteMask: "y.y.y.y"
11
   sitePublicIp: "z.z.z.z"
12
   secretName: "secret-1"
13
   siteRegion: "east"
14
15 - siteName: "site2"
   siteIp: "x.x.x.x"
16
17
   siteMask: "y.y.y.y"
18 sitePublicIp: "z.z.z."
19
     secretName: "secret-2"
     siteRegion: "west"
```

Specify the following parameters in the values.yml file.

Parameter	Description
LocalRegion	Local region where the GSLB controller is deployed.
LocalCluster	The name of the cluster in which the GSLB controller is deployed. This value is unique for
entityPrefix	The prefix for the resources on NetScaler VPX/MPX. <b>Note:</b> entityPrefix must be same
sitedata[0].siteName	The name of the first GSLB site configured in the GSLB device.
sitedata[0].siteIp	IP address for the first GSLB site. Add the IP address of the NetScaler in site1 as <i>sitedata[0].siteIp</i> .
sitedata[0].siteMask	The netmask of the first GSLB site IP address.
sitedata[0].sitePublicIp	The public IP address of the first GSLB Site.
sitedata[0].secretName	The name of the secret that contains the login credentials of the first GSLB site.
sitedata[0].siteRegion	The region of the first site.
sitedata[1].siteName	The name of the second GSLB site configured in the GSLB device.
sitedata[1].siteIp	IP address for the second GSLB site. Add the IP address of the NetScaler in site2 as <i>sitedata[0].siteIp</i>
sitedata[1].siteMask	The netmask of the second GSLB site IP address.
sitedata[1].sitePublicIp	The public IP address of the second GSLB site.
sitedata[1].secretName	The secret containing the login credentials of the second site.
sitedata[1].siteRegion	The region of the second site.

#### Note:

The order of the GSLB site information should be the same in all the clusters. The first site in the order is considered as the primary site for pushing the configuration. When that primary site goes down, the next site in the list becomes the new primary. For example, if the order of sites is site1 followed by site2 in cluster1, all other clusters should have the same order. 4. Verify the installation using the following command: kubectl get pods -l app=gslbrelease-netscaler-gslb-controller.

After the successful installation of the GSLB controller on each cluster, the ADNS service will be configured and the management access will be enabled on both the GSLB devices.

## Synchronize GSLB configuration

Run the following commands in the same order on the primary NetScaler GSLB device to enable automatic synchronization of the GSLB configuration between the primary and secondary GSLB devices.

```
set gslb parameter -automaticconfigsync enable
sync gslb config -debug
```

## Examples for global traffic policy (GTP) deployments

The GTP configuration should be the same across all the clusters.

In the following examples, an application app1 is deployed in the default namespace of the cluster1 of the east region and default namespace of cluster2 of the west region.

Note:

The destination information in the GTP yaml should be in the format servicename. namespace.region.cluster, where the servicename and namespace corresponds to the Kubernetes object of kind Service and its namespace.

You can specify the load balancing method for canary and failover deployments.

Starting from NSIC release 2.1.4, you can configure multiple monitors for services in the GSLB setup. For more information, see Multi-monitor support for GSLB.

## Example 1: Round robin deployment

Use this deployment to distribute the traffic evenly across the clusters. The following example configures a GTP for round robin deployment.

Use the weight field to direct more client requests to a specific cluster within a group. Specify a custom header that you want to add to the GSLB-endpoint monitoring traffic by adding the customHeader argument under the monitor parameter.

```
1 kubectl apply -f - <<EOF</pre>
```

```
2 apiVersion: "citrix.com/v1beta1"
```

```
3 kind: globaltrafficpolicy
4 metadata:
5 name: gtp1
6 namespace: default
7 spec:
  serviceType: 'HTTP'
8
   hosts:
9
10
   - host: 'app1.com'
11
     policy:
       trafficPolicy: 'ROUNDROBIN'
13
       targets:
       - destination: 'app1.default.east.cluster1'
14
         weight: 2
15
        - destination: 'app1.default.west.cluster2'
16
17
          weight: 5
18
        monitor:
19
        - monType: http
         uri: ''
20
          customHeader: "Host: <custom hostname>\r\n x-b3-traceid:
             afc38bae00096a96\r\n\r\n"
    respCode: 200
23 EOF
```

## **Example 2: Failover deployment**

Use this policy to configure the application in active-passive mode. In a failover deployment, the application is deployed in multiple clusters. Failover is achieved between the application instances (target destinations) in different clusters based on the weight assigned to those target destinations in the GTP policy.

The following example shows a sample GTP configuration for failover. Using the primary field, you can specify which target destination is active and which target destination is passive. The default value for the primary field is True indicating that the target destination is active. Bind a monitor to the endpoints in each cluster to probe their health.

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: "citrix.com/v1beta1"
3 kind: globaltrafficpolicy
4 metadata:
   name: gtp1
5
6 namespace: default
7 spec:
   serviceType: 'HTTP'
8
9
   hosts:
    - host: 'app1.com'
10
11
       policy:
        trafficPolicy: 'FAILOVER'
12
        secLbMethod: 'ROUNDROBIN'
13
14
        targets:
         - destination: 'app1.default.east.cluster1'
15
```

```
16
         weight: 1
         - destination: 'app1.default.west.cluster2'
17
18
           primary: false
19
           weight: 1
20
         monitor:
21
         - monType: http
           uri: ''
22
23
           respCode: 200
24 EOF
```

#### Example 3: RTT deployment

Use this policy to monitor the real-time status of the network and dynamically direct the client request to the target destination with the lowest RTT value.

Following is a sample global traffic policy for round trip time deployment.

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: "citrix.com/v1beta1"
3 kind: globaltrafficpolicy
4 metadata:
5
    name: gtp1
    namespace: default
6
7 spec:
8
   serviceType: 'HTTP'
9 hosts:
    - host: 'app1.com'
11
       policy:
12
        trafficPolicy: 'RTT'
13
        targets:
         - destination: 'app1.default.east.cluster1'
14
         - destination: 'app1.default.west.cluster2'
15
16
       monitor:
17
         - monType: tcp
18 EOF
```

## **Example 4: Canary deployment**

Use the canary deployment when you want to roll out new versions of the application to selected clusters before moving it to production.

This section describes a sample global traffic policy with Canary deployment, where a new version of an application needs to be rolled out before deploying in production.

In this example, an application is deployed in a cluster cluster2 in the west region. A new version of the application is getting deployed in cluster1 of the east region. Using the weight field you can specify how much traffic is redirected to each cluster. Here, weight is specified as 40 percent. Hence, only 40 percent of the traffic is directed to the new version. If the weight field is not

mentioned for a destination, it is considered as part of the production which takes the majority traffic. When the newer version of the application is found as stable, the new version can be rolled out to other clusters as well.

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: "citrix.com/v1beta1"
3 kind: globaltrafficpolicy
4 metadata:
5 name: gtp1
6 namespace: default
7 spec:
  serviceType: 'HTTP'
8
9
    hosts:
10 - host: 'app1.com'
11
     policy:
      trafficPolicy: 'CANARY'
12
13
       secLbMethod: 'ROUNDROBIN'
14
       targets:
15
       - destination: 'app1.default.east.cluster1'
          weight: 40
16
      - destination: 'app1.default.west.cluster2'
17
18
       monitor:
19
        - monType: http
         uri: ''
20
21
         respCode: 200
22 EOF
```

## Example 5: Static proximity

Use this policy to select the service that best matches the proximity criteria.

Following GTP is an example for static proximity deployment.

Note:

For static proximity, you need to apply the location database manually on all the GSLB devices:

```
add locationfile /var/netscaler/inbuilt_db/Citrix_Netscaler_InBuilt_GeoIP_DB
.
```

```
1 kubectl apply -f - <<EOF
2 apiVersion: "citrix.com/v1beta1"
3 kind: globaltrafficpolicy
4 metadata:
5 name: gtp1
6 namespace: default
7 spec:
8 serviceType: 'HTTP'
9 hosts:
10 - host: 'app1.com'
11 policy:
```

```
trafficPolicy: 'STATICPROXIMITY'
12
13
       targets:
        - destination: 'app1.default.east.cluster1'
14
       - destination: 'app1.default.west.cluster2'
15
16
       monitor:
17
        - monType: http
         uri: ''
18
         respCode: 200
19
20 EOF
```

## Example 6: source IP persistence

The following traffic policy is an example to enable source IP persistence by providing the parameter sourceIpPersistenceId.

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: "citrix.com/v1beta1"
3 kind: globaltrafficpolicy
4 metadata:
   name: gtp1
5
6 namespace: default
7 spec:
8 serviceType: 'HTTP'
9 hosts:
10 - host: 'app1.com'
     policy:
11
     trafficPolicy: 'ROUNDROBIN'
13
       sourceIpPersistenceId: 300
       targets:
14
       - destination: 'app1.default.east.cluster1'
15
16
         weight: 2
17
       - destination: 'app1.default.west.cluster2'
18
         weight: 5
19
       monitor:
20
        - monType: tcp
         uri: ''
21
         respCode: 200
23 EOF
```

## Example for global service entry (GSE)

GSE configuration is applied in a specific cluster based on the cluster endpoint information. The GSE name must be the same as the target destination name in the global traffic policy.

Note:

Creating GSE is optional. If GSE is not created, NetScaler Ingress Controller looks for matching ingress with host matching <svcname>.<namespace>.<region>.<cluster> format.

For a global traffic policy mentioned in the earlier section, here is the global service entry for cluster1. In this example, the global service entry name app1.default.east.cluster1 is one of the target destination names in the global traffic policy created.

```
1 kubectl apply -f - <<EOF
2 apiVersion: "citrix.com/v1beta1"
3 kind: globalserviceentry
4 metadata:
5 name: 'app1.default.east.cluster1'
6 namespace: default
7 spec:
8 endpoint:
9 ipv4address: 10.102.217.70
10 monitorPort: 33036
11 EOF
```

## **Multi-monitor support for GSLB**

In a GSLB setup, you can configure multiple monitors to monitor services of the same host. The monitors can be of different types, depending on the request protocol used to check the health of the services. For example, HTTP, HTTPS, and TCP.

In addition to configuring multiple monitors, you can define the following additional parameters for a monitor:

- Destination port: The service port that the monitor uses for the health check.
- SNI: SNI status of this monitor. Possible values are True and False.
- CN: Common name (CN) to be used in the SNI request. The common name is typically the domain name of the server being monitored.

You can also define the combination of parameters for each monitor as per your requirement. In the GTP with the multi-monitor support, you can monitor each service and if any of the services is down, the site is marked as down, and the traffic is directed to the other site. The following GTP YAML example has multiple monitors:

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: "citrix.com/v1beta1"
3 kind: globaltrafficpolicy
4 metadata:
5
    name: gtp1
   namespace: default
6
7 spec:
8 serviceType: 'HTTP'
9 hosts:
     - host: 'app1.com'
10
11
       policy:
         trafficPolicy: 'STATICPROXIMITY'
12
13
         targets:
```
14	- destination: 'app1.default.east.cluster1'
15	- destination: 'app1.default.west.cluster2'
16	monitor:
17	- monType: HTTPS
18	uri: ''
19	respCode: '200,300,400'
20	destinationPort: 1000
21	- monType: HTTP
22	uri: ''
23	respCode: '200,300,400'
24	destinationPort: 3000
25	- monType: HTTPS
26	uri: ''
27	respCode: '200,300,400'
28	destinationPort: 4000
29	- monType: TCP
30	uri: ''
31	destinationPort: 5000
32	respCode: '200'
33	- monType: HTTPS
34	uri: 'test.com'
35	sni: Irue
36	respLode: '200,300,400'
37	
38	- monlype: HITPS
39	
40	sill: llue
41	
42	destinationPort: 7000
45	ctatus:
44	status.
45	1
47	J
48	FOF
10	

# NetScaler GSLB controller for single site

November 6, 2024

# Overview

For ensuring high availability, proximity-based load balancing, and scalability, you need to deploy an application in multiple Kubernetes clusters. GSLB solution ensures better performance and reliability for your Kubernetes services that are exposed using Ingress. NetScaler GSLB controller configures NetScaler (GSLB device) to load balance services among geographically distributed locations. In a single-site GSLB solution, a GSLB device in a data center is configured by the GSLB controller deployed in each Kubernetes cluster of a data center. This GSLB device load balances services deployed in multiple clusters of the data center.

The following diagram describes the deployment topology for NetScaler GSLB controller in a data center with two Kubernetes clusters and a single GSLB site.

# Note:

The NetScaler (MPX or VPX) used for GSLB and Ingress can be the same or different. In the following diagram, the same NetScaler is used for GSLB and Ingress.



The numbers in the following steps map to the numbers in the earlier diagram.

- 1. In each cluster, the NetScaler Ingress Controller configures NetScaler using Ingress.
- 2. In each cluster, the NetScaler GSLB controller configures the GSLB device with the GSLB configuration.
- 3. A DNS query for the application URL is sent to the GSLB virtual server configured on NetScaler. The DNS resolution on the GSLB virtual server resolves to an IP address on any one of the clusters based on the configured global traffic policy (GTP).

- 4. Based on the DNS resolution, data traffic lands on either the Ingress front-end IP address or the content switching virtual server IP address of one of the clusters.
- 5. The required application is accessed through the GSLB device.

# Deploy NetScaler GSLB controller

## Prerequisites

- Access to Kubernetes clusters hosted in cloud or on-premises. The Kubernetes cluster in the cloud can be a managed Kubernetes (for example: GKE, EKS, or AKS) or a custom created Kubernetes deployment. Kubernetes clusters must be properly configured and operational. Ensure the network connectivity between NetScaler and the cluster's pod network.
- For Kubernetes deployment, you need access to Kubernetes clusters running version 1.21 or later.
- For OpenShift deployment, you need access to OpenShift clusters running version 4.11 or later.
- You have installed Helm version 3.x or later. To install Helm, see here.
- NetScaler MPX/VPX is deployed in a standalone, HA, or cluster setup depending on your specific needs.
  - For instructions to deploy NetScaler MPX, see NetScaler documentation.
  - For instructions to deploy NetScaler VPX, see Deploy a NetScaler VPX instance.
- Determine the NSIP (NetScaler IP) address using which NetScaler Ingress Controller communicates with NetScaler. The IP address might be any one of the following IP addresses depending on the type of NetScaler deployment:
  - NSIP (for standalone appliances) The management IP address of a standalone NetScaler appliance. For more information, see IP Addressing in NetScaler.
  - SNIP (for appliances in High Availability mode) The subnet IP address. For more information, see IP Addressing in NetScaler.
  - CLIP (for appliances in Cluster mode) The cluster management IP (CLIP) address for a cluster NetScaler deployment. For more information, see IP addressing for a cluster.
- A user account in NetScaler VPX or NetScaler MPX. NetScaler Ingress Controller uses a system user account in NetScaler to configure NetScaler MPX or NetScaler VPX. For instructions to create the system user account on NetScaler, see Create System User Account for NetScaler Ingress Controller in NetScaler.

You can directly pass the user name and password or use Kubernetes secrets. If you want to use Kubernetes secrets, create a secret for the user name and password using the following command:

```
1 kubectl create secret generic nslogin --from-literal=username=<
    username> --from-literal=password=<password>
```

# Steps to deploy GSLB controller for a single site

In this procedure, we deploy the same HTTPS application in two clusters of a site, deploy the GSLB controller in each cluster such that the GSLB controllers configure the GSLB device (NetScaler VPX or NetScaler MPX) to load balance services deployed across the clusters.

Note:

- Repeat the following steps in both the clusters to set up a single-site GSLB solution. There are exceptions added in some steps highlighting what must be done in a particular cluster; perform the steps accordingly.
- You must deploy ingressed resources in the application namespace.
- You must deploy GTP and GSE resources in the same namespace. We recommend you to deploy these resources in the application namespace.
- 1. Create a certificate-key pair with cnn.com as the common name in the certificate signing request (CSR) in NetScaler VPX or NetScaler MPX. For information about creating a certificate-key pair, see Create a certificate.
- 2. Deploy the CNN application using the following command.

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: apps/v1
3 kind: Deployment
4 metadata:
5
   name: cnn-website
   labels:
6
   name: cnn-website
7
8
     app: cnn-website
9 spec:
10
   selector:
11
     matchLabels:
        app: cnn-website
   replicas: 2
13
14 template:
15
   metadata:
16
       labels:
17
          name: cnn-website
18
          app: cnn-website
19
      spec:
       containers:
```

```
- name: cnn-website
22
          image: quay.io/sample-apps/cnn-website:v1.0.0
23
          ports:
           - name: http-80
24
            containerPort: 80
25
26
           - name: https-443
             containerPort: 443
27
28 ---
29 apiVersion: v1
30 kind: Service
31 metadata:
32
    name: cnn-website
33
     labels:
34
      app: cnn-website
35 spec:
    type: NodePort
37
    ports:
     - name: http-80
39
      port: 80
40
      targetPort: 80
41
    - name: https-443
    port: 443
targetPort: 443
42
43
   selector:
44
45
    name: cnn-website
46 EOF
```

#### Note:

For an OpenShift deployment, run the following command oc adm policy addscc-to-user anyuid system:serviceaccount:<namespace>:default to grant a specific Security Context Constraint (SCC) to a service account. Replace < namespace> with the actual namespace where you have deployed the CNN application.

- 3. Deploy NSIC.
  - a) Add the NetScaler Helm chart repository to your local Helm registry using the following command:

1 helm repo add netscaler https://netscaler.github.io/netscalerhelm-charts/

If the NetScaler Helm chart repository is already added to your local registry, use the following command to update the repository:

1 helm repo update netscaler

b) Update values.yaml to configure NetScaler Ingress Controller as described as following.

Example values.yaml for cluster1

```
1 license:
```

```
2 accept: yes
```

- 3 adcCredentialSecret: nslogin # K8s Secret created as part of prerequisiste
- 4 nsIP: <x.x.x> # CLIP (for appliances in Cluster mode), SNIP (
   for appliances in High Availability mode) , NSIP (for
   standalone appliances)
- 5 openshift: **false** # set to **true for** OpenShift deployments
- 6 entityPrefix: cluster1 # unique for each NSIC instance.
- 7 clusterName: cluster1
- 8 ingressClass: ['nsic-vpx'] # ingress class used in the ingress resources
- 9 # serviceClass- To use service type LB, specify the service class

Example values.yaml for cluster2



c) Install NSIC using the Helm chart by running the following command.

```
1 helm install nsic netscaler/netscaler-ingress-controller -f
values.yaml
```

For information about the mandatory and optional parameters that you can configure during NSIC installation, see Configuration.

#### Note:

If an earlier version of NSIC is already deployed in the cluster, deploy the CRD specification using the following command: kubectl apply -f https:// raw.githubusercontent.com/netscaler/netscaler-helm-charts /refs/heads/master/netscaler-ingress-controller/crds/crds .yaml.

4. Deploy the following ingress resource.

Ingress resource in cluster1.

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: networking.k8s.io/v1
3 kind: Ingress
4 metadata:
5
    name: frontend-ingress
    annotations:
6
       ingress.citrix.com/frontend-ip: <NSVIP1> # vserver IP1
7
8
       ingress.citrix.com/preconfigured-certkey: '{
   "certs": [{
9
    "name": "singlesite", "type": "default" }
10
11
    ] }
12
    ' # provide the certificate key created in step 1
13 spec:
14
     tls:
15
     - {
16
     }
17
     ingressClassName: nsic-vpx
18
19
     rules:
       - host: cnn.com
20
21
         http:
22
           paths:
23
              - path: /
24
                pathType: Prefix
25
                backend:
26
                  service:
27
                    name: cnn-website
28
                    port:
29
                      number: 80
30 EOF
```

Ingress resource in cluster2.

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: networking.k8s.io/v1
3 kind: Ingress
4 metadata:
5
    name: frontend-ingress
6
     annotations:
       ingress.citrix.com/frontend-ip: <NSVIP2> # vserver IP2
7
       ingress.citrix.com/preconfigured-certkey: '{
8
9
    "certs": [{
    "name": "singlesite", "type": "default" }
10
11
    1 }
    ' # provide the name of the certificate key created in step 1
12
13 spec:
14
     tls:
15
       - {
     }
17
18
     ingressClassName: nsic-vpx
19
     rules:
    - host: cnn.com
20
```

```
http:
22
            paths:
23
              - path: /
                pathType: Prefix
24
25
                backend:
                  service:
27
                     name: cnn-website
28
                     port:
29
                       number: 80
30 EOF
```

5. Create the secrets required for the GSLB controller to connect to GSLB devices and push the configuration from the GSLB controller.

```
1 kubectl create secret generic secret--from-literal=username=<
    username for gslb device>--from-literal=password=<password for
    gslb device>
```

#### Note:

- This secret is provided as a parameter in the GSLB controller helm install command for the respective sites. Specify the credentials of a NetScaler (GSLB device) user as username and password in the command.
- In this case, because the GSLB device and the ingress device are the same, you can use the same secret that is created in the prerequisites section.
- 6. Install GSLB controller using the Helm chart by running the following command.

```
1 helm install my-release netscaler/netscaler-gslb-controller -f
values.yaml
```

#### Note:

The chart installs the recommended RBAC roles and role bindings by default.

#### Example values.yaml file:

```
license:
1
2
       accept: yes
3
     localRegion: "east"
     localCluster: "cluster1" # use cluster2 when deploying GSLB
4
        controller in cluster2
5
    entityPrefix: "gslb" # should be same for GSLB controller in
        both clusters
    nsIP: "x.x.x.x"
6
7
    openshift: false # set to true for OpenShift deployments
8
     adcCredentialSecret: <Secret-for-NetScaler-credentials>
9
     sitedata:
       - siteName: "site1"
10
11
     siteIp: "x.x.x.x"
```

12	siteMask:	
13	sitePublici	0:
14	secretName:	"secret"
15	siteRegion:	"east"

Specify the following parameters in the YAML file.

Parameter	Description
LocalRegion	Local region where the GSLB controller is deployed. This value is the same for GSLB controller deployment across all the clusters.
LocalCluster	The name of the cluster in which the GSLB controller is deployed. This value is unique for each Kubernetes cluster.
sitedata[0].siteName	The name of the GSLB site.
sitedata[0].sitelp	IP address for the GSLB site. Add the IP address of the NetScaler in site 1 as <i>sitedata[0].siteIp</i> .
sitedata[0].siteMask	The netmask of the GSLB site IP address.
sitedata[0].sitePublicIp	The site public IP address of the GSLB site.
sitedata[0].secretName	The name of the secret that contains the login credentials of the GSLB site.
sitedata[0].siteRegion	The region of the GSLB site.
NSIP	The SNIP (subnet IP address) of the GSLB device. Add the <i>sitedata[0].sitelp</i> as SNIP on NetScaler.
adcCredentialSecret	The Kubernetes secret containing the login credentials for NetScaler VPX or MPX.

After the successful installation of a GSLB controller on each cluster, GSLB site and ADNS service are configured and management access is enabled on the GSLB site IP address.

# Note:

If an earlier version of GSLB controller is already deployed in the cluster, deploy the
GTP and GSE CRD specifications using the following commands: kubectl apply
-f https://raw.githubusercontent.com/netscaler/netscalerk8s-ingress-controller/master/gslb/Manifest/gtp-crd.yaml and
kubectl apply -f https://raw.githubusercontent.com/netscaler/
netscaler-k8s-ingress-controller/master/gslb/Manifest/gse-crd
.yaml.

7. Deploy a Global traffic policy based on your requirement.

# Notes:

- Ensure that the GTP configuration is the same across all the clusters. For information on GTP CRD and allowed values, see GTP CRD.
- The destination information in the GTP YAML must be in the format servicename .namespace.region.cluster, where the service name and namespace correspond to the Kubernetes object of type Service and its namespace, respectively.

You can specify the load balancing method for canary and failover deployments.

# • Example 1: Round robin deployment

Use this deployment to distribute the traffic evenly across the clusters. The following example configures a GTP for round robin deployment.

You can use the weight field to direct more client requests to a specific cluster within a group.

```
kubectl apply -f - <<EOF</pre>
1
    apiVersion: "citrix.com/v1beta1"
2
3 kind: globaltrafficpolicy
4 metadata:
5
    name: gtp1
6 spec:
7
     serviceType: 'HTTP'
8
     hosts:
9
      - host: 'cnn.com'
10
        policy:
          trafficPolicy: 'ROUNDROBIN'
11
          targets:
          - destination: 'cnn-website.default.east.cluster1'
13
14
            weight: 2
          - destination: 'cnn-website.default.east.cluster2'
15
16
            weight: 5
17
          monitor:
18
          - monType: http
19
            uri: ''
20
            respCode: 200
21
    EOF
```

# • Example 2: Failover deployment

Use this policy to configure the application in active-passive mode. In a failover deployment, the application is deployed in multiple clusters. Failover is achieved between the instances in target destinations based on the weight assigned to those target destinations in the GTP policy.

The following example shows a sample GTP configuration for failover. Using the primary field, you can specify which target destination is active and which target destination is passive. The

default value for the primary field is True indicating that the target destination is active. Bind a monitor to the endpoints in each cluster to probe their health.

```
kubectl apply -f - <<EOF</pre>
1
     apiVersion: "citrix.com/v1beta1"
2
     kind: globaltrafficpolicy
3
4
     metadata:
5
       name: gtp1
6
    spec:
7
       serviceType: 'HTTP'
8
       hosts:
       - host: 'cnn.com'
9
10
        policy:
           trafficPolicy: 'FAILOVER'
11
12
           secLbMethod: 'ROUNDROBIN'
13
           targets:
14
           - destination: 'cnn-website.default.east.cluster1'
15
             weight: 1
            - destination: 'cnn-website.default.east.cluster2'
16
17
              primary: false
18
             weight: 1
19
           monitor:
20
           - monType: http
             uri: ''
21
22
              respCode: 200
23
     EOF
```

#### • Example 3: RTT deployment

Use this policy to monitor the real-time status of the network and dynamically direct the client request to the target destination with the lowest RTT value.

The following example configures a GTP for RTT deployment.

```
kubectl apply -f - <<EOF</pre>
1
2
    apiVersion: "citrix.com/v1beta1"
3 kind: globaltrafficpolicy
4 metadata:
5
    name: gtp1
6 spec:
7
     serviceType: 'HTTP'
8
     hosts:
9
      - host: 'cnn.com'
10
        policy:
         trafficPolicy: 'RTT'
          targets:
          - destination: 'cnn-website.default.east.cluster1'
13
14
          - destination: 'cnn-website.default.east.cluster2'
15
          monitor:
16
          - monType: http
            uri: ''
17
18
            respCode: 200
19 EOF
```

### • Example 4: Canary deployment

Use the canary deployment when you want to roll out new versions of the application to selected clusters before moving it to production.

This section describes a sample global traffic policy with Canary deployment, where you need to roll out a newer version of an application in stages before deploying it in production.

In this example, a stable version of an application is deployed in cluster2. A new version of the application is deployed in cluster1. Using the weight field, specify how much traffic is redirected to each cluster. Here, weight is specified as 40 percent. Hence, only 40 percent of the traffic is directed to the new version. If the weight field is not mentioned for a destination, it is considered as part of the production which takes the majority traffic. When the newer version of the application is stable, the new version can be rolled out to the other clusters.

```
1
    kubectl apply -f - <<EOF</pre>
   apiVersion: "citrix.com/v1beta1"
2
   kind: globaltrafficpolicy
3
4 metadata:
5
     name: gtp1
6
     namespace: default
7
   spec:
   serviceType: 'HTTP'
8
9
     hosts:
     - host: 'cnn.com'
11
        policy:
          trafficPolicy: 'CANARY'
12
          secLbMethod: 'ROUNDROBIN'
13
14
          targets:
          - destination: 'cnn-website.default.east.cluster1'
15
16
            weight: 40
          - destination: 'cnn-website.default.east.cluster2'
17
18
          monitor:
19
          - monType: http
           uri: ''
20
21
            respCode: 200
22
    EOF
```

#### • Example 5: Static proximity

Use this policy to select the service that best matches the proximity criteria. The following traffic policy is an example for static proximity deployment.

# Note:

For static proximity, you must apply the location database manually: add locationfile /var/netscaler/inbuilt\_db/Citrix\_Netscaler\_InBuilt\_GeoIP\_DB\_IPv4

.

```
kubectl apply -f - <<EOF</pre>
2
     apiVersion: "citrix.com/v1beta1"
3
     kind: globaltrafficpolicy
4
     metadata:
5
       name: gtp1
6
       namespace: default
7
     spec:
8
       serviceType: 'HTTP'
9
       hosts:
10
       - host: 'cnn.com'
11
         policy:
           trafficPolicy: 'STATICPROXIMITY'
12
13
           targets:
            - destination: 'cnn-website.default.east.cluster1'
14
15
            - destination: 'cnn-website.default.east.cluster2'
16
            monitor:
17
            - monType: http
              uri: ''
18
19
              respCode: 200
     EOF
```

### • Example 6: Source IP persistence

The following traffic policy is an example to enable source IP persistence by providing the parameter sourceIpPersistenceId.

```
1
     kubectl apply -f - <<EOF</pre>
2
     apiVersion: "citrix.com/v1beta1"
3
     kind: globaltrafficpolicy
4
     metadata:
5
       name: gtp1
6
       namespace: default
7
     spec:
8
       serviceType: 'HTTP'
9
       hosts:
        - host: 'cnn.com'
10
11
          policy:
12
            trafficPolicy: 'ROUNDROBIN'
13
            sourceIpPersistenceId: 300
14
            targets:
15
            - destination: 'cnn-website.default.east.cluster1'
16
              weight: 2
            - destination: 'cnn-website.default.east.cluster2'
17
18
             weight: 5
19
            monitor:
            - monType: http
21
              uri: ''
              respCode: 200
23
     EOF
```

8. Deploy the GSE resource.

GSE configuration is applied in a specific cluster based on the cluster endpoint information. The GSE name must be the same as the target destination name in the global traffic policy.

# Note:

Creating a GSE is optional. If GSE is not created, NetScaler Ingress Controller looks for matching ingress with host matching <svcname>.<namespace>.<region>.< cluster> format.

For a global traffic policy mentioned in the earlier examples, the following YAML is the global service entry for cluster1. In this example, the global service entry name cnn.default.east .cluster1 is one of the target destination names in the global traffic policy.

```
1 kubectl apply -f - <<EOF
2 apiVersion: "citrix.com/v1beta1"
3 kind: globalserviceentry
4 metadata:
5 name: 'cnn-website.default.east.cluster1'
6 spec:
7 endpoint:
8 ipv4address: NSVIP1 # vserver IP1
9 monitorPort: 80
10 EOF
```

For a global traffic policy mentioned in the earlier examples, the following YAML is the global service entry for cluster2. In this example, the global service entry name cnn.default.east.cluster2 is one of the target destination names in the global traffic policy.

```
1 kubectl apply -f - <<EOF
2 apiVersion: "citrix.com/v1beta1"
3 kind: globalserviceentry
4 metadata:
5 name: 'cnn-website.default.east.cluster2'
6 spec:
7 endpoint:
8 ipv4address: NSVIP2 # vserver IP2
9 monitorPort: 80
10 EOF
```

9. To send a DNS query, use the following command.

```
1 dig@siteIP cnn.com
```

# Service Mesh lite

December 31, 2023

An Ingress solution (either hardware or virtualized or containerized) typically performs L7 proxy functions for north-south (N-S) traffic. The Service Mesh lite architecture uses the same Ingress solution to manage east-west traffic as well.

In a standard Kubernetes deployment, east-west (E-W) traffic traverses the built-in kube-proxy deployed in each node. Kube-proxy is an L4 proxy that can only perform TCP/UDP based load balancing and cannot offer the benefits provided by an L7 proxy.

NetScaler (MPX, VPX, or CPX) can provide the benefits of L7 proxy for E-W traffic such as:

- Mutual TLS and SSL offload.
- Content based routing, allow or block traffic based on HTTP and HTTPS header parameters.
- Advanced load balancing algorithms (least connections or least response time).
- Observability of east-west traffic through measuring golden signals (errors, latencies, saturation, traffic volume). NetScaler ADM Service Graph is an observability solution to monitor and debug microservices.

A Service Mesh architecture (such as Istio or LinkerD) is complex to manage. Service Mesh lite architecture is a lightweight version and much simpler to get started to achieve the same requirements.

To configure east-west communication with NetScaler CPX in a Service mesh lite architecture, you must first understand how the kube-proxy is configured to manage east-west traffic.

# East-west communication with kube-proxy

When you create a Kubernetes deployment for a microservice, Kubernetes deploys a set of pods based on the replica count. To access those pods, you create a Kubernetes service which provides an abstraction to access those pods. The abstraction is provided by assigning a Cluster IP address to the service.

Kubernetes DNS gets populated with an address record that maps the service name with the Cluster IP address. So, when an application, say tea wants to access a microservice named coffee then DNS returns the Cluster IP address of the coffee service to the tea application. The tea application initiates a connection which is then intercepted by kube-proxy to load balance it to a set of coffee pods.



# East-west communication with NetScaler CPX in Service Mesh Lite architecture

The goal is to insert the NetScaler CPX in the east-west path and use the Ingress rules to control this traffic.

Perform the following steps to configure east-west communication with NetScaler CPX.

# Step 1: Modify the coffee service definition to point to NetScaler CPX

For NetScaler CPX to manage east-west traffic, the FQDN of the microservice (for example, coffee) should point to the NetScaler CPX IP address instead of the Cluster IP of the target microservice (coffee). (This NetScaler CPX deployment can be the same as the Ingress NetScaler CPX device.) After this modification, when a pod in the Kubernetes cluster resolves the FQDN for the coffee service, the IP address of the NetScaler CPX is returned.



### Note:

If you are deploying service mesh lite to bring up the service graph in NetScaler ADM for observability, then you should add the label citrix-adc: cpx in all the services of your application which are pointing to the NetScaler CPX IP address after modifying the service.

# Step 2: Create a headless service named coffee-headless for coffee microservice pods

Since you have modified the coffee service to point to NetScaler CPX, you need to create one more service that represents coffee microservice deployment.

The following is a sample headless service resource:

```
1 apiVersion: v1
2 kind: Service
3 metadata:
   name: coffee-headless
4
5 spec:
6 #headless Service
    clusterIP: None
7
8
     ports:
9
     - name: coffee-443
       port: 443
11
       targetPort: 443
   selector:
12
       name: coffee-deployment
13
```

# Step 3: Create an Ingress resource with rules for the coffee-headless service

With the changes in the previous steps, you are now ready to create an Ingress object that configures the NetScaler CPX to control the east-west traffic to the coffee microservice pods.

The following is a sample Ingress resource:

```
apiVersion: extensions/v1beta1

kind: Ingress

metadata:

name: coffee-ingress

Spec:

rules:

- host: coffee

http:

paths:

- path: /

backend:

servicePort: 80
```

Using the usual Ingress load balancing methodology with these changes, NetScaler CPX can now load balance the east-west traffic. The following diagrams show how the NetScaler CPX Service Mesh Lite architecture provides L7 proxying for east-west communication between tea and coffee microservices using the Ingress rules:

# NetScaler ingress controller



# East-west communication with NetScaler MPX or VPX in Service Mesh lite architecture

NetScaler MPX or VPX acting as an Ingress can also load balance east-west microservice communication in a similar way as mentioned in the previous section with slight modifications. The following procedure shows how to achieve the same.

# Step 1: Create an external service resolving the coffee host name to NetScaler MPX/VPX IP address

There are two ways to do it. You can add an external service mapping a host name or by using an IP address.

# Mapping by a host name (CNAME)

- Create a domain name for the Ingress endpoint IP address(Content Switching virtual server IP address) in NetScaler MPX or VPX (for example, myadc-instance1.us-east-1. mydomain.com) and update it in your DNS server.
- Create a Kubernetes service for coffee with externalName as myadc-instance1.useast-1.mydomain.com.

 Now, when any pod looks up for the coffee microservice a CNAME(myadc-instance1.us -east-1.mydomain.com) is returned.

```
1 kind: Service
2 apiVersion: v1
3 metadata:
4 name: coffee
5 spec:
6 type: ExternalName
7 externalName: myadc - instance1.us-east-1.mydomain.com
```

**Mapping a host name to an IP address** When you want your application to use the host name **coffee** that will redirect to the virtual IP address hosted in NetScaler MPX or VPX, you can create the following:

```
1 ----
2 kind: "Service"
3 apiVersion: "v1"
4 metadata:
5 name: "coffee"
6 spec:
7
   ports:
    -
8
9
       name: "coffee"
10
        protocol: "TCP"
11
        port: 80
12 ---
13 kind: "Endpoints"
14 apiVersion: "v1"
15 metadata:
16 name: "coffee"
17 subsets:
18
    _
19
      addresses:
20
21
          ip: "1.1.1.1" # Ingress IP in MPX
       ports:
23
24
           port: 80
           name: "coffee"
25
```

## Step 2: Create a headless service for microservice pods

Since you have modified the coffee service to point to NetScaler MPX, you need to create one more service that represents coffee microservice deployment.

# Step 3: Create an Ingress resource

Create an Ingress resource using the ingress.citrix.com/frontend-ip annotation where the value matches the Ingress endpoint IP address in NetScaler MPX or VPX.

Now, you can create an Ingress object that configures the NetScaler MPX or VPX to control the eastwest traffic to the coffee microservice pods.

The following is a sample ingress resource:

```
apiVersion: extensions/v1beta1

kind: Ingress

metadata:

name: coffee-ingress

annotations:

ingress.citrix.com/frontend-ip: "1.1.1.1"

Spec:

rules:

- host: coffee

http:

paths:

- path: /

backend:

serviceName: coffee-headless

servicePort: 80
```

Using the usual ingress load balancing methodology with these changes NetScaler MPX can now load balance east-west traffic. The following diagram shows a NetScaler MPX or VPX configured as the N-S and E-W proxy using the Ingress rules.



# Automated deployment of applications in Service Mesh lite

To deploy an application in a Service Mesh lite architecture, you need to perform multiple tasks manually. However, when you want to deploy multiple applications which consist of several microservices, you have an easier way to deploy the services in a Service Mesh lite architecture. NetScaler provides you an automated way to generate ready to deploy YAMLs.

This document provides information on how to generate all the necessary YAMLs for Service Mesh lite deployment from your existing YAMLs using the NetScaler provided script.

# **NetScaler Operator**

May 2, 2025

Operator is an open source toolkit designed to package, deploy, and manage OpenShift native applications in an effective, automated, and scalable way.

For information on deploying NetScaler Operator, see Deploy NetScaler Operator.

You can deploy the following controllers by using NetScaler Operator:

Deploy NetScaler Ingress Controller

- Deploy NetScaler Ingress Controller with CPX
- Deploy NetScaler Cloud Controller
- Deploy NetScaler GSLB Controller
- Deploy NetScaler IPAM Controller
- Deploy NetScaler Observability Exporter

# **NetScaler Operator release notes**

May 2, 2025

# Version 3.2.0

# Support for NetScaler IPAM Controller

You can now deploy NetScaler IPAM Controller by using NetScaler Operator. For more information, see Deploy NetScaler IPAM Controller in OpenShift using NetScaler Operator.

# Enhanced security for communication between GSLB sites with NetScaler GSLB Controller

For NetScaler GSLB controller deployment, an additional key sitesyncpassword is supported when creating a secret that GSLB controller uses to connect to GSLB devices and push the configuration. This key enhances the security of communication between the GSLB sites. For more information, see Deploy NetScaler GSLB controller.

# Support for UpgradeImageAlways parameter

NetScaler Operator now supports UpgradeImageAlways parameter. If this param is set to True , images will be upgraded for controllers while updating the operator. If user wants to use image different than the current version provided by the operators, then this param should be set to False .

# Version 3.1.2

### What's new

**Support for GSLB Controller deployment** You can now deploy GSLB Controller by using NetScaler Operator. For more information, see Deploy NetScaler GSLB Controller in OpenShift using NetScaler Operator.

**Support to configure multiClusterPrefix parameter** You can now configure the multiClusterPrefix parameter in the NetScaler Ingress Controller deployment YAML. For information on how to set up multi-cluster ingress solution, see Multi-cluster ingress.

**NetScaler certificate validation in NetScaler Ingress Controller** You can now enable NetScaler certificate validation in NetScaler Ingress Controller. For more information, see Enable NetScaler certificate validation in NetScaler Ingress Controller.

**Internet Content Adaptation Protocol (ICAP) support** You can now configure remote content inspection using ICAP.

For more information, see Remote content inspection or content transformation service using ICAP.

Address field of the ingress resource is updated In an NSIC sidecar deployment, in which NetScaler CPX is exposed using a service of type ClusterIP, NodePort, or LoadBalancer, the Address (Status.LoadBalancer.IP) field of the ingress is updated when you specify updateIngressStatus as **true** in the deployment YAML.

For more information, see ingress status update for sidecar deployments.

# **Fixed issues**

• The service account token generated in Azure Kubernetes for Kubernetes API authentication expires every hour. As a result, the controller restarts periodically to retrieve the newly generated token for continued authentication.

# **Upgrade NetScaler Operator**

December 12, 2024

To upgrade NetScaler Operator to version 3.1.2 from earlier versions, perform the following steps:

1. Back up YAMLs of all the installed controllers, such as NSIC, GSLB, and NSOE, by using the following command:

2. Back up the citrix.com CRD instances managed by NSIC and GSLB controller to a YMAL file by using the following command:

3. Remove citrix.com CRDs by using the following command:

```
1 $(oc get crds -o jsonpath='{
2 .items[*].metadata.name }
3 ' | tr ' ' '\n' | grep 'citrix.com'); do oc delete crd $crd; done
```

4. Delete all the controller instances within NetScaler Operator.

#### Important

The **Delete all operand instances for this operator** option displayed as part of the Operator deletion workflow might not work sometimes and some residual components are not deleted. Follow the steps mentioned here to avoid conflicts.

a) Navigate to **Operators > Installed Operators** and click **NetScaler Operator**.

Home >	Installed Operators									
Operators 🗸	Instance of portion of									
OperatorHub										
Installed Operators	Name   Search by name									
Workloads 🗸 🗸	Name I	Managed Namespaces	Status	Last updated	Provided APIs					
Workloads 🗸	Name I NetScaler Operator	Managed Namespaces I	Status Succeeded	Last updated 🚱 Nov 19, 2024, 2:23 PM	Provided APIs Netscaler Cloud Controller					
Workloads 🗸 Vorkloads	Name I NetScaler Operator 3.11 provided by NetScaler	Managed Namespaces I	Status Succeeded Up to date	Last updated Nov 19, 2024, 223 PM	Provided APIs Netscaler Cloud Controller Netscaler CPX With Ingress Controller Netscaler GSLB Controller					
Workloads v Pods Deployments DeploymentConfigs	Name I NetScaler Operator 3/1 provided by NetScaler	Managed Namespaces 🛛	Status © Succeeded Up to date	Last updated	Provided APIs Netscaler Cloud Controller Netscaler CPX With Ingress Controller Netscaler (SSLB Controller Netscaler Ingress Controller View I more					

- b) Click All instances to list all the controller instances.
- c) To delete each controller instance, click the ellipses icon next to a controller and then click **Delete**.

In the following example, a NetScaler CPX with NetScaler Ingress Controller instance is deleted.

Home	>	>	NetScaler Operat	tor					Actions 💌
Operators	~		2.4.2 provided by Ne	etScaler					7100010
OperatorHub		Details	YAML S	ubscription Events	All instances	NetScaler CPX with Ingress	Controller	NetScaler Ingress Controller	NetScale
Installed Operators	4							_	
Workloads	>	All In:	stances <sup>st</sup>	now operands in: 💿 All name	espaces O Current	namespace only		Crea	ite new 🔻
Networking	>	<b>T</b> Filte	er 👻 Name 🗸	<ul> <li>Search by name</li> </ul>	Z				
Storage	>	Nam	e 1	Kind 1	Namespace 1	Status 1	Labels 🌐	Last updated	
Builds	>	NIC	cnnwebsite	NetscalerIngressController	NS cnnwebsite	Conditions: Initialized, Deployed	No labels	🚱 Oct 28, 2024, 12:38 F	PM 🚦
Observe	>	NCV	🖤 cpxcnnwebsite	NetscalerCpxWithIngressC ontroller	NS cnnwebsite	Conditions: Initialized, Deployed	No labels	Oct 28, 2024, 1:59 Pt	4 I
Compute	>	NIC	galaxy	NetscalerIngressController	NS galaxy	Conditions: Initialized, Deployed	No labels	Edit NetscalerCpxWithIngressCont	troller ontroller
User Management	>	NIC	httpbin	NetscalerIngressController	NS httpbin	Conditions: Initialized, Deployed	No labels	Oct 28, 2024, 1:48 Pf	M i

- d) Click **Delete** to confirm deleteion.
- 5. Remove **netscaler.com** CRDs of NetScaler Operator (NSIC, NSIC+CPX, and NSOE) by using the following command:



6. To uninstall NetScaler Operator, navigate to **Operators > Installed Operators**, click the ellipses icon next to NetScaler Ingress Controller Operator, and then click **Uninstall Operator**.

Home	>	Installs	d Operators					
Operators	~	Installed Op	perators are represented	by ClusterServiceVersions within t	nis Namespace. For more info	rmation, see the Understanding Operators	documentation 🖉. Or create an Operati	or and
OperatorHub		ClusterServ	viceVersion using the Op	perator SDK 🖉.				
Installed Operators								
		Name 🝷	Search by name	7				
Workloads	,	Name	I	Managed Namespaces	Status	Last updated	Provided APIs	
Networking	>	6	Grafana Operator	All Namespaces	Succeeded	Nov 8, 2024, 3:53 PM	Grafana Alert Rule Group	:
Storage	>	~	5.15.1 provided by Grafana Labs		op to date		Grafana Dashboard Grafana Datasource	
Builds	>	_	NetScaler	All Namespaces	Succeeded	Nov 8, 2024, 3:59 PM	View 3 more NetScaler CPX with Ingress	i
Observe	>	1	Operator 2.4.2 provided by		Up to date		Controller NetScaler Ingress Edit Subscript	tion
Compute	>		NetScaler				Uninstall Ope	rator

7. Click Uninstall to confirm uninstallation.

# 🛕 Uninstall Operator?

Operator NetScaler Operator will be removed from all-namespaces.

If your Operator configured off-cluster resources, these will continue to run and require manual cleanup.



- 8. Convert the controller YAMLs you backed up in step 1 to the new schema introduced in NetScaler Operator 3.1.2. To do so, for each parameter in the controller YAML, check the respective modified parameter in the new schema and update the YAMLs accordingly. For information on parameters in NetScaler Operator version 3.1.2, see NetScaler Operator.
- 9. Deploy NetScaler Operator. For information on deploying NetScaler Operator, see Deploy NetScaler Operator.

- 10. Install the controllers one after the other with the YAMLs updated in step 8. For more information, see NetScaler Operator.
- 11. Install the CRDs backed up in step 2 by using the following command:

```
1 oc apply -f all_citrix_crd_instances.yaml
```

# **Deploy NetScaler Operator**

February 26, 2025

Perform the following steps:

- 1. Log in to the OpenShift cluster console.
- 2. Navigate to **Operators > OperatorHub**, select **Certified** source in the left panel, select **NetScaler Operator**, and then click **Install**.

Home	OperatorHub	
Operators	Discover Operators from the Kubernetes community and Red Hat partners, curated by Red Hat. You can purchase commercial software	e through Red Hat Marketplace 🗗 You can install Op
Workloads	Catalog providing a seri-service experience.	
Networking	All Items All Items All Items All All All All All All All All All Al	
Storage	Application Runtime	
Builds	Cloud Provider Database Certified	
Observe	Developer Tools NetScaler Operator Development Tools provided by NetScaler	
Compute	Drivers and plugins Integration & Delivery This is an operator to install Cloud	
User Management	Logging & Tracing Native portfolio provided by NetScaler	
Administration	Monitoring	
	OpenShift Optional	
	Other	
	Security	
	Streaming & Messaging	
	Other	
	Source Red Hat (0)	
	<ul> <li>Certified (1)</li> <li>Community (0)</li> <li>Marketplace (0)</li> </ul>	

> NetScale	er Operator y NetScaler
Install	
Channel	NetScaler provides various products to empower Developer/DevOps managing
stable 🔻	OpenShift Cluster. Using this operator you can deploy below controllers for Ne
	1. NetScaler Cloud Controller
Version	2. NetScaler Ingress Controller
3.1.2 🗸	3. NetScaler CPX with Ingress Controller
	4. NetScaler GSLB Controller
Capability level	5. NetScaler Observability Exporter
Sasic Install	
Seamless Upgrades	Installation
Full Lifecycle	Refer installation instructions.
<ul> <li>Deep Insights</li> <li>Auto Pilot</li> </ul>	This operator version contains:
	1. NetScaler Ingress Controller version 2.2.10
Source	2. NetScaler CPX version 14.1-25.111
Certified	3. NetScaler GSLB Controller version 2.2.10
Provider	4. NetScaler Observability Exporter version 1.10.001
NetScaler	5. NetScaler Metrics Exporter 1.4.9
Infractructure features	Support

- 3. To subscribe to NetScaler Operator, select one of the following options:
  - All namespaces on the cluster (default): NetScaler Operator is available in all the namespaces on the OpenShift cluster. Hence, this option enables you to initiate the NetScaler instance from any namespace on the cluster.
  - A specific namespace on the cluster: NetScaler Operator is available in the selected namespace on the OpenShift cluster. Hence, this option enables you to initiate the NetScaler Operator instance on the selected namespace only.
- 4. In this case, let's select **A specific namespace on the cluster** and then Click **Install**.

	OperatorHub > Operator Installation			
Home >	Install Operator			
Operators 👻	Install your Operator by subscribing to one of the update channels to keep the Operator up to date. The strategy determines either manual or automatic updates.			
OperatorHub	Update channel *	NetScaler Operator		
Installed Operators	stable	provided by NetScaler		
Workloads 🗸	Version *	Provided APIs		
	312 *	Netscaler Cloud Controller	Netscaler CPX With Ingress	NGC Netscaler GSLB Controller
Pods		To install NetScaler Cloud Controller for	Controler	To install NetScaler GSLB Controller for
Deployments	Installation mode *	NetScaler VPX/MPX/SDX with Azure Redhat OpenShift in Azure.	To install NetScaler CPX with Ingress Controller.	NetScaler VPA/MPA/SUX
DeploymentConfigs	All namespaces on the cluster (default)     Operator will be available in all Namespaces.			
StatefulSets	A specific namespace on the cluster			
Secrets	Operator will be available in a single Namespace only.	Metropler Ingress Controller	Netronar Obrepublity Exporter	
Confictation	Installed Namespace *	The second ingreas comone	The scale observating exporter	
	dB default     ·	for NetScaler VPX/MPX/SDX.	Exporter for NetScaler.	
GronJobs				
Jobs	Update approval * ①			
DaemonSets	Automatic			
OraclianState	O Manual			
nepitadets				
NeplicationControllers				
HorizontalPodAutoscalers	Install Cancel			

Wait until NetScaler Operator is subscribed successfully.

>	NetScaler Operator netscaler-operator.v3.1.2 provided by NetScaler	⊘
Install View C	ed operator: ready for use Operator View installed Operators in Namespace openshift-operators	

5. (Optional) To deploy NetScaler Operator in all the namespaces, ensure All namespaces on the cluster (default) is selected as the installation mode and openshiftoperators is selected as the installed namespace, and then click Install.

#### Note:

Before deploying NetScaler Operator in all the namespaces, ensure to delete any existing NetScaler Operator instances in all the namespaces.

# Install Operator

Install your Operator by subscribing to one of the update channels to keep the Operator up to c

Update channel * 💿
stable 🗸 🗸
Version *
3.1.2
Installation mode *
<ul> <li>All namespaces on the cluster (default)</li> <li>Operator will be available in all Namespaces.</li> </ul>
<ul> <li>A specific namespace on the cluster</li> <li>Operator will be available in a single Namespace only.</li> </ul>
Installed Namespace *
PR openshift-operators     ▼
Update approval * ③
Automatic
O Manual
Install

Wait until NetScaler Operator is subscribed successfully.



6. Navigate to **Workloads > Pods** section and verify that the **netscaler-Operator-controllermanager** pod is up and running.

Home	Pode							Cre	eate Pod
Operators	1003								
Workloads	▼ Filter ▼ Name ▼	Search by name	7	I					
Pods	Name 1	Status 🗍	Ready 💲	Restarts 1	Owner 1	Memory 1	CPU 1	Created 1	
Deployments	P netscaler-operator-	C Running	2/2	0	RS netscaler-operator-	80.5 MiB	0.001 cores	Nov 19, 2024, 10,200 AM	:
DeploymentConfigs	65fb58f966-ksbdc				65fb58f966			10:29 AM	

You can deploy the following controllers using NetScaler Operator:

- Deploy NetScaler Ingress Controller
- Deploy NetScaler Ingress Controller with CPX
- Deploy NetScaler Cloud Controller
- Deploy NetScaler GSLB Controller
- Deploy NetScaler Observability Exporter

# Deploy NetScaler Ingress Controller in OpenShift using NetScaler Operator

November 29, 2024

In this setup, NetScaler Ingress Controller configures NetScaler MPX or NetScaler VPX residing outside the OpenShift cluster.

# Prerequisites

• Red Hat OpenShift Cluster (version 4.1 or later).

- Deploy NetScaler Operator. For information on how to deploy NetScaler Operator, see Deploy NetScaler Operator.
- Identify the IP address that NetScaler Ingress Controller needs to communicate with NetScaler. This IP address might be any one of the following IP addresses depending on the type of NetScaler deployment:
  - NSIP (for standalone appliances) The management IP address of a standalone NetScaler appliance. For more information, see IP Addressing in NetScaler.
  - SNIP (for appliances in High Availability mode) The subnet IP address. For more information, see IP Addressing in NetScaler.
  - CLIP (for appliances in Cluster mode) The cluster management IP (CLIP) address for a clustered NetScaler deployment. For more information, see IP addressing for a cluster.
- The user name and password of NetScaler VPX or NetScaler MPX used as the ingress device. NetScaler must have a system user account (non-default) with certain privileges so that the NetScaler Ingress Controller can configure NetScaler VPX or NetScaler MPX. For instructions to create a system user account on NetScaler, see Create system user account for NetScaler NetScaler Ingress Controller in NetScaler.

You can directly pass the user name and password as environment variables to the controller, or use OpenShift secrets (recommended). To create a secret for the user name and password using the following command, modify the <username> and <password> to required values:

```
1 oc create secret generic nslogin --from-literal=username=<
    username> --from-literal=password=<password>
```

• To export to any service monitor, install the service monitor CRD by using the following command:

```
1 kubectl create -f https://raw.githubusercontent.com/prometheus-
operator/prometheus-operator/refs/heads/main/example/
prometheus-operator-crd/monitoring.coreos.com_servicemonitors.
yaml
```

# Deploy NetScaler Ingress Controller as a standalone pod using NetScaler Operator

Using NetScaler Operator you can deploy NetScaler Ingress Controller as a standalone pod in an Open-Shift cluster. NetScaler Ingress Controller configures NetScaler MPX or NetScaler VPX, which is deployed as an ingress device or router for an application running in the OpenShift cluster. The following diagram explains the topology:



Perform the following steps:

- 1. Log in to the OpenShift 4.x Cluster console.
- 2. Deploy an Apache application using the console.
  - a) Navigate to **Workloads > Deployments > Create Deployment** and use the following YAML file to create the deployment.

NOTE:

The Apache application is for the demonstration purpose only. You can modify the YAML file based on your requirement.

```
___
1
  apiVersion: apps/v1
2
3 kind: Deployment
4 metadata:
5
   name: apache
6
    labels:
7
         name: apache
8 spec:
9
   selector:
10
     matchLabels:
11
         app: apache
   replicas: 2
12
13
  template:
14
       metadata:
15
         labels:
16
           app: apache
```

```
17 spec:
18 containers:
19 - name: apache
20 image: httpd:latest
21 ports:
22 - containerPort: 80
23 ---
```

b) Navigate to **Workloads > Pods** section and ensure that the Apache application pods are up and running.

Red Hat OpenShift Container Pla	tform				0 ≣	1 <b>19</b> k	ube:admin 👻			
Search	You are logged in as a temporary administrative user. Update the <u>cluster OAuth configuration</u> to allow others to log in.									
Events	Project: default ~						O Add ~			
Catalog Developer Catalog Installed Operators	Pods									
OperatorHub Operator Management	2 Running 0	Pending 0 Terminati	ng O CrashLoopBa	ckOff 0 Completed	Filt	0 Unknown				
Workloads 🗸 🗸	Select All Filters						2 of 2 Items			
Pods	NAME †	NAMESPACE	POD LABELS	NODE	STATUS	READI	NESS			
Deployments Deployment Configs Stateful Sets	P apache- 6f66d8dcb8- 2k2dn	NS default	app=apache pod-te=6f6	N ip-10-0-140- 206.ec2.intern al	₽ Running	Ready	1			
Secrets Config Maps Cron Jobs	P apache- 6f66d8dcb8- brmmg	NS default	app=apache pod-te=6f6	8 ip-10-0-140- 206.ec2.intern al	₽ Running	Ready	I			

3. Create a service for the Apache application. Navigate to **Networking > Services > Create Service** and use the following YAML file.





4. Create an ingress for the Apache application. Navigate to Networking > Ingress > Create Ingress and use the following YAML to create the ingress. Ensure that you update the VIP address of NetScaler VPX in the ingress YAML before applying it in the cluster.

```
apiVersion: networking.k8s.io/v1
 1
   kind: Ingress
2
3 metadata:
4
     annotations:
        ingress.citrix.com/frontend-ip: <NSVIP>
5
6
     name: vpx-ingress
7
     spec:
8
        ingressClassName: nsic-vpx
9
        rules:

    host: citrix-ingress-operator.com

10
11
          http:
12
            paths:
13
            - backend:
14
                service:
15
                  name: apache
16
                  port:
17
                    number: 80
18
              path: /
19
              pathType: Prefix
```



- 5. Navigate to **Operators > Installed Operators** and click **NetScaler Operator**.
- 6. Click the NetScaler NetScaler Ingress Controller tab and select Create NetScalerIngress-Controller.

≡ <mark><sup>l</sup> Red Hat</mark> OpenShift			<b>Ⅲ</b> ♣6	• •	kube:admin 🕶
Administrator	•	You are logged in as a temporary administrative user. Update the <u>skuter.Okuth.configuration</u> to allow others to log in.			
Home	,	Installed Operators   Operator details			
Operators	~	NetScaler Operator     3.1 provided by NetScaler			Actions 👻
OperatorHub Installed Operators		Details YAML Subscription Events All instances Netscaler Cloud Controller Netscaler CPX With Ingress Controller Netscaler GSLB Controller	Netscaler Ing	ress Controller	Netscaler Obser
Workloads Pods	×	NetscaleringressControllers	Create NetscaleringressController		
Deployments DeploymentConfigs StatefulSets Serrets		No operands found Operands are declarative components used to define the behavior of the application.			

The NetScaler NetScaler Ingress Controller YAML definition is displayed. Optionally, you can select Form view button and update the parameters in a form.
Red Hat OpenShift		🗰 🏚 9 🔂 😧 kube:admin 🗸
	You are logged in as a temporary administrative user. Update the cluster OAuth configu	<u>iration</u> to allow others to log in.
Home >	Project: default 🔻	
Operators × OperatorHub	Create NetscalerIngressController Create by completing the form. Default values may be provided by the Operator authors.	
Installed Operators	Configure via: O Form view	
Workloads 🗸 🗸	Alt + F1 Accessibility help   🔞 View shortcuts	NetscaleringressController ×
Pods Deployments DeploymentConfigs StatefulSets Secrets ConfigMaps CronJobs Jobs	1     kind: tetscaler.com/vl       2     aptrexion: netscaler.com/vl       3     metadata:       4     name: nic       5     name: nic       6     spc::       7     adCtedentialSecret: "'       8     affinity: ()       9     analyticscorrig:       10     distributefracing:       11     enablt: faise       12     endpoint:       13     endpoint:       14     jservice: "'       15     service: "'       16     required: faise       17     timeseries:	Schema NetscaleringressController is the Schema for the netscaleringressControllers API  • aplversion strong APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal value, and may reject turrecognized schemas to the latest internal value, and may reject turrecognized values. More infor http://jut.k.88.jocomunuity/contrbuors/devel/sig- architecture/api-conventions.md/#resources
DaemonSets ReplicaSets ReplicationControllers HorizontalPodAutoscalers PodDisruptionBudgets	18       auditios:         19       enable: false         20       events:         21       enable: false         23       enable: false         24       mode: avro         25       port: 30002         26       transactions:         27       enable: false	kind string Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the client submits requests to Cannot be updated. In CannelCase More infer. https://git.k8s.io/community/contributors/devel/sig- architecture/api-conventions.md#types-kinds
Networking 🗸	Create Cancel	• metadata object

7. Refer the following table that lists the mandatory and optional parameters that you can configure during installation.

### Note:

- Ensure that the acceptLicense parameter is set to Yes.
- Specify a unique value for entityPrefix.
- Provide the IP address of NetScaler VPX instance for nsIP parameter and Kubernetes secret created using NetScaler VPX credentials in adcCredentialSecret parameter.
- To associate NetScaler Ingress Controller with the ingress resource specified in this procedure, provide the ingress class name using the ingressClass parameter as ingressClass: ['NSIC-vpx'].

### Note:

Optionally, you can click the form view and set the parameters as required.

### You can configure other available parameters depending on your use case.

Parameters	Mandatory or Optional	Default value	Description
acceptLicense	Mandatory	no	Set this value to Yes to accept the NetScaler Ingress Controller EULA.

Parameters	Mandatory or Optional	Default value	Description
affinity	Optional	N/A	Affinity labels for pod assignment.
extraVolumes	Optional	N/A	Additional VolumeMounts to be
imagePullSecrets	Optional	N/A	List of Kubernetes secrets to be used for pulling the images from a private Docker registry or repository. For more information on how to create this secret, see Pull an Image from a Private Registry.
nodeSelector	Optional	N/A	The node label key:value pair to be used for nodeSelector option in NetScaler Ingress Controller deployment.
podAnnotations	Optional	N/A	Map of annotations to add to the pods.
tolerations	Optional	N/A	The tolerations for the NetScaler Ingress Controller deployment.
Parameter	Mandatory/Optional	Default value	Description
	analyticsConfig		
distributedTracing.enat	ole Optional	false	Set this value to <b>true</b> to enable OpenTracing in NetScaler
distributedTracing.sam	plingrate Optional	100	OpenTracing sampling rate in percentage.

Parameter	Mandatory/Optional	Default value	Description
required	Mandatory	false	Set this value to <b>true</b> if you want to configure NetScaler to send metrics and transaction records to the analytics server.
endpoint.metrics.service	Optional	N/A	The IP address or DNS address of the analytics server. Format: servicename. namespace, servicename. namespace.svc. cluster.local ,namespace/
endpoint.transactions.ser	vice Optional	N/A	servicename. An IP address or a service name with the namespace of the analytics service deployed in the Kubernetes environment. Format: namespace/
timeseries.port	Optional	30002	servicename. The port used to expose the analytics service outside the cluster for a time-series endpoint
timeseries.auditlogs.enab	le Optional	false	Set this value to <b>true</b> to export audit log
timeseries.events.enable	Optional	false	Set this value to <b>true</b> to export events from NetScaler.

Parameter	Mandatory/Optional	Default value	Description
timeseries.metrics.enab	le Optional	false	Set this value to <b>true</b> to enable sending metrics from NetScaler.
timeseries.metrics.mode	e Optional	Avro	Mode of the metric endpoint.
timeseries.metrics.expo	rtFreque <b>ល្អp</b> tional	30	Time interval for exporting time-series data. Possible values range from 30 to 300 seconds.
timeseries.metrics.scher	naFile Optional	schema.json	Name of a schema file with the required NetScaler counters to be added and configured for metrics collector. A reference schema file reference_schema .json with all the supported counters is also available in the path /var/ metrics_conf/. This schema file can be used as a reference to build a custom list of
timeseries.metrics.enab	leNativeScoppaippenal	false	Set this value to <b>true</b> for native export of metrics
transactions.port	Optional	30001	The port used to expose analytics service outside the cluster for a transaction endpoint.

Parameter	Mandatory/Optional	Default value	Description
transactions.enable	Optional	false	Set this value to <b>true</b> to export transactions from NetScaler.
Parameter	Mandatory/Optional	Default value	Description
	exporter		
extraVolumeMounts	Optional	N/A	Additional VolumeMounts to be mounted in the exporter container.
image	Optional	quay.io/netscaler/netsc adc-metrics- exporter@sha <sha value of the latest release&gt;</sha 	ːalđʰe metrics exporter image hosted on Quay.io.
required	Optional	false	Set to <b>true</b> to run the Exporter for NetScaler Stats along with NetScaler Ingress Controller to pull metrics for the NetScaler MPX or NetScaler VPX.
pullPolicy	Optional	lfNotPresent	The exporter image pull policy.
ports.containerPort	Optional	8888	The exporter container port.
resources	Optional	N/A	CPU/memory resources for a metrics exporter container.
serviceMonitorExtraLabel	s Optional	N/A	Extra labels for a service monitor when exporter is enabled.

Parameter	Mandatory/Optional	Default value	Description
	ingressController		
clusterName	Optional	N/A	A unique identifier of the Kubernetes cluster on which NetScaler Ingress Controller is deployed.
defaultSSLCertSecret	Optional	N/A	Kubernetes secret name that needs to be used as a default non-SNI certificate in NetScaler.
defaultSSLSNICertSecret	Optional	N/A	Kubernetes secret name that needs to be used as a default SNI certificate in NetScaler.
disableAPIServerCertVerify	/ Optional	false	Set this parameter to True for disabling API Server certificate verification.
disableOpenshiftRoutes	Optional	false	By default OpenShift routes are processed in the OpenShift environment. This parameter can be used to disable NetScaler Ingress Controller processing the OpenShift routes.
enableLivenessProbe	Optional	true	Set to <b>false</b> to disable liveness porbe settings for NetScaler Ingress Controller.
enableReadinessProbe	Optional	true	Set to <b>false</b> to disable readiness probe settings for NetScaler Ingress Controller.

Parameter	Mandatory/Optional	Default value	Description
entityPrefix	Optional	N/A	The prefix for the
			resources on NetScaler
			MPX or NetScaler VPX.
ignoreNodeExternalIP	Optional	False	While adding node IP
			address as a service
			group member for
			type LoadBalancer
			services or NodePort
			services, NetScaler
			Ingress Controller has
			a selection criteria
			where it chooses
			Node ExternalIP
			if available and
			Node InternalIP
			if
			Node ExternalIP
			is not available. But
			some users might
			want to use Node
			InternalIP over Node
			ExternalIP even if
			Node ExternalIP is
			present. If this variable
			is set to <b>true</b> , NSIC
			prioritizes the Node
			InternalIP to be used
			for service group
			members even if Node
			ExternalIP is present.
image	Mandatory	quay.io/	The NetScaler Ingress
		netscaler/	Controller image
		netscaler-k8s-	hosted on Quay.io.
		ingress-	
		controller@sha	
		value of latest	
		release	

Parameter	Mandatory/Optional	Default value	Description
ingressClass	Optional	N/A	Parameter to associate
			a particular ingress
			resource with
			NetScaler Ingress
			Controller. For more
			information on ingress
			class, see Ingress class
			support. For
			Kubernetes version >=
			1.19, this parameter
			creates an IngressClass
			object with the name
			specified here.
extraVolumeMounts	Optional	N/A	Additional
			VolumeMounts to be
			mounted in the
			NetScaler Ingress
			Controller container.
ipam	Optional	false	Set this paramter to
			true to use the IPAM
			controller to
			automatically allocate
			an IP address to the
			service of type
			LoadBalancer.
jsonLog	Optional	false	Set this argument to
			true if log messages
			are required in the
			JSON format.

Parameter	Mandatory/Optional	Default value	Description
kubernetesURL	Optional	N/A	The kube-apiserver url that NetScaler Ingress Controller uses to register the events. If the value is not specified, NetScaler Ingress Controller uses the internal kube-apiserver IP address
livenessProbe	Optional	N/A	LivenessPorbe settings for NetScaler Ingress
logLevel	Optional	INFO	The log level to control the logs generated by NetScaler Ingress Controller. The supported log levels are: CRITICAL, ERROR, WARNING, INFO, DEBUG, TRACE, and NONE.
logProxy	Optional	N/A	Elasticsearch or Kafka or Zipkin endpoint for NetScaler observability exporter.

Parameter	Mandatory/Optional	Default value	Description
multiClusterPrefix	Optional	mc	The prefix for the
			shared resources on
			NetScaler MPX or
			NetScaler VPX for
			multicluster ingress
			feature. NetScaler
			Ingress Controllers
			that are collaboratively
			sharing the content
			switching virtual
			server IP address must
			be configured with the
			same value for
			multiclusterPrefix. For
			more information see
			this.
namespaceLabels	Optional	N/A	Provide the
			namespace labels
			selectors to be used by
			NetScaler Ingress
			Controller for
			routeSharding in the
			OpenShift cluster.
nodeLabels	Optional	,,,,	If there are pods on
			nodes with node
			labels, NetScaler
			Ingress Controller
			configures NetScaler
			to advertise the VIP
			using BGP.

Parameter	Mandatory/Optional	Default value	Description
nodeWatch	Optional	false	Set the parameter to
			<b>true</b> if you want to
			automatically
			configure network
			route from NetScaler
			MPX or NetScaler VPX
			to the pods in the
			Kubernetes cluster.
			For more information,
			see Automatically
			configure route on the
Di u	Outlined	(-)	NetScaler instance.
nsncPbr	Optional	Talse	Set this parameter to
			true to inform
			NetScaler Ingress
			Controller that
			NetScaler Node
			Controller is
			configuring Policy
			Based Routes (PBR) or
			NetScaler. For more
			information, see
			NSNC-PBR-SUPPORT
openShift	Optional	true	Set this parameter to
			false if the
			OpenShift
			environment is not
			being used.

Parameter	Mandatory/Optional	Default value	Description
optimizeEndpointBinding	g Optional	false	Set this parameter to <b>true</b> to enable binding of back-end endpoints to a service group in a single API-call. Recommended when endpoints (pods) per application are large in number. Applicable only for NetScaler
podIPsforServiceGroupM	embersOptional	false	<pre>version &gt;=13.0-45.7. By default, NetScaler Ingress Controller adds node IP address and node port as service group members while configuring type LoadBalancer Services and NodePort services. If set to True NetScaler Ingress Controller adds pod IP address and pod port instead of node IP address and node port. You can set this parameter to true if there is a route between NetScaler and K8s clusters internal pods either using feature-node-watch argument or using NetScaler Node Controller</pre>

Parameter	Mandatory/Optional	Default value	Description
profileHttpFrontend	Optional	N/A	The front-end HTTP profile. For details, see Configuration using FRON- TEND_HTTP_PROFILE
profileSslFrontend	Optional	N/A	The front-end SSL profile. For details, see Configuration using FRON- TEND_SSL_PROFILE
profileTcpFrontend	Optional	N/A	The front-end TCP profile. For details, see Configuration using FRON- TEND_TCP_PROFILE
pullPolicy	Mandatory	IfNotPresent	The NetScaler Ingress Controller image pull policy.
rbacRole	Optional	false	<ul> <li>Set this parameter to</li> <li>true to deploy</li> <li>NetScaler Ingress</li> <li>Controller with RBAC</li> <li>role set rbacRole=true;</li> <li>by default NetScaler</li> <li>Ingress Controller gets</li> <li>installed with RBAC</li> <li>Cluster-</li> <li>Role(rbacRole=false).</li> </ul>
readinessProbe	Optional	N/A	Readiness probe settings of NetScaler Ingress Controller.
resources	Optional	N/A	CPU/memory resources for the NetScaler Ingress Controller container.

Parameter	Mandatory/Optional	Default value	Description
routeLabels	Optional	N/A	The route labels selectors to be used by NetScaler Ingress Controller for routeSharding in OpenShift cluster.
serviceClass	Optional	N/A	<ul> <li>Parameter used to</li> <li>associate only specific</li> <li>services to NetScaler</li> <li>Ingress Controller if</li> <li>multiple ingress</li> <li>controllers are</li> <li>deployed. For more</li> <li>information on service</li> <li>class, see Service class</li> <li>support.</li> </ul>
setAsDefaultIngressClass	Optional	False	Sets the IngressClass object as the default ingress class. New ingresses without an "ingressClassName" field specified will be assigned this ingress class. This parameter is applicable only for Kubernetes versions >= 1.19.

updateIngressStatus	Optional	true	Set this parameter if
			Status.
			LoadBalancer.
			Ingress field of the
			ingress resources
			managed by NetScaler
			Ingress Controller
			needs to be updated
			with allocated IP
			addresses. For more
			information see this.

Parameter	Mandatory/Optional	Default value	Description
	netscaler		
adcCredentialSecret	Mandatory	NA	Secret required for NetScaler Ingress Controller to connect to NetScaler.
nitroReadTimeout	Optional	20	The nitro read timeout in seconds.
nsConfigDnsRec	Optional	false	Set to <b>true</b> to enable DNS address record addition in NetScaler through ingress.
nsCookieVersion	Optional	0	The persistence cookie version (0 or 1).
nsDnsNameserver	Optional	N/A	DNS nameservers in NetScaler.
nsHTTP2ServerSide	Optional	OFF	Set this parameter to ON for enabling HTTP2 for NetScaler service group configurations.
nsIP	Mandatory	N/A	NetScaler IP address.
nsLbHashAlgo.required	Optional	false	The load balancing consistent hashing algorithm.

Parameter	Mandatory/0	Optional	Default value	Description
nsLbHashAlgo.hashFin	gers Optior	al	256	The number of fingers to be used for the hashing algorithm. Possible values are from 1 to 1024.
nsLbHashAlgo.hashAlg	orithm Optior	al	default	The supported algorithm. The supported algorithms are <b>default</b> , jarh, and prac.
nsPort	Optior	ial	443	The port used by NetScaler Ingress Controller to communicate with NetScaler. You can use port 80 for HTTP.
nsProtocol	The protocol NetScaler I Controlle communica NetScaler. Yo HTTP on p	used by ngress er to ite with u can use ort 80.	Optional	P
nsSNIPS	Optior	al	N/A	The list of subnet IP addresses on the NetScaler, which is used to create PBR routes instead of static routes PBR support
nsSvcLbDnsRec	Optior	ial	false	Set this parameter to <b>true</b> to enable DNS address record addition in NetScaler through yype LoadBalancer Service.
nsVIP	Optior	al	N/A	The virtual IP address on NetScaler.

1. After updating the values for the required parameters, click **Create**.

Ensure that NetScaler Ingress Controller is successfully deployed and initialized.

Home	•	Installed C	Operators > O	perator details										
Operators	~	>	NetScaler Op 311 provided by	y NetScaler										Actions
OperatorHub Installed Operators		Details	YAML	Subscription	Events	All instances	Netscaler Cloud	Controller	Netscaler CPX With	Ingress Controller	Netscaler GSLB Control	ller Netscaler Ingre	ss Controller	Netscaler (
Workloads Pods	Ť	Netse	calerIngr	essControlle	ers							I	Create Netscaler	IngressControlle
Deployments DeploymentConfigs		Name	Search by r	name	7									
StatefulSets		Name	e I		Kind	I		Status 1		Labels 1		Last updated 🕴		
Secrets ConfigMaps		NIC	netscaleringre	sscontroller-sample	Nets	calerIngressControll	er	Condition: A	Wallable	app.kubernetes.io, app.kubernetes.io,	managed-by=kustomize iname=netscaler-operator	Nov 19, 2024, 11:13 AM		1

2. Navigate to **Workloads > Pods** section and verify whether the NetScaler NetScaler Ingress Controller pod is up and running.

<b>Red Hat</b> OpenShift							∎ ♠9 €	kube:a	dmin <del>-</del>
Home >	↑ Project: default	You are logged in	as a temporary adr	ninistrative user. U	pdate the <u>cluster OAuth configura</u>	ntion to allow others to	o log in.		
Operators  OperatorHub Installed Operators	Pods T Filter • Name •	Search by name7						Cre	ate Pod
Workloads	Name †	Status 1	Ready 1	Restarts 🗍	Owner 1	Memory 1	CPU 1	Created 1	
Pods	P apache-7c646cfd49- 9f2lh	C Running	1/1	0	RS apache-7c646cfd49	27.1 MiB	0.000 cores	Jun 23, 2023, 3:35 PM	÷
Deployments DeploymentConfigs	apache-7c646cfd49- dtdv7	2 Running	1/1	0	RS apache-7c646cfd49	21.2 MiB	0.000 cores	Jun 23, 2023, 3:35 PM	:
StatefulSets Secrets	P netscaler-operator- controller-manager- 5dbf5649cf-rvrjr	2 Running	2/2	0	RS netscaler-operator- controller-manager- 5dbf5649cf	199.9 MiB	0.007 cores	Jun 23, 2023, 3:34 PM	8
ConfigMaps	Pinic-netscaler-ingress- controller-1-deploy	Completed	0/1	0	RC nic-netscaler-ingress- controller-1	-	0.000 cores	Jun 23, 2023, 3:55 PM	E
CronJobs Jobs	Pinic-netscaler-ingress- controller-1-vvvk8	C Running	1/1	0	RC nic-netscaler-ingress- controller-1	-	-	Jun 23, 2023, 3:55 PM	*
DaemonSets ReplicaSets ReplicationControllers HorizontalPodAutoscalers PodDisruptionBudgets									
Networking V	•								

### 3. Verify the deployment by sending traffic:

1 curl http://citrix-ingress-Operator.com --resolve citrix-ingress-Operator.com:80:<VIP>

The previous curl command should return the following:

```
1 <html><body><h1>It works!</h1></body></html>
```

### Note:

- Ensure that the pod network in OpenShift cluster is reachable from NetScaler VPX or NetScaler MPX if you are using a service of type ClusterIP for your application. To configure static route automatically using NetScaler Ingress Controller, see Configure static route.
- When you delete an NetScaler Ingress Controller instance, the ClusterRole and ClusterRole binding associated with the service account of the NetScaler Ingress Controller instance

are not deleted automatically. You must manually delete the ClusterRole and ClusterRole binding associated with the service account of the deleted NetScaler Ingress Controller instance.

### References

- For information about how to deploy NetScaler Observability Exporter using NetScaler Operator, see Deploy NetScaler Observability Exporter using NetScaler Operator.
- For information about how to deploy NetScaler Agent Operator, see Install a NetScaler agent operator using the OpenShift console.
- Alternatively, you can deploy NetScaler Ingress Controller using Helm charts. See Deploy the NetScaler NetScaler Ingress Controller using Helm charts.

# Deploy NetScaler Ingress Controller (NSIC) with CPX in OpenShift using NetScaler Operator

February 14, 2025

In this setup, NetScaler Ingress Controller configures NetScaler CPX deployed in the OpenShift cluster.

### **Prerequisites**

- Access to Red Hat OpenShift Cluster (version 4.1 or later).
- Deploy NetScaler Operator. For information on how to deploy NetScaler Operator, see Deploy NetScaler Operator.
- To export to any service monitor, install the service monitor CRD by using the following command:

• Install Prometheus Operator if you want to view the metrics of the NetScaler CPX collected through the direct Prometheus export.

<sup>1</sup> kubectl create -f https://raw.githubusercontent.com/prometheusoperator/prometheus-operator/refs/heads/main/example/ prometheus-operator-crd/monitoring.coreos.com\_servicemonitors. yaml

### Deploy NetScaler Ingress Controller as a sidecar with NetScaler CPX

Using the NetScaler Operator, you can deploy NetScaler CPX with the NetScaler Ingress Controller as a sidecar. The NetScaler Ingress Controller configures the NetScaler CPX which is deployed as an ingress or router for an application running in the OpenShift cluster. The following diagram explains the topology.

Perform the following steps:

- 1. Log in to the OpenShift 4.x Cluster console.
- 2. Deploy an Apache application using the console.
  - a) Navigate to **Workloads > Deployments > Create Deployment** and use the following YAML to create the deployment.

1	
2	apiVersion: apps/v1
3	kind: Deployment
4	metadata:
5	name: apache
6	labels:
7	name: apache
8	spec:
9	selector:
10	matchLabels:
11	app: apache
12	replicas: 2
13	template:
14	metadata:
15	labels:
16	app: apache
17	spec:
18	containers:
19	- name: apache
20	<pre>image: httpd:latest</pre>
21	ports:
22	- containerPort: 80
23	



Note:

The Apache application is for the demonstration purpose only. You can modify the YAML file based on your requirement.

b) Navigate to Workloads > Pods section and ensure that the Apache application pods are up and running.

Red Hat OpenShift Container P	latform				⊕ ≣	0	kube:admin 👻
Search	You are	logged in as a temporary a	dministrative user. Upda	ate the <u>cluster OAuth co</u>	nfiguration to	allow others to lo	g in.
Events	Project: default ~						O Add ~
Catalog  Developer Catalog Installed Operators OperatorHub Operator Management	Pods Create Pod	Pending 0 Terminating	0 CrashLoopBack	xOff 0 Completed	Filt	er Pods by name	-
Workloads 🗸	Select All Filters						2 of 2 Items
Pods	NAME †	NAMESPACE	POD LABELS	NODE	STATUS	READ	INESS
Deployments Deployment Configs Stateful Sets	apache- 6f66d8dcb8- 2k2dn	NS) default	app=apache pod-te=6f6	(N) ip-10-0-140- 206.ec2.intern al	2 Running	Read	y I
Secrets Config Maps Cron Jobs	e apache- 6f66d8dcb8- brmmg	(NS) default	app=apache pod-te=6f6	8 ip-10-0-140- 206.ec2.intern al	2 Running	Read	y I
lobe							

3. Create a service for the Apache application. Navigate to **Networking > Services > Create Service** and use the following YAML to create the service.

apiVersion: v1
 kind: Service
 metadata:
 name: apache

```
5 spec:
6 ports:
7 - port: 80
8 targetPort: 80
9 selector:
10 app: apache
```

Red Hat OpenShift			₩ ♠ 3 ♀ Rube:admin ▼
📽 Administrator	•	You are logged in as a temporary administrative user. Update the cluster OAuth configurat	<u>ion</u> to allow others to log in.
łome	>	Project: default 🔻	
Operators	~	Create Service Create by manually entering YAML or JSON definitions, or by dragging and dropping a file into the editor.	
OperatorHub Installed Operators		Alt • F1 Accessibility help         I View shortcuts           1 #Expose the apache web server as a Service         I I I I I I I I I I I I I I I I I I I	Service X
forkloads	>	2 aplVersion: v1 3 kind: Service 4 metadata:	Schema
etworking Services	~	5         name: apache           6         spect:           7         ports:           8         - port:80           9         i targetPort:80	Service is a named abstraction of software service (for example, mysql) consisting of local port (for example 3306) that the proxy listens on, and the selector that determines which pods will answer requests sent through
Routes Ingresses		19 selector: 11 app: apache	the proxy.  • apiVersion string
NetworkPolicies			APIVersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to the latest internal
orage	`		value, and may reject unrecognized values. More info:
ilds	>		https://git.k8s.io/community/contributors/devel/s ig-architecture/api-conventions.md#resources
oserve	>	Cranda Changel	• kind
ompute	>	Create Cancer	string

4. Create an Ingress for the Apache application. Navigate to **Networking > Ingress > Create Ingress** and use the following YAML to create the ingress.





- 5. Navigate to **Operators > Installed Operators** and select **NetScaler Operator**.
- 6. Click NetScaler CPX with Ingress Controller tab.

Instable Operators > Operator details  NetScaler Operator  12 provided by/McScaler				Actions 👻					
Details YAML Subscription Events All instances Netscaler Cloud Controller Netscaler	CPX With Ingress Controller Netscaler GSLB Controller	Netscaler Ingress Controller	Netscaler Observability Exporter						
NetscalerCpxWithIngressControllers Stow operands in:   Al namespace O Current namespace only									
No operands found									
	On assault are destaution components used to define the holowing of	the							

7. Click Create NetScalerCpxWithIngressController.

The NetScaler CPX with NetScaler Ingress Controller YAML definition is displayed.

## Create NetscalerCpxWithIngressController

Create by completing the form. Default values may be provided by the Operator authors.



- Ensure that the acceptLicense parameter is set to Yes.
- To associate NSIC with the ingress resource specified in this procedure, provide the ingress class name using the ingressClass parameter as ingressClass: [' nsic-cpx].

Parameters	Mandatory or Optional	Default value	Description
acceptLicense	Mandatory	no	Set Yes to accept the
			NetScaler Ingress
			Controller EULA.
affinity	Optional	N/A	Affinity labels for pod
			assignment.

Refer to the following table to configure other available parameters depending on your use case.

Parameters	Mandatory or Optional	Default value	Description
nodeSelector	Optional	N/A	The node label
			key:value pair to be
			used for nodeSelector
		option in NetSca	option in NetScaler
			Ingress Controller
			deployment.
podAnnotations	Optional	N/A	Map of annotations to
		add to tl	add to the pods.
replicaCount	Optional	1	The number of
			CPX-NSIC pods to be
			deployed. With
			cpxBgpRouter :
			<b>true</b> , replicaCount is
			1 because CPX is
			deployed as a
			DaemonSet.
tolerations	Optional	N/A	Specify the tolerations
			for the CPX-NSIC
			deployment.

Parameter	Mandatory/Optional	Default value	Description
	admSettings		
ADMIP	Optional	N/A	The NetScaler Console (formerly NetScaler ADM) IP address.
bandWidth	Optional	1000	The desired bandwidth capacity for NetScaler CPX in Mbps.
bandWidthLicense	Optional	false	Set this parameter to <b>true</b> if you want to use bandwidth-based licensing for NetScaler CPX.
cpxCores	Optional	1	The desired number of vCPUs for NetScaler CPX.

Parameter	Mandatory/Optional	Default value	Description
licenseEdition	Optional	PLATINUM	The possible values
			are Standard,
			Platinum, and
			Enterprise.
licenseServerIP	Optional	N/A	NetScaler Console
			(formerly NetScaler
			ADM) IP address to
			license NetScaler CPX.
			For more information,
			see Licensing.
licenseServerPort	Optional	27000	NetScaler Console
			(formerly NetScaler
			ADM) port if
			non-default port is
			used.
loginSecret	Optional	N/A	The secret key to log in
			to NetScaler Console.
platform	Optional	false	Set to <b>true</b> to enable
			CP1000 platform
			license.
vCPULicense	Optional	N/A	Set the value to <b>true</b>
			if you want to use a
			vCPU-based license for
			NetScaler CPX.
Parameter	Mandatory/Optional	Default value	Description
	analyticsConfig		
distributedTracing.enable	Optional	false	Set the value to <b>true</b>

distributedTracing.samplingrate Optional

in NetScaler. The OpenTracing sampling rate in percentage.

100

to enable OpenTracing

Parameter I	Mandatory/Optional	Default value	Description
endpoint.metrics.service	Optional	N/A	Set this value as the IP address or DNS address of the analytics server. For- mat:servicename. namespace.svc. cluster.local, namespace/ servicename.
endpoint.transactions.servi	ce Optional	N/A	Set this parameter as the IP address or service name with namespace of the analytics service deployed in the Kubernetes environment. Format: namespace/ servicename.
required	Mandatory	false	Set this parameter to <b>true</b> if you want to configure NetScaler to send metrics and transaction records to the analytics service.
timeseries.port	Optional	5563	The port used to expose analytics service for the time-series endpoint.
timeseries.metrics.enable	Optional	false	Set this value to <b>true</b> to enable sending metrics from NetScaler.
timeseries.metrics.mode	Optional	Avro	The mode of metric endpoint.

Parameter	Mandatory/Opt	ional Default value	Description
timeseries.metrics.ex	oortFreque <b>0p</b> tional	30	The time interval for exporting the time-series data. The possible values range from 30 to 300 seconds
timeseries.metrics.sch	nemaFile Optional	schema.json	The name of a schema file with the required NetScaler counters to be added and configured for the metrics collector. A reference schema file reference_schema .json with all the supported counters is also available in path /var/ metrics_conf/. This schema file can be used as a reference to build a custom list of counters.
timeseries.metrics.en	ableNative <b>Scpaipe</b> nal	false	Set this parameter to <b>true</b> for native export of metrics.
timeseries.auditlogs.e	nable Optional	false	Set this parameter to <b>true</b> to export audit log data from NetScaler.
timeseries.events.ena	ble Optional	false	Set this paramter to <b>true</b> to export events from NetScaler
transactions.enable	Optional	false	Set this parameter to <b>true</b> to export transactions from NetScaler.

Parameter	Mandatory/Optional	Default value	Description
transactions.port	Optional	5557	The port used to expose analytics service for transaction endpoint.

Parameter	Mandatory/Optional	Default value	Description
	exporter		
required	Optional	false	Set this parameter to <b>true</b> to run the <b>Exporter for NetScaler</b> Stats along with NetScaler Ingress Controller to pull metrics for NetScaler
image	Optional	quay.io/ netscaler/ netscaler-adc- metrics- exporter@sha256 :3	The metrics exporter image hosted on Quay.io.

### d54b9151c742c00117be74e1ffd54430358e4cf628

pullPolicy	Optional	lfNotPresent	The exporter image pull policy.
resources	Optional	{}	CPU/memory resources for the metrics exporter container.
ports.containerPort	Optional	8888	The exporter container port.
extraVolumes	Optional	N/A	Additional volumes for the additional volumeMounts.

Parameter	Mandatory/Optional	Default value	Description
imagePullSecrets	Optional	N/A	A list of Kubernetes secrets to be used for pulling the images from a private Docker registry or repository. For more information on how to create this secret, see Pull an Image from a Private Registry.

Parameter	Mandatory/Optional	Default value	Description
	ingressController		
defaultSSLCertSecret	Optional	N/A	The Kubernetes secret name that needs to be used as a default non-SNI certificate in NetScaler.
defaultSSLSNICertSecret	Optional	N/A	The Kubernetes secret name that needs to be used as a default SNI certificate in NetScaler.
disableAPIServerCertVerif	y Optional	false	Set this parameter to <b>true</b> for disabling API server certificate verification.
disableOpenshiftRoutes	false	By default OpenShift routes are processed in the OpenShift environment. Set this parameter to <b>true</b> to disable processing of the OpenShift routes by NetScaler Ingress	

Parameter	Mandatory/Optional	Default value	Description
enableLivenessProbe	Optional	true	Set this parameter to
			false to disable
			liveness probe settings
			for NetScaler Ingress
			Controller.
enableReadinessProbe	Optional	true	Readiness probe
			settings for NetScaler
			Ingress Controller.
entityPrefix	Optional	N/A	The prefix for the
			resources on NetScaler
			CPX.
extraVolumeMounts	Optional	N/A	
image	Mandatory	quay.io/	The NetScaler Ingress
		netscaler/	Controller image
		netscaler-k8s-	hosted on Quay.io.
		ingress-	
		controller@sha250	5
		:2	
		ff38654476234bfe2	2c3908932702b90c8e

ingressClass	Optional	N/A	Parameter to associate a particular ingress resource with NetScaler Ingress Controller. For more information on ingress class, see Ingress class support. For Kubernetes version >= 1.19, this parameter creates an IngressClass object with the name specified here.

Parameter	Mandatory/Optional	Default value	Description
ipam	Optional	false	Set this parameter to <b>true</b> to enable IP address management using the IPAM
jsonLog	Optional	false	Set this parameter to <b>true</b> if log messages are required in the ISON format
kubernetesURL	Optional	N/A	The kube-apiserver url that NetScaler Ingress Controller uses to register the events. If the value is not specified, NetScaler Ingress Controller uses the internal kube-apiserver IP address
livenessProbe.initialDelayS	second3ptional	30	The time period in seconds for which Kubernetes waits before performing the first liveness check.
livenessProbe.periodSecon	ids Optional	60	After the initial delay, Kubernetes performs a liveness check every 60 seconds.
logLevel	Optional	INFO	The log level to control the logs generated by NSIC. The supported log levels are: CRITICAL, ERROR, WARNING, INFO, DEBUG, TRACE, and NONE.

Parameter	Mandatory/Optional	Default value	Description
logProxy	Optional	N/A	Elasticsearch or Kafka or Zipkin endpoint for NetScaler observability exporter.
namespaceLabels	Optional	N/A	The namespace labels selectors to be used by NetScaler Ingress Controller for routeSharding in OpenShift cluster.
openshift	Optional	true	Set this parameter to <b>false</b> if the OpenShift environment is not being used.
optimizeEndpointBinding	g Optional	false	Set to <b>true</b> to enable binding of back-end endpoints to a service group in a single API-call. We recommended to set this parameter to true when the number of endpoints (pods) per application are large in number. Applicable only for NetScaler Version >=13.0-45.7
profileHttpFrontend	Optional	N/A	The front-end HTTP profile. For details, see Configuration using FRON- TEND_HTTP_PROFILE

Parameter	Mandatory/Optional	Default value	Description
profileSslFrontend	Optional	N/A	The front-end SSL profile. For details, see Configuration using FRON-
profileTcpFrontend	Optional	N/A	TEND_SSL_PROFILE The front-end TCP profile. For details, see Configuration using FRON-
pullPolicy	Mandatory	IfNotPresent	TEND_TCP_PROFILE The NetScaler Ingress Controller image pull
rbacRole	Optional	false	To deploy NSIC with RBAC role, set rbacRole= <b>true</b> . By default, NSIC gets installed with ClusterRole (rbacRole=false).
readinessProbe	Optional	N/A	Readiness probe settings.
resources	Optional	8	CPU/memory resources for the NetScaler Ingress Controller container
routeLabels	Optional	N/A	You can use this parameter to provide the route labels selectors to be used by NetScaler Ingress Controller for routeSharding in the OpenShift cluster.

Parameter	Mandatory/Optional	Default value	Description
setAsDefaultIngressClass	Optional	False	Sets the IngressClass object as the default ingress class. New ingresses without an "ingressClassName" field specified will be assigned this ingress class. This parameter is applicable only for Kubernetes versions >= 1.19.
updateIngressStatus	Optional	True	Set this parameter if Status. LoadBalancer. Ingress field of the ingress resources managed by NetScaler Ingress Controller needs to be updated with allocated IP addresses. For more information see this.
Parameter	Mandatory/Optional	Default value	Description
	netscalerCpx		
bgpSettings.bgpConfig	Optional	N/A	This parameter represents BGP configurations in the YAML format. For the description about individual fields, refer the documentation

Parameter	Mandatory/Optional	Default value	Description
cpxBgpRouter	Optional	false	If this parameter is set to true, CPX is deployed as a daemonset in the BGP controller mode, where the BGP advertisements are done for attracting external traffic to the Kubernetes clusters.
cpxCommands	Optional	N/A	User-provided NetScaler bootup config that is applied when CPX is instantiated. This config is not a dynamic config, and any subsequent changes to the configmap don't reflect in the CPX config unless the pod is restarted. For more info, refer this.
cpxLicenseAggregator	Optional	N/A	IP address/FQDN of the CPX License Aggregator if it is being used to license the CPX.

Parameter	Mandatory/Optional	Default value	Description		
cpxShellCommands	Optional	N/A	A user-provided		
			bootup config that is		
			applied when CPX is		
			instantiated. Note that		
			this is not a dynamic		
			config, and any		
			subsequent changes		
			to the configmap don't		
			reflect in the CPX		
			config unless the pod		
			is restarted. For more		
			info, refer the		
			documentation.		
daemonSet	Optional	false	Set this parameter to		
			true if NetScaler CPX		
			must be deployed as a		
			DaemonSet.		
enableLivenessProbe	Optional	true	Set to <b>false</b> to		
			disable liveness probe		
			settings for NetScaler		
			CPX.		
enableStartupProbe	Optional	true	Set to <b>false</b> to		
			disable the		
			startupProbe settings		
			for NetScaler CPX.		
extraVolumeMounts	Optional	N/A			
hostName	Optional	N/A	This entity is used to		
			set Hostname of the		
			CPX.		
image	Mandatory	quay.io/netscaler/net	scal <b>e</b> retScaler CPX image		
		cpx@sha256:04defadf631 <b>b250et6ada5</b> 38 <b>940</b> 8ed7afe8035k			
Parameter	Mandatory/Optional	Default value	Description		
------------------	--------------------	---------------	---		
ingressIP	Optional	N/A	External IP address to be used by ingress resources if the value is not overridden by ingress.com/ frontend-ip annotation in ingress resources. This IP address is also advertised to external		
			routers when cpxBgpRouter is set to true		
livenessProbe	Optional	N/A	livenessProbe settings for NetScaler CPX.		
mgmtHttpPort	Optional	9080	The HTTP port to be used for a NodePort CPX service.		
mgmtHttpsPort	Optional	9443	The HTTPS port to be used for a NodePort CPX service.		
nitroReadTimeout	Optional	20	The nitro read timeout in seconds.		
nsConfigDnsRec	Optional	false	To enable/disable DNS address record addition in NetScaler through ingress.		
nsCookieVersion	Optional	0	The persistence cookie version. Possible values are 0 and 1.		
nsDnsNameserver	Optional	N/A	DNS nameservers in NetScaler.		
nsEnableLabels	Optional	true	Enable labels for NetScaler CPX.		

Parameter	Mandatory/Optional	Default value	Description
nsGateway	Optional	192.168.1.1	Gateway used by CPX for internal communication when it's run in the host mode, i.e when cpxBgpRouter is set to true. If this parameter is not specified, the first IP address in the nsIP network is used as a gateway. It must be in the same
nsHTTP2ServerSide	Optional	OFF	network as nsIP. Set this parameter to ON to enable HTTP2 for NetScaler service
nsIP	Optional	192.168.1.2	NetScaler IP address used by NetScaler CPX for internal communication when it's run in the host mode, i.e when cpxBgpRouter is set to <b>true</b> . A /24 internal network is created in this IP range which is used for internal communications within the network namespace.
nsLbHashAlgo.required	Optional	false	Set this parameter to <b>true</b> to set the LoadBalancing consistent hashing algorithm.

# NetScaler ingress controller

Parameter	Mandatory	//Optional	Default value	Description
nsLbHashAlgo.hashFing	ers Opti	onal	256	The number of fingers to be used for the hashing algorithm. Possible values are from 1 to 1024.
nsLbHashAlgo.hashAlgo	rithm Opti	onal	'default'	Supported algorithms are <b>default</b> , jarh, and prac.
nsProtocol	Opti	onal	http	The protocol used for communication between NetScaler CPX and NetScaler Ingress Controller.
nsSvcLbDnsRec	Opti	onal	false	Set this parameter to <b>true</b> to enable DNS address record addition in NetScaler through the type LoadBalancer service.
prometheusCredentialS	ecret Opti	onal	N/A	The secret key required to create a read-only user for native export of metrics using Prometheus.
pullPolicy	Mand	atory	lfNotPresent	The NetScaler CPX image pull policy.
service.annotations	Opti	onal	N/A	Dictionary of annotations to be used in the CPX service. For example, see this.

Parameter	Mandatory/Optional	Default value	Description
service.spec	Optional	N/A	Specification settings
·	·	·	of NetScaler CPX. To
			expose the CPX service
			as a NodePort, specify
			service.spec.
			type as NodePort,
			and specify a secret
			resource name for the
			front-end server
			certificate using the
			service.citrix.
			com/secret
			annotation under
			service.
			annotations.For
			example,
			service.citrix.
			com/secret: tls
			-secret. The note
			following this table
			provides an example
			of how to specify the
			required fields in the
			deployment YAML to
			expose the CPX service
			as a NodePort.
startupProbe	Optional	N/A	Startup probe settings
			for NetScaler CPX.

## Note

To expose the CPX service as a NodePort, specify the service type as NodePort and provide the secret resource name for the front-end server certificate in the CPX-with-NSIC deployment YAML as following:

```
1 service:
2 annotations:
3 service.citrix.com/secret: tls-secret
4 spec:
```

#### 5 type: NodePort

8. After updating the values for the required parameters, click **Create**. Ensure that the NetScaler CPX with Ingress Controller is successfully deployed and initialized.

Installed Operator 3 Coperator details           NetScaler Operator         312 provided by NetScaler								
Details YAML Subscription Even	ts All instances Netscaler Cloud Contro	oller Netscaler CPX With Ingress Contro	ller Netscaler GSLB Controller N	Netscaler Ingress Controller Netscaler Obse	vability Exporter			
NetscalerCpxWithIngressControllers Show operands in * All namespaces offy								
Name   Search by name  /								
Name I	Kind I	Namespace I	Status I	Labels I	Last updated I			
(NGW) netscalercpxwithingresscontroller-sample	NetscalerCpxWithIngressController	NS default	-	app.kubernetes.io/managed-by*kustomize	🔮 Dec 3, 2024, 11:31 AM	1		

- 9. Attach privileged security context constraints to the service account of NetScaler CPX (as it runs as a privileged pod) by using the following command:
  - Get the service account name used by NetScaler CPX using the following command in the namespace where NetScaler CPX has been deployed: oc get sa
  - Attach privileged SCC to the service account of NetScaler CPX:

```
1 oc adm policy add-scc-to-user privileged system:serviceaccount:
    namespace>:<CPX-ServiceAccount-Name retrieved in the previous
    step>
```

10. Navigate to **Workloads > Pods** section and verify that the netscaler-cpx-with-ingress -controller pod is up and running.

<b>`</b>	Pode					
~	1003					
	▼ Filter ▼ Name ▼ Search by name	Z III				
	Name †	Status 1	Ready 1	Restarts 🗍	Owner 1	Memory 1
~	Particular controller-sample- Scc57d8576-5kiz2	C Running	2/2	0	RS netscalercpxwithingresscontroller-sample- 5cc57d8576	616.9 MiB
	P netscaler-operator-controller-manager-	C Running	2/2	0	RS netscaler-operator-controller-manager-	80.7 MiB
	68ccfdb744-m5bd4				68ccfdb744	
	> ~ ~	Pods T Filer • Name • Search by name Name 1  netscaler-operator-controller-namger- Sector/SFXF-Sly2  of metscaler-operator-controller-namager- Sector/SFXF-Sly2	Pods  Filter  Po	Name     Search by name       T     Filter       Name     1       Status     1       Ready     1       Image: College Source Sou	Name     Name     Starts     Resdy     Restarts       Name     1     Starts     Resdy     Restarts       Image: Starts     2     0       Image: Starts     C     Running     2/2     0       Image: Starts     C     Running     2/2     0       Image: Starts     C     Running     2/2     0	Name       Search by name       I       Ready       Ready       Restarts       Owner       I         Name       1       Status       1       Ready       I       Owner       I         Image:

- 11. Verify the deployment by sending traffic.
  - a) Obtain the NodePort details using the following command:

1 oc get svc

b) Use cpx-service NodePort and send the traffic as shown in the following command:

The preceding curl command should return the following output:

```
1 <html><body><h1>It works!</h1></body></html>
```

# Note:

When you delete an NSIC instance, the ClusterRole and ClusterRole binding associated with the service account of the NSIC instance are not deleted automatically. You must manually delete the ClusterRole and ClusterRole binding associated with the service account of the deleted NSIC instance.

# References

- For information about how to deploy NetScaler Observability Exporter using NetScaler Operator, see Deploy NetScaler Observability Exporter using NetScaler Operator.
- For information about how to deploy NetScaler agent operator, see Install a NetScaler agent operator using the OpenShift console.
- Alternatively, you can deploy NetScaler Ingress Controller using Helm charts. See Deploy the NetScaler Ingress Controller using Helm charts.

# Deploy NetScaler Cloud Controller using NetScaler Operator

November 27, 2024

# Introduction

NetScaler Cloud Controller creates and maintains route entries for pod networks in the appropriate Azure route table.

Note:

NetScaler Cloud Controller only creates Azure routes between NetScaler and the Azure Red Hat OpenShift (ARO) cluster that resides in the same resource group.

This section provides information about deploying, configuring, and managing NetScaler Cloud Controller within Azure and ARO environments.

# Prerequisites

1. An active Azure subscription.

Note down the **subscription ID**.

- 2. An Azure account with sufficient permissions to create and manage resources. You must have access to create resources within the resource group and manage service principals.
- 3. Note down the **TENANT\_ID**, **CLIENT\_ID**, and **CLIENT\_SECRET** for your Azure account.
- 4. A running ARO cluster in Azure (version 4.11 or later).
- 5. The ARO cluster and the subnet to which the route needs to be established must be in the same VNet.

Note down the following details: **Resource Group Name**, **VNet Name**, and **Subnet Name** and Azure location where the ARO cluster and NetScaler are located.

6. Create a secret as described in the namespace where you want to deploy NetScaler Cloud Controller.

```
1 oc create secret generic azsecret --from-literal=clientid="<
    CLIENT_ID>" --from-literal=clientsecret="" <CLIENT_SECRET> " -n
    <namespace>
```

- 7. Create a service principal and assign a custom role to the service principal. For information, see Create service principal and assign a role.
- 8. Deploy NetScaler Operator. For information on how to deploy NetScaler Operator, see Deploy NetScaler Operator.

# Create a service principal and assign a role

Create a service principal and assign a role to it to set up the authentication details for the NetScaler Cloud Controller to manage Azure resources. There are two scenarios to consider while creating a service principal and assigning a role to it described as the following.

## Case I: The subnet on which the routes are to be created is not part of the ARO cluster

Follow the steps to create a service principal with the desired roles or add the roles to an existing service principal.

Note:

If a service principal exists, skip to step 9 to assign a role to the service principal.

- 1. Log in to Microsoft Azure Portal.
- 2. Go to the App registrations service.
- 3. Click New Registration.

Dashboard > App registrations >
Register an application
" Name
The user-facing display name for this application (this can be changed later).
netscaler-cloud-controller-sp 🗸
Supported account types
Who can use this application or access this API?
<ul> <li>Accounts in this organizational directory only (Citrix only - Single tenant)</li> </ul>
Accounts in any organizational directory (Any Microsoft Entra ID tenant - Multitenant)
Accounts in any organizational directory (Any Microsoft Entra ID tenant - Multitenant) and personal Microsoft accounts (e.g. Skype, Xbox)
O Personal Microsoft accounts only
Help me choose
Redirect URI (optional)
We'll return the authentication response to this URI after successfully authenticating the user. Providing this now is optional and it can be changed later, but a value is required for most authentication scenarios.

- 4. Enter the details such as name and supported account types.
- 5. Click **Register** to create your service principal.

Note down the **clientID** from your service principal details.

職 netscaler-cloud-controller-sp 🛷 …					
P Search	« 🗐 Delete 🌐 Endpoints 🖽 Preview features				
Cverview	Got a second? We would love your feedback on M	icrosoft identity platform (previously Azure AD for developer). $ ightarrow$			
Guickstart					
💉 Integration assistant					
Manage	Display name netscaler-cloud-controller-sp	Client credentials Add a certificate or secret			
Branding & properties	Application (client) ID	Redirect URIs			
Authentication	b84ab451-fb5d-4107-b731-4867cc61fcae	Add a Redirect URI			
Certificates & secrets	Object ID c1c2e435-9a0a-4ec6-82a6-f5c032c50476	Application ID URI Add an Application ID URI			

6. Click **Certifications & Secrets** in the left pane and click **New client secret**.

1	netscaler-cloud-controller-sp   Certificates & secrets 🛷 … 🛛 🗙 ×							
٩	Search	주 Got feedback?						
	Overview							
Ouickstart     Credentials enable confidential applications to identify themselves to the authentication service when receives a service when receives the authentication devices and applications of the service and the								
*	Integration assistant	(instead of a client secret) as a credential.						
Ma	inage							
-	Branding & properties	(1) Application registration certificates, secrets and federated credentials can be found in the tabs below.						
0	Authentication							
Ŷ	Certificates & secrets	Certificates (0) Client secrets (0) Federated credentials (0)						
1.1	Token configuration	A secret string that the application uses to prove its identity when requesting a token. Also can be referred to as						
9	API permissions	application password.						
•	Expose an API	+ New client secret						
1	App roles	Description New client secret Expires Value O Secret ID						

7. Enter a description and choose an expiry.

Add a client secret		×
Description	netscaler cloud controller	
Expires	Recommended: 180 days (6 months)	$\sim$

8. Note down the secret value. This value is **clientSecret**.

Certificates (0)	Client secrets (1)	Federated credentia	ls (0)			
A secret string that the application uses to prove its identity when requesting a token. Also can be referred to as application password.						
+ New client secret						
Description	Expires	Value 🛈		Secret ID		
netscaler cloud c	ontroller 29/09/2	cOn8Q~r	nmckTLyzbVNJX 🗋	8143a0f2-c1f3-4308-9 🗈 🗎	1	

9. Go to **Resource Group → <Your Resource Group > → Access Control (IAM)**.

Select the resource group where the ARO cluster and the VNet exist.

10. Click **Add** and select **Add custom role**.



11. Add the following custom.JSON sample.

```
{
1
2
3
   "properties": {
4
5
       "roleName": "NetScaler Cloud Controller Role",
       "description": "Allows managing cloud controller role",
6
       "assignableScopes": [
7
8
            "/subscriptions/<SUBSCRIPTION_ID>/resourceGroups/<
               RESOURCE_GROUP_NAME>"
9
       ],
       "permissions": [
10
            {
11
12
13
                "actions": [
                    "Microsoft.Network/routeTables/routes/write",
14
                    "Microsoft.Network/routeTables/routes/read",
15
                    "Microsoft.Network/routeTables/routes/delete",
16
                    "Microsoft.Network/routeTables/delete",
17
18
                    "Microsoft.Network/routeTables/write",
19
                    "Microsoft.Network/routeTables/read",
                    "Microsoft.Network/virtualNetworks/subnets/read",
21
                    "Microsoft.Network/virtualNetworks/subnets/write"
22
                ],
                "notActions": [],
23
                "dataActions": [],
24
25
                "notDataActions": []
             }
27
28
       ]
29
    }
31
    }
```

12. Modify the custom\_role.json file as per your desired settings.

You must add your **Resource Group** and **Subscription ID** in the JSON file.

13. Select **Start from JSON** in **Baseline Permissions** and upload the modified custom\_role. j son file.

Basics Permissions	Assignable scopes JSON Review + create						
To create a custom role for Azure resources, fill out some basic information. Learn more							
Custom role name * ③	NetScaler Cloud Controller Role	~					
Description	Allows for managing cloud controller role						
Baseline permissions () Clone a role () Start from scratch () Start from JSON							
File *	"custom_role json"	Ч					

- 14. Click **Review + Create**.
- 15. Go to Add role assignment.

+ Add $\sim$	→	Dow	nload ro	ole assig	nments	≡≡	Edit colum	ns 💍	Refresh	I	×	Remove	I	\$ Feedback	
Add role assig Add co-admir Add custom r	gnmei nistr role	nt Add ro	le assign	ment	Roles	D	eny assign	ments	Class	ic a	dmii	nistrators			

- 16. Search for the custom role created in **step 14**.
- 17. Select the role, and press **next**.

Role Members	Role Members Conditions Review + assign							
A role definition is a collection of permissions. You can use the built-in roles or you can create your own custom roles. Learn more of								
Job function roles Privileged administrator roles								
Grant access to Azure resources based on job function, such as the ability to create virtual machines.								
		X Type : All	Category : All					
Name 14	Description $\uparrow \downarrow$		Type ↑↓	Category ↑↓	Details			
NetScaler Cloud Contr	Allows for managing cloud contro	ller role	CustomRole	None	View			

18. Under Members, select the service principal created in step 5.

Add role assig	Add role assignment ···· ×							
Role <u>Members</u> C								
Selected role NetScaler Cloud Controller Role								
Assign access to	to User, group, or service principal Managed identity							
Members	+ <u>Select members</u>							
	Name	Object ID	Туре					
	netscaler-cloud-controller-sp	115dc1cf-1283-4ec3-ba60-c704	App	1				
Description	Optional							

19. Click Review + Assign.

#### Case II: The subnet on which the routes are to be created is part of the ARO cluster

In this case, you must have access to the service principal created during the OpenShift installation. The automatically created service principal during the ARO installation has a **aro-app-** prefix.

- 1. Go to **Resource Group** → **<Your Resource Group** > where the ARO cluster and the VNet exist. Go to **Resource Group** → **<Your Resource Group** > → **Access Control (IAM)**
- 2. Click Add and select Add custom role.



3. Copy the following in a JSON file named custom\_role.json.

```
],
10
        "permissions": [
            {
11
12
13
                "actions": [
14
                    "Microsoft.Network/routeTables/routes/write",
                    "Microsoft.Network/routeTables/routes/read",
15
16
                    "Microsoft.Network/routeTables/routes/delete",
                    "Microsoft.Network/routeTables/delete",
17
18
                    "Microsoft.Network/routeTables/write",
                    "Microsoft.Network/routeTables/read",
19
                    "Microsoft.Network/virtualNetworks/subnets/read",
20
21
                    "Microsoft.Network/virtualNetworks/subnets/write"
                ],
                "notActions": [],
23
                "dataActions": [],
24
25
                "notDataActions": []
             }
26
27
28
       ]
29
    }
31
    }
```

4. Modify the custom\_role.json file as per your desired settings.

You must add your **Resource Group** and **Subscription ID** in the JSON file.

5. Select **Start from JSON** in **Baseline Permissions** and upload the modified custom\_role. j son file.

	Assignable scopes JSON Review + create						
To create a custom role for Azure resources, fill out some basic information. Learn more							
Custom role name * 🛈	NetScaler Cloud Controller Role	~					
Description	Allows for managing cloud controller role						
Baseline permissions () Clone a role () Start from scratch () Start from JSON							
File *	"custom_role_json"	В					

- 6. Click **Review + Create**.
- 7. Click Add role assignment.

+ Add $\sim$	₹	Dow	nload	l role	assig	nments	≡≡	Edit co	lumns	Ö	Refresh	I	$\times$	Remove		Ŕ	Feedback	
Add role assis Add co-admi Add custom r	gnm nistr role	ent Add r	ole ass	ignme	nt	Roles	D	Deny ass	signme	ents	Class	ic a	dmir	istrators				

8. Search for the custom role created in step 4. Select the role, and click next.



9. In the **Members** section, select the service principal that is managing the OpenShift cluster.

Add role assig	nment				×		
Selected role NetScaler Cloud Controller Role							
Assign access to	User, group, or service principal     Managed identity	al					
Members	+ <u>Select members</u>						
	Name	Object ID	Туре				
	netscaler-cloud-controller-sp	115dc1cf-1283-4ec3-ba60-c704	Арр	Î			
Description	Optional						

10. Click Review + Assign.

#### Steps to deploy NetScaler Cloud Controller instance

- 1. Log in to the 4.x OpenShift console.
- 2. Navigate to **Operators > Installed Operators** and select **NetScaler Operator**.
- 1. Go to NetScaler Operator, click the NetScaler Cloud Controller tab, and then click **Create NetScaler Cloud Controller**.

Installe	d Operators	Operator de	etails				
>	3.1.2 provid	er Operator ded by NetScal	ler			Actions	•
Subs	cription	Events	All instances	Netscaler Cloud Controller	Netscaler CPX With Ingress Controller	Netscaler GSLB Controller	Net
Net	scalerC	loudCor	ntrollers <sup>Sh</sup>	now operands in: <ul> <li>All namespaces</li> </ul>	O Current namespace only	Create NetscalerCloudControlle	er

The YAML for NetScaler Cloud Controller CRD is displayed. Optionally, you can select Form view button and update the parameters in a form.

2. Edit the .YAML file to add the following values obtained in the prerequisites section.



## Refer to the following table for the descriptions of the fields in the displayed YAML.

Parameter	Mandatory/Optional	Default value	Description
name	Mandatory	netscalercloudcont	roller-Name of the NetScaler
		sample	Cloud Controller
			instance.
acceptLicense	Mandatory	no	Set this value to Yes
			to accept the NetScaler
			Cloud Controller EULA.
affinity	Optional	N/A	Affinity labels for pod
			assignment.

# NetScaler ingress controller

Parameter	Mandatory/Optional	Default value	Description
nodeSelector	Optional	N/A	The node label key:value pair to be used for nodeSelector option in NetScaler Cloud Controller deployment.
podAnnotations	Optional	N/A	Map of annotations to add to the pods.
tolerations	Optional	N/A	The tolerations for the NetScaler Cloud Controller deployment.

Parameter	Mandatory/Optional	Default value	Description
	azure		
clientsecret	Mandatory	N/A	Secret name created in the prerequisites section.
image	Mandatory	quay.io/	The NetScaler Cloud
		netscaler/	Controller image
		netscaler-cloud	hosted on Quay.io.
		-	
		controller@shava	lue
		of latest	
		release	
location	Mandatory	N/A	Azure location.
pullPolicy	Optional	IfNotPresent	The NetScaler Cloud Controller image pull policy.
resourcegroupname	Mandatory	N/A	Resource group name where the VNet exists.
resources	Optional	N/A	CPU/memory resources for a NetScaler Cloud Controller container.

Parameter	Mandatory/Optional	Default value	Description
subnetname	Mandatory	N/A	Azure VNet subnet name from where the route is to be established to the ARO cluster. This subnet has to be in the same VNet as the same ARO Cluster.
subscriptionid	Mandatory	N/A	Subscription ID for your Azure account.
tenantid	Mandatory	N/A	Tenant ID of the Azure service principal.
vnetname	Mandatory	N/A	Azure VNet name in which the required subnet and the OpenShift Cluster exists.

- 3. After updating the values for the required parameters, click **Create**.
- 4. Navigate to the **Workloads > Pods** section and verify whether the NetScaler Cloud Controller pod is up and running.

After the controller pod is up and running, the required route table with the subnet is created.

Notes:

- The controller creates a route table and associates it with the given subnet if the route table does not exist. The format of the new table will be <SUBNET\_NAME>\_<LOCATION>\_rt.
- To avoid conflicts with routes created/updated by the NetScaler Cloud Controller from other automation processes or users, the routes managed by NetScaler Cloud Controller start with the prefix **NSCC\_**, followed by a hash value for tracking. Do not create routes with the same name prefixes to avoid conflicts.

The following roles are installed when NetScaler Cloud Controller is deployed. NetScaler Cloud Controller requires these roles to manage its CRD and get node-specific information.

API Group	Resources	Verbs
rbac.authorization.k8s.io	clusterrolebindings	create, delete, get, list, patch, update, watch
rbac.authorization.k8s.io	clusterroles	create, delete, get, list, patch, update, watch
core	nodes	All
core	serviceaccounts	create, delete, get, list, patch, update, watch
apps	deployments	create, delete, get, list, patch, update, watch
core	pods	create, delete, get, list, patch, update, watch, serviceaccounts

# Deploy NetScaler GSLB Controller in OpenShift using NetScaler Operator

## November 26, 2024

NetScaler Operator enables you to deploy NetScaler GSLB Controller in an OpenShift cluster. For information about GSLB, see GSLB overview and deployment topologies.

# **Deploy NetScaler GSLB Controller**

## Prerequisites

- Red Hat OpenShift Cluster (version 4.11 or later).
- Deploy NetScaler Operator. For information on how to deploy NetScaler Operator, see Deploy NetScaler Operator.

## Steps to deploy NetScaler GSLB Controller using NetScaler Operator

Do the following steps:

1. Navigate to **Operators > Installed Operators** and select **NetScaler Operator**.

2. Click the NetScaler GSLB Controller tab and click Create NetscalerGslbController.

Installed Operators > Operator details								
>	3.1.2 provide	Operator ed by NetScaler				Actions 💌		
ption	Events	All instances	Netscaler Cloud Controller	Netscaler CPX With Ingress Controller	Netscaler GSLB Controller	Netscaler Ingr		
NetscalerGslbControllers Show operands in:  All namespaces O Current namespace only Create NetscalerGslbController								

The NetScaler GSLB Controller instance YAML definition is displayed. Optionally, you can select Form view button and update the parameters in a form.

3. Update the values for the required parameters in the displayed NetScaler GSLB Controller instance YAML and click **Create**. For information on the parameters that you need to configure, see the following table.

# Create NetscalerGslbController

Create by completing the form. Default values may be provided by the Operator authors.

Configure via: O Fo	rm view 💿 YAML view	
	Alt + F1 Accessibility help	View shortcuts Show tooltips
1 kind: Nets	calerGslbController	All Annual An Annual Annual Annua
2 apiVersior	: netscaler.com/v1	
3 metadata:		
4 labels:		
5 app.ku	ubernetes.io/managed-by: kustomize	
6 app.ku	<pre>ubernetes.io/name: netscaler-operator</pre>	
7 name: ne	etscalergslbcontroller-sample	
8 namespac	e: default	
9 spec:		
10 affinity	/: {}	
11 acceptLi	cense: 'No'	
12 gslbCont	roller:	
13 disabl	eAPIServerCertVerify: false	
14 entity	Prefix: 🛄	
15 image:	<pre>'quay.io/netscaler/netscaler-k8s-ingress-cont</pre>	roller@sha256:2ff3865447
16 kuberr	netesURL: ''	-
17 local	luster: 🛄	
18 local	legion: 🛄	
19 logLev	vel: Debug	
20 opensh	ift: true	
21 pullPc	licy: IfNotPresent	
Create	el	🛓 Download
Parameter	Description	Mandatory/Optional
	gslbController	
image	The NetScaler GSLB Controller image.	Mandatory

Parameter	Description	Mandatory/Optional
pullPolicy	The NetScaler GSLB Controller	Optional
	image pull policy.	
imagePullSecrets	List of OpenShift secrets to be	Optional
	used for pulling the images	
	from a private Docker registry	
	or repository. For more	
	information on how to create	
	this secret, see Pull an Image	
	from a Private Registry.	
entityPrefix	The prefixes for the resources	Mandatory
	on NetScaler VPX or NetScaler	
	MPX.	
acceptLicense	Set yes to accept the GSLB	Mandatory
	Controller end user license	
	agreement.	
logLevel	The log level to control the logs	Optional
	generated by NetScaler GSLB	
	Controller. The supported log	
	levels are: CRITICAL, ERROR,	
	WARNING, INFO, DEBUG,	
	TRACE, and NONE. For more	
	information, see Logging.	
openshift	Set this argument if the	Mandatory
	OpenShift environment is	
	being used.	
isableAPIServerCertVerify	Set this parameter to true for	Optional
	disabling API Server certificate	
	verification.	
kubernetes URL	The kube-apiserver url that	Optional
	NetScaler GSLB Controller uses	
	to register the events. If the	
	value is not specified, NetScaler	
	GSLB Controller uses the	
	internal kube-apiserver IP	
	address.	

Parameter	Description	Mandatory/Optional
cleanupGSLBSiteConfig	Set this parameter to true to clean up the GSLB site configuration.	Optional
LocalRegion	Local region where the GSLB controller is deployed.	Mandatory
LocalCluster	The name of the cluster in which the GSLB controller is deployed. This value is unique for each OpenShift cluster.	Mandatory
siteData[0].siteName	The name of the first GSLB site configured in the GSLB device.	Mandatory
siteData[0].secretName	The name of the secret that contains the login credentials of the first GSLB site.	Mandatory
siteData[0].siteIp	IP address for the first GSLB site. Add the IP address of the NetScaler in site1 as <i>siteData[0].siteIp</i> .	Mandatory
sitedata[0].siteRegion	The region of the first site.	Mandatory
sitedata[0].siteMask	The netmask of the first GSLB site IP address.	Mandatory
sitedata[0].sitePublicIp	The public IP address of the first GSLB Site.	Mandatory
sitedata[1].siteName	The name of the second GSLB site configured in the GSLB device.	Mandatory
sitedata[1].secretName	The secret containing the login credentials of the second site.	Mandatory
sitedata[1].sitelp	IP address for the second GSLB site. Add the IP address of the NetScaler in site2 as <i>siteData[0].siteIp</i>	Mandatory
sitedata[1].siteRegion	The region of the second site.	Mandatory
sitedata[1].siteMask	The netmask of the second GSLB site IP address.	Mandatory
sitedata[1].sitePublicIp	The public IP address of the second GSLB Site.	Mandatory

© 1997–2025 Citrix Systems, Inc. All rights reserved.

Parameter	Description	Mandatory/Optional
	netscaler	
nsIP	NetScaler IP address.	Mandatory
adcCredentialSecret	Secret required for the GSLB controller to connect to GSLB devices and push the configuration from the GSLB	Optional
nsProtocol	The protocol used by NetScaler GSLB Controller to communicate with NetScaler. You can also use HTTP on port 80.	Optional
nsPort	The port used by NetScaler GSLB Controller to communicate with NetScaler. You can use port 80 for HTTP.	Optional
nitroReadTimeout	The duration, in seconds, that NetScaler GSLB Controller has to wait to receive a response from NetScaler before terminating the connection.	Optional
arameter	Mandatory/Optional	Description
ame	Mandatory	Name of the NetScaler GSLB

l'uluineter	Mandatory, optionat	Description
name	Mandatory	Name of the NetScaler GSLB
		Controller instance.
acceptLicense	Mandatory	Set this value to Yes to accept
		the NetScaler GSLB Controller
		EULA.
affinity	Optional	Affinity labels for pod
		assignment.
nodeSelector	Optional	The node label key:value pair
		to be used for nodeSelector
		option in NetScaler GSLB
		Controller deployment.

Parameter	Mandatory/Optional	Description
podAnnotations	Optional	Map of annotations to add to
		the pods.
tolerations	Optional	The tolerations for the
		NetScaler GSLB Controller
		deployment.

## Notes:

- If you want to configure more than two GSLB sites, include the details of all the GSLB sites in the siteData list as required.
- You must specify either adcCredentialSecret or the combination of secret-Store.username and secretStore.password to provide NetScaler VPX or NetScaler MPX credentials in the GSLB Controller instance YAML.
- 4. Navigate to the **Workloads > Pods** section and verify whether the NetScaler GSLB Controller pod is up and running.

Home >	Installe	d Operators 🔸 🕻	Operator details									
Operators 🗸	>	NetScaler O 3.11 provided b	Operator by NetScaler									Actions •
OperatorHub	Detai	is YAML	Subscription Event	s All instance	s Netscaler Cloud Controller	Netscaler CPX With Ingress	Controller Ne	tscaler GSLB Controller	Netscaler Ingress Controller	Netscaler Observability Expo	ter	
Installed Operators												
Workloads 🗸	Net	scalerGsli	bControllers								Cre	ate NetscalerGslbController
Pods Deployments	Name	• • Search by	name									
DeploymentConfigs	Na	me I		Kind		Status I		Labels I		Last updated		
StatefulSets	6	C netscalergsl	bcontroller-sample	Netso	calerGslbController			app.kuberne	rtes.lo/managed-by=kustomize	Nov 20, 2024, 2:54 Pf	(	1
ConfigMaps	-							app.kuberne	tes.xo/name*netscaler-operator			
E SpenShif	ft									<b>≜</b> 1 <b>€</b>	😯 kut	be:admin <del>-</del>
🌣 Administrator					You are logged	in as a temporary admir	istrative user	. Update the <u>cluster</u>	OAuth configuration to	allow others to log in.		
			Project: demo	•								
Home		>									_	
			Pods									Create Pod
Operators		*										
OperatorHub			▼ Filter ▼	Name	<ul> <li>Search by name</li> </ul>							
Installed Operators												
instance operators			Name 1		Status 🔱	Ready 🗍	Restarts	1 Owner 1	Memor	ry î CPU î	Created 1	
Workloads		•	P nsgc-7f bp8zm	bcc5f97d-	2 Running	1/1	0	RS nsgc-7fb	cc5f97d -	-	1 Aug 202 10:45	4, 1
Pods												
Deployments												
DeploymentConfigs												
StatefulSets												
Secrets												
ConfigMaps												
CronJobs												
Jobs												
DaemonSets												
ReplicaSets												
ReplicationController	<u>د</u>											
ReplicationControllers	s											
ReplicationControllers HorizontalPodAutosca	s alers											

Skip to References.

# References

• For information about deploying global traffic policy (GTP) and global service entry (GSE), see GTP and GSE deployment.

# Deploy NetScaler Ingress Controller as a standalone pod by using NetScaler Operator

## May 2, 2025

You can deploy NetScaler IPAM Controller as a standalone pod in an OpenShift cluster for IP address management by using NetScaler Operator. When you create a service of type LoadBalancer or Ingress, you can use the NetScaler IPAM controller to automatically allocate an IP address. Once the IPAM controller is deployed, it allocates IP address to services of type LoadBalancer or Ingress from predefined IP address ranges. The NetScaler ingress controller configures the IP address allocated to the service as virtual IP address (VIP) in NetScaler MPX or VPX that is deployed as an ingress device or router for an application running in the OpenShift cluster. The following diagram explains the topology:



## Prerequisites

The following are the prerequisites for deploying NetScaler Ingress Controller as a standalone pod using NetScaler Operator:

- Have Red Hat OpenShift Cluster (version 4.11 or later).
- Deploy NetScaler Operator. For information on how to deploy NetScaler Operator, see Deploy NetScaler Operator.
- Deploy NetScaler Ingress Controller to configure service or ingress on NetScaler VPX for which NetScaler IPAM is allocating the IP address. For information on how to deploy NetScaler Ingress Controller, see Deploy NetScaler Ingress Controller.

## **Deployment steps**

- 1. Log in to the **OpenShift 4.x Cluster** Console.
- 2. Allocate IP to an Ingress Using **NetScaler IPAM Controller**.
- 3. Allocate IP for an Ingress for Apache application.
  - a) Navigate to Networking > Ingress > Create Ingress.
  - b) Use the YAML (like the following one) to create the ingress. Ensure that you have the required *annotation* for IPAM, which allocates an IP address for the ingress. This IP address gets configured as the VIP address on NetScaler VPX by the NetScaler Ingress Controller.

```
1
  ___
2 apiVersion: networking.k8s.io/v1
3 kind: Ingress
4 metadata:
5 annotations:
6
       # IPAM uses the below annotation to select the IP Range,
          from which it will allocate IPs
7
       ingress.citrix.com/ipam-range: ingress-range
8 name: apache-ingress
9 spec:
10 ingressClassName: nsic-vpx
11 rules:
12 - host: citrix-ingress-operator.com
13
      http:
14
       paths:
15
       - backend:
16
           service:
17
               name: apache
18
               port:
19
               number: 80
20
           path: /
21
           pathType: ImplementationSpecific
22 ---
```

- 4. Allocate IP to a Service of Type LoadBalancer.
  - a) Navigate to Networking > Services > Create Service.
  - b) Use the following YAML to create the service. Ensure that you have the required *annotation* for IPAM, which allocates an IP address for the service. This IP address gets configured as the VIP address of NetScaler VPX by the NetScaler Ingress Controller.

1	
2	apiVersion: v1
3	kind: Service
4	metadata:
5	name: apache-lb
6	labels:
7	app: apache
8	annotations:
9	# CIC uses the below annotation to select services to be
	configured on NetScaler VPX/MPX
10	service.citrix.com/ <b>class:</b> 'nsic-vpx'
11	# IPAM uses the below annotation to select the IP Range,
	from which it will allocate IPs
12	service.citrix.com/ipam-range: 'service-range'
13	spec:
14	type: LoadBalancer
15	ports:
16	- name: http
17	port: 80
18	targetPort: http
19	selector:
20	app: apache
21	

5. Navigate to **Operators > Installed Operators** and click **NetScaler Operator**.

≡ <sup>♣</sup> Red Hat OpenShift							₩ <b>\$</b> 5 <b>0</b>	kuberadmin <del>v</del>
oc Administrator				You are logge	ed in as a temporary administrative user. Update the <u>cluster OAuth c</u>	onfiguration to allow others to log in.		
11		Project: c	lemo 🝷					
Home	,	Install	ed Operators					
Operators	~	Installed O	perators are represented by ClusterServiceVersions with	thin this Namespace. For more information, see the Understanding Op	perators documentation [2]. Or create an Operator and ClusterServi	ceVenion using the Operator SDK [2]		
OperatorHub								
Installed Operators		Name •	Search by name					
Workloads	*	Name		Managed Namespaces	Status	Last updated	Provided APIs	
Pods		>	NetScaler Operator	The operator is running in opershift-operators but is managing	Succeeded Up to date	Apr 24, 2025, 10:41 AM	Netscaler Cloud Controller Netscaler CPX With Ingress Controller	1
Deployments				this namespace			Netscaler GSLB Controller Netscaler Ingress Controller	
DeploymentConfigs							View 2 more	
Secrets								
ConfigMaps								
CronJobs								
Jobs								
DeemonSets								
ReplicaSets								
ReplicationControllers								
PodDisruptionBudgets								
Networking	>							
Storage	>							
Builds	•							
Observe	•							
Compute	•							
User Management	,							
Administration	,							

= <b>Red Hat</b> OpenShift						III 🌲 5 🗿 kubesadmin 🗸
Administrator *			You are logged in as a temporar	y administrative user. Update the <u>cluster OAuth.com</u>	figuration to allow others to log in.	
	Project: demo 🔻					
	Installed Operators > Operator details					î
Operators 👻	> NetScaler Operator 3.2.0 provided by NetScaler					Actions 💌
OperatorHub	Datala MAMI Coloristica S	Surata All instances Materiales (1	and Casterline Materials CDV With In-	Natural Controller Natural COLD Con	stanling – Ninternales in service Constanting	Natural In IDAM Controller - Natural in Observability Forester
Installed Operators	Jetais PARE Subscription E	Sterits Antistances Netscaler G	Netscaler CFX Mitring	pear controller Tretacaler COLD Cor	noner Hetacaler ingreas controller	Neticale in Am Controller - Neticaler Observability Exporter
Workloads 🗸	Provided APIs					Provider NetScaler
	Netscaler Cloud Controller	Netscale: CPX With Ingress	NGC Netscaler GSLB Controller	No Netscaler Ingress Controller	Netscaler IPAM Controller	Created at
Deployments	To install NetScalar Cloud Controllar for	Controller	To install NatScalar GSI B Controllar for	To instal NetScalar Increas Controller	To install NatScalar IPAM Controllar for	Apr 24, 2025, 10:40 AM
DeploymentConfigs	NetScaler VPX/MPX/SDX with Azure	To install NetScaler CPX with Ingress	NetScaler VPX/MPX/SDX	for NetScaler VPX/MPX/SDX.	NetScaler CPX/VPX/MPX/SDX.	Links
StatefulSets	Redhat OpenShift in Azure.	Controller.				Netscaler Operator https://docs.netscaler.com/en-us/netscaler-k0s-ingress-controlley/deploy/netscal
	Create instance	Create instance	Create instance	Create instance	Create instance	er-operator-openshift gf
ConfigMaps				· · · · · · · · · · · · · · · · · · ·		Maintainers
						Priyanka Sharma privanka sharmaticloud.com
CronJobs	Netscaler Observability Exporter					
	To install NetScaler Observability					Arijit Ray arijit raymoloud.com
DeemonSets	Exporter for NetScaler.					
ReplicaSets	Counte instance					
ReplicationControllers	Come interior					
HorizontalPodAutoscalers						
PodDisruptionBudgets	Description					
	NetScaler provides various products to empower	Developer/DevOps managing your applications in	OpenShift Cluster. Using this operator you can deplo	y below controllers for NetScaler appliance:		
Networking >	1. NetScaler Cloud Controller					
Storana	2. NetScaler Ingress Controller					
, and the second s	3. NetScaler CPX with Ingress Controller					
Builds >	4. NetScaler GSLB Controller					
	5. NetScaler Observability Exporter 6. NetScaler IDAM Controller					
Observe >	C. HELCON I AN CONDUCT					
Comoute >	Installation					
	Refer installation instructions.					
User Management >	This operator version contains:					
Administration	1. NetScaler Ingress Controller version 2.2.K					
	2. NetScaler CPX version 141-25111 2. NetRadia CRI R. Castralia version 2.210					

6. Click the **NetScaler IPAM Controller** tab and select **Create NetScalerIPAMController**. The NetScaler IPAM Controller YAML definition is displayed.

Optionally, you can select the **Form view** button and update the parameters in a form. Refer to the Mandatory and Optional Parameters section that lists the mandatory and optional parameters that you can configure during installation. You can configure other available parameters depending on your use case.

		🗰 🖡 5 🗢 🕢 kubesadmin <del>-</del>
• Administrator	You are logged in as a temporary administrative user. Update the <u>cluster Davis configuration</u> to allow others to log in	
	Project demo •	
Home	Installed Openitors > Openitor deads	
Operators 👻	NetSulf Uperior	Actions 📼
OperatorHub		
Installed Operators	Details YAML Subscription Events All instances Netscaler Cloud Controller Netscaler CPX With Ingress Controller Netscaler OSLB Controller Netscaler Ingress Controller	ntroller Netscaler Observability Exporter
Workloads 🗸	Netscaler/pamControllers Show operands in: * Ali nomespaces O Current namespace only	Create Netscaler/parr/Controller
Deployments	No operands found	
DeploymentConfigs	Operands are declarative components used to define the behavior of the	
StatefulSets	application.	
Secrets		
ConfigMaps		
DaemonSets		
ReplicaSets		
ReplicationControllers		
HorizontalPodAutoscalers		
PodDisruptionBudgets		
Networking >		
Storage >		

Red Hat OpenShift			III 🌲 5 🗢 O kubezadmin 🗸
et Administrator	You are logged in as a temporary administrative user. Update the <u>cluster Okuth confi</u>	figuration to allow	others to log in.
	Project demo 🔹		
Home	Create Netscaler/pamController		
Operators	Create by completing the form. Default values may be provided by the Operator authors.		
OperatorHub	Configure via: 🔿 Form view 🔹 VAML view		
Installed Operators			
Workloads	1     kind: wetscalerigaeKontroller     0     View shortcuts     0       3     wetscalerigaeKontroller     0     View shortcuts     0	Show tooltips	NetscalerlpamController X
Networking	3 estadata: 4 labela: 5 uno bioexeter folgement he jurtanta		Schema
	app.kubernetex.io/name: netscale-operator     pose_net_cale_operator		NetscaleriparnController is the Schema for the netscaleriparncontrollers API
Routes	Instructure a generative agenerative agen		aplVersion     string
Ingresses	10 acceptions: 'Yes'		APWersion defines the versioned schema of this representation of an object. Servers should convert recognized schemas to
NetworkPolicies	clusterlane: '     clusterlane: '     sage: 'aum/io/retscaler/netscaler-izam-controllen804256;7(205902cidf0e475615)7790/bca039c003eca05566(20556/0008318)'		the store memory value, and may reject intercognized values, receipting https://git.kbs.io/community/contributos/devel/sig-architecture/spi-conventions.md#resources
Storage	14 local1999: 15   vijskneg: [['ingress-range': ['i.i.i.0/26']], ("Service-range': ['2.2.2.1-2.2.2.19']]]'	-	• kind
Builds	17 podemotation: () 18 pilleolicy: (Bittereset		string Kind is a string value representing the REST resource this object represents. Servers may infer this from the endpoint the
Observe	19         resubliquessifie: true           20         tabletiginationsconds: 3153000           21         tabletiginations: []		client submits requests to Cannot be updated. In CamelCase. More info https://git.k8sio/community/contributors/devel/sig-architecture/spi-conventions.md#types-kinds
Compute	22		• metadata object
User Management			Standard object's metadata. More info: https://git.k8sio/community/contributors/devel/sig-architecture/api- conventions.mdWmetadata
Administration			View details
			• spec object
			NetscalerlpamControllerSpec defines the desired state of NetscalerlpamController
			View details
			status     shart
			Netscaler(pamControllerStatus defines the observed state of Netscaler(pamController
			View details
	Conte	*Download	

#### Notes:

- Ensure that the acceptLicense parameter is set to Yes.
- It is mandatory to provide values for the parameters on localIPAM or Infoblox.
- Provide clusterName if you want to use Infoblox.
- 7. After updating the values for the required parameters, click Create
- 8. Ensure that NetScaler IPAM Controller is successfully deployed and initialized. Navigate to **Workloads > Pods section** and verify whether NetScaler IPAM Controller pod is up and running.
- 9. Verify the deployment by checking the logs. Also, for **ingress**, you can check the ingress status to confirm that the IP address is allocated to ingress by NetScaler IPAM Controller.

## **Ingress Status**



**For service type LB**, you can check the service External-IP field to confirm that IP is allocated to the service by NetScaler IPAM Controller.

## **Service Status**

```
1 root@master:~/Citrix/openshift# oc get svc
```

2	NAME	TYPE	CLUSTER-I	Р	EXTERNAL-IP	PORT(S)
		AGE				
~						
3	apache-lb	LoadBalancer	172.30.24	3.206	2.2.2.2	
	00.22400					
	00.32490	5/ICF 15				
4	rootAmaster	••~/Citrix/open	$shift = 0 \sigma$	et vin		
-	100 centas cer	• / cr cr r x / open		cc vip		
5	NAME		VTP	AGE		
0	10.012			/ OL		
6	service-ana	iche-1b	2.2.2.2	165		
0	ocivice apo		2.2.2.2	±00		

# Parameters for NetScaler IPAM Controller

Parameter	Mandatory/Optional	Default Value	Description			
image	Mandatory	quay.io/netscaler/netscal@he NetScaler IPAM				
		ipam-	Controller image			
		controller@sha256	:7c25b <b>d@aaie1.</b> df0e4f5615a779a9bea9390			
pullPolicy	Mandatory	IfNotPresent	The NetScaler IPAM			
			Controller image pull			
			policy.			
reuseIngressVip	Optional	True	Allows using the same			
			IP for all ingresses			
			using the same			
			vipRange.			
clusterName	Mandatory if	N/A	Identifies the cluster in			
	infoblox.enabled is		which the IPAM			
	true		controller is deployed.			
tokenExpirationSeconds	Mandatory	31536000	Time in seconds when			
•	2		the token of			
			serviceAccount			
			expires.			

# Mandatory and Optional Parameters

## Local IPAM Parameters

Parameter	Mandatory/Optional	Default Value	Description
vipRange	Mandatory	N/A	Provides the IPAM VIP
			Range.

## **Infoblox Parameters**

Parameter	Mandatory/Optional	Default Value	Description
gridHost	Mandatory if infoblox.enabled is true	N/A	Infoblox grid host IP or FQDN.
credentialSecret	Mandatory if infoblox.enabled is true	N/A	Infoblox credentials.
httpTimeout	Optional	10	Infoblox client HTTP timeout in seconds.
maxRetries	Optional	3	Infoblox client max retries on failure.
netView	Optional	default	Infoblox NetView.
vipRange	Mandatory	N/A	Infoblox IPAM VIP Range.

# Deploy NetScaler Observability Exporter (NSOE) by using NetScaler Operator

#### November 29, 2024

NetScaler Observability Exporter is a container that collects metrics and transactions from NetScaler and transforms them to suitable formats, such as JSON and AVRO. You can export the data collected by NSOE to any desired endpoint for analysis and get valuable insights at the microservices level for applications proxied by NetScaler devices.

# Prerequisites

- Red Hat OpenShift Cluster (version 4.1 or later).
- Deploy NetScaler Operator. See Deploy NetScaler Operator.
- Because NSOE operates via any User ID (uid), deploy the following security context constraints (SCC) for the namespace in which NSOE is deployed.

```
1 oc adm policy add-scc-to-user anyuid system:serviceaccount:<
    namespace>:default
```

## Deploy NetScaler Observability Exporter using NetScaler Operator

Perform the following steps:

- 1. Log in to the OpenShift 4.x Cluster console.
- 2. Navigate to **Operators > Installed Operators** and select **NetScaler Operator**.
- 3. Click NetScaler Observability Exporter and select Create NetScalerObservabilityExporter .



The NetScaler Observability Exporter YAML definition is displayed. Optionally, you can select Form view button and update the parameters in a form.

#### Create NetscalerObservabilityExporter Create by completing the form. Default values may be provided by the Operator authors. Configure via: O Form view O YAML view View shortcuts Alt + F1 Accessibility help Show tooltips kind: NetscalerObservabilityExporter apiVersion: netscaler.com/v1 metadata: 4 labels: app.kubernetes.io/managed-by: kustomize app.kubernetes.io/name: netscaler-op name: netscalerobservabilityexporter-sample 8 namespace: default 9 spec: 10 acceptLicense: 'No' 11 affinity: {} elasticsearch: 12 enabled: false 14 server: 'http://elasticsearch:9200' headless: false image: 'quay.io/netscaler/netscaler-observability-exporter@sha256:626e3a954c67t 16 jsonTransRateLimiting: 18 enabled: false 19 limit: 100 20 queueLimit: 1000 window: 5 📩 Download Create Cancel

4. Refer the following table that lists the mandatory and optional parameters and their default values that you can configure during installation.

# NetScaler ingress controller

Parameters	Mandatory or Optional	Default value	Description
acceptLicense	Mandatory	no	Set this value to Yes to
			accept the NSOE EULA.
affinity	Optional	N/A	Affinity labels for pod
			assignment.
nodePortRequired	Optional	false	Set this parameter to
			true to create a NSOE
			nodeport service.
nodeSelector	Optional	N/A	The node label
			key:value pair to be
			used for the
			nodeSelector option in
			an NSOE deployment.
podAnnotations	Optional	N/A	Map of annotations to
			add to the pods.
pullPolicy	Optional	IfNotPresent	The NSOE image pull
			policy.
replicas	Optional	1	The number of NSOE
			pods that run at any
			given time.
resources	Optional	N/A	CPU/memory
			resources for an NSOE
			container.
tolerations	Optional	N/A	The tolerations for an
			NSOE deployment.
headless	Optional	false	Set this parameter to
			<b>true</b> to create a
			headless service.
image	Mandatory	quay.io/netscaler/netsc	al <b>ē</b> he NSOE image
		observability-	hosted on Quay.io.
		exporter@shavalue of	
		latest NSOE image	

Parameters	Mandatory or Optional	Default value	Description	
nsoeLogLevel	Optional	INFO	The log level to contro the logs generated by NSOE. The supported log levels are: CRITICAL, ERROR, WARNING, INFO, DEBUG, TRACE, and NONE.	
Parameters	Mandatory or Optional	Default value	Description	
	elasticsearch			
enabled	Optional	false	Set this parameter to <b>true</b> to enable an Elasticsearch endpoint	
server	Optional	http://elasticsearch:9200	for transactions. D The Elasticsearch endpoint.	
Parameters	Mandatory or Optional	Default value	Description	
	jsonTransRateLimiting			
enabled	Optional	false	Set this parameter to <b>true</b> to enable rate-limiting of transactions for JSON-based endpoints, such as Splunk, ElasticSearch, and Zipkin.	
limit	Optional	100	The rate limit for JSON transactions. Hundred amounts to approximately 800 TPS.	

Parameters	Mandatory or Optional	Default value	Description
queueLimit	Optional	1000	The amount of transactional data that can be saved before NSOE starts shedding them. For Zipkin, 1000 amounts to approximately 64 MB of data. For Splunk and ElasticSearch, 1000 amounts to approximately 32 MB
Window	Optional	5	The recalculation window in seconds. The lower the window size (any value greater than 0), the more effective is the rate limiting. However, setting a lower window size results in CPU overhead.
Parameters	Mandatory or Optional	Default value	Description
	kafka		
auditLogs	Optional	no	Set this parameter to yes to export the audit logs to Kafka.
broker	Optional	N/A	The kafka broker IP address.
dataFormat	Optional	AVRO	The format of the data exported to Kafka. Possible values are JSON and AVRO.

Parameters	Mandatory or Optional	Default value	Description
enabled	Optional	false	Set this parameter to <b>true</b> to enable sending transaction data to the kafka
			server
events	Optional	no	Set this parameter to
			yes to export events
			to Kafka.
topic	Optional	HTTP	The kafka topic details
			to upload data.
Parameters	Mandatory or Optional	Default value	Description
	nstracing		
enabled	Optional	false	Set this parameter to
	- F		true to enable
			sending trace data to a
			tracing server.
server	Optional	zipkin:9411/api/v1/spa	ans The tracing server API
			endpoint.
Parameters	Mandatory or Optional	Default value	Description
	splunk		
authToken	Optional	N/A	The authentication
			token for Splunk.
enabled	Optional	false	Set this parameter to
			<b>true</b> to enable
			data to the Splunk
			server.
indexPrefix	Optional	N/A	The Splunk index
	•	·	prefix to upload the
			transactions.
server	Optional	N/A	The Splunk server API
			endpoint.

Parameters	Mandatory or Optional	Default value	Description
	timeseries		
enabled	Optional	false	Set this parameter to <b>true</b> to enable metrics data upload in Prometheus format. Currently, Prometheus is the only metrics endpoint supported.

## Notes:

NSOE can be deployed in multiple namespaces. Also, multiple instances of NSOE can also be deployed in the same namespace, provided the deployment name is different for each instance. Before deploying, ensure that the prerequisite any-uid SCC is deployed for the target namespace.

## 5. After updating the values for the required parameters, click **Create**.

Home	•	Installed O	perators > Op	perator details							
Operators	~	>	NetScaler Op 3.1.1 provided by	erator / NetScaler							Actions 👻
OperatorHub Installed Operators		Details	YAML	Subscription	Events	All instances	Netscaler Clo	ud Controller	Netscaler CPX V	Vith Ingress Controller	Netscaler GSL
Workloads Pods	×	Netsc	alerObs	ervabilityEx	porters					Create NetscalerObser	vabilityExporter
Deployments		Name 🔻	Search by n	name	Z						
DeploymentConfigs StatefulSets		Name	I	Kind	1	Sta	tus 1	Labels 🗘		Last updated	
Secrets		NOE	netscalerobse orter-sample	rvabilityexp Nets	alerObservabi	lityExporter -		app.kuberr app.kuberr	etes.io/mana=kusto etes.i =netscaler-op	Nov 20, 2024, 3:27 PM	:
ConfigMaps											

Ensure that the NetScaler Observability Exporter is successfully deployed.

6. Navigate to **Workloads > Pods** section and verify that the NetScaler Observability Exporter pod is up and running.
| E <b>Red Hat</b><br>OpenShift                   |   |  |               |                             |                      |  |                         | <b>III 4</b> 11 | 🕈 🕜 kube                 | admin <del>v</del> |
|---|---|--|---------------|-----------------------------|----------------------|--|-------------------------|-----------------|--------------------------|--------------------|
| ✿ Administrator                                 | • | Project: default 👻   | You           | are logged in as a temporar | administrative user. | Jpdate the <u>cluster OAuth configuration</u> to     | allow others to log in. |                 |                          |                    |
| Home  | > | Pods   |               |                             |                      |  |                         |                 | C                        | reate Pod          |
| Operators<br>OperatorHub<br>Installed Operators | ř | ▼ Filter ▼ Name ▼ Sea                                      | tch by name / | Ready 1                     | Restarts 💲           | Owner 💈  | Memory I                | CPU 1           | Created 1                |                    |
| Workloads                                       | ~ | Papache-7c646cfd49-9f2lh                                   | 2 Running     | 1/1                         | 0                    | RS apache-7c646cfd49                                 | 27.1 MiB                | 0.000 cores     | Jun 23, 2023, 3:35<br>PM | I                  |
| Pods  |   | apache-7c646cfd49-dtdv7                                    | 2 Running     | 1/1                         | 0                    | RS apache-7c646cfd49                                 | 23.2 MIB                | 0.000 cores     | Jun 23, 2023, 3:35<br>PM | i                  |
| Deployments                                     |   | netscaler-operator-controller-<br>manager-5dbf5649cf-rvrjr | 2 Running     | 2/2                         | 0                    | netscaler-operator-controller-<br>manager-5dbf5649cf | 217.3 MiB               | 0.004 cores     | Jun 23, 2023, 3:34<br>PM | ÷                  |
| StatefulSets<br>Secrets                         |   | noe-945ddc7c9-5p98g  | C Running     | 1/1                         | 0                    | noe-945ddc7c9  | -                       | -               | Jun 23, 2023, 5:40<br>PM | :                  |
| ConfigMaps                                      |   |  |               |                             |                      |  |                         |                 |                          |                    |
| Cron Jobs<br>Jobs                               |   |  |               |                             |                      |  |                         |                 |                          |                    |
| DaemonSets<br>ReplicaSets                       |   |  |               |                             |                      |  |                         |                 |                          |                    |
| ReplicationControllers                          |   |  |               |                             |                      |  |                         |                 |                          |                    |

## Deploy the NetScaler Ingress Controller as an OpenShift router plug-in

## October 16, 2024

In an OpenShift cluster, external clients need a way to access the services provided by pods. OpenShift provides two resources for communicating with services running in the cluster: routes and Ingress.

In an OpenShift cluster, a route exposes a service on a given domain name or associates a domain name with a service. OpenShift routers route external requests to services inside the OpenShift cluster according to the rules specified in routes. When you use the OpenShift router, you must also configure the external DNS to make sure that the traffic is landing on the router.

The NetScaler Ingress Controller can be deployed as a router plug-in in the OpenShift cluster to integrate with NetScalers deployed in your environment. The NetScaler Ingress Controller enables you to use the advanced load balancing and traffic management capabilities of NetScaler with your Open-Shift cluster.

OpenShift routes can be secured or unsecured. Secured routes specify the TLS termination of the route.

The NetScaler Ingress Controller supports the following OpenShift routes:

- **Unsecured Routes**: For Unsecured routes, HTTP traffic is not encrypted.
- **Edge Termination**: For edge termination, TLS is terminated at the router. Traffic from the router to the endpoints over the internal network is not encrypted.
- **Passthrough Termination**: With passthrough termination, the router is not involved in TLS offloading and encrypted traffic is sent straight to the destination.

• **Re-encryption Termination**: In re-encryption termination, the router terminates the TLS connection but then establishes another TLS connection to the endpoint.

For detailed information on routes, see the OpenShift documentation.

You can either deploy a NetScaler MPX or VPX appliance outside the OpenShift cluster or deploy NetScaler CPXs as pods inside the cluster. The NetScaler Ingress Controller integrates NetScalers with the OpenShift cluster and automatically configures NetScalers based on rules specified in routes.

Based on how you want to use NetScaler, there are two ways to deploy the NetScaler Ingress Controller as a router plug-in in the OpenShift cluster:

- As a sidecar container alongside NetScaler CPX in the same pod: In this mode, the NetScaler Ingress Controller configures the NetScaler CPX.
- As a standalone pod in the OpenShift cluster: In this mode, you can control the NetScaler MPX or VPX appliance deployed outside the cluster.

For information on deploying the NetScaler Ingress Controller to control the OpenShift ingress, see the NetScaler Ingress Controller for Kubernetes.

You can use NetScaler for load balancing OpenShift control plane (master nodes). NetScaler provides a solution to automate the configuration of NetScaler using Terraform instead of manually configuring the NetScaler. For more information, see NetScaler as a load balancer for the OpenShift control plane.

## **Alternate Backend Support**

OpenShift Alternate backends is now supported by NetScaler Ingress Controller.

NetScaler is configured according to the weights provided in the routes definition and traffic is distributed among the service pods based on those weights.

The following is an example of a route manifest with alternate backend:

```
kind: Route
1
2
    apiVersion: route.openshift.io/v1
3 metadata:
4
     name: r1
5
     labels:
6
        name: apache
7
     annotations:
8
        ingress.citrix.com/frontend-ip: "<Frontend-ip>"
9
  spec:
10
     host: some.alternate-backends.com
11
      to:
12
        kind: Service
13
        name: apache-1
14
        weight: 30
```

15	alternateBackends:
16	- kind: Service
17	name: apache-2
18	weight: 20
19	- kind: Service
20	name: apache-3
21	weight: 50
22	port:
23	targetPort: 80
24	wildcardPolicy: None

For this route, 30 percent of the traffic is sent to the service apache-1 and 20 percent is sent to the service apache2 and 50 percent to the service apache-3 based on weights provided in the route manifest

## Supported Citrix components on OpenShift

Citrix components	Versions
NetScaler Ingress Controller	Latest
NetScaler VPX	12.1 50.x and later
NetScaler CPX	13.0–36.28

## Note:

CRDs provided for the NetScaler Ingress Controller is not supported for OpenShift routes. You can use OpenShift ingress to use CRDs.

## Deploy NetScaler CPX as a router within the OpenShift cluster

In this deployment, you can use the NetScaler CPX instance for load balancing the North-South traffic to microservices in your OpenShift cluster. The NetScaler Ingress Controller is deployed as a sidecar alongside the NetScaler CPX container in the same pod using the cpx\_cic\_side\_car.yaml file.

**Before you begin**: When you deploy NetScaler CPX as a router, port conflicts can arise with the default router in OpenShift. You should remove the default router in OpenShift before deploying NetScaler CPX as a router. To remove the default router in OpenShift, perform the following steps:

1. Back up the default router configuration using the following command.

```
1 oc get -o yaml dc/router clusterrolebinding/router-router-role
serviceaccount/router > default-router-backup.yaml
```

2. Delete the default router using the following command.

1 oc delete -f **default**-router-backup.yaml

Perform the following steps to deploy NetScaler CPX as a router with the NetScaler Ingress Controller as a sidecar.

1. Download the cpx\_cic\_side\_car.yaml file using the following command:

```
1 wget https://raw.githubusercontent.com/citrix/citrix-k8s-ingress-
controller/master/deployment/openshift/manifest/
cpx_cic_side_car.yaml
```

2. Add the service account to privileged security context constraints (SCC) of OpenShift.

```
1 oc adm policy add-scc-to-user privileged system:serviceaccount:
    default:citrix
```

3. Deploy the NetScaler Ingress Controller using the following command:

1 oc create -f cpx\_cic\_side\_car.yaml

4. Verify if the NetScaler Ingress Controller is deployed successfully using the following command:

1 oc get pods --all-namespaces

## Deploy NetScaler MPX/VPX as a router outside the OpenShift cluster

In this deployment, the NetScaler Ingress Controller which runs as a stand-alone pod allows you to control the NetScaler MPX, or VPX appliance from the OpenShift cluster. You can use the cic.yaml file for this deployment.

**Note:** The NetScaler MPX or VPX can be deployed in *standalone*, *high-availability*, or *clustered* modes.

## Prerequisites

- Determine the IP address needed by the NetScaler Ingress Controller to communicate with the NetScaler appliance. The IP address might be any one of the following depending on the type of NetScaler deployment:
  - NSIP (for standalone appliances): The management IP address of a standalone NetScaler appliance. For more information, see IP Addressing in NetScaler.
  - SNIP (for appliances in High Availability mode): The subnet IP address. For more information, see IP Addressing in NetScaler.

- CLIP (for appliances in clustered mode): The cluster management IP (CLIP) address for a clustered NetScaler deployment. For more information, see IP addressing for a cluster.
- The user name and password of NetScaler VPX or NetScaler MPX used as the Ingress device. If you are not using the default credentials, NetScaler must have a system user account with certain privileges so that NetScaler Ingress Controller can configure NetScaler MPX or NetScaler VPX. To create a system user account on NetScaler, see Create system user account for NetScaler Ingress Controller in NetScaler.

You can directly pass the user name and password as environment variables to the NetScaler Ingress Controller or use OpenShift secrets (recommended). If you want to use OpenShift secrets, create a secret for the user name and password using the following command:

```
1 oc create secret generic nslogin --from-literal=username=<
    username> --from-literal=password=<password>
```

**Create a system user account for the NetScaler Ingress Controller in NetScaler** The NetScaler Ingress Controller configures a NetScaler appliance (MPX or VPX) using a system user account of the NetScaler appliance. The system user account must have the permissions to configure the following tasks on the NetScaler:

- Add, Delete, or View Content Switching (CS) virtual server
- Configure CS policies and actions
- Configure Load Balancing (LB) virtual server
- Configure Service groups
- Cofigure SSL certkeys
- Configure routes
- Configure user monitors
- Add system file (for uploading SSL testkeys from OpenShift)
- Configure Virtual IP address (VIP)
- Check the status of the NetScaler appliance
- Configure SSL actions and policies
- Configure SSL vServer
- Configure responder actions and policies

To create the system user account, perform the following:

- 1. Log on to the NetScaler appliance using the following steps:
  - a) Use an SSH client, such as PuTTy, to open an SSH connection to the NetScaler appliance.
  - b) Log on to the appliance by using the administrator credentials.
- 2. Create the system user account using the following command:

1 add system user <username> <password>

#### For example:

1 add system user cic mypassword

3. Create a policy to provide required permissions to the system user account. Use the following command:

1	<pre>add cmdpolicy cic-policy ALLOW '^\(\?!shell)\(\?!sftp)\(\?!scp) \(\?!batch)\(\?!source)\(\?!.\*superuser)\(\?!.\*nsroot)\(\?! install)\(\?!show\s+system\s+\(user cmdPolicy file))\(\?!\(set  add rm create export kill)\s+system)\(\?!\(unbind bind)\s+ system\s+\(user group))\(\?!diff\s+ns\s+config)\(\?!\(set unset  add rm bind unbind switch)\s+ns\s+partition).\* \(^install\s \*\(wi wf)) \(^\S+\s+system\s+file)^\(\?!shell)\(\?!sftp)\(\?! scp)\(\?!batch)\(\?!source)\(\?!.\*superuser)\(\?!.\*nsroot) \(\?!install)\(\?!show\s+system\s+\(user cmdPolicy file)) \(\?!\(set add rm create export kill)\s+system)\(\?!\(unbind  bind)\s+system\s+\(user group))\(\?!diff\s+ns\s+config)\(\?!\(</pre>
	<pre>bind)\s+system\s+\(user group))\(\?!diff\s+ns\s+config)\(\?!\( set unset add rm bind unbind switch)\s+ns\s+partition).\* \(^ install\s\*\(wi wf)) \(^\S+\s+system\s+file)'</pre>

**Note**: The system user account would have privileges based on the command policy that you define.

The command policy mentioned in *step 3* is similar to the built-in sysAdmin command policy with another permission to upload files.

In the command policy specification provided, special characters which need to be escaped are already omitted to easily copy-paste into the NetScaler command line.

For configuring the command policy from the NetScaler configuration wizard (GUI), use the following command policy specification.



4. Bind the policy to the system user account using the following command:

```
1 bind system user cic cic-policy \mathbf{0}
```

## Deploy the NetScaler Ingress Controller as a pod in an OpenShift cluster

Perform the following steps to deploy the NetScaler Ingress Controller as a pod:

1. Download the cic.yaml file using the following command:

```
1 wget https://raw.githubusercontent.com/citrix/citrix-k8s-ingress-
controller/master/deployment/openshift/manifest/cic.yaml
```

2. Edit the cic.yaml file and enter the values for the following environmental variables:

Environment Variable	Mandatory or Optional	Description
NS_IP	Mandatory	The IP address of the NetScaler appliance. For more details, see Prerequisites.
NS_USER and NS_PASSWORD	Mandatory	The user name and password of the NetScaler VPX or MPX appliance used as the Ingress device. For more details, see Prerequisites.
EULA	Mandatory	The End User License Agreement. Specify the value as Yes.
NS_VIP	Optional	NetScaler Ingress Controller uses the IP address provided in this environment variable to configure a virtual IP address to the NetScaler that receives Ingress traffic. <b>Note:</b> NS_VIP acts as a fallback when the frontend-ip annotation is not provided in Ingress or Route yaml. Not supported for Type Loadbalancer service

## 3. Add the service account to privileged security context constraints (SCC) of OpenShift.

4. Save the edited cic.yaml file and deploy it using the following command:

```
1 oc create -f cic.yaml
```

5. Verify if the NetScaler Ingress Controller is deployed successfully using the following command:

1 oc create get pods --all-namespaces

- 6. Configure static routes on NetScaler VPX or MPX to reach the pods inside the OpenShift cluster.
  - a) Use the following command to get the information about host names, host IP addresses, and subnets for static route configuration.

1 oc get hostsubnet

- b) Log on to the NetScaler instance.
- c) Add the route on the NetScaler instance using the following command.

```
1 add route <pod_network> <netmask> <gateway>
```

```
<B>Example:</b>
1
2
3
         oc get hostsubnet
4
                          HOST
5
         NAME
                                         HOST IP
                                                         SUBNET
6
         os.example.com os.example.com 192.168.122.46 192.1.1.0/24
7
8
    From the output of the `oc get hostsubnet` command:
9
10
         <pod_network> = 192.1.1.0
11
         Value for subnet = 192.1.1.0/x where x = 24 that means <
             netmask>= 255.255.255.0
12
         <gateway> = 192.168.122.46
13
14
    The required static route to add on NetScaler is as follows:
15
16
         add route 10.1.1.0 255.255.255.0 192.168.122.46
```

## Example: Deploy the NetScaler Ingress Controller as a router plug-in in an OpenShift cluster

In this example, the NetScaler Ingress Controller is deployed as a router plug-in in the OpenShift cluster to load balance an application.

1. Deploy a sample application (apache.yaml) in your OpenShift cluster and expose it as a service in your cluster using the following command.

```
1 oc create -f https://raw.githubusercontent.com/citrix/citrix-k8s-
ingress-controller/master/deployment/openshift/manifest/apache.
yaml
```

When you deploy a normal Apache pod in OpenShift, it may fail as Apache pod always runs as a root pod. OpenShift has strict security checks which block running a pod as root or binding to port 80. As a workaround, you can add the default service account of the pod to the privileged security context of OpenShift by using the following commands:

```
    oc adm policy add-scc-to-user privileged system:serviceaccount
:default:default
    oc adm policy add-scc-to-group anyuid system:authenticated
```

The content of the apache.yaml file is given as follows.

```
1 ----
  apiVersion: apps/v1
2
3 kind: Deployment
4 metadata:
5
   labels:
6
      name: apache-only-http
    name: apache-only-http
7
8 spec:
9
    replicas: 4
   selector:
10
    matchLabels:
11
12
       app: apache-only-http
13 template:
14 metadata:
15
       labels:
16
          app: apache-only-http
     spec:
17
18
       containers:
19
         - image: raghulc/apache-multiport-http:1.0.0
20
          imagePullPolicy: IfNotPresent
21
         name: apache-only-http
22
          ports:
23
          - containerPort: 80
24
            protocol: TCP
25
          - containerPort: 5080
            protocol: TCP
27
          - containerPort: 5081
28
            protocol: TCP
29
          - containerPort: 5082
30
           protocol: TCP
31 ---
32 apiVersion: apps/v1
33 kind: Deployment
34 metadata:
35 labels:
      name: apache-only-ssl
36
37
   name: apache-only-ssl
38 spec:
39
   replicas: 4
40 selector:
```

```
41 matchLabels:
42
        app: apache-only-ssl
43
   template:
     metadata:
44
45
       labels:
46
          app: apache-only-ssl
     spec:
47
48
        containers:
49
         - image: raghulc/apache-multiport-ssl:1.0.0
50
          imagePullPolicy: IfNotPresent
51
         name: apache-only-ssl
52
          ports:
53
          - containerPort: 443
54
            protocol: TCP
55
          - containerPort: 5443
            protocol: TCP
57
         - containerPort: 5444
            protocol: TCP
58
59
          - containerPort: 5445
            protocol: TCP
61 ---
62 apiVersion: v1
63 kind: Service
64 metadata:
65 name: svc-apache-multi-http
66 spec:
67
    ports:
68
    - name: apache-http-6080
     port: 6080
69
      targetPort: 5080
71
   - name: apache-http-6081
     port: 6081
72
      targetPort: 5081
74
   - name: apache-http-6082
75
      port: 6082
76
     targetPort: 5082
77
     selector:
78
     app: apache-only-http
  ____
79
80 apiVersion: v1
81 kind: Service
82 metadata:
   name: svc-apache-multi-ssl
83
84 spec:
85
   ports:
86
     - name: apache-ssl-6443
     port: 6443
87
88
      targetPort: 5443
89 - name: apache-ssl-6444
    port: 6444
targetPort: 5444
90
91
92
     - name: apache-ssl-6445
93 port: 6445
```

```
94 targetPort: 5445
95 selector:
96 app: apache-only-ssl
97 ---
```

2. Deploy the NetScaler Ingress Controller for NetScaler VPX as a stand-alone pod in the OpenShift cluster using the steps in Deploy the NetScaler Ingress Controller as a pod.

```
1 oc create -f cic.yaml
```

#### Note:

To deploy the NetScaler Ingress Controller with NetScaler CPX in the OpenShift cluster, perform the steps in Deploy the NetScaler Ingress Controller as a sidecar with NetScaler CPX.

- 3. Create an OpenShift route for exposing the application.
  - For creating an unsecured OpenShift route (unsecured-route.yaml), use the following command:

1 oc create -f unsecured-route.yaml

• For creating a secured OpenShift route with edge termination (secured-edge-route.yaml), use the following command.

```
1 oc create -f secured-route-edge.yaml
```

• For creating a secured OpenShift route with passthrough termination (secured-passthrough-route.yaml), use the following command.

```
1 oc create -f secured-passthrough-route.yaml
```

• For creating a secured OpenShift route with re-encryption termination (secured-reencryptroute.yaml), use the following command.

1 oc create -f secured-reencrypt-route.yaml

To see the contents of the YAML files for OpenShift routes in this example, see YAML files for routes.

Note:

For a secured OpenShift route with passthrough termination, you must include the default certificate.

4. Access the application using the following command.

## **YAML files for routes**

This section contains YAML files for unsecured and secured routes specified in the example.

## Note:

Keys used in this example are for testing purpose only. You must create your own keys for the actual deployment.

The contents of the unsecured-route.yaml file is given as follows:

```
1 apiVersion: v1
2 kind: Route
3 metadata:
    name: unsecured-route
4
5 spec:
    host: unsecured-route.openshift.citrix-cic.com
6
7
    path: "/"
   to:
8
9
       kind: Service
       name: svc-apache-multi-http
10
```

See, secured-edge-route for the contents of the secured-edge-route.yaml file.

The contents of the secured-passthrough-route is given as follows:

```
1 apiVersion: v1
2 kind: Route
3 metadata:
    name: secured-passthrough-route
4
5 spec:
    host: secured-passthrough-route.openshift.citrix-cic.com
6
7
    to:
       kind: Service
8
9
       name: svc-apache-multi-ssl
    tls:
       termination: passthrough
11
```

See, secured-reencrypt-route for the contents of the secured-reencrypt-route.yaml file.

# Deploy the NetScaler Ingress Controller with OpenShift router sharding support

## December 31, 2023

OpenShift router sharding allows distributing a set of routes among multiple OpenShift routers. By default, an OpenShift router selects all routes from all namespaces. In router sharding, labels are

added to routes or namespaces and label selectors to routers for filtering routes. Each router shard selects only routes with specific labels that match its label selection parameters.

NetScaler can be integrated with OpenShift in two ways and both deployments support OpenShift router sharding.

- NetScaler CPX deployed as an OpenShift router along with NetScaler Ingress Controller inside the cluster
- NetScaler Ingress Controller as a router plug-in for NetScaler MPX or VPX deployed outside the cluster

To configure router sharding for a NetScaler deployment on OpenShift, a NetScaler Ingress Controller instance is required per shard. The NetScaler Ingress Controller instance is deployed with route or namespace labels or both as environment variables depending on the criteria required for sharding. When the NetScaler Ingress Controller processes a route, it compares the route's labels or route's namespace labels with the selection criteria configured on it. If the route satisfies the criteria, the appropriate configuration is applied to NetScaler, otherwise it does not apply the configuration.

In router sharding, selecting a subset of routes from the entire pool of routes is based on selection expressions. Selection expressions are a combination of multiple values and operations.

For example, consider there are some routes with various labels for service level agreement(sla), geographical location (geo), hardware requirements (hw), department (dept), type, and frequency as shown in the following table.

Label	Values
sla	high, medium, low
geo	east, west
hw	modest, strong
dept	finance, dev, ops
type	static, dynamic
frequency	high, weekly

The following table shows selectors for route labels or namespace labels and a few sample selection expressions based on labels in the example. Route selection criteria is configured on the NetScaler Ingress Controller by using environment variables ROUTE\_LABELS and NAMESPACE\_LABLES.

Type of selector	Example
OR operation	ROUTE_LABELS='dept in (dev, ops)'
AND operation	ROUTE_LABELS=' hw=strong,type=dynamic,geo=west'
NOT operation	ROUTE_LABELS='dept!= finance'
Exact match	NAMESPACE_LABELS='frequency=weekly'
Exact match with both route and namespace labels Key based matching independent of value	NAMESPACE_LABELS='frequency=weekly' ROUTE_LABELS='sla=low' NAMESPACE_LABELS='name'
NOT operation with key based matching independent of value	NAMESPACE_LABELS='!name'

The label selectors use the language supported by Kubernetes labels.

If you want, you can change route or namespace labels by editing them later. Once you change the labels, router shard is revalidated and based on the change the NetScaler Ingress Controller updates the configuration on NetScaler.

## Deploy NetScaler CPX with OpenShift router sharding

To deploy CPX with OpenShift router sharding support, perform the following steps:

1. Download the cpx\_cic\_side\_car.yaml file using the following command:

```
1 wget https://raw.githubusercontent.com/citrix/citrix-k8s-ingress-
controller/master/deployment/openshift/manifest/
cpx_cic_side_car.yaml
```

2. Edit the cpx\_cic\_side\_car.yaml file and specify the route labels and namespace label selectors as environment variables.

The following example shows how to specify a sample route label and namespace label in the cpx\_cic\_side\_car.yaml file. This example selects routes with label "name" values as either abc or xyz and with namespace label as frequency=high.

1	er	1V:
2	-	name: "ROUTE_LABELS"
3		<pre>value: "name in (abc,xyz)"</pre>
4	-	name: "NAMESPACE_LABELS"
5		value: "frequency=high"

3. Deploy the NetScaler Ingress Controller using the following command.

```
1 oc create -f cpx_cic_side_car.yaml
```

## Deploy the NetScaler Ingress Controller router plug-in with OpenShift router sharding support

To deploy a NetScaler Ingress Controller router plug-in with router sharding, perform the following steps:

1. Download the cic.yaml file using the following command:

```
1 wget https://raw.githubusercontent.com/citrix/citrix-k8s-ingress-
controller/master/deployment/openshift/manifest/cic.yaml
```

2. Edit the cic.yaml file and specify the route labels and namespace label selectors as environment variables.

The following example shows how to specify a sample route label and namespace label in the cic.yaml file. This example selects routes with label "name" values as either abc or xyz and with namespace label as frequency=high.

1	env:
2	<pre>- name: "ROUTE_LABELS"</pre>
3	<pre>value: "name in (abc,xyz)"</pre>
4	<pre>- name: "NAMESPACE_LABELS"</pre>
5	value: "frequency=high"

3. Deploy the NetScaler Ingress Controller using the following command.

1 oc create -f cic.yaml

## Example: Create an OpenShift route and verify the route configuration on NetScaler VPX

This example shows how to create an OpenShift route with labels and verify the router shard configuration.

In this example, route configuration is verified on a NetScaler VPX deployment.

Perform the following steps to create a sample route with labels.

 Define the route in a YAML file. Following is an example for a sample route named as route. yaml.

```
1 apiVersion: v1
2 kind: Route
```

```
3 metadata:
4
   name: web-backend-route
    namespace: default
5
6
     labels:
7
          sla: low
8
          name: abc
9 spec:
     host: web-frontend.cpx-lab.org
10
11
      path: "/web-backend"
      port:
13
          targetPort: 80
14
       to:
15
          kind: Service
          name: web-backend
16
```

2. Use the following command to deploy the route.

1 oc create -f route.yaml

3. Add labels to the namespace where you create the route.

1 oc **label** namespace **default** 'frequency=high'

## Verify route configuration

You can verify the OpenShift route configuration on a NetScaler VPX by performing the following steps:

- 1. Log on to NetScaler VPX by performing the following:
  - Use an SSH client such as PuTTy, to open an SSH connection to NetScaler VPX.
  - Log on to NetScaler VPX by using administrator credentials.
- 2. Check if the service group is created using the following command.

```
1 show serviceGroup
```

3. Verify the route configuration on NetScaler VPX in the show serviceGroup command output.

Following is a sample route configuration from the show serviceGroup command output.

```
1 > show serviceGroup
2 k8s-web-backend-route_default_80_k8s-web-backend_default_80_svc -
HTTP
3 State: ENABLED Effective State: DOWN Monitor Threshold : 0
4 Max Conn: 0 Max Req: 0 Max Bandwidth: 0 kbits
5 Use Source IP: N0
6 Client Keepalive(CKA): N0
7 TCP Buffering(TCPB): N0
```

```
8 HTTP Compression(CMP): NO
9 Idle timeout: Client: 180 sec Server: 360 sec
10 Client IP: DISABLED
11 Cacheable: NO
12 SC: OFF
13 SP: OFF
14 Down state flush: ENABLED
15 Monitor Connection Close : NONE
16 Appflow logging: ENABLED
17 ContentInspection profile name: ???
18 Process Local: DISABLED
19 Traffic Domain: 0
```

## Deploy NetScaler CPX as an Ingress device in an Azure Kubernetes Service cluster

December 31, 2023

This topic explains how to deploy NetScaler CPX as an ingress device in an Azure Kubernetes Service (AKS) cluster. NetScaler CPX supports both the Advanced Networking (Azure CNI) and Basic Networking (Kubenet) mode of AKS.

Note:

If you want to use Azure repository images for NetScaler CPX or the NetScaler Ingress Controller instead of the default quay.io images, then see Deploy NetScaler CPX as an Ingress device in an AKS cluster using Azure repository images.

## Deploy NetScaler CPX as an ingress device in an AKS cluster

Perform the following steps to deploy NetScaler CPX as an ingress device in an AKS cluster.

Note:

In this procedure, Apache web server is used as the sample application.

1. Deploy the required application in your Kubernetes cluster and expose it as a service in your cluster using the following command.

```
1 kubectl create -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/deployment/azure/manifest/apache.
yaml
```

In this example, apache.yaml is used. You should use the specific YAML file for your application.

2. Deploy NetScaler CPX as an ingress device in the cluster using the following command.

```
1 kubectl create -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/deployment/azure/manifest/
standalone_cpx.yaml
```

3. Create the ingress resource using the following command.

```
1 kubectl create -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/deployment/azure/manifest/
cpx_ingress.yaml
```

4. Create a service of type LoadBalancer for accessing the NetScaler CPX by using the following command.

```
1 kubectl create -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/deployment/azure/manifest/
cpx_service.yaml
```

This command creates an Azure load balancer with an external IP for receiving traffic.

5. Verify the service and check whether the load balancer has created an external IP. Wait for some time if the external IP is not created.

1	kubectl get svc
2	
3	NAME TYPE CLUSTER-IP EXTERNAL-IP PORT (S) AGE
4	
5	apache  ClusterIP 10.0.103.3 none  80/TCP   2m
6	cpx-ingress  LoadBalancer  10.0.37.255   pending  80:32258/TCP
	,443:32084/TCP  2m
7	Kubernetes ClusterIP   10.0.0.1  none   443/TCP   22h

6. Once the external IP for the load-balancer is available as follows, you can access your resources using the external IP for the load balancer.

The health check for the cloud load-balancer is obtained from the readinessProbe configured in the NetScaler CPX deployment yaml file. If the health check fails, you should check the readinessProbe configured for NetScaler CPX. For more information, see readinessProbe and external Load balancer.

7. Access the application using the following command.

```
1 curl http://<External-ip-of-loadbalancer>/ -H 'Host: citrix-
ingress.com
```

## **Quick Deploy**

For the ease of deployment, you can just deploy a single all-in-one manifest that would combine the steps explained in the previous topic.

1. Deploy a NetScaler CPX ingress with in built NetScaler Ingress Controller in your Kubernetes cluster using the all-in-one.yaml.

```
1 kubectl create -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/deployment/azure/manifest/all-in-
one.yaml
```

2. Access the application using the following command.

```
1 curl http://<External-ip-of-loadbalancer>/ -H 'Host: citrix-
ingress.com'
```

## Note:

To delete the deployment, use the kubectl delete -f all-in-one.yaml command.

# Deploy NetScaler Ingress Controller in an Azure Kubernetes Service cluster with NetScaler VPX

December 31, 2023

This topic explains how to deploy the NetScaler Ingress Controller with NetScaler VPX in an Azure Kubernetes Service (AKS) cluster. You can also configure the Kubernetes cluster on Azure VMs and then deploy the NetScaler Ingress Controller with NetScaler VPX.

The procedure to deploy for both AKS and Azure VM is the same. However, if you are configuring Kubernetes on Azure VMs you need to deploy the CNI plug-in for the Kubernetes cluster.

## Prerequisites

You should complete the following tasks before performing the steps in the procedure.

• Ensure that you have a Kubernetes cluster up and running.

Note:

For more information on creating a Kubernetes cluster in AKS, see Guide to create an AKS cluster.

## Topology

The following is the sample topology used in this deployment.



## Get a NetScaler VPX instance from Azure Marketplace

You can create NetScaler VPX from the Azure Marketplace.

For more information on how to create a NetScaler VPX instance from Azure Marketplace, see Get NetScaler VPX from Azure Marketplace.

## Get the NetScaler Ingress Controller from Azure Marketplace

To deploy the NetScaler Ingress Controller, an image registry should be created on Azure and the corresponding image URL should be used to fetch the NetScaler Ingress Controller image.

For more information on how to create a registry and get the image URL, see Get NetScaler Ingress Controller from Azure Marketplace.

Once a registry is created, the NetScaler Ingress Controller registry name should be attached to the AKS cluster used for deployment.

```
1 az aks update -n <cluster-name> -g <resource-group-where-aks-
deployed> --attach-acr <cic-registry>
```

#### **Deploy NetScaler Ingress Controller**

Perform the following steps to deploy the NetScaler Ingress Controller.

1. Create NetScaler VPX login credentials using Kubernetes secret.

```
1 kubectl create secret generic nslogin --from-literal=username='<
    azure-vpx-instance-username>' --from-literal=password='<azure-
    vpx-instance-password>'
```

## Note:

The NetScaler VPX user name and password should be the same as the credentials set while creating NetScaler VPX on Azure.

2. Using SSH, configure a SNIP in the NetScaler VPX, which is the secondary IP address of the NetScaler VPX. This step is required for the NetScaler to interact with pods inside the Kubernetes cluster.

1 add ns ip <snip-vpx-instance-private-ip> <vpx-instance-primary-ipsubnet>

- snip-vpx-instance-private-ip is the dynamic private IP address assigned while adding a SNIP during the NetScaler VPX instance creation.
- vpx-instance-primary-ip-subnet is the subnet of the primary private IP address of the NetScaler VPX instance.

To verify the subnet of the private IP address, SSH into the NetScaler VPX instance and use the following command.

```
1 show ip <primary-private-ip-addess>
```

- 3. Update the NetScaler VPX image URL, management IP, and VIP in the NetScaler Ingress Controller YAML file.
  - a) Download the NetScaler Ingress Controller YAML file.

```
1 wget https://raw.githubusercontent.com/citrix/citrix-k8s-
ingress-controller/master/deployment/azure/manifest/
azurecic/cic.yaml
```

If you do not have wget installed, you can use the fetch or curl command.

b) Update the NetScaler Ingress Controller image with the Azure image URL in the cic.yaml file.

```
1 - name: cic-k8s-ingress-controller
2 # CIC Image from Azure
3 image: "<azure-cic-image-url>"
```

c) Update the primary IP address of the NetScaler VPX in the cic.yaml in the following field with the primary private IP address of the Azure VPX instance.

d) Update the NetScaler VPX VIP in the cic.yaml in the following field with the private IP address of the VIP assigned during VPX Azure instance creation.

```
1 # Set NetScaler VIP for the data traffic
2 - name: "NS_VIP"
3 value: "X.X.X.X"
```

4. Once you have configured the NetScaler Ingress Controller with the required values, deploy the NetScaler Ingress Controller using the following command.

1 kubectl create -f cic.yaml

## Verify the deployment using a sample application

1. Deploy the required application in your Kubernetes cluster and expose it as a service in your cluster using the following command.

```
1 kubectl create -f https://raw.githubusercontent.com/citrix/citrix
        -k8s-ingress-controller/master/deployment/azure/manifest/
        azurecic/apache.yaml
```

2. Create the Ingress resource using the following command.

```
1 kubectl create -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/deployment/azure/manifest/
azurecic/ingress.yaml
```

3. To validate your deployment, use the following command.

```
1 $ curl --resolve citrix-ingress.com:80:<Public-ip-address-of-VIP>
http://citrix-ingress.com
2 <html><body><h1>It works!</h1></body></html>
```

The response is received from the sample microservice (Apache) which is inside the Kubernetes cluster. NetScaler VPX has load-balanced the request.

## Deploy NetScaler CPX as an Ingress device in Google Cloud Platform

December 31, 2023

This topic explains how to deploy NetScaler CPX as an ingress device in Google Kubernetes Engine (GKE)

## Prerequisites

You should complete the following tasks before performing the steps in the procedure.

- Ensure that you have a Kubernetes Cluster up and running.
- If you are running your cluster in GKE, ensure that you have configured a cluster-admin role binding.

You can use the following command to configure cluster-admin role binding.

You can get your Google account details using the following command.

1 gcloud info | grep Account

## Deploy NetScaler CPX as an ingress device in Google Cloud Platform

1. Deploy the required application in your Kubernetes cluster and expose it as a service in your cluster using the following command.

```
1 kubectl create -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/deployment/gcp/manifest/apache.
yaml
```

In this example, apache.yaml is used. You should use the specific YAML file for your application.

2. Deploy NetScaler CPX as an ingress device in the cluster using the following command.

```
1 kubectl create -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/deployment/gcp/manifest/
standalone_cpx.yaml
```

3. Create the ingress resource using the following command.

```
kubectl create -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/deployment/gcp/manifest/
cpx_ingress.yaml
```

4. Create a service of type LoadBalancer for accessing the NetScaler CPX by using the following command.

```
1 kubectl create -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/deployment/gcp/manifest/
cpx_service.yaml
```

#### Note:

This command creates a load balancer with an external IP for receiving traffic.

1. Verify the service and check whether the load balancer has created an external IP. Wait for some time if the external IP is not created.

2. Once the external IP for the load-balancer is available as follows, you can access your resources using the external IP for the load balancer.

```
1 kubectl get svc
2
3 |Name | Type | Cluster-IP | External IP| Port\(s) | Age |
```

```
4 |-----| -----| -----| -----| -----|
5 |apache| ClusterIP|10.7.248.216|none|80/TCP |3m|
6 |cpx-ingress|LoadBalancer|10.7.241.6|EXTERNAL-IP CREATED|80:32258/
        TCP,443:32084/TCP|3m|
7 |kubernetes| ClusterIP| 10.7.240.1|none|443/TCP|22h|`
```

The health check for the cloud load-balancer is obtained from the readinessProbe configured in the NetScaler CPX service YAML file. If the health check fails, you should check the readinessProbe configured for NetScaler CPX.

For more information, see readinessProbe and external Load balancer.

3. Access the application using the following command.

```
1 curl http://<External-ip-of-loadbalancer>/ -H 'Host: citrix-
ingress.com'
```

## **Quick Deploy**

For the ease of deployment, you can just deploy a single all-in-one manifest that would combine the steps explained in the previous topic.

1. Deploy a NetScaler CPX ingress with in built NetScaler Ingress Controller in your Kubernetes cluster using the all-in-one.yaml.

```
1 kubectl create -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/deployment/gcp/manifest/all-in-
one.yaml
```

2. Access the application using the following command.

```
1 curl http://<External-ip-of-loadbalancer>/ -H 'Host: citrix-
ingress.com'
```

## Note:

To delete the deployment, use the kubectl delete -f all-in-one.yaml command.

## **Deploy the NetScaler Ingress Controller in Anthos**

January 21, 2025

Anthos is a hybrid and multi cloud platform that lets you run your applications on existing on-prem hardware or in the public cloud. It provides a consistent development and operation experience for cloud and on-premises environments.

NetScaler Ingress Controller can be deployed in Anthos GKE on-premises using the following deployment modes:

- Exposing NetScaler CPX with the sidecar ingress controller as a service of type LoadBalancer.
- Dual-tier Ingress deployment

## Expose NetScaler CPX as a service of type LoadBalancer in Anthos GKE on-prem

In this deployment, NetScaler VPX or MPX is deployed outside the cluster at Tier-1 and NetScaler CPX at Tier-2 inside the Anthos cluster similar to a dual-tier deployment. However instead of using Ingress, the NetScaler CPX is exposed using the Kubernetes service of type LoadBalancer.

The NetScaler Ingress Controller automates the process of configuring the IP address provided in the LoadBalancerIP field of the service specification.

## Prerequisites

- You must deploy a Tier-1 NetScaler VPX or MPX in the same subnet as the Anthos GKE on-prem user cluster.
- You must configure a subnet IP address (SNIP) on the Tier-1 NetScaler and Anthos GKE on-prem cluster nodes should be reachable using the IP address.
- To use a NetScaler VPX or MPX from a different network, use node controller to enable communication between the NetScaler and the Anthos GKE on-prem cluster.
- You must set aside a virtual IP address (VIP) to be used as a Load Balancer IP address.

## Deploy NetScaler CPX as service of type LoadBalancer in Anthos GKE on-premises

Perform the following steps to deploy NetScaler CPX as a service of type LoadBalancer in Anthos GKE on-premises.

1. Deploy the required application in your Kubernetes cluster and expose it as a service in your cluster using the following command.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig create -f https://
raw.githubusercontent.com/citrix/citrix-k8s-ingress-controller/
master/deployment/anthos/manifest/service-type-lb/apache.yaml
```

In this example, apache.yaml is used. You should use the specific YAML file for your application.

2. Deploy NetScaler CPX with the sidecar NetScaler Ingress Controller as Tier-2 Ingress device using the cpx-cic.yaml file.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig create -f https://
raw.githubusercontent.com/citrix/citrix-k8s-ingress-controller/
master/deployment/anthos/manifest/service-type-lb/cpx-cic.yaml
```

- 3. (Optional) Create a self-signed SSL certificate and a key to be used with the Ingress for TLS configuration.
  - 1 openssl req -subj '/CN=anthos-citrix-ingress.com/0=Citrix Systems Inc/C=IN' -new -newkey rsa:2048 -days 5794 -nodes -x509 -keyout \$PWD/anthos-citrix-certificate.key -out \$PWD/anthos-citrixcertificate.crt;openssl rsa -in \$PWD/anthos-citrix-certificate. key -out \$PWD/anthos-citrix-certificate.key

## Note:

If you already have an SSL certificate, you can create a Kubernetes secret using the same. This is just an example command to create a self-signed certificate and also this command assumes the host name of the application to be anthos-citrix-ingress.com.

4. Create a Kubernetes secret with the created SSL cert-key pair.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig create secret tls
anthos-citrix --cert=$PWD/anthos-citrix-certificate.crt --key=
$PWD/anthos-citrix-certificate.key
```

5. Create an Ingress resource for Tier-2 using the tier-2-ingress.yaml file.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig create -f https://
    raw.githubusercontent.com/citrix/citrix-k8s-ingress-controller/
    master/deployment/anthos/manifest/service-type-lb/tier-2-
    ingress.yaml
```

## 6. Create a Kubernetes secret for the Tier-1 NetScaler.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig create secret
generic nslogin --from-literal=username='citrix-adc-username'
--from-literal=password='citrix-adc-password'
```

- 7. Deploy the NetScaler Ingress Controller as a Tier-1 ingress controller.
  - a) Download the cic.yaml file.

- b) Enter the management IP address of NetScaler. Update the Tier-1 NetScaler's management IP address in the placeholder Tier-1-Citrix-ADC-IP specified in the cic. yaml file.
- c) Save and deploy the cic.yaml using the following command.

- 8. Expose NetScaler CPX as a Kubernetes service of type LoadBalancer.
  - a) Download the cpx-service-type-lb.yaml file.
  - b) Edit the YAML file and specify the value of VIP-**for**-accessing-microservices as the VIP address which is to be used for accessing the applications inside the cluster. This VIP address is the one set aside to be used as a Load Balancer IP address.
  - c) Save and deploy the cpx-service-type-lb.yaml file using the following command.

9. Update the DNS records with the IP address of VIP-**for**-accessing-microservices for accessing the microservice. In this example, to access the Apache microservice, you must have the following DNS entry.

```
1 `<VIP-for-accessing-microservices> anthos-citrix-ingress.com`
```

10. Use the following command to access the application.

```
1 curl -k --resolve anthos-citrix-ingress.com:443:<VIP-for-
accessing-microservices> https://anthos-citrix-ingress.
com/ <html><body><h1>It works!</h1></body></html>
```

## Note:

```
In this command, --resolve anthos-citrix-ingress.com:443:<VIP-for-accessing-microservices> is used to override the DNS configuration part in step 9 for demonstration purpose.
```

## Clean up the installation: Expose NetScaler CPX as service of type LoadBalancer

To clean up the installation, use the kubectl --kubeconfig delete command to delete each deployment.

To delete the NetScaler CPX service deployment (CPX+CIC service) use the following command:

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig delete -f cpx-
service-type-lb.yaml
```

To delete the Tier-2 Ingress object, use the following command.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig delete -f tier-2-
ingress.yaml
```

To delete the NetScaler CPX deployment along with the sidecar NetScaler Ingress Controller, use the following command.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig delete -f cpx-cic.
yaml
```

To delete the stand-alone NetScaler Ingress Controller, use the following command.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig delete -f cic.yaml
```

To delete the Apache microservice, use the following command.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig delete -f apache.
yaml
```

To delete the Kubernetes secret, use the following command.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig delete secret anthos
        -citrix
```

To delete the nslogin secret, use the following command.

## **Dual tier Ingress deployment**

In a dual-tier Ingress deployment, NetScaler VPX or MPX is deployed outside the Kubernetes cluster (Tier-1) and NetScaler CPXs are deployed inside the Kubernetes cluster (Tier-2).

NetScaler MPX or VPX devices in Tier-1 proxy the traffic (North-South) from the client to NetScaler CPXs in Tier-2. The Tier-2 NetScaler CPX then routes the traffic to the microservices in the Kubernetes cluster. The NetScaler Ingress Controller deployed as a standalone pod configures the Tier-1 NetScaler. The sidecar NetScaler Ingress Controller in one or more NetScaler CPX pods configures the associated NetScaler CPX in the same pod.

## Prerequisites

• You must deploy a Tier-1 NetScaler VPX or MPX in the same subnet as the Anthos GKE on-prem user cluster.

- You must configure a subnet IP address (SNIP) on the Tier-1 NetScaler and Anthos GKE on-prem cluster nodes should be reachable using the IP address.
- To use a NetScaler VPX or MPX from a different network, use the node controller to enable communication between the NetScaler and the Anthos GKE on-prem cluster.
- You must set aside a virtual IP address to be used as a front-end IP address in the Tier-1 Ingress manifest.

## Dual-tier Ingress deployment in Anthos GKE on-prem

Perform the following steps to deploy a dual-tier Ingress deployment of NetScaler in Anthos GKE onprem.

1. Deploy the required application in your Kubernetes cluster and expose it as a service in your cluster using the following command.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig create -f https://
raw.githubusercontent.com/citrix/citrix-k8s-ingress-controller/
master/deployment/anthos/manifest/dual-tiered-ingress/apache.
yaml
```

## Note:

In this example, apache.yaml is used. You should use the specific YAML file for your application.

2. Deploy NetScaler CPX with the NetScaler Ingress Controller as Tier-2 Ingress using the cpx-cic.yaml file.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig create -f https://
raw.githubusercontent.com/citrix/citrix-k8s-ingress-controller/
master/deployment/anthos/manifest/dual-tiered-ingress/cpx-cic.
yaml
```

3. Expose NetScaler CPX as a Kubernetes service using the cpx-service.yaml file.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig create -f https://
    raw.githubusercontent.com/citrix/citrix-k8s-ingress-controller/
    master/deployment/anthos/manifest/dual-tiered-ingress/cpx-
    service.yaml
```

4. (Optional) Create a self-signed SSL certificate and a key to be used with the Ingress for TLS configuration.

Note:

If you already have an SSL certificate, you can create a Kubernetes secret using the same.

```
1 openssl req -subj '/CN=anthos-citrix-ingress.com/0=Citrix Systems
Inc/C=IN' -new -newkey rsa:2048 -days 5794 -nodes -x509 -keyout
$PWD/anthos-citrix-certificate.key -out $PWD/anthos-citrix-
certificate.crt;openssl rsa -in $PWD/anthos-citrix-certificate.
key -out $PWD/anthos-citrix-certificate.key
```

## Note:

This is just an example command to create a self-signed certificate and also this command assumes that the hostname of the application to be anthos-citrix-ingress.com.

5. Create a Kubernetes secret with the created SSL cert-key pair.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig create secret tls
anthos-citrix --cert=$PWD/anthos-citrix-certificate.crt --key=
$PWD/anthos-citrix-certificate.key
```

6. Create an Ingress resource for Tier-2 using the tier-2-ingress.yaml file.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig create -f https://
    raw.githubusercontent.com/citrix/citrix-k8s-ingress-controller/
    master/deployment/anthos/manifest/dual-tiered-ingress/tier-2-
    ingress.yaml
```

7. Create a Kubernetes secret for the Tier-1 NetScaler.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig create secret
generic nslogin --from-literal=username='citrix-adc-username'
--from-literal=password='citrix-adc-password'
```

- 8. Deploy the NetScaler Ingress Controller as a Tier-1 ingress controller.
  - a) Download the cic.yaml file.
  - b) Enter the management IP address of NetScaler. Update the Tier-1 NetScaler's management IP address in the placeholder Tier-1-Citrix-ADC-IP specified in the cic. yaml file.
  - c) Save and deploy the cic.yaml using the following command.

- 9. Create an Ingress resource for Tier-1 using the tier-1-ingress.yaml file.
  - a) Download the tier-1-ingress.yaml file.
  - b) Edit the YAML file and replace VIP-Citrix-ADC with the VIP address which was set aside.

c) Save and deploy the tier-1-ingress.yaml file using the following command.

10. Update the DNS records with the IP address of VIP-Citrix-ADC for accessing the microservice. In this example, to access the Apache microservice, you must have the following DNS entry.

11. Use the following command to access the application.

```
1 curl -k --resolve anthos-citrix-ingress.com:443:<VIP-Citrix-ADC
2 https://anthos-citrix-ingress.com/
2 <html><body><h1>It works!</h1></body></html>
```

Note:

```
In this command, --resolve anthos-citrix-ingress.com:443:<VIP-for-accessing-microservices> is used to override the DNS configuration part.
```

## Clean up the installation: Dual tier Ingress

To clean up the installation, use the kubectl --kubeconfig delete command to delete each deployment.

To delete the Tier-1 Ingress object, use the following command.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig delete -f tier-1-
ingress.yaml
```

To delete the Tier-2 Ingress object, use the following command.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig delete -f tier-2-
ingress.yaml`
```

To delete the NetScaler CPX deployment along with the sidecar NetScaler Ingress Controller, use the following command.

1 kubectl --kubeconfig user-cluster-1-kubeconfig delete -f cpx-cic. yaml

To delete the NetScaler CPX service deployment, use the following command:

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig delete -f cpx-
service.yaml
```

To delete the stand-alone NetScaler Ingress Controller use the following command:

1 kubectl --kubeconfig user-cluster-1-kubeconfig delete -f cic.yaml

To delete the Apache microservice, use the following command.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig delete -f apache.
yaml
```

To delete the Kubernetes secret, use the following command.

```
1 kubectl --kubeconfig user-cluster-1-kubeconfig delete secret anthos
        -citrix
```

To delete the nslogin secret, use the following command.

```
kubectl --kubeconfig user-cluster-1-kubeconfig delete secret
    nslogin`
```

## Deploy NetScaler VPX in active-active high availability in EKS environment using Amazon ELB and NetScaler Ingress Controller

## December 31, 2023

1

The topic covers a solution to deploy NetScaler VPX in active-active high availability mode on multiple availability zones in AWS Elastic Container Service (EKS) platform. The solution combines AWS Elastic load balancing (ELB) and NetScaler VPX to load balance the Ingress traffic to the microservices deployed in EKS cluster. AWS ELB handles the Layer 4 traffic and the NetScaler VPXs provides advanced Layer 7 functionalities such as, advanced load balancing, caching, and content-based routing.

## **Solution overview**

A basic architecture of an EKS cluster would include three public subnet and three private subnets deployed across three availability zones as shown in the following diagram:



With the solution, the architecture of the EKS cluster would be as shown in the following diagram:



In the AWS cloud, AWS Elastic Load Balancing handles the Layer 4 TCP connections and load balances the traffic using a flow hash routing algorithm. The ELB can be either Network Load Balancer or a Classic Load Balancer.

AWS ELB listens for incoming connections as defined by its listeners. Each listener forwards a new connection to one of the available NetScaler VPX instances. The NetScaler VPX instance load balances the traffic to the EKS pods. It also performs other Layer 7 functionalities such as, rewrite policy, responder policy, SSL offloading and so on provided by NetScaler VPX. A NetScaler Ingress Controller is deployed in the EKS cluster for each NetScaler VPX instance. The NetScaler Ingress Controllers are configured with the same ingress class. And, it configures the Ingress objects in the EKS cluster on the respective NetScaler VPX instances.

AWS Elastic Load Balancing (ELB) has a DNS name to which an IP address is assigned dynamically. The DNS name can be added as Alias A record for your domain in Route53 to access the application hosted in the EKS cluster.

## **Deployment process**

Perform the following to deploy the solution:

- 1. Deploy NetScaler VPX Instances.
- 2. Deploy NetScaler Ingress Controller.
- 3. Set up Amazon Elastic Load Balancing. You can either set up Network Load Balancer or Classic Load Balancer.
- 4. Verify the solution.

## **Deploy NetScaler VPX instances**

NetScaler VPX is available as CloudFormation Template. The CloudFormation template deploys an instance of NetScaler VPX with single ENI on a given subnet. It also configures the NSIP, VIP, and SNIP for the NetScaler VPX instance.

For this solution you need to deploy two instances of NetScaler VPX. Deploy the NetScaler VPX instances on two availability zones by specifying the same NetScaler VPX and different public subnet.

After you deploy the NetScaler VPX instances, you can verify the deployment by reviewing the output of the CloudFormation template as shown in the following screenshot. The output must show the various IP addresses (VIP, SNIP, and NSIP) configured for the NetScaler VPX instances:

Overview	Outputs	Resources	Events	Template	Parameters	Tags	Stack Policy	Change Sets	Rollback Triggers	85	
Key				Value				Description		Export Name	
PrivateNSIP 1				10.0.165.10	10.0.165.108			Private IP (NS IP) used for management			
SNIP 1			10.0.160.14	10.0.160.143				Private IP address of the VPX instance associate			
PrivateVIP			10.0.161.33	10.0.161.33			Private IP address of the VPX instance associate				
PublicNSIP			54.68.242.2	54.68.242.236			Public IP (NS IP) used for management				
ManagementURL			https://54.6	https://54.68.242.236			HTTPS URL to the Management GUI (uses self-si				
InstanceIdNS			i-08687cfee6f5c3433			Instance Id of newly created VPX instance					
SecurityGroup		sg-0bd10f1	sg-0bd10f17e64d6f194			Security group id that the VPX belongs to					
PublicipVIP			54.202.202	54.202.202.189			Elastic IP address of the VPX instance associated				
ManagementURL2			http://54.68.242.236			HTTP URL to the Management GUI					

The CloudFormation template deploys the NetScaler VPX instance with primary IP address of the NetScaler VPX EC2 instance as VIP and secondary IP address as management IP address.

After the NetScaler VPX instances are successfully deployed, you must edit the security groups to allow traffic from EKS node group security group. Also, you must change the EKS node group security group to allow traffic from VPX instances.

## **Deploy NetScaler Ingress Controller**

Deploy separate instance of NetScaler Ingress Controller for each NetScaler VPX instance. Follow the deployment instructions to deploy NetScaler Ingress Controller.

After the NetScaler VPX instance is up, you must set up a system user account on the NetScaler VPX instances. The system user account is used by NetScaler Ingress Controller to log into the NetScaler VPX instances. For instruction to set up the system user account, see Create System User Account for CIC in NetScaler.

1. Edit the NetScaler Ingress Controller deployment YAML (citrix-ingress-controller.yaml).

Replace NS\_IP with the Private NSIP address of the respective NetScaler VPX instance. Also, provide the system user account user name and password that you have created on the NetScaler VPX instance. Once you edited the citrix-ingress-controller.yaml file, deploy the updated YAML file using the following command:

1 kubectl apply -f citrix-ingress-controller .yaml

- 2. Perform Step 1 on the second NetScaler Ingress Controller instance.
- 3. Ensure that both the pods are UP and running. Also, verify if NetScaler Ingress Controller is able to connect to the respective NetScaler VPX instance using the logs:

```
1 kubectl logs <cic_pod_name>
```

After the NetScaler Ingress Controller pods are deployed and running in the EKS cluster. Any, Kubernetes Ingress resource configured with the citrix ingress class is automatically configured on both the NetScaler VPX instances.

## Setup elastic load balancing

Depending upon your requirement you can configure any of the following load balancers:

- Network Load Balancers
- Classic Load Balancers
**Set up network load balancer** Network Load Balancer (NLB) is a good option for handling TCP connection load balancing. In this solution, NLB is used to accept the incoming traffic and route it to one of the NetScaler VPX instances. NLB load balances using the flow hash algorithm based on the protocol, source IP address, source port, destination IP address, destination port, and TCP sequence number.

To set up NLB:

- 1. Log on to the AWS Management Console for EC2.
- In the left navigation bar, click **Target Group**. Create two different target groups. One target group (*Target-Group-80*) for routing traffic on port 80 and the other target group (*Target-Group-443*) for routing traffic on 443 respectively.

			_
AMIs 4	Create target group		×
Bundle Tasks	Mary load balance mides so wet		
ELASTIC BLOCK STORE Volumes	settings that you specify.	to the targets in a target group using the target group settings that you specify, and performs nearch checks on the targets using the nearch check	
Snapshots Lifecycle Manager	Target group name 👔	Target-Group-80	
NETWORK & SECURITY     Security Groups     Flastic IPs	Target type	Instance     IP     Lambda function	
Placement Groups Key Pairs	Protocol (j)	(TCP :	
Network Interfaces	Port (j)	80	
LOAD BALANCING	VPC ()	vpc-0fb43ae7a368d730d (10.0.0.0/16)   kurr \$	
Target Groups			
AUTO SCALING     Launch Configurations	Health check settings		
Auto Scaling Groups	Protocol (j)	TCP \$	
SYSTEMS MANAGER	Advanced health che	ck settings	
Run Command			
State Manager			-
Configuration Compliance		Cancel Create	

- 3. Create a target group named, **\*Target-Group-80**. Perform the following:
  - a) In the Target group name field, enter the target group name as Target-Group-80.
  - b) In the Target type field, select Instance.
  - c) From the **Protocol** list, select **TCP**.
  - d) In the **Port** field, enter **80**.
  - e) From the VCP list, select you VPC where you deployed your EKS cluster.
  - f) In the Health check settings section, use TCP for health check.
  - g) Optional. You can modify the Advance health check settings to configure health checks.

5 4	Create target group		
dle Tasks	Your load balancer routes requests settings that you specify.	to the targets in a target group using the target group settings that you specify, and performs health checks	on the targets using the health check
pshots typie Manager	Target group name (j)	Target-Group-80	
WORK & SECURITY unity Groups tic IPs	Target type	Instance     IP     Lambda function	
ement Groups Pairs	Protocol (j)	(TOP E)	
vork Interfaces	Port (j)	80	
d Balancers	VPC ()	vpc-0fb43ae7a368d730d (10.0.0.0/16)   kurr \$	
) SCALING	Health check settings		
ch Configurations Scaling Groups	Protocol (j)	(TCP ¢)	
rems manager nces Command	Advanced health chee	:k settings	
e Manager			Create
figuration			Cancel

- 4. Create a target group named, **\*Target-Group-443**. Perform the following:
  - a) In the **Target group name** field, enter the target group name as **Target-Group-443**.
  - b) In the **Target type** field, select Instance.
  - c) From the **Protocol** list, select **TCP**.
  - d) In the **Port** field, enter **443**.
  - e) From the **VCP** list, select you VPC where you deployed your EKS cluster.
  - f) In the **Health check settings** section, use TCP for health check.
  - g) Optional. You can modify the Advance health check settings to configure health checks.

Capacity Heservations			
IMAGES			
AMIs Crea	ite target group		×
Bundle Tasks	ad balancer routes requests	to the targets in a target group using the target group settings that you specify and performs health checks on the targets using the health check	
ELASTIC BLOCK STORE setting	s that you specify.	to no targeto in a target group doing no target group outlings that you speeny, and performs notice encode on the targets doing no notice encode	
Volumes			
Snapshots	arget group name	Taroet-Group-443	
Lifecycle Manager		ter gar an an in a	
NETWORK & SECURITY	Target type	O Instance	
Security Groups		IP     Interface     Inte	
Elastic IPs			
Placement Groups	Protocol (i)	TCP	
Key Pairs			
Network Interfaces	Port (i)	443	
Load Balancers	VPC (I)	vnc-0fb43ae7a368d730d (10.0.0.0/16)   kurr 1	
Heal	th check settings		
AUTO SCALING			
Launch Configurations	Protocol (i)	TCP	
Auto Scaling Groups			
SYSTEMS MANAGER		E or Manage	
Run Command	ivanced nearth cheo	ck settings	
State Manager			_
Configuration		Cancel Create	
Compliance		· · · · · · · · · · · · · · · · · · ·	_

- 5. Once you have created the target groups, you must register the target instances.
  - a) Select the created target group in the list page, click the **Target** tab, and select **edit**.

- b) In the Instances tab, select the two NetScaler VPX instances and click Add to registered.
- 6. Repeat **Step 5** for the other target group that you have created.
- 7. Create Network Load Balancer.
  - a) In the left navigation bar, select Load Balancers, then click Create Load Balancer.
  - b) In the Select load balancer type window, click **Create** in the Network Load balancer panel.

Select load balancer type		
Elastic Load Balancing supports three types of load balancers: Application about which load balancer is right for you	Load Balancers, Network Load Balancers (new), and Classic Load Balance	rs. Choose the load balancer type that meets your needs. Learn more
Application Load Balancer	Network Load Balancer	Classic Load Balancer
HTTP	TCP TLS	PREVIOUS GENERATION for HTTP, HTTPB, and TCP
Create Choose an Application Load Balancer when you need a flexible feature set for your web applications with HTTP and HTTPS traffic. Operating at the request level, Application Load Balancers provide advanced routing and visibility features targeted at application architectures, including microservices and containers. Learn more >	Create Choose a Network Load Balancer when you need ultra-high performance, the ability to terminate TLS connections at scala, centralize certificate deployment, and static IP addresses for your application. Operating at the connection level, Network Load Balancers are capable of handling millions of requests per second securely while maintaining ultra-low latencies. Learn more >	Create Choose a Classic Load Balancer when you have an existing application running in the EC2-Classic network. Learn more >

- 8. In the **Configure Load Balancer** page, do the following:
  - a) In the Name field, enter a name for the load balancer.
  - b) In the Scheme field, select internet-facing.
  - c) In the Listeners section, click **Add listener** and add two entries with TCP as the load balancer protocol and 80 and 443 as the load balancer port respectively as shown in the following image:

1. Configure Load Balancer	2. Configure Security Settings 3.	Configure Routing	4. Register Targets	5. Review			
Step 1: Configure	Load Balancer						
Basic Configuration							
To configure your load balance receives TCP traffic on port 80.	r, provide a name, select a scher	ne, specify one or mo	re listeners, and sele	ct a network. The default	configuration is an Internet-	facing load balancer in the sele	ected network with a listener that
Name ()	VPX-HA-NLB						
Scheme ()	o internet-facing						
	Internal						
Listeners							
A listener is a process that che	cks for connection requests, usir	ng the protocol and po	ort that you configure	ed.			
Load Balancer Protocol				Load Balancer Po	rt		
TCP \$				80			۵
TCP \$				443			0
Add listener							

d) In the Availability Zones section, select the VPC, availability zones, and subnets where the NetScaler VPX instances are deployed.

1. Configure Load Balancer	Configure Security Settings 3. Configure Routing 4. Register Targets 5. Review
Step 1: Configure	oad Balancer
Availability Zones	
Specify the Availability Zones IP per Availability Zone if you	anable for your load balancer. The load balancer routes traffic to the targets in these Availability Zones only. You can specify only one subnet per Availability Zone. You may also add one Elas' h to have specific addresses for your load balancer.
Click here to manage your El	c IPs.
VPC ()	vpc-0fb43ae7a368d730d (10.0.0.0/16)   kumar-eks-test-VPC \$
Availability Zones	z us-west-2a subnet-0fa540d4e2c406466 (Public subnet 1)
	IPv4 address ① Assigned by AWS +
	Us-west-2b Select a subnet \$
	2 us-west-2c [subnet-01d6027e6842b828b (Public subnet 3) \$
	IPv4 address () Assigned by AWS \$

- 9. In the **Configure routing** page. do the following:
  - a) In the Target group list, click Existing target group.
  - b) In the Name field, enter Target-Group-80.
  - c) In the **Target type** field, select **Instance**.
  - d) In the **Protocol** list, select **TCP**.
  - e) In the **Port** field, enter **80**.
  - f) Select TCP from the Protocol list in the Health checks section as shown in the following image:

1. Configure Load Balancer	2. Configure Security Settings	3. Configure Routing 4. Regis	ter Targets 5. Review			
Step 3: Configur Your load balancer routes re be associated with only one	re Routing equests to the targets in this ta load balancer.	rget group using the protocol and	f port that you specify, and	performs health checks on th	e targets using these health c	heck settings. Note that each target group can
Target group						
Target group	Existing target group	p \$				
Name	(i) Target-Group-80	\$				
Target ty	ype   Instance  IP					
Protocol	() TCP	\$				
Port	(i) 80					
Health checks						
Protocol	() TCP	\$				
<ul> <li>Advanced health</li> </ul>	check settings					

10. In the **Review** page, review your configuration and click **Create**.

ease review the load b	V alancer details before continuing		
Load balancer			Ed
	Name VPX-HA-NLB		
	Scheme Internet-facing		
	Listeners Double - Protocoll CP		
	For Upc-0fbd3ae7a38kd730d (kumar-eks-test-VPCStack-1.JP0DKCKL.JOMW)		
	Subnets subnet-0fa540d4e2c406466 (Public subnet 1), subnet-01d6027e6842b828b (Public subnet 3) A		
	Tags		
De tiere			-
Routing			E
Ta	rget group Existing target group		
Target g	oup name Target-Group-80		
	Port 80		
	arget type installate Destroad TCP		
Health chec	Endlood TOP		
Health	heck port traffic port		
Healthy	threshold 3		
Unhealthy	threshold 3		
	Interval 30		
Targete			Ec
largets			

11. After the Network load balancer is created, select the load balancer that you have created for the list page. Select **Listeners** tab, select **TCP : 444** and then click **Edit**.

	Placement Groups Key Pairs	Descri	iption	Listeners	Mon	itoring Int	tegrated services	Tags						
-	Network Interfaces	A liste listene	ner chec ers and lis	ks for conne stener rules.	ction req	uests using its	s configured protoco	and po	rt, and the load balancer uses t	he listener rules to	o route request	s to targets. You ca	in add, remove,	or update
1	Load Balancers													
	Target Groups	Add	listener	Edit	Delete									
-	AUTO SCALING Launch Configurations		Listene	er ID		Security poli	icy SSL Certifica	ite	Default action					
	Auto Scaling Groups		TCP:8 arn27	0 7f3704d1970	)f87 <del>+</del>	N/A	N/A		Forward to Target-Group-80					
-	SERVICES Run Command		TCP:4 arne3	<b>43</b> 305eee3da4	ef5a <del>*</del>	N/A	N/A		Forward to Target-Group-80					

12. In the Listeners page, delete the default action and then select Target-Group-443 in the Forward to list.

<	Listeners				VPX-HA-NLB   TCP:443	~ <b>2</b>	Θ
	View/edit listener. Each listen	er must include one action of type forward	đ.			Upd	ate
	VPX-HA-NLB   TCP Listeners belonging to Netwo routed. Learn more	: 443 rk Load Balancers check for connection re	equests using the protocol and port y	rou configure. Each listener n	nust include a default actio	n to ensure all reques	its are
	ARN arn:aws:elasticloadbalancing	us-west-2:957176903625:listener/net/VP	X-HA-NLB/7cd5a04dfa79d9d1/e330	5eee3da4ef5a			
	Protocol : port Select the protocol for connect TCP • : 443 Default action(s)	ons from the client to your load balancer, and	I enter a port number from which to liste	n to for traffic.			
	1. Forward to	io comp	Û				
	Target-Group-443	$\odot$	•				
	+ Add action		~				

13. Click Update.

**Set up classic load balancer** Alternative to Amazon Network load balancer, you can set up Classic Load Balancer (CLB) as Tier 1 TCP load balancer.

- 1. Log on to the AWS Management Console for EC2.
- 2. In the left navigation bar, select Load Balancers, then click Create Load Balancer.

	aws	Services	*	Resource Groups	*	*		\$
-	IMAGES AMIs		Crea	ate Load Balancer	Actio	ons 👻		
	Bundle Tasks	4	Q	search : env:%2520ty	/pe 🙁	Add filter		
-	ELASTIC BLOCK STO Volumes	DRE		Name		<ul> <li>DNS name</li> </ul>	- State	VPC ID
	Snapshots Lifecycle Manager						No results for	und. Please alter your search.
-	NETWORK & SECURI	TY						
	Elastic IPs							
	Placement Groups Key Pairs							
-	Network Interfaces	;	Sele	ct a load balancer				0.0.0
Ī	Load Balancers Target Groups							

3. In the **Select load balancer type** window, click **Create** on the Classic Load balancer panel.

lastic Load Balancing supports three types of load balancers: Application bout which load balancer is right for you	Load Balancers, Network Load Balancers (new), and Classic Load Balance	rs. Choose the load balancer type that meets your needs. Learn more
Application Load Balancer	Network Load Balancer	Classic Load Balancer
HTTP HTTPS	TCP TLS	PREVIOUS GENERATION for HTTP, HTTPB, and TOP
Create Choose an Application Load Balancer when you need a flexible feature set for your web applications with HTTP and HTTPS traffic. Operating at the request level, Application Load Balancers provide advanced routing and visibility features targeted at application architectures, including microservices and containers. Learn more >	Create Choose a Network Load Balancer when you need ultra-high performance, the ability to terminate TLS connections at scale, centralize certificate deployment, and static IP addresses for your application. Operating at the connection level, Network Load Balancers are capable of handling millions of nequests per second securely while maintaining ultra-low latencies. Learn more >	Create Choose a Classic Load Balancer when you have an existing application running in the EC2-Classic network. Learn more >

- 4. In the **Define Load Balancer** page, do the following:
  - a) In the Load Balancer name field, enter a name for the load balancer.
  - b) In the Create LB Inside list, select your NetScaler VPX.
  - c) In the Listener Configuration section, click Add and add two entries with TCP as the load balancer protocol and 80 and 443 as the load balancer port respectively. Also, select TCP as instance protocol and 80 and 443 as the instance port respectively as shown in the following image:

1. Define Load Balancer	2. Assign	Security Groups	3. Configure Security Settings	4. Configure Health Check	5. Add EC2 Instances	6. Add Tags	7. Review	
Step 1: Define L	oad E	Balancer						
Basic Configurati	ion							
This wizard will walk you th ports and protocols for you port 80.	rough sett ir load bali	ting up a new loa ancer. Traffic fror	ad balancer. Begin by giving y m your clients can be routed fr	our new load balancer a uni rom any load balancer port t	que name so that you c to any port on your EC2	can identify it fro 2 instances. By	om other load balancers you might create. You v default, we've configured your load balancer wi	will also need to configure th a standard web server on
Load Balance	r name:	VPX-HA						
Create LE	Inside:	EC2-Classic		\$				
Create an internal load be	alancer:	(what's this?)						
Listener Config	uration:							
Load Balancer Protocol			Load Balancer Port		Instance Protocol		Instance Port	
ТСР	\$		80		TCP	\$	80	8
TCP	\$		443		TCP	\$	443	8
Add								

d) In the **Select Subnets** section, select two public subnets in two different availability zones for the Classic Load balancer to route the traffic. These subnets are same as where you have deployed the NetScaler VPX instances.

lability Zone where you wish traffic to be rout ad balancer.	ed by your load balancer. If you have instance	es in only one Availability Zone, please select	at least two Subnets in different Availability
kumar-eks-test-VPCStack-1JP0DKCKLJOM	N		
Availability Zone	Subnet ID	Subnet CIDR	Name
us-west-2a	subnet-09fa77ee0fc512855	10.0.176.0/24	Management subnet- Primary
us-west-2a	subnet-05e67a9a5d2fb64ba	10.0.0/19	Private subnet 1A
us-west-2b	subnet-08b3b9d70f2e73056	10.0.32.0/19	Private subnet 2A
us-west-2b	subnet-02c7729d3657184e5	10.0.144.0/20	Public subnet 2
us-west-2c	subnet-0a23c4d432639d440	10.0.64.0/19	Private subnet 3A
us-west-2c	subnet-0f55057eb3ac9eaa8	10.0.177.0/24	Management subnet-secondary
Availability Zone	Subnet ID	Subnet CIDR	Name
us-west-2a	subnet-0fa540d4e2c406466	10.0.128.0/20	Public subnet 1
us-west-2c	subnet-01d6027e6842b828b	10.0.160.0/20	Public subnet 3
	ability Zone where you wish traffic to be rout ad balancer. auman-eks-test-VPCStack-1JP0DKCKLJOM Availability Zone us-west-2a us-west-2b us-west-2b us-west-2c us-west-2c Availability Zone us-west-2a us-west-2a us-west-2a us-west-2a	Availability Zone     Subnet ID       Availability Zone     \$ubmet-Dfis/77e0fc512855       us-west-2a     submet-Ofis/77e0fc512855       us-west-2b     submet-Ofis/77e0fc512855       us-west-2b     submet-Ofis/77e0fc512855       us-west-2c     submet-Ofis/77e0fc512855       us-west-2c     submet-Ofis/77e0fc512855       us-west-2c     submet-Ofis/77e0fc512855       wavest-2c     submet-Ofis/77e0fc512855       wavest-2c     submet-Ofis/76fc7e03ac59aa4432       us-west-2c     submet-Ofis/76fc7e03ac59aa4432	Availability Zone where you wish traffic to be routed by your load balancer. If you have instances in only one Availability Zone, please select databancer.           Availability Zone         Subnet ID         Subnet CIDR           us-west-2a         subnet-09627/9605612855         10.0.176.0/24           us-west-2b         subnet-05667/963620664ba         10.0.0.019           us-west-2c         subnet-02627/293086718465         10.0.144.0/20           us-west-2c         subnet-026276203654382036440         10.0.84.0/19           us-west-2c         subnet-0055067603a269aa84         10.0.170.0/24           us-west-2c         subnet-01650576b3a269aa8440         10.0.84.0/19           us-west-2c         subnet-01650576b3a269aa84         10.0.170.0/24           us-west-2c         subnet-01650576b3a269aa84         10.0.170.0/24           us-west-2c         subnet-01650576b3a269aa88         10.0.170.0/24           us-west-2c         subnet-01650576b3a269aa88         10.0.120.0/20

e) In the **Assign Security Groups** page, select a security group for the ELB instance. The security group can be same as the security group attached to NetScaler VPX ENI or it can be a new security group.

If you are using a new security group, make sure that you allow traffic to the NetScaler VPX security group from the ELB security group and conversely.

Step 2: Assign Security Groups

As	sign a security group:	Create a new security group Select an existing security group	Filter VPC security groups \$
	Security Group ID	Name	Description
	sg-0ac7db042182514d1	default	default VPC security group
	sg-0ed7dfa631128ff24	eks3-SecurityGroup-1Q96QG85RUL8B	Allow http/s and ssh to ENI from Internet
	sg-0dbdcade8db496ca2	kumar-3nic-SecurityGroup-1A7YWQXDBLBL4	Enable SSH access via port 22, HTTP with port 80 and HTTPS
	sg-0ec55d354e8476b90	kumar-eks-test-EKSStack-1NKQYWILV3SL2-BastionStack-HR38FFXQGX6D-BastionSecurityGroup-16NLJM3V8DX70	Enables SSH Access to Bastion Hosts
	sg-0ec25b1bed98b81ed	kumar-eks-test-EKSStack-1NKQYWILV3SL2-ControlPlaneSecurityGroup-T1YIR95SK2VU	Cluster communication
	sg-081311750dfd5ca8f	kumar-eks-test-EKSStack-1NKQYWILV3SL2-FunctionStack-N0VGJSBCJJFW-EKSLambdaSecurityGroup-VM16XQSRSIOQ	Security group for lambda to communicate with cluster API
	sg-055be021f07d76076	kumar-eks-test-EKSStack-1NKQYWILV3SL2-NodeGroupStack-1ESSD2MUSV1VW-NodeSecurityGroup-186A7LDNXLOVT	Security group for all nodes in the node group
	sg-0e4412ee40203eae4	kumar-nic1-modified-SecurityGroup-SBV01CRVLZAF	Allow http/s and ssh to ENI from Internet
	sg-0ec9d8396ed7882c7	kumar-nic1-modified2-SecurityGroup-1PIJ7IKZRM4BL	Allow http/s and ssh to ENI from Internet
	sg-0a1dcdf9d89899fa1	kumar-single-nic-SecurityGroup-KY0IOS0LVAFB	Allow http/s and ssh to ENI from Internet
	sg-0949674c9858284e2	mod9-SecurityGroup-T7UFDGNZ0F7W	Allow http/s and ssh to ENI from Internet
	sg-00bbff620ff6e05b1	VPX VIP	VPX VIP policy

f) In the Configure Health Check page, select the configuration for the health check. By default health check is set as TCP on port 80, optionally you can do the health check on

### port 443 as well.

Step 4: Configure Hea Your load balancer will automatically p balancer. Customize the health check	1 Check form health checks on your EC2 instances and only route traffic to instances that pass the health check. If an instance fails the health check, it is automatically removed from the loar meet your specific needs.
Ping Protocol	TCP \$
Ping Port	80
Advanced Details	
Response Timeout	5 seconds
Interval (i)	30 seconds
Unhealthy threshold (j)	2 \$
Healthy threshold (i)	10 \$

g) In the **Add EC2 Instances** page, select two NetScaler VPX instances that were deployed earlier.

```
1 ![Classic ADD EC2 Instances](/en-us/netscaler-k8s-ingress-
controller/media/classic-add-ec2.png)
```

- h) In the Add Tags page, add tags as per your requirement.
- i) In the Review page, review your configurations.
- j) Click Create.

### Verify the solution

After you have successfully deployed NetScaler VPX, AWS ELB, and NetScaler Ingress Controller, you can verify the solution using a sample service.

Perform the following:

1. Deploy a sample service and ingress using app.yaml.

```
1 kubectl apply -f app.yaml
```

- 2. Log on to the NetScaler VPX instance and verify if the Content Switching vserver are successfully configured on both the NetScaler VPX instance. Do the following:
  - a) Log on to the NetScaler VPX instance. Perform the following:
    - i. Use an SSH client, such as PuTTy, to open an SSH connection to the NetScaler VPX instance.
    - ii. Log on to the instance by using the administrator credentials.
  - b) Verify if the Content Switching (cs) vserver is configured on the instance using the following command:

1 sh cs vserver

**Output:** 

1	1) k8s-10.0.139.87:80:http (10.0.139.87:80) - HTTP Type:
	CONTENT
2	State: UP
3	Last state change was at Fri Apr 12 14:24:13 2019
4	Time since last state change: 3 days, 03:09:18.920
5	Client Idle Timeout: 180 sec
6	Down state flush: ENABLED
7	Disable Primary Vserver On Down : DISABLED
8	Comment: uid=
	NNJRYQ54VM2KWCXOERK6HRJHR4VEQYRI7U3W4BNFQLTIAENMTHW/
	====
9	Appflow logging: ENABLED
10	Port Rewrite : DISABLED
11	State Update: DISABLED
12	Default: Content Precedence: RULE
13	Vserver IP and Port insertion: OFF
14	L2Conn: OFF Case Sensitivity: ON
15	Authentication: OFF
16	401 Based Authentication: OFF
17	Push: DISABLED Push VServer:
18	Push Label Rule: none
19	Listen Policy: NONE
20	IcmpResponse: PASSIVE
21	RHIstate: PASSIVE
22	Traffic Domain: 0

c) Access the application test.example.com using the DNS name of the ELB instance.

1 # curl -H 'Host: test.example.com' <DNS name of the ELB>

### Example:

1 % curl -H 'Host: test.example.com' http://VPX-HA -829787521.us-west-2.elb.amazonaws.com

d) To delete the deployment, use the following command:

1 kubectl delete -f app.yaml

# Troubleshooting

Problem	Resolution
CloudFormation stack failure	Ensure that the IAM user or role has sufficient privilege to create EC2 instances and Lambda configurations. Ensure that you haven't exceeded the resource quota.

Problem	Resolution
NetScaler Ingress Controller unable to communicate with the NetScaler VPX instances.	Ensure that user name and password is correct in citrix-ingress-controller.yaml file.
	Ensure that the NetScaler VPX security group allows the traffic on port 80 and 443 from the EKS node group security group.
The services are DOWN in the NetScaler VPX instances.	Ensure that the NetScaler VPX traffic can reach the EKS cluster. Modify the security group of EKS node group to allow traffic from NetScaler VPX security group.
Traffic not routing to NetScaler VPX instance from ELB.	Ensure that security group of NetScaler VPX allows traffic from the ELB security group.

# Deploy the NetScaler Ingress Controller for NetScaler with admin partitions

### June 24, 2025

NetScaler Ingress Controller is used to automatically configure one or more NetScaler based on the Ingress resource configuration. The ingress NetScaler appliance (MPX or VPX) can be partitioned into logical entities called admin partitions, where each partition can be configured and used as a separate NetScaler appliance. For more information, see Admin Partition. NetScaler Ingress Controller can also be deployed to configure NetScaler with admin partitions.

For NetScaler with admin partitions, you must deploy a single instance of NetScaler Ingress Controller for each partition. And, the partition must be associated with a partition user specific to the NetScaler Ingress Controller instance.

Note:

NetScaler Metrics Exporter supports exporting metrics from the admin partitions of NetScaler.

# Prerequisites

Ensure that:

• Admin partitions are configured on the NetScaler appliance. For instructions see, Configure admin partitions.

• Create a partition user specifically for the NetScaler Ingress Controller. NetScaler Ingress Controller configures the NetScaler using this partition user account. Ensure that you do not associate this partition user to other partitions in the NetScaler appliance.

### Note:

For SSL-related use cases in the admin partition, ensure that you use NetScaler version 12.0–56.8 and above.

# To deploy the NetScaler Ingress Controller for NetScaler with admin partitions

1. Download the citrix-k8s-ingress-controller.yaml using the following command:

```
1 wget https://raw.githubusercontent.com/citrix/citrix-k8s-ingress-
controller/master/deployment/baremetal/citrix-k8s-ingress-
controller.yaml
```

2. Edit the citrix-k8s-ingress-controller.yaml file and enter the values for the following environmental variables:

Environment Variable	Mandatory or Optional	Description
NS_IP	Mandatory	The IP address of the NetScaler appliance. For more details, see Prerequisites.
NS_USER and NS_PASSWORD	Mandatory	The user name and password of the partition user that you have created for the NetScaler Ingress Controller. For more details, see Prerequisites.
NS_VIP	Mandatory	NetScaler Ingress Controller uses the IP address provided in this environment variable to configure a virtual IP address to the NetScaler that receives the Ingress traffic. <b>Note:</b> NS_VIP acts as a fallback when the frontend-ip annotation is not provided in Ingress YAML. Only Supported for Ingress.

Environment Variable	Mandatory or Optional	Description
NS_SNIPS	Optional	Specifies the SNIP addresses on the NetScaler appliance or the SNIP addresses on a specific admin partition on the NetScaler appliance.
NS_ENABLE_MONITORING	Mandatory	Set the value Yes to monitor NetScaler. <b>Note:</b> Ensure that you disable NetScaler monitoring for NetScaler with admin partitions. Set the value
EULA	Mandatory	The End User License Agreement. Specify the value as Yes.
Kubernetes_url	Optional	The kube-apiserver url that NetScaler Ingress Controller uses to register the events. If the value is not specified, NetScaler Ingress Controller uses the internal kube-apiserver IP address
LOGLEVEL	Optional	The log levels to control the logs generated by NetScaler Ingress Controller. By default, the value is set to DEBUG. The supported values are: CRITICAL, ERROR, WARNING, INFO, and DEBUG. For more information, see Log Levels
NS_PROTOCOL and NS_PORT	Optional	Defines the protocol and port that must be used by the NetScaler Ingress Controller to communicate with NetScaler. By default, the NetScaler Ingress Controller uses HTTPS on port 443. You can also use HTTP on port 80.

Environment Variable	Mandatory or Optional	Description
ingress-classes	Optional	If multiple ingress load
		balancers are used to load
		balance different ingress
		resources. You can use this
		environment variable to specify
		the NetScaler Ingress
		Controller to configure
		NetScaler associated with a
		specific ingress class. For
		information on Ingress classes,
		see Ingress class support

3. Once you update the environment variables, save the YAML file and deploy it using the following command:

1 kubectl create -f citrix-k8s-ingress-controller.yaml

4. Verify if the NetScaler Ingress Controller is deployed successfully using the following command:

1 kubectl get pods --all-namespaces

# Use case: How to securely deliver multitenant microservice-based applications using NetScaler admin partitions

You can isolate ingress traffic between different microservice based applications with the NetScaler admin partition using NetScaler Ingress Controller. NetScaler admin partition enables multitenancy at the software level in a single NetScaler instance. Each partition has its own control plane and network plane.

You can deploy one instance of NetScaler Ingress Controller in each namespace in a cluster.

For example, imagine you have two namespaces in a Kubernetes cluster and you want to isolate these namespaces from each other under two different admins. You can use the admin partition feature to separate these two namespaces. Create namespace 1 and namespace 2 and deploy NetScaler Ingress Controller separately in both of these namespaces.

NetScaler Ingress Controller instances provide configuration instructions to the respective NetScaler partitions using the system user account specified in the YAML manifest.



In this example, apache and guestbook sample applications are deployed in two different namespaces (namespace 1 and namespace 2 respectively) in a Kubernetes cluster. Both apache and guestbook application teams want to manage their workload independently and do not want to share resources. NetScaler admin partition helps to achieve multitenancy and in this example, two partitions (default, partition1) are used to manage both application workload separately.

The following prerequisites apply:

- Ensure that you have configured admin partitions on the NetScaler appliance. For instructions see, Configure admin partitions.
- Ensure that you create a partition user account specifically for the NetScaler Ingress Controller. NetScaler Ingress Controller configures the NetScaler using this partition user account. Ensure that you do not associate this partition user to other partitions in the NetScaler appliance.

# Example

The following example scenario shows how to deploy different applications within different namespaces in a Kubernetes cluster and how the request can be isolated from ADC using the admin partition.

In this example, two sample applications are deployed in two different namespaces in a Kubernetes cluster. In this example, it is used a default partition in NetScaler for the apache application and the admin partition p1 for the guestbook application.

### **Create namespaces**

Create two namespaces ns1 and ns2 using the following commands:

```
    kubectl create namespace ns1
    kubectl create namespace ns2
```

### Configurations in namespace ns1

1. Deploy the apache application in ns1.

```
1 apiVersion: v1
2 kind: Namespace
3 metadata:
4
   name: nsl
5
6 ---
7 apiVersion: apps/v1
8 kind: Deployment
9 metadata:
10 labels:
11
    app: apache-ns1
12
   name: apache-ns1
13 namespace: nsl
14 spec:
15
   replicas: 2
   selector:
17
     matchLabels:
18
       app: apache-ns1
   template:
19
20 metadata:
21
      labels:
22
         app: apache-ns1
     aı
spec:
23
24
       containers:
25
        - image: httpd
26
         name: httpd
27 ---
28
29 apiVersion: v1
30 kind: Service
31 metadata:
32 creationTimestamp: null
```

```
33 labels:
34
     app: apache-ns1
35 name: apache-ns1
36 namespace: ns1
37 spec:
38
    ports:
39
     - port: 80
    protocol: TCP
40
41
      targetPort: 80
42 selector:
43
     app: apache-ns1
```

2. Deploy NetScaler Ingress Controller in ns1.

You can use the YAML file to deploy NetScaler Ingress Controller or use the Helm chart.

Ensure that you use the user credentials that are bound to the default partition.

1 helm install cic-def-part-ns1 citrix/citrix-ingress-controller -set nsIP=<nsIP of ADC>,license.accept=yes,adcCredentialSecret= nslogin,ingressClass[0]=citrix-def-part-ns1 --namespace ns1

3. Deploy the Ingress resource.

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
    name: ingress-apache-ns1
4
5
    namespace: nsl
    annotations:
6
     ingress.citrix.com/frontend-ip: "< ADC VIP IP >"
7
8 spec:
    ingressClassName: "citrix-def-part-ns1"
9
10
     rules:
     - host: apache-ns1.com
11
12
      http:
13
         paths:
         - backend:
14
15
             service:
              name: apache-ns1
17
               port:
18
                 number: 80
19
           pathType: Prefix
20
           path: /index.html
```

4. NetScaler Ingress Controller in ns1 configures the ADC entities in the default partition.

### Configurations in namespace ns2

- 1. Deploy guestbook application in ns2.
  - 1 apiVersion: v1

```
2 kind: Namespace
3 metadata:
4 name: ns2
5 ---
6 apiVersion: v1
7 kind: Service
8 metadata:
   name: redis-master
9
   namespace: ns2
11 labels:
12
     app: redis
13tier: backend14role: master
15 spec:
   ports:
16
    - port: 6379
17
      targetPort: 6379
18
19 selector:
    app: redis
20
      tier: backend
21
22
     role: master
23 ---
24 apiVersion: apps/v1 # for k8s versions before 1.9.0 use apps/
      v1beta2 and before 1.8.0 use extensions/v1beta1
25 kind: Deployment
26 metadata:
27
    name: redis-master
28
    namespace: ns2
29 spec:
   selector:
     matchLabels:
31
       app: redis
32
        role: master
34
        tier: backend
35 replicas: 1
   template:
    metadata:
37
38
       labels:
39
           app: redis
40
           role: master
41
          tier: backend
42
      spec:
43
        containers:
44
         - name: master
45
          image: k8s.gcr.io/redis:e2e # or just image: redis
46
          resources:
47
            requests:
48
              cpu: 100m
49
              memory: 100Mi
50
           ports:
51
           - containerPort: 6379
52 ---
53 apiVersion: v1
```

```
54 kind: Service
55 metadata:
56
    name: redis-slave
57
    namespace: ns2
58
     labels:
59
       app: redis
       tier: backend
       role: slave
61
62 spec:
63
    ports:
64
     - port: 6379
    selector:
       app: redis
66
       tier: backend
67
       role: slave
68
69 ---
70 apiVersion: apps/v1 # for k8s versions before 1.9.0 use apps/
      v1beta2 and before 1.8.0 use extensions/v1beta1
71 kind: Deployment
72 metadata:
73
    name: redis-slave
74
     namespace: ns2
75 spec:
76
     selector:
      matchLabels:
77
78
         app: redis
79
         role: slave
         tier: backend
80
81
     replicas: 2
82
     template:
83
       metadata:
84
         labels:
           app: redis
85
86
           role: slave
87
           tier: backend
88
       spec:
89
         containers:
90
         - name: slave
91
           image: gcr.io/google_samples/gb-redisslave:v1
           resources:
92
              requests:
94
                cpu: 100m
95
                memory: 100Mi
96
           env:
           - name: GET_HOSTS_FROM
97
             value: dns
              # If your cluster config does not include a dns service,
99
                  then to
             # instead access an environment variable to find the
                 master
              # service's host, comment out the 'value: dns' line
                 above, and
              # uncomment the line below:
```

```
# value: env
104
           ports:
105
            - containerPort: 6379
106 ---
107 apiVersion: v1
108 kind: Service
109 metadata:
110
    name: frontend
    namespace: ns2
111
112 labels:
113 app: guestbook
114 tier: frontend
115 spec:
     # if your cluster supports it, uncomment the following to
116
         automatically create
117
      # an external load-balanced IP for the frontend service.
118 # type: LoadBalancer
119
    ports:
120
     - port: 80
121
      selector:
122
      app: guestbook
        tier: frontend
124
125 apiVersion: apps/v1 # for k8s versions before 1.9.0 use apps/
       vlbeta2 and before 1.8.0 use extensions/vlbeta1
126 kind: Deployment
127 metadata:
128
     name: frontend
129
     namespace: ns2
130 spec:
131
    selector:
132
      matchLabels:
133
        app: guestbook
134
         tier: frontend
135 replicas: 3
136
    template:
     metadata:
137
138
        labels:
139
            app: guestbook
140
            tier: frontend
141
       spec:
142
         containers:
          - name: php-redis
143
144
            image: gcr.io/google-samples/gb-frontend:v4
145
            resources:
146
              requests:
147
                cpu: 100m
148
                memory: 100Mi
149
           env:
            - name: GET_HOSTS_FROM
151
              value: dns
              # If your cluster config does not include a dns service,
152
                  then to
```

153	<pre># instead access environment variables to find service     host</pre>
154	<pre># info, comment out the 'value: dns' line above, and uncomment the</pre>
155	# line below:
156	# value: env
157	ports:
158	- containerPort: 80

2. Deploy NetScaler Ingress Controller in namespace ns2.

Ensure that you use the user credentials that are bound to the partition p1.

```
1 helm install cic-adm-part-p1 citrix/citrix-ingress-controller --
set nsIP=<nsIP of ADC>,nsSNIPS='[<SNIPs in partition p1>]',
license.accept=yes,adcCredentialSecret=admin-part-user-p1,
ingressClass[0]=citrix-adm-part-ns2 --namespace ns2
```

3. Deploy ingress for the guestbook application.

```
apiVersion: networking.k8s.io/v1
1
2 kind: Ingress
3 metadata:
4
    annotations:
5
     ingress.citrix.com/frontend-ip: "<VIP in partition 1>"
6
    name: guestbook-ingress
7
    namespace: ns2
8 spec:
    ingressClassName: citrix-adm-part-ns2
9
10
    rules:
     - host: www.guestbook.com
12
      http:
13
        paths:
14
         - backend:
15
            service:
16
              name: frontend
17
               port:
18
                 number: 80
19
           path: /
           pathType: Prefix
```

4. NetScaler Ingress Controller in ns2 configures the ADC entities in partition p1.

# Deploy Citrix solution for service of type LoadBalancer in AWS

### December 31, 2023

A service of type LoadBalancer is a simpler and faster way to expose a microservice running in a Kubernetes cluster to the external world. In cloud deployments, when you create a service of type LoadBalancer, a cloud managed load balancer is assigned to the service. The service is, then, exposed using the load balancer. For more information about services of type LoadBalancer, see Services of type LoadBalancer.

With the Citrix solution for service of type LoadBalancer, you can use NetScaler to directly load balance and expose a service instead of the cloud managed load balancer. NetScaler provides this solution for service of type LoadBalancer for on-prem and cloud. Services of type LoadBalancer are natively supported in Kubernetes deployments on public clouds such as AWS, GCP, and Azure.

When you deploy a service in AWS, a load balancer is created automatically and the IP address is allocated to the external field of the service. In this Citrix solution, allocates the IP address and that IP address is the VIP of NetScaler VPX. NetScaler Ingress Controller, deployed in a Kubernetes cluster, configures a NetScaler deployed outside the cluster to load balance the incoming traffic. So, the service is accessed through NetScaler VPX instead of the cloud load balancer.

You need to specify the service type as LoadBalancer in the service definition. Setting the type field to LoadBalancer provisions a load balancer for your service on AWS.

is used to automatically allocate IP addresses to services of type LoadBalancer from a specified range of IP addresses. For more information about the Citrix solution for services of type LoadBalancer, see Expose services of type LoadBalancer.

You can deploy the Citrix solution for service of type LoadBalancer in AWS using Helm charts or YAML files.

# Prerequisites

- Ensure that the Elastic Kubernetes Service (EKS) cluster version 1.18 or later is running.
- Ensure that NetScaler VPX and EKS are deployed and running in the same VPC. For information about creating NetScaler VPX in AWS, see Create a NetScaler VPX instance from AWS Market-place.

# Deploy Citrix solution for service of type LoadBalancer in AWS using Helm charts

Perform the following steps to configure the Citrix solution for service of type LoadBalancer using Helm charts.

- 1. Download the unified-lb-values.yaml file and edit the YAML file for specifying the following details:
  - NetScaler VPX NSIP. For more information, see NetScaler Ingress Controller Helm chart.
  - Secret created using the NetScaler VPX credentials. For more information, see NetScaler Ingress Controller Helm chart.

- List of VIPs to be used in IPAM controller. For more information, see IPAM Helm chart.
- 2. Deploy and NetScaler Ingress Controller on your Amazon EKS cluster using the edited YAML file. Use the following commands:

```
1 helm repo add citrix https://citrix.github.io/citrix-helm-charts/
2
3 helm install serviceLB citrix/citrix-cloud-native -f values.yaml
```

- 3. Deploy the application and service in Amazon EKS:
  - a) Add the following annotation in the service manifest:

1 beta.kubernetes.io/aws-load-balancer-type: "external"

 b) Deploy the application and service with the modified annotation using the following command:

```
1 kubectl create -f https://github.com/citrix/citrix-k8s-ingress
        -controller/blob/master/docs/how-to/typeLB/aws/guestbook-
        all-in-one-lb.yaml
```

**Note**: The guestbook microservice is a sample used in this procedure. You can deploy an application of your choice. Ensure that the service should be of type LoadBalancer and the service manifest should contain the annotation.

- c) Associate an elastic IP address with the VIP of NetScaler VPX.
- d) Access the application using a browser. For example, http://EIP-associatedwith-vip.

### Deploy Citrix solution for service of type LoadBalancer in AWS using YAML

Perform the following steps to deploy the Citrix solution for service of type LoadBalancer using YAML.

- 1. Download the citrix-k8s-ingress-controller.yaml file and specify the following details.
  - NetScaler VPX NSIP
  - Secret created using the NetScaler VPX credentials. For information about creating the secret, see Create a secret.
  - Specify the argument for :

```
1 args:
2 - --ipam
3 citrix-ipam-controller
```

2. Deploy the NetScaler Ingress Controller using the modified YAML.

1 kubectl create -f citrix-k8s-ingress-controller.yaml

3. Deploy the NetScaler VIP CRD which enables communication between the NetScaler Ingress Controller and the IPAM controller using the following command.

1 kubectl create -f https://raw.githubusercontent.com/citrix/citrixk8s-ingress-controller/master/crd/vip/vip.yaml

For more information about deploying NetScaler VIP CRD, see Deploy the VIP CRD.

4. Deploy the IPAM controller. For information about deploying the IPAM controller, see Deploy the IPAM controller.

Note:

Specify the list of NetScaler VPX VIPs in the VIP\_RANGE field of the IPAM deployment YAML file.

- 5. Deploy the application with service type LoadBalancer in Amazon EKS using the following steps:
  - a) Add the following annotation in the service manifest.

1 beta.kubernetes.io/aws-load-balancer-type: "external"

b) Deploy the application and service with the modified annotation using the following command.

```
1 kubectl create -f https://github.com/citrix/citrix-k8s-ingress
        -controller/blob/master/docs/how-to/typeLB/aws/guestbook-
        all-in-one-lb.yaml
```

#### Note:

The guestbook microservice is a sample used in this procedure. You can deploy an application of your choice. Ensure that the service should be of type LoadBalancer and the service manifest should contain the annotation.

- c) Associate an elastic IP address with the VIP of NetScaler VPX.
- d) Access the application using a browser. For example, http://EIP-associatedwith-vip.

# Multi-cloud and GSLB solution with Amazon EKS and Microsoft AKS clusters

### December 31, 2023

You can deploy multiple instances of the same application across multiple clouds provided by different cloud providers. This multi-cloud strategy helps you to ensure resiliency, high availability, and proximity. A multi-cloud approach also allows you to take advantage of the best of each cloud provider by reducing the risks such as vendor lock-in and cloud outages.

NetScaler with the help of the NetScaler Ingress Controller can perform multi-cloud load balancing. NetScaler can direct traffic to clusters hosted on different cloud provider sites. The solution performs load balancing by distributing the traffic intelligently between the workloads running on Amazon EKS (Elastic Kubernetes Service) and Microsoft AKS (Azure Kubernetes Service) clusters.

You can deploy the multi-cloud and GSLB solution with Amazon EKS and Microsoft AKS.

# **Deployment topology**

The following diagram explains a deployment topology of the multi-cloud ingress and load balancing solution for Kubernetes service provided by Amazon EKS and Microsoft AKS.



# Prerequisites

- You should be familiar with AWS and Azure.
- You should be familiar with NetScaler and NetScaler networking.
- Instances of the same application must be deployed in Kubernetes clusters on Amazon EKS and Microsoft AKS.

To deploy the multi-cloud and GSLB solution, you must perform the following tasks.

- 1. Deploy NetScaler VPX in AWS.
- 2. Deploy NetScaler VPX in Azure.
- 3. Configure ADNS service on NetScaler VPX deployed in AWS and AKS.
- 4. Configure GSLB service on NetScaler VPX deployed in AWS and AKS.
- 5. Apply GTP and GSE CRDs on AWS and Azure Kubernetes clusters.
- 6. Deploy the GSLB controller.

# **Deploying NetScaler VPX in AWS**

You must ensure that the NetScaler VPX instances are installed in the same virtual private cloud (VPC) on the EKS cluster. It enables NetScaler VPX to communicate with EKS workloads. You can use an existing EKS subnet or create a subnet to install the NetScaler VPX instances.

Also, you can install the NetScaler VPX instances in a different VPC. In that case, you must ensure that the VPC for EKS can communicate using VPC peering. For more information about VPC peering, see VPC peering documentation.

For high availability (HA), you can install two instances of NetScaler VPX in HA mode.

1. Install NetScaler VPX in AWS. For information on installing NetScaler VPX in AWS, see Deploy NetScaler VPX instance on AWS.

NetScaler VPX requires a secondary public IP address other than the NSIP to run GSLB service sync and ADNS service.

2. Open the AWS console and choose EC2 > Network Interfaces > VPX primary ENI ID > Manage IP addresses. Click Assign new IP Address.

lanage IP addresse	S Info
sign or unassign IPv4 and IPv	vb addresses to or from a network interface.
<ul> <li>To assign additional and associate them</li> </ul>	al public IPv4 addresses to this network interface, you must allocate Elastic IP addresses
eth0: eni-03bb2124d3	3ba6646d - Primary network interface - 192.168.192.0/19
eth0: eni-03bb2124d3	3ba6646d - Primary network interface - 192.168.192.0/19
eth0: eni-03bb2124d2	3ba6646d - Primary network interface - 192.168.192.0/19
<ul> <li>eth0: eni-03bb2124d3</li> <li>IPv4 addresses</li> <li>Private IP address</li> </ul>	3ba6646d - Primary network interface - 192.168.192.0/19 Public IP address
<ul> <li>eth0: eni-03bb2124d3</li> <li>IPv4 addresses</li> <li>Private IP address</li> <li>192.168.211.73</li> </ul>	3ba6646d - Primary network interface - 192.168.192.0/19         Public IP address         Unassign
<ul> <li>v eth0: eni-03bb2124d2</li> <li>IPv4 addresses</li> <li>Private IP address</li> <li>192.168.211.73</li> <li>192.168.201.3</li> </ul>	3ba6646d - Primary network interface - 192.168.192.0/19         Public IP address         Unassign         Unassign

After the secondary public IP address has been assigned to the VPX ENI, associate an elastic IP address to it.

3. Choose **EC2** > **Network Interfaces** > **VPX ENI ID** - **Actions**, click **Associate IP Address**. Select an elastic IP address for the secondary IP address and click **Associate**.

ociate an Elastic IP address with one of the private IPv4 addresses for the network interface.		
Association details		
Network interface		
eni-03bb2124d3ba6646d		
Elastic IP address		
3.128.55.34	•	
Private IPv4 address		
192.168.211.73	•	
Allow reassociation		
$\hfill \square$ Allow the Elastic IP address to be reassociated with this network interface		

4. Log in to the NetScaler VPX instance and add the secondary IP address as SNIP and enable the management access using the following command:

```
1 add ip 192.168.211.73 255.255.224.0 -mgmtAccess ENABLED -type SNIP
```

Note:

- To log in to NetScaler VPX using SSH, you must enable the SSH port in the security group. Route tables must have an internet gateway configured for the default traffic and the NACL must allow the SSH port.
- If you are running the NetScaler VPX in High Availability (HA) mode, you must perform this configuration in both of the NetScaler VPX instances.
- 5. Enable Content Switching (CS), Load Balancing (LB), Global Server Load Balancing(GSLB), and SSL features in NetScaler VPX using the following command:

1 enable feature \*feature\*

Note:

To enable GSLB, you must have an additional license.

6. Enable port 53 for UDP and TCP in the VPX security group for NetScaler VPX to receive DNS traffic. Also enable the TCP port 22 for SSH and the TCP port range 3008–3011 for GSLB metric exchange. For information on adding rules to the security group, see Adding rules to a security group.

7. Add a nameserver to NetScaler VPX using the following command:

```
1 add nameserver *nameserver IP*
```

### **Deploying NetScaler VPX in Azure**

You can run a standalone NetScaler VPX instance on an AKS cluster or run two NetScaler VPX instances in High Availability mode on the AKS cluster.

While installing, ensure that the AKS cluster must have connectivity with the VPX instances. To ensure the connectivity, you can install the NetScaler VPX in the same virtual network (VNet) on the AKS cluster in a different resource group.

While installing the NetScaler VPX, select the VNet where the AKS cluster is installed. Alternatively, you can use VNet peering to ensure the connectivity between AKS and NetScaler VPX if the VPX is deployed in a different VNet other than the AKS cluster.

1. Install NetScaler VPX in AWS. For information on installing NetScaler VPX in AKS, see Deploy a NetScaler VPX instance on Microsoft Azure.

You must have a SNIP with public IP for GSLB sync and ADNS service. If SNIP already exists, associate a public IP address with it.

2. To associate, choose **Home** > **Resource group** > **VPX instance** > **VPX NIC instance**. Associate a public IP address as shown in the following image. Click **Save** to save the changes.

Home > Multicluster-VPX > ns-vpx-express > ns-vpx-express-nic1 >	
snip	
ns-vpx-express-nic1	
🔚 Save 🗙 Discard	
Public IP address settings	
Public IP address	
Disassociate Associate	
Public IP address *	
(New) gslb	$\sim$
Create new	
Private IP address settings	
Virtual network/subnet	
aks-vnet-39725228/aks-subnet	
Assignment	
Dynamic Static	
IP address	
10.240.0.11	

3. Log in to the Azure NetScaler VPX instance and add the secondary IP as SNIP with the management access enabled using the following command:

1 add ip 10.240.0.11 255.255.0.0 -type SNIP -mgmtAccess ENABLED

If the resource exists, you can use the following command to set the management access enabled on the existing resource.

```
1 set ip 10.240.0.11 -mgmtAccess ENABLED
```

4. Enable CS, LB, SSL, and GSLB features in the NetScaler VPX using the following command:

1 enable feature \*feature\*

To access the NetScaler VPX instance through SSH, you must enable the inbound port rule for the SSH port in the Azure network security group that is attached to the NetScaler VPX primary interface.

- 5. Enable the inbound rule for the following ports in the network security group on the Azure portal.
  - TCP: 3008–3011 for GSLB metric exchange
  - TCP: 22 for SSH

- TCP and UDP: 53 for DNS
- 6. Add a nameserver to NetScaler VPX using the following command:

```
1 add nameserver *nameserver IP*
```

### Configure ADNS service in NetScaler VPX deployed in AWS and Azure

The ADNS service in NetScaler VPX acts as an authoritative DNS for your domain. For more information on the ADNS service, see Authoritative DNS service.

1. Log in to AWS NetScaler VPX and configure the ADNS service on the secondary IP address and port 53 using the following command:

1 add service Service-ADNS-1 192.168.211.73 ADNS 53

Verify the configuration using the following command:

1 show service Service-ADNS-1

2. Log in to Azure NetScaler VPX and configure the ADNS service on the secondary IP address and port 53 using the following command:

1 add service Service-ADNS-1 10.240.0.8 ADNS 53

Verify the configuration using the following command:

1 show service Service-ADNS-1

3. After creating two ADNS service for the domain, update the NS record of the domain to point to the ADNS services in the domain registrar.

For example, create an 'A'record nsl.domain.com pointing to the ADNS service public IP address. NS record for the domain must point to nsl.domain.com.

### Configure GSLB service in NetScaler VPX deployed in AWS and Azure

You must create GSLB sites on NetScaler VPX deployed on AWS and Azure.

 Log in to AWS NetScaler VPX and configure GSLB sites on the secondary IP address using the following command. Also, specify the public IP address using the *-publicIP* argument. For example:

```
1 add gslb site aws_site 192.168.197.18 -publicIP 3.139.156.175
2
3 add gslb site azure_site 10.240.0.11 -publicIP 23.100.28.121
```

2. Log in to Azure NetScaler VPX and configure GSLB sites. For example:

```
1 add gslb site aws_site 192.168.197.18 -publicIP 3.139.156.175
2
3 add gslb site azure_site 10.240.0.11 -publicIP 23.100.28.121
```

3. Verify that the GSLB sync is successful by initiating a sync from any of the sites using the following command:

```
1 sync gslb config - debug
```

### Note:

If the initial sync fails, review the security groups on both AWS and Azure to allow the required ports.

# Apply GTP and GSE CRDs on AWS and Azure Kubernetes clusters

The global traffic policy (GTP) and global service entry (GSE) CRDs help to configure NetScaler for performing GSLB in Kubernetes applications. These CRDs are designed for configuring NetScaler GSLB controller for applications deployed in distributed Kubernetes clusters.

# GTP CRD

The GTP CRD accepts the parameters for configuring GSLB on the NetScaler including deployment type (canary, failover, and local-first), GSLB domain, health monitor for the ingress, and service type.

For GTP CRD definition, see the GTP CRD. Apply the GTP CRD definition on AWS and Azure Kubernetes clusters using the following command:

```
1 kubectl apply -f https://raw.githubusercontent.com/citrix/citrix-k8s-
ingress-controller/master/gslb/Manifest/gtp-crd.yaml
```

### GSE CRD

The GSE CRD specifies the endpoint information (information about any Kubernetes object that routes traffic into the cluster) in each cluster. The global service entry automatically picks the external IP address of the application, which routes traffic into the cluster. If the external IP address of the routes change, the global service entry picks a newly assigned IP address and configure the GSLB endpoints of NetScalers accordingly.

For the GSE CRD definition, see the GSE CRD. Apply the GSE CRD definition on AWS and Azure Kubernetes clusters using the following command:

```
1 kubectl apply -f https://raw.githubusercontent.com/citrix/citrix-k8s-
ingress-controller/master/gslb/Manifest/gse-crd.yaml
```

### **Deploy GSLB controller**

GSLB controller helps you to ensure the high availability of the applications across clusters in a multicloud environment.

You can install the GSLB controller on the AWS and Azure clusters. GSLB controller listens to GTP and GSE CRDs and configures the NetScaler for GSLB that provides high availability across multiple regions in a multi-cloud environment.

To deploy the GSLB controller, perform the following steps:

1. Create an RBAC for the GSLB controller on the AWS and Azure Kubernetes clusters.

1 kubectl apply -f https://raw.githubusercontent.com/citrix/citrixk8s-ingress-controller/master/gslb/Manifest/gslb-rbac.yaml

2. Create the secrets on the AWS and Azure clusters using the following command:

```
Note:
```

Secrets enable the GSLB controller to connect and push the configuration to the GSLB devices.

1 kubectl create secret generic secret-1 --from-literal=username=< username> --from-literal=password=<password>

#### Note:

You can add a user to NetScaler using the add system user command.

- 3. Download the GSLB controller YAML file from gslb-controller.yaml.
- 4. Apply the gslb-controller.yaml in an AWS cluster using the following command:

1 kubectl apply -f gslb-controller.yaml

For the AWS environment, edit the gslb-controller.yaml to define the LOCAL\_REGION, LOCAL\_CLUSTER, and SITENAMES environment variables.

The following example defines the environment variable LOCAL\_REGION as *us-east-* 2 and LOCAL\_CLUSTER as *eks-cluster* and the SITENAMES environment variable as *aws\_site,azure\_site*.

```
1 name: "LOCAL_REGION"
2 value: "us-east-2"
3 name: "LOCAL_CLUSTER"
4 value: "eks-cluster"
5 name: "SITENAMES"
6 value: "aws_site,azure_site"
7 name: "aws_site_ip"
8 value: "NSIP of aws VPX(internal IP)"
9 name: "aws_site_region"
10 value: "us-east-2"
11 name: "azure_site_ip"
12 value: "NSIP of azure_VPX(public IP)"
13 name: "azure_site_region"
14 value: "central-india"
15 name: "azure_site_username"
16 valueFrom:
17
   secretKeyRef:
18
     name: secret-1
19
     key: username
20 name: "azure_site_password"
21 valueFrom:
22 secretKeyRef:
23
    name: secret-1
24 key: password
25 name: "aws_site_username"
26 valueFrom:
27
   secretKeyRef:
    name: secret-1
28
29
     key: username
30 name: "aws_site_password"
31 valueFrom:
32
   secretKeyRef:
33
     name: secret-1
34
   key: password
```

Apply the gslb-controller.yaml in the Azure cluster using the following command:

1 kubectl apply -f gslb-controller.yaml

5. For the Azure site, edit the gslb-controller.yaml to define LOCAL\_REGION, LOCAL\_CLUSTER, and SITENAMES environment variables.

The following example defines the environment variable LOCAL\_REGION as *central-india*, LOCAL\_CLUSTER as *azure-cluster*, and SITENAMES as *aws\_site*, *azure\_site*.

```
1 name: "LOCAL_REGION"
2 value: "central-india"
3 name: "LOCAL_CLUSTER"
4 value: "aks-cluster"
5 name: "SITENAMES"
6 value: "aws_site,azure_site"
7 name: "aws_site_ip"
```

```
8 value: "NSIP of AWS VPX(public IP)"
9 name: "aws_site_region"
10 value: "us-east-2"
11 name: "azure_site_ip"
12 value: "NSIP of azure VPX(internal IP)"
13 name: "azure_site_region"
14 value: "central-india"
15 name: "azure_site_username"
16 valueFrom:
   secretKeyRef:
17
18
    name: secret-1
19
     key: username
20 name: "azure_site_password"
21 valueFrom:
22
   secretKeyRef:
23
     name: secret-1
24
     key: password
25 name: "aws_site_username"
26 valueFrom:
27
   secretKeyRef:
28
     name: secret-1
29
      key: username
30 name: "aws_site_password"
31 valueFrom:
32
   secretKeyRef:
    name: secret-1
34
      key: password
```

### Note:

The order of the GSLB site information should be the same in all clusters. The first site in the order is considered as the master site for pushing the configuration. Whenever the master site goes down, the next site in the list becomes the new master. Hence, the order of the sites should be the same in all Kubernetes clusters.

### **Deploy a sample application**

In this example application deployment scenario, an https image of apache is used. However, you can choose the sample application of your choice.

The application is exposed as type LoadBalancer in both AWS and Azure clusters. You must run the commands in both AWS and Azure Kubernetes clusters.

1. Create a deployment of a sample apache application using the following command:

1 kubectl create deploy apache --image=httpd:latest port=80

2. Expose the apache application as service of type LoadBalancer using the following command:

1 kubectl expose deploy apache --type=LoadBalancer --port=80

3. Verify that an external IP address is allocated for the service of type LoadBalancer using the following command:

After deploying the application on AWS and Azure clusters, you must configure the GTE custom resource to configure high availability in the multi-cloud clusters.

Create a GTP YAML resource gtp\_isntance.yaml as shown in the following example.

```
1 apiVersion: "citrix.com/v1beta1"
    kind: globaltrafficpolicy
2
3
    metadata:
4
      name: gtp-sample-app
5
      namespace: default
6
   spec:
7
      serviceType: 'HTTP'
8
      hosts:
9
      - host: <domain name>
10
        policy:
          trafficPolicy: 'FAILOVER'
11
          secLbMethod: 'ROUNDROBIN'
12
13
          targets:
14
          - destination: 'apache.default.us-east-2.eks-cluster'
15
            weight: 1
          - destination: 'apache.default.central-india.aks-cluster'
16
            primary: false
17
18
             weight: 1
19
          monitor:
          - monType: http
            uri: ''
21
22
             respCode: 200
23
      status:
24
         {
25
     }
```

In this example, traffic policy is configured as FAILOVER. However, the multi-cluster controller supports multiple traffic policies. For more information, see the documentation for the traffic policies.

Apply the GTP resource in both the clusters using the following command:

1 kubectl apply -f gtp\_instance.yaml

You can verify that the GSE resource is automatically created in both of the clusters with the required endpoint information derived from the service status. Verify using the following command:

```
    kubectl get gse
    kubectl get gse *name* -o yaml
```

Also, log in to NetScaler VPX and verify that the GSLB configuration is successfully created using the following command:

1 show gslb runningconfig

As the GTP CRD is configured for the traffic policy as FAILOVER, NetScaler VPX instances serve the traffic from the primary cluster (EKS cluster in this example).

1 curl -v http://\*domain\_name\*

However, if an endpoint is not available in the EKS cluster, applications are automatically served from the Azure cluster. You can ensure it by setting the replica count to 0 in the primary cluster.

# NetScaler VPX as ingress and GSLB device for Amazon EKS and Microsoft AKS clusters

You can deploy the multi-cloud and multi-cluster ingress and load balancing solution with Amazon EKS and Microsoft AKS with NetScaler VPX as GSLB and the same NetScaler VPX as ingress device too.

To deploy the multi-cloud multi-cluster ingress and load balancing with NetScaler VPX as the ingress device, you must complete the following tasks described in the previous sections:

- 1. Deploy NetScaler VPX in AWS
- 2. Deploy NetScaler VPX in Azure
- 3. Configure ADNS service on NetScaler VPX deployed in AWS and AKS
- 4. Configure GSLB service on NetScaler VPX deployed in AWS and AKS
- 5. Apply GTP and GSE CRDs on AWS and Azure Kubernetes clusters
- 6. Deploy the GSLB controller

After completing the preceding tasks, perform the following tasks:

- 1. Configure NetScaler VPX as Ingress Device for AWS
- 2. Configure NetScaler VPX as Ingress Device for Azure

### **Configure NetScaler VPX as Ingress device for AWS**

Perform the following steps:

1. Create NetScaler VPX login credentials using Kubernetes secret

The NetScaler VPX password is usually the instance-id of the VPX if you have not changed it.
2. Configure SNIP in the NetScaler VPX by connecting to the NetScaler VPX using SSH. SNIP is the secondary IP address of Citrix a VPX to which the elastic IP address is not assigned.

1 add ns ip 192.168.84.93 255.255.224.0

This step is required for NetScaler to interact with the pods inside the Kubernetes cluster.

3. Update the NetScaler VPX management IP address and VIP in the NetScaler Ingress Controller manifest.

```
1 wget https://raw.githubusercontent.com/citrix/citrix-k8s-ingress-
controller/master/deployment/aws/quick-deploy-cic/manifest/cic.
yaml
```

Note:

If you do not have wget installed, you can use fetch or curl.

4. Update the primary IP address of NetScaler VPX in the cic.yaml in the following field.

```
1 # Set NetScaler NSIP/SNIP, SNIP in case of HA (mgmt has to be
enabled)
2 - name: "NS_IP"
3 value: "X.X.X.X"
```

5. Update the NetScaler VPX VIP in the cic.yaml in the following field. This is the private IP address to which you have assigned an elastic IP address

```
1 # Set NetScaler VIP for the data traffic
2 - name: "NS_VIP"
3 value: "X.X.X.X"
```

6. Once you have edited the YAML file with the required values deploy NetScaler Ingress Controller.

```
1 kubectl create -f cic.yaml
```

#### Configure NetScaler VPX as Ingress device for Azure

Perform the following steps:

1. Create NetScaler VPX login credentials using Kubernetes secrets.

```
1 kubectl create secret generic nslogin --from-literal=username='<
azure-vpx-instance-username>' --from-literal=password='<azure-
vpx-instance-password>'
```

#### Note:

The NetScaler VPX user name and password should be the same as the credentials set

while creating NetScaler VPX on Azure.

2. Using SSH, configure a SNIP in the NetScaler VPX, which is the secondary IP address of the NetScaler VPX. This step is required for the NetScaler to interact with pods inside the Kubernetes cluster.

```
1 add ns ip <snip-vpx-instance-private-ip> <vpx-instance-primary-ip-
subnet>
```

- snip-vpx-instance-private-ip is the dynamic private IP address assigned while adding a SNIP during the NetScaler VPX instance creation.
- vpx-instance-primary-ip-subnet is the subnet of the primary private IP address of the NetScaler VPX instance.

To verify the subnet of the private IP address, SSH into the NetScaler VPX instance and use the following command.

```
show ip <primary-private-ip-addess>
```

- 3. Update the NetScaler VPX image URL, management IP address, and VIP in the NetScaler Ingress Controller YAML file.
  - a) Download the NetScaler Ingress Controller YAML file.

```
1 wget https://raw.githubusercontent.com/citrix/citrix-k8s-
ingress-controller/master/deployment/azure/manifest/
azurecic/cic.yaml
```

Note:

If you do not have wget installed, you can use the fetch or curl command.

b) Update the NetScaler Ingress Controller image with the Azure image URL in the cic.yaml file.

```
1 - name: cic-k8s-ingress-controller
2 # CIC Image from Azure
3 image: "<azure-cic-image-url>"
```

c) Update the primary IP address of the NetScaler VPX in the cic.yaml with the primary private IP address of the Azure VPX instance.

```
1 # Set NetScaler NSIP/SNIP, SNIP in case of HA (mgmt has to be
enabled)
2 - name: "NS_IP"
3 value: "X.X.X.X"
```

d) Update the NetScaler VPX VIP in the cic.yaml with the private IP address of the VIP assigned during VPX Azure instance creation.

```
1 # Set NetScaler VIP for the data traffic
2 - name: "NS_VIP"
3 value: "X.X.X.X"
```

4. Once you have configured NetScaler Ingress Controller with the required values, deploy the NetScaler Ingress Controller using the following command.

1 kubectl create -f cic.yaml

# Annotations

June 26, 2025

### **Ingress annotations**

The following ingress annotations are supported by NetScaler:

Annotations	Туре	Required	Description	Default	Possible value
	_				Possible
Annotations	Туре	Required	Description	Default	value
ingress. citrix.com / multicluste -policy- priority- order	String	Optional	DescriptionWhen loadbalancingdifferentapplicationsin themulti-clusteringress setup,separatecontentswitchingpolicies arecreated foreachapplication.In such cases,if you requirea particularsequence forpolicybinding, youmust assign aprioritynumber to thecontentswitchingpolicies byusing theingress.citrix.com/multicluster-policy-priority-orderannotation.For more	NA	ingress.citrix.com/multiclu policy- priority-order: '{"frontend": { "80": "3", "9443": "1"}, "backend": "2"}'

# NetScaler ingress controller

					Possible
Annotations	Туре	Required	Description	Default	value
ingress.	String	Optional	Specify an IP	NA	Numeric IP
citrix.com	oung	optionat	address that		address. For
/frontend- ingress. ip citrix.com	String	Optional	needs to be Specify the used as the IPSET name content	NA	example, NetScaler 1.2.3.4 IPSET entity
/frontend-			that needs to		name
iβge€shame	String	Optional	beed indeo virtual server	http	http,tcp,
citrix.com			Pre content		udp,
/insecure-					sip_udp,or
service-			HTTTP/TCP/JP-		any
type			Bis HULDP/any Ways to Annotateon configure awing With configure awing With wintual server.		
ingress.	String	Optional	switching Configure the	80	Valid port
citrix.com			virtual server		number
/insecure-			lP address content		
port			such as IPAM Note: The configuration, iPSE actault nsVIP, that you and so on specify in the		
			p/sip <sub>tation</sub> /any		
			protagols.		
ingress.	String	Optional	Spesify be	ssl	ssl,
citrix.com			Pontigurad in		ssl_tcp
/secure-			NetStaler.		
service-			SSL/SSL_TCP		
type			as the		
			protocol for		
			content		
			switching		
			virtual server.		

Annotations	Туре	Required	Description	Default	Possible value
ingress. citrix.com /secure- port	String	Optional	Configure the port for content switching virtual server for HTTPS	443	Valid port number
ingress. citrix.com /insecure- termination	String	Optional	traffic. Configure the behavior for HTTP traffic. Use allow to allow HTTP traffic; use redirect to redirect the HTTP request to HTTPS; or use disallow if you want to drop the	disallow	allow, redirect, ordisallow
ingress.	String	Optional	HTTP traffic. Configure	NA	Possible HTTP
citrix.com /default- ingress. response- citrix.com code /secure- Baskena.	String/JSON String/JSON	Optional Optional	NetScaler to trigger an Specify if you HTTP want a secure response HTTPS code when a conneyton request lands	False NA	response codes are 404 As string: and 503. True/False. As JSON: As string: "
/backend- secret			Setwicate on NetScaler NetScale/Wand and if any of the following Rackendream- conditions are abbilication, if met for all the packend between backend backend		SUPOffee_Rame >secrets", .:{."} '. Service_Name >": "< kubernetes secret>", } '
© 1997–2025 Cit	rix Systems, Inc	. All rights reser	ved the ingress content YAML, If you switching want.different policies settings for		294

Annotations	Туре	Required	Description	Default	Possible value
		0			
ingress.	String/JSON	Optional	Specify the CA	NA	As string: "
citrix.com			certificate		kubernetes
/backend-			that you want		secret",
ca-secret			to use for		As JSON:
			backend com-		'{ "<
			munication		Service_Name
			between		>": "<
			NetScaler and		kubernetes
			Kubernetes		secret>",
			pods.		} '
ingress.	JSON	Optional	Specify	NA	One or more
citrix.com			already		NetScaler
1			existing SSL		sslcertkey
ingress preconfigur	-e <b>ł</b> \$ON	Optional	Configure the certificate	NA	Valid entity names
citrix.com -certkey			settings/para- keys on		NetScaler with
/lbvserver			meters of NetScaler that		entity
nigress.	JSON	Optional	NetScaler	NA	parameter in
citrix.com			LBVserver		key:value
/ ingress.			entity.		tonnet.
servicegrou	IP JSON	Optional	NetScaler Example: figure	NA	pärämeter in NetScaler
/monitor			Servicegeoup		key:value
/ 110111 201			entrial server.		format.
			EXample: "		
			Ingress ingress		key.vatue
			ġġţŗţĸ ebm		format.
			/ NetScaler		
			BEFOREESEUC	gd	
ingress.	String	Optional		NA	dsr
citrix.com			serve Example:	HASH	
/ ingress.	o		(DSR) jngress Specify the	<b>c</b> •	prefixor
deployment	String	Optional	Dolfeigenrausen Bathe Leduur	pretix	exact
/nath-			DaskerScater.×		exact.
ingress.	String	Ontional	SKHENLE Y/	NΔ	Value
citrix.com	String	optionat	applications satercalarity or		matching any
/ipam-			address ange		of the range
ingress. range	JSON	Optional	<b>Brample</b> of		names
citrix.com			denta franse		configured in
/external-			Specified to ""		IPAM
service			<b>MSServer</b> yon		controller
			Referencese		controller.
			Balanters."		
© 1997–2025 Ci	trix Systems, Inc	. All rights reser	ved.		295
			Example: is		
			han sing see		
			Backward com		
			7 i namection		

Annotations	Туре	Required	Description	Default	Possible value
ingress.	String	Optional	Specify the		
/canary- ingress. weight citrix.com	String	Optional	traffic to be Provide an directed to HTTP header	NA	
/canary-by ingress. -header citrix.com	List	Optional	key to direct Provide hffp traffic to the Header values canary	NA	List of header values as
/canary-by ingress. -header- citrix.com	String/JSON	Optional	Example: Versione See Traffic to the Panary X. com	NA	strings. As a string: " CRD_Instance_Name
/bot_crd ingress. citrix.com	String/JSON	Optional	version. See Boochon See Boochon this section. policies com	NA	". As JSON: Asa string: " CRD_Instance_Name Service_Name
ratelimit_c	erd		application'sy Refeatinity CRD citrix.com Batancing /canary-by Provise Server pader- See this value: '[" batancing value: '["		<pre>&gt;";" CRD_Instance_Name Service_Name &gt;":" CRD_Instance_Name " } '.</pre>
			/bot_crd:		
			botdefense " binds the policy to all		
			the services in the ingress or ingress.		
			citrix.com /bot_crd: '{ "		
			appname":		
			<pre>botdefense " } ' binds the policy to</pre>		
			only the front-end		

Annotations	Туре	Required	Description Default	Possible value
			<pre>Example: ingress. citrix.com / ratelimit_crd : " ratelimitexample " binds the policy to all the services in the ingress or ingress. citrix.com / ratelimit_crd : '{ " appname": "</pre>	
			<pre>ratelimitexample " } ' binds the policy to only frontend service.</pre>	
ingress. citrix.com /auth_crd	String/JSON	Optional	Bind the policies created byNAAuth CRD to the application's load balancing virtual server.See this section.	As a string: " CRD_Instance_Name ", As JSON: { "< Service_Name >":" CRD_Instance_Name "}

Annotations	Туре	Required	Description	Default	Possible value
			<pre>Example: ingress. citrix.com /auth_crd: " authexample " binds the policy to all the services in the ingress or ingress. citrix.com /auth_crd: '{ " appname": "</pre>		
			<pre>authexample " } ' binds the policy to only the front-end service.</pre>		
ingress. citrix.com /waf_crd	String/JSON	Optional	Bind the policies created by WAF CRD to the application's load balancing virtual server. See this section.	NA	As a string: "CRD_Instance_Name" , As JSON: '{ "< Service_Name >":" CRD_Instance_Name " } '

					Possible
Annotations	Туре	Required	Description	Default	value
			Example:		
			ingress.		
			citrix.com		
			/waf_crd:		
			"wafbasic"		
			binds the		
			policy to all		
			the services in		
			the ingress or		
			ingress.		
			citrix.com		
			/waf_crd:		
			'{ "		
			appname":		
			"wafbasic"		
			} ' binds		
			the policy to		
			only the		
			front-end		
			service		
ingress.	String/JSON	Optional	Bind the	NA	As a string: "
cıtrıx.com			policies		CRD_Instance_Name
/rewrite-			created by		", As JSON:
responder_c	rd		Rewrite-		'{ "<
			Responder		Service_Name
			CRD to the		>";"
			application's		
			load		· } · ·
			balancing		
			virtual server.		
			See this		
			section.		

	Toma	Demined	Description	Defeult	Possible
Annotations	Гуре	Required	Description	Default	value
			Example: ingress. citrix.com /rewrite- responder_co : " blockurlpol "Binds the policy to all the services in the ingress or ingress. citrix.com /rewrite- responder_co : '{ " appname": " blockurlpol	erd Licy erd	
			" } ' binds the policy to only the front-end service.	licy	
<pre>ingress. citrix.com /rewrite- responder_d</pre>	String/JSON	Optional	Bind the policies created by rewrite- responder CRD to the application's load balancing virtual server. See this section.	NA	As a string: " CRD_Instance_Name ". As JSON: '{ "< Service_Name >":" CRD_Instance_Name " } '.

					Possible
Annotations	Туре	Required	Description	Default	value
			Example:		
			ingress.		
			citrix.com		
			/rewrite-		
			responder_	crd	
			: "		
			blockurlpo	licy	
			" binds the		
			policy to all		
			the services in		
			the ingress or		
			ingress.		
			citrix.com		
			/rewrite-		
			responder_	crd	
			: '{ "		
			appname":		
			blockurlpo	licy	
			" } ' binds		
			the policy to		
			only the		
			front-end		
			service.		

#### **Service annotations**

The following are the service annotations supported by NetScaler.

In service annotations, index is the ordered index of the ports in a service specification file. For example, if there are two ports in the service specification, then the index for the first port is zero and for the second port is one.

Annotations	Туре	Required	Description	Default	Possible value
service. citrix.com /frontend- ip	String	Optional	Specify an IP adress that needs to be used as	NA	Numeric IP address, for example, '1,2,3,4'
© 1997–2025 Ci	trix Systems,	Inc. All rights reser	ved content switching virtual server IP address.		301

Annotations	Туре	Required	Description	Default	Possible value
service.	String	Optional	Select a	NA	Value
citrix.com			particular IP		matching any
/ipam- service.		Outload	address range Redirect HTTP		of the range
range citrix.com	JSON	Optional	from a set of traffic to a	NA	names
/insecure-			ranges secure port.		configured in
service redirect citrix.com	String	Optional	Specificente Example: SetScaler Service.	NA	EBUE and REEMOREPT
/ssl- service.	String	Optional	tennination.	НТТР	TCP, HTTP,
citrix.com			protocol fore-		SSL,UDP,ANY
/service- service. type-< citrix.com index>	String	Optional	the viscal of the second secon	NA	, SSL_TCP, Certificate and Data in PEM SIP_UDP. Format
service certificate citrix.com	String	Optional	value in the E ' value in the E '	NA	Key data in PEM Format
/ssl-key-			zample:		
data-<	String	Optional	BEMVORDEAL.	NA	data in DEM
index>					Gormat
service.	String	Optional	Specify the CA	NA	CA certificate
citrix.com			fer light cate		data in PEM
/ssl-			value to verify		format Kubernetes
Băckend-ca	String	Optional	fleserver \  REMIormat.	NA	secret Name
/secret			certificate eate		secretitame
sertitecate	String	Optional	Fré hack end resource for	NA	Kubernetes
ਫ਼ਖ਼ੑੑਗ਼ੑਫ਼ਖ਼ਙਫ਼ਖ਼	C	·	the triffeate for the front-end		secret Name
709 <sup>e</sup> šècret			Millie ARACHICA-		Kubernetes
citrix.com	String	Optional	sentitistion Pte	NA	secret Name
/backend-			hentication. Hierback eale		
service. secret	String	Optional	Change server	NA	Kubernetes
citrix.com	-		<b>Settlente</b> aste		secret Name
/backend-			setsteres.		NetScaler
că-secret	String	Optional		NA	ssicertkey
/					entity name
, service preconfigur	eString	Optional		NA	NetScaler
citrix.com -certkey		·	Sector Stephesete		sslcertkey
/			cantieure di for		entity name
preconfigur	ed				
©9997–2025 Ci	trix Systems,	Inc. All rights reser	ved to the for		302
certkey			by periaeo las		
			certificate, to		
			pind multiple		

certkey

					Possible
Annotations	Туре	Required	Description	Default	value
service.	String	Optional	Specify the	NA	NetScaler
/ service. preconfigur citrix.com	e\$tring	Optional	configured Specify the certificate key name of a pre-	NA	entity name NetScaler sslcertkey
-backend- / certkey preconfigur	ed		in NetScaler configured CA to be bound certificate key to the		entity name
-backend- Smart annota ca-certkey	tions for HTTI	P, TCP, or SSL	In NetScaler profiles-end SSL to bound to service group. the back-end		
Annotations	Туре	Required	SSL service SSL service Description Protector	Default	Possible value
ingress. citrix.com /frontend- httpprofile	String/JSON	Optional	server authen- server authen- server authen- server authen- thorsend Famospariation famospariati	NA red	<pre>Example: ingress. citrix.com /frontend- httpprofile : '{ " dropinvalreq ":"enabled ", " websocket" : "</pre>
			certkey: ' coffee-ca- cert'		enabled" } '

Annotations	Туре	Required	Description	Default	Possible value
ingress. citrix.com /backend- httpprofile	String/JSON	Optional	Create the back-end HTTP profile (Server Plane).	NA	<pre>Example: ingress. citrix.com /backend- httpprofile : '{ "app -1": { " dropinvalreqs ":"enabled ", " websocket" : " enabled" } } '</pre>
ingress. citrix.com /frontend- tcpprofile	String/JSON	Optional	Create the front-end TCP profile (Client Plane)	NA	<pre>Example: ingress. citrix.com /frontend- tcpprofile : '{ "ws ":"enabled ", "sack" : "enabled " } '</pre>
ingress. citrix.com /backend- tcpprofile	String/JSON	Optional	Create the back-end TCP profile (Server Plane)	NA	<pre>Example:ingress .citrix. com/ backend- tcpprofile : '{ " citrix-svc ":{ "ws":" enabled", "sack" : " enabled" } } '</pre>

	_				Possible
Annotations	Туре	Required	Description	Default	value
ingress. citrix.com /frontend- sslprofile	String/JSON	Optional	Create the front-end SSL profile (Client Plane). The front and SSL	NA	Example: ingress. citrix.com /frontend-
			profile is required only if you have enabled TLS on the Client Plane.		: '{ "hsts ":"enabled ", "tls12" : " enabled" }
ingress. citrix.com /backend- sslprofile	String/JSON	Optional	Create the back-end SSL profile (Server Plane). The SSL back-end profile is required only if you use ingress. citrix.com /secure- backend.	NA	<pre>Example: ingress. citrix.com /backend- sslprofile : '{ " citrix-svc ":{ "hsts ":"enabled ", "tls1" : "enabled " } } '</pre>

## **Smart annotations for Ingress**

Smart annotation is an option provided by NetScaler Ingress Controller to efficiently enable NetScaler features using the NetScaler entity name. The NetScaler Ingress Controller converts the Ingress in Kubernetes to a set of NetScaler objects. You can efficiently control these objects using smart annotations.

Note

To use smart annotations, you must have a good understanding of NetScaler features and their respective entity names. For more information about NetScaler features and entity names, see NetScaler documentation.

Smart annotation takes JSON format as input. The key and value that you pass in the JSON format

must match the NetScaler NITRO format. For more information about the NetScaler NITRO API, see NetScaler REST APIs - NITRO documentation.

For example, if you want to enable the SRCIPDESTIPHASH based lb method, you must use the corresponding NITRO key and value format lbmethod, SRCIPDESTIPHASH respectively.

The following table details the smart annotations provided by NetScaler Ingress Controller:

NetScaler Entity Name	Smart Annotation	Example
lbvserver	ingress.citrix.com/ lbvserver	<pre>ingress.citrix.com/ lbvserver: '{ " appname":{ "lbmethod ":"SRCIPDESTIPHASH" } } '</pre>
servicegroup	ingress.citrix.com/ servicegroup	<pre>ingress.citrix.com/ servicegroup: '{ " appname":{ "cip": " Enabled","cipHeader ":"X-Forwarded-For" } } '</pre>
monitor	ingress.citrix.com/ monitor	<pre>ingress.citrix.com/ monitor: '{ "appname ":{ "type":"http" } } '</pre>
csvserver	ingress.citrix.com/ csvserver	<pre>ingress.citrix.com/ csvserver: '{ " stateupdate": " ENABLED" } '</pre>

For information on smart annotations for HTTP, TCP, and SSL profiles, see Configure HTTP, TCP, or SSL profiles on NetScaler.

#### Sample ingress YAML with smart annotations

The following sample Ingress YAML includes smart annotations to enable NetScaler features using the entities such as, lbvserver, servicegroup, and monitor:

```
1 kubectl apply -f - <<EOF
2 apiVersion: networking.k8s.io/v1
3 kind: Ingress</pre>
```

#### NetScaler ingress controller

```
4 metadata:
5
   annotations:
       ingress.citrix.com/frontend-ip: 192.168.1.1
6
7
       ingress.citrix.com/insecure-port: "80"
       ingress.citrix.com/lbvserver: '{
8
   "appname":{
9
    "lbmethod":"LEASTCONNECTION", "persistenceType":"SOURCEIP" }
10
11
    }
12
    1
13
       ingress.citrix.com/monitor: '{
    "appname":{
14
15
    "type":"http" }
16
    }
    ۲.
17
18
       ingress.citrix.com/servicegroup: '{
19
    "appname":{
20
    "usip":"yes" }
21
    }
    1
22
23
    name: citrix
24 spec:
25
    rules:
26
     - host: citrix.org
27
       http:
28
         paths:
29
         - backend:
             service:
31
               name: appname
32
               port:
33
                 number: 80
34
           path: /
35
           pathType: Prefix
36 EOF
```

The sample Ingress YAML includes use cases related to the service, citrix-svc, and the following table explains the smart annotations used in the sample:

Smart Annotation	Description
<pre>ingress.citrix.com/lbvserver: '{ "appname":{ "lbmethod":" LEASTCONNECTION", " persistenceType":"SOURCEIP" } } '</pre>	Sets the load balancing method as Least Connection and also configures Source IP address persistence.
<pre>ingress.citrix.com/servicegroup: '{ "appname":{ "usip":"yes" } } '</pre>	Enables Use Source IP Mode (USIP) on NetScaler device. When you enable USIP on NetScaler, it uses the client's IP address for communication with the back-end pods.

NetScaler ingress controller	
Smart Annotation	Description
<pre>ingress.citrix.com/monitor: '{ " appname":{ "type":"http" } } '</pre>	Creates a custom HTTP monitor for the service group.

#### Note:

When multiple ingresses are sharing the same front-end IP address and port, you cannot have conflicting configurations provided through multiple ingress configurations.

By default, the content switching virtual server does not depend on the state of the target load balancing virtual servers bound to it. The annotation ingress.citrix.com/csvserver: '{ " stateupdate": "ENABLED" } ' sets the content switching virtual server to consider its state based on the state of the load balancing virtual server bound to it using the content switching policies.

#### Smart annotations for routes

Similar to Ingress, you can also use smart annotations with OpenShift routes. NetScaler Ingress Controller converts the routes in OpenShift to a set of NetScaler objects.

The following table details the smart annotations provided by NetScaler Ingress Controller:

NetScaler entity name	Smart annotation	Example
lbvserver	route.citrix.comp(DDDBESTEPHASH)	<pre>route.citrix.com/lbvserver: '{ ' } } '</pre>
servicegroup	route.citraixplocand/eselty/iX-egonowapro	<pre>route.citrix.com/servicegroup: led-For" } } '</pre>
monitor	route.citrix.com/monitor	<pre>route.citrix.com/monitor: '{ "</pre>

#### Sample route manifest with smart annotations

The following example is a route YAML file.

```
1 kubectl apply -f - <<EOF
2 apiVersion: route.openshift.io/v1
3 kind: Route
4 metadata:
5 name: citrix
6 annotations:
7 route.citrix.com/lbvserver: '{
```

```
NetScaler ingress controller
```

```
"appname":{
8
9
    "lbmethod":"LEASTCONNECTION", "persistenceType":"SOURCEIP" }
10
    }
    1
11
       route.citrix.com/servicegroup: '{
12
13
    "appname":{
    "usip":"yes" }
14
15
    }
    1
16
17
       route.citrix.com/monitor: '{
    "appname":{
18
19
    "type":"http" }
    }
    ŧ.
21
22 spec:
23
   host: citrix.org
24
   port:
25
       targetPort: 80
26 to:
27
       kind: Service
28
      name: appname
29
      weight: 100
   wildcardPolicy: None
31 EOF
```

The sample route manifest includes use cases related to the service citrix-svc and the following table explains the smart annotations used in the sample route:

Smart annotation	Description
<pre>route.citrix.com/lbvserver: '{ " appname":{ "lbmethod":" LEASTCONNECTION", " persistenceType":"SOURCEIP" } } '</pre>	Sets the load balancing method as Least Connection and also configures Source IP address persistence.
<pre>route.citrix.com/servicegroup: '{   "appname":{ "usip":"yes" } } '</pre>	Enables Use Source IP Mode (USIP) on NetScaler. When you enable USIP on the NetScaler, it uses the IP address of the client for communication with the back-end pods.
<pre>route.citrix.com/monitor: '{ " appname":{ "type":"http" } } '</pre>	Creates a custom HTTP monitor for the service group.

## Sample YAML with the service annotation to redirect insecure traffic

This example shows how to redirect traffic from clients making requests on an insecure port 80 to the secure port 443.

The following annotation is specified in the service YAML file to redirect traffic:

```
1 service.citrix.com/insecure-redirect: '{
2 "port-443": 80 }
3 '
```

Sample service definition:

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: v1
3 kind: Service
4 metadata:
5
    name: frontend-service
6
   annotations:
7
      service.citrix.com/service-type-0: SSL
8
      service.citrix.com/frontend-ip: '192.2.170.26'
9
      service.citrix.com/secret: '{
10 "port-443": "web-ingress-secret" }
11
       service.citrix.com/ssl-termination-0: 'EDGE'
12
13
       service.citrix.com/insecure-redirect: '{
   "port-443": 80 }
14
   1.1
16 spec:
17
   type: LoadBalancer
18
  selector:
19
     app: frontend
20 ports:
    - port: 443
21
      targetPort: 80
22
23
      name: port-443
24 EOF
```

## Smart annotations for services

Smart annotations for services are used to configure NetScaler with custom values for NetScaler configuration parameters. The annotations are used for services of type LoadBalancer and for the services in NetScaler CPX used for East-West traffic.

Note:

If you have configured a service with NodePort or ClusterIP for the North-South traffic, then NetScaler is configured using the applicable ingress smart annotations rather than service annotations.

Smart annotations for services take JSON format as input. The key and value that you pass in the JSON format must match the NetScaler NITRO format. For more information about the NetScaler NITRO API, see NetScaler REST APIs - NITRO Documentation.

Example smart annotation for services:

```
1 service.citrix.com/lbvserver: '{
2 "80-tcp":{
3 "lbmethod":"SRCIPDESTIPHASH" }
4 }
5 '
```

This annotation sets the load balancing method as SRCIPDESTIPHASH in the load balancing virtual server for the 80-tcp port of the given service.

NetScaler Entity Name	Smart Annotation for Service	Example
lbvserver	service.citrix.com/ lbvserver	<pre>service.citrix.com/ lbvserver: '{ "80-tcp ":{ "lbmethod":" SRCIPDESTIPHASH" } } '</pre>
csvserver	service.citrix.com/ csvserver	<pre>service.citrix.com/ csvserver: '{ "l2conn ":"on" } '</pre>
servicegroup	service.citrix.com/ servicegroup	<pre>service.citrix.com/ servicegroup: '{ "80- tcp":{ "usip":"yes" } } '</pre>
monitor	service.citrix.com/ monitor	<pre>service.citrix.com/ monitor: '{ "80-tcp ":{ "type":"http" } } ', service.citrix.com/ monitor: '{ "80-tcp":    "<pre-configured- monitor1-on-netscaler="">", "443-tcp": "<pre- configured-monitor2-="" on-netscaler="">" } '</pre-></pre-configured-></pre>

The following table describes the smart annotations for services:

NetScaler Entity Name	Smart Annotation for Service	Example
analyticsprofile	service.citrix.com/ analyticsprofile	<pre>service.citrix.com/ analyticsprofile: '{ "80-tcp":{ " webinsight": { " httpurl":"ENABLED", " httpuseragent":" ENABLED" } } '</pre>

You can use the smart annotations for services as follows:

- By providing the port-protocol value in the annotation: In the service definition, if you provide the port-protocol value in the annotation then the annotation is restricted to the particular port of that service.
- By not providing the port-protocol value in the annotation: If you do not provide the port
   -protocol value in the annotation, then the annotation is applicable to all the ports used by
   the service.

#### Sample ingress YAML with smart annotations for services

The following YAML is a sample deployment and service definition for a basic apache web-server based application. It includes smart annotations for services to enable NetScaler features using the entities such as lbvserver, csvserver, servicegroup, monitor, and analyticsprofile:

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: apps/v1
3 kind: Deployment
4 metadata:
5
   name: apache
6
    labels:
7
         name: apache
8 spec:
9 selector:
10
     matchLabels:
11
        app: apache
   replicas: 8
13
   template:
14
     metadata:
15
        labels:
16
           app: apache
17
      spec:
18
        containers:
19
         - name: apache
20
           image: httpd:latest
```

```
ports:
21
22
           - name: http
23
            containerPort: 80
24
           imagePullPolicy: IfNotPresent
25
26 ---
27 #Expose the apache web server as a service
28 apiVersion: apps/v1
29 kind: Service
30 metadata:
31 name: apache
32 annotations:
33
      service.citrix.com/csvserver: '{
   "l2conn":"on" }
34
   1.1
       service.citrix.com/lbvserver: '{
37
    "80-tcp":{
    "lbmethod":"SRCIPDESTIPHASH" }
38
39
    }
    1
40
41
       service.citrix.com/servicegroup: '{
    "80-tcp":{
42
    "usip":"yes" }
43
44
    }
    \mathbf{t}_{i}
45
46
       service.citrix.com/monitor: '{
47
    "80-tcp":{
    "type":"http" }
48
49
    }
50
51
       service.citrix.com/frontend-ip: '10.217.212.16'
52
       service.citrix.com/analyticsprofile: '{
    "80-tcp":{
53
54
    "webinsight": {
    "httpurl":"ENABLED", "httpuseragent":"ENABLED" }
55
56
    }
57
    }
    ŧ.
58
59
       NETSCALER_VPORT: '80'
     labels:
61
       name: apache
62 spec:
    externalTrafficPolicy: Local
64 type: LoadBalancer
   selector:
65
66
      name: apache
   name
ports:
67
68
   - name: http
69
      port: 80
70
      targetPort: http
71
    selector:
72
     app: apache
73
```

74 EOF

### Examples

#### Sample ingress YAML for SIP\_UDP support in insecure service type annotation

The following sample ingress YAML includes the configuration for enabling SIP over UDP support using the ingress.citrix.com/insecure-service-type annotation.

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: networking.k8s.io/v1
3 kind: Ingress
4 metadata:
5 annotations:
       ingress.citrix.com/frontend-ip: 1.1.1.1
6
       ingress.citrix.com/insecure-port: "5060"
7
8
       ingress.citrix.com/insecure-service-type: sip_udp
9
       ingress.citrix.com/lbvserver: '{
   "asterisk17":{
10
    "lbmethod":"CALLIDHASH","persistenceType":"CALLID" }
11
12
    }
    t.
13
14
   name: sip-ingress
15 spec:
    defaultBackend:
16
17
       service:
18
         name: asterisk17
19
         port:
          number: 5060
21 EOF
```

# ConfigMap support for the NetScaler Ingress Controller

December 31, 2023

The ConfigMap API resource holds key-value pairs of configuration data that can be consumed in pods or to store configuration data for system components such as controllers.

ConfigMaps allow you to separate your configurations from your pods and make your workloads portable. Using ConfigMaps, you can easily change and manage your workload configurations and reduce the need to hardcode configuration data to pod specifications.

The NetScaler Ingress Controller supports the configuration command line arguments, and environment variables mentioned in deploying the NetScaler Ingress Controller. But, you cannot update these configurations at runtime without rebooting the NetScaler Ingress Controller pod. With ConfigMap support, you can update the configuration automatically while keeping the NetScaler Ingress Controller pod running. You do not need to restart the pod after the update.

# Supported environment variables in the NetScaler Ingress Controller

The values for the following environment variables in the NetScaler Ingress Controller can be specified in a ConfigMap.

- LOGLEVEL: Specifies the log levels to control the logs generated by the NetScaler Ingress Controller (debug, info, critical, and so on). The default value is debug.
- NS\_HTTP2\_SERVER\_SIDE: Enables HTTP2 for NetScaler service group configurations with possible values as ON or OFF.
- NS\_PROTOCOL: Specifies the protocol to establish the ADC session (HTTP/HTTPS). The default value is http.
- NS\_PORT: Specifies the port to establish a session. The default value is 80.
- NS\_COOKIE\_VERSION: Specifies the persistence cookie version (0 or 1). The default value is 0.
- NS\_DNS\_NAMESERVER: Enables adding DNS nameservers on NetScaler VPX.
- POD\_IPS\_FOR\_SERVICEGROUP\_MEMBERS: Specifies to add the IP address of the pod and port as service group members instead of NodeIP and NodePort while configuring services of type LoadBalancer or NodePort on an external tier-1 NetScaler.
- IGNORE\_NODE\_EXTERNAL\_IP: Specifies to ignore an external IP address and add an internal IP address for NodeIP while configuring NodeIP for services of type LoadBalancer or NodePort on an external tier-1 NetScaler.
- FRONTEND\_HTTP\_PROFILE: Sets the HTTP options for the front-end virtual server (client plane), unless overridden by the ingress.citrix.com/frontend-httpprofile smart annotation in the ingress definition.
- FRONTEND\_TCP\_PROFILE: Sets the TCP options for the front-end virtual server (client side), unless overridden by the ingress.citrix.com/frontend-tcpprofile smart annotation in the ingress definition.
- FRONTEND\_SSL\_PROFILE: Sets the SSL options for the front-end virtual server (client side) unless overridden by the ingress.citrix.com/frontend-sslprofile smart annotation in the ingress definition.
- JSONLOG: Set this argument to true if log messages are required in JSON format.
- NS\_ADNS\_IPS: Enables configuring NetScaler as an ADNS server.

For more information about profile environment variables (FRONTEND\_HTTP\_PROFILE, FRON-TEND\_TCP\_PROFILE, and FRONTEND\_SSL\_PROFILE), see Configure HTTP, TCP, or SSL profiles on NetScaler.

Note:

This is an initial version of the ConfigMap support and currently supports only a few parameters. Earlier, these parameters were configurable through environment variables except the NS\_HTTP2\_SERVER\_SIDE parameter.

# Configuring ConfigMap support for the NetScaler Ingress Controller

This example shows how to create a ConfigMap and apply the ConfigMap to the NetScaler Ingress Controller. It also shows how to reapply the ConfigMap after you make changes. You can also optionally delete the changes.

Perform the following to configure ConfigMap support for the NetScaler Ingress Controller.

1. Create a YAML file cic-configmap.yaml with the required key-value pairs in the ConfigMap.

```
1 apiVersion: v1
2 kind: ConfigMap
3 metadata:
4
    name: cic-configmap
5
     labels:
       app: citrix-ingress-controller
6
7 data:
8 LOGLEVEL: 'info'
9 NS_PROTOCOL: 'http'
10 NS PORT: '80'
    NS_COOKIE_VERSION: '0'
11
    NS_HTTP2_SERVER_SIDE: 'ON'
```

2. Deploy the cic-configmap.yaml using the following command.

kubectl create -f cic-configmap.yaml

3. Edit the cic.yaml file for deploying the NetScaler Ingress Controller as a stand-alone pod and specify the following:

```
1 Args:
2 ---configmap
3 default/cic-configmap
```

#### Note:

It is mandatory to specify the namespace. If the namespace is not specified, ConfigMap is not considered.

Following is a sample YAML file for deploying the NetScaler Ingress Controller with the ConfigMap configuration. For the complete YAML file, see citrix-k8s-ingress-controller.yaml.

```
apiVersion: apps/v1
1
2
  kind: Deployment
3 metadata:
    name: cic-k8s-ingress-controller
4
5 spec:
    selector:
6
7
      matchLabels:
8
         app: cic-k8s-ingress-controller
9
      replicas: 1
10
      template:
        metadata:
11
12
          name: cic-k8s-ingress-controller
13
           labels:
14
             app: cic-k8s-ingress-controller
15
         annotations:
16
         spec:
17
           serviceAccountName: cic-k8s-role
18
           containers:
19
           - name: cic-k8s-ingress-controller
20
             image: "quay.io/citrix/citrix-k8s-ingress-controller
                 :1.36.5"
21
             env:
             # Set NetScaler NSIP/SNIP, SNIP in case of HA (mgmt has
22
                 to be enabled)
23
             - name: "NS_IP"
               value: "x.x.x.x"
24
25
             - name: "EULA"
26
               value: "yes"
27
             args:
28
               - --ingress-classes
29
                  citrix
               - --feature-node-watch
31
                  false
32
                - --configmap
                  default/cic-configmap
34
             imagePullPolicy: Always
```

4. Deploy the NetScaler Ingress Controller as a stand-alone pod by applying the YAML.

1 kubectl apply -f cic.yaml

5. If you want to change the value of an environment variable, edit the values in the ConfigMap. In this example, the value of NS\_HTTP2\_SERVER\_SIDE is changed to 'OFF'.

```
1 apiVersion: v1
2 kind: ConfigMap
3 metadata:
4 name: cic-configmap
5 labels:
6 app: citrix-ingress-controller
```

```
7 data:
8 LOGLEVEL: 'info'
9 NS_PROTOCOL: 'http'
10 NS_PORT: '80'
11 NS_COOKIE_VERSION: '0'
12 NS_HTTP2_SERVER_SIDE: 'OFF'
```

6. Reapply the ConfigMap using the following command.

1 kubectl apply -f cic-configmap.yaml

7. (Optional) If you need to delete the ConfigMap, use the following command.

```
1 kubectl delete -f cic-configmap.yaml
```

When you delete the ConfigMap, the environment variable configuration falls back as per the following order of precedence:

ConfigMap configuration > environment variable configuration > default

(Optional) In case, you want to define all keys in a ConfigMap as environment variables in the NetScaler Ingress Controller, use the following in the NetScaler Ingress Controller deployment YAML file.

```
    envFrom:
    configMapRef:
    name: cic-configmap
```

# **Ingress configurations**

#### December 31, 2023

Kubernetes Ingress provides you a way to route requests to services based on the request host or path, centralizing a number of services into a single entry point.

NetScaler Ingress Controller is built around the Kubernetes Ingress and automatically configures one or more NetScaler based on the Ingress resource configuration.

#### Host name based routing

The following sample Ingress definition demonstrates how to set up an Ingress to route the traffic based on the host name:

```
    apiVersion: networking.k8s.io/v1
    kind: Ingress
    metadata:
    name: virtual-host-ingress
```

```
5 namespace: default
6 spec:
7
   rules:
8 - host: foo.bar.com
9
     http:
      paths:
11
        - backend:
12
            service:
13
             name: service1
14
              port:
15
               number: 80
          pathType: Prefix
16
17
          path: /
   - host: bar.foo.com
18
     http:
19
     paths:
20
        - backend:
21
22
           service:
23
             name: service2
              port:
24
              number: 80
25
26
          pathType: Prefix
27
          path: /
```

After the sample Ingress definition is deployed, all the HTTP request with a host header is load balanced by NetScaler to service1. And, the HTTP request with a host header is load balancer by NetScaler to service2.

# Path based routing

The following sample Ingress definition demonstrates how to set up an Ingress to route the traffic based on URL path:

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
4 name: path-ingress
   namespace: default
5
6 spec:
7
    rules:
     - host: test.example.com
8
     http:
9
         paths:
         - backend:
11
12
            service:
13
             name: service1
14
               port:
                number: 80
15
           path: /foo
16
17
           pathType: Prefix
         - backend:
18
```

```
19service:20name: service221port:22number: 8023path: /24pathType: Prefix
```

After the sample Ingress definition is deployed, any HTTP requests with host test.example.com and URL path with prefix / foo, NetScaler routes the request to service1 and all other requests are routed to service2.

NetScaler Ingress Controller follows first match policy to evaluate paths. For effective matching, NetScaler Ingress Controller orders the paths based on descending order of the path's length. It also orders the paths that belong to same hosts across multiple ingress resources.

# Wildcard host routing

The following sample Ingress definition demonstrates how to set up an ingress with wildcard host.

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
4
   name: wildcard-ingress
5 namespace: default
6 spec:
7
   rules:
8
    - host: '*.example.com'
9
       http:
10
         paths:
11
         - backend:
12
             service:
13
              name: service1
14
               port:
15
                 number: 80
           path: /
16
17
           pathType: Prefix
```

After the sample Ingress definition is deployed, HTTP requests to all the subdomains of example. com is routed to service1 by NetScaler.

#### Note:

Rules with non-wildcard hosts are given higher priority than wildcard hosts. Among different wildcard hosts, rules are ordered on the descending order of length of hosts.

# **Exact path matching**

Ingresses belonging to networking.k8s.io/v1 APIversion can make use of PathType: Exact to consider the path for the exact match.

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
   name: Path-exact-Ingress
4
   namespace: default
5
6 spec:
7
   rules:
    - host: test.example.com
8
9
     http:
10
         paths:
         - backend:
11
12
             service:
13
               name: service1
14
               port:
15
                 name: 80
16
           path: /exact
17
           pathType: Exact
```

(Deprecated as of Kubernetes 1.22+) By default for Ingresses belonging to extension/vlbeta1 , paths are treated as Prefix expressions. Using the annotation ingress.citrix.com/path -match-method: "exact" in the ingress definition defines the NetScaler Ingress Controller to consider the path for the exact match.

The following sample Ingress definition demonstrates how to set up Ingress for exact path matching:

```
apiVersion: extension/v1beta1
1
2 kind: Ingress
3 metadata:
4 name: path-exact-ingress
5 namespace: default
6 annotations:
7
      ingress.citrix.com/path-match-method: "exact"
8 spec:
9
   rules:
    - host:test.example.com
10
11
     http:
12
        paths:
13
         - path: /exact
14
           backend:
15
            serviceName: service1
16
             servicePort: 80
```

After the sample Ingress definition is deployed, HTTP requests with path /exact is routed by NetScaler to service1 but not to /exact/somepath.

## **Non-Hostname routing**

Following example shows path based routing for the default traffic that does not match any host based routes. This ingress rule applies to all inbound HTTP traffic through the specified IP address.

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
    name: default-path-ingress
4
5
    namespace: default
6 spec:
    rules:
7
     - http:
8
         paths:
9
10
         - backend:
11
             service:
12
               name: service1
13
               port:
14
                 number: 80
15
           path: /foo
           pathType: Prefix
16
17
         - backend:
18
            service:
19
               name: service2
20
               port:
21
                 number: 80
22
           path: /
23
           pathType: Prefix
```

All incoming traffic that does not match the ingress rules with host name is matched here for the paths for routing.

## **Default back end**

Default back end is a service that handles all traffic that is not matched against any of the Ingress rules.

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
4
   name: default-ingress
  namespace: default
5
6 spec:
   defaultBackend:
7
8
     service:
9
       name: testsvc
10
         port:
11
           number: 80
```

# Note:

A global default back end can be specified if NetScaler CPX is load balancing the traffic. You can create a default back end per frontend-ip:port combination in case of NetScaler VPX or MPX is the ingress device.

# Ingress class support

July 2, 2025

# What is Ingress class?

In a Kubernetes cluster, there might be multiple ingress controllers and you need to have a way to associate a particular ingress resource with an ingress controller.

You can specify the ingress controller that should handle the ingress resource by using ingressClassName in your ingress resource definition.

# **NetScaler Ingress Controller and Ingress classes**

The NetScaler Ingress Controller supports accepting multiple ingress resources, which have spec. ingressClassName. Each ingress resource can be associated with only one ingress.class. However, the Ingress Controller might need to handle various ingress resources from different classes.

You can associate the Ingress Controller with multiple ingress classes using the --ingressclasses argument under the spec section of the YAML file.

If ingress-classes is not specified for the Ingress Controller, then it accepts all ingress resources irrespective of the presence of ingressClassName in the ingress object.

If ingress-classes is specified, then the Ingress Controller accepts only those ingress resources that match ingressClassName. The Ingress controller does not process an Ingress resource without ingress.class in such a case.

**Note**: Ingress class names are case-insensitive.

## Sample YAML configurations with Ingress classes

Following is the snippet from a sample YAML file to associate ingress-classes with the Ingress Controller. This configuration works in both cases where the Ingress Controller runs as a standalone
pod or runs as a sidecar with NetScaler CPX. In the given YAML snippet, the following ingress classes are associated with the Ingress Controller.

- my-custom-class
- Citrix

```
1 spec:
       serviceAccountName: cic-k8s-role
2
3
       containers:
4
       - name: cic-k8s-ingress-controller
        image:"quay.io/citrix/citrix-k8s-ingress-controller:latest"
5
6
         # specify the ingress classes names to be supportedbyIngress
            Controller in args section.
         # First line should be --ingress-classes, andeverysubsequent line
7
             should be
         # the name of allowed ingress class. In the givenexampletwo
8
            classes named
9
         # "citrix" and "my-custom-class" are accepted. Thiswill be case-
            insensitive.
         args:
11

    --ingress-classes

12
             Citrix
13
             my-custom-class
```

Following is the snippet from an Ingress YAML file where the Ingress class association is depicted. In the given example, an Ingress resource named web-ingress is associated with the ingress class my -custom-class. If the NetScaler Ingress Controller is configured to accept my-custom-class, it processes this Ingress resource.

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
4 annotations:
5 name: web-ingress
6 spec:
7 ingressClassName: my-custom-class
```

#### **Ingress V1 and IngressClass support**

With the Kubernetes version 1.19, the Ingress resource is generally available.

As a part of this change, a new resource named as IngressClass is added to the ingress API. Using this resource, you can associate specific Ingress controllers to Ingresses. For more information on the IngressClass resource, see the Kubernetes documentation.

The following is a sample IngressClass resource.

```
1 apiVersion: networking.k8s.io/v1
```

```
2 kind: IngressClass
```

```
3 metadata:
4 name: citrix
5 spec:
6 controller: citrix.com/ingress-controller
```

An IngressClassresource must refer to the ingress class associated with the controller that should implement the Ingress rules as shown as follows:

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
4 name: minimal-ingress
5 spec:
6
   ingressClassName: citrix
7
    rules:
8
    - host: abc.com
9
       http:
10
         paths:
         - path: /
11
12
           pathType: Prefix
13
           backend:
14
             service:
15
               name: test
16
               port:
17
                 number: 80
```

The NetScaler Ingress Controller uses the following rules to match the Ingresses.

- If the NetScaler Ingress Controller is started without specifying the --ingress-classes argument:
  - If the Kubernetes version is lesser than 1.19 (IngressClass V1 resource is supported)
    - \* Matches any ingress object
  - If the Kubernetes version is greater than or equal to 1.19 (IngressClass V1 resource is supported)
    - \* Matches any ingress object in which the spec.ingressClassName field is not set.
    - \* Matches any ingress if the spec.ingressClassName field of the Ingress object is set and a v1.IngressClass resource exists with the same name and the spec. controller field of the resource is citrix.com/ingress-controller.
- If the NetScaler Ingress Controller is started with one or more ingress classes set using the -ingress-classes argument.
  - If the Kubernetes version is lesser than 1.19 (IngressClass V1 resource is supported)
    - \* Matches any ingress with the ingress class annotation kubernetes.io/ingress .class matching to that of the configured ingress classes.

- If the Kubernetes version is greater than or equal to 1.19 (IngressClass V1 resource is supported).
  - Matches any ingress in which the ingress class annotation kubernetes.io/ ingress.class matches with the configured ingress classes. This annotation is deprecated but it has higher precedence over the spec.IngressClassName field to support backward compatibility.
  - \* Matches any ingress object, if a v1.IngressClass resource exists with the following attributes:
    - The name of the resource matches the --ingress-classes argument value.
    - The spec.controller field of the resource is set as the citrix.com/ ingress-controller.
    - The name of the resource matches with the spec.ingressClassName field of the Ingress object.
  - \* Matches any ingress object where the spec.ingressClassName field is not set and if a v1.IngressClass resource exists with the following attributes:
    - The name of the resources matches the --ingress-classes argument value.
    - The spec.controller field of the resource is set as citrix.com/ingress -controller.
    - The resource is configured as the default class using the ingressclass. kubernetes.io/is-**default-class** annotation. For more information, see the Kubernetes documentation.

#### Note:

- If both the annotation and spec.ingressClassName is defined, the annotation is matched before the spec.ingressClassName. If the annotation does not match, the matching operation for the spec.ingressClassName field is not performed.
- When you are using Helm charts to install the NetScaler Ingress Controller, if the IngressClass resource is supported and the NetScaler Ingress Controller is deployed with the --ingress-classes argument, the v1.IngressClass resource is created by default.

#### Updating the Ingress status for the Ingress resources with the specified IP address

To update the Status.LoadBalancer.Ingress field of the Ingress resources managed by NetScaler Ingress Controller with the allocated IP addresses, specify the command line argument --update-ingress-status yes when you start NetScaler Ingress Controller.

## Note:

For Helm deployments, the Helm chart equivalent parameter for update-ingress-status is updateIngressStatus, which needs to be set to **true**.

This feature is supported on both standalone and sidecar NSIC deployments. For an NSIC deployed as a sidecar, this feature is supported from NSIC version 2.2.10.

## Ingress status update for sidecar deployments

In Kubernetes, Ingress can be used as a single entry point for exposing multiple applications to the outside world. The Ingress would have an Address (Status.LoadBalancer.IP) field which is updated after the successful ingress creation. This field is updated with a public IP address or host name through which the Kubernetes application can be reached. In cloud deployments, this field can also be the IP address or host name of a cloud load-balancer.

In cloud deployments, NetScaler CPX along with the ingress controller is exposed using a service of type LoadBalancer which in turn creates a cloud load-balancer. The cloud load balancer then exposes the NetScaler CPX along with the ingress controller. So, the Ingress resources exposed with the NetScaler CPX should be updated using the public IP address or host name of the cloud load balancer.

This update is applicable on on-prem deployments too.

- In dual-tier ingress deployments, in which NetScaler CPX is exposed using a service of type LoadBalancer to the tier-1 NetScaler VPX ingress, the Address (Status.LoadBalancer
   .IP) field of the ingress resources operated by NetScaler CPX is updated with the VIP address.
- In dual-tier ingress deployments, in which NetScaler CPX is exposed using a service of type ClusterIP or NodePort, the Address (Status.LoadBalancer.IP) field of the ingress resources operated by NetScaler CPX is updated with the cluster IP address of the CPX service.

#### Sample ingress output after an ingress status update

The following is a sample ingress output after the ingress status update:

1	\$ kubectl get ingress				
2					
3	NAME	HOSTS	ADDRESS		
		PORTS	AGE		
4	sample-ingress	sample.citrix.com	<pre>sample.abc.somexampledomain.</pre>		
	com 80	1d			

# Service class for services of type LoadBalancer

#### December 31, 2023

When services of type LoadBalancer are deployed, all such services are processed by the NetScaler Ingress Controller and configured on NetScalers. However, there may be situations where you want to associate only specific services to a NetScaler Ingress Controller if multiple Ingress controllers are deployed.

For Ingress resources this functionality is already available using the Ingress class feature. Similar to the Ingress class functionality for Ingress resources, service class functionality is now added for services of type LoadBalancer.

You can associate a NetScaler Ingress Controller with multiple service classes using the --service -classes argument under the spec section of the YAML file. If a service class is not specified for the ingress controller, then it accepts all services of the type LoadBalancer irrespective of the presence of the service.citrix.com/class annotation in the service.

If the service class is specified to the NetScaler Ingress Controller, then it accepts only those services of the type LoadBalancer that match the service.citrix.com/class annotation. In this case, the NetScaler Ingress Controller does not process a type LoadBalancer service if it is not associated with the service.citrix.com/class annotation.

## Sample YAML configurations with service classes

Following is a snippet from a sample YAML file to associate service-classes with the Ingress Controller. In this snippet, the following service classes are associated with the Ingress Controller.

- svc-class1
- svc-class2

```
1 spec:
2 serviceAccountName: cic-k8s-role
3 containers:
4 - name: cic-k8s-ingress-controller
5
    # specify the service classes to be supported by NetScaler Ingress
        Controller in args section.
6
    # First line should be --service-classes, and every subsequent line
       should be
    # the name of allowed service class. In the given example two classes
7
        named
8
    # "svc-class1" and "svc-class2" are accepted. This will be case-
       insensitive.
9
    args:
      - --service-classes
```

11svc-class112svc-class2

Following is a snippet from a type LoadBalancer service definition YAML file where the service class association is depicted. In this example, an Apache service is associated with the service class svc-class1. If the NetScaler Ingress Controller is configured to accept svc-class1, it configures the service on the NetScaler.

```
1 apiVersion: v1
2 kind: Service
3 metadata:
4
        name: apache
5
        annotations:
6
            service.citrix.com/class: 'svc-class1'
7
        labels:
8
           name: apache
9 spec:
    type: LoadBalancer
11
      selector:
12
          name: apache
13
    ports:
14
      - name: http
15
       port: 80
16
        targetPort: http
17
      selector:
18
      app: apache
```

# Configure HTTP, TCP, or SSL profiles on NetScaler

#### June 24, 2025

Configurations such as, HTTP, TCP, or SSL for a NetScaler appliance can be specified using individual entities such as HTTP profile, TCP profile, or SSL profile respectively. The profile is a collection of settings pertaining to the individual protocols, for example, HTTP profile is a collection of HTTP settings. It offers ease of configuration and flexibility. Instead of configuring the settings on each entity you can configure them in a profile and bind the profile to all the entities that the settings apply to.

NetScaler Ingress Controller enables you to configure HTTP, TCP, or SSL related configuration on the Ingress NetScaler using profiles.

## **Understand NetScaler configuration in Kubernetes environment**

In a Kubernetes environment, the Ingress NetScaler uses Content Switching (CS) virtual server as the front end for external traffic. That is, it is the entity that receives the requests from the client. After

processing the request, the CS virtual server passes the request data to a load balancing (LB) entity. The LB virtual server and the associated service group processes the request data and then forwards it to the appropriate app (microservice).

You need to have a separate front end configuration for the entities that receive the traffic from the client (highlighted as Client Plane in the diagram) and a back end configuration for the entities that forward the traffic from the NetScaler to the microservices in Kubernetes (highlighted as Server Plane in the diagram).



The NetScaler Ingress Controller provides individual smart annotations for the front end and back-end configurations that you can use based on your requirement.

## **HTTP profile**

An HTTP profile is a collection of HTTP settings. A default HTTP profile (nshttp\_default\_profile ) is configured to set the HTTP configurations that are applied by default, globally to all services and virtual servers.

The NetScaler Ingress Controller provides the following two smart annotations for HTTP profile. You can use these annotations to define the HTTP settings for the NetScaler. When you deploy an ingress that includes these annotations, the NetScaler Ingress Controller creates an HTTP profile derived from the default HTTP profile (nshttp\_default\_profile) configured on the NetScaler. Then, it applies the parameters that you have provided in the annotations to the new HTTP profile and applies the profile to the NetScaler.

Smart annotation	Description	Sample
<pre>ingress.citrix.com/ frontend-httpprofile</pre>	Use this annotation to create the front-end HTTP profile (Client Plane)	<pre>ingress.citrix.com/ frontend-httpprofile:     '{ "dropinvalreqs":" enabled", "websocket"     : "enabled" } '</pre>
<pre>ingress.citrix.com/ backend-httpprofile</pre>	gress.citrix.com/ Use this annotation to create ckend-httpprofile the back-end HTTP profile (Server Plane).	
	Note: Ensure that you manually enable the HTTP related global parameters on the NetScaler. For example, to use HTTP2 at the back end (Server Plane), ensure that you can enable HTTP2Serverside global parameter in the NetScaler. For more information, see Configurating HTTP2.	

## **TCP profile**

A TCP profile is a collection of TCP settings. A default TCP profile (nstcp\_default\_profile) is configured to set the TCP configurations that is applied by default, globally to all services and virtual servers.

The NetScaler Ingress Controller provides the following two smart annotations for TCP profile. You can use these annotations to define the TCP settings for the NetScaler. When you deploy an ingress that includes these annotations, the NetScaler Ingress Controller creates a TCP profile derived from the default TCP profile (nstcp\_default\_profile) configured on the NetScaler. Then, it applies the parameters that you have provided in the annotations to the new TCP profile and applies the profile to the NetScaler.

Smart annotation	Description	Sample
<pre>ingress.citrix.com/ frontend-tcpprofile</pre>	Use this annotation to create the front-end TCP profile (Client Plane)	<pre>ingress.citrix.com/ frontend-tcpprofile: '{ "ws":"enabled", " sack" : "enabled" } '</pre>
ingress.citrix.com/ backend-tcpprofile	Use this annotation to create the back-end TCP profile (Server Plane)	<pre>ingress.citrix.com/ backend-tcpprofile: ' { "citrix-svc":{ "ws ":"enabled", "sack" :   "enabled" } } '</pre>

## SSL profile

An SSL profile is a collection of settings for SSL entities. It offers ease of configuration and flexibility. Instead of configuring the settings on each entity, you can configure them in a profile and bind the profile to all the entities that the settings apply to.

## Prerequisites

On the NetScaler, by default, SSL profile is not enable on the Ingress NetScaler. Ensure that you manually enable SSL profile on the NetScaler. Enabling the SSL profile overrides all the existing SSL related setting on the NetScaler, for detailed information on SSL profiles, see <u>SSL profiles</u>.

SSL profiles are classified into two categories:

- Front end profiles: containing parameters applicable to the front-end entity. That is, they apply to the entity that receives requests from a client.
- Back-end profiles: containing parameters applicable to the back-end entity. That is, they apply to the entity that sends client requests to a server.

Once you enable SSL profiles on the NetScaler, a default front end profile (ns\_default\_ssl\_profile\_front ) is applied to the SSL virtual server and a default back-end profile (ns\_default\_ssl\_profile\_backend ) is applied to the service or service group on the NetScaler.

The NetScaler Ingress Controller provides the following two smart annotations for SSL profile. You can use these annotations to customize the default front end profile (ns\_default\_ssl\_profile\_frontend) and back-end profile (ns\_default\_ssl\_profile\_backend) based on your requirement:

Smart annotation	Description	Sample	
ingress.citrix.com/	Use this annotation to create	ingress.citrix.com/	
frontend-sslprofile	the front end SSL profile (Client	frontend-sslprofile:	
	Plane). The front end SSL	'{ "hsts":"enabled",	
	profile is required only if you	"tls12" : "enabled" }	
	have enabled TLS on the Client	1. C. S.	
	Plane.		
ingress.citrix.com/	Use this annotation to create	ingress.citrix.com/	
backend-sslprofile	the back-end SSL profile	backend-sslprofile: '	
	(Server Plane). The SSL back	{ "citrix-svc":{ "	
	end profile is required only if	hsts":"enabled", "	
	you use the	tls1" : "enabled" }	
	ingress.citrix.com/secure-	1. B. S.	
	backend annotation for the		
	back-end.		

Important: SSL profile does not enable you to configure SSL certificate.

## Front-end profile configuration using annotations

HTTP, TCP, and SSL front-end profiles are attached to the client-side content switching virtual server or SSL virtual server. Since there can be multiple ingresses that use the same frontend-ip and also use the same content switching virtual server in the front-end, there can be possible conflicts that can arise from the front-end profiles annotation specified in multiple ingresses that share the front-end IP address.

The following are the guidelines for front-end profiles annotations for HTTP, TCP, and SSL.

- For all ingresses with the same front-end IP address, it is recommended to have the same value for the front-end profile is specified in all ingresses.
- If there are multiple ingresses that share front-end IP address, one can also create a separate ingress for each front-end IP address with empty rules (referred as the front-end ingress) where one can specify the front-end IP annotation as shown in the following example. You do not need to specify the front-end profile annotation in each ingress definition.
  - To create a front-end ingress for an HTTP type virtual server, see the following example:
    - 1 #Sample ingress manifest **for** the front-end configuration **for** an HTTP virtual server
      - 2 #The values for the parameters are for demonstration purpose only.

```
4
5
    apiVersion: networking.k8s.io/v1
   kind: Ingress
6
7
   metadata:
8
     name: frontend-ingress
9
      annotations:
      # /* The CS virtual server is derived from the combination
         of
   insecure-port/secure-port, frontend-ip, and
11
12
    secure-service-type/insecure-service-type annotations. */
        ingress.citrix.com/insecure-port: "80"
13
14
        ingress.citrix.com/frontend-ip: "x.x.x.x"
        ingress.citrix.com/frontend-httpprofile:'{
15
    "dropinvalreqs":"enabled", "markconnreqInval" : "enabled" }
16
17
    1
18
        ingress.citrix.com/frontend-tcpprofile: '{
    "ws":"enabled", "sack" :
19
    "enabled" }
20
    1
21
22
   spec:
23
      rules:
24
       # Empty rule
25
      - host:
```

- To create a front-end ingress for SSL type service, see the following example:

```
#Sample ingress manifest for the front-end configuration for
1
       an SSL virtual server
    #The values for the parameters are for demonstration purpose
       only.
3
4
    apiVersion: networking.k8s.io/v1
5
6
   kind: Ingress
7
   metadata:
     name: frontend-ingress
8
9
      annotations:
     #The CS virtual server is derived from the combination of
10
11
      #insecure-port/secure-port, frontend-ip, and
      #secure-service-type/insecure-service-type annotations.
12
        ingress.citrix.com/insecure-port: "80"
13
14
        ingress.citrix.com/secure-port: "443"
15
        ingress.citrix.com/frontend-ip: "x.x.x.x"
16
        ingress.citrix.com/frontend-sslprofile:
    '{
17
    "tls13":"enabled", "hsts" : "enabled" }
18
    1
19
        ingress.citrix.com/frontend-tcpprofile: '{
21
    "ws":"enabled", "sack" :
    "enabled" }
    1
23
24
    spec:
```

```
25 rules:
26 - host:
27 #Presense of tls is considered as a secure service
28 tls:
29 - hosts:
```

- If there are different values for the same front-end profile annotations in multiple ingresses, the following order is used to bind the profiles to the virtual server.
  - If any ingress definition has a front-end annotation with pre-configured profiles, that is bound to the virtual server.
  - Merge all the (key, values) from different ingresses of the same front-end IP address and use the resultant (key, value) for the front-end profiles smart annotation.
  - If there is a conflict for the same key due to different values from different ingresses, a value is randomly chosen and other values are ignored. You must avoid having conflicting values.
- If there is no front-end profiles annotation specified in any of the ingresses which share the frontend IP address, then the global values from the ConfigMap that is FRONTEND\_HTTP\_PROFILE
   , FRONTEND\_TCP\_PROFILE, or FRONTEND\_SSL\_PROFILE is used for the HTTP, TCP, and SSL front-end profiles respectively.

## Global front-end profile configuration using ConfigMap variables

The ConfigMap variable is used for the front-end profile if it is not overridden by front-end profiles smart annotation in one or more ingresses that shares a front-end IP address. If you need to enable or disable a feature using any front-end profile for all ingresses, you can use the variables FRONTEND\_HTTP\_PROFILE, FRONTEND\_TCP\_PROFILE, or FRONTEND\_SSL\_PROFILE for HTTP, TCP, and SSL profiles respectively. For example, if you want to enable TLS 1.3 for all SSL ingresses, you can use FRONTEND\_SSL\_PROFILE to set this value instead of using the smart annotation in each ingress definition. See the ConfigMap documentation to know how to use ConfigMap with NetScaler Ingress Controller.

#### Configuration using FRONTEND\_HTTP\_PROFILE

The FRONTEND\_HTTP\_PROFILE variable is used for setting the HTTP options for the front-end virtual server (client plane), unless overridden by the ingress.citrix.com/frontend-httpprofile smart annotation in the ingress definition.

To use an existing profile on NetScaler or use a built-in HTTP profile.

```
1 apiVersion: v1
```

```
2 kind: ConfigMap
```

```
3 metadata:
4 name: cic-configmap
5 labels:
6 app: citrix-ingress-controller
7 data:
8 FRONTEND_HTTP_PROFILE: |
9 preconfigured: my_http_profile
```

In this example, my\_http\_profile is a pre-existing HTTP profile in NetScaler.

Alternatively, you can set the profile parameters as specified as follows. See the HTTP profile NITRO documentation for all possible key-values.

```
1 apiVersion: v1
2 kind: ConfigMap
3 metadata:
   name: cic-configmap
4
   labels:
5
   app: citrix-ingress-controller
6
7 data:
8 FRONTEND_HTTP_PROFILE:
9
    config:
        dropinvalreqs: 'ENABLED'
10
        websocket: 'ENABLED'
11
```

## Configuration using FRONTEND\_TCP\_PROFILE

The FRONTEND\_TCP\_PROFILE variable is used for setting the TCP options for the front-end virtual server (client side), unless overridden by the ingress.citrix.com/frontend-tcpprofile smart annotation in the ingress definition.

To use an existing profile on NetScaler or use a built-in TCP profile:

```
1 apiVersion: v1
2 kind: ConfigMap
3 metadata:
4 name: cic-configmap
5 labels:
6 app: citrix-ingress-controller
7 data:
8 FRONTEND_TCP_PROFILE: |
9 preconfigured: my_tcp_profile
```

In this example, my\_tcp\_profile is a pre-existing TCP profile in NetScaler.

Alternatively, you can set the profile parameters as follows. See the NetScaler TCP profile NITRO documentation for all possible key values.

```
1 apiVersion: v1
2 kind: ConfigMap
```

```
3 metadata:
4 name: cic-configmap
5 labels:
6 app: citrix-ingress-controller
7 data:
8 FRONTEND_TCP_PROFILE: |
9 config:
10 sack: 'ENABLED'
11 nagle: 'ENABLED'
```

#### Configuration using FRONTEND\_SSL\_PROFILE

The FRONTEND\_SSL\_PROFILE variable is used for setting the SSL options for the front-end virtual server (client side) unless overridden by the ingress.citrix.com/frontend-sslprofile smart annotation in the ingress definition.

#### Note:

For the SSL profile to work correctly, you must enable the default profile in NetScaler using the set ssl parameter -defaultProfile ENABLED command. Make sure that NetScaler Ingress Controller is restarted after enabling the default profile. The default profile is automatically enabled when NetScaler CPX is used as an ingress device. For more information about the SSL default profile, see the SSL profile documentation.

To use an existing profile on NetScaler or use a built-in SSL profile,

```
1 apiVersion: v1
2 kind: ConfigMap
3 metadata:
4 name: cic-configmap
5 labels:
6 app: citrix-ingress-controller
7 data:
8 FRONTEND_SSL_PROFILE: |
9 preconfigured: my_ssl_profile
```

In this example, my\_ssl\_profile is the pre-existing SSL profile in NetScaler.

#### Note:

Default front end profile (ns\_default\_ssl\_profile\_frontend) is not supported using the FRONTEND\_SSL\_PROFILE.preconfigured variable.

Alternatively, you can set the profile parameters as shown in the following example. See the SSL profile NITRO documentation for information on all possible key-values.

```
1 apiVersion: v1
2 kind: ConfigMap
3 metadata:
```

```
4 name: cic-configmap
5 labels:
6 app: citrix-ingress-controller
7 data:
8 FRONTEND_SSL_PROFILE: |
9 config:
10 tls13: 'ENABLED'
11 hsts: 'ENABLED'
```

The following example shows binding SSL cipher groups to the SSL profile. The order is as specified in the list with the higher priority is provided to the first in the list and so on. You can use any SSL ciphers available in NetScaler or user-created cipher groups in this field. For information about the list of cyphers available in the NetScaler, see Ciphers in NetScaler.

```
1 apiVersion: v1
2 kind: ConfigMap
3 metadata:
   name: cic-configmap
4
   labels:
5
    app: citrix-ingress-controller
6
7 data:
8 FRONTEND_SSL_PROFILE:
9
      config:
        tls13: 'ENABLED'
10
11
        ciphers:
12
         - TLS1.3-AES256-GCM-SHA384
13
        - TLS1.3-CHACHA20-POLY1305-SHA256
```

## **Back-end configuration**

Any ingress definition that includes service details, spec:rules:host, spec:backend entry, and so on are considered as back-end configuration.

#### Sample backend ingress manifest without TLS configuration

```
1 #The values for the parameters are for demonstration purpose only.
2
   apiVersion: networking.k8s.io/v1
3 kind: Ingress
4 metadata:
5
   annotations:
6
     # /* The CS virtual server is derived from the combination of
        insecure-port/secure-port, frontend-ip, and secure-service-type/
        insecure-service-type annotations. */
7
       ingress.citrix.com/backend-httpprofile: '{
   "apache":{
8
    "markhttp09inval": "disabled" }
9
10
    }
11
```

```
ingress.citrix.com/backend-tcpprofile: '{
12
    "apache":{
13
    "sack":"enabled" }
14
15
     }
    1
16
17
       ingress.citrix.com/frontend-ip: 'VIP_IP'
       ingress.citrix.com/insecure-port: "80"
18
19
     name: apache-ingress
20
   spec:
21
    rules:
     - host: www.apachetest.com
23
       http:
24
          paths:
25
          - backend:
26
              service:
27
                name: apache
28
                port:
29
                  number: 80
            path: /
31
            pathType: Prefix
```

#### Sample backend ingress manifest with TLS configuration

```
1 #The values for the parameters are for demonstration purpose only.
2
3 apiVersion: networking.k8s.io/v1
4 kind: Ingress
5 metadata:
6
     annotations:
7
     # /* The CS virtual server is derived from the combination of
        insecure-port/secure-port, frontend-ip, and secure-service-type/
        insecure-service-type annotations. */
       ingress.citrix.com/backend-httpprofile: '{
8
9
    "hotdrink":{
    "markhttp09inval": "disabled" }
10
11
    }
    1
12
13
       ingress.citrix.com/backend-sslprofile: '{
    "hotdrink":{
14
15
    "snienable": "enabled" }
16
    }
    r.
17
       ingress.citrix.com/backend-tcpprofile: '{
18
    "hotdrink":{
19
    "sack":"enabled" }
21
    }
22
23
       ingress.citrix.com/frontend-ip: 'VIP_IP'
       ingress.citrix.com/secure-backend: '{
24
25
    "hotdrink":"true" }
26
```

```
ingress.citrix.com/secure-port: "443"
   name: hotdrink-ingress
28
29 spec:
    rules:
31
     - host: hotdrinks.beverages.com
32
       http:
         paths:
34
         - backend:
35
             service:
               name: hotdrink
37
               port:
38
                 number: 443
           path: /
39
40
           pathType: Prefix
41
     tls:
42
     - secretName: hotdrink.secret
```

## Using built-in or existing user-defined profiles on the Ingress NetScaler

You can use the individual smart annotations to configure the built-in profiles or existing user-defined profiles on the Ingress NetScaler for the front end and back-end configurations based on your requirement. For more information on built-in profiles, see Built-in TCP Profiles and Built-in HTTP profiles.

For the front end configuration, you can provide the name of the built-in or existing user-defined profiles on the Ingress NetScaler. The following is a sample ingress annotation:

1 ingress.citrix.com/frontend-httpprofile: "http\_preconf\_profile1"

Where, 'http\_preconf\_profile1' is the profile that exists on the Ingress NetScaler.

For the back-end configuration, you must provide the name of the built-in or existing profile on the Ingress NetScaler and the back-end service name. The following is a sample ingress annotation:

```
1 ingress.citrix.com/backend-httpprofile: '{
2 "citrix-svc": "http_preconf_profile1" }
3 '
```

Where, 'http\_preconf\_profile1' is the profile that exists on the Ingress NetScaler and citrix-svc is the back-end service name.

#### Sample HTTP profile

```
ingress.citrix.com/frontend-httpprofile: "http_preconf_profile"
ingress.citrix.com/backend-httpprofile: '{
    "citrix-svc": "http_preconf_profile" }
    '
```

#### Sample TCP profile

```
ingress.citrix.com/frontend-tcpprofile: "tcp_preconf_profile"
ingress.citrix.com/backend-tcpprofile: '{
    "citrix-svc":"tcp_preconf_profile" }
    '
```

#### Sample SSL profile

```
ingress.citrix.com/frontend-sslprofile: "ssl_preconf_profile"
ingress.citrix.com/backend-sslprofile: '{
    "citrix-svc":"ssl_preconf_profile" }
    '
```

## Example for applying HTTP, SSL, and TCP profiles

This example shows how to apply HTTP, SSL, or TCP profiles.

To create SSL, TCP, and HTTP profiles and bind them to the defined Ingress resource, perform the following steps:

1. Define the front-end ingress resource with the required profiles. In this Ingress resource, backend and TLS is not defined.

A sample YAML (ingress1.yaml) is provided as follows:

```
1
     apiVersion: networking.k8s.io/v1
     kind: Ingress
2
3
     metadata:
4
     name: ingress-vpx1
5
     annotations:
         ingress.citrix.com/insecure-termination: "allow"
6
7
         ingress.citrix.com/frontend-ip: "10.221.36.190"
8
         ingress.citrix.com/frontend-tcpprofile: '{
    "ws":"disabled", "sack" : "disabled" }
9
10
    1
11
         ingress.citrix.com/frontend-httpprofile: '{
    "dropinvalreqs":"enabled", "markconnreqInval" : "enabled" }
12
13
    1.1
14
         ingress.citrix.com/frontend-sslprofile: '{
    "hsts":"enabled", "tls13" : "enabled" }
15
    1
16
17
    spec:
      ingressClassName: "vpx"
18
19
      tls:
20
      - hosts:
21
      rules:
22
       - host:
```

2. Deploy the front-end ingress resource.

kubectl create -f ingress1.yaml

3. Define the secondary ingress resource with the same front-end IP address and TLS and the backend defined which creates the load balancing resource definition.

A sample YAML (ingress2.yaml) is provided as follows:

```
1
     apiVersion: networking.k8s.io/v1
2
     kind: Ingress
3
     metadata:
4
     name: ingress-vpx2
5
      annotations:
       ingress.citrix.com/insecure-termination: "allow"
6
7
       ingress.citrix.com/frontend-ip: "10.221.36.190"
8
    spec:
9
       ingressClassName: "vpx"
10
       tls:
       - secretName: <hotdrink-secret>
11
12
      rules:
13
       - host: hotdrink.beverages.com
14
         http:
15
           paths:
16
           - path:
17
             backend:
               serviceName: frontend-hotdrinks
18
19
               servicePort: 80
```

4. Deploy the back-end ingress resource.

1

kubectl create -f ingress2.yaml

5. Once the YAMLs are applied the corresponding entities, profiles, and ingress resources are created and they were bound to the ingress resource.

```
1 # show cs vserver <k8s150-10.221.36.190_443_ssl>
2
3
     k8s150-10.221.36.190_443_ssl (10.221.36.190:443) - SSL Type:
        CONTENT
4
     State: UP
5
     Last state change was at Thu Apr 22 20:14:44 2021
6
     Time since last state change: 0 days, 00:10:56.850
7
     Client Idle Timeout: 180 sec
8
     Down state flush: ENABLED
9
     Disable Primary Vserver On Down : DISABLED
10
     Comment: uid=
        QEYQI2LDW5WR4A6P3NSZ37XICK0JKV4HPEM2H4PSK4HWA3JQWCLQ====
11
     TCP profile name: k8s150-10.221.36.190_443_ssl
12
     HTTP profile name: k8s150-10.221.36.190_443_ssl
13
     Appflow logging: ENABLED
     State Update: DISABLED
14
     Default: Content Precedence: RULE
15
```

```
16 Vserver IP and Port insertion: OFF
17
    L2Conn: OFF Case Sensitivity: ON
18 Authentication: OFF
    401 Based Authentication: OFF
19
20
   Push: DISABLED Push VServer:
21
    Push Label Rule: none
22
   Persistence: NONE
23
    Listen Policy: NONE
24
    IcmpResponse: PASSIVE
25
    RHIstate: PASSIVE
26
   Traffic Domain: 0
27
28
     1) Content-Switching Policy: k8s150-ingress-vpx1_tier-2-
        adc_443_k8s150-frontend-hotdrinks_tier-2-adc_80_svc
        Priority: 200000004 Hits: 0
29
     Done
```

#### Example: Adding SNI certificate to an SSL virtual server

This example shows how to add a single SNI certificate.

Note:

For the SSL profile to work correctly, you must enable the default profile in NetScaler using the set ssl parameter -defaultProfile ENABLED command. Make sure that NetScaler Ingress Controller is restarted after enabling default profile. For more information about the SSL default profile, see documentation.

1. Define the front-end ingress resource with the required profiles. In this Ingress resource, backend and TLS is not defined.

A sample YAML (ingress1.yaml) is provided as follows:

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
   name: ingress-vpx1
4
   annotations:
5
    ingress.citrix.com/insecure-termination: "allow"
6
7
     ingress.citrix.com/frontend-ip: "10.221.36.190"
     ingress.citrix.com/frontend-tcpprofile: '{
8
   "ws":"disabled", "sack" : "disabled" }
9
10
     ingress.citrix.com/frontend-httpprofile: '{
11
    "dropinvalreqs":"enabled", "markconnreqInval" : "enabled" }
12
13
14
     ingress.citrix.com/frontend-sslprofile: '{
   "snienable": "enabled", "hsts":"enabled", "tls13" : "enabled" }
15
16
17 spec:
```

```
18 ingressClassName: "vpx"
19 tls:
20 - hosts:
21 rules:
22 - host:
```

2. Deploy the front-end ingress resource.

```
1 kubectl create -f ingress1.yaml
```

 Define the secondary ingress resource with the same front-end IP address defining back-end as well as SNI certificates. If hosts are specified then the certkey specified as the secret name is added as the SNI certificate.

A sample YAML (ingress2.yaml) is provided as follows:

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
4 name: ingress-vpx2
5
   annotations:
     ingress.citrix.com/insecure-termination: "allow"
6
     ingress.citrix.com/frontend-ip: "10.221.36.190"
7
8 spec:
    ingressClassName: "vpx"
9
     tls:
11
    - hosts:
12

    hotdrink.beverages.com

13
       secretName: hotdrink-secret
14 rules:

    host: hotdrink.beverages.com

15
16
      http:
17
         paths:
18
         - path: /
19
           backend:
20
             serviceName: web
21
             servicePort: 80
```

4. Deploy the secondary ingress resource.

1 kubectl create -f ingress2.yaml

If multiple SNI certificates need to be bound to the front-end VIP, following is a sample YAML file.

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
4   name: ingress-vpx-frontend
5   ingress.citrix.com/insecure-termination: "allow"
6   ingress.citrix.com/frontend-ip: "10.221.36.190"
7 spec:
8   ingressClassName: "vpx"
9   tls:
```

```
10 - hosts:
   - hotdrink.beverse
secretName: hotdrink-secret
11

    hotdrink.beverages.com

12
13 - hosts:
14
         - frontend.agiledevelopers.com
    secretName: <frontend-secret>
15
   rules:
16

    host: hotdrink.beverages.com

17
18
     http:
19
       paths:
20
         - path: /
21
          backend:
22
             serviceName: web
23
             servicePort: 80
24
    - host: frontend.agiledevelopers.com
25
      http:
26
       paths:
27
         - path: /
28
          backend:
29
            serviceName: frontend-developers
             servicePort: 80
```

#### Example: Binding SSL cipher group

This example shows how to bind SSL cipher group.

Note:

For the SSL profile to work correctly, you must enable the default profile in NetScaler using the set ssl parameter -defaultProfile ENABLED command. Make sure that NetScaler Ingress Controller is restarted after enabling default profile.

Set default SSL profile on NetScaler using the command set ssl parameter -defaultProfile ENABLED before deploying NetScaler Ingress Controller. If you have already deployed NetScaler Ingress Controller, then redeploy it. For more information about the SSL default profile, see documentation.

For information on supported Ciphers on the NetScaler appliances, see Ciphers available on the NetScaler appliances.

For information about securing cipher, see securing cipher.

A sample YAML (cat frontend\_ingress.yaml) is provided as follows:

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
4 name: ingress-vpx
5 annotations:
```

```
ingress.citrix.com/insecure-termination: "allow"
6
7
      ingress.citrix.com/frontend-ip: "10.221.36.190"
      ingress.citrix.com/frontend-tcpprofile: '{
8
9 "ws":"disabled", "sack" : "disabled" }
10
11
      ingress.citrix.com/frontend-httpprofile: '{
    "dropinvalreqs":"enabled", "markconnreqInval" : "enabled" }
12
13
14
      ingress.citrix.com/frontend-sslprofile: '{
    "snienable": "enabled", "hsts":"enabled", "tls13" : "enabled", "
15
       ciphers" : [{
    "ciphername": "test", "cipherpriority" :"1" }
16
17
    ] }
18
19 spec:
20
    ingressClassName: "citrix"
21 tls:
    - hosts:
22
23 rules:
24
      - host:
```

## Log levels

December 31, 2023

The logs generated by NetScaler Ingress Controller are available as part of kubernetes logs. You can specify NetScaler Ingress Controller to log in the following log levels:

- CRITICAL
- ERROR
- WARNING
- INFO
- DEBUG

By default, NetScaler Ingress Controller is set to log in INFO log level. If you want to specify NetScaler Ingress Controller to log in a particular log level then you need to specify the log level in the NetScaler Ingress Controller deployment YAML file before deploying the NetScaler Ingress Controller. You can specify the log level in the spec section of the YAML file as follows:

```
1 apiVersion: v1
2 kind: Pod
3 metadata:
4 name: citrixingresscontroller
5 labels:
6 app: citrixingresscontroller
7 spec:
```

```
8
         serviceAccountName: cpx
9
         containers:
         - name: citrixingresscontroller
10
           image: "quay.io/citrix/citrix-k8s-ingress-controller:1.36.5"
11
12
           env:
13
           # Set kube api-server URL
           - name: "kubernetes_url"
14
             value: "https://10.x.x.x:6443"
15
           # Set NetScaler Management IP
16
17
           - name: "NS_IP"
18
             value: "10.x.x.x"
           # Set log level
19
           - name: "LOGLEVEL"
20
             value: "DEBUG"
21
           - name: "EULA"
22
23
             value: "yes"
24
           args:
25
           - -- feature-node-watch
26
             true
27
           imagePullPolicy: Always
```

#### Modify the log levels

To modify the log level configured on the NetScaler Ingress Controller instance, you need to delete the instance and update the log level value in the following section and redeploy the NetScaler Ingress Controller instance:

```
1 # Set log level
2 - name: "LOGLEVEL"
3 value: "XXXX"
```

Once you update the log level, save the YAML file and deploy it using the following command:

```
1 kubectl create -f citrix-k8s-ingress-controller.yaml
```

# SSL certificate for services of type LoadBalancer through the Kubernetes secret resource

October 8, 2024

This section provides information on how to use the SSL certificate stored as a Kubernetes secret with services of type LoadBalancer. The certificate is applied if the annotation service.citrix.com/service-type is SSL or SSL\_TCP.

## Using the NetScaler Ingress Controller default certificate

If the SSL certificate is not provided, you can use the default NetScaler Ingress Controller certificate to configure SSL and SSL SNI certificates. You can use default-ssl-certificate and default-ssl-snicertificate arguments to provide a secret to configure non-SNI and SNI certificates respectively.

```
1 `--default-ssl-certificate <NAMESPACE>/<SECRET_NAME>`
2 `--default-ssl-sni-certificate <NAMESPACE>/<SECRET_NAME>`
```

## Service annotations for SSL certificate as Kubernetes secrets

NetScaler Ingress Controller provides the following service annotations to use SSL certificates stored as Kubernetes secrets for services of type LoadBalancer.

Service annotation	Description
service.citrix.com/secret	Use this annotation to specify the name of the secret resource for the front-end server certificate. It must contain a certificate and key. You can also provide a list of intermediate CA certificates in the certificate section followed by the server certificate. These intermediate CAs are automatically linked and sent to the client during the SSL handshake. To bind multiple front-end server certificates, provide a list of comma-separated secrets configured for certificates. For example, service.citrix.com/secret: apache
	-secret1,apache-secret2.
service.citrix.com/ca-secret	Use this annotation to provide a CA certificate for client certificate authentication. This certificate is bound to the front-end SSL virtual server in NetScaler.
service.citrix.com/backend-secret	Use this annotation if the back-end communication between NetScaler and your workload is on an encrypted channel, and you need the client authentication in your workload. This certificate is sent to the server during the SSL handshake and it is bound to the back end SSL service group.

Service annotation	Description
service.citrix.com/backend-ca- secret	Use this annotation to enable server authentication which authenticates the back-end server certificate. This configuration binds the CA certificate of the server to the SSL
service.citrix.com/preconfigured- certkey	service on the NetScaler. Use this annotation to specify the name of an already existing cert key in the NetScaler to be used as a front-end server certificate. To bind multiple front-end server certificates, provide a list of comma-separated cert keys that are already configured for certificates. For example, service.citrix.com/preconfigured- certkey: preconfcert1, preconfcert2.
<pre>service.citrix.com/preconfigured- ca-certkey</pre>	Use this annotation to specify the name of the preconfigured cert key in the NetScaler to be used as a CA certificate for client certificate authentication. This certificate is bound to the
service.citrix.com/preconfigured- backend-certkey	front-end SSL virtual server in NetScaler. Use this annotation to specify the name of the preconfigured cert key in the NetScaler to be bound to the back-end SSL service group. This
service.citrix.com/preconfigured- backend-ca-certkey	certificate is sent to the server during the SSL handshake for server authentication. Use this annotation to specify the name of the preconfigured CA cert key in the NetScaler to bound to back-end SSL service group for server authentication.

#### **Examples: Front-end secret and Front-end CA secret**

Following are some examples for the service.citrix.com/secret annotation:

The following annotation is applicable to all ports in the service.

service.citrix.com/secret: hotdrink-secret

You can use the following notation to specify the certificate applicable to specific ports by giving either portname or port-protocol as key.

1

1

```
# port-protocol : secret
1
2
3
           service.citrix.com/secret: '{
4
    "443-tcp": "hotdrink-secret", "8443-tcp": "hotdrink-secret" }
5
6
7
            # portname: secret
8
9
           service.citrix.com/secret: '{
    "https": "hotdrink-secret" }
10
    1.1
11
```

Following are some examples for the service.citrix.com/ca-secret annotation.

You need to specify the following annotation to attach the generated CA secret which is used for client certificate authentication for a service deployed in Kubernetes.

The following annotation is applicable to all ports in the service.

service.citrix.com/ca-secret: hotdrink-ca-secret

You can use the following notation to specify the certificate applicable to specific ports by giving either portname or port-protocol as key.

```
1
    # port-protocol: secret
    service.citrix.com/ca-secret: '{
2
   "443-tcp": "hotdrink-ca-secret", "8443-tcp": "hotdrink-ca-secret" }
3
4
5
6
   # portname: secret
    service.citrix.com/ca-secret: '{
8
    "https": "hotdrink-ca-secret" }
9
10
    1
```

#### Examples: back-end secret and back-end CA secret

Following are some examples for the service.citrix.com/backend-secret annotation.

```
1 # port-protocol: secret
2
    service.citrix.com/backend-secret: '{
    "443-tcp": "hotdrink-secret", "8443-tcp": "hotdrink-secret" }
3
4
5
6
   # portname: secret
7
8
    service.citrix.com/backend-secret: '{
9
    "tea-443": "hotdrink-secret", "tea-8443": "hotdrink-secret" }
10
11
```

```
12 # applicable to all ports
13
14 service.citrix.com/backend-secret: "hotdrink-secret"
```

Following are some examples for the service.citrix.com/backend-ca-secret annotation.

```
1
     # port-proto: secret
2
     service.citrix.com/backend-ca-secret: '{
    "443-tcp": "coffee-ca", "8443-tcp": "tea-ca" }
3
4
5
6
     # portname: secret
     service.citrix.com/backend-ca-secret: '{
8
    "coffee-443": "coffee-ca", "tea-8443": "tea-ca" }
9
10
11
     # applicable to all ports
12
13
14
     service.citrix.com/backend-ca-secret: "hotdrink-ca-secret"
```

# BGP advertisement of external IP addresses for type LoadBalancer services and Ingresses using NetScaler CPX

December 31, 2023

Kubernetes service of type LoadBalancer support is provided by cloud load balancers in a cloud environment.

Cloud service providers enable this support by automatically creates a load balancer and assign an IP address which is displayed as part of the service status. Any traffic destined to the external IP address is load balanced on NodeIP and NodePort by the cloud load balancer. Once the traffic reaches the Kubernetes cluster, kube-proxy performs the routing to the actual application pods using iptables or IP virtual server rules. However, for on-prem environments the cloud load balancer auto configuration is not available.

You can expose the services of type LoadBalancer using the NetScaler Ingress Controller and Tier-1 NetScaler devices such as NetScaler VPX or MPX. The NetScaler VPX or MPX residing outside the Kubernetes cluster load balances the incoming traffic to the Kubernetes services. For more information on such a deployment, see Expose services of type LoadBalancer.

However, it may not be always feasible to use an external ADC device to expose the service of type LoadBalancer in an on-prem environment. Some times, it is desirable to manage all related resources from the Kubernetes cluster itself without any external component. The NetScaler Ingress Controller

provides a way to expose the service of type LoadBalancer using NetScaler CPX that runs within the Kubernetes cluster. The existing BGP fabric to route the traffic to the Kubernetes nodes is leveraged to implement this solution.

In this deployment, NetScaler CPX is deployed as a daemonset on the Kubernetes nodes in host mode. NetScaler CPX establishes a BGP peering session with your network routers, and uses that peering session to advertise the IP addresses of external cluster services. If your routers have ECMP capability, the traffic is load-balanced to multiple CPX instances by the upstream router, which in turn load-balances to actual application pods. When you deploy the NetScaler CPX with this mode, NetScaler CPX adds iptables rules for each service of type LoadBalancer on Kubernetes nodes. The traffic destined to the external IP address is routed to NetScaler CPX pods.

The following diagram explains a deployment where NetScaler CPX is exposing a service of type Load-Balancer:



As shown in the diagram, NetScaler CPX runs as a daemon set and runs a BGP session over port 179 on the node IP address pointed by the Kubernetes node resource. For every service of type LoadBalancer added to the Kubernetes API server, the NetScaler Ingress Controller configures the NetScaler CPX to advertise the external IP address to the BGP router configured. A /32 prefix is used to advertise the routes to the external router and the node IP address is used as a gateway to reach the external IP address. Once the traffic reaches to the Kubernetes node, the iptables rule steers the traffic to NetScaler CPX which in turn load balance to the actual service pods.

With this deployment, you can also use Kubernetes ingress resources and advertise the Ingress virtual IP (VIP) address to the router. You can specify the NS\_VIP environment variable while deploying the NetScaler Ingress Controller which acts as the VIP for all ingress resources. When an Ingress resource is added, NetScaler CPX advertises the NS\_VIP to external routers through BGP to attract the traffic. Once traffic comes to the NS\_VIP, NetScaler CPX performs the content switching and load balancing

#### as specified in the ingress resource.

Note:

For this solution to work, the NetScaler Ingress Controller must run as a root user and must have the NET\_ADMIN capability.

#### Deploy NetScaler CPX solution for services of type LoadBalancer

This procedure explains how to deploy NetScaler CPX as a daemonset in the host network to expose services of type LoadBalancer.

This configuration includes the following tasks:

- Deploy NetScaler CPX with the NetScaler Ingress Controller as sidecar
- BGP configuration
- Service configuration

#### Prerequisites

- You must configure the upstream router for BGP routing with ECMP support and add Kubernetes nodes as neighbors.
- If the router supports load balancing, it is better to use a stable ECMP hashing algorithm for load-balancing with a higher entropy for even load-balancing.

Perform the following:

1. Download the rbac.yaml file and deploy the RBAC rules for NetScaler CPX and the NetScaler Ingress Controller.

```
1 kubectl apply -f rbac.yaml
```

2. Download the citrix-k8s-cpx-ingress.yml using the following command.

```
1 wget https://raw.githubusercontent.com/citrix/citrix-k8s-ingress-
controller/master/docs/configure/cpx-bgp-router/citrix-k8s-cpx-
ingress.yml
```

- 3. Edit the citrix-k8s-cpx-ingress.yaml file and specify the required values.
  - The argument configmap specifies the ConfigMap location for the NetScaler Ingress Controller in the form of namespace/name.
  - The argument -- ipam citrix-ipam-controller can be specified if you are running the for automatic IP address allocation.

- (Optional) nodeSelector to select the nodes where you need to run the NetScaler CPX daemonset. By default, it is run on all worker nodes.
- 4. Apply the citrix-k8s-cpx-ingress.yaml file to create a daemonset which starts NetScaler CPX and the NetScaler Ingress Controller.

1 kubectl apply -f citrix-k8s-cpx-ingress.yml

5. Create a ConfigMap (configmap.yaml) with the BGP configuration which is passed as an argument to the NetScaler Ingress Controller. For detailed information on BGP configuration, see BGP configuration.

You must have the following information to configure BGP routing:

- The router IP address for NetScaler CPX to connect
- The autonomous system (AS number) of the router
- The AS number for NetScaler CPX

Following is a sample ConfigMap with the BGP configuration.

```
1 apiVersion: v1
2 kind: ConfigMap
3 metadata:
4
    name: config
5
    labels:
6
      app: cic
7 data:
    NS\_BGP\_CONFIG:
8
9
       bgpConfig:
10
       - bgpRouter:
           localAS: 100
11
12
           neighbor:
13
           - address: 10.102.33.33
14
             remoteAS: 100
15
             advertisementInterval: 10
16
             ASOriginationInterval: 10
```

6. Apply the ConfigMap created in step 5 to apply the BGP configuration.

kubectl apply -f configmap.yaml

7. Create a YAML file with the required configuration for service of type LoadBalancer.

Note:

For detailed information, see service configuration. The service configuration section explains different ways to get an external IP address for the service and also how to use the service annotation provided by NetScaler to configure different NetScaler functionalities.

Following is an example for configuration of service of type LoadBalancer.

```
1 apiVersion: v1
 2 kind: Service
3 metadata:
4
    name: kuard-service
5
    annotations:
      # This uses IPAM to allocate an IP from range 'Dev'
6
7
      # service.citrix.com/ipam-range: 'Dev'
      service.citrix.com/frontend-ip: 172.217.163.17
8
9
      service.citrix.com/service-type-0: 'HTTP'
      service.citrix.com/service-type-1: 'SSL'
     service.citrix.com/lbvserver: '{
11
   "80-tcp":{
12
    "lbmethod":"ROUNDROBIN" }
13
14
    }
15
    1
16
       service.citrix.com/servicegroup: '{
   "80-tcp":{
17
    "usip":"yes" }
18
19
     }
20
       service.citrix.com/ssl-termination: edge
21
       service.citrix.com/monitor: '{
22
23
    "80-tcp":{
    "type":"http" }
24
25
    }
27
       service.citrix.com/frontend-httpprofile: '{
    "dropinvalreqs":"enabled", "websocket" : "enabled" }
28
29
       service.citrix.com/backend-httpprofile: '{
31
    "dropinvalreqs":"enabled", "websocket" : "enabled" }
32
    1
33
      service.citrix.com/frontend-tcpprofile: '{
    "ws":"enabled", "sack" : "enabled" }
34
    1
       service.citrix.com/backend-tcpprofile: '{
    "ws":"enabled", "sack" : "enabled" }
37
    1
38
      service.citrix.com/frontend-sslprofile: '{
    "hsts":"enabled", "tls12" : "enabled" }
40
41
42
       service.citrix.com/backend-sslprofile: '{
    "tls12" : "enabled" }
43
44
45
       service.citrix.com/ssl-certificate-data-1: |
46
         ----BEGIN----
47
              [...]
         ----END----
48
49
      service.citrix.com/ssl-key-data-1: |
50 spec:
51
    type: LoadBalancer
52
     selector:
53 app: kuard
```

```
      54
      ports:

      55
      - port: 80

      56
      targetPort: 8080

      57
      name: http

      58
      - port: 443

      59
      targetPort: 8443

      60
      name: https
```

8. Apply the service of type LoadBalancer.

```
1 kubectl apply -f service-example.yaml
```

Once the service is applied, the NetScaler Ingress Controller creates a load balancing virtual server with BGP route health injection enabled. If the load balancing virtual server state is UP, the route for the external IP address is advertised to the neighbor router with a /32 prefix with the node IP address as the gateway.

#### **BGP configuration**

BGP configuration is performed using the ConfigMap which is passed as an argument to the NetScaler Ingress Controller.

You must have the following information to configure BGP routing:

- The router IP address so that NetScaler CPX can connect to it
- The autonomous system (AS number) of the router
- The AS number for NetScaler CPX

In the following ConfigMap for the BGP configuration, the bgpConfig field represents the BGP configuration.

```
1 apiVersion: v1
2 kind: ConfigMap
3 metadata:
   name: config
4
5
    labels:
6
     app: cic
7 data:
  NS_BGP_CONFIG:
8
9
     bgpConfig:
10
       - bgpRouter:
          localAS: 100
11
12
          neighbor:
13
           - address: x.x.x.x
14
             remoteAS: 100
15
             advertisementInterval: 10
16
             ASOriginationInterval: 10
```

The following table explains the various fields of the bgpConfig field.

Field	Description	Туре	Default value	Required			
nodeSelector	If the	string		No			
	nodeSeclector						
	field is present,	field is present,					
	then the BGP						
	router						
	configuration is						
	applicable for						
	nodes which						
	matches the						
	nodeSelector						
	field.						
	nodeSelector						
	accepts comma						
	separated						
	key=value						
	pairs where each						
	key represents a						
	label name and						
	the value is the						
	label value. For						
	example:						
	nodeSelector:						
	datacenter=ds1,rac	ck-					
	rack1						
bgpRouter	Specifies the BGP	bgpRouter		Yes			
	configuration.						
	For information						
	on different fields						
	of the						
	bgpRouter,see						
	the following						
	table.						

The following table explains the fields for the  ${\tt bgpRouter}$  field.

#### NetScaler ingress controller

Field	Description	Туре	Default value	Required	
localAS	AS number for the NetScaler CPX	integer		Yes	
neighbor	Neighbor router BGP configuration.	neighbor		Yes	

The following table explains the neighbor field.

Field	Description	Туре	Default value	Required	
address	IP address for the neighbor router.	string		Yes	
remoteAS	AS number of the neighbor router.	integer			Yes
advertisemen	nthis field sets a minimum interval between the sending of BGP routing updates (in seconds).	integer	10 seconds	Yes	
ASOriginatio	orhis field sets the interval of sending AS origination routing updates (in seconds).	integer	10 seconds	Yes	

# Different neighbors for different nodes

By default, every node in the cluster connects to all the neighbors listed in the configuration. But, if the Kubernetes cluster is spread across different data centers or different networks, different neighbor

configurations for different nodes may be required. You can use the nodeSelector field to select the nodes required for the BGP routing configurations.



An example ConfigMap with the nodeSelector configuration is given as follows:



In this example, the router with the IP address 10.102.33.44 is used as a neighbor by nodes with the label datacenter=ds1. The router with the IP address 10.102.28.12 is used by the nodes with the
label datacenter=ds2.

## Service configuration

#### **External IP address configuration**

An external IP address for the service of type LoadBalancer can be obtained by using one of the following methods.

• Specifying the service.citrix.com/frontend-ip annotation in the service specification as follows.

```
1 metadata:
2 annotations:
3 service.citrix.com/frontend-ip: 172.217.163.17
```

• Specifying an IP address in the spec.loadBalancerIP field of the service specification as follows.

spec:
loadBalancerIP: 172.217.163.17

• By automatically assigning a virtual IP address to the service using the IPAM controller provided by NetScaler. If one of the other two methods is specified, then that method takes precedence over the IPAM controller. The IPAM solution is designed in such a way that you can easily integrate the solution with ExternalDNS providers such as Infoblox. For more information, see Interoperability with ExternalDNS. For deploying and using the , see the documentation.

#### Service annotation configuration

The NetScaler Ingress Controller provides many service annotations to leverage the various functionalities of the NetScaler. For example, the default service type for the load balancing virtual server is TCP, but you can override this configuration by the service.citrix.com/service-type annotation.

```
1 metadata:
2 annotations:
3 service.citrix.com/service-type-0: 'HTTP'
4 service.citrix.com/service-type-1: 'SSL'
```

With the help of various annotations provided by the NetScaler Ingress Controller, you can leverage various ADC functionalities like SSL offloading, HTTP rewrite and responder policies, and other custom resource definitions (CRDs).

For more information on all annotations for service of type LoadBalancer, see service annotations.

For using secret resources for SSL certificates for Type LoadBalancer services, see SSL certificate for services of type LoadBalancer.

#### **External traffic policy configuration**

By default, the NetScaler Ingress Controller adds all the service pods as a back-end for the load balancing virtual service in NetScaler CPX. This step ensures better high availability and equal distribution to the service pod instances. All nodes running NetScaler CPX advertises the routes to the upstream server and attracts the traffic from the router. This behavior can be changed by setting the spec .externalTrafficPolicy of the service to Local. When the external traffic policy is set to Local, only the pods running in the same node is added as a back-end for the load balancing virtual server as shown in the following diagram. In this mode, only those nodes which have the service pods advertise the external IP address to the router and CPX sends the traffic only to the local pods. If you do not want the traffic hopping across the nodes for performance reasons, you can use this feature.



#### **Using Ingress resources**

The NetScaler Ingress Controller provides an nt variable NS\_VIP, which is the external IP Address for all ingress resources. Whenever an ingress resource is added, NetScaler CPX advertises the ingress IP address to the external routers.

The NetScaler Ingress Controller provides various annotations for ingress. For more information, see the Ingress annotation documentation.

Perform the following steps for the Ingress Configuration:

1. Download the rbac.yaml file and deploy the RBAC rules for NetScaler CPX and the NetScaler Ingress Controller.

kubectl apply -f rbac.yaml

2. Download the citrix-k8s-cpx-ingress.yml using the following command.

```
wget https://raw.githubusercontent.com/citrix/citrix-k8s-ingress-
controller/master/docs/configure/cpx-bgp-router/citrix-k8s-cpx-
ingress.yml
```

- 3. Edit the citrix-k8s-cpx-ingress.yml file and specify the required values.
  - The argument configmap specifies the ConfigMap location for the NetScaler Ingress Controller in the form of namespace or name.
  - The environment variable NS\_VIP to specify the external IP to be used for all Ingress resources. (This is a required parameter).
- 4. Apply the citrix-k8s-cpx-ingress.yml file to create a daemonset which starts NetScaler CPX and the NetScaler Ingress Controller.

1 kubectl apply -f citrix-k8s-cpx-ingress.yml

- 5. Configure BGP using ConfigMap as shown in the previous section.
- 6. Deploy a sample ingress resource as follows. This step advertises the IP address specified in the NS\_VIP environment variable to the external router configured in ConfigMap.

```
1 kubectl apply -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/docs/configure/cpx-bgp-router/
ingress-example.yaml
```

7. Access the application using NS\_VIP: <port>. By default, Ingress uses port 80 for insecure communication and port 443 for secure communication (If TLS section is provided).

**Note:** Currently, the ingress.citrix.com/frontend-ip annotation is not supported for BGP advertisements.

#### **Helm Installation**

You can use Helm charts to install the NetScaler CPX as BGP router. For more information, see the Citrix Helm chart documentation.

### Troubleshooting

- By default. NetScaler CPX uses the IP address range range 192.168.1.0/24 for internal communication, the IP address 192.168.1.1 as internal gateway to the host, and the IP address IP address 192.168.1.2 as NSIP. The ports 9080 and 9443 are used as management ports between the NetScaler Ingress Controller and NetScaler CPX for HTTP and HTTPS. If the 192.168.1.0/24 network falls within the range of PodCIDR, you can allocate a different set of IP addresses for internal communication. The NS\_IP and NS\_GATEWAY environment variables control which IP address is used by NetScaler CPX for NSIP and gateway respectively. The same IP address must also be specified as part of the NetScaler Ingress Controller environment variable NS\_IP to establish the communication between the NetScaler Ingress Controller and NetScaler CPX.
- By default, BGP on NetScaler CPX runs on port 179 and all the BGP traffic coming to the TCP port 179 is handled by NetScaler CPX. If there is a conflict, for example if you are using Calico's external BGP peering capability to advertise your cluster prefixes over BGP, you can change the BGP port with the environment variable to the NetScaler Ingress Controller BGP\_PORT.
- Use source IP (USIP) mode of NetScaler does not work due to the constraints in Kubernetes. If the source IP address is required by the service, you can enable the CIP (client IP header) feature on the HTTP/SSL service-type services by using the following annotations.

service.citrix.com/servicegroup: '{"cip":"ENABLED", "cipheader":"x-forwarded-for"}'

# NetScaler CPX integration with MetalLB in layer 2 mode for on-premises Kubernetes clusters

#### December 31, 2023

Kubernetes service of type LoadBalancer support is provided by cloud load balancers in a cloud environment. Cloud service providers enable this support by automatically creates a load balancer and assign an IP address which is displayed as part of the service status. Any traffic destined to the external IP address is load balanced on NodeIP and NodePort by the cloud load balancer.

NetScaler provides different options to support the type LoadBalancer services in an on-premises environment including:

• Using an external NetScaler VPX or NetScaler MPX as a tier-1 load balancer to load balance the incoming traffic to Kubernetes services.

For more information on such a deployment, see Expose services of type LoadBalancer.

- Expose applications running in a Kubernetes cluster using the NetScaler CPX daemonset running inside the Kubernetes cluster along with a router supporting ECMP over BGP. ECMP router load balances the traffic to multiple NetScaler CPX instances. NetScaler CPX instances load balances the actual application pods. For more information on such a deployment, see BGP advertisement of external IP addresses for type LoadBalancer services and Ingresses using NetScaler CPX.
- Expose the NetScaler CPX services as an external IP service with a node external IP address. You can use this option if an external ADC as tier-1 is not feasible, and a BGP router does not exist. In this deployment, Kubernetes routes the traffic coming to the spec.externalIP of the NetScaler CPX service on service ports to NetScaler CPX pods. Ingress resources can be configured using the NetScaler Ingress Controller to perform SSL (Secure Sockets Layer) offloading and load balancing applications. However, this deployment has the major drawback of not being reliable if there is a node failure.
- Use MetalLB which is a load-balancer implementation for bare metal Kubernetes clusters in the layer 2 mode with NetScaler CPX to achieve ingress capability.

This documentation shows how you can leverage MetalLB along with NetScaler CPX to achieve ingress capability in bare-metal clusters when the other solutions are not feasible. MetalLB in layer 2 mode configures one node to send all the traffic to the NetScaler CPX service. MetalB automatically moves the IP address to a different node if there is a node failure. Thus providing better reliability than the ExternalIP service.

**Note:** MetalLB is still in the beta version. See the official documentation to know about the project maturity and any limitations.

Perform the following steps to deploy NetScaler CPX integration with MetalLB in layer 2 mode for onpremises Kubernetes clusters.

- 1. Install and configure MetalLB
- 2. Configure MetalLB configuration for layer 2
- 3. Install NetScaler CPX service

## Install and configure MetalLB

First, you should install MetalLB in layer 2 mode. For more information on different types of installations for MetalLB, see the MetalLB documentation.

Perform the following steps to install MetalLB:

1. Create a namespace for deploying MetalLB.

```
1 kubectl apply -f https://raw.githubusercontent.com/metallb/metallb
/v0.9.5/manifests/namespace.yaml
```

2. Deploy MetalLB using the following command.

```
1 kubectl apply -f https://raw.githubusercontent.com/metallb/metallb
/v0.9.5/manifests/metallb.yaml
```

3. Perform the following step if you are performing the installation for the first time.

```
1 kubectl create secret generic -n metallb-system memberlist --from-
literal=secretkey="$(openssl rand -base64 128)"
```

4. Verify the MetalLB installation and ensure that the speaker and controller is in the running state using the following command:

1 kubectl get pods -n metallb-system

These steps deploy MetalLB to your cluster, under the metallb-system namespace.

The MetalLB deployment YAML file contains the following components:

- The metallb-system/controller deployment: This component is the cluster-wide controller that handles IP address assignments.
- The metallb-system/speaker daemonset. This component communicates using protocols of your choice to make the services reachable.
- Service accounts for the controller and speaker, along with the RBAC permissions that the components need to function.

#### **MetalLB configuration for Layer 2**

Once MetalLB is installed, you should configure the MetalLB for layer 2 mode. MetalLB takes a range of IP addresses to be allocated to the type LoadBalancer services as external IP. In this deployment, a NetScaler CPX service acts as a front-end for all other applications. Hence, a single IP address is sufficient.

Create a ConfigMap for MetalLB using the following command where metallb-config.yaml is the YAML file with the MetalLB configuration.

```
1 kubectl create - f metallb-config.yaml
```

Following is a sample MetalLB configuration for layer2 mode. In this example, 192.168.1.240-192.168.1.240 is specified as the IP address range.

```
    apiVersion: v1
    kind: ConfigMap
    metadata:
    namespace: metallb-system
    name: config
```

```
6 data:
7 config: |
8 address-pools:
9 - name: default
10 protocol: layer2
11 addresses:
12 - 192.168.1.240-192.168.1.240
```

### **NetScaler CPX service installation**

Once the metal LB is successfully installed, you can install the NetScaler CPX deployment and a service of type LoadBalancer.

To install NetScaler CPX, you can either use the YAML file or Helm charts.

To install NetScaler CPX using the YAML file, perform the following steps:

1. Download the NetScaler CPX deployment manifests.

```
1 wget https://github.com/citrix/citrix-k8s-ingress-controller/blob/
master/deployment/baremetal/citrix-k8s-cpx-ingress.yml
```

- 2. Edit the NetScaler CPX deployment YAML:
  - Set the replica count as needed. It is better to have more than one replica for high availability.
  - Change the service type to LoadBalancer.
- 3. Apply the edited YAML file using the Kubectl command.

```
1 kubectl apply - f citrix-k8s-cpx-ingress.yaml
```

4. View the service using the following command:

1 kubectl get svc cpx-service -output yaml

You can see that MetalLB allocates an external IP address to the NetScaler CPX service as follows:

```
1 apiVersion: v1
2 kind: Service
3 metadata:
    name: cpx-service
4
5
    namespace: default
6 spec:
7
   clusterIP: 10.107.136.241
8 externalTrafficPolicy: Cluster
9 healthCheckNodePort: 31916
     ports:
10
     - name: http
11
       nodePort: 31528
```

13	port: 80
14	protocol: TCP
15	targetPort: 80
16	- name: https
17	nodePort: 31137
18	port: 443
19	protocol: TCP
20	targetPort: 443
21	selector:
22	app: cpx-ingress
23	sessionAffinity: None
24	type: LoadBalancer
25	status:
26	loadBalancer:
27	ingress:
28	- ip: 192.168.1.240

#### **Deploy a sample application**

Perform the following steps to deploy a sample application and verify the deployment.

1. Create a sample deployment using the sample-deployment.yaml file.

```
1 kubectl create - f sample-deployment.yaml
```

2. Expose the application with a service using the sample-service.yaml file.

1 kubectl create - f sample-service.yaml

3. Once the service is created, you can add an ingress resource using the sample-ingress.yaml.

```
1 kubectl create - f sample-ingress.yaml
```

You can test the Ingress by accessing the application using a cpx-service external IP address as follows:

```
1 curl -v http://192.168.1.240 -H 'host: testdomain.com'
```

#### **Additional references**

For more information on configuration and troubleshooting for MetalLB see the following links:

- Metal LB troubleshooting
- Configuring routing for metal LB in layer 2 mode

# Advanced content routing for Kubernetes Ingress using the HTTPRoute CRD

#### June 26, 2025

Kubernetes native Ingress offers basic host and path-based routing which is supported by the NetScaler Ingress Controller.

NetScaler also provides an alternative approach using content routing CRDs for supporting advanced routing capabilities. Content Routing CRDs include Listener CRD and HTTPRoute CRD. These CRDs provide advanced content routing features such as regex based expression and content switching based on query parameters, cookies, HTTP headers, and other NetScaler custom expressions.

With the Ingress version networking.k8s.io/v1, Kubernetes introduces support for resource backends. A resource backend is an ObjectRef to another Kubernetes resource within the same namespace as an Ingress object.

Now, NetScaler supports configuring the HTTP route CRD resource as a resource backend in Ingress. By default, Ingress supports only limited content routing capabilities like path and host-based routing. With this feature, you can extend advanced content routing capabilities to Ingress and configure various content switching options. For a given domain, you can use the HTTPRoute custom resource to configure content switching without losing the third party compatibility support of the Kubernetes Ingress API.



Note:

- This feature supports the Kubernetes Ingress version networking.k8s.io/v1 that is available on Kubernetes 1.19 and later versions.
- If the Ingress path routing and HTTPRoute are used for the same domain, all the content routing policies from the HTTPRoute resource get lower priority than the Ingress based

content routing policies. So, it is recommended to configure all the content switching policies of theHTTPRoute resource for a given domain if advanced content routing is required.

## Configure advanced content routing for Kubernetes Ingress using the HTTPRoute CRD

This procedure shows how to deploy an HTTPRoute resource as a resource backend to support advanced content routing.

### Prerequisites

- Ensure that the ingress API version networking.k8s.io/v1 is available in the Kubernetes cluster.
- Ensure that the HTTPRoute CRD is deployed.

### Deploy the Ingress resource

Define the Ingress resource with the resource back-end pointing to a HTTPRoute custom resource in a YAML file. Specify all the front-end configurations such as certificates, front-end profiles, front-end IP address, and ingress class as part of the Ingress resource.

Following is a sample Ingress resource named as sample-ingress.yaml.

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
   name: kuard-ingress
4
   annotations:
5
       ingress.citrix.com/frontend-ip: "x.x.x.x"
6
7
       ingress.citrix.com/insecure-termination: "redirect"
8 spec:
9
    ingressClassName: citrix
    tls:
11
     - secretName: web-ingress-secret
12
    rules:
13
    - host: kuard.example.com
14
      http:
15
         paths:
         - pathType: ImplementationSpecific
16
17
           backend:
18
             resource:
19
               apiGroup: citrix.com
20
               kind: HTTPRoute
21
               name: kuard-example-route
```

After defining the Ingress resource in a YAML file, deploy the YAML file using the following command. Here, sample-ingress.yaml is the YAML file definition.

```
1 kubectl apply -f sample-ingress.yaml
```

In this example, content switching policies for the domain kuard.example.com are defined as part of the HTTPRoute custom resource called kuard-example-route. Certificates, frontend-ip, and ingress **class** are specified as part of the Ingress resource. Back-end annotations such as load balancing method and service group configurations are specified as part of the HTTPRoute custom resource.

#### **Deploy the HTTPRoute resource**

Define the HTTP route configuration in a YAML file. In the YAML file, use HTTPRoute in the kind field and in the spec section add the HTTPRoute CRD attributes based on your requirement for the HTTP route configuration.

For more information about API description and examples, see the HTTPRoute documentation.

Following is a sample HTTPRoute resource configuration. This example shows how to use query parameters based content switching for the various Kubernetes back-end microservices.

```
1 apiVersion: citrix.com/v1
2 kind: HTTPRoute
3 metadata:
4
  name: kuard-example-route
5 spec:
6
    hostname:
7
     - kuard.example.com
8
     rules:
9
     - name: kuard-blue
      match:
10
11
       - queryParams:
         - name: version
12
13
           contains: v2
14
       action:
15
        backend:
           kube:
17
             service: kuard-blue
             port: 80
18
19
     - name: kuard-green
20
      match:
21
       - queryParams:
22
         - name: version
23
           contains: v3
24
       action:
25
         backend:
26
           kube:
27
             service: kuard-green
28
             port: 80
29
     - name: kuard-default
       match:
```

```
31 - path:
32 prefix: /
33 action:
34 backend:
35 kube:
36 service: kuard-purple
37 port: 80
```

After you have defined the HTTP routes in the YAML file, deploy the YAML file. In this example, httproute is the YAML definition.

```
1 kubectl apply -f httproute.yaml
```

## Profile support for the Listener CRD

#### December 31, 2023

You can use individual entities such as HTTP profile, TCP profile, and SSL profile to configure HTTP, TCP, and SSL respectively for the Listener CRD. Profile support for the Listener CRD helps you to customize the default protocol behavior. You can also select the SSL ciphers for the SSL virtual server.

## **HTTP profile**

An HTTP profile is a collection of HTTP settings. A default HTTP profile called nshttp\_default\_profile is configured to set the HTTP configurations. These configurations are applied, by default, globally to all services and virtual servers. You can customize the HTTP configurations for a Listener resource by specifying spec.policies.httpprofile. If specified, NetScaler Ingress Controller creates a new HTTP profile with the default values derived from the default HTTP profile and configures the values specified.

It helps to derive the default values from the default HTTP profile and configures the values specified.

The following example YAML shows how to enable websocket for a given front-end virtual server.

```
1
       apiVersion: citrix.com/v1
2
       kind: Listener
3
       metadata:
4
         name: test-listener
5
         namespace: default
6
       spec:
7
         vip: x.x.x.x
8
         port: 80
9
         protocol: http
10
         policies:
```

```
11httpprofile:12config:13websocket: "ENABLED"
```

For information about all the possible key-value pairs for the HTTP profile see, HTTP profile.

Note:

The 'name' is auto-generated.

You can also specify a built-in HTTP profile or a pre-configured HTTP profile and bind it to the frontend virtual server as shown in the following example.

```
1 apiVersion: citrix.com/v1
2 kind: Listener
3 metadata:
   name: test-listener
4
5 namespace: default
6 spec:
7
   vip: x.x.x.x
8
    port: 80
   protocol: http
9
     policies:
       httpprofile:
11
12
         preconfigured: 'nshttp_default_strict_validation'
```

## **TCP profile**

A TCP profile is a collection of TCP settings. A default TCP profile called nstcp\_default\_profile is configured to set the TCP configurations. These configurations are applied, by default, globally to all services and virtual servers. You can customize the TCP settings by specifying spec.policies.tcpprofile. When you specify spec.policies.tcpprofile, NetScaler Ingress Controller creates a TCP profile that is derived from the default TCP profile and applies the values provided in the specification, and binds it to the front-end virtual server.

For information about all the possible key-value pairs for a TCP profile, see TCP profile.

Note:

The name is auto-generated.

The following example shows how to enable tcpfastopen and HyStart for the front-end virtual server.

```
1 apiVersion: citrix.com/v1
2 kind: Listener
3 metadata:
4 name: test-listener
```

```
5 namespace: default
6 spec:
7
   vip: x.x.x.x
8 port: 80
9
  protocol: http
10
     policies:
11
       tcpprofile:
12
        config:
          tcpfastopen: "ENABLED"
13
14
           hystart: "ENABLED"
```

You can also specify a built-in TCP profile or a pre-configured TCP profile name as shown in the following example:

```
1 apiVersion: citrix.com/v1
2 kind: Listener
3 metadata:
4 name: test-listener
5
   namespace: default
6 spec:
7
    vip: x.x.x.x
   port: 80
8
   protocol: http
9
10 policies:
     tcpprofile:
         preconfigured: 'nstcp_default_Mobile_profile'
12
```

## SSL profile

An SSL profile is a collection of settings for SSL entities. SSL profile makes configuration easier and flexible. You can configure the settings in a profile and bind that profile to a virtual server instead of configuring the settings on each entity. An SSL profile allows you to customize many SSL parameters such as TLS protocol and ciphers. For more information about SSL profile, see SSL profile infrastructure.

#### Note:

By default, NetScaler creates a legacy SSL profile. The legacy SSL profile has many drawbacks including non-support for advanced protocols such as SSLv3. Hence, it is recommended to enable the default SSL profiles in NetScaler before NetScaler Ingress Controller is launched.

To enable the advanced SSL profile, use the following command in the NetScaler command line:

set ssl parameter -defaultProfile ENABLED

The command enables the default SSL profile for all the existing SSL virtual servers and the SSL service groups.

You can specify spec.policies.sslprofile to customize the SSL profile. When specified, NetScaler Ingress Controller creates an SSL profile derived from the default SSL front-end profile: ns\_default\_ssl\_profile\_frontend.

For information about key-value pairs supported in the SSL profile, see SSL profile.

Note:

The **name** is auto-generated.

The following example shows how to enable TLS1.3 and HSTS for the front-end virtual server.

```
1 apiVersion: citrix.com/v1
2 kind: Listener
3 metadata:
4
    name: test-listener
5 namespace: default
6 spec:
7
    vip: x.x.x.x
8
   port: 443
9
     certificates:
10
     - secret:
11
         name: my-cert
11 name: my-ce
12 protocol: https
13 policies:
14
     sslprofile:
15
         config:
           tls13: "ENABLED"
16
           hsts: "ENABLED"
17
```

You can specify a built-in or pre-configured SSL profile name as shown in the following example:

```
1 apiVersion: citrix.com/v1
2 kind: Listener
3 metadata:
   name: test-listener
4
   namespace: default
5
6 spec:
7
   vip: x.x.x.x
   port: 443
8
9
    certificates:
    - secret:
         name: my-cert
11
   protocol: https
12
13
   policies:
      sslprofile:
14
         preconfigured: 'ns_default_ssl_profile_secure_frontend'
15
```

## SSL ciphers

The Ingress NetScaler has built-in cipher groups. By default, virtual servers use a DEFAULT cipher group for an SSL transaction. To use ciphers which are not part of the DEFAULT cipher group, you must explicitly bind them to an SSL profile. You can use spec.policies.sslciphers to provide a list of ciphers, list of built-in cipher groups, or the list of user-defined cipher groups.

Note:

The order of priority of ciphers is the same order defined in the list. The first one in the list gets the first priority and likewise.

The following example shows how to provide a list of built-in cipher suites.

```
1
     apiVersion: citrix.com/v1
2
     kind: Listener
   metadata:
3
4
       name: test-listener
       namespace: default
5
     spec:
6
7
       vip: x.x.x.x
8
       port: 443
9
       certificates:
       - secret:
           name: my-cert
11
12
      protocol: https
13
       policies:
         sslciphers:
14
15
         - 'TLS1.2-ECDHE-RSA-AES128-GCM-SHA256'
         - 'TLS1.2-ECDHE-RSA-AES256-GCM-SHA384'
16
         - 'TLS1.2-ECDHE-RSA-AES-128-SHA256'
17
18
         - 'TLS1.2-ECDHE-RSA-AES-256-SHA384'
```

For information about the list of cipher suites available in NetScaler, see SSL profile infrastructure.

Ensure that NetScaler has a user-defined cipher group for using a user-defined cipher group. Perform the following steps to configure a user-defined cipher group:

- 1. Create a user-defined cipher group. For example, MY-CUSTOM-GROUP.
- 2. Bind all the required ciphers to the user-defined cipher group.
- 3. Note down the user-defined cipher group name.

For detailed instructions, see Configure a user-defined cipher group.

**Note:** The order of priority of ciphers is the same order defined in the list. The first one in the list gets the first priority and likewise.

The following example shows how to provide a list of built-in cipher groups and/or user defined cipher group. The user-defined cipher groups must be present in NetScaler before you apply it to Listener.

```
apiVersion: citrix.com/v1
1
    kind: Listener
2
3
    metadata:
4
      name: test-listener
5
      namespace: default
6 spec:
7
     vip: x.x.x.x
     port: 443
8
9
      certificates:
      - secret:
11
           name: my-cert
     protocol: https
12
13
     policies:
       sslciphers:
14
15
        - 'SECURE'
        - 'HIGH'
16
        - 'MY-CUSTOM-CIPHERS'
17
```

In the preceding example, SECURE and HIGH are built-in cipher groups in NetScaler. MY-CUSTOM-CIPHERS is the pre-configured user-defined cipher groups.

**Note:** If you have specified the pre-configured SSL profile, you must bind the ciphers manually through NetScaler and spec.policies.sslciphers is not applied on the pre-configured SSL profile.

**Note:** The built-in cipher groups can be used in Tier-1 and Tier-2 NetScaler. The user-defined cipher group can be used only in a Tier-1 NetScaler.

## **Analytics profile**

Analytics profile enables NetScaler to export the type of transactions or data to an external platform. If you are using NetScaler Observability Exporter to collect metrics and transactions data and export it to endpoints such Elasticsearch or Prometheus, you can configure the analytics profile to select the type of data that needs to be exported.

Note:

For the Analytics profile to be functional, you must configure the NetScaler Observability Exporter. Analytics configuration support using ConfigMap.

The following example shows how to enable webinsight and tcpinsight in the analytics profile.

```
1 apiVersion: citrix.com/v1
2 kind: Listener
3 metadata:
4 name: test-listener
5 namespace: default
```

```
spec:
6
7
      vip: x.x.x.x
8
     port: 443
9
      certificates:
10
       - secret:
           name: my-cert
11
12
       protocol: https
     policies:
13
14
       analyticsprofile:
         config:
16
          - type: webinsight
17
          - type: tcpinsight
```

The following example shows how to select the additional parameters for the type of webinsight which you want to be exported to NetScaler Observability Exporter. For information about the valid key-value pair, see Analytics Profile.

```
1 apiVersion: citrix.com/v1
   kind: Listener
2
3 metadata:
4
     name: test-listener
5
     namespace: default
  spec:
6
    vip: x.x.x.x
7
     port: 443
8
9
     certificates:
10
      - secret:
11
          name: my-cert
12
     protocol: https
13
      policies:
14
       analyticsprofile:
         config:
16
          - type: webinsight
17
            parameters:
             httpdomainname: "ENABLED"
18
19
              httplocation: "ENABLED"
```

The following example shows how to use pre-configured analytics profiles.

```
apiVersion: citrix.com/v1
1
2
     kind: Listener
3
     metadata:
      name: test-listener
4
5
      namespace: default
   spec:
6
7
      vip: x.x.x.x
8
     port: 443
9
      certificates:
10
       - secret:
           name: my-cert
11
12
       protocol: https
13
       policies:
      analyticsprofile:
14
```

```
15 preconfigured:
16 - 'custom-websingiht-analytics-profile'
17 - 'custom-tcpinsight-analytics-profile'
```

## IP address management using the IPAM controller

#### December 10, 2024

The IPAM controller is an application provided by NetScaler for IP address management and it runs in parallel to NetScaler Ingress Controller in the Kubernetes cluster. You can assign IP addresses to ingress resources, services of type LoadBalaIncer, and listener resources from a specified IP address range using the IPAM controller. NetScaler Ingress Controller configures an IP address allocated to a service, or ingress, or listener resource as a virtual IP address (VIP) in NetScaler MPX or NetScaler VPX.

The IPAM controller requires the VIP CustomResourceDefinition (CRD) provided by NetScaler. The VIP CRD is used for internal communication between NetScaler Ingress Controller and the IPAM controller.

## Infoblox integration with IPAM controller

With Infoblox integration, the IPAM controller assigns IP addresses to services of type LoadBalancer, ingress, or listener resources from Infoblox.

The following annotations are supported in an ingress resource, listener resource, and service of type LoadBalancer respectively: ingress.citrix.com/ipam-range, listeners.citrix.com / ipam-range, and service.citrix.com / ipam-range. You can define the IP address range for an IPAM controller by using these annotations. The IP address range must either be an IPAM VIP range or an Infoblox VIP range.

#### Note:

- Infoblox integration is supported for versions up to 2.12.3, starting from IPAM Controller version 2.2.10.
- You must not update or delete the IP address allocated by IPAM using Infoblox.

The IPAM controller poses the following challenges in a multi-cluster setup:

- An IP address must be provided manually.
- A common IP address must be assigned for each application deployed across multiple clusters.

• Limitations in sharing the IP address range between IPAM controllers in other clusters, adding complexity to the task.

Infoblox integration with IPAM helps with the following tasks:

- Assign an available IP address from the specified IP address range using Infoblox network management.
- Request the IP address associated with a domain name, ensuring the retrieval of a pre-existing IP address.
- In a multi-cluster ingress setup, the application deployed across various clusters can be accessed using a single, consistent IP address.

During initialization, the IPAM controller connects with Infoblox and creates a network view if it's not already available. This network view creation is based on the network range in the DNS view that you provide. If the network range in the DNS view is not provided, the IPAM controller creates a network view.

## Assign an IP address using the IPAM controller

To configure an ingress resource or listener resource or service of type LoadBalancer with an IP address from the IPAM controller, perform the following steps:

- 1. Deploy the VIP CRD
- 2. Deploy the NetScaler Ingress Controller
- 3. Deploy the IPAM controller
- 4. Deploy the application and ingress resource

#### Step 1: Deploy the VIP CRD

Perform the following step to deploy the NetScaler VIP CRD which enables communication between the NetScaler Ingress Controller and the IPAM controller.

```
1 kubectl create -f https://raw.githubusercontent.com/netscaler/netscaler
-k8s-ingress-controller/master/crd/vip/vip.yaml
```

For more information on VIP CRD, see the VIP CustomResourceDefinition.

#### Step 2: Deploy NetScaler Ingress Controller

#### Prerequisites

• Kubernetes cluster and a kubectl command-line tool to communicate with the cluster.

• Create a secret using NetScaler VPX or NetScaler MPX credentials by using the following command:

```
1 kubectl create secret generic nslogin --from-literal=username
=<username> --from-literal=password=<password>
```

• Add the NetScaler Helm chart repository by using the following command:

```
1 helm repo add netscaler https://netscaler.github.io/netscaler-
helm-charts/
```

#### **Deploy NSIC**

1. Update the NetScaler Helm chart repository by using the following command:

```
1 helm repo update netscaler
```

2. Update values.yaml to configure NetScaler Ingress Controller as described as following.

Sample values.yaml:

```
1
    license:
2
      accept: yes
3
    adcCredentialSecret: nslogin # K8s Secret Name
    nsIP: <x.x.x> # CLIP (for appliances in Cluster mode), SNIP (for
4
         appliances in High Availability mode), NSIP (for standalone
         appliances)
5
    openshift: false # set to true for OpenShift deployments
    entityPrefix: cluster1 # unique for each NSIC instance.
6
7
    ipam: true
    ingressClass: ['cic-vpx'] # ingress class used in the ingress
8
        resources
9
    serviceClass: ['netscaler'] # To use service type LB, specify
        the service class
```

3. Install NetScaler Ingress Controller for your NetScaler VPX or using the following Helm command:

```
1 helm install nsic netscaler/netscaler-ingress-controller -f values
    .yaml
```

For detailed information about deploying and configuring NetScaler Ingress Controller using Helm charts, see the Helm chart repository.

#### **Step 3: Deploy IPAM controller**

Prerequisites

• For Infoblox integration, ensure that you create a user role in Infoblox with the following permissions:

#### - IP Address Management:

- \* Allocate and manage IP addresses.
- \* Create, update, and delete host records.
- Extensible Attributes Management:
  - \* Create, update, and delete extensible attributes.
- For Infoblox integration, create a Kubernetes secret with Infoblox user credentials by running the following command:

```
1 kubectl create secret generic infobloxlogin --from-literal=
username=<Infoblox_username> --from-literal=password=<
Infoblox_password>
```

In this procedure, let's deploy an IPAM controller and integrate Infoblox with the IPAM controller.

1. Add the NetScaler Helm chart repository to your local registry by using the following command:

```
1 helm repo add netscaler https://netscaler.github.io/netscaler-helm
        -charts/
```

2. Install NetScaler IPAM controller by using the following command:

```
1 helm install netscaler-ipam-controller netscaler/netscaler-ipam-
controller -f values.yaml
```

Sample values.yaml:



For information about all the configurable parameters while installing the IPAM controller by using Helm charts, see the Helm chart repository.

#### **Step 4: Deploy Ingress resources**

Perform the following steps to deploy a sample application and ingress resource.

1. Deploy the CNN application by using the following command:

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: apps/v1
3 kind: Deployment
4 metadata:
5
    name: cnn-website
6
    labels:
7
      name: cnn-website
8
       app: cnn-website
9 spec:
10
   selector:
11
     matchLabels:
12
        app: cnn-website
13
   replicas: 2
14
   template:
15
      metadata:
16
        labels:
17
          name: cnn-website
18
           app: cnn-website
19
      spec:
20
        containers:
21
         - name: cnn-website
22
          image: quay.io/sample-apps/cnn-website:v1.0.0
23
          ports:
24
           - name: http-80
            containerPort: 80
25
26
           - name: https-443
27
            containerPort: 443
28 ---
29 apiVersion: v1
30 kind: Service
31 metadata:
32
    name: cnn-website
33
     labels:
34
       app: cnn-website
35 spec:
    type: NodePort
37
     ports:
38
     - name: http-80
     port: 80
39
40
      targetPort: 80
41
     - name: https-443
42
      port: 443
43
       targetPort: 443
44
    selector:
45
      name: cnn-website
46 EOF
```

2. Deploy the ingress resource to send traffic to the CNN application by using the following command:

```
1
     kubectl apply -f - <<EOF
2
     apiVersion: networking.k8s.io/v1
3
     kind: Ingress
4
     metadata:
5
       name: cnn-ingress
6
       annotations:
7
         ingress.citrix.com/ipam-range: "<IPAM-vip-range or Infoblox-
             vip-range>"
8
     spec:
9
       ingressClassName: cic-vpx
10
       rules:
11
       - host: www.cnn.com
12
         http:
13
           paths:
            - path: /
14
15
              pathType: Prefix
16
              backend:
                service:
17
18
                  name: cnn-website
19
                  port:
                    number: 80
21
     EOF
```

For the preceding ingress example, the order of IP assignment is as follows:

- If the default NS\_VIP environment variable is provided, NetScaler Ingress Controller makes a request to IPAM controller only if the range-name (ingress.citrix.com/ipam-range:) is provided in the ingress. If the annotation is not provided, NS\_VIP is used for that ingress.
- If the default NS\_VIP environment variable is not provided, NetScaler Ingress Controller always make a request to IPAM controller for IP assignment.

#### **IP address allocations**

- For services of type LoadBalancer, a unique IP address is allocated to each service from the VIP range or Infoblox VIP range.
- For an ingress resource, an IP address in the specified IP range is allocated. When more ingress resources refer to the same VIP range, the IP address allocated to the first ingress resource is allocated to all the other ingress resources.
- Services of type LoadBalancer, ingress resources, and listener resources can use NetScaler IPAM controller for IP address allocations at the same time. If an IP address is allocated to any one resource type, it is not available for another resource type. But, the same IP address can be allocated to multiple ingress resources.

**IP address range associated with a unique name** You can assign a unique name to the IP address range and define the range in the VIP\_RANGE environment variable. This way of assigning the IP address range enables you to differentiate between the IP address ranges. When you create the services of type LoadBalancer, you can use the service.citrix.com/ipam-range annotation in the service definition to specify the IP address range to use for IP address allocation.

#### **Multiple IP address allocations**

For ingress resources, an IP address can be allocated multiple times by specifying the Helm chart parameter reuseIngressVip as **true**, because multiple ingress resources might be handled by a single content switching virtual server. If the specified IP range has only a single IP address, it is allocated multiple times. And, if the IP range consists of multiple IP addresses, only one of them is allocated repeatedly.

To facilitate multiple allocations, the IPAM controller tracks the allocated IP addresses. The IPAM controller places an IP address into the free pool only when all allocations of that IP address by ingress resources are released.

## **Apply CRDs through annotations**

#### December 31, 2023

You can now apply CRDs such as Rewrite and Responder, Ratelimit, Auth, WAF, and Bot for ingress resources and services of type load balancer by referring them using annotations. Using this feature, when there are multiple services in an Ingress resource, you can apply the rewrite and responder policy for a specific service or all the services based on your requirements.

The following are the two benefits of this feature:

- You can apply a CRD at a per-ingress, per-service level. For example, the same service referred through an internal VIP may have different set of rewrite-responder policies compared to the one exposed outside.
- Operations team can create CRD instances without specifying the service names. The application developers can choose the right policies based on their requirements.

#### Note:

CRD instances should be created without service names.

## Ingress annotation for referring CRDs

An Ingress resource can refer a Rewrite and Responder CRD directly using the ingress.citrix. com/rewrite-responder annotation.

The following are different ways of referring the rewrite-responder CRD using annotations.

• You can apply the Rewrite and Responder CRD for all the services referred in the given ingress using the following format:

```
ingress.citrix.com/rewrite-responder_crd: <Rewritepolicy Custom-
resoure-instance-name>
```

Example:

ingress.citrix.com/rewrite-responder\_crd: "blockurlpolicy"

In this example, the Rewrite and Responder policy is applied for all the services referred in the given ingress.

• You can apply the Rewrite and Responder CRD to a specified Kubernetes service in an Ingress resource using the following format:

Example:

```
ingress.citrix.com/rewrite-responder_crd: '{
    "frontendsvc": "blockurlpolicy", "backendsvc": "
        addresponseheaders" }
    '
```

In this example, the rewrite policy blockurlpolicy is applied on the traffic coming to the frontendsvc service and the addresponseheaders policy is applied to the backendsvc service coming through the current ingress.

You can also apply the Auth, Bot, WAF, and Ratelimit CRDs using ingress annotations:

The following table explains the annotations and examples for Auth, Bot, WAF, and Ratelimit CRDs.

Annotation	Examples	Description
<pre>ingress.citrix.com/ bot crd</pre>	<pre>ingress.citrix.com/ bot crd: !{ "frontend</pre>	Applies the botdefense
	": "botdefense" } '	the front-end service.

#### NetScaler ingress controller

Annotation	Examples	Description
ingress.citrix.com/ auth_crd	<pre>ingress.citrix.com/ auth_crd: '{ " frontend": " authexample" } '</pre>	Applies the authexample policy to the front-end service.
<pre>ingress.citrix.com/ waf_crd</pre>	<pre>ingress.citrix.com/ waf_crd: "wafbasic"</pre>	Applies the WAF policy wafbasic to all services in the Ingress
<pre>ingress.citrix.com/ ratelimit_crd</pre>	<pre>ingress.citrix.com/ ratelimit_crd: " throttlecoffeeperclie "</pre>	Applies the rate limit policy throttlecoffeeperclientip ntito all services in the Ingress.

## Service of type LoadBalancer annotation for referring Rewrite and Responder CRD

A service of type LoadBalancer can refer a Rewrite and Responder CRD using annotations.

The following is the format for the annotation:

```
1 service.citrix.com/rewrite-responder: <Rewritepolicy Custom-resoure-
instance-name>
```

## Listener CRD support for Ingress through annotation

#### July 7, 2025

Ingress is a standard Kubernetes resource that specifies HTTP routing capability to back-end Kubernetes services. NetScaler Ingress Controller provides various annotations to fine-tune the Ingress parameters for both front-end and back-end configurations. For example, using the ingress.citrix .com/frontend-ip annotation you can specify the front-end listener IP address configured in NetScaler by NetScaler Ingress Controller. Similarly, there are other front-end annotations to finetune HTTP and SSL parameters. When there are multiple Ingress resources and if they share front-end IP and port, specifying these annotations in each Ingress resource is difficult.

Sometimes, there is a separation of responsibility between network operations professionals (NetOps) and developers. NetOps are responsible for coming up with front-end configurations like frontend IP, certificates, and SSL parameters. Developers are responsible for the HTTP routing and backend configurations. NetScaler Ingress Controller already provides content routing CRDs such as listener CRD for front-end configurations and HTTProute for back-end routing logic. Now, Listener CRD can be applied for Ingress resources using an annotation provided by NetScaler.

Through this feature, you can use the Listener CRD for your Ingress resource and separate the creation of the front-end configuration from the Ingress definition. Hence, NetOps can separately define the Listener resource to configure front-end IP, certificates, and other front-end parameters (TCP, HTTP, and SSL). Any configuration changes can be applied to the listener resources without changing each Ingress resource. In NetScaler, a listener resource corresponds to content switching virtual servers, SSL virtual servers, certkeys and front-end HTTP, SSL, and TCP profiles.

Notes:

- For the Analytics profile to be functional, you must configure the NetScaler Observability Exporter. For more information, see Analytics configuration support using ConfigMap.
- While using this feature, you must ensure that all ingresses with the same front-end IP and port refer to the same Listener resource. For Ingresses that use the same front-end IP and port combinations, one Ingress referring to a listener resource and another Ingress referring to the ingress.citrix.com/frontend-ip annotation is not supported.

## Restrictions

When Listener is used for the front-end configurations, the following annotations are ignored and there may not be any effect:

- ingress.citrix.com/frontend-ip
- Ingress.citrix.com/frontend-ipset-name
- ingress.citrix.com/secure-port
- ingress.citrix.com/insecure-port
- ingress.citrix.com/insecure-termination
- ingress.citrix.com/secure-service-type
- ingress.citrix.com/insecure-service-type
- ingress.citrix.com/csvserver
- ingress.citrix.com/frontend-tcpprofile
- ingress.citrix.com/frontend-sslprofile
- ingress.citrix.com/frontend-httpprofile

## **Deploying a Listener CRD resource for Ingress**

Using the ingress.citrix.com/listener annotation, you can specify the name and namespace of the Listener resource for the ingress in the form of namespace/name. The namespace is not required if the Listener resource is in the same namespace as that of Ingress. Following is an example for the annotation:

1 ingress.citrix.com/listener: default/listener1

Here, **default** is the namespace of the Listener resource and <code>listener1</code> is the name of the Listener resource which specifies the front-end parameters.

Perform the following steps to deploy a Listener resource for the Ingress:

1. Create a Listener resource (listener.yaml) as follows:

```
1 apiVersion: citrix.com/v1
2 kind: Listener
3 metadata:
4
    name: my-listener
5
     namespace: default
6 spec:
7
     ingressclass: citrix
8
     vip: '192.168.0.1' # Virtual IP address to be used, not required
         when CPX is used as ingress device
9
    port: 443
10
     protocol: https
11
     redirectPort: 80
12
     secondaryVips:
     - "10.0.0.1"
13
14
     - "1.1.1.1"
15
    policies:
16
      httpprofile:
         config:
17
18
           websocket: "ENABLED"
19
       tcpprofile:
         config:
           sack: "ENABLED"
21
       sslprofile:
23
         config:
24
           ssl3: "ENABLED"
25
       sslciphers:
26
       - SECURE
       - MEDIUM
27
28
       analyticsprofile:
         config:
29
30
         - type: webinsight
31
           parameters:
32
             allhttpheaders: "ENABLED"
       csvserverConfig:
         rhistate: 'ACTIVE'
34
```

Here, the Listener resource my-listener in the default namespace specifies the front-end configuration such as VIP, secondary VIPs, HTTP profile, TCP profile, SSL profile, and SSL ciphers. It creates a content switching virtual server in NetScaler on port 443 for HTTPS traffic, and all HTTP traffic on port 80 is redirected to HTTPS.

### Note:

The vip field in the Listener resource is not required when NetScaler CPX is used as an ingress device. For NetScaler VPX, VIP is the same as the pod IP address which is automatically configured by NetScaler Ingress Controller.

#### 2. Apply the Listener resource.

1 kubectl apply -f listener.yaml

3. Create an Ingress resource (ingress.yaml) by referring to the Listener resource.

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
    name: my-ingress
4
5
    namespace: default
6
    annotations:
7
      ingress.citrix.com/listener: my-listener
8 spec:
9
    ingressClassName: citrix
10 tls:
11
     - secretName: my-secret
      hosts:
12
     - example.com
13
   rules:
14

    host: example.com

15
16
      http:
17
         paths:
18
         - path: /
19
           pathType: Prefix
           backend:
21
             service:
               name: kuard
23
               port:
24
                 number: 80
```

Here, the ingress resource my-ingress refers to the Listener resource my-listener in the default namespace for front-end configurations.

4. Apply the ingress resource.

```
1 kubectl apply -f ingress.yaml
```

## **Certificate management**

There are two ways in which you can specify the certificates for Ingress resources. You can specify the certificates as part of the Ingress resource or provide the certificates as part of the Listener resource.

#### **Certificate management through Ingress resource**

In this approach, all certificates are specified as part of the regular ingress resource as follows. A Listener resource does not specify certificates. In this mode, you need to specify certificates as part of the Ingress resource.

```
1
     apiVersion: networking.k8s.io/v1
2
   kind: Ingress
3 metadata:
4
     name: my-ingress
5
      namespace: default
     annotations:
6
7
      ingress.citrix.com/listener: my-listener
8 spec:
9
      ingressClassName: citrix
10
     tls:
11
      - secretName: my-secret
12
       hosts:
      - example.com
13
     rules:
14
15
     - host: example.com
16
       http:
17
          paths:
          - path: /
18
19
            pathType: Prefix
20
            backend:
              service:
22
                name: kuard
23
                port:
24
                  number: 80
```

## **Certificate management through Listener resource**

In this approach, certificates are provided as part of the Listener resource. You do not have to specify certificates as part of the Ingress resource.

The following Listener resource example shows certificates.

```
1 apiVersion: citrix.com/v1
2 kind: Listener
3 metadata:
   name: my-listener
4
  namespace: default
5
6 spec:
7 ingressclass: citrix
8 certificates:
9 - secret:
10
       name: my-secret
     # Secret named 'my-secret' in current namespace bound as default
11
         certificate
```

```
12
   default: true
13
     - secret:
         # Secret 'other-secret' in demo namespace bound as SNI
14
            certificate
15
         name: other-secret
         namespace: demo
16
     vip: '192.168.0.1' # Virtual IP address to be used, not required when
17
         CPX is used as ingress device
     port: 443
18
     protocol: https
19
20
     redirectPort: 80
```

In the Ingress resource, secrets are not specified as shown in the following example.

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
4 name: my-ingress
5 namespace: default
6
    annotations:
      ingress.citrix.com/listener: my-listener
7
8 spec:
9 ingressClassName: citrix
10 tls:
   # TLS field is empty as the certs are specified in Listener
11
12
    rules:
     - host: example.com
13
14
       http:
15
         paths:
16
         - path: /
           pathType: Prefix
17
18
           backend:
19
             service:
20
              name: kuard
21
               port:
22
                 number: 80
```

# Configuring consistent hashing algorithm using NetScaler Ingress Controller

#### December 31, 2023

Load balancing algorithms define the criteria that the NetScaler appliance uses to select the service to which to redirect each client request. Different load balancing algorithms use different criteria and consistent hashing is one the load balancing algorithms supported by NetScaler.

Consistent hashing algorithms are often used to load balance when the back-end is a caching server to achieve stateless persistency.

Consistent hashing can ensure that when a cache server is removed, only the requests cached in that specific server is rehashed and the rest of the requests are not affected. For more information on the consistent hashing algorithm, see the NetScaler documentation.

You can now configure the consistent hashing algorithm on NetScaler using NetScaler Ingress Controller. This configuration is enabled with in the NetScaler Ingress Controller using a ConfigMap.

## **Configure hashing algorithm**

A new parameter NS\_LB\_HASH\_ALGO is introduced in the NetScaler Ingress Controller ConfigMap for hashing algorithm support.

Supported environment variables for consistent hashing algorithm using ConfigMap under the NS\_LB\_HASH\_ALGO parameter:

- hashFingers: Specifies the number of fingers to be used for the hashing algorithm. Possible values are from 1 to 1024. Increasing the number of fingers provides better distribution of traffic at the expense of extra memory.
- hashAlgorithm: Specifies the supported algorithm. Supported algorithms are **default**, jarh, prac.

The following example shows a sample ConfigMap for configuring consistent hashing algorithm using NetScaler Ingress Controller. In this example, the hashing algorithm is used as Prime Re-Shuffled Assisted CARP (PRAC) and the number of fingers to be used in PRAC is set as 50.

```
1
       apiVersion: v1
2
       kind: ConfigMap
3
       metadata:
4
       name: cic-configmap
5
       labels:
           app: citrix-ingress-controller
6
7
       data:
8
       NS\_LB\_HASH\_ALGO:
9
           hashFingers: 50
10
           hashAlgorithm: 'prac'
```

## Add DNS records using NetScaler Ingress Controller

December 31, 2023

A DNS address record is a mapping of the domain name to the IP address.

When you want to use NetScaler as a DNS resolver, you can add the DNS records on NetScaler using NetScaler Ingress Controller.

For more information on creating DNS records on NetScaler, see the NetScaler documentation.

#### **Adding DNS records for Ingress resources**

You need to enable the following environment variable during the NetScaler Ingress Controller deployment to add DNS records for an Ingress resource.

NS\_CONFIG\_DNS\_REC: This variable is configured at the boot time and cannot be changed at runtime. Possible values are **true** or **false**. The default value is false and you need to set it as true to enable the DNS server configuration. When you set the value as **true**, an address record is created on NetScaler.

### Adding DNS records for services of type LoadBalancer

You need to perform the following tasks to add DNS records for services of type LoadBalancer:

- Enable the NS\_SVC\_LB\_DNS\_REC environment variable by setting the value as True for adding DNS records for a service of type LoadBalancer.
- Specify the DNS host name using the service.citrix.com/dns-hostname annotation.

When you create a service of type LoadBalancer with the service.citrix.com/dns-hostname annotation, NetScaler Ingress Controller adds the DNS record on NetScaler. The DNS record is configured using the domain name specified in the annotation and the external IP address assigned to the service.

When you delete a service of type LoadBalancer with the service.citrix.com/dns-hostname annotation, NetScaler Ingress Controller removes the DNS records from the NetScaler.

NetScaler Ingress Controller also removes the stale entries of DNS records during boot up if the service is not available.

The following example shows a sample service of type LoadBalancer with the annotation configuration to add DNS records to NetScaler:

```
1 apiVersion: v1
2 kind: Service
3 metadata:
4
    name: guestbook
5
   annotations:
        service.citrix.com/dns-hostname: "guestbook.com"
6
7 spec:
    loadBalancerIP: "192.2.212.16"
8
9 type: LoadBalancer
10
     ports:
     - port: 9006
      targetPort: 80
13
       protocol: TCP
```

14 selector: 15 app: guestbook

## **Open policy agent support for Kubernetes with NetScaler**

December 31, 2023

Open policy agent (OPA) is an open source, general-purpose policy engine that unifies policy enforcement across different technologies and systems. OPA provides a high-level declarative language that lets you specify policy as code and simple APIs to offload policy decision-making from your software. Using OPA, you can decouple policy decision-making from policy enforcement. You can use OPA to enforce policies through NetScaler in a Kubernetes environment.

With OPA, you can create a centralized policy-decision making system for an environment involving multiple NetScalers or multiple devices which are distributed. The advantage of this approach is you have to make changes only on the OPA server for any decision specific changes applicable to multiple devices.

For more information on OPA, see the OPA documentation.

The OPA integration on NetScaler can be supported through HTTP callout, where OPA can be used with or without authentication. An HTTP callout is an HTTP or HTTPS request that the NetScaler appliance generates and sends to an external application as part of the policy evaluation.

For more information on the HTTP callout support, see the HTTP callout documentation.

For more information regarding authentication support, see the Authentication and authorization policies for Kubernetes with NetScaler.

The following diagram provides an overview of how to integrate OPA with the NetScaler cloud native solution.

OPA integration

In the OPA integration diagram, each number represents the corresponding task in the following list:

- 1. Creating the required Kubernetes objects using Kubernetes commands. This step should include creating the CRD to send the HTTP callout to the OPA server.
- 2. Configuring NetScaler. NetScaler is automatically configured by NetScaler Ingress Controller based on the created Kubernetes objects.
- 3. Sending user request for resources from client. The user might get authenticated if authentication CRDs are created.

- 4. Sending HTTP callout to OPA server in JSON format from NetScaler carrying authorization parameters.
- 5. Sending authorization decision from OPA server based on the rules defined in REGO, the policy language for OPA.
- 6. Sending response to the client based on the authorization decision.

### **Example use cases**

#### Example 1: Allow or deny access to resources based on the client source IP address

Following is an example HTTP callout policy to the OPA server using rewrite policy CRD to allow or deny access to resources based on the client source IP address and the corresponding OPA rules.

In the example, the OPA server responds with "result": **true** if the client source IP address is 192.2.162.0/24, else it responds with "result": **false**.

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4
    name: calloutexample
5 spec:
6 responder-policies:
7
       - servicenames:
8
           - frontend
9
         responder-policy:
           respondwith:
             http-payload-string: '"HTTP/1.1 401 Access denied\r\n\r\n"' #
11
                 Access is denied if the respose from OPA server contains
                 false.
           respond-criteria: 'sys.http_callout("callout_name").CONTAINS("
12
               false")'
13
           comment: 'Invalid access'
14
     httpcallout_policy:
15
       - name: callout_name
         server_ip: "192.2.156.160" #OPA Server IP
16
17
         server_port: 8181 #OPA Server Port
18
         http_method: 'POST'
19
         host_expr: "\"192.2.156.160\""
         url_stem_expr: "\"/v1/data/example/allow\"" #URL stem expression
20
            to be used
21
         body_expr: '"{
    \"input\": {
22
    \"clientinfo\": [{
23
    \"id\": \"ci\", \"ip\": [\""+ CLIENT.IP.SRC +"\"] }
24
25
    ] }
26
    }
    "' #JSON to OPA server carrying client IP
27
28
         headers:
```
```
29 - name: Content-Type
30 expr: '"application/json"'
31 return_type: TEXT
32 result_expr: "HTTP.RES.BODY(100)"
```

Following are the rules defined through the Rego policy language on the OPA server for the HTTP callout policy for this example:

```
1
       package example
2
       default allow = false
                                                              # unless
3
           otherwise defined, allow is false
4
5
       allow = true {
                                            # allow is true if...
6
           count(violation) != 0
                                                              # the ip
               matches regex.
        }
8
9
10
11
       violation[client.id] {
12
                                    # a client is in the violation set if...
           client := input.clientinfo[_]
13
14
            regex.match("192.2.162.", client.ip[_])
                                                              # the client is
                not part of 192.2.162.0/24 network.
        }
15
```

#### Example 2: Allow or deny access based on user group after authentication

Following is an example HTTP callout policy to the OPA server using rewrite policy CRD to allow or deny access to resources based on user group after authentication and the corresponding OPA rules.

In this example, the OPA server responds with "result": **true** if the user is part of the beverages group, else it responds with "result": **false**.

Following is the HTTP callout policy to the OPA server through the rewrite policy CRD.

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4
    name: calloutexample
5 spec:
   responder-policies:
6
7
       - servicenames:
           - frontend
8
9
         responder-policy:
           respondwith:
10
             http-payload-string: '"HTTP/1.1 401 Access denied\r\n\r\n"' #
11
                Access is denied if the respose from OPA server contains
                 false.
```

```
respond-criteria: 'sys.http_callout("callout_name").CONTAINS("
12
               false")'
            comment: 'Invalid access'
13
14
15
     httpcallout_policy:
16
       - name: callout_name
         server_ip: "192.2.156.160" #OPA Server IP
17
         server_port: 8181 #OPA Server Port
18
19
         http_method: 'POST'
         host_expr: "\"192.2.156.160\""
20
         url_stem_expr: "\"/v1/data/example/allow\"" #URL stem expression
             to be used
22
         body_expr: '"{
    \"input\": {
23
    \"users\": [{
24
25
    \"name\": \""+ AAA.USER.NAME +"\", \"group\": [\""+ AAA.USER.GROUPS
        +"\"] }
26
    ] }
27
     }
    "' #JSON to OPA server carrying username and group information
28
29
         headers:
         - name: Content-Type
31
            expr: '"application/json"'
32
         return_type: TEXT
33
         result_expr: "HTTP.RES.BODY(100)"
```

Following are the rules defined through the Rego language on the OPA server for this example:

```
1
       package example
2
       default allow = false
                                                              # unless
3
           otherwise defined, allow is false
4
5
       allow = true {
                                            # allow is true if...
6
7
           count(isbeveragesuser) != 0
                                                              # the user is
               part of beverages group.
        }
8
9
10
11
       isbeveragesuser[user.name] {
12
                             # a user is beverages user...
13
           user := input.users[_]
           user.group[_] == "beverages"
                                                              # if it is part
14
                of beverages group.
15
        }
```

You can perform authentication using the request header (401 based) or through forms based.

Following is a sample authentication policy using request header-based authentication. In this policy, local authentication is used.

```
1 apiVersion: citrix.com/v1beta1
```

```
2 kind: authpolicy
3 metadata:
4
   name: localauth
5 spec:
6
       servicenames:
7
       - frontend
8
9
       authentication_mechanism:
         using_request_header: 'ON'
11
       authentication_providers:
13
14
            - name: "local-auth-provider"
              basic_local_db:
15
16
                  use_local_auth: 'YES'
17
18
       authentication_policies:
19
20
            - resource:
                path: []
21
22
                method: []
              provider: ["local-auth-provider"]
23
24
25
       authorization_policies:
26
27
            - resource:
28
                path: []
29
                method: []
30
                claims: []
```

Following is a sample authentication policy using form-based authentication. In this policy, localbased authentication is used.

```
1 apiVersion: citrix.com/v1beta1
2 kind: authpolicy
3 metadata:
    name: localauth
4
5 spec:
      servicenames:
6
7
       - frontend
8
9
       authentication mechanism:
         using_forms:
           authentication_host: "fqdn_authenticaton_host"
11
           authentication_host_cert:
13
             tls_secret: authhost-tls-cert-secret
14
           vip: "192.2.156.156"
16
       authentication_providers:
17
18
           - name: "local-auth-provider"
19
             basic_local_db:
20
                 use_local_auth: 'YES'
```

```
21
22
        authentication_policies:
23
24
            - resource:
25
                path: []
26
                method: []
              provider: ["local-auth-provider"]
27
28
29
        authorization_policies:
32
            - resource:
33
                path: []
34
                method: []
                claims: []
```

# Example 3: Allow or deny access based on authentication attributes obtained during authentication

Following is an example HTTP callout policy to the OPA server using the rewrite policy CRD to allow or deny access based on authentication attributes obtained during authentication and the corresponding OPA rules.

In the example, the OPA server responds with "result":**true**' if the user member of attribute contains grp1, else it responds with "result":**false**.

The following is the sample HTTP callout policy to the OPA server through the rewrite policy CRD:

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4
   name: calloutexample
5 spec:
6 responder-policies:
7
       - servicenames:
8
           - frontend
         responder-policy:
9
10
           respondwith:
             http-payload-string: '"HTTP/1.1 401 Access denied\r\n\r\n"' #
11
                 Access is denied if the respose from OPA server contains
                 false.
           respond-criteria: 'sys.http_callout("callout_name").CONTAINS("
12
              false")'
           comment: 'Invalid access'
13
14
15
     httpcallout_policy:
16
       - name: callout_name
17
         server_ip: "192.2.156.160" #OPA Server IP
18
         server_port: 8181 #OPA Server Port
19
         http_method: 'POST'
```

```
host_expr: "\"192.2.156.160\""
20
21
         url_stem_expr: "\"/v1/data/example/allow\"" #URL stem expression
             to be used
         body_expr: '"{
22
    \"input\": {
23
24
    \"users\": [{
    \"name\": \""+ AAA.USER.NAME +"\", \"attr\": [\""+ aaa.user.attribute
25
        ("memberof") +"\"] }
26
    ] }
27
    }
    "' #JSON to OPA server carrying username and "memberof" attribute
28
       information
29
         headers:
         - name: Content-Type
           expr: '"application/json"'
31
32
         return_type: TEXT
         result_expr: "HTTP.RES.BODY(100)"
```

Following are the rules defined through the Rego language on the OPA server for this example:

```
1
       package example
2
       default allow = false
                                                              # unless
3
           otherwise defined, allow is false
4
5
       allow = true {
                                            # allow is true if...
6
           count(isbeveragesuser) != 0
7
                                                              # the user is
               part of grp1.
8
        }
9
10
11
       isbeveragesuser[user.name] {
                             # a user is part of allow group...
12
13
           user := input.users[
                                 1
           regex.match("CN=grp1", user.attr[_])
14
                                                              # if it is part
                of grp1 group. }
```

You can perform authentication using request header (401 based) or through forms based. In this example, LDAP authentication is used, where the user member of attribute is obtained from the LDAP server during authentication.

Following is a sample authentication policy using request header-based authentication.

```
1 apiVersion: citrix.com/vlbetal
2 kind: authpolicy
3 metadata:
4 name: ldapauth
5 spec:
6 servicenames:
7 - frontend
8
9 authentication_mechanism:
```

```
10
         using_request_header: 'ON'
11
12
       authentication_providers:
13
            - name: "ldap-auth-provider"
14
              ldap:
                  server_ip: "192.2.156.160"
15
                  base: 'dc=aaa,dc=local'
16
                  login_name: accountname
17
18
                  sub_attribute_name: CN
                  server_login_credentials: ldapcredential
19
20
                  attributes_to_save: memberof #memberof attribute to be
                      obtained from LDAP server for user
21
       authentication_policies:
23
            - resource:
24
                path: []
25
                method: []
              provider: ["ldap-auth-provider"]
26
27
       authorization_policies:
28
29
            - resource:
                 path: []
                 method: []
31
                 claims: []
```

Following is a sample authentication policy using form-based authentication.

```
1 apiVersion: citrix.com/v1beta1
2 kind: authpolicy
3 metadata:
4 name: authhotdrinks
5 spec:
6
       servicenames:
7
       - frontend
8
       authentication_mechanism:
9
10
         using_forms:
           authentication_host: "fqdn_authenticaton_host"
11
12
           authentication_host_cert:
13
             tls_secret: authhost-tls-cert-secret
           vip: "192.2.156.156"
14
15
16
       authentication_providers:
           - name: "ldap-auth-provider"
17
18
             ldap:
                 server_ip: "192.2.156.160"
19
                 base: 'dc=aaa,dc=local'
20
21
                 login_name: accountname
                 sub_attribute_name: CN
23
                 server_login_credentials: ldapcredential
24
                 attributes_to_save: memberof #memberof attribute to be
                     obtained from LDAP server for user
25
```

```
26 authentication_policies:
27
28 - resource:
29 path: []
30 method: []
31 provider: ["ldap-auth-provider"]
```

### **Exporting metrics directly to Prometheus**

#### May 16, 2024

NetScaler Ingress Controller now supports exporting metrics directly from NetScaler to Prometheus. With NetScaler Ingress Controller, you can automate the configurations required on NetScaler for exporting metrics directly.

Once you export the metrics, you can visualize the exported NetScaler metrics for easier interpretation and understanding using tools such as Grafana.

#### Configuring direct export of metrics from NetScaler CPX to Prometheus

To enable NetScaler Ingress Controller to configure NetScaler CPX to support direct export of metrics to Prometheus, you need to perform the following steps:

1. Create a Kubernetes secret to enable read-only access for a user. This step is required for NetScaler CPX to export metrics to Prometheus.

```
1 kubectl create secret generic prom-user --from-literal=username=<
    prometheus-username> --from-literal=password=<prometheus-
    password>
```

2. Deploy NetScaler CPX with NetScaler Ingress Controller using the following Helm commands:

```
1 helm repo add netscaler https://netscaler.github.io/netscaler-helm
        -charts/
2
3 helm install my-release netscaler/netscaler-cpx-with-ingress-
        controller --set license.accept=yes,nsic.
        prometheusCredentialSecret=<Secret-for-read-only-user-creation
        >,analyticsConfig.required=true,analyticsConfig.timeseries.
        metrics.enable=true,analyticsConfig.timeseries.port=5563,
        analyticsConfig.timeseries.metrics.mode=prometheus,
        analyticsConfig.timeseries.metrics.enableNativeScrape=true
```

The new parameters specified in the command are explained as follows:

- nsic.prometheusCredentialSecret: Specifies the Kubernetes secret name for creating the read only user for native Prometheus support.
- analyticsConfig.timeseries.metrics.enableNativeScrape: Set this value to **true** for directly exporting metrics to Prometheus
- 3. Add the appropriate Prometheus scrape job under scrape\_configs in the Prometheus configuration depending on the Prometheus deployment.
  - If your Prometheus server is outside the Kubernetes cluster, add a scrape job under scrape\_configs in the Prometheus configuration. For a sample Prometheus scrape job, see the Prometheus integration documentation.
  - If your Prometheus server is within the same Kubernetes cluster, add a new Prometheus job to configure Prometheus for directly exporting from a NetScaler CPX pod. For more information, see kubernetes\_sd\_config. A sample Prometheus job is given as follows:

```
- job_name: 'kubernetes-cpx'
2
     scheme: http
3
     metrics_path: /nitro/v1/config/systemfile
4
     params:
5
       args: ['filename:metrics_prom_ns_analytics_time_series_profile
           .log,filelocation:/var/nslog']
6
       format: ['prometheus']
7
     basic_auth:
8
                  # Prometheus username set in nsic.
       username:
           prometheusCredentialSecret
9
       password: # Prometheus password set in nsic.
           prometheusCredentialSecret
10
     scrape_interval: 30s
     kubernetes_sd_configs:
11
12
     - role: pod
13
     relabel_configs:
14
     - source_labels: [
        __meta_kubernetes_pod_annotation_netscaler_prometheus_scrape]
       action: keep
15
16
       regex: true
17
     - source_labels: [__address__,
         __meta_kubernetes_pod_annotation_netscaler_prometheus_port]
18
       action: replace
       regex: ([^:]+)(?::\d+)?;(\d+)
19
       replacement: $1:$2
       target_label: __address__
21
22
     - source_labels: [__meta_kubernetes_namespace]
23
      action: replace
24
       target_label: kubernetes_namespace
     - source_labels: [__meta_kubernetes_pod_name]
25
26
       action: replace
       target_label: kubernetes_pod_name
27
```

#### Note:

For more information on Prometheus integration, see the NetScaler Prometheus integration documentation.

# Configuring direct export of metrics from NetScaler VPX or NetScaler MPX to Prometheus

To enable NetScaler Ingress Controller to configure NetScaler VPX or NetScaler MPX to support direct export of metrics to Prometheus, you need to perform the following steps:

1. Deploy NetScaler Ingress Controller as a stand-alone pod using the Helm command:

- 2. Create a system user with read only access for NetScaler VPX. For more details on the user creation, see the NetScaler Prometheus integration documentation.
- 3. Add a scrape job under scrape\_configs in the prometheus configuration for enabling Prometheus to scrape from NetScaler VPX. For a sample Prometheus scrape job, see Prometheus configuration.

#### Note:

The scrape configuration section specifies a set of targets and configuration parameters describing how to scrape them. For more information on NetScaler specific parameters used in the configuration, see the NetScaler documentation.

# Configure static route on NetScaler VPX or MPX

#### December 13, 2024

In a Kubernetes cluster, pods run on an overlay network such as Flannel, Calico, or Weave. The pods in the cluster are assigned an IP address from the overlay network, which is different from the host network. NetScaler MPX or NetScaler VPX outside the Kubernetes cluster receives all the ingress traffic to the microservices deployed in the Kubernetes cluster. You need to establish network connectivity between NetScaler instance and the pods for the ingress traffic to reach the microservices.

One of the ways to achieve network connectivity between pods and NetScaler MPX or NetScaler VPX instance outside the Kubernetes cluster is to configure routes on NetScaler instance to the overlay network. You can either do this manually or NetScaler Ingress Controller provides an option to automatically configure the network.

#### Note

Ensure that NetScaler MPX or NetScaler VPX has SNIP configured on the host network. The host network is the network on which the Kubernetes nodes communicate with each other.

#### Manually configure a route on NetScaler

1. On the master node in the Kubernetes cluster, get the podCIDR using the following command:

```
kubectl get nodes -o jsonpath="{ range .items[*] } { 'podNetwork
: ' } { .spec.podCIDR } { '\t' } { 'gateway: ' } { .status.
addresses[0].address } { '\n' } { end } "
```

1	<pre>podNetwork:</pre>	10.244.0.0/24	gateway:	10.106.162.108
2	<pre>podNetwork:</pre>	10.244.2.0/24	gateway:	10.106.162.109
3	<pre>podNetwork:</pre>	10.244.1.0/24	gateway:	10.106.162.106

- 2. Log on to NetScaler instance.
- 3. Add a route on NetScaler using the podCIDR information. Use the following command:

add route <pod\_network> <podCIDR\_netmask> <gateway>
Examples:
add route 192.244.0.0 255.255.255.0 192.106.162.108
add route 192.244.2.0 255.255.255.0 192.106.162.109
add route 192.244.1.0 255.255.255.0 192.106.162.106

#### Automatically configure a route on NetScaler

In the NSIC deployment, you can use the parameter nodeWatch to automatically configure a route on the associated NetScaler instance.

Specify the value of nodeWatch as **true** to enable automatic route configuration. For example:

#### 1 helm install my-release netscaler/netscaler-ingress-controller --set nsIP=<NSIP>,license.accept=yes,adcCredentialSecret=<Secret-for-NetScaler-credentials>,nodeWatch=true

#### Note:

By default, the nodeWatch argument is set to **false**; set the argument to **true** to enable the automatic route configuration.

# Establish network between Kubernetes nodes and Ingress NetScaler using node controller

December 31, 2023

In Kubernetes environments, when you expose the services for external access through the Ingress device you need to appropriately configure the network between the Kubernetes nodes and the Ingress device.

Configuring the network is challenging as the pods use private IP addresses based on the CNI framework. Without proper network configuration, the Ingress device cannot access these private IP addresses. Also, manually configuring the network to ensure such reachability is cumbersome in Kubernetes environments.

Also, if the Kubernetes cluster and the Ingress NetScaler are in different subnets, you cannot establish a route between them using <u>Static routing</u>. This scenario requires an overlay mechanism to establish a route between the Kubernetes cluster and the Ingress NetScaler.

NetScaler provides a node controller that you can use to create a VXLAN based overlay network between the Kubernetes nodes and the Ingress NetScaler as shown in the following diagram:



#### Note:

NetScaler Node Controller does not work in a setup where a NetScaler cluster is configured as an ingress device. NetScaler Node Controller requires to establish a Virtual Extensible LAN (VXLAN) tunnel between NetScaler and Kubernetes nodes to configure routes and creating a VXLAN on a NetScaler cluster is not supported.

#### To establish network connectivity using node controller:

- 1. Deploy the NetScaler Ingress Controller. Perform the following steps:
  - a) Download the citrix-k8s-ingress-controller.yaml using the following command:

```
1 wget https://raw.githubusercontent.com/citrix/citrix-k8s-
ingress-controller/master/deployment/baremetal/citrix-k8s-
ingress-controller.yaml
```

- b) Edit the citrix-k8s-ingress-controller.yaml file and enter the values for the environmental variables. For more information, see Deploy the NetScaler Ingress Controller.
- c) Once you update the environment variables, save the YAML file and deploy it using the following command:

```
1 kubectl create -f citrix-k8s-ingress-controller.yaml
```

d) Verify if the NetScaler Ingress Controller is deployed successfully using the following command:

```
1 kubectl get pods --all-namespaces
```

2. Deploy the node controller. For information on how to deploy the node controller, see Deploy the Citrix k8s node controller.

# **Expose Service of type NodePort using Ingress**

#### December 31, 2023

In a single-tier deployment, the Ingress NetScaler (VPX or MPX) outside the Kubernetes cluster receives all the Ingress traffic to the microservices deployed in the Kubernetes cluster. For the Ingress traffic to reach the microservices, you need to establish network connectivity between the Ingress NetScaler instance and pods.

As pods run on overlay network, the pod IP addresses are private IP addresses and the Ingress NetScaler instance cannot reach the microservices running within the pods. To make the service accessible from outside of the cluster, you can create the service of type NodePort. The NetScaler instance load balances the Ingress traffic to the nodes that contain the pods.

To create the service of type NodePort, in your service definition file, specify spec.type:NodePort and optionally specify a port in the range 30000–32767.

#### Sample deployment

Consider a scenario wherein you are using a NodePort based service, for example, an apache app and want to expose the app to North-South traffic using an Ingress. In this case, you need to create the apache app deployment, define the service of type NodePort, and create an Ingress definition to configure Ingress NetScaler to send the North-South traffic to the nodeport of the apache app.

In this example, you create a deployment named apache, and deploy it in your Kubernetes cluster.

1. Create a manifest for the deployment named apache-deployment.yaml.

1	#	If using <b>this</b> on GKE
2	#	Make sure you have cluster-admin role <b>for</b> your account
3	#	kubectl create clusterrolebinding citrix-cluster-admin
		clusterrole=cluster-adminuser= <username google<="" of="" th="" your=""></username>
		account>
4	#	
5		

```
6 #For illustration a basic apache web server is used as a
      application
7 apiVersion: apps/v1
8 kind: Deployment
9 metadata:
10 name: apache
11
     labels:
12
        name: apache
13 spec:
14 selector:
15
     matchLabels:
16
        app: apache
17
   replicas: 4
18
   template:
19
     metadata:
20
       labels:
21
          app: apache
22
     spec:
23
       containers:
24
         - name: apache
          image: httpd:latest
25
26
           ports:
27
           - name: http
28
             containerPort: 80
29
           imagePullPolicy: IfNotPresent
```

Containers in this deployment listen on port 80.

2. Create the deployment using the following command:

```
1 kubectl create -f apache-deployment.yaml
```

3. Verify that four pods are running using the following:

1 kubectl get pods

4. Once you verify that pods are up and running, create a service of type NodePort. The following is a manifest for the service:

```
1 #Expose the apache web server as a Service
2 apiVersion: v1
3 kind: Service
4 metadata:
5
   name: apache
6
   labels:
7
      name: apache
8 spec:
9
    type: NodePort
10
   ports:
    - name: http
11
     port: 80
12
      targetPort: http
13
14 selector:
```

15 app: apache

5. Copy the manifest to a file named apache-service.yaml and create the service using the following command:

1 kubectl create -f apache-service.yaml

The sample deploys and exposes the Apache web server as a service. You can access the service using the <NodeIP>:<NodePort> address.

6. After you have deployed the service, create an Ingress resource to configure the Ingress NetScaler to send the North-South traffic to the nodeport of the apache app. The following is a manifest for the Ingress definition named as vpx-ingress.yaml.

```
apiVersion: networking.k8s.io/v1
1
2 kind: Ingress
3 metadata:
4
    annotations:
5
       ingress.citrix.com/frontend-ip: xx.xxx.xx
6
    name: vpx-ingress
7 spec:
8
    defaultBackend:
9
     service:
10
        name: apache
11
         port:
12
           number: 80
```

7. Deploy the Ingress object.

```
1 kubectl create -f vpx-ingress.yaml
```

### Configure pod to pod communication using Calico

December 31, 2023

Configuring a network in Kubernetes is a challenge. It requires you to deal with many nodes and pods in a cluster system. There are four problems you need to address while configuring the network:

- · Container to container (which collectively provides a service) communication
- Pod to pod communication
- Pod to service communication
- External to service communication

#### Pod to pod communication

By default, docker creates a virtual bridge called docker0 on the host machine and it assigns a private network range to it. For each container that is created, a virtual Ethernet device is attached to this bridge. The virtual Ethernet device is then mapped to eth0 inside the container, with an IP from the network range. This process happens for each host that is running docker. There is no coordination between these hosts therefore the network ranges might collide.

Because of this, containers can only communicate with containers that are connected to the same virtual bridge. To communicate with other containers on other hosts, they must rely on port mapping. That means, you need to assign a port on the host machine to each container and then forward all the traffic on that port to that container.

Since the local IP address of the application is translated to the host IP address and port on the host machine, Kubernetes assumes that all nodes can communicate with each other without NAT. It also assumes that the IP address that a container sees for itself is the same IP address that the other containers see for the container. This approach also enables you to port applications easily from virtual machines to containers.

Calico is one of the many different networking options that offer these capabilities for Kubernetes.

#### Calico

Calico is designed to simplify, scale, and secure cloud networks. The open source framework enables Kubernetes networking and network policy for clusters across the cloud. Within the Kubernetes ecosystem, Calico is starting to emerge as one of the most popularly used network frameworks or plug-ins, with many enterprises using it at scale.

Calico uses a pure IP networking fabric to deliver high performance Kubernetes networking, and its policy engine enforces developer intent for high-level network policy management. Calico provides Layer 3 networking capabilities and associates a virtual router with each node. It enables host to host and pod to pod networking. Calico allows establishment of zone boundaries through BGP or encapsulation through IP on IP or VXLAN methods.

#### Integration between Kubernetes and Calico

Calico integrates with Kubernetes through a CNI plug-in built on a fully distributed, layer 3 architecture. Hence, it scales smoothly from a single laptop to large enterprise. It relies on an IP layer and it is relatively easy to debug with existing tools.

#### **Configure the network with Calico**

First, bring up a Kubernetes cluster with Calico using the following commands:

```
1 > kubeadm init --pod-network-cidr=192.168.0.0/16
2 > export KUBECONFIG=/etc/kubernetes/admin.conf
3 > kubectl apply -f calico.yaml
```

A master node is created with Calico as the CNI. After the master node is up and running, you can join the other nodes to the master using the join command.

Calico processes that are part of the Kubernetes master node are:

Calico etcd

kube-system calico-etcd-j4rwc 1/1 Running

Calico controller

kube-system calico-kube-controllers-679568f47c-vz69g 1/1 Running

Calico nodes

kube-system calico-node-ct6c9 2/2 Running

Note:

When you join a node to the Kubernetes cluster, a new *Calico node* is initiated on the Kubernetes node.

#### **Configure BGP peer with Ingress NetScaler**

Whenever you deploy an application after establishing the Calico network in the cluster, Kubernetes assigns an IP address from the IP address pool of Calico to the service associated with the application.

Border Gateway Protocol (BGP) uses autonomous system number (AS number) to identify the remote nodes. The AS number is a special number assigned by IANA used primarily with BGP to identify a network under a single network administration that uses unique routing policy.

**Configure BGP on Kubernetes using Ingress NetScaler** Using a YAML file, you can apply BGP configuration of a remote node using the kubectl create command. In the YAML file, you need to add the peer IP address and the AS number. The peer IP address is the Ingress NetScaler IP address and the AS number that is used in the Ingress NetScaler.

**Obtain the AS Number of the cluster** Using the calicoctl command, you can obtain the AS number that is used by Calico BGP in the Kubernetes cluster as shown in the following image:

root@ubuntu194:~/kubeCluster# ETCD\_ENDPOINTS=http://10.102.33.194:6666 ./calicoctl.1 get bgpConfiguration
NAME LOGSEVERITY MESHENABLED ASNUMBER
default Info false 64512

**Configure global BGP peer** Using the calicoctl utility, you can peer Calico nodes with global BGP speakers. This kind of peers is called global peers.

Create a YAML definition file called bgp.yml with the following definition:

```
1 apiVersion: projectcalico.org/v3 # This is the version of Calico
2 kind: BGPPeer # BGPPeer specifies that its Global peering.
3 metadata:
4 name: bgppeer-global-3040 # The name of the configuration
5 spec:
6 peerIP: 10.102.33.208 # IP address of the Ingress NetScaler
7 asNumber: 500 # AS number configured on the Ingress NetScaler
```

Deploy the definition file using the following command:

1 > kubectl create -f bgp.yml

Add the BGP configurations on the Ingress NetScaler Perform the following:

- 1. Log on to the NetScaler command-line interface.
- 2. Enable the BGP feature using the following command:

```
1 > en feature bgp
2 Done
```

3. Type vtysh and press Enter.

```
1 > vtysh
2 ns#
```

4. Change to config terminal using the conf t command:

```
    ns#conf t
    Enter configuration commands, one per line. End with CNTL/Z.
    ns(config)#
```

5. Add the BGP route with the AS number as 500 for demonstration purpose. You can use any number as the AS number.

```
1 ns(config)# router bgp 500
2 ns(config-router)#
```

6. Add neighbors using the following command:

```
1 ns(config-router)# Neighbor 10.102.33.198 remote-as 64512
2 ns(config-router)# Neighbor 10.102.22.202 remote-as 64512
```

7. Review the running configuration using the following command:

```
1 ns(config-router)#show running-config
2 !
3 log syslog
4 log record-priority
5
6 ns route-install bgp
7
  1
8 interface lo0
9
   ip adress 127.0.0.1/8
10 ipv6 address fe80: :1/64
11
  ipv6 address : :1/128
12
13 interface vlan0
   ip address 10.102.33.208/24
14
   ipv6 address fe80::2cf6:beff:fe94:9f63/64
15
16
17 router bgp 500
18 max-paths ebgp 8
19
   max-paths ibgp 8
   neighbor 10.102.33.198 remote-as 64512
20
21
   neighbor 10.102.33.202 remote-as 64512
22 !
23 end
24 ns(config-router)# In the sample, the AS number of Calico is
      64512, you can change this number as per your requirement.
```

8. Install the BGP routes to NetScaler routing table using the following command:

```
1 ns(config)# ns route-install bgp
2 ns(config)#
3 exit
4 ns#exit
5 Done
```

9. Verify the route and add to the routing table using the following command:

>sh r	oute					
	Network	Netmask	Gateway/OwnedIP	State	Traffic Domain	Туре
1)	0.0.0.0	0.0.0.0	10.102.33.1	UP	0	STATIC
2)	127.0.0.0	255.0.0.0	127.0.0.1	UP	0	PERMANENT
3)	10.102.33.0	255.255.255.0	10.102.33.208	UP	0	DIRECT
4)	192.168.1.0	255.255.255.0	10.102.33.198	UP	0	BGP
5)	192.168.43.128	255.255.255.192	10.102.33.198	UP	0	BGP
6)	192.168.71.64	255.255.255.192	10.102.33.202	UP	0	BGP
Done						
_						

Once the route is installed, the NetScaler is able to communicate with services that are present in the Kubernetes cluster:

#### NetScaler ingress controller



#### Troubleshooting

You can verify BGP configurations on the master node in the Kubernetes cluster using the calicoctl script.

#### View the peer IP address and AS number configurations

You can view the peer IP address and AS number configurations using the following command:

```
1 >./calicoctl.1 get bgpPeer
2 NAME PEERIP NODE ASN
3 bgppeer-global-3040 10.102.33.208 (global) 500
```

#### View the BGP node status

You can view the status of a BGP node using the following command:

# Enhancements for Kubernetes service of type LoadBalancer support in the NetScaler Ingress Controller

#### February 8, 2024

Kubernetes service of type LoadBalancer support in the NetScaler Ingress Controller is enhanced with the following features:

- BGP route health injection (RHI) support
- Advertise or recall load balancer IP addresses (VIPs) based on the availability of service's pods in a set of nodes (zones) defined by node's labels

#### Support for automatic configuration of BGP RHI on NetScaler

Route health injection (RHI) allows the NetScaler to advertise the availability of a VIP as a host route throughout the network using BGP. However, you had to manually perform the configuration on NetScaler to support RHI. Using NetScaler Ingress Controllers deployed in a Kubernetes environment, you can automate the configuration on NetScalers to advertise VIPs.

When a service of type LoadBalancer is created, the NetScaler Ingress Controller configures a VIP on the NetScaler for the service. If BGP RHI support is enabled for the NetScaler Ingress Controller, it automatically configures NetScaler to advertise the VIP to the BGP network. Using the service .citrix.com/vipparams annotation, you can enable IP parameters for the VIP. For example, see the service.YAML file in the step 5 of Configuring BGP RHI on NetScalers using the NetScaler Ingress Controller. For information on the supported IP parameters, see nsip configuration.

#### Advertise and recall VIPs based on the availability of pods

In the topology as shown in the following diagram, nodes in a Kubernetes cluster are physically distributed across three different racks. They are logically grouped into three zones. Each zone has a NetScaler MPX as the Tier-1 ADC and a NetScaler Ingress Controller on the same in the Kubernetes cluster. NetScaler Ingress Controllers in all zones listen to the same Kubernetes API server. So, whenever a service of type LoadBalancer is created, all NetScalers in the cluster advertises the same IP address to the BGP fabric. Even, if there is no workload on a zone, the NetScaler in that zone still advertises the IP address.



NetScaler provides a solution to advertise or recall the VIP based on the availability of pods in a zone. You need to label the nodes on each zone so that the NetScaler Ingress Controller can identify nodes belonging to the same zone. The NetScaler Ingress Controller on each zone performs a check to see if there are pods on nodes in the zone. If there are pods on nodes in the zone, it advertises the VIP. Otherwise, it revokes the advertisement of VIP from the NetScaler on the zone.

### Configuring BGP RHI on NetScalers using the NetScaler Ingress Controller

This topic provides information on how to configure BGP RHI on NetScalers using the NetScaler Ingress Controller based on a sample topology. In this topology, nodes in a Kubernetes cluster are deployed across two zones. Each zone has a NetScaler VPX or MPX as the Tier-1 ADC and a NetScaler Ingress Controller for configuring ADC in the Kubernetes cluster. The ADCs are peered using BGP with the upstream router.



#### Prerequisites

• Configure NetScaler MPX or VPX as a BGP peer with the upstream routers.

Perform the following steps to configure BGP RHI support based on the sample topology.

1. Label nodes in each zone using the following command:

For zone 1:

```
1 kubectl label nodes node1 rack=rack-1
2 kubectl label nodes node2 rack=rack-1
```

For zone 2:

```
1 kubectl label nodes node3 rack=rack-2
2 kubectl label nodes node4 rack=rack-2
```

2. Configure the following environmental variables in the NetScaler Ingress Controller configuration YAML files as follows:

For zone 1:

```
1 - name: "NODE_LABELS"
2 value: "rack-1"
3 - name: "BGP_ADVERTISEMENT"
4 value: "True"
```

For zone 2:

```
1 - name: "NODE_LABELS"
2 value: "rack-2"
3 - name: "BGP_ADVERTISEMENT"
4 value: "True"
```

A sample cic.yaml file for deploying the NetScaler Ingress Controller on zone 1 is provided as follows:

```
1 apiVersion: v1
2 kind: Pod
3 metadata:
4
    name: cic-k8s-ingress-controller-1
5
     labels:
       app: cic-k8s-ingress-controller-1
6
7 spec:
     serviceAccountName: cic-k8s-role
8
9
     containers:
     - name: cic-k8s-ingress-controller
10
11
       image: "quay.io/citrix/citrix-k8s-ingress-controller:1.36.5"
12
13
       env:
14
      # Set NetScaler NSIP/SNIP, SNIP in case of HA (mgmt has to be
          enabled)
       - name: "NS_IP"
15
         value: "10.217.212.24"
17
      # Set username for Nitro
18
       - name: "NS_USER"
19
         valueFrom:
           secretKeyRef:
             name: nslogin
22
             key: username
23
      # Set user password for Nitro
      - name: "NS_PASSWORD"
24
25
         valueFrom:
26
           secretKeyRef:
27
           name: nslogin
28
           key: password
29
      - name: "EULA"
        value: "yes"
       - name: "NODE_LABELS"
31
         value: "rack=rack-1"
32
33
       - name: "BGP_ADVERTISEMENT"
         value: "True"
34
35 args:
    - --ipam
37
       citrix-ipam-controller
38 imagePullPolicy: Always
```

3. Deploy the NetScaler Ingress Controller using the following command.

#### Note:

You need to deploy the NetScaler Ingress Controller on both racks (per zone).

1 Kubectl create -f cic.yaml

4. Deploy a sample application using the web-frontend-lb.yaml file.

Kubectl create -f web-frontend-lb.yaml

The content of the web-frontend-lb.yaml is as follows:

```
apiVersion: v1
1
2 kind: Deployment
3 metadata:
    name: web-frontend
4
5 spec:
6
    selector:
7
      matchLabels:
8
        app: web-frontend
9
   replicas: 4
10 template:
11
      metadata:
12
         labels:
13
           app: web-frontend
     spec:
14
15
        containers:
16
         - name: web-frontend
17
           image: 10.217.6.101:5000/web-test:latest
18
           ports:
19
             - containerPort: 80
20
           imagePullPolicy: Always
```

5. Create a service of type LoadBalancer for exposing the application.

1 Kubectl create -f web-frontend-lb-service.yaml

The content of the web-frontend-lb-service.yaml is as follows:

```
1 apiVersion: v1
2 kind: Service
3 metadata:
    name: web-frontend
4
5
    annotations:
      service.citrix.com/class: 'cic-vpx'
6
7
       service.citrix.com/frontend-ip: 1.1.1.1
8
      service.citrix.com/vipparams: '{
9
   "vserverrhilevel": "ONE_VSERVER", "hostroute": "ENABLED", "metric
       ": 10 }
    1
10
    labels:
11
       app: web-frontend
13 spec:
```

```
14 type: LoadBalancer
15 ports:
16 - port: 80
17 protocol: TCP
18 name: http
19 selector:
20 app: web-frontend
```

6. Verify the service group creation on NetScalers using the following command.

```
1 show servicegroup <service-group-name>
```

Following is a sample output for the command.

```
1 # show servicegroup k8s-web-frontend_default_80_svc_k8s-web-
      frontend_default_80_svc
2
   k8s-web-frontend_default_80_svc_k8s-web-frontend_default_80_svc -
3
      TCP
4 State: ENABLED Effective State: UP Monitor Threshold : 0
5 Max Conn: 0 Max Req: 0 Max Bandwidth: 0 kbits
6 Use Source IP: NO
7 Client Keepalive(CKA): NO
8 TCP Buffering(TCPB): NO
9 HTTP Compression(CMP): NO
10 Idle timeout: Client: 9000 sec Server: 9000 sec
11 Client IP: DISABLED
12 Cacheable: NO
13 SC: OFF
14 SP: OFF
15 Down state flush: ENABLED
16 Monitor Connection Close : NONE
17 Appflow logging: ENABLED
18 ContentInspection profile name: ???
19 Process Local: DISABLED
20 Traffic Domain: 0
21
22
23 1)
        10.217.212.23:30126 State: UP Server Name: 10.217.212.23
      Server ID: None Weight: 1
24
     Last state change was at Wed Jan 22 23:35:11 2020
     Time since last state change: 5 days, 00:45:09.760
26
27
     Monitor Name: tcp-default State: UP Passive: 0
28
     Probes: 86941 Failed [Total: 0 Current: 0]
29
     Last response: Success - TCP syn+ack received.
30
     Response Time: 0 millisec
31
32 2)
        10.217.212.22:30126 State: UP Server Name: 10.217.212.22
      Server ID: None Weight: 1
     Last state change was at Wed Jan 22 23:35:11 2020
34
     Time since last state change: 5 days, 00:45:09.790
35
```

```
Monitor Name: tcp-default State: UP Passive: 0
Probes: 86941 Failed [Total: 0 Current: 0]
Last response: Success - TCP syn+ack received.
```

7. Verify the VIP advertisement on the BGP router using the following command.

# **TLS certificates handling in NetScaler Ingress Controller**

#### April 11, 2024

NetScaler Ingress Controller provides option to configure TLS certificates for NetScaler SSL-based virtual servers. The SSL virtual server intercepts SSL traffic, decrypts it and processes it before sending it to services that are bound to the virtual server.

By default, SSL virtual server can bind to one default certificate and the application receives the traffic based on the policy bound to the certificate. However, you have the Server Name Indication (SNI) option to bind multiple certificates to a single virtual server. NetScaler determines which certificate to present to the client based on the domain name in the TLS handshake.

NetScaler Ingress Controller handles the certificates in the following three ways:

- NetScaler Ingress Controller default Certificate
- Preconfigured certificates
- TLS section in the Ingress YAML

#### Prerequisite

For handling TLS certificates using NetScaler Ingress Controller, you need to enable TLS support in NetScaler for the application and also if you are using certificates in your Kubernetes deployment then you need to generate Kubernetes secret using the certificate.

#### Enable TLS support in NetScaler for the application

NetScaler Ingress Controller uses the **TLS** section in the ingress definition as an enabler for TLS support with NetScaler.

#### Note:

If there is a default certificate or if there are preconfigured certificates, you need to add an empty secret in the **spec.tls.secretname** field in your ingress definition to enable TLS.

#### The following sample snippet of the ingress definition:

```
1 spec:
2 tls:
3 - secretName:
```

#### Generate Kubernetes secret

To generate Kubernetes secret for an existing certificate, use the following kubectl command:

The command creates a Kubernetes secret with a PEM formatted certificate under tls.crt key and a PEM formatted private key under tls.key key.

Alternatively, you can also generate the Kubernetes secret using the following YAML definition:

```
1 apiVersion: v1
2 kind: Secret
3 metadata:
4 name: k8s-secret
5 data:
6 tls.crt: base64 encoded cert
7 tls.key: base64 encoded key
```

Deploy the YAML using the kubectl -create <file-name> command. It creates a Kubernetes secret with a PEM formatted certificate under tls.crt key and a PEM formatted private key under tls.key key.

#### NetScaler Ingress Controller default certificate

The default secrets provided in NetScaler Ingress Controller can be used to configure SSL and SSL SNI certificates in NetScaler.

You can use **default**-sni-certificate and **default**-ssl-sni-certificate arguments to provide a secret to configure non-SNI and SNI certificates respectively. When you specify the arguments in the NetScaler Ingress Controller deployment YAML file, provide the secret name and the namespace where the secret has been deployed in the cluster as following:

- Argument to add in the YAML file to use the default certificate as a non-SNI certificate:
   --default-ssl-certificate <NAMESPACE>/<SECRET\_NAME>.
- Argument to add in the YAML file to use the default certificate as an SNI certificate: --default
   -ssl-sni-certificate <NAMESPACE>/<SECRET\_NAME>

#### Note:

If you deploy NetScaler Ingress Controller using a Helm chart or operators, both the default secrets must be deployed in the same namespace where NetScaler Ingress Controller is being deployed. Provide the non-SNI secret name in the defaultSSLCertSecret parameter and the SNI secret name in the defaultSSLSNICertSecret parameter during the NetScaler Ingress Controller installation. For example, defaultSSLCertSecret: <non-SNI secret name>; defaultSSLSNICertSecret: <SNI secret name>.

The following is a sample NetScaler Ingress Controller YAML definition file that contains a TLS secret (hotdrink.secret) picked from the ssl namespace and provided as the NetScaler Ingress Controller default certificate.

#### Note:

Namespace is mandatory along with a valid SECRET\_NAME.

```
1 apiVersion: v1
2 kind: Pod
3 metadata:
4
   name: cic
    labels:
5
6
       app: cic
7 spec:
8
   serviceAccountName: cpx
9
   containers:
10
     - name: cic
11
      image: "xxxx"
12
      imagePullPolicy: Always
13
     args:
14
      - -- default-ssl-certificate
       ssl/hotdrink.secret
15
16
     env:
17
      # Set NetScaler ADM Management IP
      - name: "NS_IP"
18
       value: "xx.xx.xx.xx"
19
20
      # Set port for Nitro
      - name: "NS_PORT"
21
       value: "xx"
22
23
       # Set Protocol for Nitro
24
       - name: "NS_PROTOCOL"
       value: "HTTP"
25
26
       # Set username for Nitro
```

```
27 - name: "NS_USER"
28 value: "nsroot"
29 # Set user password for Nitro
30 - name: "NS_PASSWORD"
31 value: "nsroot"
```

For information about the behaviour of NetScaler Ingress Controller in different scenarios related to Kubernetes ingress in Kubernetes cluster, see the following table.

Default SSL	Default SSL SNI				
Secret in	Secret in				
NetScaler Ingress	NetScaler Ingress		Co cust in la succe	A	
Controller	Controller	Host in ingress	Secret in ingress	Actions	
Yes	No	Not provided	Not provided	Bind non-SNI	
				secret as a	
				non-SNI	
				certificate in SSL	
				virtual server.	
No	Yes	Not provided	Not Provided	Bind default SNI	
				secret as an SNI	
				certificate in SSL	
Vec	Vez	Not Drowided	Net www.ided	virtual server. • Bind	
Yes	Yes	Not Provided	Not provided	non-SNI	
				secret as a	
				non-SNI	
Yes	No	Provided	Not provided	Bind nontificate.	
				secretBind SNI	
				non-SAPcret as an	
				certifiente in SSL	
				virtualesetificate.	
No	Yes	Provided	Not provided	Bind SANsoebietlas	
				an SN <b>hoentSNd</b> ate	
				in SSLseictreablas	
				serverSNI since	
Yes	Yes	Provided	Not provided	• Rillfliple secrets are secrets are secrets as a B8DrSNI certificate	
				in SSL	
				virtual	
				server.	
@ 1007, 2025 Citain Customer Inc. All rights recorded					
© 1997–2025 Citrix	systems, Inc. All righ	nts reserved.		425 secret as an	
				SNI	
				certificate	
				in SSL	

Default SSL	Default SSL SNI				
Secret in	Secret in				
NetScaler Ingress	NetScaler Ingress				
Controller	Controller	Host in Ingress	Secret in Ingress	Actions	
Anything	Anything	Not provided	Provided	Bind secret provided as a non-SNI certificate in SSL virtual server.	
Anything	Anything	Provided	Provided	Bind secret provided as an SNI certificate in SSL virtual server.	
Anything	Anything	Provided/Not provided	Multiple	Bind all the secrets provided as SNI certificates in SSL virtual server.	

For information about behaviour of NetScaler Ingress Controller in different scenarios related to Open-Shift route in OpenShift cluster, see the following table.

Default SSL Secret in	Default SSL SNI Secret in			
NetScaler Ingress	NetScaler Ingress		Key and Cert in	
Controller	Controller	Route type	Route	Actions
Yes	No	Edge	Not Provided	Bind non-SNI secret as a non-SNI certificate in SSL virtual server.
No	Yes	Edge	Not Provided	Bind default SNI secret as an SNI certificate in SSL virtual server.
Yes	Yes	Edge	Not Provided	<ul> <li>Bind non-SNI secret as a non-SNI certificate.</li> <li>Bind SNI</li> </ul>
© 1997–2025 Citrix	Systems, Inc. All rigi	nts reserved.		secret asła6 SNI certificate. • Also, bind

Default SSL Secret in NetScaler Ingress Controller	Default SSL SNI Secret in NetScaler Ingress Controller	Route type	Key and Cert in Route	Actions
Yes	No	Reencrypt	Not Provided	Bind non-SNI secret as a non-SNI certificate in SSL
No	Yes	Reencrypt	Not provided	virtual server. Bind SNI secret as a SNI certificate in
Yes	Yes	Reencrypt	Not provided	SSL virtual server. • Bind non-SNI secret as a non-SNI
Yes	No	Passthrough	Not provided	Bind nontsincate secret <sup>i</sup> as Sal SNI
No	Yes	Passthrough	Not provided	Bind Stervecret as an SN BioctiSitate in SSLscicteates an
Yes	Yes	Passthrough	Not provided	serverSNI • Bindficate. • AGD, SNId
Anything	Anything	Edge	Provided	SNI certificate in Multiple
Anything	Anything	Reencrypt	Provided	SSL Virtual server. Secrets are getting provided as an Bound m SNI certificate in SSL Virtual SSL Virtual
Anything	Anything	Passthrough	OpenShift doesn' t allow	NA server. NA certificate in SSL virtual
				server.

#### **Preconfigured certificates**

NetScaler Ingress Controller allows you to use the certkeys that are already configured on the NetScaler. You must provide the details about the certificate using the following annotation in your ingress definition:

```
1 ingress.citrix.com/preconfigured-certkey : '{
2 \"certs\": \[ {
3 \"name\": \"<name>\", \"type\": \"default|sni|ca\" }
4 ] }
5 '
```

You can provide details about multiple certificates as a list within the annotation. Also, you can define the way the certificate is treated. In the following sample annotation, certkey1 is used as a non-SNI certificate and certkey2 is used as an SNI certificate:

```
ingress.citrix.com/preconfigured-certkey : '{
    "certs": [ {
    "name": "certkey1", "type": "default" }
    , {
    "name": "certkey2", "type": "sni" }
    ] }
    7
```

If the type parameter is not provided with the name of a certificate, then it is considered as the default (non-SNI) type.

Note:

Ensure that you use this feature in cases where you want to reuse the certificates that are present on the NetScaler and bind them to the applications that are managed by NetScaler Ingress Controller. NetScaler Ingress Controller does not manage the life cycle of the certificates. That is, it does not create or delete the certificates, but only binds them to the necessary applications.

#### **TLS section in the ingress YAML**

Kubernetes allows you to provide the TLS secrets in the spec: section of an ingress definition. This section describes how the NetScaler Ingress Controller uses these secrets.

#### With the host section

If the secret name is provided with the host section, NetScaler Ingress Controller binds the secret as an SNI certificate.

1 spec: 2 tls:

```
3 - secretName: fruitjuice.secret
4 hosts:
5 - items.fruit.juice
```

#### Without the host section

If the secret name is provided without the host section, NetScaler Ingress Controller binds the secret as a default certificate.

```
1 spec:
2 tls:
3 - secretName: colddrink.secret
```

#### Note:

If there are more than one secret given then NetScaler Ingress Controller binds all the certificates as SNI enabled certificates.

#### **Points to note**

- 1. When, multiple secrets are provided to the NetScaler Ingress Controller, the following precedence is followed:
  - a) preconfigured-default-certkey or non-host tls secret
  - b) default-ssl-certificate
- 2. If there is a conflict in precedence among the same grade certificates (for example, two ingress files configure a non-host TLS secret each, as default/non-SNI type), then the NetScaler Ingress Controller binds the NetScaler Ingress Controller default certificate as the non-SNI certificate and uses all other certificates with SNI.
- 3. Certificate used for a secret given under the TLS section must have a CN name. Otherwise, it does not bind to NetScaler.
- 4. If SNI enabled for SSL virtual server then:
  - Non-SNI (Default) certificate is used for the following HTTPs requests:

```
1 curl -1 -v -k https://1.1.1.1/
2
3 curl -1 -v -k -H 'HOST:*.colddrink.beverages' https://
1.1.1.1/
```

• SNI enabled certificate is used for a request with full domain name:

```
1 curl -1 -v -k https://items.colddrink.beverages/
```

If any request is received that does not match with certificates, CN name fails.

# TLS client authentication support in NetScaler

#### December 31, 2023

In TLS client authentication, a server requests a valid certificate from the client for authentication and ensures that it is only accessible by authorized machines and users.

You can enable TLS client authentication using NetScaler SSL-based virtual servers. With client authentication enabled on a NetScaler SSL virtual server, the NetScaler asks for the client certificate during the SSL handshake. The appliance checks the certificate presented by the client for normal constraints, such as the issuer signature and expiration date.



The following diagram explains the TLS client authentication feature on NetScaler.

TLS client authentication can be set to mandatory, or optional.

If the SSL client authentication is set as mandatory and the SSL Client does not provide a valid client certificate, then the connection is dropped. A valid client certificate means that it is signed or issued by a specific Certificate Authority, and not expired or revoked.

If it is marked as optional, then the NetScaler requests the client certificate, but the connection is not dropped. The NetScaler proceeds with the SSL transaction even if the client does not present a certificate or the certificate is invalid. The optional configuration is useful for authentication scenarios like two-factor authentication.

# **Configuring TLS client authentication**

Perform the following steps to configure TLS client authentication.

1. Enable the TLS support in NetScaler.

The NetScaler Ingress Controller uses the **TLS** section in the Ingress definition as an enabler for TLS support with NetScaler.

The following is a sample snippet of the Ingress definition:

```
1 spec:
2 tls:
3 - secretName:
```

2. Apply a CA certificate to the Kubernetes environment.

To generate a Kubernetes secret for an existing certificate, use the following kubectl command:

Note:

You must specify 'tls.crt='while creating a secret. This file is used by the NetScaler Ingress Controller while parsing a CA secret.

3. Configure Ingress to enable client authentication.

You need to specify the following annotation to attach the generated CA secret which is used for client certificate authentication for a service deployed in Kubernetes.

```
1 ingress.citrix.com/ca-secret: '{
2 "frontend-hotdrinks": "hotdrink-ca-secret" }
3 '
```

By default, client certificate authentication is set to mandatory but you can configure it to optional using the frontend\_sslprofile annotation in the front end configuration.

```
ingress.citrix.com/frontend_sslprofile: '{
    "clientauth":"ENABLED", "clientcert": "optional" }
    '
```

#### Note:

The frontend\_sslprofile only supports the front end Ingress configuration. For more information, see front end configuration.

# TLS server authentication support in NetScaler using the NetScaler Ingress Controller

December 31, 2023
Server authentication allows a client to verify the authenticity of the web server that it is accessing. Usually, the NetScaler device performs SSL offload and acceleration on behalf of a web server and does not authenticate the certificate of the Web server. However, you can authenticate the server in deployments that require end-to-end SSL encryption.

In such a situation, the NetScaler device becomes the SSL client and performs the following:

- carries out a secure transaction with the SSL server
- verifies that a CA whose certificate is bound to the SSL service has signed the server certificate
- checks the validity of the server certificate.

To authenticate the server, you must first enable server authentication and then bind the certificate of the CA that signed the certificate of the server to the SSL service on the NetScaler appliance. When you bind the certificate, you must specify the bind as a CA option.

#### **Configuring TLS server authentication**

Perform the following steps to configure TLS server authentication.

1. Enable the TLS support in NetScaler.

The NetScaler Ingress Controller uses the **TLS** section in the Ingress definition as an enabler for TLS support with NetScaler.

The following is a sample snippet of the Ingress definition:

1 spec: 2 tls: 3 - secretName:

- 2. To generate a Kubernetes secret for an existing certificate, perform the following.
  - a) Generate a client certificate to be used with the service.

b) Generate a secret for an existing CA certificate. This certificate is required to sign the back end server certificate.

#### Note:

You must specify tls.crt= while creating a secret. This file is used by the NetScaler Ingress Controller while parsing a CA secret.

3. Enable secure back end communication to the service using the following annotation in the Ingress configuration.

```
ingress.citrix.com/secure-backend: "True"
```

4. Use the following annotation to bind the certificate to SSL service. This certificate is used when the NetScaler acts as a client to send the request to the back end server.

```
1 ingress.citrix.com/backend-secret: '{
2 "tea-beverage": "tea-beverage", "coffee-beverage": "coffee-
beverage" }
3 '
```

5. To enable server authentication which authenticates the back end server certificate, you can use the following annotation. This configuration binds the CA certificate of the server to the SSL service on the NetScaler.



# Install, link, and update certificates on a NetScaler using the NetScaler Ingress Controller

June 26, 2025

On the Ingress NetScaler, you can install, link, and update certificates. Many server certificates are signed by multiple hierarchical certificate authorities (CAs). This means that certificates form a chain.

A certificate chain is an ordered list of certificates containing an SSL certificate and certificate authority (CA) certificates. It enables the receiver to verify that the sender and all CAs are trustworthy. The chain or path begins with the SSL certificate, and each certificate in the chain is signed by the entity identified by the next certificate in the chain.

Any certificate that sits between the SSL certificate and the root certificate is called a chain or intermediate certificate. The intermediate certificate is the signer or issuer of the SSL certificate. The root CA certificate is the signer or issuer of the intermediate certificate.

If the intermediate certificate is not installed on the server (where the SSL certificate is installed) it may prevent some browsers, mobile devices, and applications from trusting the SSL certificate. To make the SSL certificate compatible with all clients, it is necessary that the intermediate certificate is installed.



# Certificates linking in Kubernetes

The NetScaler Ingress Controller supports automatic provisioning and renewal of TLS certificates using the Kubernetes cert-manager. The cert-manager issues certificates from different sources, such as Let's Encrypt and HashiCorp Vault and converts them to Kubernetes secrets.

The following diagram explains how the cert-manager performs certificate management.



When you create a Kubernetes secret from a PEM certificate embedded with multiple CA certificates, you need to link the server certificates with the associated CAs.

While applying the Kubernetes secret, you can link the server certificates with all the associated CAs using the Ingress NetScaler. Linking the server certificates and CAs enable the receiver to verify if the sender and CAs are trustworthy.

The following is a sample Ingress definition:

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
4
   name: frontendssl
5 spec:
6
   rules:
     - host: frontend.com
7
       http:
8
         paths:
9
         - backend:
11
             service:
12
               name: frontend
13
               port:
14
                 number: 443
```

```
15 path: /web-frontend/frontend.php
16 pathType: Prefix
17 tls:
18 - secretName: certchain1
```

On the NetScaler, you can verify if certificates are added to the NetScaler. Perform the following:

- 1. Log on to the NetScaler command-line interface.
- 2. Verify if certificates are added to the NetScaler using the following command:

1 >show certkey

For sample outputs, see the NetScaler documentation.

3. Verify that the server certificate and CAs are linked using the following command:

1 >show certlink

#### Output:

#### 1 1) Cert Name: k8s-3KC24EQYHG6ZKEDAY5Y3SG26MT2 CA Cert Name: k8s -3KC24EQYHG6ZKEDAY5Y3SG2\_ic1 2 3 2) Cert Name: k8s-3KC24EQYHG6ZKEDAY5Y3SG2\_ic1 CA Cert Name: k8s -3KC24EQYHG6ZKEDAY5Y3SG2\_ic2

# Certificate key bundle in NetScaler by using the NetScaler Ingress Controller

NetScaler Ingress Controller now supports certificate bundle (certkeybundle) functionality, which is supported on NetScaler starting from release 14.1 build 21.x. With this functionality, the issue with the certificate chain and the additional handling that is required when two certificates share an intermediate CA are resolved. For more information on certificate key bundle support in NetScaler, see Support for SSL certificate key bundle.

NetScaler Ingress Controller creates a certificate bundle for certificates provided in:

- The TLS section of Ingress
- service.citrix.com/secret annotation in serviceTypeLB
- Secret in the Listener

To enable this feature, a new environment variable CERT\_BUNDLE is added in the NetScaler Ingress Controller, which can be set by using the certBundle argument from the Helm Charts. For more information on setting the certBundle, see Helm Charts of NetScaler Ingress Controller. NetScaler Ingress Controller adds the certificate bundle to the content switching virtual server of type SSL.

Certificate Bundle creation in NetScaler does not work in the following scenarios:

- A secure back end (service group of type SSL) is configured.
- Self-signed certificates are present.
- preconfigured-certkey annotation is set in the ingress, or servicetypeLB, or Listener.

#### Points to note

- 1. Certificate bundle creation in NetScaler fails if the following order is not met:
  - Server certificate (SC) must be placed at the top of the bundle file.
  - IC[1-9] are intermediate certificates. IC[i] is issued by IC[i+1]. The certificates must be placed in a sequence, and all the intermediate certificates must be present in the bundle.
  - Certificates must be of PEM format only.
  - Server certificate key (SCK) can be placed anywhere in the bundle.
  - A maximum of 9 intermediate certificates are supported
- 2. If you upgrade the Controllers with the CERT\_BUNDLE feature, the previous certificate key bindings get removed and a new certkeybundle gets created. The new certkeybundle gets bound to the context switching virtual server of type SSL.

# **Configure SSL passthrough using Kubernetes Ingress**

#### June 5, 2025

SSL passthrough feature allows you to pass incoming security sockets layer (SSL) requests directly to a server for decryption rather than decrypting the request using a load balancer. SSL passthrough is widely used for web application security and it uses the TCP mode to pass encrypted data to servers.

The proxy SSL passthrough configuration does not require the installation of an SSL certificate on the load balancer. SSL certificates are installed on the back end server as they handle the SSL connection instead of the load balancer.

The following diagram explains the SSL passthrough feature.



As shown in this diagram, SSL traffic is not terminated at the NetScaler and SSL traffic is passed through the NetScaler to the back end server. SSL certificate at the back end server is used for the SSL handshake.

The NetScaler Ingress Controller provides the following Ingress annotation that you can use to enable SSL passthrough on the Ingress NetScaler:

```
ingress.citrix.com/ssl-passthrough: 'True|False'
```

The default value of the annotation is False.

SSL passthrough is enabled for all services or host names provided in the Ingress definition. SSL passthrough uses host name (wildcard host name is also supported) and ignores paths given in Ingress.

Note:

The NetScaler Ingress Controller does not support SSL passthrough for non-hostname based Ingress. Also, SSL passthrough is not valid for default back end Ingress.

To configure SSL passthrough on the Ingress NetScaler, you must define the ingress.citrix.com /ssl-passthrough: as shown in the following sample Ingress definition. You must also enable TLS for the host as shown in the example.

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
4 annotations:
5
       ingress.citrix.com/frontend-ip: x.x.x.x
       ingress.citrix.com/insecure-termination: redirect
6
7
       ingress.citrix.com/secure-backend: "True"
       ingress.citrix.com/ssl-passthrough: "True"
8
   name: hotdrinks-ingress
9
10 spec:
11
    ingressClassName: citrix
12
   rules:
13
     - host: hotdrinks.beverages.com
14
       http:
15
         paths:
16
         - backend:
17
             service:
18
               name: frontend-hotdrinks
19
               port:
20
                 number: 443
21
           path: /
           pathType: Prefix
22
23
     tls:
24
     - secretName: beverages
```

# Automated certificate management with cert-manager

December 31, 2023

NetScaler Ingress Controller supports automatic provisioning and renewal of TLS certificates using cert-manager. The cert-manager is a native Kubernetes certificate management controller. It issues certificates from different sources, such as Let's Encrypt and HashiCorp Vault.

As shown in the following diagram, cert-manager interacts with the external Certificate Authorities (CA) to sign the certificates and converts it to Kubernetes secrets. These secrets are used by NetScaler Ingress Controller to configure SSL virtual server on the NetScaler.



For detailed configurations, refer:

- Deploying HTTPS web applications on Kubernetes with NetScaler Ingress Controller and Let's Encrypt using cert-manager
- Deploying HTTPS web application on Kubernetes with NetScaler Ingress Controller and HashiCorp Vault using cert-manager

# Deploy HTTPS web application on Kubernetes with the NetScaler Ingress Controller and Let's Encrypt using cert-manager

June 5, 2025

Let's Encrypt and the ACME (Automatic Certificate Management Environment) protocol enables you to set up an HTTPS server and automatically obtain a browser-trusted certificate. To get a certificate for your website's domain from Let's Encrypt, you have to demonstrate control over the domain by

accomplishing certain challenges. A challenge is one among a list of specified tasks that only someone who controls the domain can accomplish.

Currently there are two types of challenges:

- HTTP-01 challenge: HTTP-01 challenges are completed by posting a specified file in a specified location on a website. Let's Encrypt CA verifies the file by making an HTTP request on the HTTP URI to satisfy the challenge.
- DNS-01 challenge: DNS01 challenges are completed by providing a computed key that is present at a DNS TXT record. Once this TXT record has been propagated across the internet, the ACME server can successfully retrieve this key via a DNS lookup. The ACME server can validate that the client owns the domain for the requested certificate. With the correct permissions, cert-manager automatically presents this TXT record for your specified DNS provider.

On successful validation of the challenge, a certificate is granted for the domain.

This topic provides information on how to securely deploy an HTTPS web application on a Kubernetes cluster, using:

- The NetScaler Ingress Controller
- JetStack's cert-manager to provision TLS certificates from the Let's Encrypt project.

# Prerequisites

Ensure that you have:

- The domain for which the certificate is requested is publicly accessible.
- Enabled RBAC on your Kubernetes cluster.
- Deployed NetScaler MPX, VPX, or CPX deployed in Tier 1 or Tier 2 deployment model.

In the Tier 1 deployment model, NetScaler MPX or VPX is used as an Application Delivery Controller (ADC). The NetScaler Ingress Controller running in Kubernetes cluster configures the virtual services for services running on Kubernetes cluster. NetScaler runs the virtual service on the publicly routable IP address and offloads SSL for client traffic with the help of the Let's Encrypt generated certificate.

In the Tier 2 deployment model, a TCP service is configured on the NetScaler (VPX/MPX) running outside the Kubernetes cluster. This service is created to forward the traffic to NetScaler CPX instances running in the Kubernetes cluster. NetScaler CPX ends the SSL session and loadbalances the traffic to actual service pods.

• Deployed the NetScaler Ingress Controller. Click here for various deployment scenarios.

- Opened port 80 for the virtual IP address on the firewall for the Let's Encrypt CA to validate the domain for HTTP01 challenge.
- A DNS domain that you control, where you host your web application for the ACME DNS01 challenge.
- Administrator permissions for all deployment steps. If you encounter failures due to permissions, make sure you have administrator permissions.

# Install cert-manager

To install cert-manager, see the cert-manager installation documentation.

You can install cert-manager either using manifest files or Helm chart.

Once you install the cert-manager, verify that cert-manager is up and running as explained verifying the installation.

#### Deploy a sample web application

#### Perform the following to deploy a sample web application:

Note:

Kuard, a Kubernetes demo application is used for reference in this topic.

1. Create a deployment YAML file (kuard-deployment.yaml) for Kuard with the following configuration:

1	apiVersion: apps/v1
2	kind: Deployment
3	metadata:
4	labels:
5	app: kuard
6	name: kuard
7	spec:
8	replicas: 1
9	selector:
10	matchLabels:
11	app: kuard
12	template:
13	metadata:
14	labels:
15	app: kuard
16	spec:
17	containers:
18	- image: gcr.io/kuar-demo/kuard-amd64:1
19	<pre>imagePullPolicy: Always</pre>
20	name: kuard

21	ports:
22	- containerPort: 8080
23	protocol: TCP

2. Deploy the Kuard deployment file (kuard-deployment.yaml) to your cluster, using the following commands:

% kubectl create -f kuard-deployment.yaml

```
1 deployment.extensions/kuard created
```

% kubectl get pod -l app=kuard

1	NAME	READY	STATUS	RESTARTS	AGE
2					
3	kuard-6fc4d89bfb-djljt	1/1	Running	Θ	24s

3. Create a service for the deployment. Create a file called service.yaml with the following configuration:

1	apiVersion: v1
2	kind: Service
3	metadata:
4	name: kuard
5	spec:
6	ports:
7	- port: 80
8	targetPort: 8080
9	protocol: TCP
10	selector:
11	app: kuard

4. Deploy and verify the service using the following commands:

```
% kubectl create -f service.yaml
1
3 service/kuard created
4 % kubectl get svc kuard
5 NAME
          TYPE
                      CLUSTER-IP
                                      EXTERNAL-IP
                                                    PORT(S)
                                                              AGE
6 kuard
          ClusterIP
                      10.103.49.171
                                                    80/TCP
                                      <none>
                                                              13s
```

5. Expose this service to outside world by creating an Ingress that is deployed on NetScaler CPX or VPX as Content switching virtual server.

Note:

Ensures that you change the value of ingressClassName: citrix to your ingress class on which the NetScaler Ingress Controller is started.

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
```

```
3
       metadata:
4
          annotations:
5
         name: kuard
6
       spec:
7
         ingressClassName: citrix
8
          rules:
          - host: kuard.example.com
9
10
            http:
11
              paths:
12
              - backend:
13
                  service:
14
                     name: kuard
15
                     port:
16
                       number: 80
17
                path: /
18
                pathType: Prefix
```

#### Note:

You must change the value of spec.rules.host to the domain that you control. Ensure that a DNS entry exists to route the traffic to NetScaler CPX or VPX.

#### 1. Deploy the Ingress using the following command:

```
1 % kubectl apply -f ingress.yml
2 ingress.extensions/kuard created
3
4 root@ubuntu-:~/cert-manager# kubectl get ingress
5 NAME HOSTS ADDRESS PORTS AGE
6 kuard kuard.example.com 80 7s
```

2. Verify that the Ingress is configured on NetScaler CPX or VPX by using the following command:

```
1 $ kubectl exec -it cpx-ingress-5b85d7c69d-ngd72 /bin/bash
2
3 root@cpx-ingress-55c88788fd-qd4rg:/# cli_script.sh 'show cs
      vserver'
4 exec: show cs vserver
5 1) k8s-192.168.8.178_80_http (192.168.8.178:80) - HTTP Type:
      CONTENT
6
     State: UP
7
    Last state change was at Sat Jan 4 13:36:14 2020
8
     Time since last state change: 0 days, 00:18:01.950
9
     Client Idle Timeout: 180 sec
     Down state flush: ENABLED
     Disable Primary Vserver On Down : DISABLED
11
12
    Comment: uid=
        MPPL57E3AFY6NMNDGDKN2VT57HEZV0V53Z7DWKH44X2SGLIH4ZWQ====
13
     Appflow logging: ENABLED
     Port Rewrite : DISABLED
14
     State Update: DISABLED
15
     Default: Content Precedence: RULE
16
     Vserver IP and Port insertion: OFF
17
```

```
18 L2Conn: OFF Case Sensitivity: ON
19
    Authentication: OFF
20
    401 Based Authentication: OFF
21
   Push: DISABLED Push VServer:
22
   Push Label Rule: none
23
   Persistence: NONE
24
   Listen Policy: NONE
25
    IcmpResponse: PASSIVE
   RHIstate: PASSIVE
27
    Traffic Domain: 0
28 Done
29
30 root@cpx-ingress-55c88788fd-qd4rg/# exit
31 exit
```

3. Verify that the webpage is correctly being served when requested using the curl command.

```
1 % curl -sS -D - kuard.example.com -o /dev/null
2 HTTP/1.1 200 OK
3 Content-Length: 1458
4 Content-Type: text/html
5 Date: Thu, 21 Feb 2019 09:09:05 GMT
```

#### Configure issuing ACME certificate using the HTTP challenge

This section describes a way to issue the ACME certificate using the HTTP validation. If you want to use the DNS validation, skip this section and proceed to the next section.

The HTTP validation using cert-manager is a simple way of getting a certificate from Let's Encrypt for your domain. In this method, you prove ownership of a domain by ensuring that a particular file is present at the domain. It is assumed that you control the domain if you are able to publish the given file under a given path.

#### Deploy the Let's Encrypt ClusterIssuer with the HTTP01 challenge provider

The cert-manager supports two different CRDs for configuration, an Issuer, scoped to a single namespace, and a ClusterIssuer, with cluster-wide scope.

For the NetScaler Ingress Controller to use the Ingress from any namespace, use ClusterIssuer. Alternatively, you can also create an Issuer for each namespace on which you are creating an Ingress resource.

For more information, see cert-manager documentation for HTTP validation.

1. Create a file called issuer-letsencrypt-staging.yaml with the following configuration:

```
1 apiVersion: cert-manager.io/v1alpha2
2 kind: ClusterIssuer
3 metadata:
4
    name: letsencrypt-staging
5 spec:
6
    acme:
7
      # You must replace this email address with your own.
      # Let's Encrypt will use this to contact you about expiring
8
9
      # certificates, and issues related to your account.
      email: user@example.com
11
       server: https://acme-staging-v02.api.letsencrypt.org/directory
      privateKeySecretRef:
12
13
        # Secret resource used to store the account's private key.
14
        name: example-issuer-account-key
15
      # Add a single challenge solver, HTTP01 using citrix
16
      solvers:
17
       - http01:
18
           ingress:
19
             class: citrix
```

spec.acme.solvers\[].http01.ingress.class refers to the Ingress class of NetScaler Ingress Controller. If the NetScaler Ingress Controller has no ingress class, you do not need to specify this field.

#### Note:

This is a sample Clusterissuer of cert-manager.io/v1alpha2 resource. For more information, see cert-manager http01 documentation.

The staging Let's Encrypt server issues fake certificate, but it is not bound by the API rate limits of the production server. This approach lets you set up and test your environment without worrying about rate limits. You can repeat the same step for the Let's Encrypt production server.

2. After you edit and save the file, deploy the file using the following command:

```
    % kubectl apply -f issuer-letsencrypt-staging.yaml
    clusterissuer "letsencrypt-staging" created
```

3. Verify that the issuer is created and registered to the ACME server.

```
    % kubectl get issuer
    NAME AGE
    letsencrypt-staging 8d
```

4. Verify that the ClusterIssuer is properly registered using the command kubectl describe issuer letsencrypt-staging:

```
    % kubectl describe issuer letsencrypt-staging
    3 Status:
    4 Acme:
```

5	Uri: https://acme-sta	ging-v02.api.letsencrypt.org/acme/acct
	/8200869	
6	Conditions:	
7	Last Transition Time:	2019-02-11T12:06:31Z
8	Message:	The ACME account was registered with
	the ACME server	
9	Reason:	ACMEAccountRegistered
10	Status:	True
11	Type:	Ready

#### Issue certificate for the Ingress object

Once ClusterIssuer is successfully registered, you can get a certificate for the Ingress domain 'kuard.example.com'.

You can request a certificate for the specified Ingress resource using the following methods:

- Adding Ingress-shim annotations to the ingress object.
- Creating a certificate CRD object.

The first method is quick and simple, but if you need more customization and granularity in terms of certificate renewal, you can choose the second method. You can choose the method according to your needs.

Adding Ingress-shim annotations to the Ingress object In this approach, you add the following two annotations to the Ingress object for which you request a certificate from the ACME server.

1 certmanager.io/cluster-issuer: "letsencrypt-staging"

Note:

You can find all supported annotations from cert-manager for Ingress-shim, at supportedannotations.

Also, modify the ingress.yaml to use TLS by specifying a secret.

```
apiVersion: networking.k8s.io/v1
1
2
   kind: Ingress
3 metadata:
4
      annotations:
5
         certmanager.io/cluster-issuer: letsencrypt-staging
     name: kuard
6
   spec:
7
      ingressClassName: citrix
8
       rules:
9
10
       - host: kuard.example.com
11
         http:
```

12	paths:
13	- backend:
14	service:
15	name: kuard
16	port:
17	number: 80
18	pathType: Prefix
19	path: /
20	tls:
21	- hosts:
22	<ul> <li>kuard.example.com</li> </ul>
23	<pre>secretName: kuard-example-tls</pre>

The cert-manager.io/cluster-issuer: \"letsencrypt-staging\" annotation tells cert-manager to use the letsencrypt-staging cluster-wide issuer to request a certificate from Let's Encrypt's staging servers. Cert-manager creates a certificate object that is used to manage the lifecycle of the certificate for kuard.example.com. The value for the domain name and challenge method for the certificate object is derived from the ingress object. Cert-manager manages the contents of the secret as long as the Ingress is present in your cluster.

Deploy the ingress.yaml file using the following command:

```
1 % kubectl apply -f ingress.yml
2
3 ingress.extensions/kuard configured
4 % kubectl get ingress kuard
5 NAME HOSTS ADDRESS PORTS AGE
6 kuard kuard.example.com 80, 443 4h39m
```

**Create a certificate CRD resource** Alternatively, you can deploy a certificate CRD object independent of the Ingress object. Documentation of "certificate"CRD can be found at HTTP validation.

1. Create the certificate.yaml file with the following configuration:

```
apiVersion: cert-manager.io/v1alpha2
           kind: Certificate
3
           metadata:
4
             name: example-com
5
             namespace: default
           spec:
6
             secretName: kuard-example-tls
7
8
             issuerRef:
9
              name: letsencrypt-staging
10
             commonName: kuard.example.com
11
             dnsNames:
12
             - www.kuard.example.com
```

The spec.secretName key is the name of the secret where the certificate is stored on successfully issuing the certificate.

1. Deploy the certificate.yaml file on the Kubernetes cluster:

```
1 kubectl create -f certificate.yaml
2 certificate.cert-manager.io/example-com created
```

2. Verify that certificate custom resource is created by the cert-manager which represents the certificate specified in the Ingress. After few minutes, if ACME validation goes well, certificate 'READY'status is set to true.

```
% kubectl get certificates.cert-manager.io kuard-example-tls
1
2
   NAME
                       READY
                               SECRET
                                                   AGE
3
   kuard-example-tls
                       True
                               kuard-example-tls
                                                   3m44s
4
5
6 % kubectl get certificates.cert-manager.io kuard-example-tls
7 Name:
                kuard-example-tls
                default
8 Namespace:
9 Labels:
                 <none>
10 Annotations: <none>
11 API Version: cert-manager.io/v1alpha2
12 Kind:
               Certificate
13 Metadata:
     Creation Timestamp:
                          2020-01-04T17:36:26Z
14
15
     Generation:
                          1
     Owner References:
16
17
      API Version:
                              extensions/v1beta1
       Block Owner Deletion: true
18
19
      Controller:
                              true
20
      Kind:
                              Ingress
21
     Name:
                              kuard
                              2cafa1b4-2ef7-11ea-8ba9-06bea3f4b04a
22
      UID:
23
     Resource Version:
                              81263
                              /apis/cert-manager.io/v1alpha2/
24
     Self Link:
        namespaces/default/certificates/kuard-example-tls
25
    UID:
                              bbfa5e51-2f18-11ea-8ba9-06bea3f4b04a
26 Spec:
27
    Dns Names:
      acme.cloudpst.net
28
29
     Issuer Ref:
      Group:
                   cert-manager.io
31
       Kind:
                   ClusterIssuer
32
       Name:
                   letsencrypt-staging
     Secret Name: kuard-example-tls
33
34 Status:
     Conditions:
       Last Transition Time: 2020-01-04T17:36:28Z
37
                              Certificate is up to date and has not
       Message:
          expired
38
       Reason:
                              Ready
39
       Status:
                              True
40
                              Ready
       Type:
                              2020-04-03T16:36:27Z
41
     Not After:
42 Events:
```

43	Туре	Reason	Age	From	Message
44					
45	Normal	GeneratedKey	24m	cert-manager	Generated a <b>new</b>
	priv	<b>ate</b> key			
46	Normal	Requested	24m	cert-manager	Created <b>new</b>
	Cert	ificateRequest	resou	rce "kuard-exa	mple-tls-3030465986"
47	Normal	Issued	24m	cert-manager	Certificate issued
	succ	essfully			

#### 3. Verify that the secret resource is created.

1	% kubectl get secret	kuard-example-tls		
2	NAME	TYPE	DATA	AGE
3	kuard-example-tls	kubernetes.io/tls	3	3m13s

#### Issuing an ACME certificate using the DNS challenge

This section describes a way to use the DNS validation to get the ACME certificate from Let'sEncrypt CA. With a DNS-01 challenge, you prove the ownership of a domain by proving that you control its DNS records. This is done by creating a TXT record with specific content that proves you have control of the domain's DNS records. For detailed explanation of DNS challenge and best security practices in deploying DNS challenge, see A Technical Deep Dive: Securing the Automation of ACME DNS Challenge Validation.

Note:

In this procedure, route53 is used as the DNS provider. For other providers, see cert-manager documentation of DNS validation.

#### Deploy the Let's Encrypt ClusterIssuer with the DNS01 challenge provider

Perform the following to deploy the Let's Encrypt ClusterIssuer with the DNS01 challenge provider:

- 1. Create an AWS IAM user account and download the secret access key ID and secret access key.
- 2. Grant the following IAM policy to your user:

Route53 access policy

- 3. Create a Kubernetes secret acme-route53 in kube-system namespace.
  - 1 % kubectl create secret generic acme-route53 --from-literal secret -access-key=<secret\_access\_key>
- 4. Create an Issuer or ClusterIssuer with the DNS01 challenge provider.

You can provide multiple providers under DNS01, and specify which provider to be used at the time of certificate creation.

You must have access to the DNS provider for cert-manager to create a TXT record. Credentials are stored in the Kubernetes secret specified in spec.dns01.secretAccessKeySecretRef . For detailed instructions on how to obtain credentials, see the DNS provider documentation.

1	apiVersion: cert-manager.io/v1alpha2
2	kind: ClusterIssuer
3	metadata:
4	name: letsencrypt-staging
5	spec:
6	acme:
7	# You must replace <b>this</b> email address with your own.
8	<pre># Let's Encrypt will use this to contact you about</pre>
	expiring
9	<pre># certificates, and issues related to your account.</pre>
10	email: user@example.com
11	<pre>server: https://acme-staging-v02.api.letsencrypt.org/</pre>
	directory
12	privateKeySecretRef:
13	name: example-issuer-account-key
14	solvers:
15	- dns01:
16	route53:
17	region: us-west-2
18	accessKeyID: <iamkey></iamkey>
19	secretAccessKeySecretRef:
20	name: acme-route53
21	key: secret-access-key

#### Note:

Replace user@example.com with your email address. For each domain mentioned in a DNS01 stanza, cert-manager uses the provider's credentials from the referenced Issuer to create a TXT record called \\\_acme-challenge. This record is then verified by the ACME server to issue the certificate. For more information about the DNS provider configuration, and the list of supported providers, see DNS01 reference doc.

5. After you edit and save the file, deploy the file using the following command:

```
1 % kubectl apply -f acme_clusterissuer_dns.yaml
2 clusterissuer "letsencrypt-staging" created
```

6. Verify if the issuer is created and registered to the ACME server using the following command:

```
1% kubectl get issuer2NAMEAGE3letsencrypt-staging8d
```

7. Verify if the ClusterIssuer is properly registered using the command kubectl

```
describe issuer letsencrypt-staging:
```

```
1 Status:
2
   Acme:
     Uri: https://acme-staging-v02.api.letsencrypt.org/acme/acct
3
         /8200869
4 Conditions:
5
     Last Transition Time:
                             2019-02-11T12:06:31Z
                             The ACME account was registered with
6
      Message:
         the ACME server
                             ACMEAccountRegistered
7
      Reason:
8
                             True
      Status:
9
      Type:
                             Ready
```

#### Issue certificate for the Ingress resource

Once the issuer is successfully registered, you can get a certificate for the ingress domain kuard. example.com. Similar to HTTP01 challenge, there are two ways you can request the certificate for a specified Ingress resource:

- Adding Ingress-shim annotations to the Ingress object.
- Creating a certificate CRD object. For detailed instructions, see Create a Certificate CRD resource.

Adding Ingress-shim annotations to the ingress object Add the following annotation to the Ingress object along with the spec.tls section:

```
1 certmanager.io/cluster-issuer: "letsencrypt-staging"
```

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
4
   annotations:
       cert-manager.io/cluster-issuer: letsencrypt-staging
5
6 name: kuard
7 spec:
8
    ingressClassName: citrix
9 rules:
     - host: kuard.example.com
10
11
      http:
         paths:
13
         - backend:
14
             service:
15
              name: kuard
16
               port:
17
                 number: 80
18
           pathType: Prefix
19
           path: /
```

```
20 tls:
21 - hosts:
22 - kuard.example.com
23 secretName: kuard-example-tls
```

The cert-manager creates a Certificate CRD resource with the DNS01 challenge. It uses credentials specified in the ClusterIssuer to create a TXT record in the DNS server for the domain you own. Then, Let's Encypt CA validates the content of the TXT record to complete the challenge.

Adding a Certificate CRD resource\*\* Alternatively, you can explicitly create a certificate custom resource definition resource to trigger automatic generation of certificates.

1. Create the certificate.yaml file with the following configuration:

```
apiVersion: cert-manager.io/v1alpha2
1
2
           kind: Certificate
3
           metadata:
4
             name: example-com
5
             namespace: default
6
           spec:
             secretName: kuard-example-tls
7
8
             issuerRef:
9
               name: letsencrypt-staging
             commonName: kuard.example.com
11
             dnsNames:
12
             - www.kuard.example.com
```

After successful validation of the domain name, certificate READY status is set to True.

2. Verify that the certificate is issued.

```
1% kubectl get certificate kuard-example-tls23NAMEREADYSECRETAGE4-example-tlsTruekuard-example-tls10m
```

You can watch the progress of the certificate as it is issued, using the following command:

```
% kubectl describe certificates kuard-example-tls | tail -n 6
```

1	Not After:			2020-04-04T13:34:23Z		
2	Events:					
3	Туре	Reason	Age	From	Message	
4						
5	Normal	Requested	11m	cert-manager	Created <b>new</b>	
	Cert	tificateRequ	lest res	source "kuard-e	example-tls-3030465986"	
6	Normal	Issued	7m21s	cert-manager	Certificate issued	
	suco	cessfully				

# Verify certificate in NetScaler

Letsencrypt CA successfully validated the domain and issued a new certificate for the domain. A kubernetes.io/tls secret is created with the secretName specified in the tls: field of the Ingress. Also, cert-manager automatically initiates a renewal, 30 days before the expiry.

For HTTP challenge, cert-manager creates a temporary Ingress resource to route the Let's Encrypt CA generated traffic to cert-manager pods. On successful validations of the domain, this temporary Ingress is deleted.

1. Verify that the secret is created using the following command:

1	% kubectl get secre	t kuard-example-tls		
2				
3	NAME	ТҮРЕ	DATA	AGE
4	kuard-example-tls	kubernetes.io/tls	3	30m

The NetScaler Ingress Controller picks up the secret and binds the certificate to the content switching virtual server on the NetScaler CPX. If there are any intermediate CA certificates, it is automatically linked to the server certificate and presented to the client during SSL negotiation.

2. Log on to NetScaler CPX and verify if the certificate is bound to the SSL virtual server.

```
1 % kubectl exec -it cpx-ingress-55c88788fd-n2x9r bash -c cpx-
      ingress
2 Defaulting container name to cpx-ingress.
3 Use 'kubectl describe pod/cpx-ingress-55c88788fd-n2x9r -n default'
       to see all of the containers in this pod.
4
5 % cli_script.sh 'sh ssl vs k8s-192.168.8.178_443_ssl'
6 exec: sh ssl vs k8s-192.168.8.178_443_ssl
7
8
     Advanced SSL configuration for VServer k8s-192.168.8.178_443_ssl
        :
9
     DH: DISABLED
     DH Private-Key Exponent Size Limit: DISABLED Ephemeral RSA:
10
        ENABLED Refresh Count: 0
11
    Session Reuse: ENABLED Timeout: 120 seconds
     Cipher Redirect: DISABLED
13
     ClearText Port: 0
     Client Auth: DISABLED
14
15
     SSL Redirect: DISABLED
16
     Non FIPS Ciphers: DISABLED
17
     SNI: ENABLED
18
     OCSP Stapling: DISABLED
19
    HSTS: DISABLED
20
    HSTS IncludeSubDomains: NO
21
    HSTS Max-Age: 0
22
    HSTS Preload: NO
23
     SSLv3: ENABLED TLSv1.0: ENABLED TLSv1.1: ENABLED TLSv1.2:
        ENABLED TLSv1.3: DISABLED
```

```
Push Encryption Trigger: Always
24
     Send Close-Notify: YES
     Strict Sig-Digest Check: DISABLED
     Zero RTT Early Data: DISABLED
27
28
     DHE Key Exchange With PSK: NO
29
     Tickets Per Authentication Context: 1
   , P_256, P_384, P_224, P_5216) CertKey Name: k8s-
      GVWNYGVZKKRHKF7MZVTLOAEZYBS Server Certificate for SNI
   7) Cipher Name: DEFAULT
     Description: Default cipher list with encryption strength >= 128
        bit
34 Done
36 % cli_script.sh 'sh certkey'
37
  1) Name: k8s-GVWNYGVZKKRHKF7MZVTLOAEZYBS
     Cert Path: k8s-GVWNYGVZKKRHKF7MZVTLOAEZYBS.crt
     Key Path: k8s-GVWNYGVZKKRHKF7MZVTLOAEZYBS.key
40
     Format: PEM
41
     Status: Valid,
                      Days to expiration:89
42
     Certificate Expiry Monitor: ENABLED
43
     Expiry Notification period: 30 days
44
     Certificate Type: "Client Certificate" "Server Certificate"
45
     Version: 3
46
     Serial Number: 03B2B57EA9E61A93F1D05EA3272FA95203C2
47
     Signature Algorithm: sha256WithRSAEncryption
     Issuer: C=US,O=Let's Encrypt,CN=Let's Encrypt Authority X3
48
49
     Validity
50
       Not Before: Jan 5 13:34:23 2020 GMT
       Not After : Apr 4 13:34:23 2020 GMT
51
52
     Subject: CN=acme.cloudpst.net
     Public Key Algorithm: rsaEncryption
54
     Public Key size: 2048
     Ocsp Response Status: NONE
56 2) Name: k8s-GVWNYGVZKKRHKF7MZVTLOAEZYBS_ic1
57
     Cert Path: k8s-GVWNYGVZKKRHKF7MZVTLOAEZYBS.crt_ic1
58
     Format: PEM
59
     Status: Valid,
                      Days to expiration:437
     Certificate Expiry Monitor: ENABLED
61
     Expiry Notification period: 30 days
62
     Certificate Type: "Intermediate CA"
63
     Version: 3
     Serial Number: 0A0141420000015385736A0B85ECA708
64
65
     Signature Algorithm: sha256WithRSAEncryption
     Issuer: O=Digital Signature Trust Co., CN=DST Root CA X3
     Validity
       Not Before: Mar 17 16:40:46 2016 GMT
68
69
       Not After : Mar 17 16:40:46 2021 GMT
     Subject: C=US,O=Let's Encrypt,CN=Let's Encrypt Authority X3
71
     Public Key Algorithm: rsaEncryption
72
     Public Key size: 2048
73
     Ocsp Response Status: NONE
74 Done
```

The HTTPS webserver is now UP with a fake LE signed certificate. Next step is to move to production with the actual Let's Encrypt certificates.

# **Move to production**

After successfully testing with Let's Encrypt-staging, you can get the actual Let's Encrypt certificate.

You need to change Let's Encrypt endpoint from https:acme-staging-v02.api.letsencrypt .org/directory to https:acme-v02.api.letsencrypt.org/directory

Then, change the name of the ClusterIssuer from letsencrypt-staging to letsencryptproduction

```
1 apiVersion: cert-manager.io/v1alpha2
2 kind: ClusterIssuer
3 metadata:
    name: letsencrypt-prod
4
5 spec:
6
    acme:
7
       # You must replace this email address with your own.
       # Let's Encrypt will use this to contact you about expiring
8
9
       # certificates, and issues related to your account.
10
       email: user@example.com
       server: https://acme-v02.api.letsencrypt.org/directory
11
       privateKeySecretRef:
12
13
         # Secret resource used to store the account's private key.
14
         name: example-issuer-account-key
       # Add a single challenge solver, HTTP01 using citrix
15
16
       solvers:
       - http01:
17
18
           ingress:
19
             class: citrix
```

#### Note:

Replace user@example.com with your email address.

Deploy the file using the following command:

```
1 % kubectl apply -f letsencrypt-prod.yaml
2
3 clusterissuer "letsencrypt-prod" created
```

Now, repeat the procedure of modifying the annotation in Ingress or creating a CRD certificate which triggers the generation of new certificate.

Note

Ensure that you delete the old secret so that cert-manager starts a fresh challenge with the production CA.

```
1 % kubectl delete secret kuard-example-tls
2
3 secret "kuard-example-tls" deleted
```

Once the HTTP website is up, you can redirect the traffic from HTTP to HTTPS using the annotation ingress.citrix.com/insecure-termination: redirect in the ingress object.

# Troubleshooting

Since the certificate generation involves multiple components, this section summarizes the troubleshooting techniques that you can use if there was failures.

#### Verify the status of certificate generation

The certificate CRD object defines the life cycle management of generation and renewal of certificates. You can view the status of the certificate using the kubectl describe command as follows.

```
1 % kubectl get certificate
2
                      READY
                              SECRET
                                                 AGE
3 NAME
4 kuard-example-tls
                    False kuard-example-tls
                                                 9s
5
6 % kubectl describe certificate kuard-example-tls
7
8 Status:
9
   Conditions:
      Last Transition Time: 2019-03-05T09:50:29Z
10
11
      Message:
                            Certificate does not exist
12
      Reason:
                            NotFound
13
      Status:
                            False
                            Ready
14
     Type:
15 Events:
                          Age From
16
   Type Reason
                                             Message
                           ____
17
     ____
                                ____
    Normal OrderCreated 22s cert-manager Created Order resource "
18
        kuard-example-tls-1754626579"
```

Also you can view the major certificate events using the kubectl events command:

```
1 kubectl get events
2
3 LAST SEEN TYPE REASON KIND MESSAGE
4 36s Normal Started Challenge Challenge
scheduled for processing
```

5	36s	Normal	Created	0rder	Created
	Challeng	e resourc	e "kuard-example-tls	-1754626579-0"	for domain "acme
	.cloudps	t.net"			
6	38s	Normal	OrderCreated	Certificate	Created Order
	resource	"kuard-e	xample-tls-175462657	9"	
7	38s	Normal	CreateCertificate	Ingress	Successfully
	created	Certifica	te "kuard-example-tl	s"	

#### Analyze logs from cert-manager

If there is a failure, first step is to analyze logs from the cert-manager component. Identify the certmanager pod using the following command:

1	% kubectl get po -n cert-manager			
2				
3	NAME	READY	STATUS	RESTARTS
	AGE			
4	cert-manager- <b>76</b> d48d47bf-5w4vx	1/1	Running	Θ
	23h			
5	cert-manager-webhook-67cfb86d56-6qtxr	1/1	Running	Θ
	23h			
6	cert-manager-webhook-ca-sync-x4q6f	0/1	Completed	4
	23h			

Here cert-manager-76d48d47bf-5w4vx is the main cert-manager pod, and other two pods are cert-manager webhook pods.

Get the logs of the cert-manager using the following command:

1 % kubectl logs -f cert-manager-76d48d47bf-5w4vx -n cert-manager

If there is any failure to get the certificate, the ERROR logs give details about the failure.

#### Check the Kubernetes secret

Use the kubectl describe command to verify if both certificates and key are populated in Kubernetes secret.

```
1 % kubectl describe secret kuard-example-tls
2
3 Name: kuard-example-tls
4 Namespace: default
5 Labels: certmanager.k8s.io/certificate-name=kuard-example-tls
6 Annotations: certmanager.k8s.io/alt-names: acme.cloudpst.net
7 certmanager.k8s.io/issuer-kind: ClusterIssuer
9 certmanager.k8s.io/issuer-name: letsencrypt-staging
10
11 Type: kubernetes.io/tls
```

12		
13	Data	
14	====	
15	tls.crt:	3553 bytes
16	tls.key:	1679 bytes
17	ca.crt:	0 bytes

If both tls.crt and tls.key are populated in the Kubernetes secret, certificate generation is complete. If only tls.key is present, certificate generation is incomplete. Analyze the cert-manager logs for more details about the issue.

#### Analyze logs from the NetScaler Ingress Controller

If a Kubernetes secret is generated and complete, but it is not uploaded to the NetScaler, you can analyze the logs from the NetScaler Ingress Controller using the following command.

1 % kubectl logs -f cpx-ingress-685c8bc976-zgz8q

# Deploy an HTTPS web application on Kubernetes with NetScaler Ingress Controller and HashiCorp Vault using cert-manager

June 5, 2025

For ingress resources deployed with the NetScaler Ingress Controller, you can automate TLS certificate provisioning, revocation, and renewal using cert-manager and HashiCorp Vault. This topic provides a sample workflow that uses HashiCorp Vault as a self-signed certificate authority for certificate signing requests from cert-manager.

Specifically, the workflow uses the Vault PKI Secrets Engine to create a certificate authority (CA). This tutorial assumes that you have a Vault server installed and reachable from the Kubernetes cluster. The PKI secrets engine of Vault is suitable for internal applications. For external facing applications that require public trust, see automating TLS certificates using Let's Encrypt CA.

The workflow uses a Vault secret engine and authentication methods. For the full list of Vault features, see the following Vault documentation:

- Vault Secrets Engines
- Vault Authentication Methods

This topic provides you information on how to deploy an HTTPS web application on a Kubernetes cluster, using:

- NetScaler Ingress Controller
- JetStack's cert-manager to provision TLS certificates from HashiCorp Vault
- HashiCorp Vault

#### Prerequisites

Ensure that you have:

- The Vault server is installed, unsealed, and is reachable from the Kubernetes cluster. For information on installing the Vault server, see the Vault installation documentation.
- Enabled RBAC on your Kubernetes cluster.
- Deployed NetScaler MPX, VPX, or CPX in Tier 1 or Tier 2 deployment model.

In the Tier 1 deployment model, NetScaler MPX or VPX is used as an Application Delivery Controller (ADC). The NetScaler Ingress Controller running in the Kubernetes cluster configures the virtual services for the services running on the Kubernetes cluster. NetScaler runs the virtual service on the publicly routable IP address and offloads SSL for client traffic with the help of the Let's Encrypt generated certificate.

In the Tier 2 deployment, a TCP service is configured on the NetScaler (VPX/MPX) running outside the Kubernetes cluster to forward the traffic to NetScaler CPX instances running in the Kubernetes cluster. NetScaler CPX ends the SSL session and load-balances the traffic to actual service pods.

- Deployed NetScaler Ingress Controller. See Deployment Topologies for various deployment scenarios.
- Administrator permissions for all the deployment steps. If you encounter failures due to permissions, make sure that you have the administrator permission.

#### Note:

The following procedure shows steps to configure Vault as a certificate authority with NetScaler CPX used as the ingress device. When a NetScaler VPX or MPX is used as the ingress device, the steps are the same except the steps to verify the ingress configuration in the NetScaler.

# Deploy cert-manager using the manifest file

Perform the following steps to deploy cert-manager using the supplied YAML manifest file.

1. Install cert-manager. For information on installing cert-manager, see the cert-manager documentation. 1 kubectl apply -f https://github.com/jetstack/cert-manager/releases /download/vx.x.x/cert-manager.yaml

You can also install cert-manager with Helm. For more information, see the cert-manager documentation.

2. Verify that cert-manager is up and running using the following command.

1 2	% kubectl -n cert-manager get all NAME	RE	ADY	STATUS	
	RESTARTS AGE				
3	pod/cert-manager-77fd74fb64-d68v7 4m41s	1/:	1	Running	Θ
4	<pre>pod/cert-manager-webhook-67bf86d45-k77j 4m41s</pre>	j 1/:	1	Running	Θ
5					
6	NAME TYPE	CLU	JSTER-	IP	
	EXTERNAL-IP PORT(S) AGE				
7	service/cert-manager-webhook ClusterI 443/TCP 13d	P 10	.108.1	61.154	<none></none>
8					
9	NAME	READY	UP-T	O-DATE	
1.0	AVAILABLE AGE	1 / 1	1		1
TO	13d deployment.apps/cert-manager	1/1	T		T
11	deployment.apps/cert-manager-webhook	1/1	1		1
12	130				
13	NAME		DE	SIRED	CURRENT
	READY AGE				
14	replicaset.apps/cert-manager-77fd74fb64 1 13d		1		1
15	replicaset.apps/cert-manager-webhook-67	bf86d4	51		1
1.0	1 13d				
16 17	NAME				TONS
1	DURATION AGE				TONS
18	iob.batch/cert-manager-webhook-ca-svnc			1/1	
	22s 13d			,	
19	job.batch/cert-manager-webhook-ca-sync- 21s 10d	1549750	6800	1/1	
20	job.batch/cert-manager-webhook-ca-sync- 19s 3d8h	155036	1600	1/1	
21					
22	NAME		SCHEDU	LE SUS	SPEND
23	ACTIVE LAST SCHEDULE AGE cronjob.batch/cert-manager-webhook-ca-sy 0 3d8h 13d	ync (	@weekl	y Fal	lse

# Deploy a sample web application

Perform the following steps to deploy a sample web application.

Note:

Kuard, a Kubernetes demo application is used for reference in this topic.

1. Create a deployment YAML file (kuard-deployment.yaml) for Kuard with the following configuration.

1	apiVersion: apps/v1
2	kind: Deployment
3	metadata:
4	name: kuard
5	spec:
6	replicas: 1
7	selector:
8	matchLabels:
9	app: kuard
10	template:
11	metadata:
12	labels:
13	app: kuard
14	spec:
15	containers:
16	- image: gcr.io/kuar-demo/kuard-amd64:1
17	imagePullPolicy: Always
18	name: kuard
19	ports:
20	- containerPort: 8080

2. Deploy the Kuard deployment file (kuard-deployment.yaml) to your cluster, using the following commands.

1	% kubectl create -f kuar	d-deploy	/ment.yaml		
2	deployment.extensions/kuard created				
3	% kubectl get pod -l app	=kuard			
4	NAME	READY	STATUS	RESTARTS	AGE
5	kuard-6fc4d89bfb-djljt	1/1	Running	Θ	24s

3. Create a service for the deployment. Create a file called service.yaml with the following configuration.

```
apiVersion: v1
1
2
         kind: Service
3
         metadata:
4
          name: kuard
5
         spec:
6
           ports:
7
           - port: 80
8
             targetPort: 8080
```

```
9 protocol: TCP
10 selector:
11 app: kuard
```

4. Deploy and verify the service using the following command.

```
    % kubectl create -f service.yaml
    service/kuard created
    % kubectl get svc kuard
    NAME TYPE CLUSTER-IP EXTERNAL-IP PORT(S) AGE
    kuard ClusterIP 10.103.49.171 <none> 80/TCP 13s
```

5. Expose this service to the outside world by creating an Ingress that is deployed on NetScaler CPX or VPX as Content switching virtual server.

#### Note:

Ensure that you change ingressClassName: citrix to your ingress class on which NetScaler Ingress Controller is started.

```
1
        apiVersion: networking.k8s.io/v1
2
        kind: Ingress
3
       metadata:
4
         annotations:
5
         name: kuard
6
       spec:
7
         ingressClassName: citrix
8
          rules:
          - host: kuard.example.com
9
            http:
11
              paths:
12
              - backend:
13
                  service:
14
                    name: kuard
15
                    port:
16
                      number: 80
17
                path: /
18
                pathType: Prefix
```

#### Note:

Change the value of spec.rules.host to the domain that you control. Ensure that a DNS entry exists to route the traffic to NetScaler CPX or VPX.

#### 6. Deploy the Ingress using the following command.

```
    % kubectl apply -f ingress.yml
    ingress.extensions/kuard created
    root@ubuntu-vivek-225:~/cert-manager# kubectl get ingress
    NAME HOSTS ADDRESS PORTS AGE
    kuard kuard.example.com 80 7s
```

7. Verify if the ingress is configured on NetScaler CPX or VPX using the following command.

```
1 kubectl exec -it cpx-ingress-5b85d7c69d-ngd72 /bin/bash
2 root@cpx-ingress-5b85d7c69d-ngd72:/# cli_script.sh 'sh cs vs'
3 exec: sh cs vs
4 1) k8s-10.244.1.50:80:http (10.244.1.50:80) - HTTP Type: CONTENT
5
     State: UP
     Last state change was at Thu Feb 21 09:02:14 2019
6
7
     Time since last state change: 0 days, 00:00:41.140
8
     Client Idle Timeout: 180 sec
9
    Down state flush: ENABLED
10
    Disable Primary Vserver On Down : DISABLED
11
    Comment: uid=75
        VBGF07NZXV7SCI4LSDJML2Q5X6FSNK6NXQPWGMD0YGBW2IM0GQ====
12
    Appflow logging: ENABLED
     Port Rewrite : DISABLED
13
     State Update: DISABLED
14
     Default: Content Precedence: RULE
15
     Vserver IP and Port insertion: OFF
16
    L2Conn: OFF Case Sensitivity: ON
18
    Authentication: OFF
19
    401 Based Authentication: OFF
20
    Push: DISABLED Push VServer:
21
    Push Label Rule: none
22
    Listen Policy: NONE
23
     IcmpResponse: PASSIVE
24
   RHIstate: PASSIVE
25
    Traffic Domain: 0
26 Done
27 root@cpx-ingress-5b85d7c69d-ngd72:/# exit
28 exit
```

8. Verify if the page is correctly being served when requested using the curl command.

```
1 % curl -sS -D - kuard.example.com -o /dev/null
2 HTTP/1.1 200 OK
3 Content-Length: 1458
4 Content-Type: text/html
5 Date: Thu, 21 Feb 2019 09:09:05 GMT
```

Once you have deployed the sample HTTP application, you can proceed to make the application available over HTTPS. Here the Vault server signs the CSR generated by the cert-manager and a server certificate is automatically generated for the application.

In the following procedure, you use the configured Vault as a certificate authority and configure the cert-manager to use the Vault as signing authority for the CSR.

# **Configure HashiCorp Vault as Certificate Authority**

In this procedure, you set up an intermediate CA certificate signing request using HashiCorp Vault. This Vault endpoint is used by the cert-manager to sign the certificate for the ingress resources.

Note:

Ensure that you have installed the  $\mathbf{j}\mathbf{q}$  utility before performing these steps.

#### Create a root CA

For the sample workflow you can generate your own Root Certificate Authority within the Vault. In a production environment, you should use an external Root CA to sign the intermediate CA that Vault uses to generate certificates. If you have a root CA generated elsewhere, skip this step.

Note:

PKI\_ROOT is a path where you mount the root CA, typically it is pki. \${DOMAIN} in this procedure is example.com

```
1 % export DOMAIN=example.com
2 % export PKI_ROOT=pki
3
4 % vault secrets enable -path="${
5
   PKI_ROOT }
   " pki
6
7
8 # Set the max TTL for the root CA to 10 years
9 % vault secrets tune -max-lease-ttl=87600h "${
10 PKI_ROOT }
   11
11
12
13 % vault write -format=json "${
   PKI_ROOT }
14
   "/root/generate/internal \
15
16
  common_name="${
17 DOMAIN }
18
   CA root" ttl=87600h | tee \
19 >(jq -r .data.certificate > ca.pem) \
20 >(jq -r .data.issuing_ca > issuing_ca.pem) \
21 >(jq -r .data.private_key > ca-key.pem)
23 #Configure the CA and CRL URLs:
24
25 % vault write "${
26
  PKI_ROOT }
27
   "/config/urls \
         issuing_certificates="${
28
   VAULT_ADDR }
29
30 /v1/${
```

#### **Generate an intermediate CA**

After creating the root CA, perform the following steps to create an intermediate CSR using the root CA.

1. Enable pki from a different path PKI\_INT from root CA, typically pki\\\_int. Use the following command:

```
1
      % export PKI_INT=pki_int
      % vault secrets enable -path=${
2
3
   PKI_INT }
4
   pki
5
6
     # Set the max TTL to 3 year
7
8
      % vault secrets tune -max-lease-ttl=26280h ${
9
   PKI_INT }
```

2. Generate CSR for \$ { DOMAIN } that needs to be signed by the root CA. The key is stored internally to the Vault. Use the following command:

```
% vault write -format=json "${
1
2
   PKI_INT }
3
   "/intermediate/generate/internal \
4
        common_name="${
5
   DOMAIN }
   CA intermediate" ttl=26280h | tee \
6
        >(jq -r .data.csr > pki_int.csr) \
7
8
        >(jq -r .data.private_key > pki_int.pem)
```

3. Generate and sign the \${ DOMAIN } certificate as an intermediate CA using root CA, store it as intermediate.cert.pem. Use the following command:

```
1 % vault write -format=json "${
2 PKI_ROOT }
3 "/root/sign-intermediate csr=@pki_int.csr
4 format=pem_bundle ttl=26280h \
5 | jq -r '.data.certificate' > intermediate.cert.pem
```

If you are using an external root CA, skip the preceding step and sign the CSR manually using the root CA.

4. Once the CSR is signed and the root CA returns a certificate, it needs to be added back into the Vault using the following command:

```
1 % vault write "${
2 PKI_INT }
3 "/intermediate/set-signed certificate=@intermediate.cert.pem
```

5. Set the CA and CRL location using the following command.

```
1
       vault write "${
2
    PKI_INT }
3
    "/config/urls issuing_certificates="${
4
   VAULT_ADDR }
5
    /v1/${
6
   PKI_INT }
7
    /ca" crl_distribution_points="${
8
   VAULT_ADDR }
   /v1/${
9
10 PKI_INT }
11
   /crl"
```

An intermediate CA is set up and can be used to sign certificates for ingress resources.

#### **Configure a role**

A role is a logical name which maps to policies. An administrator can control the certificate generation through the roles.

Create a role for the intermediate CA that provides a set of policies for issuing or signing the certificates using this CA.

There are many configurations that can be configured when creating roles. For more information, see the Vault role documentation.

For the workflow, create a role kube-ingress that allows you to sign certificates of \${ DOMAIN } and its subdomains with a TTL of 90 days.

```
# with a Max TTL of 90 days
1
2
      vault write ${
   PKI_INT }
3
   /roles/kube-ingress \
4
                 allowed_domains=${
5
   DOMAIN }
6
7
    8
                 allow_subdomains=true \
9
                 max_ttl="2160h" \
                 require_cn=false
```
#### **Create Approle based authentication**

After configuring an intermediate CA to sign the certificates, you need to provide an authentication mechanism for the cert-manager to use the Vault for signing the certificates. Cert-manager supports Approle authentication method which provides a way for the applications to access the Vault defined roles.

An AppRole represents a set of Vault policies and login constraints that must be met to receive a token with those policies. For more information on this authentication method, see the Approle documentation.

#### **Create an Approle**

Create an Approle named Kube-role. The secret\_id for the cert-manager should not be expired to use this Approle for authentication. Hence, do not set a TTL or set it to 0.

```
1 % vault auth enable approle
2
3 % vault write auth/approle/role/kube-role token_ttl=0
```

#### Associate a policy with the Approle

Perform the following steps to associate a policy with an Approle.

1. Create a file pki\_int.hcl with the following configuration to allow the signing endpoints of the intermediate CA.

2. Add the file to a new policy called <a href="https://www.signusing.command">kube\_allow\_signusing.command</a>.

```
1 vault policy write kube-allow-sign pki_int.hcl
```

3. Update this policy to the Approle using the following command.

```
1 vault write auth/approle/role/kube-role policies=kube-allow-sign
```

The kube-role approle allows you to sign the CSR with intermediate CA.

#### Generate the role ID and secret ID

The role ID and secret ID are used by the cert-manager to authenticate with the Vault.

Generate the role ID and secret ID and encode the secret ID with Base64. Perform the following:

```
1 % vault read auth/approle/role/kube-role/role-id
2 role_id db02de05-fa39-4855-059b-67221c5c2f63
3
4 % vault write -f auth/approle/role/kube-role/secret-id
5 secret_id 6a174c20-f6de-a53c-74d2-6018fcceff64
6 secret_id_accessor c454f7e5-996e-7230-6074-6ef26b7bcf86
7
8 # encode secret_id with base64
9 % echo 6a174c20-f6de-a53c-74d2-6018fcceff64 | base64
10 NmExNzRjMjAtZjZkZS1hNTNjLTc0ZDItNjAx0GZjY2VmZjY0Cg==
```

#### **Configure issuing certificates in Kubernetes**

After you have configured Vault as the intermediate CA, and the Approle authentication method for the cert-manager to access Vault, you need to configure the certificate for the ingress.

#### Create a secret with the Approle secret ID

Perform the following to create a secret with the Approle secret ID.

1. Create a secret file called secretid.yaml with the following configuration.

#### Note:

The secret ID data.secretId is the base64 encoded secret ID generated in Generate the role id and secret id. If you are using an Issuer resource in the next step, the secret must be in the same namespace as the Issuer. For ClusterIssuer, the secret must be in the cert-manager namespace.

2. Deploy the secret file (secretid.yaml) using the following command.

```
1 % kubectl create -f secretid.yaml
```

#### **Deploy the Vault cluster issuer**

The cert-manager supports two different CRDs for configuration, an Issuer, which is scoped to a single namespace, and a ClusterIssuer, which is cluster-wide. For the workflow, you need to use ClusterIssuer.

Perform the following steps to deploy the Vault cluster issuer.

1. Create a file called issuer-vault.yaml with the following configuration.

```
1 apiVersion: cert-manager.io/v1
2 kind: ClusterIssuer
3 metadata:
4
    name: vault-issuer
5 spec:
6
    vault:
       path: pki_int/sign/kube-ingress
7
8
      server: <vault-server-url>
9
       #caBundle: <base64 encoded caBundle PEM file>
10
       auth:
11
         appRole:
12
           path: approle
           roleId: "db02de05-fa39-4855-059b-67221c5c2f63"
13
14
           secretRef:
15
             name: cert-manager-vault-approle
16
             key: secretId
```

SecretRef is the Kubernetes secret name created in the previous step. Replace roleId with the role\_id retrieved from the Vault.

An optional base64 encoded caBundle in the PEM format can be provided to validate the TLS connection to the Vault Server. When caBundle is set it replaces the CA bundle inside the container running the cert-manager. This parameter has no effect if the connection used is in plain HTTP.

2. Deploy the file (issuer-vault.yaml) using the following command.

```
1 % kubectl create -f issuer-vault.yaml
```

3. Using the following command verify if the Vault cluster issuer is successfully authenticated with the Vault.

```
% kubectl describe clusterIssuer vault-issuer | tail -n 7
1
2
    Conditions:
3
      Last Transition Time:
                              2019-02-26T06:18:40Z
      Message:
                              Vault verified
4
5
                              VaultVerified
      Reason:
6
      Status:
                              True
7
      Type:
                              Ready
8 Events:
                              <none>
```

Now, you have successfully setup the cert-manager for Vault as the CA. The next step is securing the ingress by generating the server certificate. There are two different options for securing your ingress. You can proceed with one of the approaches to secure your ingresses.

- Ingress Shim approach
- Manually creating the certificate CRD object for the certificate.

#### Ingress-shim approach

In this approach, you modify the ingress annotation for the cert-manager to automatically generate the certificate for the given host name and store it in the specified secret.

1. Modify the ingress with the tls section specifying a host name and secret. Also, specify the cert-manager annotation cert-manager.io/cluster-issuer as follows.

1	apiVersion: networking.k8s.io/v1
2	kind: Ingress
3	metadata:
4	annotations:
5	cert-manager.io/cluster-issuer: vault-issuer
6	name: kuard
7	spec:
8	ingressClassName: citrix
9	rules:
10	- host: kuard.example.com
11	http:
12	paths:
13	- backend:
14	service:
15	name: kuard-service
16	port:
17	number: 80
18	path: /
19	pathType: Prefix
20	tls:
21	- hosts:
22	- kuard.example.com
23	secretName: kuard-example-tls

#### 2. Deploy the modified ingress as follows.

```
% kubectl apply -f ingress.yml
1
2
    ingress.extensions/kuard created
3
4
    % kubectl get ingress kuard
5
                                ADDRESS
                                          PORTS
                                                    AGE
    NAME HOSTS
6
    kuard
           kuard.example.com
                                          80, 443
                                                    12s
```

This step triggers a certificate object by the cert-manager which creates a certificate signing request (CSR) for the domain kuard.example.com. On successful signing of CSR, the certificate is stored in the secret name kuard-example-tls specified in the ingress.

1. Verify that the certificate is successfully issued using the following command.

```
% kubectl describe certificates kuard-example-tls | grep -A5
1
       Events
2
    Events:
3
                       Age From
    Type Reason
                                          Message
                            ____
                       ____
4
    ____
           _____
                                          _____
    Normal CertIssued 48s cert-manager Certificate issued
5
       successfully
```

#### Create a certificate CRD object for the certificate

Once the issuer is successfully registered, you need to get the certificate for the ingress domain kuard .example.com.

You need to create a certificate resource with the commonName and dnsNames. For more information, see cert-manager documentation. You can specify multiple dnsNames which are used for the SAN field in the certificate.

To create a "certificate" CRD object for the certificate, perform the following:

1. Create a file called certificate.yaml with the following configuration.

```
apiVersion: cert-manager.io/v1
1
2 kind: Certificate
3 metadata:
   name: kuard-example-tls
4
   namespace: default
5
6 spec:
7
    secretName: kuard-example-tls
8 issuerRef:
     kind: ClusterIssuer
9
      name: vault-issuer
10
   commonName: kuard.example.com
12
   duration: 720h
13 #Renew before 7 days of expiry
14 renewBefore: 168h
15
   commonName: kuard.example.com
    dnsNames:
17
    - www.kuard.example.com
```

The certificate has CN=kuard.example.com and SAN=Kuard.example.com,www.kuard.example.com.

spec.secretName is the name of the secret where the certificate is stored after the certificate is issued successfully.

2. Deploy the file (certificate.yaml) on the Kubernetes cluster using the following command.

% kubectl create -f certificate.yaml certificate.certmanager.k8s.io/kuard-example-tls created

#### Verify if the certificate is issued

You can watch the progress of the certificate as it is issued using the following command:

```
1 % kubectl describe certificates kuard-example-tls | grep -A5 Events
2 Events:
3
   Туре
           Reason
                      Age From
                                          Message
4
                       -----
    ____
           _____
                                          ____
   Normal CertIssued 48s cert-manager Certificate issued
5
       successfully > **Note** > You may encounter some errors due to
       the Vault policies. If you encounter any such errors, return to
       the Vault and fix it.
```

After successful signing, a kubernetes.io/tls secret is created with the secretName specified in the Certificate resource.

```
1% kubectl get secret kuard-example-tls2NAMETYPEDATAAGE3kuard-exmaple-tlskubernetes.io/tls34m20s
```

#### Modify the ingress to use the generated secret

Perform the following steps to modify the ingress to use the generated secret.

 Edit the original ingress and add a spec.tls section specifying the secret kuard-example -tls as follows.

1	apiVersion: networking.k8s.io/v1
2	kind: Ingress
3	metadata:
4	annotations:
5	name: kuard
6	spec:
7	<pre>ingressClassName: citrix</pre>
8	rules:
9	- host: kuard.example.com
10	http:
11	paths:
12	- backend:
13	service:
14	name: kuard
15	port:
16	number: 80
17	pathType: Prefix
18	path: /
19	tls:

```
20 - hosts:
21 - kuard.example.com
22 secretName: kuard-example-tls
```

2. Deploy the ingress using the following command.

```
1 % kubectl apply -f ingress.yml
2 ingress.extensions/kuard created
3
4 % kubectl get ingress kuard
5 NAME HOSTS ADDRESS PORTS AGE
6 kuard kuard.example.com 80, 443 12s
```

#### Verify the Ingress configuration in NetScaler

Once the certificate is successfully generated, NetScaler Ingress Controller uses this certificate for configuring the front-end SSL virtual server. You can verify it with the following steps.

1. Log on to NetScaler CPX and verify if the Certificate is bound to the SSL virtual server.

```
1 % kubectl exec -it cpx-ingress-668bf6695f-4fwh8 bash
2 cli_script.sh 'shsslvs'
3 exec: shsslvs
4 1) Vserver Name: k8s-10.244.3.148:443:ssl
5
     DH: DISABLED
6
     DH Private-Key Exponent Size Limit: DISABLED Ephemeral RSA:
        ENABLED Refresh Count: 0
7
     Session Reuse: ENABLED Timeout: 120 seconds
     Cipher Redirect: DISABLED
8
9
     SSLv2 Redirect: DISABLED
     ClearText Port: 0
11
     Client Auth: DISABLED
12
     SSL Redirect: DISABLED
13
    Non FIPS Ciphers: DISABLED
14
     SNI: ENABLED
15
     OCSP Stapling: DISABLED
    HSTS: DISABLED
16
17
    HSTS IncludeSubDomains: NO
18
    HSTS Max-Age: 0
19
     SSLv2: DISABLED SSLv3: ENABLED TLSv1.0: ENABLED TLSv1.1:
        ENABLED TLSv1.2: ENABLED TLSv1.3: DISABLED
20
     Push Encryption Trigger: Always
     Send Close-Notify: YES
22
     Strict Sig-Digest Check: DISABLED
23
     Zero RTT Early Data: DISABLED
24
     DHE Key Exchange With PSK: NO
25
     Tickets Per Authentication Context: 1
26 Done
27
28 root@cpx-ingress-668bf6695f-4fwh8:/# cli_script.sh 'shsslvs k8s
      -10.244.3.148:443:ssl'
```

```
29 exec: shsslvs k8s-10.244.3.148:443:ssl
     Advanced SSL configuration for VServer k8s-10.244.3.148:443:ssl:
31
32
     DH: DISABLED
33
     DH Private-Key Exponent Size Limit: DISABLED Ephemeral RSA:
        ENABLED Refresh Count: 0
     Session Reuse: ENABLED Timeout: 120 seconds
34
35
     Cipher Redirect: DISABLED
     SSLv2 Redirect: DISABLED
37
     ClearText Port: 0
38
     Client Auth: DISABLED
     SSL Redirect: DISABLED
39
     Non FIPS Ciphers: DISABLED
40
     SNI: ENABLED
41
     OCSP Stapling: DISABLED
42
43
     HSTS: DISABLED
44
     HSTS IncludeSubDomains: NO
45
     HSTS Max-Age: 0
46
     SSLv2: DISABLED SSLv3: ENABLED TLSv1.0: ENABLED TLSv1.1:
        ENABLED TLSv1.2: ENABLED TLSv1.3: DISABLED
47
     Push Encryption Trigger: Always
48
     Send Close-Notify: YES
49
     Strict Sig-Digest Check: DISABLED
     Zero RTT Early Data: DISABLED
51
     DHE Key Exchange With PSK: NO
52
     Tickets Per Authentication Context: 1
53 , P_256, P_384, P_224, P_5216) CertKey Name: k8s-
      LM0303U6KC6WXKCBJAQY6K6X6J0 Server Certificate for SNI
54
55 7) Cipher Name: DEFAULT
     Description: Default cipher list with encryption strength >= 128
        bit
57 Done
58
59 root@cpx-ingress-668bf6695f-4fwh8:/# cli_script.sh 'sh certkey k8s
       -LM0303U6KC6WXKCBJAQY6K6X6J0'
60 exec: sh certkey k8s-LMO3O3U6KC6WXKCBJAQY6K6X6JO
     Name: k8s-LM0303U6KC6WXKCBJAQY6K6X6J0 Status: Valid,
                                                             Days to
        expiration:0
     Version: 3
62
63
     Serial Number: 524C1D9306F784A2F5277C05C2A120D5258D9A2F
64
     Signature Algorithm: sha256WithRSAEncryption
     Issuer: CN=example.com CA intermediate
     Validity
67
       Not Before: Feb 26 06:48:39 2019 GMT
68
       Not After : Feb 27 06:49:09 2019 GMT
     Certificate Type: "Client Certificate" "Server Certificate"
69
70
     Subject: CN=kuard.example.com
71
     Public Key Algorithm: rsaEncryption
72
     Public Key size: 2048
73
     Ocsp Response Status: NONE
74
     2) URI:http://127.0.0.1:8200/v1/pki_int/crl
     3) VServer name: k8s-10.244.3.148:443:ssl Server Certificate for
```

76 Done

SNI

The HTTPS webserver is up with the vault signed certificate. Cert-manager automatically renews the certificate as specified in the RenewBefore parameter in the certificate, before expiry of the certificate.

Note:

The Vault signing of the certificate fails if the expiry of a certificate is beyond the expiry of the root CA or intermediate CA. You should ensure that the CA certificates are renewed in advance before the expiry.

2. Verify that the application is accessible using the HTTPS protocol.

```
1 % curl -sS -D - https://kuard.example.com -k -o /dev/null
2 HTTP/1.1 200 OK
3 Content-Length: 1472
4 Content-Type: text/html
5 Date: Tue, 11 May 2021 20:39:23 GMT
```

# Enable NetScaler certificate validation in the NetScaler Ingress Controller

#### December 31, 2023

The NetScaler Ingress Controller provides an option to ensure secure communication between the NetScaler Ingress Controller and NetScaler by using the HTTPS protocol. You can achieve this by using pre-loaded certificates in the NetScaler. As an extra measure to avoid any possible man-in-the-middle (MITM) attack, the NetScaler Ingress Controller also allows you to validate the SSL server certificate provided by the NetScaler.

To enable certificate signature and common name validation of the ADC server certificate by the NetScaler Ingress Controller, security administrators can optionally install signed (or self-signed) certificates in the NetScaler and configure the NetScaler Ingress Controller with the corresponding CA certificate bundle. Once the validation is enabled and CA certificate bundles are configured, the NetScaler Ingress Controller starts validating the certificate (including certificate name validation). If the validation fails, the NetScaler Ingress Controller logs the same and none of the configurations are used on an unsecure channel.

This validation is turned off by default and an administrator can chose to enable the validation in the NetScaler Ingress Controller as follows.

#### Prerequisites

- For enabling certificate validation, you must configure a NetScaler with proper SSL server certificates (with proper server name or IP address in certificate subject). For more information, see NetScaler documentation.
- The CA certificate for the installed server certificate-key pair is used to configure the NetScaler Ingress Controller to enable validation of these certificates.

#### Configure the NetScaler Ingress Controller for certificate validation

To make a CA certificate available for configuration, you need to configure the CA certificate as a Kubernetes secret so that the NetScaler Ingress Controller can access it on a mounted storage volume.

To generate a Kubernetes secret for an existing certificate, use the following kubectl command:

Alternatively, you can also generate the Kubernetes secret using the following YAML definition:

```
1 apiVersion: v1
2 kind: Secret
3 metadata:
4 name: ciccacert
5 data:
6 myCA.pem: <base64 encoded cert>
```

The following is a sample YAML file with the NetScaler Ingress Controller configuration for enabling certificate validation.

```
1 kind: Pod
2 metadata:
3
   name: cic
   labels:
4
5
      app: cic
6 spec:
7
    serviceAccountName: cpx
   # Make secret available as a volume
8
9
    volumes:
    - name: certs
11
     secret:
        secretName: ciccacert
12
13 containers:
14
    - name: cic
      image: "xxxx"
15
       imagePullPolicy: Always
16
17
      args: []
```

```
18
       # Mounting certs in a volume path
19
       volumeMounts:
20
       - name: certs
         mountPath: <Path to mount the certificate>
21
22
         readOnly: true
23
       env:
24
       # Set NetScaler ADM Management IP
25
       - name: "NS_IP"
26
         value: "xx.xx.xx.xx"
27
       # Set port for Nitro
28
       - name: "NS_PORT"
29
       value: "xx"
       # Set Protocol for Nitro
       - name: "NS_PROTOCOL"
31
32
         # Enable HTTPS protocol for secure communication
33
         value: "HTTPS"
34
       # Set username for Nitro
       - name: "NS_USER"
        value: "nsroot"
37
       # Set user password for Nitro
38
       - name: "NS_PASSWORD"
         value: "nsroot"
39
40
       # Certificate validation configurations
       - name: "NS_VALIDATE_CERT"
41
         value: "yes"
42
       - name: "NS_CACERT_PATH"
43
44
         value: " <Mounted volume path>/myCA.pem"
```

As specified in the example YAML file, following are the specific changes required for enabling certificate validation in the NetScaler Ingress Controller.

#### **Configure Kubernetes secret as a volume**

• Configure a volume section declared with secret as the source. Here, secretName should match the Kubernetes secret name created for the CA certificate.

#### Configure a volume mount location for the CA certificate

- Configure a volumeMounts section with the same name as that of secretName in the volume section
- Declare a mountPath directory to mount the CA certificate
- Set the volume as ReadOnly

#### Configure secure communication

• Set the environment variable NS\_PROTOCOL as HTTPS

• Set the environment variable NS\_PORT as ADC HTTPS port

#### Enable and configure CA validation and certificate path

- Set the environment variable NS\_VALIDATE\_CERT to yes (no for disabling)
- Set the environment variable NS\_CACERT\_PATH as the mount path (volumeMounts->mountPath)/ PEM file name (used while creating the secret).

### **Disable API server certificate verification**

December 31, 2023

While communicating with the API server from NetScaler Ingress Controller or GSLB ingress, you have the option to disable the API server certificate verification on NetScaler Ingress Controller.

# Disable API server certificate verification on NetScaler Ingress Controller or GSLB ingress

When you deploy NetScaler Ingress Controller using YAML, you can disable the API server certificate verification by providing the following argument in the NetScaler Ingress Controller deployment YAML file.

```
1 args:
2 - --disable-apiserver-cert-verify
3 true
```

When you deploy NetScaler Ingress Controller using Helm charts, the parameter disableAPIServerCertVeri can be mentioned as True in the Helm values file as follows:

```
1 disableAPIServerCertVerify: True
```

# Create a self-signed certificate and linking into Kubernetes secret

June 5, 2025

Use the steps in the procedure to create a self-signed certificate using OpenSSL and link into Kubernetes secret. You can use this secret to secure your Ingress.

#### Create a self-signed certificate

You can create a TLS secret by using the following steps. In this procedure, a self-signed certificate and key are created.

You can link it to the Kubernetes secret and use that secret in the Ingress for securing the Ingress.

Note:

Here, example.com is used for reference. You must replace example.com with the required domain name.

In the example, the generated certificate has a validity of one year as the days are mentioned as 365 days.

#### Linking the certificate to a Kubernetes secret

Perform the following steps to link the certificate to the Kubernetes secret.

1. Run the following command to create a Kubernetes secret based on the TLS certificate that you have created.

2. Run the following command to view the secret that contains the TLS certificate information:

```
1 kubectl get secret tls-secret
```

#### **Deploy the Ingress**

Create and apply the Ingress configuration. The following YAML can be used for reference.

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
4 name: ingress-demo
5 namespace: netscaler
6 annotations:
7 spec:
8 ingressClassName: citrix
9 tls:
```

```
10
     - secretName: tls-secret
11
      hosts:
       - "example.com"
12
13 rules:
14
     - host: "example.com"
15
       http:
16
         paths:
         - path: /
17
           pathType: Prefix
18
19
           backend:
20
             service:
21
               name: service-test
22
               port:
23
                 number: 80
```

# Authentication and authorization policies for Kubernetes with NetScaler

#### October 17, 2024

Authentication and authorization policies are used to enforce access restrictions to the resources hosted by an application or API server. While you can verify the identity using the authentication policies, authorization policies are used to verify whether a specified request has the necessary permissions to access a resource.

NetScaler provides a Kubernetes CustomResourceDefinition (CRD) called the **Auth CRD** that you can use with NetScaler Ingress Controller to define authentication policies on NetScaler.

#### **Auth CRD definition**

The Auth CRD is available in the NetScaler Ingress Controller GitHub repo at: auth-crd.yaml. The Auth CRD provides attributes for the various options that are required to define the authentication policies on NetScaler.

#### Auth CRD attributes

The Auth CRD provides the following attributes that you use to define the authentication policies:

- servicenames
- authentication\_mechanism
- authentication\_providers
- authentication\_policies

#### authorization\_policies

#### Servicenames

The name of the services for which the authentication and authorization policies need to be applied.

#### **Authentication mechanism**

The following authentication mechanisms are supported:

• Using request headers:

Enables user authentication using the request header. You can use this mechanism when the credentials or API keys are passed in a header (typically Authorization header). For example, you can use authentication using request headers for basic, digest, bearer authentication, or API keys.

• Using forms:

You can use this mechanism with user or web authentication including the relying party configuration for OpenID connect and the service provider configuration for SAML.

When the authentication mechanism is not specified, the default is authentication using the request header.

The following are the attributes for forms based authentication.

Attribute	Description
authentication_host	Specifies a fully qualified domain name (FQDN) to which the user must be redirected for ADC authentication service. This FQDN should be unique and should resolve to the front-end IP address of NetScaler with Ingress/service type LoadBalancer or the VIP address of the Listener
authentication_host_cert	CRD. Specifies the name of the SSL certificate to be used with the authentication_host. This certificate is mandatory while performing
ingress_name	authentication using the form. Specifies the Ingress name for which the authentication using forms is applicable.

Attribute	Description
ingressclass	Specifies the ingress class so that only the
	ingress controller associated with the specified
	ingress class processes the resource. Otherwise,
	all the controllers in the cluster will process this resource.
lb_service_name	Specifies the name of the service of type
	LoadBalancer for which the authentication using
	forms is applicable.
listener_name	The name of the Listener CRD for which the
	authentication using forms is applicable.
vip	Specifies the front-end IP address of the Ingress
	for which the authentication using forms is
	applicable. This attribute refers to the
	frontend-ip address provided with the
	Ingress. If there is more than one Ingress
	resource which uses the same frontend-ip, it is
	recommended to use vip.

#### Note:

- While using forms, authentication can be enabled for all types of traffic. Currently, granular authentication is not supported.
- Depending on the resource to which you need to apply form based authentication, you can use one of the ingress\_name, lb\_service\_name, listener\_name, or vip attributes to specify the resource.

#### Authentication providers

The **providers** define the authentication mechanism and parameters that are required for the authentication mechanism.

**Basic authentication** Specifies that local authentication is used with the HTTP basic authentication scheme. To use basic authentication, you must create user accounts on the ingress NetScaler.

Note:

The user for local auth must be created on NetScaler VPX or NetScaler MPX. To create a user

for local auth on NetScaler using the CLI, run the following command: add aaa user < username> -password <password>. For information on how to create a user for local auth on NetScaler using the GUI, see Configuring Local Users.

**OAuth authentication** The OAuth authentication mechanism, requires an external identity provider to authenticate the client using oAuth2 and issue an Access token. When the client presents the Access token to a NetScaler as an access credential, the NetScaler validates the token using the configured values. If the token validation is successful then NetScaler grants access to the client.

Attribute	Description
Issuer	The identity (usually a URL) of the server whose tokens need to be accepted for authentication.
jwks_uri	The URL of the endpoint that contains JWKs (JSON Web Key) for JWT (JSON Web Token)
audience	The identity of the service or application for which the token is applicable.
token_in_hdr	The custom header name where the token is present. The default value is the Authorization header. Note: You can specify more than one header.
token_in_param	The query parameter where the token is present.
signature_algorithms	Specifies the list of signature algorithms which are allowed. By default HS256, RS256, and RS512 algorithms are allowed.
introspect_url	The URL of the introspection endpoint of the authentication server (IdP). If the access token presented is an opaque token, introspection is
client_credentials	used for the token verification. The name of the Kubernetes secrets object that contains the client id and client secret required
claims_to_save	to authenticate with the authentication server. The list of claims to be saved. Claims are used to create authorization policies.

**OAuth authentication attributes** The following are the attributes for OAuth authentication:

OpenID Connect (OIDC) is a simple identity layer on top of the OAuth 2.0 protocol. OIDC allows clients

to verify the identity of the end-user based on the authentication performed by an authorization server, as well as to obtain basic profile information about the end-user. In addition to the OAuth attributes, you can use the following attributes to configure OIDC.

Attribute	Description
metadata_url	Specifies the URL that is used to get OAUTH or
	OIDC provider metadata.
user_field	Specifies the attribute in the token from which
	the user name should be extracted. By default,
	NetScaler examines the email attribute for user
	ID.
default_group	Specifies the group assigned to the request if
	authentication succeeds. This group is in
	addition to any extracted groups from the token.
grant_type	Specifies the type of flow to the token end point.
	The default value is CODE.
pkce	Specifies whether to enable Proof Key for Code
	Exchange (PKCE). The default value is ENABLED.
token_ep_auth_method	Specifies the authentication method to be used
	with the token end point. The default value is
	client_secret_post.

**SAML authentication** Security assertion markup language (SAML) is an XML-based open standard which enables authentication of users across products or organizations. The SAML authentication mechanism, requires an external identity provider to authenticate the client. SAML works by transferring the client identity from the identity provider to the NetScaler. On successful validation of the client identity, the NetScaler grants access to the client.

The following are the attributes for SAML authentication.

Attribute	Description
metadata_url	Specifies the URL used for obtaining SAML metadata.
<pre>metadata_refresh_interval</pre>	Specifies the interval in minutes for fetching metadata from the specified metadata URL.
signing_cert	Specifies the SSL certificate to sign requests from the service provider (SP) to the identity provider (IdP).

Attribute	Description
audience	Specifies the identity of the service or
	application for which the token is applicable.
issuer_name	Specifies the name used in requests sent from
	SP to IdP to identify the NetScaler.
binding	Specifies the transport mechanism of the SAML
	message. The default value is POST.
artifact_resolution_service_url	Specifies the URL of the artifact resolution
	service on IdP.
logout_binding	Specifies the transport mechanism of the SAML
	logout. The default value is POST.
reject_unsigned_assertion	Rejects unsigned SAML assertions. If this value is
	ON, it rejects assertion without signature.
user_field	Specifies the SAML user ID specified in the SAML
	assertion
default_authentication_group	Specifies the default group that is chosen when
	the authentication succeeds in addition to
	extracted groups.
skewtime	Specifies the allowed clock skew time in minutes
	on an incoming SAML assertion.
attributes_to_save	Specifies the list of attribute names separated by
	commas which needs to be extracted and stored
	as key-value pairs for the session on NetScaler.

**LDAP authentication** LDAP (Lightweight Directory Access Protocol) is an open, vendor-neutral, industry standard application protocol for accessing and maintaining distributed directory information services over an Internet Protocol (IP) network. A common use of LDAP is to provide a central place to store user names and passwords. LDAP allows many different applications and services to connect to the LDAP server to validate users.

Note:

LDAP authentication is supported through both the authentication mechanisms using the request header or using forms.

The following are the attributes for LDAP authentication.

Attribute	Description
server_ip	Specifies the IP address assigned to the LDAP server.
server_name	Specifies the LDAP server name as an FQDN.
server_port	Specifies the port on which the LDAP server accepts connections. The default value is 389.
base	Specifies the base node on which to start LDAP searches. If the LDAP server is running locally, the default value of base is dc=netscaler,
server_login_credentials	Specifies the Kubernetes secret object providing credentials to log in to the LDAP server. The secret data should have user name and password.
login_name	Specifies the <b>LDAP login name</b> attribute. The NetScaler uses the LDAP login name to query external LDAP servers or Active Directories.
security_type	Specifies the type of security used for communications between the NetScaler and the LDAP server. The default is TLS.
validate_server_cert	Validates LDAP server certificates. The default value is NO.
hostname	Specifies the host name for the LDAP server. If validate_server_cert is ON, this value must be the host name on the certificate from the LDAP. A host name mismatch causes a connection failure.
<pre>sub_attribute_name</pre>	Specifies the LDAP group subattribute name. This attribute is used for group extraction from the LDAP server.
group_attribute_name	Specifies the LDAP group attribute name. This attribute is used for group extraction on the LDAP server.

Attribute	Description
search_filter	Specifies the string to be combined with the
	default LDAP user search string to form the
	search value. For example, if the search filter
	"vpnallowed=true"is combined with the LDAP
	login name "samaccount"and the user-supplied
	user name is "bob", the result is the LDAP search
	string ""
	(&(vpnallowed=true)(samaccount=bob)'".
	Enclose the search string in two sets of double
	quotation marks.
auth_timeout	Specifies the number of seconds the NetScaler
	waits for a response from the server. The default
	value is 3.
password_change	Allows password change requests. The default
	value is DISABLED.
attributes_to_save	List of attribute names separated by comma
	which needs to be fetched from the LDAP server
	and stored as key-value pairs for the session on
	NetScaler.

#### **Authentication policies**

The **authentication\_policies** allow you to define the traffic selection criteria to apply the authentication mechanism and also to specify the provider that you want to use for the selected traffic.

Authentication policy supports two formats through which you can specify authentication rules:

- resource format
- expression format

The following are the attributes for policies with resource format:

Attribute	Description
path	An array of URL path prefixes that refer to a specific API endpoint. For example, /api/v1/products/.

Attribute	Description
method	An array of HTTP methods. Allowed values are
	GET, PUT, POST, or DELETE. The traffic is selected
	if the incoming request URI matches with any of
	the paths and any of the listed methods. If the
	method is not specified then the path alone is
	used for the traffic selection criteria.
provider	Specifies the authentication mechanism that
	needs to be used. If the authentication
	mechanism is not provided, then authentication
	is not performed.

The following attributes are for authentication policies with expression format:

Attribute	Description
expression	Specifies NetScaler expression to be evaluated
	based on authentication
provider	Specifies the authentication mechanism that
	needs to be used. If the authentication
	mechanism is not provided, then authentication
	is not performed.

#### Note:

If you want to skip authentication for a specific end point, create a policy with the provider attribute set as empty list. Otherwise, the request is denied.

#### Authorization policies

Authorization policies allow you to define the traffic selection criteria to apply the authorization requirements for the selected traffic.

Authorization policy supports two formats through which the you can specify the authorization rules:

- resource format
- expression format

The following are the attributes for authorization policies with resource format:

Attribute	Description
path	An array of URL path prefixes that refer to a
	specific API endpoint. For example,
	/api/v1/products/.
method	An array of HTTP methods. Allowed values are
	GET, PUT, POST, or DELETE.
claims	Specifies the claims required to access a specific
	API endpoint. name indicates the claim name
	and values indicate the required permissions.
	You can have more than one claim. If an empty
	list is specified, it implies that authorization is
	not required.
	<b>Note:</b> Any claim that needs to be used for
	authorization, should be saved as part of
	authentication.

The following are the attributes for authorization policies with expression format:

Attribute	Description
expression	Specifies an expression to be evaluated for
	authorization.

#### Note:

NetScaler requires both authentication and authorization policies for the API traffic. Therefore, you must configure an authorization policy with an authentication policy. Even if you do not have any authorization checks, you must create an authorization policy with empty claims. Otherwise, the request is denied with a 403 error.

#### Note:

Authorization would be successful if the incoming request matches a policy (path, method, and claims). All policies are tried until there is a match. If it is required to selectively bypass authorization for a specific end point, an explicit policy needs to be created.

#### **Deploy the Auth CRD**

Perform the following to deploy the Auth CRD:

- 1. Download the CRD (auth-crd.yaml).
- 2. Deploy the Auth CRD using the following command:

```
1 kubectl create -f auth-crd.yaml
```

For example:

#### How to write authentication and authorization policies

After you have deployed the CRD provided by NetScaler in the Kubernetes cluster, you can define the authentication policy configuration in a .yaml file. In the .yaml file, use authpolicy in the kind field and in the spec section add the **Auth CRD** attributes based on your requirement for the policy configuration.

After you deploy the .yaml file, NetScaler Ingress Controller applies the authentication policy configuration on NetScaler.

#### Local auth provider

The following is a sample authentication and authorization policy definition for the local-auth-provider type (local\_auth.yaml).

```
1
     apiVersion: citrix.com/v1beta1
2
     kind: authpolicy
    metadata:
3
4
       name: authexample
5
    spec:
         servicenames:
6
7
         - frontend
8
         authentication_providers:
9
             - name: "local-auth-provider"
11
               basic_local_db:
12
                    use_local_auth: 'YES'
13
         authentication_policies:
14
             - resource:
16
                  path:
17
                    - '/orders/'
                    - '/shipping/'
18
                  method: [GET, POST]
19
20
                provider: ["local-auth-provider"]
```

21	
22	<pre># skip authentication for this</pre>
23	- resource:
24	path:
25	- '/products/'
26	method: [GET]
27	provider: []
28	
29	authorization_policies:
30	<pre># skip authorization</pre>
31	- resource:
32	path: []
33	<pre>method: []</pre>
34	claims: []

The sample policy definition performs the following:

- NetScaler performs the local authentication on the requests to the following:
  - **GET** or **POST** operation on orders and shipping end points.
- NetScaler does not perform the authentication for **GET** operation on the **products** endpoint.
- NetScaler does not apply any authorization permissions.

#### Note:

The user for local auth must be created on NetScaler VPX or NetScaler MPX. To create a user for local auth on NetScaler using the CLI, run the following command: add aaa user < username> -password <password>. For information on how to create a user for local auth on NetScaler using the GUI, see Configuring Local Users.

#### oAuth JWT verification

The following is a sample authentication and authorization policy definition for oAuth JWT verification (oauth\_jwt\_auth.yaml).

```
1 apiVersion: citrix.com/v1beta1
2 kind: authpolicy
3 metadata:
4 name: authexample
5 spec:
6
      servicenames:
       - frontend
7
8
9
       authentication_providers:
         - name: "jwt-auth-provider"
10
11
           oauth:
             issuer: "https://sts.windows.net/tenant1/"
12
             jwks_uri: "https://login.microsoftonline.com/tenant1/
13
                discovery/v2.0/keys"
```

```
audience : ["https://api.service.net"]
14
15
              claims_to_save : ["scope"]
16
17
       authentication_policies:
18
            - resource:
19
                path:
                  - '/orders/'
20
                  - '/shipping/'
21
22
                method: [GET, POST]
23
              provider: ["jwt-auth-provider"]
24
25
            # skip authentication for this
26
            - resource:
27
                path:
                  - '/products/'
28
29
                method: [GET]
30
              provider: []
31
32
       authorization_policies:
33
            - resource:
                path:
34
                  - '/orders/'
                  - '/shipping/'
                method: [POST]
37
38
                claims:
39
                  - name: "scope"
40
                    values: ["read", "write"]
41
            - resource:
42
                path:
                  - '/orders/'
43
44
                method: [GET]
45
                claims:
                  - name: "scope"
46
47
                    values: ["read"]
            # skip authorization, no claims required
48
49
            - resource:
50
                path:
                  - '/shipping/'
51
                method: [GET]
52
53
                claims: []
```

The sample policy definition performs the following:

- NetScaler performs JWT verification on the requests to the following:
  - The GET or POST operation on orders and shipping endpoints.
- NetScaler skips authentication for the **GET** operation on the **products** endpoint.
- NetScaler requires the scope claim with read and write permissions for **POST** operation on **orders** and **shipping** endpoints.

- NetScaler requires the scope claim with the read permission for GET operation on the orders endpoint.
- NetScaler does not need any permissions for **GET** operation on the **shipping** end point.

For OAuth, if the token is present in a custom header, it can be specified using the token\_in\_hdr attribute as follows:

```
1 oauth:
2 issuer: "https://sts.windows.net/tenant1/"
3 jwks_uri: "https://login.microsoftonline.com/tenant1/discovery/
v2.0/keys"
4 audience : ["https://vault.azure.net"]
5 token_in_hdr : [ " custom-hdr1 " ]
```

Similarly, if the token is present in a query parameter, it can be specified using the token\_in\_param attribute as follows:

```
1 oauth:
2 issuer: "https://sts.windows.net/tenant1/"
3 jwks_uri: "https://login.microsoftonline.com/tenant1/discovery/
v2.0keys"
4 audience : ["https://vault.azure.net"]
5 token_in_param : [ " query-param1 " ]
```

#### oAuth Introspection

The following is a sample authentication and authorization policy definition for oAuth JWT verification. (oauth\_intro\_auth.yaml)

```
apiVersion: citrix.com/v1beta1
1
2
     kind: authpolicy
3
    metadata:
4
       name: authexample
5
    spec:
6
         servicenames:
7
         - frontend
8
         authentication_providers:
9
             - name: "introspect-provider"
11
               oauth:
                 issuer: "ns-idp"
12
                 jwks_uri: "https://idp.aaa/oauth/idp/certs"
13
                  audience : ["https://api.service.net"]
14
                  client_credentials: "oauthsecret"
15
16
                  introspect_url: https://idp.aaa/oauth/idp/introspect
17
                  claims_to_save : ["scope"]
18
19
         authentication_policies:
20
             - resource:
```

```
21
                  path: []
22
                  method: []
23
                provider: ["introspect-provider"]
24
25
          authorization_policies:
26
              - resource:
27
                  path: []
28
                  method: [POST]
29
                  claims:
                  - name: "scope"
                    values: ["read", "write"]
32
              - resource:
33
                  path: []
34
                  method: [GET]
                  claims:
                  - name: "scope"
37
                    values: ["read"]
```

The sample policy definition performs the following:

- NetScaler performs the oAuth introspection as specified in the provider introspectprovider for all requests.
- NetScaler requires the scope claim with read and write permissions for all **POST** requests.
- NetScaler requires the scope claim with the read permission for all **GET** requests.

#### Creating a secrets object with client credentials for introspection

A Kubernetes secrets object is needed for configuring the OAuth introspection. You can create a secret object in a similar way as shown in the following example:

```
1 apiVersion: v1
2 kind: Secret
3 metadata:
4 name: oauthsecret
5 type: Opaque
6 stringData:
7 client_id: "nsintro"
8 client_secret: "nssintro"
```

#### Note:

Keys of the opaque secret object must be client\_id and client\_secret. A user can set the values for them as desired.

#### SAML authentication using forms

The following is an example for SAML authentication using forms. In the example, authhost-tlscert-secret and saml-tls-cert-secret are Kubernetes TLS secrets referring to certificate and key.

Note:

When certkey.cert and certkey.key are certificate and key respectively for the authentication host, then the authhost-tls-cert-secret can be formed using the following command:

```
1 kubectl create secret tls authhost-tls-cert-secret --key="certkey.
key" --cert="certkey.cert
```

Similarly, you can use this command to form saml-tls-cert-secret with the required certificate and key.

```
1 apiVersion: citrix.com/v1beta1
2 kind: authpolicy
3 metadata:
   name: samlexample
4
5 spec:
       servicenames:
6
       - frontend
7
8
9
       authentication_mechanism:
10
         using_forms:
           authentication_host: "fqdn_authenticaton_host"
11
           authentication_host_cert:
12
13
              tls_secret: authhost-tls-cert-secret
14
           listener_name: "example-listener"
15
       authentication_providers:
16
           - name: "saml-auth-provider"
17
18
             saml:
19
                 metadata_url: "https://idp.aaa/metadata/samlidp/aaa"
20
                  signing_cert:
21
                      tls_secret: saml-tls-cert-secret
22
23
       authentication_policies:
24
25
           - resource:
               path: []
26
27
               method: []
             provider: ["saml-auth-provider"]
28
29
       authorization_policies:
31
32
           - resource:
               path: []
```

 34
 method: []

 35
 claims: []

The sample policy definition performs the following:

• NetScaler performs SAML authentication as specified in the provider saml-auth-provider for all requests.

Note: Granular authentication is not supported for the forms mechanism.

- NetScaler requires the group claim with admin permission for all **POST** requests.
- NetScaler does not require any specific permission for **GET** requests.

#### **OpenID Connect authentication using forms**

The following is an example for creating OpenID Connect authentication to configure NetScaler in a Relaying Party (RP) role to authenticate users for an external identity provider. The authentication\_mechanism must be set to using\_forms to trigger the OpenID Connect procedures.

```
1 apiVersion: citrix.com/v1beta1
2 kind: authpolicy
3 metadata:
4
   name: authoidc
5
   spec:
       servicenames:
6
7
       - frontend
       authentication_mechanism:
8
9
           using_forms:
                authentication_host: "10.221.35.213"
10
                authentication_host_cert:
11
                     tls_secret: "oidc-tls-secret"
                ingress_name: "example-ingress"
13
14
15
       authentication_providers:
16
           - name: "oidc-provider"
17
18
             oauth:
               audience : ["https://app1.citrix.com"]
19
20
               client_credentials: "oidcsecret"
               metadata_url: "https://10.221.35.214/oauth/idp/.well-known/
                   openid-configuration"
               default_group: "groupA"
22
23
               user_field: "sub"
24
                pkce: "ENABLED"
25
                token_ep_auth_method: "client_secret_post"
26
27
       authentication_policies:
28
29
            - resource:
```

```
path: []
31
                method: []
              provider: ["oidc-provider"]
32
33
34
       authorization_policies:
35
            #default - no authorization requirements
37
            - resource:
38
                path: []
                method: []
40
                claims: []
```

The sample policy definition performs the following:

 NetScaler performs OIDC authentication (relying party) as specified in the provider oidcprovider for all requests.

Note: Granular authentication is not supported for the forms mechanism.

• NetScaler does not require any authorization permissions.

#### LDAP authentication using the request header

The following is an example for LDAP authentication using the request header.

In this example, ldapcredential is the Kubernetes secret referring to the LDAP server credentials. See the ldap\_secret.yaml file for information on how to create LDAP server credentials.

```
1 apiVersion: citrix.com/v1beta1
2 kind: authpolicy
3 metadata:
4 name: ldapexample
5 spec:
       servicenames:
6
7
       - frontend
8
       authentication_providers:
9
10
           - name: "ldap-auth-provider"
11
             ldap:
                 server_ip: "192.2.156.160"
                 base: 'dc=aaa,dc=local'
13
                 login_name: accountname
14
15
                 sub_attribute_name: CN
16
                  server_login_credentials: ldapcredential
17
           - name: "local-auth-provider"
18
19
             basic_local_db:
20
                 use_local_auth: 'YES'
21
22
       authentication_policies:
23
```

```
24
            - resource:
25
                path: []
26
                method: []
              provider: ["ldap-auth-provider"]
27
28
29
        authorization_policies:
31
32
            - resource:
33
                path: []
34
                method: []
                claims: []
```

**Note:** With the request header based authentication mechanism, granular authentication based on traffic is supported.

#### LDAP authentication using forms

In the example authhost-tls-cert-secret is the Kubernetes TLS secret referring to certificate and key.

When certkey.cert and certkey.key are certificate and key respectively for the authentication host, then the authhost-tls-cert-secret can be formed using the following command:

```
1 kubectl create secret tls authhost-tls-cert-secret --key="certkey.
key" --cert="certkey.cert
```

In this example, ldapcredential is the Kubernetes secret referring to the LDAP server credentials. See the ldap\_secret.yaml file for information on how to create LDAP server credentials.

```
1 apiVersion: citrix.com/v1beta1
2 kind: authpolicy
3 metadata:
4 name: ldapexample
5 spec:
       servicenames:
6
       - frontend
7
8
       authentication_mechanism:
9
10
         using_forms:
11
           authentication_host: "fqdn_authenticaton_host"
12
           authentication_host_cert:
13
             tls_secret: authhost-tls-cert-secret
           vip: "192.2.156.156"
14
15
16
       authentication_providers:
17
           - name: "ldap-auth-provider"
18
             ldap:
19
                  server_ip: "192.2.156.160"
```

```
base: 'dc=aaa,dc=local'
20
21
                  login_name: accountname
22
                  sub_attribute_name: CN
23
                  server_login_credentials: ldapcredential
24
25
       authentication_policies:
26
27
            - resource:
28
                path: []
29
                method: []
              provider: ["ldap-auth-provider"]
31
32
       authorization_policies:
33
34
            - resource:
                path: []
                method: []
37
                claims: []
```

The sample policy definition performs the following:

- NetScaler performs the LDAP authentication for entire traffic (all requests).
- NetScaler does not apply any authorization permission.

The following is an example for LDAP\_secret.yaml.

```
1 apiVersion: v1
2 kind: Secret
3 metadata:
4 name: ldapcredential
5 type: Opaque
6 stringData:
7 username: 'ldap_server_username'
8 password: 'ldap_server_password'
```

#### Example for NetScaler expression support with Auth CRD

This example shows how you can specify NetScaler expressions along with authentication and authorization policies:

```
1 apiVersion: citrix.com/v1beta1
2
   kind: authpolicy
3
  metadata:
4
     name: authexample
5
  spec:
6
     servicenames:
7

    frontend

8
        authentication_mechanism:
9
          using_request_header: 'ON'
10
11
```

12	authentication_providers:
13	<pre>- name: "ldap-auth-provider"</pre>
14	ldap:
15	server_ip: "192.2.156.160"
16	base: 'dc=aaa,dc=local'
17	login_name: accountname
18	<pre>sub_attribute_name: CN</pre>
19	<pre>server_login_credentials: ldapcredential</pre>
20	<pre># "memberof" attribute details are extracted from LDAP</pre>
	server.
21	attributes_to_save: memberof
22	
23	authentication_policies:
24	<pre># Perform LDAP authentication for the host hotdrink.beverages</pre>
25	- expression: 'HTTP RED HOSTNAME SET TEXT MODE(IGNORECASE) ED
23	("hotdrink.beverages.com")'
26	provider: ["ldap-auth-provider"]
27	
28	
29	authorization_policies:
30	<pre># ALLOW the session only if the authenticated user is     associated with attribute "memberof" having value "grp4"</pre>
31	<pre>- expression: 'aaa.user.attribute("memberof").contains("grp4 ")'</pre>

# **Rate limiting in Kubernetes using NetScaler**

#### October 17, 2024

In a Kubernetes deployment, you can rate limit the requests to the resources on the back end server or services using rate limiting feature provided by the ingress NetScaler.

NetScaler provides a Kubernetes CustomResourceDefinitions (CRDs) called the **Rate limit CRD** that you can use with the NetScaler Ingress Controller to configure the rate limiting configurations on the NetScalers used as Ingress devices.

Apart from rate limiting the requests to the services in a Kubernetes environment, you can use the Rate limit CRD for API security as well. The Rate limit CRD allows you to limit the REST API request to API servers or specific API endpoints on the API servers. It monitors and keeps track of the requests to the API server or endpoints against the allowed limit per time slice and hence protects from attacks such as the DDoS attack.

You can enable logging for observability with the rate limit CRD. Logs are stored on NetScaler which can be viewed by checking the logs using the shell command. The file location is based on the syslog configuration. For example, /var/logs/ns.log.

#### Rate limit CRD definition

The Rate limit CRD spec is available in the NetScaler Ingress Controller GitHub repo at: ratelimitcrd.yaml. The **Rate limit CRD provides** attributes for the various options that are required to define the rate limit policies on the Ingress NetScaler that acts as an API gateway.

#### **Rate limit CRD attributes**

The following table lists the various attributes provided in the Rate limit CRD:

CRD attribute	Description
ingressclass	Specifies the ingress class so that only the ingress controller associated with the specified ingress class processes the resource. Otherwise, all the controllers in the cluster will process this resource
servicename	The list of Kubernetes services to which you want to apply the rate limit policies.
selector_keys	The traffic selector keys that filter the traffic to identify the API requests against which the throttling is applied and monitored. <b>Note:</b> The selector_keys is an optional attribute. You can choose to configure zero, one or more of the selector keys. If more than one selector keys are configured then it is considered as a logical AND expression.
	<b>path:</b> An array of URL path prefixes that refer to a specific API endpoint. For example, /api/v1/products/.
	<b>method:</b> An array of HTTP methods. Allowed values are GET, PUT, POST, or DELETE.
	<b>header_name:</b> HTTP header that has the unique API client or user identifier. For example, X-apikey which comes with a unique API-key that identifies the API client sending the request.

CRD attribute	Description
	<b>per_client_ip:</b> Allows you to monitor and apply the configured threshold to each API request received per unique client IP address.
req_threshold	The maximum number of requests that are allowed in the given time slice (request rate).
timeslice	The time interval specified in microseconds (multiple of 10 s), during which the requests are monitored against the configured limits. If not specified it defaults to 1000 milliseconds.
limittype	It allows you to configure the type of throttling algorithms that you want to use to apply the limit. Supported algorithms are burst and smooth. The default is the <b>burst</b> mode.
throttle_action	It allows you to define the throttle action that needs to be taken on the traffic that is throttled for crossing the configured threshold.
	<b>DROP:</b> Drops the requests above the configured traffic limits.
	<b>RESET:</b> Resets the connection for the requests crossing the configured limit.
	<b>RESPOND:</b> Responds with the standard " <b>429</b> <b>Too many requests</b> "response.
redirect_url	This attribute is an optional attribute that is required only if throttle_action is configured with the value REDIRECT.
logpackets	Enables audit logs.
logexpression	Specifies the default-syntax expression that defines the format and content of the log message.
loglevel	Specifies the severity level of the log message that is generated.
# **Deploy the Rate limit CRD**

Perform the following to deploy the Rate limit CRD:

- 1. Download the CRD (ratelimit-crd.yaml).
- 2. Deploy the Rate limit CRD using the following command:

```
1 kubectl create -f ratelimit-crd.yaml
```

For example,

# How to write a rate-based policy configuration

After you have deployed the CRD provided by NetScaler in the Kubernetes cluster, you can define the rate-based policy configuration in a .yaml file. In the .yaml file, use ratelimit in the kind field and in the spec section add the Rate limit CRD attributes based on your requirement for the policy configuration.

After you deploy the .yaml file, the NetScaler Ingress Controller applies the rate-based policy configuration on the Ingress NetScaler device.

Following are some examples for rate limit policy configurations.

## Limit API requests to configured API endpoint prefixes

Consider a scenario wherein you want to define a rate-based policy in NetScaler to limit the API requests to 15 requests per minute from each unique client IP address to the configured API endpoint prefixes. Create a .yaml file called ratelimit-example1.yaml and use the appropriate CRD attributes to define the rate-based policy as follows:

```
1 apiVersion: citrix.com/v1beta1
2 kind: ratelimit
3 metadata:
4 name: throttle-req-per-clientip
5 spec:
6 servicenames:
```

```
7
      - frontend
8
   selector_keys:
9
    basic:
10
     path:
11
       - "/api/v1/products"
12
       - "/api/v1/orders/"
       per_client_ip: true
13
    req_threshold: 15
14
15
   timeslice: 60000
    throttle_action: "RESPOND"
16
17
    logpackets:
       logexpression: "http.req.url"
18
19
       loglevel: "INFORMATIONAL"
```

#### Note:

You can initiate multiple Kubernetes objects for different paths that require different rate limit configurations.

After you have defined the policy configuration, deploy the .yaml file using the following command:

1 root@master:~#kubectl create -f ratelimit-example1.yaml
2 ratelimit.citrix.com/throttle-req-per-clientip created

The NetScaler Ingress Controller applies the policy configuration on the Ingress NetScaler device.

#### **Limit API requests to calender APIs**

Consider a scenario wherein you want to define a rate-based policy in a NetScaler to limit the API requests (GET or POST) to five requests from each API client identified using the HTTP header X-API -Key to the calender APIs. Create a .yaml file called ratelimit-example2.yaml and use the appropriate CRD attributes to define the rate-based policy as follows:

```
1 apiVersion: citrix.com/v1beta1
2 kind: ratelimit
3 metadata:
4 name: throttle-calendarapi-perapikey
5 spec:
   servicenames:
6
7

    frontend

8
   selector_keys:
    basic:
9
10
         path:
11
           - "/api/v1/calender"
         method:
12
          - "GET"
13
           - "POST"
14
15
         header_name: "X-API-Key"
```

After you have defined the policy configuration, deploy the .yaml file using the following command:

```
1 root@master:~#kubectl create -f ratelimit-example2.yaml
2 ratelimit.citrix.com/throttle-req-per-clientip created
```

The NetScaler Ingress Controller applies the policy configuration on the Ingress NetScaler device.

# **Use Rewrite and Responder policies in Kubernetes**

#### February 29, 2024

In a Kubernetes environment, to deploy specific layer 7 policies to handle scenarios such as:

- Redirecting HTTP traffic to a specific URL
- Blocking a set of IP addresses to mitigate DDoS attacks
- Imposing HTTP to HTTPS

Requires you to add appropriate libraries within the microservices and manually configure the policies. Instead, you can use the Rewrite and Responder features provided by the Ingress NetScaler device to deploy these policies.

NetScaler provides Kubernetes CustomResourceDefinitions (CRDs) that you can use with the NetScaler Ingress Controller to automate the configurations and deployment of these policies on the NetScalers used as Ingress devices.

The Rewrite and Responder CRD provided by NetScaler is designed to expose a set of tools used in front-line NetScalers. Using these functionalities you can rewrite the header and payload of ingress and egress HTTP traffic as well as respond to HTTP traffic on behalf of a microservice.

Once you deploy the Rewrite and Responder CRD in the Kubernetes cluster. You can define extensive rewrite and responder policies using datasets, pat sets, and string maps and also enable audit logs for statistics on the ingress device. For more information on the rewrite and responder policy feature provided by NetScaler ADC, see Rewrite policy and Responder policy.

Note:

The Rewrite and Responder CRD is not supported for OpenShift routes. You can use OpenShift

ingress to use Rewrite and Responder CRD.

# **Deploy the NetScaler Rewrite and Responder CRD**

The NetScaler Rewrite and Responder CRD deployment YAML file: rewrite-responder-policiesdeployment.yaml.

Note:

Ensure that you do not modify the deployment YAML file.

#### Deploy the CRD, using the following command:

```
1 kubectl create -f rewrite-responder-policies-deployment.yaml
```

For example,

```
    root@master:~# kubectl create -f rewrite-responder-policies-deployment.
yaml
    customresourcedefinition.apiextensions.k8s.io/rewritepolicies.citrix.
com created
```

# **Rewrite and Responder CRD attributes**

The **CRD** provides attributes for the various options required to define the rewrite and responder policies. Also, it provides attributes for dataset, pat set, string map, and audit logs to use within the rewrite and responder policies. These CRD attributes correspond to **NetScaler** command and attribute respectively.

## **Rewrite policy**

The following table lists the **CRD** attributes that you can use to define a rewrite policy. Also, the table lists the corresponding NetScaler command and attributes.

CRD attribute	NetScaler command	NetScaler attribute	
rewrite-criteria	Add rewrite policy	rule	
default-action	Add rewrite policy	undefAction	
operation	Add rewrite action	type	
target	Add rewrite action	target	
modify-expression	Add rewrite action	stringBuilderExpr	

CRD attribute	NetScaler command	NetScaler attribute
multiple-occurence-modify	Add rewrite action	Search
additional-multiple-occurence- modify	Add rewrite action	RefineSearch
Direction	Bind lb vserver	Туре

# **Responder policy**

The following table lists the **CRD** attributes that you can use to define a responder policy. Also, the table lists the corresponding NetScaler command and attributes.

CRD attribute	NetScaler command	NetScaler attribute
Redirect	Add responder action	Type (the value of type)
url	Add responder action	Target
redirect-status-code	Add responder action	responseStatusCode
redirect-reason	Add responder action	reasonPhrase
Respond-with	Add responder action	Type (the value of type)
http-payload-string	Add responder action	Target
Noop	Add responder policy	Action (the value of action)
Reset	Add responder policy	Action (the value of action)
Drop	Add responder policy	Action (the value of action)
Respond-criteria	Add responder policy	Rule
Default-action	Add responder policy	undefAction

## Audit log

The following table lists the **CRD** attributes provide to enable audit log within the rewrite or responder policies. Also, the table lists the corresponding NetScaler command and attributes.

CRD attribute	attribute NetScaler command	
Logexpression	Add audit message action	stringBuilderExpr

RD attribute NetScaler command		NetScaler attribute
Loglevel	Add audit message action	Loglevel

#### Dataset

The following table lists the **CRD** attributes for dataset that you can use within the rewrite or responder policies. Also, the table lists the corresponding NetScaler command and attributes.

CRD attribute	NetScaler command	NetScaler attribute
Name	Add policy dataset	Name
Туре	Add policy dataset	Туре
Values	Bind policy dataset	Value

#### Patset

CRD attribute	NetScaler command	NetScaler attribute
Name	Add policy patset	Name
Values	Bind policy patset	string

## String map

CRD attribute	NetScaler command	NetScaler attribute
Name	Add policy stringmap	Name
Кеу	Bind policy stringmap	Key
Value	Bind policy stringmap	Value

# **Goto-priority-expression**

The following table provides information about the **goto**-priority-expression attribute, which is a CRD attribute for binding a group of multiple consecutive policies to services.

CRD attribute	NetScaler command	NetScaler attribute	Supported values	Default value
goto-priroty- expression	Bind lb vserver	gotoPriorityExp	pressi <b>&amp;r</b> EXT and END	End

For more information on how to use the **goto**-priority-expression attribute, see the example Modify strings and host name in the requested URL

# How to write a policy configuration

After you have deployed the CRD provided by NetScaler in the Kubernetes cluster, you can define the policy configuration in a.yaml file. In the .yaml file, use rewritepolicy in the kind field and based on your requirement add any of the following individual sections in spec for policy configuration.

- rewrite-policy To define rewrite policy configuration.
- responder-policy To define responder policy configuration.
- logpackets To enable audit logs.
- dataset To use a data set for extensive policy configuration.
- patset To use a pat set for extensive policy configuration.
- stringmaps To use string maps for extensive policy configuration.

In these sections, you need to use the CRD attributes provided for the respective policy configuration (rewrite or responder) to define the policy.

Also, in the spec section, you need to include a rewrite-policies section to specify the service or services to which the policy must be applied. For more information, see Sample policy configurations.

After you deploy the .yaml file, the NetScaler Ingress Controller applies the policy configuration on the Ingress NetScaler device.

## Guidelines for the policy configuration

- If the CRD is associated with a namespace then, by default, the policy is applied to the services associated with the namespace. For example, if you have the same service name associated with multiple namespaces, then the policy is applied to the service that belongs to the namespace associated with the CRD.
- If you have defined multiple policies in a single . yaml file then the first policy configuration defined in the file takes priority and the subsequent policy configurations is applied as per the

sequence. If you have multiple policies defined in different files then the first policy configuration defined in the file that you deployed first takes priority.

## Guidelines for the usage of Goto-priority-expression

- The rewrite and responder policies can be combined as multiple groups using the NEXT keyword within the **goto**-priority-expression field.
- When the **goto**-priority-expression field is NEXT within the current policy and if the current policy evaluates to True, the next policy in the group is executed and the flow moves to the next consecutive policies unless the **goto**-priority-expression field points to END.
- When the current policy evaluates to FALSE, the **goto**-priority-expression has no impact, as the policy execution stops at the current policy.
- The rewrite or responder policy group within the rewrite or responder policies begins with the policy assigned with **goto**-priority-expression as NEXT and includes all the consecutive policies until the **goto**-priority-expression field is assigned to END.
- When you group rewrite or responder policies using **goto**-priority-expression, the service names bound to the policies within the group should be the same.
- The last policy within the rewrite-policies or responder-policies should always have the **goto**-priority-expression as END.
- If the **goto**-priority-expression field is not specified for a policy, the default value of END is assigned to **goto**-priority-expression.

## Note:

For more information on how to use the **goto**-priority-expression field, see the example Modify strings and host name in the requested URL.

# Create and verify a rewrite and responder policy

Consider a scenario where you want to define a policy in NetScaler to rewrite all the incoming URLs to **new**-url-**for**-the-application and send it to the microservices. Create a .yaml file called target-url-rewrite.yaml and use the appropriate CRD attributes to define the rewrite policy as follows:

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4 name: targeturlrewrite
5 spec:
6 rewrite-policies:
```

```
- servicenames:
7
           - citrix-svc
8
9
         logpackets:
           logexpression: "http.req.url"
10
11
           loglevel: INFORMATIONAL
         rewrite-policy:
12
13
           operation: replace
           target: 'http.req.url'
14
           modify-expression: '"new-url-for-the-application"'
15
           comment: 'Target URL Rewrite - rewrite the url of the HTTP
               request'
17
           direction: REQUEST
            rewrite-criteria: 'http.req.is_valid'
18
```

After you have defined the policy configuration, deploy the .yaml file using the following command:

```
1 kubectl create -f target-url-rewrite.yaml
```

After you deploy the .yaml file, the NetScaler Ingress Controller applies the policy configuration on the Ingress NetScaler device.

On the master node in the Kubernetes cluster, you can verify the status of the applied rewrite policy CRD using the following command:

```
1 Kubectl get rewritepolicies.citrix.com targeturlrewrite
```

You can view the status as follows:

```
    kubectl get rewritepolicies.citrix.com targeturlrewrite
    NAME STATUS MESSAGE
    targeturlrewrite Success CRD Activated
```

If there are issues while creating or applying the CRD, the same can be debugged using the citrix-k8singress-controller logs.

1 kubectl logs citrixingresscontroller

Also, you can verify whether the configuration is applied on the NetScaler by using the following steps.

- 1. Log on to the NetScaler command-line.
- 2. Use the following command to verify if the configuration is applied to the NetScaler:

```
1 show run | grep `lb vserver`
2 add lb vserver k8s-citrix_default_80_k8s-citrix-svc_default_80_svc
HTTP 0.0.0.0 0 -persistenceType NONE -cltTimeout 180
3 bind lb vserver k8s-citrix_default_80_k8s-citrix-
svc_default_80_svc k8s-citrix_default_80_k8s-citrix-
svc_default_80_svc
```

```
4 bind lb vserver k8s-citrix_default_80_k8s-citrix-
    svc_default_80_svc -policyName
    k8s_crd_rewritepolicy_rwpolicy_targeturlrewrite_0_default -
    priority 100300076 -gotoPriorityExpression END -type REQUEST
```

You can verify that the policy k8s\_crd\_rewritepolicy\_rwpolicy\_targeturlrewrite\_0\_defa is bound to the load balancing virtual server.

# Sample policy configurations

#### **Responder policy configuration**

Following is a sample responder policy configuration (block-list-urls.yaml)

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4 name: blocklisturls
5 spec:
6 responder-policies:
7
     - servicenames:
           - citrix-svc
8
9
         responder-policy:
          respondwith:
11
            http-payload-string: '"HTTP/1.1 401 Access denied"'
12
           respond-criteria: 'http.req.url.equals_any("blocklistUrls")'
13
           comment: 'Blocklist certain Urls'
14
15
   patset:
16
17
       - name: blocklistUrls
18
         values:
19
          - '/app1'
20
           - '/app2'
21
           - '/app3'
```

In this example, if NetScaler receives any URL that matches the /app1, /app2, or /app3 strings defined in the patset, NetScaler blocks the URL.

#### Policy with audit logs enabled

Following is a sample policy with audit logs enabled (block-list-urls-audit-log.yaml).

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4 name: blocklisturls
5 spec:
```

```
6
     responder-policies:
7
       - servicenames:
           - citrix-svc
8
9
         logpackets:
           logexpression: "http.req.url"
10
11
           loglevel: INFORMATIONAL
         responder-policy:
12
           respondwith:
13
14
             http-payload-string: '"HTTP/1.1 401 Access denied"'
15
           respond-criteria: 'http.req.url.equals_any("blocklistUrls")'
16
           comment: 'Blocklist certain Urls'
17
18
19
     patset:
       - name: blocklistUrls
20
21
         values:
22
           - '/app1'
           - '/app2'
23
           - '/app3'
24
```

#### **Multiple policy configurations**

You can add multiple policy configurations in a single .yaml file and apply the policies to the NetScaler device. You need add separate sections for each policy configuration (multi-policy-config.yaml).

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4
   name: multipolicy
5 spec:
6 responder-policies:
     - servicenames:
7
8
           - citrix-svc
         responder-policy:
9
           redirect:
            url: '"www.citrix.com"'
12
           respond-criteria: 'client.ip.src.TYPECAST_text_t.equals_any("
              redirectIPs")'
           comment: 'Redirect IPs to citrix.com'
13
14
       - servicenames:
15
           - citrix-svc
16
         responder-policy:
17
           redirect:
             url: 'HTTP.REQ.HOSTNAME+http.req.url.
18
                MAP_STRING_DEFAULT_TO_KEY("modifyurls")'
           respond-criteria: 'http.req.is_valid'
19
20
           comment: 'modify specific URLs'
21
22
     rewrite-policies:
23
       - servicenames:
```

```
- citrix-svc
24
25
         rewrite-policy:
26
           operation: insert_http_header
            target: 'sessionID'
27
            modify-expression: '"48592th42gl24456284536tgt2"'
28
29
           comment: 'insert SessionID in header'
            direction: RESPONSE
31
            rewrite-criteria: 'http.res.is_valid'
32
33
34
     dataset:
       - name: redirectIPs
37
         type: ipv4
38
         values:
           - 10.1.1.100
39
40
            - 1.1.1.1 - 1.1.1.100
41
            - 2.2.2/10
42
43
     stringmap:
       - name: modifyurls
44
         comment: Urls to be modified string
45
46
         values:
            - key: '"/app1/"'
47
             value: '"/internal-app1/"'
48
            - key: '"/app2/"'
49
50
             value: '"/internal-app2/"'
```

The example contains two responder policies and a rewrite policy, based on these policies NetScaler performs the following:

- Any requests that match the client IP addresses specified in the redirectIPs dataset, that is, 10.1.1.100, IP addresses in the range 1.1.1.1 1.1.1.100 and IP addresses in the subnet 2.2.2/10, is redirected to www.citrix.com.
- Any incoming URL with strings provided in the modifyurls stringmap is modified to the value provided in the stringmap. For example, if the incoming URL has the string /app1/ is modified to /internal-app1/
- Adds a session ID as a new header in the response to the client.

## Example use cases

## Add response headers

When the requested URL from the client contains /citrix-app/, you can add the following headers in the HTTP response from the microservices to the client using a rewrite policy:

• Client source port to the header

- Server destination IP address
- random HTTP header

The following sample rewrite policy definition adds these headers to the HTTP response from the microservices to the client:

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4 name: addresponseheaders
5 spec:
6 rewrite-policies:
    - servicenames:
7
         - frontend
8
9
       rewrite-policy:
10
         operation: insert_before_all
11
         target: http.res.full_header
         modify-expression: '"\r\nx-port: "+client.tcp.srcport+"\r\nx-ip
12
            :"+client.ip.dst+"\r\nx-new-dummy-header: Sending_a_gift"
         multiple-occurence-modify: 'text("\r\n\r\n")'
13
         comment: 'Response header rewrite'
14
15
         direction: RESPONSE
         rewrite-criteria: 'http.req.url.contains("/citrix-app/")'
```

Create a YAML file (add\_response\_headers.yaml) with the rewrite policy definition and deploy the YAML file using the following command:

1 kubectl create -f add\_response\_headers.yaml

You can verify the HTTP header added to the response as follows:

```
1 $ curl -vvv http://app.cic-citrix.org/citrix-app/
2 * Trying 10.102.33.176...
3 * TCP_NODELAY set
4 * Connected to app.cic-citrix.org (10.102.33.176) port 80 (#0)
5 > GET /citrix-app/ HTTP/1.1
6 > Host: app.cic-citrix.org
7
  > User-Agent: curl/7.54.0
8 > Accept: */*
9
  >
10 < HTTP/1.1 200 OK
11 < Server: nginx/1.8.1</pre>
12 < Date: Fri, 29 Mar 2019 11:14:04 GMT
13 < Content-Type: text/html</pre>
14 < Transfer-Encoding: chunked
15 < Connection: keep-alive
16 < X-Powered-By: PHP/5.5.9-1ubuntu4.14</pre>
17 < x-port: 22481 ========> NEW RESPONSE HEADER
18 < x-ip:10.102.33.176 ========> NEW RESPONSE HEADER
19
  < x-new-dummy-header: Sending_a_gift ==========> NEW RESPONSE
      HEADER
20 <
```

```
21 <html>
22 <head>
23 <title> Front End App - v1 </title>
24
25
26 TRIMMED
27 .....
```

## Add custom header to the HTTP response packet

Using a rewrite policy, you can add custom headers in the HTTP response from the microservices to the client.

The following sample rewrite policy definition adds a custom header to the HTTP response from the microservices to the client:

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4 name: addcustomheaders
5 spec:
6 rewrite-policies:
7
       - servicenames:
           - frontend
8
9
         rewrite-policy:
10
           operation: insert_before_all
           target: http.res.full_header
11
           modify-expression: '"\r\nx-request-time:"+sys.time+"\r\nx-using
12
              -citrix-ingress-controller: true"'
           multiple-occurence-modify: 'text("\r\n\r\n")'
13
           comment: 'Adding custom headers'
14
           direction: RESPONSE
15
16
           rewrite-criteria: 'http.req.is_valid'
```

Create a YAML file (add\_custom\_headers.yaml) with the rewrite policy definition and deploy the YAML file using the following command:

1 kubectl create -f add\_custom\_headers.yaml

You can verify the custom HTTP header added to the response as follows:

```
1 $ curl -vvv http://app.cic-citrix.org/
2 * Trying 10.102.33.176...
3 * TCP_NODELAY set
4 * Connected to app.cic-citrix.org (10.102.33.176) port 80 (#0)
5 > GET / HTTP/1.1
6 > Host: app.cic-citrix.org
7 > User-Agent: curl/7.54.0
8 > Accept: */*
9 >
```

```
10 < HTTP/1.1 200 OK
11 < Server: nginx/1.8.1</pre>
12 < Date: Fri, 29 Mar 2019 12:15:09 GMT
13 < Content-Type: text/html</pre>
14 < Transfer-Encoding: chunked
15 < Connection: keep-alive
16 < X-Powered-By: PHP/5.5.9-1ubuntu4.14</pre>
17 < x-request-time:Fri, 29 Mar 2019 13:27:40 GMT =======> NEW
      HEADER ADDED
18 < x-using-citrix-ingress-controller: true =========> NEW HEADER
      ADDED
19 <
20 <html>
21 <head>
22 <title> Front End App - v1 </title>
23 <style>
24
25 TRIMMED
26 .....
```

#### **Replace host name in the request**

You can define a rewrite policy as shown in the following example YAML (http\_request\_modify\_prefixasp .yaml) to replace the host name in an HTTP request as per your requirement:

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
   name: httpheadermodifyretainprefix
4
5 spec:
6 rewrite-policies:
7
     - servicenames:
8
           - frontend
9
         rewrite-policy:
           operation: replace_all
           target: 'http.req.header("host")'
12
           modify-expression: '"citrix-service-app"'
13
           multiple-occurence-modify: 'text("app.cic-citrix.org")'
           comment: 'HTTP header rewrite of hostname'
14
15
           direction: REQUEST
           rewrite-criteria: 'http.req.is_valid'
16
```

Create a YAML file (http\_request\_modify\_prefixasprefix.yaml) with the rewrite policy definition and deploy the YAML file using the following command:

1 kubectl create -f http\_request\_modify\_prefixasprefix.yaml

You can verify the policy definition using the curl command. The host name in the request is replaced with the defined host name.

#### 1 curl http://app.cic-citrix.org/prefix/foo/bar

Output:
▼ Hypertext Transfer Protocol
<pre>GET /prefix/foo/bar HTTP/1.1\r\n</pre>
[Expert Info (Chat/Sequence): GET /prefix/foo/bar HTTP/1.1\r\n]
Request Method: GET
Request URI: /prefix/foo/bar
Request Version: HTTP/1.1
Host: citrix-service-app\r\n
User-Agent: curl/7.54.0\r\n
Accept: */*\r\n
\r\n
<pre>[Full request URI: http://citrix-service-app/prefix/foo/bar]</pre>
[HTTP request 1/1]

#### Modify the application root

You can define a rewrite policy to modify the application root if the existing application root is /.

The following sample rewrite policy modifies / to /citrix-approot/ in the request URL:

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4 name: httpapprootrequestmodify
5 spec:
   rewrite-policies:
6
7
     - servicenames:
          - frontend
8
        rewrite-policy:
9
10
         operation: replace
          target: http.req.url
11
          modify-expression: '"/citrix-approot/"'
12
          comment: 'HTTP app root request modify'
13
14
          direction: REQUEST
15
          rewrite-criteria: http.req.url.eq("/")
```

Create a YAML file (http\_approot\_request\_modify.yaml) with the rewrite policy definition and deploy the YAML file using the following command:

kubectl create -f http\_approot\_request\_modify.yaml

Using the curl command, you can verify if the application root is modified as per your requirement:

```
1 curl -vvv http://app.cic-citrix.org/
```

Output:

🔻 Hypertext Transfer Protocol
▶ GET /citrix-approot/ HTTP/1.1\r\n
Host: app.cic-citrix.org\r\n
User-Agent: curl/7.54.0\r\n
Accept: */*\r\n
\r\n
<pre>[Full request URI: http://app.cic-citrix.org/citrix-approot/]</pre>
[HTTP request 1/1]
[Response in frame: 126]

#### Modify the strings in the requested URL

You can define a rewrite policy to modify the strings in the requested URL as per your requirement.

The following sample rewrite policy replaces the strings something to simple in the requested URL:

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4 name: httpurlreplacestring
5 spec:
6 rewrite-policies:
   - servicenames:
7
        - frontend
8
      rewrite-policy:
9
        operation: replace_all
10
11
        target: http.req.url
        modify-expression: '"/"'
12
        multiple-occurence-modify: 'regex(re~((^(\/something\/))))(^\/
13
           something$))~)'
        comment: 'HTTP url replace string'
14
15
        direction: REQUEST
16
        rewrite-criteria: http.req.is_valid
```

Create a YAML file (http\_url\_replace\_string.yaml) with the rewrite policy definition and deploy the YAML using the following command:

1 kubectl create -f http\_url\_replace\_string.yaml

You can verify the policy definition using a curl request with the string something. The string something is replaced with the string simple as shown in the following examples:

Example 1:

1 curl http://app.cic-citrix.org/something/simple/citrix

Output:

```
Hypertext Transfer Protocol
GET /simple/citrix HTTP/1.1\r\n
Host: app.cic-citrix.org\r\n
User-Agent: curl/7.54.0\r\n
Accept: */*\r\n
\r\n
[Full request URI: http://app.cic-citrix.org/simple/citrix]
[HTTP request 1/1]
[Response in frame: 17]
```

#### Example 2:

```
1 curl http://app.cic-citrix.org/something
```

Or,

```
1 curl http://app.cic-citrix.org/something/
```

Output:

```
    Hypertext Transfer Protocol
    GET / HTTP/1.1\r\n
    Host: app.cic-citrix.org\r\n
    User-Agent: curl/7.54.0\r\n
    Accept: */*\r\n
    \r\n
    [Full request URI: http://app.cic-citrix.org/]
    [HTTP request 1/1]
    [Response in frame: 32]
```

#### Add the X-Forwarded-For header within an HTTP request

You can define a rewrite policy as shown in the following example YAML (http\_x\_forwarded\_for\_insert .yaml) to add the X-Forwarded-For header within an HTTP request:

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4 name: httpxforwardedforaddition
5 spec:
   rewrite-policies:
6
7
      - servicenames:
           - frontend
8
9
         rewrite-policy:
10
           operation: insert_http_header
           target: X-Forwarded-For
           modify-expression: client.ip.src
12
13
           comment: 'HTTP Initial X-Forwarded-For header add'
14
           direction: REQUEST
```

15	rewrite-criteria: 'HTTP.REQ.HEADER("X-Forwarded-For").EXISTS.
	NOTI
	NOT
16	
10	
17	- servicenames:
10	frontond
TQ	- Troncend
19	rewrite-policy:
20	operation: replace
21	target. HTTP REO HEADER("X-Eorwarded-Eor")
~ 1	cargee. Intra Reg. newber( x for warded for )
22	<pre>modify-expression: 'HTTP.REQ.HEADER("X-Forwarded-For").APPEND</pre>
	(II II) ADDEND (CLIENT ID CDC) I
	(",").APPEND(CLIENT.IP.SRC)
23	comment: 'HTTP Append X-Forwarded-For TPs'
20	
24	direction: REQUEST
25	rowrite eriteria. LUTTO DEC NEADED(NY Ferwarded FerN) EVICES
20	rewrite=criteria: "nile.key.HEADER("X=Forwarded=For").EXISTS"

Create a YAML file (http\_x\_forwarded\_for\_insert.yaml) with the rewrite policy definition and deploy the YAML file using the following command:

1 kubectl create -f http\_x\_forwarded\_for\_insert.yaml

Using the curl command you can verify the HTTP packet with and without the X-Forwarded-For header.

Example: Output of the HTTP request packet without X-Forwarded-For header:

```
1 curl http://app.cic-citrix.org/
```

Output:

```
    Hypertext Transfer Protocol
    GET / HTTP/1.1\r\n
    Host: app.cic-citrix.org\r\n
    User-Agent: curl/7.54.0\r\n
    Accept: */*\r\n
    X-Forwarded-For: 10.150.16.22\r\n
    \r\n
    [Full request URI: http://app.cic-citrix.org/]
    [HTTP request 1/1]
    [Response in frame: 7]
```

Example: Output of the HTTP request packet with X-Forwarded-For header:

```
1 curl curl --header "X-Forwarded-For: 1.1.1.1" http://app.cic-citrix.
org/
```

Output:

```
    Hypertext Transfer Protocol
    GET / HTTP/1.1\r\n
    Host: app.cic-citrix.org\r\n
    User-Agent: curl/7.54.0\r\n
    Accept: */*\r\n
    X-Forwarded-For: 1.1.1.1,10.150.16.22\r\n
    \r\n
    [Full request URI: http://app.cic-citrix.org/]
    [HTTP request 1/1]
    [Response in frame: 65]
```

# Redirect HTTP request to HTTPS request using responder policy

You can define a responder policy definition as shown in the following example YAML(http\_to\_https\_redirect.yaml) to redirect HTTP requests to HTTPS request:

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
  name: httptohttps
4
5 spec:
6 responder-policies:
       - servicenames:
7
           - frontend
8
9
         responder-policy:
           redirect:
             url: '"https://" +http.req.HOSTNAME.SERVER+":"+"443"+http.req
11
                .url'
12
           respond-criteria: 'http.req.is_valid'
           comment: 'http to https'
13
```

Create a YAML file (http\_to\_https\_redirect.yaml) with the responder policy definition and deploy the YAML file using the following command:

1 kubectl create -f http\_to\_https\_redirect.yaml

You can verify if the HTTP request is redirected to HTTPS as follows:

Example 1:



Example 2:

1 \$ curl -vvv http://app.cic-citrix.org/simple
2 \* Trying 10.102.33.176...

```
3 * TCP_NODELAY set
4 * Connected to app.cic-citrix.org (10.102.33.176) port 80 (#0)
5 > GET /simple HTTP/1.1
6 > Host: app.cic-citrix.org
7 > User-Agent: curl/7.54.0
8 > Accept: */*
9 >
10 < HTTP/1.1 302 Found : Moved Temporarily
11 < Location: https://app.cic-citrix.org:443/simple ======>
Redirected to HTTPS
12 < Connection: close
13 < Cache-Control: no-cache
14 < Pragma: no-cache
15 <
16 * Closing connection 0
```

# Modify strings and host name in the requested URL

This example shows the usage of the **goto**-priority-expression attribute. The guidelines for usage of the **goto**-priority-expression field can be found at [How to write a policy configuration. This example modifies the URL http://www.citrite.org/something/simple/citrix to http://app.cic-citrix.org/simple/citrix.

Two rewrite policies are written to modify the URL:

- Rewrite policy 1: This policy is used to modify the host name www.citrite.org to app.cic -citrix.org.
- Rewrite Policy 2: This policy is used to modify the url /something/simple/citrix to / simple/citrix

You can bind the two policies using the **goto**-priority-expression attribute as shown in the following YAML:

```
1 apiVersion: citrix.com/v1
2
  kind: rewritepolicy
3 metadata:
     name: hostnameurlrewrite
4
5
     spec:
    rewrite-policies:
6
7
         - servicenames:
8
             - citrix-svc
9
           goto-priority-expression: NEXT
10
           rewrite-policy:
11
             operation: replace_all
             target: 'http.req.header("host")'
13
             modify-expression: '"app.cic-citrix.org"'
14
             multiple-occurence-modify: 'text("www.citrite.org")'
15
             comment: 'HTTP header rewrite of hostname'
16
             direction: REQUEST
```

```
rewrite-criteria: 'http.req.is_valid.and(HTTP.REQ.HOSTNAME.EQ
17
                 ("www.citrite.org"))'
18
         - servicenames:
19
             - citrix-svc
20
           goto-priority-expression: END
21
           rewrite-policy:
             operation: replace_all
22
             target: http.req.url
23
             modify-expression: '"/"'
24
             multiple-occurence-modify: 'regex(re~((^(\/something\/))))(^\/
                 something$))~)'
             comment: 'HTTP url replace string'
26
             direction: REQUEST
              rewrite-criteria: 'http.req.is_valid.and(HTTP.REQ.HOSTNAME.EQ
28
                 ("www.citrite.org"))'
```

**Verification** You can verify whether the following curl request http://www.citrite.org /something/simple/citrix is modified to http://app.cic-citrix.org/simple/ citrix.

Example: Modifying the requested URL

1 curl http://www.citrite.org/something/simple/citrix

Modified host name and URL for the requested URL is present in the image shown as follows:

```
• Hypertext Transfer Protocol
• GET /simple/citrix HTTP/1.1\r\n
Host: app.cic-citrix.org\r\n
User-Agent: curl/7.47.0\r\n
Accept: */*\r\n
\r\n
[Full request URI: http://app.cic-citrix.org/simple/citrix]
[HTTP request 1/1]
```

## **HTTP callout**

An HTTP callout allows the NetScaler to generate and send an HTTP or HTTPS request to an external server as part of the policy evaluation and take appropriate action based on the response obtained from the external server. You can use the rewrite and responder CRD to initiate HTTP callout requests from the NetScaler. For more information, see the HTTP callout documentation.

# **Related** articles

• Feature Documentation

- NetScaler Rewrite Feature Documentation
- NetScaler Responder Feature Documentation
- Developer Documentation
  - NetScaler Rewrite Policy
  - NetScaler Rewrite Action
  - NetScaler Responder Policy
  - NetScaler Responder Action
  - NetScaler Audit Message Action
  - NetScaler Policy Dataset

# Remote content inspection or content transformation service using ICAP

#### November 13, 2024

The Internet Content Adaptation Protocol (ICAP) is a simple lightweight protocol for running a valueadded transformation service on HTTP messages. In a typical scenario, an ICAP client forwards HTTP requests and responses to one or more ICAP servers for processing. The ICAP servers perform content transformation on the requests and send back responses with an appropriate action to take on the request or response.

In a NetScaler setup, NetScaler acts as an ICAP client interoperating with third-party ICAP servers, such as antimalware and Data Loss Protection (DLP). When NetScaler receives incoming web traffic, it intercepts the traffic and uses a Content Inspection policy to evaluate if the HTTP request needs an ICAP processing. If yes, NetScaler decrypts and sends the message as a plain text to the ICAP servers. The ICAP servers run the content transformation service on the request message and send back a response to NetScaler. The modified messages can either be an HTTP request or an HTTP response. If NetScaler interoperates with multiple ICAP servers, NetScaler performs load balancing of ICAP servers. This setup helps when one ICAP server is not sufficient to handle all the traffic load. After the ICAP servers return a modified message, NetScaler forwards the modified message to the back-end origin server.

NetScaler also provides a secured ICAP service if the incoming traffic is of HTTPS type. NetScaler uses an SSL based TCP service to establish a secured connection between NetScaler and the ICAP servers.

In a Kubernetes environment, to enable ICAP on NetScaler through NetScaler Ingress Controller, NetScaler provides the ICAP Custom Resource Definition (CRD). By enabling ICAP, you can perform the following actions:

- Block URLs with a specified string
- Block a set of IP addresses to mitigate DDoS attacks
- Mandate HTTP to HTTPS

After you deploy the ICAP CRD in the Kubernetes cluster, you can define ICAP policies by using the ICAP CRD attributes and also enable audit logs for statistics on NetScaler. For more information on the ICAP feature provided by NetScaler, see ICAP for remote content inspection.

ICAP policies can be broadly classified as request modification and response modification policies.

**Request modification**: In the request modification (REQMOD) mode, NetScaler forwards the HTTP request received from its client to the ICAP server. For an example request modification CRD resource, see sample policy configurations.

**Response modification**: In the response modification (RESPMOD) mode, NetScaler sends an HTTP response to the ICAP server (the response sent by NetScaler is typically the response sent by the origin server). For an example response modification CRD resource, see sample policy configurations.

For detailed information on request and response modification policies, see ICAP for remote content inspection.

# **Deploy the NetScaler ICAP CRD**

The NetScaler ICAP CRD deployment YAML file is available in GitHub at the following location: icapcrd-deployment.yaml.

Note:

Ensure that you do not modify the deployment YAML file.

## Deploy the ICAP CRD by using the following command:

1 kubectl create -f icap-crd-deployment.yaml

For example,

# **ICAP CRD attributes**

The **ICAP CRD** provides attributes to enable content inspection and configure various options required to define the ICAP policies. These CRD attributes correspond to **NetScaler** commands and attributes respectively.

The following table lists the **CRD** attributes that you can use to define an ICAP policy. Also, the table lists the corresponding NetScaler commands and attributes.

# **Adding ICAP servers**

You can specify the ICAP servers under the icap-servers object in the ICAP CRD spec. You can specify either one or multiple ICAP servers based on your requirement.

CRD attribute	NetScaler command	NetScaler attribute	Description
ip (Mandatory)	add service < name> <ip> &lt; serviceType&gt; &lt; port&gt;</ip>	IP	ICAP server IP address
port (Mandatory)	add service < name> <ip> &lt; serviceType&gt; &lt; port&gt;</ip>	port	The port on which the ICAP server communicates
server-type (Mandatory)	add service < name> <ip> &lt; serviceType&gt; &lt; port&gt;</ip>	serviceType	The type of ICAP server. Possible values are TCP and SSL_TCP.

# Specifying the back-end services and ingress class

Specify the Services and ingressclass attributes to do the following configuration:

- Specify the back-end services for which the ICAP policy needs to be enabled.
- Specify the ICAP resources that need to be processed by the ingress controller.

CRD attribute

NetScaler command

Description

Services (Mandatory)	NA	Lists the back-end services for
		which the ICAP policy must be
		enabled.
ingressclass (Mandatory)	NA	Specifies the ingress class so
		that only the ingress controller
		associated with the specified
		ingress class processes the
		resource. Otherwise, all the
		controllers in the cluster
		process this CRD resource.

# Adding an ICAP profile

ICAP configurations for NetScaler are specified in an entity called the ICAP profile. The profile has a collection of the ICAP settings. The settings include parameters to dynamically generate an ICAP request, receive the ICAP response, and log content inspection data.

CRD attribute	NetScaler command	NetScaler attribute	Description
preconfigured-profile (Optional)	NA	NA	Names of the preconfigured ICAP profile.
direction (Mandatory)	add ns icapProfile < name> -uri < string> -Mode ( REQMOD \  RESPMOD )	Mode	ICAP mode of operation. Possible values are REQUEST and RESPONSE.
uri (Mandatory)	add ns icapProfile < name> -uri < string> -Mode ( REQMOD \  RESPMOD )	uri	URI representing the ICAP service.

CRD attribute	NetScaler command	NetScaler attribute	Description
preview	add ns icapProfile < name> -uri < string> -Mode ( REQMOD \  RESPMOD )- preview ( ENABLED \  DISABLED )]	preview	Enable or disable the preview header with ICAP request. This feature allows an ICAP server to see the beginning of a transaction. Then decide if it wants to opt out of the transaction early, instead of receiving the remainder of the request message.
preview-length	add ns icapProfile < name> -uri < string> -Mode ( REQMOD \  RESPMOD )- previewLength < positive_integer >	previewLength	Value of the Preview Header field. NetScaler uses the smaller value between this set value and the preview size received on OPTIONS.
host-header	add ns icapProfile < name> -uri < string> -Mode ( REQMOD \  RESPMOD )- hostHeader < string>	hostHeader	ICAP host header.

CRD attribute	NetScaler command	NetScaler attribute	Description
user-agent-header	add ns icapProfile < name> -uri < string> -Mode ( REQMOD \  RESPMOD )- userAgent < string>	userAgent	ICAP user agent header.
query-params	add ns icapProfile < name> -uri < string> -Mode ( REQMOD \  RESPMOD )- queryParams < string>	queryParams	Query parameters to be included with the ICAP request URI. Entered values must be in the arg=value format. For more than one parameter, add & between the parameters. For example, arg1=val1 &arg2=val2.
connection-keep-alive	add ns icapProfile < name> -uri < string> -Mode ( REQMOD \  RESPMOD )- connectionKeepA ( ENABLED \  DISABLED )	connectionKeepAlive live	Enable or disable sending Allow: 204 header in ICAP request.

CRD attribute	NetScaler command	NetScaler attribute	Description
insert-icap-headers	<pre>add ns icapProfile &lt; name&gt; -uri &lt; string&gt; -Mode ( REQMOD \  RESPMOD )- insertICAPHeaders <string></string></pre>	insertICAPHeaders	Insert custom ICAP headers in the ICAP request to send to the ICAP server. The headers can be static or can be dynamically constructed by using PI policy expression. For example, to send static user agent and client's IP addresses, the expression can be specified as "User-Agent: NS -ICAP-Client/V1 .0r0-Client-IP: "+CLIENT.IP. SRC+"r0. NetScaler does not check the validity of the specified header name-value. You must manually validate the specified header syntax.

CRD attribute	NetScaler command	NetScaler attribute	Description
insert-http-request	<pre>add ns icapProfile &lt; name&gt; -uri &lt; string&gt; -Mode ( REQMOD \  RESPMOD )- insertHTTPRequest <string></string></pre>	insertHTTPRequest	Exact HTTP request in the form of an expression, which NetScaler encapsulates and sends to the ICAP server. If you set this parameter, the ICAP request is sent using this header only. This attribute can be used when the HTTP header is not available to send or the ICAP server only needs part of the incoming HTTP request. NetScaler does not check the validity of this request. You must manually
req-timeout	<pre>add ns icapProfile &lt; name&gt; -uri &lt; string&gt; -Mode ( REQMOD \  RESPMOD )- reqTimeout &lt; positive_integer &gt;</pre>	reqTimeout	Time, in seconds, within which the remote server should respond to the ICAP request. If NetScaler does not receive the full response within this time, the specified request timeout action is performed. When the value is configured as zero, the functionality is disabled.

CRD attribute	NetScaler command	NetScaler attribute	Description
req-timeout-action	<pre>add ns icapProfile &lt; name&gt; -uri &lt; string&gt; -Mode ( REQMOD \  RESPMOD )- reqTimeoutAction</pre>	reqTimeoutAction	Action to perform if the virtual server representing the remote service does not respond within the timeout value configured. The supported actions are: <b>BYPASS</b> - Ignores the remote server response and sends the request/response to the client/server. If the ICAP response with encapsulated headers is not received within the request-timeout value configured, this option ignores the remote ICAP server response and sends the full request/response to server/client. <b>DROP</b> - Drops the request without sending a response to the user. <b>RESET</b> - Resets the client connection by closing it. The client may then resend the

# Content inspection policy and action

After you enable the content inspection feature, you must add an ICAP action for handling the ICAP request information. The ICAP profile and services, or load balancing virtual server that are created

are bound to the ICAP action.

CRD attribute	NetScaler command	NetScaler attribute	Description
content-inspection- criteria	add contentInspection policy <name> -rule &lt; expression&gt; - action <string></string></name>	rule	Expression that the policy uses to determine whether to perform the specified action.
default-action	<pre>add contentInspection action <name> -type ICAP - serverName &lt; string&gt; - icapProfileName <string> - ifserverdown &lt; if-server-down&gt;</string></name></pre>	undefAction	Action to perform if the result of the policy evaluation is undefined (UNDEF). An UNDEF event indicates an internal error condition. Only the built-in actions before this can be used.
log-action	<pre>add contentinspection policy <name> -rule &lt; expression&gt; - action <string> -logAction &lt; string&gt;</string></name></pre>	logAction	Name of the messagelog action to use for requests that match this policy.

CRD attribute	NetScaler command	NetScaler attribute	Description
operation	add ContentInspection action <name> -type ICAP - serverip <ip> - serverport &lt; port&gt; - icapProfileName <string></string></ip></name>	type	The type of operation this ICAP action performs. The following actions are available to configure: <b>ICAP</b> - Forwards the incoming request or response to an ICAP server for modification <b>INLINEINSPECTION</b> - Forwards the incoming or outgoing packets to IPS server for intrusion prevention. <b>MIRROR</b> - Forwards cloned packets for intrusion detection. <b>NOINSPECTION</b> - Does not forward incoming and outgoing packets to

current and further incoming packets on this transaction.

CRD attribute	NetScaler command	NetScaler attribute	Description
server-failure-action	add contentInspection action <name> -type ICAP - serverName &lt; string&gt; - icapProfileName <string> - ifserverdown &lt;&gt;</string></name>	ifserverdown	<ul> <li>Action to perform if the virtual server representing the remote service is not up. The supported actions are: <b>RESET</b> - Resets the client connection by closing it. The client program, such as a browser, handles this action and may inform the user. The client may then resend the request if desired.</li> <li><b>DROP</b> - Drops the request without sending a response to the user. <b>CONTINUE</b> - Bypasses the content inspection and resumes the traffic flow to client or server</li> </ul>

# Goto-priority-expression

The following table provides information about the **goto**-priority-expression attribute, which is a CRD attribute for binding a group of policies to services.

CRD attribute	NetScaler command	NetScaler attribute	Supported values	Default value
goto-priority- expression	Bind lb vserver	gotoPriorityExpress	i <b>ðrE</b> XT and END	End

# How to write a policy configuration

After deploying the ICAP CRD specification provided by NetScaler in the Kubernetes cluster, you can define the policy configuration in a .yaml file. In the .yaml file, use icappolicy in the kind field and based on your requirement specify values to the attributes for policy configuration. For information on mandatory and optional parameters and their descriptions, see the preceding section.

After you deploy the .yaml file, NetScaler Ingress Controller applies the policy configuration on NetScaler.

# Guidelines for the policy configuration

- If the CRD is associated with a namespace then, by default, the policy is applied to the services associated with the namespace. For example, if you have the same service name associated with multiple namespaces, then the policy is applied to the service that belongs to the namespace associated with the CRD.
- If you have defined multiple policies in a single . yaml file, then the first policy configuration defined in the file takes priority and the subsequent policy configurations are applied as per the sequence. If you have multiple policies defined in different files then the first policy configuration defined in the file that you deployed first takes priority.

## Guidelines for the usage of Goto-priority-expression

- The ICAP policies can be combined as multiple groups by using the NEXT keyword within the **goto**-priority-expression field.
- When the **goto**-priority-expression field is NEXT within the current policy and if the current policy evaluates to True, the next policy in the group is executed. This process continues with the subsequent policies until a policy with the **goto**-priority-expression field is set to END.
- When the current policy evaluates to FALSE, the **goto**-priority-expression has no impact, as the policy execution stops at the current policy.
- The ICAP policy group within the ICAP policies begins with the policy assigned with gotopriority-expression as NEXT and includes all the consecutive policies until the goto -priority-expression field is assigned to END.
- When you group ICAP policies by using **goto**-priority-expression, the service names bound to the policies within the group should be the same.
- The last policy within the ICAP should always have the **goto**-priority-expression as END.

• If the **goto**-priority-expression field is not specified for a policy, the default value of END is assigned to **goto**-priority-expression.

# Create and verify ICAP policy

Consider a scenario where you want to define a policy in NetScaler to drop all the incoming URLs to a microservice containing the string example. Create a .yaml file called exampleicappolicy. yaml and use the appropriate CRD attributes to define the ICAP policy as follows:

```
1 apiVersion: citrix.com/v1beta1
2 kind: icappolicy
3 metadata:
    name: exampleicappolicy
4
5 spec:
6
    ingressclass: "cic-vpx"
7
    services:
      - "frontend"
8
9 icap-servers:
10
       servers:
         - ip: "192.168.1.1"
11
           port: 1344
12
13
       server-type: "TCP"
14
     icap:
     - direction: "REQUEST"
15
       profile:
16
         preview: "ENABLED"
17
18
         preview-length: 1024
19
         uri: "http://icap.example.com/reqmod"
         host-header: "icap.example.com"
20
         user-agent-header: "testAgent"
21
22
         query-params: "arg1=val1&arg2=val2"
23
         connection-keep-alive: "ENABLED"
24
        req-timeout: 30
25
         req-timeout-action: "BYPASS"
26
       content-inspection-criteria: "HTTP.REQ.URL.CONTAINS(\"example\")"
27
       default-action: "DROP"
28
       goto-priority-expression: "END"
       operation: "ICAP"
29
       server-failure-action: "CONTINUE"
```

After you deploy the .yaml file by running the command kubectl create -f exampleicappolicy .yaml, NetScaler Ingress Controller applies the policy configuration on NetScaler.

On the master node in the Kubernetes cluster, you can verify the status of the applied ICAP policy CRD by running the following command:

1 kubectl get icappolicies.citrix.com exampleicappolicy

You can view the status by running the following command:
1

```
    kubectl get icappolicies.citrix.com exampleicappolicy
    NAME STATUS MESSAGE
    exampleicappolicy Success CRD Activated
```

If there are issues while creating or applying the CRD, the same can be debugged by using the ingress controller logs.

```
kubectl logs <nsic pod name>
```

Also, you can verify whether the configuration is applied on the NetScaler by using the following steps.

- 1. Log on to the NetScaler CLI.
- 2. Use the following command to verify if the configuration is applied to NetScaler:

```
show ns icapProfile | grep exampleicappolicy
show icapProfile <icapProfile-name>
show run | grep -i icap
```



#### Sample policy configurations

#### ICAP policy in response mode

In this example, the ICAP policy is applied to the response traffic. In other words, the policy will inspect and potentially modify the responses sent from the ICAP server to NetScaler.

```
1 kubectl apply -f - <<EOF
2 apiVersion: citrix.com/v1beta1
3 kind: icappolicy
4 metadata:</pre>
```

```
5 name: exampleicappolicy
6 spec:
7
   ingressclass: "cic-vpx"
8 services:
      - "frontend"
9
10 icap-servers:
11
       servers:
12
         - ip: "192.168.1.1"
13
           port: 1344
       server-type: "TCP"
14
15
   icap:
   - direction: "RESPONSE"
16
17
       profile:
         preview: "ENABLED"
18
19
         preview-length: 1024
20
         uri: "http://icap.example.com"
21
        user-agent-header: "testAgent"
22
       query-params: "arg1=val1&arg2=val2"
23
       connection-keep-alive: "ENABLED"
24
       req-timeout: 30
25
         req-timeout-action: "BYPASS"
       content-inspection-criteria: "HTTP.RES.HEADER(\"Location\").
26
          CONTAINS(\"example\")"
       default-action: "DROP"
27
28
       goto-priority-expression: "END"
29
       operation: "ICAP"
       server-failure-action: "CONTINUE"
31 EOF
```

#### **Pipelining multiple ICAP servers**

In this example, we use multiple ICAP servers for processing requests.

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: citrix.com/v1beta1
3 kind: icappolicy
4 metadata:
5
  name: exampleicappolicy
6 spec:
7 ingressclass: "cic-vpx"
8
  services:
     - "frontend"
9
10 icap-servers:
11
      servers:
        - ip: "192.168.1.1"
13
          port: 1344
14
        - ip: "192.168.1.2"
15
          port: 1344
         - ip: "192.168.1.3"
16
           port: 1344
17
       server-type: "TCP"
18
19
     icap:
```

```
- direction: "RESPONSE"
20
21
       profile:
         preview: "ENABLED"
22
         preview-length: 1024
23
         uri: "http://icap.example.com"
24
25
         user-agent-header: "testAgent"
         query-params: "arg1=val1&arg2=val2"
26
27
         connection-keep-alive: "ENABLED"
28
         req-timeout: 30
29
         req-timeout-action: "BYPASS"
       content-inspection-criteria: "HTTP.RES.HEADER(\"Location\").
           CONTAINS(\"example\")"
       default-action: "DROP"
31
       goto-priority-expression: "END"
       operation: "ICAP"
33
34
       server-failure-action: "CONTINUE"
35 EOF
```

#### Multiple policy configurations

You can add multiple policy configurations in a single .yaml file and apply the policies to NetScaler. You need to add separate sections for each policy configuration as shown in the example here.

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: citrix.com/v1beta1
3 kind: icappolicy
4 metadata:
5
  name: exampleicappolicy1
6 spec:
7
  ingressclass: "cic-vpx"
   services:
8
     - "frontend"
9
10 icap-servers:
11
      servers:
         - ip: "192.168.1.1"
13
           port: 1344
14
       server-type: "TCP"
15
   icap:
       - direction: "RESPONSE"
16
17
         profile:
           preview: "ENABLED"
18
19
           preview-length: 1024
20
           uri: "http://icap.example.com"
           user-agent-header: "testAgent"
21
           query-params: "arg1=val1&arg2=val2"
22
           connection-keep-alive: "ENABLED"
23
24
           req-timeout: 30
25
           req-timeout-action: "BYPASS"
26
         content-inspection-criteria: "HTTP.RES.HEADER(\"Location\").
            CONTAINS(\"example\")"
         default-action: "DROP"
27
```

```
goto-priority-expression: "NEXT"
28
         operation: "ICAP"
29
         server-failure-action: "CONTINUE"
31 ---
32 apiVersion: citrix.com/v1beta1
33 kind: icappolicy
34 metadata:
35
   name: exampleicappolicy2
36 spec:
37
    ingressclass: "cic-vpx"
38
   services:
      - "frontend"
39
40
   icap-servers:
41
       servers:
         - ip: "192.168.1.1"
42
43
           port: 1344
44
       server-type: "TCP"
  icap:
45
46
       - direction: "RESPONSE"
47
         profile:
48
           preview: "ENABLED"
49
           preview-length: 1024
           uri: "http://icap.example.com"
           user-agent-header: "testAgent"
51
           query-params: "arg1=val1&arg2=val2"
52
53
           connection-keep-alive: "ENABLED"
54
           req-timeout: 30
55
           req-timeout-action: "BYPASS"
         content-inspection-criteria: "HTTP.RES.HEADER(\"Location\").
            CONTAINS(\"sample\")"
57
         default-action: "DROP"
         goto-priority-expression: "END"
58
         operation: "ICAP"
59
         server-failure-action: "CONTINUE"
61 EOF
```

## VIP custom resource definition (CRD)

#### July 29, 2024

NetScaler provides a Custom Resource Definition (CRD) called **VIP** for asynchronous communication between NetScaler **IPAM controller** and NetScaler Ingress Controller.

NetScaler provides **IPAM controller** for IP address management. NetScaler **IPAM controller** allocates an IP address to a service from a defined IP address range. NetScaler Ingress Controller configures the IP address allocated to the service as the virtual IP (VIP) address in NetScaler VPX and the service is exposed using this VIP address. When a new service is created, NetScaler Ingress Controller creates a CRD object for the service without an IP address. NetScaler **IPAM controller** checks for addition, deletion, or modification to CRD object and updates the CRD object with an IP address. After the CRD object is updated, NetScaler Ingress Controller automatically configures NetScaler-specific configuration in the tier-1 NetScaler VPX.

#### Note:

The VIP CRD is not supported for OpenShift routes. You can use OpenShift ingress to use VIP CRD.

## References

• Expose services of type LoadBalancer using an IP address from the NetScaler IPAM controller

## HTTPRoute

#### July 9, 2024

HTTPRoute is a custom resource which defines the routing decision for content switching. Currently HTTPRoute supports routing based on the following:

- Host name based routing
- Path based routing
- HTTP header Name based routing
- HTTP header value based routing.
- Cookie based routing
- Query parameter based routing
- HTTP method based routing
- Routing using NetScaler policy expressions

You can define one or more rules as part of an HTTPRoute object with each rule acting as a matching criteria for routing. An action is defined for each rule when the matching criteria is met for the incoming HTTP request. An action could be one of 'backend'in which the traffic is load balanced to the backend service or 'redirect'where the redirect response is sent back to the client. 'Backend'action creates Content switching policies in ADC and 'redirect'action creates responder policies in ADC.

There are three different ways of matching criteria as explained in the following table:

Matching criteria	Description
exact	Exactly matches the incoming request. This criteria is case insensitive.

Matching criteria	Description
prefix	Matches prefix of the incoming request. This
	criteria is case insensitive. For example: /a
	matches to $/a/b$ and $/a/c$ , but not to $/c/a$ .
contains	Matches if the incoming request contains the
	specified keyword. This criteria is case sensitive.

This topic contains a sample HTTPRoute CRD object and also explains the various attributes of the HTTPRoute CRD. For the complete CRD definition, see HTTPRoute.yaml.

## **HTTP CRD object example**

The following is a sample HTTP CRD object.

```
1 apiVersion: citrix.com/v1
 2 kind: HTTPRoute
 3 metadata:
 4
   name: test-route
 5
     labels:
     domain: abc.com
 6
 7 spec:
 8
    hostname:
 9 - abc.com
10 rules:
     - name: exactpath
11
12 match:
13 - path:
14 exaction:
16
           exact: /resources
16
       backend:
17
          kube:
18
             service: resource
19
             port: 80
   – name: prefixpath
20
   match:
21
       - path:
23
           prefix: /cart
24
      action:
25
       backend:
26
           kube:
27
             service: cart
28
             port: 80
29
   - name: header
      match:
31
       - headers:
32
         - headerName:
33
             contains: Mobile
```

34	action:
35	backend:
36	kube:
37	service: mobile
38	port: 443
39	backendConfig:
40	secureBackend: true
41	lbConfig:
42	lbmethod: ROUNDROBIN

For more examples, see HTTP Route Examples.

## **HTTPRoute.spec**

HTTPRoute custom resource defines a spec field which represents the HTTP routing specification which has a list of rules with an action for each rule defined.

The following table explains the various fields in the HTTPRoute.spec attribute.

Field	Description	Туре	Required
hostname	Specifies the list of host names of the server. The host name must be a valid subdomain as defined in RFC 1123, such as test.example.com. A wildcard host name in the form of *.example.com is also valid. In that case, any subdomain of example.com is considered for the matching. The default value is * which means to match all incoming HTTP requests.	string	Yes

## NetScaler ingress controller

Field	Description	Туре	Required
rules	Specifies the list of rules with a matching routing criteria associated with an action.	[] rules	No

#### **HTTPRoute.rules**

The following table explains the various fields in the HTTPRoute.rules attribute.

Field	Description	Туре	Required
name	Specifies a name to represent the rule. This field is used as an identifier in the content routing policy name in NetScaler. <b>Note:</b> For each rule, the name must be unique.	string	Yes
		<b>Note:</b> For each rule, the name must be unique.	
action	Specifies an action for the matching rule.	rules.action	Yes
match	List of matching routes with the same action. If more than one entry is present, this matching rule treated as an OR condition and the same action is chosen for any match.	[] rules.match	No

## HTTPRoute.rules.match

The following table explains the various fields in the HTTPRoute.rules.match attribute.

Field	Description	Туре	Required
path	Specifies URL path based routing rules.	HTTPRoute.rules.match.path	
headers	Specifies the list of header based matches	[ ] HTTPRoute.rules.matcl	No 1.headers
	for content routing. If	-	
	there is more than one		
	rule, this matching criteria is treated as an		
	AND condition and all		
	rules must match.		
cookies	Specifies the list of	[]	No
	cookie based matches	HTTPRoute.rules.match	.cookies
	for content routing. If		
	there is more than one		
	rule, this matching		
	AND condition and all		
	rules must match.		
queryParams	Specifies the list of	[]	No
	query parameters for	HTTPRoute.rules.match	.queryParams
	content routing. If		
	there is more than one		
	rule, this matching		
	criteria is treated as an		
	rules must match		
method	Specifies HTTP	string	No
	method based routing	Ũ	
	rules. Possible options		
	are: GET, POST, PUT,		
	and so on. An action is		
	chosen for the HTTP		
	request with the		
	matching method.		

Field	Description	Туре	Required
policyExpression	Specifies NetScaler policy expression based routing rules. Any custom NetScaler policy expression can be specified for content routing rules. The NetScaler Ingress Controller does not check the correctness of the expression. Hence, you must check the correctness of the expression. For more information on policy expression, see Expression Prefix. For example: HTTP.REQ.URL. PATH.GET(1).EQ( "foo")	string	No

#### HTTPRoute.rules.match.path

This attribute specifies the path based matching for content routing.

Following is an example for the HTTPRoute.rules.match.path attribute.

```
match:
1
2 - path:

3 pret

4 action:
         prefix: /resources
5
      backend:
6
          kube:
7
            service: resource
8
            port: 80
9 ---
10 match:
11
      - path:
12
          regex: '/foo/[A-Z0-9]{
13 3 }
```

```
14 '
15 action:
16 backend:
17 kube:
18 service: resource
19 port: 80
```

The following table explains the various fields in the HTTPRoute.rules.match.path attribute.

Field	Description	Туре	Required
prefix	Specifies the prefix expression of paths as a matching criteria. If the beginning path of an HTTP request matches the specified path, perform a match. For example, /a matches URLs /a and /a/b.	string	No
exact	Specifies the exact path as a matching criteria. Performs a match only if the request path exactly matches the specified path.	string	No

Field	Description	Туре	Required
regex	Specifies regular	string	No
	matching criteria for		
	paths. Performs a		
	match if the specified		
	regular expression		
	matches with the		
	incoming request.		
	Only regular		
	expressions in the Perl		
	Compatible Regular		
	Expression (PCRE)		
	format are supported.		
	For more information		
	on regular expressions		
	supported by		
	NetScaler, see Regular		
	Expressions.		

#### HTTPRoute.rules.match.headers

This attribute represents the header based matching for content routing.

The following table explains the various fields in the HTTPRoute.rules.match.headers attribute.

Field	Description	Туре	Required
headerName	Specifies the header name as a matching criteria for content routing. If the header exists, it is used for matching. In this case, header value is not considered for matching.	HTTPRoute.rules.match.	h <b>⊌a</b> ders.headerName

Field	Description	Туре	Required
headerValue	Specifies the header name and value as the matching criteria for content routing. For name, exact name is the matching criteria and matching criteria for value can be specified as exact, regex, or contains expression.	HTTPRoute.rules.match.	h <b>⊌ø</b> ders.headerValue

#### HTTPRoute.rules.match.headers.headerName

This attribute represents the header name based matching for content routing.

Following example shows sample snippets for the HTTPRoute.rules.match.headers. headerName attribute configuration.

```
match:
1
2
       - headers:
3
         - headerName:
4
           exact: mobile
5 action:
6
       backend:
7
           kube:
             service: mobile-service
8
9
             port: 80
10 ---
  match:
11
       - headers:
12
13
          - headerName:
14
            regex: "Header-[a-z]{
15
   1 }
   11
16
17
       action:
18
         backend:
19
           kube:
20
             service: resource-service
21
             port: 80
```

The following table explains the various fields in the HTTPRoute.rules.match.headers. headerName attribute.

Field	Description	Туре	Required
exact	Specifies the exact header name as matching criteria for routing.	string	No
contains	Specifies the string that is designated in the contains string as a matching criteria for the header name.	string	No
regex	Specifies the regular expression as a matching criteria. Performs a match if the header name matches the specified regular expression. Only regular expressions in the PCRE format are supported.	string	No
not	The default value for this attribute is false. If this value is true, the header name must not exist in the incoming request.	boolean	No

## HTTPRoute.rules.match.headers.headerValue

This attribute represents the header name and value matching for content routing. The header name is matched exactly and value is matched according to the fields specified.

The following example shows sample snippets for the HTTPRoute.rules.match.headers. headerValue attribute configuration.

```
1 match:
2 - headers:
3 - headerValue:
4 name: Origin
```

```
5
              exact: mobile
6
              not: true
7
        action:
8
        backend:
9
           kube:
10
              service: mobile
11
              port: 80
12
   ___
13
        match:
14
        - headers:
         - headerValue:
15
16
              name: Origin
              prefix: header1
17
18
       action:
19
         backend:
20
            kube:
21
              service: service1
22
              port: 80
23 ---
24
        match:
25
        - headers:
          - headerValue:
26
27
              name: Origin
28
              regex: "[a-z]{
29
   1 }
   .....
31
        action:
32
          backend:
33
            kube:
              service: example
34
35
              port: 80
```

The following table explains the various fields in the HTTPRoute.rules.match.headers. headerValue attribute.

Field	Description	Туре	Required
name	Specifies the name of the header that must match a value. The exact , contains , and regex fields are	string	Yes
	used for matching the header value. If none of the exact , contains, and regex field is present, any value for		
exact	the name is matched. Matches if the value of the HTTP header with the name field matches exactly.	string	No
contains	Matches if the value of the HTTP header with the name field contains the designated string.	string	No
regex	Matches if the value of the HTTP header with the name field matches the regular expression. Only regular expressions in the PCRE format are	string	No
not	supported. The default value for this attribute is false. If this value is true, the header name must not match the value.	boolean	No

## HTTPRoute.rules.match.cookies

This attribute represents the cookie based matching for content routing. The cookie header in the HTTP request is used for matching. A cookie with a name field is matched against a value if the value is present. If a value is not specified, it is matched for any value.

The following example shows sample snippets for the HTTPRoute.rules.match.cookies attribute configuration.

```
1
        match:
2
        - cookies:
3
          - name: version
            contains: v1
4
5
        action:
         backend:
6
7
            kube:
              service: v1-app
8
9
               port: 80
10 ---
11
        match:
12
        - cookies:
13
          - name: version
14
            exact: v1
15
        action:
16
         backend:
            kube:
17
18
              service: v1-app
19
               port: 80
20 ---
21
        match:
22
        - cookies:
23
          - name: version
24
            regex: '[a-z]{
25
    1 }
    1.1
27
        action:
28
          backend:
29
            kube:
               service: v1-app
31
               port: 80
```

The following table explains the various fields in the HTTPRoute.rules.match.cookies attribute.

Field	Description	Туре	Required
name	Specifies the name of the cookie whose value is used for matching. If none of the matching criteria like exact, regex and contains is present for the cookie name, any value for the cookie name is matched if the name is present.	string	Yes
exact	Matches if the value of the cookie with name field matches exactly.	string	Νο
contains	Matches if the value of the cookie with name field contains the string specified.	string	No
regex	Matches if the value of the cookie with name field matches the regular expression. Only regular expressions in the PCRE format are supported.	string	No
not	The default value for this attribute is false. If this value is true, the cookie with name must exist, but must not match the value.	boolean	No

## HTTPRoute.rules.match.queryParams

This attribute represents the HTTP query parameters in the URL matching for content routing.

```
1
      match:
2
       - queryParams:
3
          - name: version
4
            contains: v1
5
     action:
6
        backend:
          kube:
7
8
             service: v1-app
9
             port: 80
10 ---
       match:
11
12
       - queryParams:
13
         - name: version
14
           regex: '[a-z]{
15
   1 }
   1
16
17
       action:
18
        backend:
19
          kube:
20
             service: v1-app
21
             port: 80
22 ---
23
        match:
24
        - queryParams:
25
         - name: version
26
           exact: v1
27
           not: true
28
        action:
29
        backend:
           kube:
31
             service: mobile
32
             port: 80
```

The following example shows sample snippets for the HTTPRoute.rules.match.cookies attribute configuration.

Field	Description	Туре	Required
name	Specifies the name of the query parameter whose value is matched against. If none of the criteria like exact, regex and contains is present, any value for the query parameters name is matched if the name is present.	string	Yes
exact	Matches if the value of the query parameter with the name field matches exactly.	string	No
contains	Matches if the value of the query parameter with the name field contains the string specified.	string	No
regex	Matches if the value of the query parameter with the name field matches the regular expression. Only PCRE format regular expression is supported	string	No
not	The default value for this attribute is false. If this value is true, the query parameter with name must exist, but must not match the value.	boolean	No

#### **HTTPRoute.rules.action**

This attribute represents the action for matching rules.

The following table explains the various fields in the HTTPRoute.rules.action attribute.

The default action for this field is to send the traffic to a back-end service. Either the back end or the	rules.action.backend	No
redirect is required. The default action is to redirect the traffic. Either the back end or	rules.action.redirect	No
	The default action for this field is to send the traffic to a back-end service. Either the back end or the redirect is required. The default action is to redirect the traffic. Either the back end or redirect is required.	The default action for rules.action.backend this field is to send the traffic to a back-end service. Either the back end or the redirect is required. The default action is to rules.action.redirect redirect the traffic. Either the back end or redirect is required.

## HTTPRoute.rules.action.backend

This attribute represents routing the traffic to back-end service.

The following table explains the various fields in the HTTPRoute.rules.action.backend attribute.

Field	Description	Туре	Required
kube	Specifies the Kubernetes service information for the back end service.	action.backend.kube	

## HTTPRoute.rules.action.backend.kube

This attribute represents the Kubernetes service for the default back end. Service must belong to the same namespace as HTTPRoute resource. If the service is of type NodePort or Loadbalancer, the list of node IP addresses and NodePort of those nodes with pods is used as back-end service in 7.

Following is an example for the HTTPRoute.rules.action.backend.kube attribute.

1 kube:

```
service: service
2
3
        namespace: default
4
        port: 80
5
        backendConfig:
        lbConfig:
6
7
            lbmethod: ROUNDROBIN
8
          servicegroupConfig:
9
            clttimeout: '20'
```

The following table explains the various fields in the HTTPRoute.rules.action.backend. kube attribute.

Field	Description	Туре	Required
service	Specifies the name of the Kubernetes service for the default back end	string	Yes
port	Specifies the port number of the Kubernetes service for	integer	Yes
backendConfig	the default back end. Specifies the back-end configurations for the default back end.	BackendConfig	No

## BackendConfig

This attribute represents the back end configurations of NetScaler. Following is an example for the BackendConfig attribute configuration.

```
1 backendConfig:
2 sercureBackend: true
3 lbConfig:
4 lbmethod: ROUNDROBIN
5 servicegroupConfig:
6 clttimeout: '20'
```

The following table explains the various fields in the BackendConfig attribute.

Field	Description	Туре	Required
secureBackend	Specifies whether the communication is secure or not. If the value of secureBackend field is <b>true</b> secure communication is used to communicate with the back end. The default value is <b>false</b> , that means HTTP is used for the back end communication		
lbConfig	Specifies the NetScaler load balancing virtual server configurations for the given back end. One can specify key-value pairs as shown in the example which sets the LBVserver configurations for the back end. For all the valid configurations, see LB virtual server configurations.	object	No

#### NetScaler ingress controller

Field	Description	Туре	Required
servicegroupConfi	gSpecifies the NetScaler service group configurations for the given back end. One can specify the key-value pairs as shown in the example which sets the service group configurations for the back end. For all the valid configurations, see service group configurations.	object	No

## HTTPRoute.rules.action.redirect

This attribute represents the redirect action.

```
1 action:
2 redirect:
3 httpsRedirect: true
4 responseCode: 302
```

The following table explains the various fields in the HTTPRoute.rules.action.redirect attribute.

# NetScaler ingress controller

Field	Description	Туре	Required
httpsRedirect	Redirects the HTTP traffic to HTTPS if this field is set to yes. Only the scheme is changed to HTTPS without modifying the other URL part. Either httpsRedirect , hostRedirect or targetExpressio is required.	n	No
hostRedirect	Rewrites the host name part of the URL to the value set in this attribute and redirect the traffic. Other part of the URL is not modified during	string	No
targetExpression	redirection. Specifies the NetScaler expression for redirection. For example, to redirect traffic to HTTPS from HTTP, the following expression can be used: ""https://" +HTTP.REQ.HOSTNAME +	string	No
responseCode	HTTP.REQ.URL.HTTP_UF Specifies the response code. The default response code is 302, which can be customized using this attribute.	RL_SAFE" Integer	No

## Advanced content routing for Kubernetes with NetScaler

#### June 25, 2024

Kubernetes native Ingress offers basic host and path based routing. But, other advanced routing techniques like routing based on header values or query strings is not supported in the Ingress structure. You can expose these features on the Kubernetes Ingress through Ingress annotations, but annotations are complex to manage and validate.

You can expose the advanced content routing abilities provided by NetScaler as a custom resource definition (CRD) API for Kubernetes.

Using content routing CRDs, you can route the traffic based on the following parameters:

- Hostname
- URL path
- HTTP headers
- Cookie
- Query parameters
- HTTP method
- NetScaler policy expression

#### Note:

An Ingress resource and content routing CRDs cannot co-exist for the same service (IP address and port). The usage of content routing CRDs with Ingress is not supported.

#### The advanced content routing feature is exposed in Kubernetes with the following CRDs:

- Listener
- HTTPRoute



## Listener CRD

A Listener CRD object represents the end-point information like virtual IP address, port, certificates, and other front-end configurations. It also defines the default actions like sending the default traffic to a back end or redirecting the traffic. A Listener CRD object can refer to HTTPRoute CRD objects which represent HTTP routing logic for the incoming HTTP request.

For the full CRD definition, see the Listener CRD.

For complete information on all attributes of the Listener CRD, see Listener CRD documentation.

Listener CRD supports HTTP, SSL, and TCP profiles. Using these profiles, you can customize the default protocol behavior. Listener CRD also supports the analytics profile which enables NetScaler to export the type of transactions or data to different endpoints.

For more information about profile support for Listener CRD, see the Profile support for the Listener CRD.

#### **Deploy the Listener CRD**

- 1. Download the Listener CRD.
- 2. Deploy the listener CRD with following command.

```
1 Kubectl create -f Listener.yaml
```

Example:

#### How to write Listener CRD objects

After you have deployed the CRD provided by NetScaler in the Kubernetes cluster, you can define the listener configuration in a YAML file. In the YAML file, use Listener in the kind field and in the spec section add the listener CRD attributes based on your requirement for the listener configuration. After you deploy the YAML file, the NetScaler Ingress Controller applies the listener configuration on NetScaler.

Following is a sample Listener CRD object definition named Listener-crd.yaml.

```
1 apiVersion: citrix.com/v1
2 kind: Listener
3 metadata:
    name: my-listener
4
5
    namespace: default
6 spec:
   certificates:
7
8 - secret:
9
        name: my-secret
       # Secret named 'my-secret' in current namespace bound as default
10
          certificate
      default: true
11
12
     - secret:
         # Secret 'other-secret' in demo namespace bound as SNI
13
            certificate
14
         name: other-secret
         namespace: demo
16
    - preconfigured: second-secret
       # preconfigured certkey name in ADC
17
   vip: '192.168.0.1' # Virtual IP address to be used, not required when
18
         CPX is used as ingress device
19
     port: 443
     protocol: https
20
21
    redirectPort: 80
22
     secondaryVips:
23
     - "10.0.0.1"
```

24	- "1.1.1.1"
25	policies:
26	httpprofile:
27	config:
28	websocket: "ENABLED"
29	tcpprofile:
30	config:
31	sack: "ENABLED"
32	sslprofile:
33	config:
34	ssl3: "ENABLED"
35	sslciphers:
36	- SECURE
37	- MEDIUM
38	analyticsprofile:
39	config:
40	- type: webinsight
41	parameters:
42	allhttpheaders: "ENABLED"
43	csvserverConfig:
44	rhistate: 'ACTIVE'
45	routes:
46	# Attach the policies from the below Routes
47	- name: domain1-route
48	namespace: default
49	- name: domain2-route
50	namespace: default
51	- labelSelector:
52	# Attach all HTTPRoutes with <b>label</b> route=my-route
53	route: my-route
54	# Default action when traffic matches none of the policies in the
	HIIPRoute
55	defaultAction:
56	backend:
57	KUDE:
58	namespace: detault
59	port: 80
60	service: detault-service
61	
62	toconing:
64	# Use round robin Lb method <b>for detault</b> service
65	servicegroupConfig.
66	# Client timeout of 20 seconds
67	$\pi$ clittimeout • "20"
01	

In this example, a listener is exposing an HTTPS endpoint. Under the certificates section, SSL certificates for the endpoint are configured using Kubernetes secrets named my-secret and other -secret and a default ADC preconfigured certificate with certkey named second-secret. The default action for the listener is configured as a Kubernetes service. Routes are attached with the listener using both label selectors and individual route references using name and namespace. After you have defined the Listener CRD object in the YAML file, deploy the YAML file using the following command. In this example, Listener-crd.yaml is the YAML definition.

```
1 Kubectl create -f Listener-crd.yaml
```

#### **HTTPRoute CRD**

An HTTPRoute CRD object represents the HTTP routing logic for the incoming HTTP requests. You can use a combination of various HTTP parameters like host name, path, headers, query parameters, and cookies to route the incoming traffic to a back-end service. An HTTPRoute object can be attached to one or more Listener objects which represent the end point information. You can have one or more rules in an HTTPRoute object, with each rule specifying an action associated with it. Order of evaluation of the rules within an HTTPRoute object is the same as the order mentioned in the object. For example, if there are two rules with the order rule1 and rule2, with rule1 is written before rule2, rule1 is evaluated first before rule2.

HTTPRoute CRD definition is available at HTTPRoute.yaml. For complete information on the attributes for HTTP Route CRD, see HTTPRoute CRD documentation.

Now, NetScaler supports configuring the HTTP route CRD resource as a resource backend in the Ingress with Kubernetes Ingress version networking.k8s.io/v1.With this feature, you can extend advanced content routing capabilities to Ingress. For more information, see Advanced content routing for Kubernetes Ingress using HTTPRoute CRD.

## **Deploy the HTTPRoute CRD**

Perform the following to deploy the HTTPRoute CRD:

- 1. Download the HTTPRoute.yaml.
- 2. Apply the HTTPRoute CRD in your cluster using the following command.

Kubectl apply -f HTTPRoute.yaml

Example:

## How to write HTTPRoute CRD objects

Once you have deployed the HTTPRoute CRD, you can define the HTTP route configuration in a YAML file. In the YAML file, use HTTPRoute in the kind field and in the spec section add the HTTPRoute

CRD attributes based on your requirement for the HTTP route configuration.

Following is a sample HTTPRoute CRD object definition named as Route-crd.yaml.

```
1 apiVersion: citrix.com/v1
2 kind: HTTPRoute
3 metadata:
4
     name: test-route
5 spec:
6
    hostname:
7
     - host1.com
8
     rules:
9
     - name: header-routing
10
       match:
       - headers:
         - headerName:
13
             exact: my-header
14
       action:
15
         backend:
16
           kube:
17
             service: mobile-app
18
             port: 80
19
             backendConfig:
20
                secureBackend: true
21
                lbConfig:
22
                  lbmethod: ROUNDROBIN
23
     - name: path-routing
24
       match:
25
       - path:
26
           prefix: /
27
       action:
28
         backend:
           kube:
29
              service: default-app
              port: 80
```

In this example, any request with a header name matching my-header is routed to the mobile-app service and all other traffic is routed to the default-app service.

For detailed explanations and API specifications of HTTPRoute, see HTTPRoute CRD.

After you have defined the HTTP routes in the YAML file, deploy the YAML file for HTTPRoute CRD object using the following command. In this example, Route-crd.yaml is the YAML definition.

1 Kubectl create -f Route-crd.yaml

Once you deploy the YAML file, the NetScaler Ingress Controller applies the HTTP route configuration on the Ingress NetScaler device.

#### Attaching HTTPRoute CRD objects to a Listener CRD object

You can attach HTTPRoute CRD objects to a Listener CRD object in two ways:

- Using name and namespace
- Using labels and selector

#### Attaching HTTPRoute CRD objects using name and namespace

In this approach, a Listener CRD object explicitly refers to one or more HTTPRoute objects by specifying the name and namespace in the routes section.

The order of evaluation of HTTPRoute objects is the same as the order specified in the Listener CRD object with the first HTTPRoute object evaluated first and so on.

For example, a snippet of the Listener CRD object is shown as follows.

```
1 routes:
2 - name: route-1
3 namespace: default
4 - name: route-2
5 namespace: default
```

In this example, the HTTPRoute CRD object named route1 is evaluated before the HTTPRoute named route2.

#### Attaching an HTTPRoute CRD object using labels and selector

You can also attach HTTPRoute objects to a Listener object by using labels and selector. You can specify one or more labels in the Listener CRD object. Any HTTPRoute objects which match the labels are automatically linked to the Listener object and the rules are created in NetScaler. When you use this approach, there is no particular order of evaluation between multiple HTTPRoute objects. Only exception is an HTTPRoute object with a default route (a route with just a host name or a '/'path) which is evaluated as the last object.

For example, snippet of a listener resource is as follows:

```
1 routes:
2 - labelSelector:
3 team: team1
```

## Listener

October 17, 2024

The Listener CRD represents the endpoint information of the content switching load balancing virtual server. This topic contains a sample Listener CRD object and also explains the various attributes of the Listener CRD. For the complete CRD definition, see Listener.yaml.

Note:

An ingress resource and content routing CRDs (Listener CRD and HTTPRoute CRD) cannot co-exist for the same endpoint (IP address and port). The usage of content routing CRDs with Ingress is not supported.

## Listener CRD object example

The following is an example of a Listener CRD object.

```
1 apiVersion: citrix.com/v1
2 kind: Listener
3 metadata:
    name: my-listener
4
5 namespace: default
6 spec:
7
    certificates:
8
     - secret:
9
        name: my-secret
10
       # Secret named 'my-secret' in current namespace bound as default
          certificate
11
      default: true
     - secret:
12
         # Secret 'other-secret' in demo namespace bound as SNI
13
            certificate
14
         name: other-secret
         namespace: demo
16 - preconfigured: second-secret
      # preconfigured certkey name in ADC
17
     vip: '192.168.0.1' # Virtual IP address to be used, not required when
18
         CPX is used as ingress device
19
     port: 443
   protocol: https
20
21
  redirectPort: 80
22 secondaryVips:
   - "10.0.0.1"
23
     - "1.1.1.1"
24
25
    policies:
26
      httpprofile:
27
         config:
28
           websocket: "ENABLED"
29
       tcpprofile:
         config:
           sack: "ENABLED"
31
       sslprofile:
32
         config:
           ssl3: "ENABLED"
34
```

35	sslciphers:
36	- SECURE
37	- MEDIUM
38	analyticsprofile:
39	config:
40	- type: webinsight
41	parameters:
42	allhttpheaders: "ENABLED"
43	csvserverConfig:
44	rhistate: 'ACTIVE'
45	routes:
46	# Attach the policies from the below Routes
47	- name: domain1-route
48	namespace: default
49	- name: domain2-route
50	namespace: default
51	- labelSelector:
52	# Attach all HTTPRoutes with <b>label</b> route=my-route
53	route: my-route
54	# Default action when traffic matches none of the policies in the
	HTTPRoute
55	defaultAction:
56	backend:
57	kube:
58	namespace: default
59	port: 80
60	service: <b>default</b> -service
61	backendConfig:
62	lbConfig:
63	# Use round robin LB method <b>for default</b> service
64	lbmethod: ROUNDROBIN
65	servicegroupConfig:
66	# Client timeout of 20 seconds
67	clttimeout: "20"

For more examples, see Listener examples.

## Listener.spec

The Listener.spec attribute defines the Listener custom resource specification. The following table explains the various fields in the Listener.spec attribute.

# NetScaler ingress controller

Field	Description	Туре	Required
protocol	Specifies the protocol of the load balancing content switching virtual server. Allowed values are: http and https.	string	yes
port	Specifies the port number of the load balancing content switching virtual server. The default port number for the HTTP protocol is 80 and the HTTPS protocol is 443.	integer	No
routes	Specifies the list of HTTPRoute resources that is to be attached to the Listener resource. The order of evaluation is as per the order of the list. That is, if multiple entries are present first route specified in the list has highest priority and so on.	[] routes	No
ingressclass	Specifies the ingress class so that only the ingress controller associated with the specified ingress class processes the resource. Otherwise, all the controllers in the cluster will process this resource.	string	No

Field	Description	Туре	Required
certificates	Specifies the list of certificates for the SSL virtual server if the protocol is HTTPS. This field is required if the protocol is HTTPS.	[] certificates	No
vip	Specifies the endpoint IP Address for the load balancing content switching virtual server. This address is required for NetScaler VPX and MPX devices, but not required for NetScaler CPXs present in the Kubernetes cluster. For NetScaler CPX, vip is same as the primary IP address of the NetScaler CPX allocated by the CNI.	string	No
defaultAction	Specifies the default action to take if none of the HTTPRoute resources specified in routes match the traffic.	action	No
policies	Specifies the option that enables you to customize HTTP, TCP, and SSL policies associated with the front-end virtual server.	[] Listener.policies	No
Field	Description	Туре	Required
---------------	--	-----------	----------
redirectPort	Specifies that the HTTP traffic on this port is redirected to the HTTPS port.	Integer	No
secondaryVips	Specifies a set of IP addresses which are used as VIPs with the primary VIP. An IPset is created and these VIPs are added to the IPset.	[] string	No

## Listener.certificates

The Listener.certificates attribute defines the TLS certificate related information for the SSL virtual server.

Following is an example for the Listener.certificates attribute.

```
    certificates:
    secret:
    name: my-secret
    namespace: demo
    default: true
    preconfigured: configured-secret
```

The following table explains the various fields in the Listener.certificates attribute.

Field	Description	Туре	Required
secret	Specifies TLS certificates specified through the Kubernetes secret resource. The secret must contain keys tls .crt and tls.key. These keys contain the certificate and private key. Either the secret or the preconfigured field is required. All certificates are bound to the SSL virtual server as SNI certificates.	certificates.secret	No

Field	Description	Туре	Required
preconfigured	Specifies the name of the preconfigured TLS <b>certkey</b> in NetScaler, and this field is applicable only for Tier-1 VPX and MPX devices. The <b>certkey</b> must be present before the actual deployment of the Listener resource and otherwise deployment of the resource fails with an error. The NetScaler Ingress Controller does not manage the life cycle of <b>certkey</b> . So, you have to manage any addition or deletion of <b>certkey</b> manually. Either the secret or the preconfigured field is required.	string	No

Field	Description	Туре	Required
default	Specifies the default certificate. Only one of the certificates can be marked as default. The default certificate is presented if virtual server receives the traffic without an SNI field. This certificate can be used to access the HTTPS application using the IP Address. Applicable values are <b>true</b> and <b>false</b>	boolean	No

# Listener.certificates.Secret

This attribute represents the Kubernetes secret resource for the TLS certificates that has to be bound to the SSL virtual server.

The following table explains the various fields in the Listener.certificate.Secret attribute.

Field	Description	Туре	Required
name	Specifies the name of	string	yes
	the Kubernetes secret		
	resource. The secret		
	must contain keys		
	named tls.crt		
	and tls.key. These		
	keys contain the		
	certificate and private		
	key. If more than		
	one tls.crt field is		
	present in the secret		
	object, then the first		
	certificate is		
	considered as a server		
	certificate and		
	remaining certificates		
	are considered as		
	intermediate CA		
	certificates. Also,		
	certificates are linked		
	recursively to each		
	other starting from the		
	server certificate.		
namespace	Specifies the	string	No
	namespace of the		
	Kubernetes secret		
	resource. If this value		
	is not specified, the		
	namespace is		
	considered as same as		
	the Listener resource.		

## Listener.routes

This attribute represents the list of HTTPRoute objects that are attached to the Listener resource. The following table explains the various fields in the Listener.routes attribute.

Field	Description	Туре	Required
name	Specifies the name of the HTTPRoute resource evaluated for the routing decision to the back end server. Either the name or the labelSelector is required.	string	No
namespace	Specifies the namespace of the HTTPRoute resource. The default value is the name space of the Listener resource.	string	No

Field	Description	Туре	Required
labelSelector	Specifies the label selector of the	object	No
	HTTPRoute resource.		
	This field provides		
	another way to attach		
	HTTPRoute resources.		
	HTTPRoute objects		
	with label keys and		
	values matching this		
	selector are		
	automatically		
	attached to the listener		
	resource. If routes get		
	attached through the		
	labelSelector, routes		
	are attached without		
	any specific order.		
	Exception for this rule		
	is a route with a		
	default path ('/') which		
	is always attached at		
	the end. As shown in		
	the example, any		
	HTTPRoute objects		
	with labels and are		
	attached to the		
	listener object. For		
	more information on		
	labels and selectors,		
	see the Kubernetes		
	Documentation.		

## Listener.action

This attribute represents the default action if a request to the load balancing virtual server does not match any of the route objects presented in the Listener.routes field.

The following table explains the various fields in the Listener.action attribute.

#### NetScaler ingress controller

Field	Description	Туре	Required
backend	The default action for this field is to send the traffic to a back-end service. Either the back end or the redirect is required.	action.backend	No
redirect	The default action is to redirect the traffic. Either the back end or redirect is required.	action.redirect	No

## Listener.action.backend

This attribute specifies the back end service for the default action. The following table explains the various fields in the Listener.action.backend attribute.

Field	Description	Туре	Required
kube	Specifies the Kubernetes service information for the back end service.	action.backend.kube	

### Listener.action.backend.kube

This attribute represents the Kubernetes back end service for the default back end. If the service is of type NodePort or Loadbalancer, the node IP address and NodePort are used to send the traffic to the back end.

Following is an example for the Listener.action.backend.kube attribute.

```
1
           kube:
2
             service: default-service
             namespace: default
3
4
             port: 80
5
             backendConfig:
6
               lbConfig:
                 lbmethod: ROUNDROBIN
7
               servicegroupConfig:
8
                 clttimeout: '20'
9
```

The following table explains the various fields in the Listener.action.backend.kube attribute.

Field	Description	Туре	Required
service	Specifies the name of the Kubernetes service for the default back end.	string	yes
namespace	Specifies the namespace of the Kubernetes service for the default back end.	string	yes
port	Specifies the port number of the Kubernetes service for the default back end.	integer	yes
backendConfig	Specifies the back-end configurations for the default back end.	BackendConfig	no

# BackendConfig

This attribute represents the back end configurations of NetScaler. Following is an example for the BackendConfig attribute.

```
    backendConfig:
    sercureBackend: true
    lbConfig:
    lbmethod: ROUNDROBIN
    servicegroupConfig:
    clttimeout: '20'
```

The following table explains the various fields in the BackendConfig attribute.

Field	Description	Туре	Required
secureBackend	Specifies whether the communication is secure or not. If the value of the secureBackend field is <b>true</b> secure communication is used to communicate with the back end. The default value is <b>false</b> , that means HTTP is used for the back end		
lbConfig	communication. Specifies the NetScaler load balancing virtual server configurations for the given back end. One can specify key-value pairs as shown in the example which sets the LBVserver configurations for the back end. For all the valid configurations, see LB virtual server configurations.	object	No

#### NetScaler ingress controller

Field	Description	Туре	Required
servicegroupConfi	gSpecifies the NetScaler service group configurations for the given back end. One can specify the key-value pairs as shown in the example which sets the service group configurations for the back end. For all the valid configurations, see service group configurations.	object	No

# Listener.action.redirect

```
    defaultAction:
    redirect:
    httpsRedirect: true
    responseCode: 302
```

The following table explains the various fields in the Listener.action.redirect attribute.

Field	Description	Туре	Required
httpsRedirect	Redirects the HTTP traffic to HTTPS if this field is set to yes. Only the scheme is changed to HTTPS without modifying the other URL part. Either httpsRedirect , hostRedirect ortargetExpressio is required.	boolean	No

Field	Description	Туре	Required
hostRedirect	Rewrites the host name part of the URL to the value set here and redirect the traffic. Other part of the URL is not modified during redirection	string	No
targetExpression	Specifies the NetScaler expression for redirection. For example, to redirect traffic to HTTPS from HTTP, the following expression can be used: ""https://" +HTTP.REQ.HOSTNAME +	string	No
	HTTP.REQ.URL.HTTP_URL	_SAFE"	
responseCode	Specifies the response code. The default response code is 302, which can be customized using this attribute.	Integer	No

# Listener.policies

This attribute represents the default policies which are used for the Listener when policies are not specified. By using Listener.policies, you can customize the TCP, HTTP, and SSL behavior.

Following is an example for the Listener.policies attribute.

```
1 policies:
2 httpprofile:
3 config:
4 websocket: "ENABLED"
5 tcpprofile:
6 config:
7 sack: "ENABLED"
```

8	sslprofile:
9	config:
10	ssl3: "ENABLED"
11	sslciphers:
12	- HIGH
13	- MEDIUM
14	analyticsprofile:
15	config:
16	- type: webinsight
17	parameters:
18	allhttpheaders: "ENABLED"
19	csvserverConfig:
20	rhistate: 'ACTIVE'
21	stateupdate: 'ENABLED'

The following table explains the various fields in the Listener.policies attribute.

Field	Description	Туре	Required
httpprofile	Specifies the HTTP configuration for the front end virtual server	Listener.policies.httppro	fi <b>N</b> eo
tcpprofile	Specifies the TCP configuration for the front-end virtual server.	Listener.policies.tcpprofi	l€NO
sslprofile	Specifies the SSL configuration for the front-end virtual server	Listener.policies.sslprofil	eNo

Field	Description	Туре	Required
sslciphers	Specifies the list of ciphers which are to be bound to the SSL profile. The order is as specified in the list with the higher priority is provided to the first in the list and so on. You can use any SSL ciphers available in NetScaler or user created cipher groups in this field. For information about the list of ciphers available in the NetScaler, see Ciphers in NetScaler.	[] string	No
analyticsprofile	Specifies the analytics profile configuration for the front-end virtual server	Listener.policies.analytic	csøœfile
csvserverConfig	Specifies the front-end CS virtual server configuration for the Listener. You can specify the key value pair as shown in the example which sets the CS virtual server configuration for the front-end.	Object	No

# Listener.policies.tcpprofile

This attribute represents the TCP profile settings for the front-end CS virtual server.

Following is an example for the Listener.policies.tcpprofile attribute.

```
1
 policies:
2
      tcpprofile:
      config:
3
       sack: "ENABLED"
4
        nagle: "ENABLED"
5
6 ---
7
      policies:
8
      tcpprofile:
       preconfigured: test-tcp-profile
9
```

The following table explains the various fields in the Listener.policies.tcpprofile attribute.

Field	Description	Туре	Required
preconfigured	Specifies the name of the preconfigured TCP profile that is to be used for the front-end CS virtual server. This profile must be present in the NetScaler before applying the policy. Otherwise, the Listener resource fails to apply. Either preconfigured or config is required.	string	No
config	Specifies the TCP profile settings for the front-end virtual server. You can specify the key-value pair as shown in the example to tune the TCP characteristics of the virtual server.	Object	No

### Listener.policies.httpprofile

This attribute represents the HTTP configuration for the front-end CS virtual server.

Following is an example for the Listener.policies.httpprofile attribute.



The following table explains the various fields in the Listener.policies.httpprofile attribute.

Field	Description	Туре	Required
preconfigured	Specifies the name of the preconfigured HTTP profile that is to be used for the front end CS virtual server. This profile must be present in the NetScaler before applying the policy. Otherwise, Listener resource fails to apply. Either preconfigured or config is required.	string	No
config	Specifies the HTTP profile settings for the front-end virtual server. You can specify the key-value pair as shown in the example to tune the HTTP protocol characteristics of the virtual server.	Object	No

# Listener.policies.sslprofile

This attribute represents the SSL profile for the front-end CS virtual server.

Following is an example for the Listener.policies.sslprofile attribute.

```
policies:
1
2
       sslprofile:
         config:
3
          ssl3: "ENABLED"
4
5
  ___
      policies:
6
7
        sslprofile:
         preconfigured: test-ssl-profile
8
```

The following table explains the various fields in the Listener.policies.sslprofile attribute.

Field	Description	Туре	Required
preconfigured	Specifies the name of the preconfigured SSL profile that is to be used for the front-end SSL virtual server. This profile must be present in the NetScaler before applying the policy. Otherwise, the Listener resource fails to apply. Either preconfigured or config is required.	string	No
config	Specifies the SSL profile configuration for the front-end virtual server. You can specify the key-value pair as shown in the example to tune the SSL characteristics of the virtual server.	Object	No

Field	Description	Туре	Required
	Note: You must en	able the default profil	es using the
	set ssl param	eter -defaultPr	ofile ENABLED command in
	NetScaler for using	; the advanced SSL fea	atures.

## Listener.policies.analyticsprofile

This attribute represents the analytics profile that is used to export counters and metrics to NetScaler Observability Exporter. By configuring this attribute, you can choose what is to be exported by creating and binding the analytics profile.

Following is an example for the Listener.policies.analyticsprofile attribute.

```
1
       policies:
2
        analyticsprofile:
3
         config:
         - type: webinsight
4
5
           parameters:
            allhttpheaders: "ENABLED"
6
7
   ___
       policies:
8
        analyticsprofile:
9
         preconfigured:
10
11
         - test-analytics-profile
         - test2-analytics-profile
12
```

The following table explains the various fields in the Listener.policies.analyticsprofile attribute.

Field	Description	Туре	Required
preconfigured	Specifies the list of preconfigured analytics profiles that needs to be bound to the front-end virtual server. These profiles must be present in the NetScaler before applying the policy. Otherwise, the Listener resource fails to apply. Either preconfigured or config is required.	[] string	No
config	Specifies the list of analytics profiles which is to be bound to the front-end virtual server. This determines the fields to be exported to NetScaler Observability Exporter.	[] Lis- tener.policies.analyticsp	No profile.config

# Listener.policies.analyticsprofile.config

This attribute represents the analytics profile configuration for the different types of insights and HTTP header parameters which need to be exported.

The following table explains the various fields in the Listener.policies.analyticsprofile .config attribute.

Field	Description	Туре	Required	
type	Specifies the type that determines the type of analytics profile to be enabled. You can enable one or more of the following types: webinsight, tcpinsight, securityinsight, videoinsight, hdxinsight, gatewayinsight, timeseries, lsninsight, and botinsight	string	Yes	
parameters	Specifies the additional parameters to be enabled as part of the analytics profile. You can specify the key-value pair as shown in the example. For example, using this field, you can select the HTTP parameters to be exported as part of webinsight	Object	No	

# Configure web application firewall policies with the NetScaler Ingress Controller

July 2, 2025

NetScaler provides a Custom Resource Definition (CRD) called the WAF CRD for Kubernetes. You can use the WAF CRD to configure the web application firewall policies with the NetScaler Ingress Con-

troller on the NetScaler VPX, MPX, SDX, and CPX. The WAF CRD enables communication between the NetScaler Ingress Controller and NetScaler for enforcing web application firewall policies.

In a Kubernetes deployment, you can enforce a web application firewall policy to protect the server using the WAF CRD. For more information about web application firewall, see Web application security.

With the WAF CRD, you can configure the firewall security policy to enforce the following types of security checks for Kubernetes native applications.

## **Common protections**

- Buffer overflow
- Content type
- Allow URL
- Block URL
- Cookie consistency
- Credit card

## **HTML protections**

- CSRF (cross side request forgery) form tagging
- Field formats
- Form field consistency
- File upload types
- HTML cross-site scripting
- HTML SQL injection

## **JSON protections**

- JSON denial of service
- JSON SQL injection
- JSON cross-site scripting

#### XML protections

- XML web services interoperability
- XML attachment
- XML cross-site scripting
- XML denial of service

- XML format
- XML message validation
- XML SOAP fault filtering
- XML SQL injection

Based on the type of security checks, you can specify the metadata and use the CRD attributes in the WAF CRD .yaml file to define the WAF policy.

## **WAF CRD definition**

The WAF CRD is available in the NetScaler Ingress Controller GitHub repository at waf-crd.yaml. The WAF CRD provides attributes for the various options that are required to define the web application firewall policies on NetScaler.

## **WAF CRD attributes**

The following table lists the various attributes provided in the WAF CRD:

CRD attribute	Description
commonchecks	Specifies a list of common security checks, which are applied irrespective of the content type.
block_urls	Protects URLs.
buffer_overflow	Protects buffer overflow.
content_type	Protects content type.
htmlchecks	Specifies a list of security checks to be applied for HTML content types.
cross_site_scripting	Prevents cross site scripting attacks.
sql_injection	Prevents SQL injection attacks.
form_field_consistency	Prevents form tampering.
csrf	Prevents cross side request forgery (CSRF) attacks.
cookie_consistency	Prevents cookie tampering or session takeover.
field_format	Validates the form submission.
fileupload_type	Prevents malicious file uploads.
jsonchecks	Specifies security checks for JSON content types.

CRD attribute	Description
xmlchecks	Specifies security checks for XML content types.
wsi	Protects web services interoperability.
redirect_url	Redirects URL when block is enabled on protection.
servicenames	Specifies the services to which the WAF policies are applied.
application_type	Protects application types.
signatures	Specifies the location of the external signature file.
html_error_object	Specifies the location of the customized error page to respond when HTML or common violations are attempted.
xml_error_object	Specifies the location of the customized error page to respond when XML violations are attempted.
json_error_object	Specifies the location of the customized error page to respond when JSON violations are attempted.
ip_reputation	Enables the IP reputation feature.
target	Determines the traffic to be inspected by the WAF. If you do not specify the traffic targeted, all traffic is inspected by default.
target.path	Specifies the list of HTTP URLs to be inspected.
target.method	Specifies the list of HTTP methods to be inspected.
target.header	Specifies the list of HTTP headers to be inspected.
exclude	Determines the traffic to be excluded by the WAF. If you do not specify the traffic to be excluded, all traffic is inspected by default.
exclude.path	Specifies the list of HTTP URLs to be excluded.
exclude.method	Specifies the list of HTTP methods to be excluded.
exclude.header	Specifies the list of HTTP headers to be excluded.

## **Deploy WAF CRD**

Perform the following steps to deploy the WAF CRD:

- 1. Download the CRD (waf-crd.yaml).
- 2. Deploy the WAF CRD using the following command:

```
1 kubectl create -f waf-crd.yaml
```

For example,

## How to write a WAF configuration

After you have deployed the WAF CRD provided by NetScaler in the Kubernetes cluster, you can define the web application firewall policy configuration in a .yaml file. In the .yaml file, use waf in the kind field. In the spec section add the WAF CRD attributes based on your requirements for the policy configuration.

After you deploy the .yaml file, the NetScaler Ingress Controller applies the WAF configuration on the Ingress NetScaler device.

The following are some examples for writing web appliction firewall policies.

#### Enable protection for cross-site scripting and SQL injection attacks

Consider a scenario in which you want to define and specify a web application firewall policy in the NetScaler to enable protection for the cross-site scripting and SQL injection attacks. You can create a .yaml file called wafhtmlxssql.yaml and use the appropriate CRD attributes to define the WAF policy as follows:

```
1 apiVersion: citrix.com/v1
2 kind: waf
3 metadata:
4
      name: wafhtmlxsssql
5 spec:
6
       servicenames:
7
           - frontend
8
       application_type:
           - "HTML"
9
       html_error_object: "http://x.x.x.x/crd/error_page.html"
11
       security_checks:
12
           html:
```

```
13cross_site_scripting: "on"14sql_injection: "on"
```

#### Apply rules to allow only known content types

Consider a scenario in which you want to define a web application firewall policy that specifies rules to allow only known content types and block unknown content types. Create a .yaml file called wafcontenttype.yaml and use the appropriate CRD attributes to define the WAF policy as follows:

```
1 apiVersion: citrix.com/v1
2 kind: waf
3 metadata:
   name: wafcontenttype
4
5 spec:
6
       servicenames:
7

    frontend

8
       application_type:
           - "HTML"
9
10
       html_error_object: "http://x.x.x/crd/error_page.html"
11
       security_checks:
           common:
12
13
             content_type: "on"
14
       relaxations:
15
           common:
16
             content_type:
17
               types:
18

    custom_cnt_type

19
                    - image/crd
```

#### Protect against known attacks

The following is an example of a WAF CRD configuration for applying external signatures. You can copy the latest WAF signatures from Signature Location to the local web server and provide the location of the copied file as *signature\_url*.

```
1 apiVersion: citrix.com/v1
2 kind: waf
3 metadata:
4
       name: wafhtmlsigxsssql
5 spec:
6 servicenames:
7

    frontend

8
       application_type:
9
          - "HTML"
       signatures: "http://x.x.x.x/crd/sig.xml"
       html_error_object: "http://x.x.x.x/crd/error_page.html"
11
12
       security_checks:
13
           html:
```

```
14cross_site_scripting: "on"15sql_injection: "on"
```

#### Protect from header buffer overflow attacks and block multiple headers

The following is an example of a WAF CRD configuration for protecting buffer overflow.

```
1 apiVersion: citrix.com/v1
2 kind: waf
3 metadata:
4
      name: wafhdrbufferoverflow
5 spec:
6 servicenames:
7

    frontend

       application_type:
8
9
           - "HTML"
       html_error_object: "http://x.x.x.x/crd/error_page.html"
11
       security_checks:
           common:
13
             buffer_overflow: "on"
             multiple_headers:
14
               action: ["block", "log"]
16
       settings:
17
           common:
             buffer_overflow:
18
19
               max_cookie_len: 409
20
               max_header_len: 4096
21
               max_url_len: 1024
```

#### Prevent repeated attempts to access random URLs on a web site

The following is an example of a WAF CRD configuration for providing URL filter rules. You can add URLs to permit under *allow\_url* and URLs to deny under *block\_url*. The URL can be a regular expression also. The target and exclude options let you specify which URLs, headers, and methods the WAF policy must inspect or ignore. If you do not set target or exclude, the WAF inspects all URLs by default.

```
apiVersion: citrix.com/v1
1
2
  kind: waf
3 metadata:
4
       name: wafurlchecks
5 spec:
6
    servicenames:
7

    frontend

8
       application_type:
9
           - "HTML"
       html_error_object: "http://x.x.x.x/crd/error_page.html"
11
       target:
```

```
12
            path:
13
                - /
14
            method:
15
                - GET
                - POST
16
17
            header:
18
                - Host
19
        exclude:
20
            path:
                - index.html
              method:
23
                - GET
24
        security_checks:
25
            common:
26
              allow_url: "on"
27
              block_url: "on"
28
        relaxations:
29
            common:
              allow_url:
31
                urls:
32
                     - payment.php
                     - cover.php
33
34
        enforcements:
35
            common:
36
              block_url:
37
                urls:
38
                     - "^[^?]*(passwd|passwords?)([.][^/?]*)?([?].*)?$"
                     - "^[^?]*(htaccess|access_log)([.][^/?]*)?([~])?([?].*)
39
                        ?$"
```

#### Prevent leakage of sensitive data

Data breaches involve leakage of sensitive data such as credit card and social security number (SSN). You can add custom regexes for the sensitive data in the *Enforcements safe objects* section.

The following is an example of a WAF CRD configuration for preventing leakage of sensitive data.

```
1 apiVersion: citrix.com/v1
2 kind: waf
3 metadata:
4
      name: wafdataleak
5 spec:
      servicenames:
6
7
           - frontend
8
      application_type:
           - "HTML"
9
       html_error_object: "http://x.x.x.x/crd/error_page.html"
10
11
       security_checks:
           common:
13
             credit_card: "on"
14
       settings:
```

```
common:
15
16
            credit_card:
               card_type: ["visa","amex"]
17
                max_allowed: 1
18
                card_xout: "on"
19
                secure_logging: "on"
21
       enforcements:
22
          common:
23
             safe_object:
24
               - rule:
25
                    name: aadhar
26
                    expression: "[1-9]{
27
    4,4 }
28
    \s[1-9]{
29
    4,4 }
    \s[1-9]{
31
    4,4 }
    11
32
33
                    max_match_len: 19
                    action: ["log","block"]
34
```

### Protect HTML forms from CSRF and form attacks

The following is an example of a WAF CRD configuration for protecting HTML forms from CSRF and form attacks.

```
1 apiVersion: citrix.com/v1
2 kind: waf
3 metadata:
4 name: wafforms
5 spec:
6
    servicenames:
7

    frontend

8
      application_type:
          - "HTML"
9
     html_error_object: "http://x.x.x.x/crd/error_page.html"
11
       security_checks:
12
          html:
            cross_site_scripting: "on"
13
            sql_injection: "on"
14
            form_field_consistency:
15
              action: ["log","block"]
16
             csrf: "on"
17
```

#### **Protect forms and headers**

The following is an example of a WAF CRD configuration for protecting both forms and headers.

```
1 apiVersion: citrix.com/v1
```

```
2 kind: waf
3 metadata:
4 name: wafhdrforms
5 spec:
6
      servicenames:
7

    frontend

8
      application_type:
9
           - "HTML"
       html_page_url: "http://x.x.x.x/crd/error_page.html"
11
       security_checks:
           common:
13
             buffer_overflow: "on"
14
             multiple_headers:
               action: ["block", "log"]
15
16
           html:
17
             cross_site_scripting: "on"
             sql_injection: "on"
18
             form_field_consistency:
19
20
               action: ["log","block"]
             csrf: "on"
21
22
       settings:
23
           common:
24
             buffer_overflow:
25
               max_cookie_len: 409
26
               max_header_len: 4096
27
               max_url_len: 1024
28
       ip_reputation: on
```

#### **Enable basic WAF security checks**

The basic security checks are required to protect any application with minimal effect on performance. It does not require any sessionization. The following is an example of a WAF CRD configuration for enabling basic WAF security checks.

```
1 apiVersion: citrix.com/v1
2 kind: waf
3 metadata:
4
       name: wafbasic
5 spec:
6
       servicenames:
7
           - frontend
8
       security_checks:
9
           common:
10
             allow_url: "on"
             block_url: "on
11
             buffer_overflow: "on"
13
             multiple_headers:
               action: ["block", "log"]
14
15
           html:
16
             cross_site_scripting: "on"
17
              field_format: "on"
```

```
sql_injection: "on"
18
19
              fileupload_type: "on"
20
           json:
              dos: "on"
21
22
              sql_injection: "on"
23
              cross_site_scripting: "on"
24
           xml:
25
              dos: "on"
26
              wsi: "on"
27
              attachment: "on"
28
              format: "on"
29
       relaxations:
            common:
              allow_url:
32
                urls:
                      "^[^?]+[.](html?|shtml|js|gif|jpg|jpeg|png|swf|pif|
                        pdf css csv)$"
                    - "^[^?]+[.](cgi|aspx?|jsp|php|pl)([?].*)?$"
```

#### Enable advanced WAF security check

Advanced security checks such as cookie consistency, allow URL closure, field consistency, and CSRF are resource-intensive (CPU and memory) as they require WAF sessionization. For example, when a form is protected by the WAF, form field information in the response is retained in the system memory. When the client submits the form in the next request, it is checked for inconsistencies before the information is sent to the web server. This process is known as sessionization. The following is an example of a WAF CRD configuration for enabling WAF advanced security checks.

```
1 apiVersion: citrix.com/v1
2 kind: waf
3 metadata:
       name: wafadvanced
4
5 spec:
6
       servicenames:
7
           - frontend
8
       security_checks:
9
           common:
             allow_url: "on"
10
             block_url: "on"
11
             buffer_overflow: "on"
12
13
             content_type: "on"
             cookie_consistency: "on"
14
15
             multiple_headers:
               action: ["log"]
16
17
           html:
             cross_site_scripting: "on"
18
19
             field_format: "on"
             sql_injection: "on"
20
21
             form_field_consistency: "on"
             csrf: "on"
```

```
fileupload_type: "on"
23
24
           json:
25
             dos: "on"
26
             sql_injection: "on"
27
            cross_site_scripting: "on"
28
           xml:
            dos: "on"
29
             wsi: "on"
31
             validation: "on"
             attachment: "on"
33
             format: "on"
34
      settings:
35
          common:
            allow_url:
               closure: "on"
37
```

#### **Enable IP reputation**

The following is an example of a WAF CRD configuration for enabling IP reputation to reject requests that come from IP addresses with bad reputation.

```
1 apiVersion: citrix.com/v1
2 kind: waf
3 metadata:
4 name: wafiprep
5 spec:
6 application_type:
7 - "HTML"
8 servicenames:
9 - frontend
10 ip_reputation: "on"
```

#### Enable IP reputation to reject requests of a particular category

The following is an example of a WAF CRD configuration for enabling IP reputation to reject requests from particular threat categories.

```
1 apiVersion: citrix.com/v1
2 kind: waf
3 metadata:
  name: wafiprepcategory
4
5 spec:
6 application_type:
         - "HTML"
7
     servicenames:
8
9

    frontend

10
      ip_reputation:
11
          action: block
12
          threat-categories:
```

13	- SPAM_SOURCES	
14	- WINDOWS_EXPLOITS	
15	- WEB_ATTACKS	
16	- BOTNETS	
17	- SCANNERS	
18	- DOS	
19	- REPUTATION	
20	- PHISHING	
21	- PROXY	
22	- NETWORK	
23	- CLOUD_PROVIDERS	
24	- MOBILE_THREATS	

#### Protect JSON applications from denial of service attacks

The following is an example of a WAF CRD configuration for protecting the JSON applications from denial of service attacks.

```
1 metadata:
2
   name: wafjsondos
3 spec:
4 servicenames:
5

    frontend

     application_type: JSON
6
7
      json_error_object: "http://x.x.x.x/crd/error_page.json"
8
      security_checks:
9
           json:
            dos: "on"
11
       settings:
12
          json:
13
            dos:
14
               container:
15
                max_depth: 2
16
               document:
17
                max_len: 20000000
18
               array:
19
                max_len: 5
20
               key:
               max_count: 10000
21
22
                max_len: 12
23
               string:
24
                max_len: 1000000
```

#### **Protect RESTful APIs**

The following is an example of a WAF CRD configuration for protecting RESTful APIs from SQL injection, cross-site scripting, and denial of service attacks.

Here, the back-end application or service is purely based on RESTful APIs.

```
1 apiVersion: citrix.com/v1
2 kind: waf
3 metadata:
4
   name: wafjson
5 spec:
      servicenames:
6
7
           - frontend
8
       application_type: JSON
9
       json_error_object: "http://x.x.x.x/crd/error_page.json"
10
       security_checks:
11
           json:
12
             dos: "on"
13
             sql_injection:
14
               action: ["block"]
15
             cross_site_scripting: "on"
16
       settings:
           json:
17
18
             dos:
19
               container:
20
                 max_depth: 5
               document:
21
22
                 max_len: 20000000
23
               array:
24
                 max_len: 10000
25
               key:
26
                 max_count: 10000
27
                 max_len: 128
               string:
28
29
                 max_len: 1000000
```

#### Protect XML applications from denial of service attacks

The following is an example of a WAF CRD configuration for protecting the XML applications from denial of service attacks.

```
1 apiVersion: citrix.com/v1
2 kind: waf
3 metadata:
      name: wafxmldos
4
5 spec:
       servicenames:
6
7
           - frontend
       application_type: XML
8
       xml_error_object: "http://x.x.x.x/crd/error_page.xml"
9
10
       security_checks:
          xml:
11
             dos: "on"
12
13
       settings:
           xml:
14
15
             dos:
16
               attribute:
```

17	max attributes: 1024
18	max name len: 128
19	max value len: 128
20	element:
21	max_elements: 1024
22	max_children: 128
23	max_depth: 128
24	file:
25	max_size: 2123
26	min_size: 9
27	entity:
28	max_expansions: 512
29	<pre>max_expansions_depth: 9</pre>
30	namespace:
31	<pre>max_namespaces: 16</pre>
32	<pre>max_uri_len: 256</pre>
33	soaparray:
34	max_size: 1111
35	cdata:
36	max_size: 65

### Protect XML applications from security attacks

This example provides a WAF CRD configuration for protecting XML applications from the following security attacks:

- SQL injection
- Cross-site scripting
- Validation (schema or message)
- Format
- Denial of service
- Web service interoperability (WSI)

```
1 apiVersion: citrix.com/v1
2 kind: waf
3 metadata:
4 name: wafxml
5 spec:
6 servicenames:
          - frontend
7
  application_type: XML
8
     xml_error_object: "http://x.x.x.x/crd/error_page.json"
9
     security_checks:
11
          xml:
12
            dos: "on"
            sql_injection: "on"
13
            cross_site_scripting: "off"
14
15
            wsi:
              action: ["block"]
16
```

```
17
              validation: "on"
              attachment: "on"
18
19
              format:
20
                action: ["block"]
21
        settings:
22
          xml:
23
              dos:
24
                attribute:
25
                    max_attributes: 1024
                    max_name_len: 128
27
                    max_value_len: 128
28
                element:
29
                    max_elements: 1024
30
                    max_children: 128
31
                    max_depth: 128
32
                file:
33
                    max_size: 2123
34
                    min_size: 9
                entity:
                    max_expansions: 512
37
                    max_expansions_depth: 9
                namespace:
39
                    max_namespaces: 16
40
                    max_uri_len: 256
41
                soaparray:
42
                    max_size: 1111
43
                cdata:
44
                    max_size: 65
45
              wsi:
                checks: ["R1000","R1003"]
46
47
              validation:
48
                soap_envelope: "on"
                validate_response: "on"
49
              attachment:
51
                url:
                    max_size: 1111
52
53
                content_type:
                    value: "crd_test"
54
```

# Configure bot management policies with the NetScaler Ingress Controller

October 17, 2024

A bot is a software application that automates manual tasks. Using Bot management policies you can allow useful bots to access your cloud native environment and block the malicious bots.

Custom Resource Definitions (CRDs) are the primary way of configuring policies in cloud native deploy-

ments. Using the Bot CRD provided by NetScaler, you can configure the bot management policies with the NetScaler Ingress Controller on the NetScaler VPX. The Bot CRD enables communication between the NetScaler Ingress Controller and NetScaler for enforcing bot management policies.

In a Kubernetes deployment, you can enforce bot management policy on the requests and responses from and to the server using the Bot CRD. For more information on security vulnerabilities, see Bot Detection.

With the Bot CRD, you can configure the bot management security policy for the following types of security vulnerabilities for the Kubernetes-native applications:

- Allow list
- Block list
- Device Fingerprint (DFP)
- Bot TPS
- Trap insertion
- IP reputation
- Rate limit

Based on the type of protections required, you can specify the metadata and use the CRD attributes in the Bot CRD . yaml file to define the bot policy.

## **Bot CRD definition**

The Bot CRD is available in the NetScaler Ingress Controller GitHub repo at bot-crd.yaml. The Bot CRD provides attributes for the various options that are required to define the bot management policies on NetScaler.

## **Bot CRD attributes**

The following table lists the various attributes provided in the Bot CRD:

CRD attribute	Description
security_checks	List of security checks to be applied for incoming
allow_list	List of allowed IP, subnet, and policy
block_list	Expressions. List of disallowed IP, subnet, and policy
	expressions.
CRD attribute	Description
--------------------	--
device_fingerprint	Inserts javascript and collects the client browser and device parameters.
trap	Inserts hidden URLs in the response.
tps	Prevents bots which cause unusual spike in requests based on configured percentage change in transactions.
reputation	Prevents access to bad IPs based on configured reputation categories.
ratelimit	Prevents bots based on rate limit.
redirect_url	Redirect URL when block is enabled on protection.
ingressclass	Specifies the ingress class so that only the ingress controller associated with the specified ingress class processes the resource. Otherwise, all the controllers in the cluster will process this
servicenames	resource. Name of the services to which the bot policies are applied.
signatures	Location of external bot signature file.
target	Determines which traffic to be inspected by the bot. If you do not specify the traffic targeted, every traffic is inspected by default.
paths	List of HTTP URLs to be inspected.
method	List of HTTP methods to be inspected.
header	List of HTTP headers to be inspected.

## **Deploy the Bot CRD**

Perform the following steps to deploy the Bot CRD:

- 1. Download the bot-crd.yaml.
- 2. Deploy the Bot CRD using the following command:

## kubectl create -f bot-crd.yaml

For example,

```
1 root@master:~# kubectl create -f bot-crd.yaml
2 customresourcedefinition.apiextensions.k8s.io/bots.citrix.com created
```

## How to write a Bot configuration

After you have deployed the Bot CRD provided by NetScaler in the Kubernetes cluster, you can define the bot management policy configuration in a YAML file. In the YAML file, specify bot in the kind field. In the spec section, add the Bot CRD attributes based on your requirements for the policy configuration.

After you deploy the YAML file, the NetScaler Ingress Controller applies the bot configuration on the Ingress NetScaler device.

Following are some examples for bot policy configurations:

## Block malicious traffic using known IP, subnet, or ADC policy expressions

When you want to define and employ a web bot management policy in NetScaler to enable bot for blocking malicious traffic, you can create a YAML file called botblocklist.yaml and use the appropriate CRD attributes to define the bot policy as follows:

```
1 apiVersion: citrix.com/v1
2 kind: bot
3 metadata:
4
       name: botblocklist
5
   spec:
6
       servicenames:
           - frontend
7
8
       security_checks:
9
            block_list: "ON"
10
       bindings:
           block_list:
11
                - subnet:
13
                    value:
14
                        - 172.16.1.0/12
15
                        - 172.16.2.0/12
16
                        - 172.16.3.0/12
                        - 172.16.4.0/12
17
18
                    action:
                        - "drop"
19
20
                - ip:
21
                    value: 10.102.30.40
22
                - expression:
                    value: http.req.url.contains("/robots.txt")
23
24
                    action:
25
                         - "reset"
26
                        - "log"
```

## Allow known traffic without bot security checks\*\*

When you want to avoid security checks for certain traffic such as staging or trusted traffic, you can avoid such traffic from security checks. You can create a YAML file called botallowlist.yaml and use the appropriate CRD attributes to define the bot policy as follows:

```
apiVersion: citrix.com/v1
1
2 kind: bot
3 metadata:
4
      name: botallowlist
5 spec:
6
      servicenames:
           - frontend
7
       security_checks:
8
           allow_list: "ON"
9
10
       bindings:
           allow_list:
11
                - subnet:
12
13
                    value:
14
                        - 172.16.1.0/12
15
                        - 172.16.2.0/12
16
                        - 172.16.3.0/12
17
                        - 172.16.4.0/12
                    action:
18
                        - "log"
19
20
                - ip:
21
                    value: 10.102.30.40
22
                - expression:
23
                    value: http.req.url.contains("index.html")
24
                    action:
25
                        - "log"
```

## **Enable bot signatures to detect bots**

NetScaler provides thousands of inbuilt signatures to detect bots based on user agents. Citrix threat intelligence team keeps on updating and releasing new bot signatures in every two weeks. The latest bot signature file is available at: Bot signatures. You can create a YAML file called botsignatures .yaml and use the appropriate CRD attributes to define the bot policy as follows:

## Enable the bot device fingerprint and customize the action

Device fingerprinting involves inserting a JavaScript snippet in the HTML response to the client. This JavaScript snippet, when invoked by the browser on the client, collects the attributes of the browser and client. And sends a POST request to NetScaler with that information. These attributes are examined to determine whether the connection is requested from a bot or a human being. You can create a YAML file called botdfp.yaml and use the appropriate CRD attributes to define the bot policy as follows:

```
1 apiVersion: citrix.com/v1
2 kind: bot
3 metadata:
4
     name: botdfp
5 spec:
6
     servicenames:
7

    frontend

8
      redirect_url: "/error_page.html"
9
       security_checks:
10
          device_fingerprint:
11
              action:
12
                  - "log"
                  - "drop"
13
```

## Enable the bot TPS and customize the action

If the bot TPS is configured, it detects incoming traffic as bots if the maximum number of requests or increase in requests exceeds the configured time interval. You can configure the TPS limits as per *geolocation, host, source IP*, and *URL* in the *bindings* section. You can create a YAML file called bottps .yaml and use the appropriate CRD attributes to define the bot policy as follows:

```
1 apiVersion: citrix.com/v1
2 kind: bot
3 metadata:
4
   name: bottps
5 spec:
6 servicenames:
7

    frontend

8
     redirect_url: "/error_page.html"
9
       security_checks:
          tps: "ON"
       bindings:
11
12
          tps:
13
               geolocation:
14
                   threshold: 101
                   percentage: 100
16
               host:
17
                   threshold: 10
18
                   percentage: 100
```

19	action:
20	- "log"
21	- "mitigation"

#### Enable the trap insertion protection and customize the action

Detects and blocks automated bots by advertising a trap URL in the client response. The URL is invisible and not accessible to the client, if it is human. The detection method is effective in blocking attacks from automated bots. Insertion of the trap URL in the URL responses is random. You can enforce the trap URL insertion to a particular URL response by configuring the trap bindings. You can create a YAML file called trapinsertion.yaml and use the appropriate CRD attributes to define the bot policy as follows:

```
1 apiVersion: citrix.com/v1
2 kind: bot
3 metadata:
4 name: trapinsertion
5 spec:
6
     servicenames:
7
           - frontend
8
      redirect_url: "/error_page.html"
9
       security_checks:
         trap:
11
            action:
             - "log"
12
              - "drop"
13
14
       bindings:
       trapinsertion:
16
          urls:
17
            - "/index.html"
             - "/submit.php"
18
             - "/login.html"
19
```

## Enable IP reputation to reject requests of a particular category

The following is an example of a Bot CRD configuration for enabling only specific threat categories of IP reputation that are suitable for the user environment. You can create a YAML file called botiprepcategory.yaml and use the appropriate CRD attributes to define the bot policy as follows:

```
1 apiVersion: citrix.com/v1
2 kind: bot
3 metadata:
4 name: botiprepcategory
5 spec:
6 servicenames:
7 - frontend
```

```
redirect_url: "/error_page.html"
8
9
       security_checks:
10
           reputation: "ON"
11
       bindings:
12
         reputation:
13
           categories:
                - SPAM_SOURCES:
14
15
                    action:
                         - "log"
16
                         - "redirect"
17
18
                - MOBILE_THREATS
                - SPAM_SOURCES
19
```

#### Enable rate limit to control request rate

The following is an example of a Bot CRD configuration for enforcing the request rate limit using the parameters: URL, cookies, and IP. You can create a YAML file called botratelimit.yaml and use the appropriate CRD attributes to define the bot policy as follows:

```
1 apiVersion: citrix.com/v1
2 kind: bot
3 metadata:
4
   name: botratelimit
5 spec:
6 servicenames:
7

    frontend

8 redirect_url: "/error_page.html"
9 security_checks:
       ratelimit: "ON"
   bindings:
11
12
      ratelimit:
13
         - url:
             value: index.html
14
15
             rate: 2000
             timeslice: 1000
16
17
         - cookie:
18
             value: citrix_bot_id
             rate: 2000
19
20
             timeslice: 1000
21
         - ip:
             rate: 2000
22
23
             timeslice: 1000
24
             action:
                 - "log"
25
                 - "reset"
26
```

# Configure cross-origin resource sharing policies with NetScaler Ingress Controller

## October 17, 2024

NetScaler provides a Custom Resource Definition (CRD) called the CORS CRD for Kubernetes. You can use the CORS CRD to configure the cross-origin resource sharing (CORS) policies with NetScaler Ingress Controller on the NetScaler.

## What is CORS

Cross-Origin resource sharing is a mechanism that allows the browser to determine whether a specific web application can share resources with another web application from a different origin. It allows users request resources (For example, images, fonts, and videos) from domains outside the original domain.

## **CORS pre-flight**

Before a web browser allowing Javascript to issue a POST to a URL, it performs a pre-flight request. A pre-flight request is a simple request to the server with the same URL using the method OPTIONS rather than POST. The web browser checks the HTTP headers for CORS related headers to determine if POST operation on behalf of the user is allowed.



## **CORS CRD definition**

The CORS CRD is available in the NetScaler Ingress Controller GitHub repo at: cors-crd.yaml. The CORS CRD provides attributes for the various options that are required to define the CORS policy on the Ingress NetScaler that acts as an API gateway. The required attributes include: servicenames, allow\_origin, allow\_methods, and allow\_headers.

Attribute Description Specifies the ingress class so that only the ingressclass ingress controller associated with the specified ingress class processes the resource. Otherwise, all the controllers in the cluster will process this resource. servicenames Specifies the list of Kubernetes services to which you want to apply the CORS policies. allow\_origin Specifies the list of allowed origins. Incoming origin is screened against this list. allow\_methods Specifies the list of allowed methods as part of the CORS protocol. allow headers Specifies the list of allowed headers as part of the CORS protocol. Specifies the number of seconds the information max age provided by the Access-Control-Allow-Methods and Access-Control-Allow-Headers headers can be cached. The default value is 86400. allow\_credentials Specifies whether the response can be shared when the credentials mode of the request is "include". The default value is 'true'.

The following are the attributes provided in the CORS CRD:

## Deploy the CORS CRD

Perform the following to deploy the CORS CRD:

- 1. Download the CORS CRD.
- 2. Deploy the CORS CRD using the following command:

```
1 kubectl create -f cors-crd.yaml
```

#### For example:

#### How to write a CORS policy configuration

After you have deployed the CORS CRD provided by NetScaler in the Kubernetes cluster, you can define the CORS policy configuration in a .yaml file. In the .yaml file, use corspolicy in the kind field and in the spec section add the CORS CRD attributes based on your requirement for the policy configuration.

The following YAML file applies the configured policy to the services listed in the servicenames field. NetScaler responds with a 200 OK response code for the pre-flight request if the origin is one of the allow\_origins ["random1234.com", "hotdrink.beverages.com"]. The response includes configured allow\_methods, allow\_headers, and max\_age.

```
1 apiVersion: citrix.com/v1beta1
2 kind: corspolicy
3 metadata:
4 name: corspolicy-example
5 spec:
6 servicenames:
7
    - "cors-service"
8 allow_origin:
     - "random1234.com"
9
     - "hotdrink.beverages.com"
10
11
     allow_methods:
    - "POST"
12
      - "GET"
13
     - "OPTIONS"
14
15 allow_headers:
     - "Origin"
      - "X-Requested-With"
17

    "Content-Type"

18
      - "Accept"
19
      - "X-PINGOTHER"
     max_age: 86400
22
     allow_credentials: true
```

After you have defined the policy configuration, deploy the .yaml file using the following commands:

```
1 user@master:~/cors$ kubectl create -f corspolicy-example.yaml
2 corspolicy.citrix.com/corspolicy-example created
```

The NetScaler Ingress Controller applies the policy configuration on the Ingress NetScaler device.

# Enable request retry feature using AppQoE for NetScaler Ingress Controller

#### December 31, 2023

When a NetScaler appliance receives an HTTP request and forwards it to a back-end server, sometimes there may be connection failures with the back-end server. You can configure the request-retry feature on NetScaler to forward the request to the next available server, instead of sending the reset to the client. Hence, the client saves round trip time when NetScaler initiates the same request to the next available service. For more information request retry feature, see the NetScaler documentation

Now, you can configure request retry on NetScaler with NetScaler Ingress Controller. Custom Resource Definitions (CRDs) are the primary way of configuring policies in cloud native deployments. Using the AppQoE CRD provided by NetScaler, you can configure request-retry policies on NetScaler with the NetScaler Ingress Controller. The AppQoE CRD enables communication between the NetScaler Ingress Controller and NetScaler for enforcing AppQoE policies.

## **AppQoE CRD definition**

The AppQoE CRD is available in the NetScaler Ingress Controller GitHub repo at: appqoe-crd.yaml. The AppQoE CRD provides attributes for the various options that are required to define the AppQoE policy on NetScaler.

Attribute	Description
servicenames	Specifies the list of Kubernetes services to which you want to apply the AppQoE policies.
on-reset	Specifies whether to set retry on connection Reset or Not
on-timeout	Specifies the time in milliseconds for retry
number-of-retries	Specifies the number of retries

The following are the attributes provided in the AppQoE CRD:

Attribute	Description
appqoe-criteria	Specifies the expression for evaluating traffic.
direction	Specifies the bind point for binding the AppQoE
	policy.

## Deploy the AppQoE CRD

Perform the following to deploy the AppQoE CRD:

- 1. Download the AppQoE CRD.
- 2. Deploy the AppQoE CRD using the following command:

```
1 kubectl create -f appqoe-crd.yaml
```

## How to write a AppQoE policy configuration

After you have deployed the AppQoE CRD provided by NetScaler in the Kubernetes cluster, you can define the AppQoE policy configuration in a .yaml file. In the .yaml file, use appqoepolicy in the kind field and in the spec section add the AppQoE CRD attributes based on your requirement for the policy configuration.

The following YAML file applies the AppQoE policy to the services listed in the servicenames field. You must configure the AppQoE action to retry on timeout and define the number of retry attempts.

```
apiVersion: citrix.com/v1
1
2
     kind: appgoepolicy
3
    metadata:
4
       name: targeturlappqoe
5
     spec:
       appgoe-policies:
6
7
         - servicenames:
8
             - apache
9
           appqoe-policy:
10
             operation-retry:
               onReset: 'YES'
11
               onTimeout: 33
13
             number-of-retries: 2
14
             appqoe-criteria: 'HTTP.REQ.HEADER("User-Agent").CONTAINS("
                 Android")'
15
             direction: REQUEST
```

After you have defined the policy configuration, deploy the .yaml file using the following commands:

\$ kubectl create -f appqoe-example.yaml

# Configuring wildcard DNS domains through NetScaler Ingress Controller

## December 31, 2023

Wildcard DNS domains are used to handle requests for non-existent domains and subdomains. In a DNS zone, you can use wildcard domains to redirect queries for all non-existent domains or subdomains to a particular server, instead of creating a separate Resource Record (RR) for each domain. The most common use of a wildcard DNS domain is to create a zone that can be used to forward mail from the internet to some other mail system.

For more information on wildcard DNS domains, see the NetScaler documentation.

Now, you can configure wildcard DNS domains on a NetScaler with NetScaler Ingress Controller. Custom Resource Definitions (CRDs) are the primary way of configuring policies in cloud native deployments. Using the Wildcard DNS CRD provided by NetScaler, you can configure wildcard DNS domains on a NetScaler with the NetScaler Ingress Controller. The Wildcard DNS CRD enables communication between NetScaler Ingress Controller and NetScaler for supporting wild card domains.

## **Usage guidelines and restrictions**

- For fully qualified domain names (FQDNs), there are multiple ways to add DNS records. You can either enable the NS\_CONFIG\_DNS\_REC variable for NetScaler Ingress Controller for the Ingress resource or use the wildcard DNS CRD. However, you should make sure that they are configured through either CRD or ingress in order to avoid multiple IP mappings to the same domain.
- It is recommended to use the Wildcard DNS CRD for the wildcard DNS configurations.
- You cannot configure wildcard DNS entries in the DNS address record through ingress if the NS\_CONFIG\_DNS\_REC is enabled for NetScaler Ingress Controller.

## Wildcard DNS CRD definition

The Wildcard DNS CRD is available in the NetScaler Ingress Controller GitHub repo at wildcarddnsentry.yaml. The **Wildcard DNS CRD provides** attributes for the various options that are required to configure wildcard DNS entries on NetScaler.

The following are the attributes provided in the Wildcard DNS CRD:

Attribute	Description
domain	Specifies the wild card domain name configured
	for the zone.
dnsaddrec	Specifies the DNS Address record with the IPv4
	address of the wildcard domain.
dnsaaaarec	Specifies the DNS AAAA record with the IPV6
	address of the wildcard domain.
soarec	Specifies the SOA record configuration details.
nsrec	Specifies the name server configuration details.

## **Deploy the Wildcard DNS CRD**

Perform the following to deploy the Wildcard DNS CRD:

- 1. Download the Wildcard DNS CRD.
- 2. Deploy the Wildcard DNS CD using the following command:

```
1 kubectl create -f wildcarddnsentry.yaml
```

## How to write a Wildcard DNS configuration policy

After you have deployed the Wildcard DNS CRD provided by NetScaler in the Kubernetes cluster, you can define the wildcard DNS related configuration in a yaml file. In the .yaml file, use wildcarddnsentry in the kind field and in the spec section add the Wildcard DNS CRD attributes based on your requirement for the policy configuration.

The following is a sample YAML file definition that configures a SOA record, NS record, DNS zone, and address and AAAA Records on NetScaler.

```
1 apiVersion:
2 citrix.com/v1
3 kind: wildcarddnsentry
4 metadata:
5
  name: sample-config
6 spec:
7 zone:
8
     domain: configexample
9
     dnsaddrec:
10
        domain-ip: 1.1.1.1
11
        ttl: 3600
12
      dnsaaaarec:
13
        domain-ip: '2001::.1'
```

```
14
         ttl: 3600
15
       soarec:
         origin-server: n2.configexample.com
16
         contact: admin.configexample.com
17
18
         serial: 100
         refresh: 3600
19
20
         retry: 3
21
         expire: 3600
22
       nsrec:
23
         nameserver: n1.configexample.com
24
         ttl: 3600
```

After you have defined the DNS configuration, deploy the wildcarddns-example.yaml file using the following command.

1 kubectl create -f wildcarddns-example.yaml

## View metrics of NetScalers using Prometheus and Grafana

#### December 31, 2023

You can use the NetScaler Metrics Exporter and Prometheus-Operator to monitor NetScaler VPX or CPX ingress devices and NetScaler CPX (east-west) devices.

#### **NetScaler Metrics Exporter**

NetScaler Metrics Exporter is a simple server that collects NetScaler stats and exports them to Prometheus using HTTP. You can then add Prometheus as a data source to Grafana and graphically view the NetScaler stats. For more information see, NetScaler Metrics Exporter.

Note:

NetScaler Metrics Exporter supports exporting metrics from the admin partitions of NetScaler.

#### Launch prometheus operator

The Prometheus Operator has an expansive method of monitoring services on Kubernetes. To get started, this topic uses kube-prometheus and its manifest files. The manifest files help you to deploy a basic working model. Deploy the Prometheus Operator in your Kubernetes environment using the following commands:

```
1 git clone https://github.com/coreos/kube-prometheus.git
2 kubectl create -f kube-prometheus/manifests/setup/
```

## 3 kubectl create -f kube-prometheus/manifests/

Once you deploy Prometheus-Operator, several pods and services are deployed. From the deployed pods, the prometheus-k8s-xx pods are for metrics aggregation and timestamping, and the grafana pods are for visualization. If you view all the container images running in the cluster, you can see the following output:

1	\$ kubectl get pods -n monitoring				
2	NAME	READY	STATUS	RESTARTS	
	AGE				
3	alertmanager-main-0	2/2	Running	Θ	2
4	n	2/2	Durantan	0	2
4	alertmanager-main-1	2/2	Running	0	2
5	alertmanager-main-2	2/2	Running	Θ	2
J	h	2/2	Running	0	2
6	grafana-5b68464b84-5fvxg	1/1	Running	Θ	2
	h		0		
7	kube-state-metrics-6588b6b755-d6ftg	4/4	Running	Θ	2
	h				
8	node-exporter-4hbcp	2/2	Running	Θ	2
	h				
9	node-exporter-kn9dg	2/2	Running	Θ	2
1.0	h nada avecatas tavias	2 / 2	Duranian	0	2
10	hode-exporter-tpxnp	2/2	Running	0	2
11	nrometheus-k8s-0	3/3	Running	1	2
± ±	h	5/5	Running	1	2
12	prometheus-k8s-1	3/3	Running	1	2
	h	-,-			
13	prometheus-operator-7d9fd546c4-m8t7v	1/1	Running	Θ	2
	h				

## Note:

The files in the manifests folder are interdependent and hence the order in which they are created is important. In certain scenarios the manifest files might be created out of order and this leads to an error messages from Kubernetes.

To resolve this scenario, re-execute the kubectl create -f kube-prometheus/ manifests/ command. Any YAML files that were not created the first time due to unmet dependencies, are created now.

It is recommended to expose the Prometheus and Grafana pods through NodePorts. To do so, you need to modify the prometheus-service.yaml and grafana-service.yaml files as follows:

## **Modify Prometheus service**

```
1 apiVersion: v1
2 kind: Service
3 metadata:
4
    labels:
5
      prometheus: k8s
6 name: prometheus-k8s
7 namespace: monitoring
8 spec:
9
   type: NodePort
    ports:
11
    - name: web
12
      port: 9090
13
      targetPort: web
14 selector:
15
    app: prometheus
16
      prometheus: k8s
```

After you modify the prometheus-service.yamlfile, apply the changes to the Kubernetes cluster using the following command:

```
1 kubectl apply -f prometheus-service.yaml
```

#### **Modify Grafana service**

```
1 apiVersion: v1
2 kind: Service
3 metadata:
4
   name: grafana
   namespace: monitoring
5
6 spec:
7
   type: NodePort
8 ports:
9
    - name: http
10
      port: 3000
11
      targetPort: http
12
   selector:
13
      app: grafana
```

After you modify the grafana-service.yamlfile, apply the changes to the Kubernetes cluster using the following command:

1 kubectl apply -f grafana-service.yaml

## **Configure NetScaler Metrics Exporter**

This topic describes how to integrate the NetScaler Metrics Exporter with NetScaler VPX or CPX ingress or NetScaler CPX (east-west) devices.

## Configure NetScaler Metrics Exporter for NetScaler VPX Ingress device

To monitor an ingress NetScaler VPX device, the NetScaler Metrics Exporter is run as a pod within the Kubernetes cluster. The IP address of the NetScaler VPX ingress device is provided as an argument to the NetScaler Metrics Exporter. To provide the login credentials to access ADC, create a secret and mount the volume at mountpath "/mnt/nslogin".

```
1 kubectl create secret generic nslogin --from-literal=username=<citrix-
adc-user> --from-literal=password=<citrix-adc-password> -n <
namespace>
```

The following is a sample YAML file to deploy the exporter:

```
1 apiVersion: v1
2 kind: Pod
3 metadata:
    name: exporter-vpx-ingress
4
5
    labels:
6
      app: exporter-vpx-ingress
7 spec:
8 containers:
9
      - name: exporter
         image: "quay.io/citrix/citrix-adc-metrics-exporter:1.4.8"
10
         imagePullPolicy: IfNotPresent
11
12
         args:
13
          - "--target-nsip=<IP_of_VPX>"
           - "--port=8888"
14
15
         volumeMounts:
16
         - name: nslogin
         mountPath: "/mnt/nslogin"
17
18
          readOnly: true
19
         securityContext:
           readOnlyRootFilesystem: true
20
21
   volumes:
    - name: nslogin
22
23
       secret:
24
         secretName: nslogin
25 ---
26 kind: Service
27 apiVersion: v1
28 metadata:
29
   name: exporter-vpx-ingress
     labels:
31
       service-type: citrix-adc-monitor
32 spec:
33
  selector:
34
      app: exporter-vpx-ingress
35
   ports:
      - name: exporter-port
37
         port: 8888
38
        targetPort: 8888
```

The IP address and the port of the NetScaler VPX device needs to be provided in the --target-nsip parameter. For example, --target-nsip=10.0.0.20.

## Configure NetScaler Metrics Exporter for NetScaler CPX Ingress device

To monitor a NetScaler CPX ingress device, the NetScaler Metrics Exporter is added as a sidecar to the NetScaler CPX.The following is a sample YAML file of a NetScaler CPX ingress device with the exporter as a side car:

1	
2	apiVersion: apps/v1
3	kind: Deployment
4	metadata:
5	labels:
6	app: cpx-ingress
7	name: cpx-ingress
8	spec:
9	replicas: 1
10	selector:
11	matchLabels:
12	app: cpx-ingress
13	template:
14	metadata:
15	annotations:
16	NETSCALER_AS_APP: "True"
17	labels:
18	app: cpx-ingress
19	spec:
20	containers:
21	- env:
22	- name: EULA
23	value: "YES"
24	- name: NS_PROTOCOL
25	value: HITP
26	- name: NS_PORT
27	value: "9080"
28	#Define the NIIRU port here
29	<pre>image: quay.io/citrix/citrix-k8s-cpx-ingress:13.0-52.24 image.pullDaliaut.ffNatDuraget</pre>
30	imagePullPolicy: ITNotPresent
31	name: cpx-ingress
32	ports:
33	- containerPort: 80
34 25	name: nup
30	protocol: TCP
30	- containerPort: 443
20	name. https protocol: TCP
20	- containerPort: 0000
39	name: nitro-http
40	protocol. TCP
41 42	- containerPort: 9442
42	Container Fort. 3443

```
43
           name: nitro-https
44
            protocol: TCP
45
          securityContext:
           privileged: true
46
47
         # Adding exporter as a sidecar
48
        - args:
49
          - --target-nsip=192.0.0.2
           - --port=8888
50
51
          - --secure=no
52
          env:
          - name: NS_USER
54
            value: nsroot
55
          - name: NS_PASSWORD
56
            value: nsroot
          image: quay.io/citrix/citrix-adc-metrics-exporter:1.4.8
57
58
          imagePullPolicy: IfNotPresent
59
         name: exporter
         securityContext:
61
         readOnlyRootFilesystem: true
62
         serviceAccountName: cpx
63 ---
64 kind: Service
65 apiVersion: v1
66 metadata:
67 name: exporter-cpx-ingress
68 labels:
69
    service-type: citrix-adc-monitor
70 spec:
   selector:
71
72
      app: cpx-ingress
73
   ports:
74 - name: exporter-port
75
       port: 8888
76
       targetPort: 8888
```

Here, the exporter uses the local IP address (192.0.0.2) to fetch metrics from the NetScaler CPX.

## Configure NetScaler Metrics Exporter for NetScaler CPX (east-west) device

To monitor a NetScaler CPX (east-west) device, the NetScaler Metrics Exporter is added as a sidecar to the NetScalerCPX.The following is a sample YAML file of a NetScaler CPX (east-west) device with the exporter as a side car:

```
1 apiVersion: apps/v1
2 kind: DaemonSet
3 metadata:
4 annotations:
5 deprecated.daemonset.template.generation: "0"
6 labels:
7 app: cpx-ew
8 name: cpx-ew
```

```
9 spec:
10 selector:
    matchLabels:
11
12
      app: cpx-ew
13 template:
    metadata:
14
15
       annotations:
16
           NETSCALER_AS_APP: "True"
         labels:
17
18
           app: cpx-ew
19
        name: cpx-ew
20
      spec:
21
        containers:
         - env:
23
           - name: EULA
24
            value: "yes"
25
           - name: NS_NETMODE
            value: HOST
26
           #- name: "kubernetes_url"
27
           # value: "https://10..xx.xx:6443"
28
29
           image: quay.io/citrix/citrix-k8s-cpx-ingress:13.0-52.24
           imagePullPolicy: IfNotPresent
31
           name: cpx
32
          securityContext:
33
             privileged: true
34
         # Add exporter as a sidecar
35
         - args:
          - --target-nsip=192.168.0.2
           - --port=8888
37
           - --secure=no
38
39
          env:
40
           - name: NS_USER
41
            value: nsroot
42
           - name: NS_PASSWORD
43
            value: nsroot
44
           image: quay.io/citrix/citrix-adc-metrics-exporter:1.4.8
45
           imagePullPolicy: IfNotPresent
46
           name: exporter
47
           securityContext:
48
             readOnlyRootFilesystem: true
49
         serviceAccountName: cpx
50 ---
51 kind: Service
52 apiVersion: v1
53 metadata:
   name: exporter-cpx-ew
54
   labels:
55
56
     service-type: citrix-adc-monitor
57 spec:
58 selector:
  app: cpx-ew
59
     ports:
61
   - name: exporter-port
```

62 port: 8888 63 targetPort: 8888

Here, the exporter uses the local IP (192.168.0.2) to fetch metrics from the NetScaler CPX (east-west) device.

## ServiceMonitors to detect NetScaler

The NetScaler Metrics Exporter helps collect data from the NetScaler VPX or CPX ingress and NetScaler CPX (east-west) devices. The Prometheus Operator needs to detect these exporters so that the metrics can be timestamped, stored, and exposed for visualization on Grafana. The Prometheus Operator uses the concept of ServiceMonitors to detect pods that belong to a service, using the labels attached to that service.

The following example YAML file detects all the exporter services (given in the sample YAML files) which have the label service-type: citrix-adc-monitor associated with them.

```
apiVersion: monitoring.coreos.com/v1
2
  kind: ServiceMonitor
3 metadata:
   name: citrix-adc-servicemonitor
4
5
   labels:
      servicemonitor: citrix-adc
6
7 spec:
8
    endpoints:
    - interval: 30s
9
       port: exporter-port
11 selector:
12
     matchLabels:
13
        service-type: citrix-adc-monitor
14 namespaceSelector:
15
      matchNames:
16
       - monitoring
17
       - default
```

The ServiceMonitor directs Prometheus to detect Exporters in the **default** and monitoring namespaces only. To detect Exporters from other namespaces add the names of those namespaces under the namespaceSelector: field.

Note:

If the Exporter that needs to be monitored exists in a namespace other than the **default** or **monitoring** namespace, then additional RBAC privileges must be provided to Prometheus to access those namespaces. The following is sample YAML (prometheus-clusterRole. yaml) file the provides Prometheus full access to resources across the namespaces:

1 apiVersion: rbac.authorization.k8s.io/v1

```
2 kind: ClusterRole
3 metadata:
4 name: prometheus-k8s
5 rules:
6 - apiGroups:
    _____
7
8
    resources:
     - nodes/metrics
9
10
    - namespaces
    - services
    - endpoints
    - pods
13
   verbs: ["*"]
14
15 - nonResourceURLs:
16
     - /metrics
17 verbs: ["*"]
```

To provide additional privileges Prometheus, deploy the sample YAML using the following command:

1 kubectl apply -f prometheus-clusterRole.yaml

#### View the metrics in grafana

The NetScaler instances that are detected for monitoring appears in the **Targets** page of the prometheus container. You can be access the **Targets** page using the following URL: http://<k8s\_cluster\_ip>:<prometheus\_nodeport>/targets:

Prometheus Alerts Graph Status - Help					
Targets          All       Unhealthy         citrix-adc/citrix-adc-servicemonitor/0 (2/2 up)       show less					
Endpoint	State	Labels	Last Scrape	Scrape Duration	Error
http://10.244.1.97:8888/metrics	UP	endpoint="exporter.port" instance="10.244.1.97:8888" job="ex porter.kibana" namespace="citrix.adc" pod="cpx.ingress.kiban a.77f8bb4d5f-96dp8" service="exporter.kibana"	26.683s ago	176.9ms	
http://10.244.2.117:8888/metrics	UP	endpoint-"exporter-port" instance-"10.244.2,117:8888" job-"e xporter-customerportal" namespace="citrix-adc" pod="cpx-ingr ess-customerportal-6bcc664498.kt6m8" service-"exporter-custom erportal"	4.941s ago	193ms	

To view the metrics graphically:

- 1. Log into grafana using http://<k8s\_cluster\_ip>:<grafafa\_nodeport> with default credentials *admin:admin*
- 2. On the left panel, select + and click **Import** to import the sample grafana dashboard.



A dashboard containing the graphs similar to the following appears:



You can further enhance the dashboard using Grafana's documentation or demo videos.

# Analytics and observability

December 31, 2023

Analytics from NetScaler instances provides you deep-level insights about application performance which helps you to quickly identify issues and take any necessary action.

## Enabling analytics using annotations in the NetScaler Ingress Controller YAML file

You can enable analytics using the analytics profile which is defined as a smart annotation in Ingress or service of type LoadBalancer configuration. You can define the specific parameters you need to monitor by specifying them in the Ingress or service configuration of the application. The following is a sample Ingress annotation with analytics profile for HTTP records:

```
ingress.citrix.com/analyticsprofile: '{ "webinsight": { "httpurl":"
ENABLED", "httpuseragent":"ENABLED", "httpHost":"ENABLED", "httpMethod
":"ENABLED", "httpContentType":"ENABLED" } } '
```

The following is a sample Ingress configuration with the analytics profile for a web application.

```
1 apiVersion: networking.k8s.io/v1
2 kind: Ingress
3 metadata:
    annotations:
4
5
       ingress.citrix.com/analyticsprofile: '{
   "webinsight": {
6
   "httpurl":"ENABLED", "httpuseragent":"ENABLED",
7
         "httphost":"ENABLED", "httpmethod":"ENABLED", "httpcontenttype":"
8
             ENABLED" }
9
    }
    1
       ingress.citrix.com/insecure-termination: allow
11
     name: webserver-ingress
12
13 spec:
14
    rules:
15
     - http:
         paths:
16
17
         - backend:
18
             service:
19
               name: webserver
20
               port:
21
                 number: 80
22
           path: /
23
           pathType: Prefix
24
     tls:
25
     - secretName: name
```

The following is a service annotation:

```
service.citrix.com/analyticsprofile: '{ "80-tcp":{ "webinsight": { "
httpurl":"ENABLED", "httpuseragent":"ENABLED" } } '
```

The following is a sample service configuration with the analytics profile which exposes an Apache web application.

```
1 apiVersion: v1
2 kind: Service
3 metadata:
4
   name: apache
5
   annotations:
      service.citrix.com/csvserver: '{
6
7 "l2conn":"on" }
   1
8
9
       service.citrix.com/lbvserver: '{
10
   "80-tcp":{
    "lbmethod":"SRCIPDESTIPHASH" }
11
12
    }
    1
13
14
       service.citrix.com/servicegroup: '{
15
    "80-tcp":{
16
    "usip":"yes" }
17
    }
    i.
18
       service.citrix.com/monitor: '{
19
20
    "80-tcp":{
21
    "type":"http" }
22
    }
    1
23
24
       service.citrix.com/frontend-ip: "192.0.2.16"
25
       service.citrix.com/analyticsprofile: '{
26
    "80-tcp":{
27
    "webinsight": {
    "httpurl":"ENABLED", "httpuseragent":"ENABLED" }
28
29
    }
    }
    1
31
32
       NETSCALER_VPORT: "80"
33
    labels:
34
      name: apache
35 spec:
   externalTrafficPolicy: Local
37
   type: LoadBalancer
38 selector:
39
     name: apache
40
     ports:
41
     - name: http
42
       port: 80
43
       targetPort: http
44
     selector:
45
       app: apache
```

For information about annotations, see the annotation documentation.

## **Analytics using NetScaler ADM**

NetScaler ADM provides a comprehensive observability solution including analytics on various events happening in the system and a service graph for monitoring services in an easy to use user interface.

NetScaler ADM analytics provide an easy and scalable way to get various insights out of the data from NetScaler instances to describe, predict, and improve the application performance. You can use one or more analytics features simultaneously on NetScaler ADM. For more information on the service graph, see the service graph documentation.

To use the ADM analytics or service graph:

- You must install an ADM agent and ensure the communication between NetScaler ADM and Kubernetes cluster or managed instances in your data center or cloud. It makes NetScaler instances discoverable by NetScaler ADM.
- Ensure that an appropriate license is available and auto licensing is enabled on ADM.

## Analytics with open source tools

NetScaler can be integrated with various open source tools for observability using NetScaler Observability Exporter. NetScaler Observability Exporter is a container which collects metrics and transactions from NetScalers and transforms them to suitable formats (such as JSON, AVRO) for supported endpoints. You can export the collected data to the desired endpoint. By analyzing the data, you can get valuable insights at a microservice level for applications proxied by NetScalers. For more information on NetScaler Observability Exporter, see the NetScaler Observability Exporter

documentation.

# Analytics configuration support using ConfigMap

## August 21, 2024

You can use NetScaler Observability Exporter to export metrics and transactions from NetScaler CPX, MPX, or VPX and analyze the exported data to get meaningful insights. You can now enable the NetScaler Observability Exporter configuration within the NetScaler Ingress Controller using a ConfigMap.

## Schema for NS\_ANALYTICS\_CONFIG

The schema for NS\_ANALYTICS\_CONFIG is as follows. The attributes in the ConfigMap must conform to this schema.

1	type: map
2	mapping:
3	NS_ANALYTICS_CONFIG:
4	required: false
5	type: map
6	mapping:
7	endpoint:
8	required: false
9	type: map
10	mapping:
11	metrics:
12	required: false
13	type: map
14	mapping:
15	service:
16	required: false
17	type: str
18	transactions:
19	required: false
20	type: map
21	mapping:
22	service:
23	required: false
24	type: str
25	distributed_tracing:
26	required: false
27	type: map
28	mapping:
29	enable:
30	required: true
31	type: str
32	enum:
33	- · true·
34 25	
30	samptingrate:
27	type: int
20	range.
30	max: 100
40	min: 0
41	timeseries.
42	required: false
43	type: map
44	mapping:
45	port:
46	required: false
47	type: int
48	metrics:
49	required: false
50	type: map
51	mapping:
52	enable:
53	required: true

55	type: str
00	enum:
56	- 'true'
57	- 'false'
58	mode:
59	required: <b>true</b>
60	type str
61	cype. sci
62	enum.
62	- prometneus
63	- avro
64	- influx
65	export_frequency:
66	required: false
67	type: int
68	range:
69	max: 300
70	min: 30
71	schema_file:
72	required: <b>false</b>
73	type: str
74	enable native scrape:
75	required: <b>false</b>
76	type: str
77	enum.
78	
79	- 'false'
00	auditlogs!
01	roquirod: falso
01	type: man
02	cype: map
83	mapping:
84	enable:
85	required: <b>true</b>
0.0	
86	type: str
86 87	type: str enum:
86 87 88	type: str enum: _ 'true'
86 87 88 89	type: str enum: – 'true' – 'false'
86 87 88 89 90	type: str enum: – 'true' – 'false' events:
86 87 88 89 90 91	type: str enum: - 'true' - 'false' events: required: <b>false</b>
86 87 88 89 90 91 92	type: str enum: - 'true' - 'false' events: required: <b>false</b> type: map
86 87 88 89 90 91 92 93	type: str enum: – 'true' – 'false' events: required: <b>false</b> type: map mapping:
86 87 88 89 90 91 92 93 94	type: str enum: - 'true' - 'false' events: required: <b>false</b> type: map mapping: enable:
86 87 88 89 90 91 92 93 94 95	type: str enum: - 'true' - 'false' events: required: false type: map mapping: enable: required: true
86 87 88 89 90 91 92 93 94 95 96	type: str enum: - 'true' - 'false' events: required: <b>false</b> type: map mapping: enable: required: <b>true</b> type: str
86 87 88 90 91 92 93 94 95 96 97	<pre>type: str enum: - 'true' - 'false' events: required: false type: map mapping: enable: required: true type: str enum:</pre>
86 87 88 89 90 91 92 93 94 95 96 97 98	<pre>type: str enum: - 'true' - 'false' events: required: false type: map mapping: enable: required: true type: str enum: - 'true'</pre>
86 87 88 89 90 91 92 93 94 95 96 97 98 99	<pre>type: str enum: - 'true' - 'false' events: required: false type: map mapping: enable: required: true type: str enum: - 'true' - 'false'</pre>
86 87 88 89 90 91 92 93 94 95 96 97 98 99	<pre>type: str enum: - 'true' - 'false' events: required: false type: map mapping: enable: required: true type: str enum: - 'true' - 'false' transactions:</pre>
86 87 88 89 90 91 92 93 94 95 96 97 98 99 100	<pre>type: str enum: - 'true' - 'false' events: required: false type: map mapping: enable: required: true type: str enum: - 'true' - 'false' transactions: required: false</pre>
86 87 88 89 90 91 92 93 94 95 96 97 98 99 100 101	<pre>type: str enum: - 'true' - 'false' events: required: false type: map mapping: enable: required: true type: str enum: - 'true' - 'false' transactions: required: false type: map</pre>
86 87 88 89 90 91 92 93 94 95 96 97 98 99 100 101 102	<pre>type: str enum: - 'true' - 'false' events: required: false type: map mapping: enable: required: true type: str enum: - 'true' - 'false' transactions: required: false type: map mapping: enable: required: false type: map mapping: enable: required: false type: map</pre>
86 87 88 89 90 91 92 93 94 95 96 97 98 99 100 101 102 103	<pre>type: str enum: - 'true' - 'false' events: required: false type: map mapping: enable: required: true type: str enum: - 'true' - 'false' transactions: required: false type: map mapping: areal.e.</pre>
86 87 88 89 90 91 92 93 94 95 96 97 98 99 100 101 102 103 104	<pre>type: str enum: - 'true' - 'false' events: required: false type: map mapping: enable: required: true type: str enum: - 'true' - 'false' transactions: required: false type: map mapping: enable: } }</pre>
86 87 88 89 90 91 92 93 94 95 96 97 98 99 100 101 102 103 104	<pre>type: str enum: - 'true' - 'false' events: required: false type: map mapping: enable: required: true type: str enum: - 'true' - 'false' transactions: required: false type: map mapping: enable: required: true</pre>

107	enum:
108	- 'true'
109	- 'false'
110	port:
111	required: false
112	type: int

#### Notes:

- Starting from NSIC release 2.0.x, the endpoint.server parameter has been replaced with the endpoint.metrics.service parameter.
- Starting from NSIC release 2.0.x, the endpoint.service parameter has been replaced with the endpoint.transactions.service parameter.
- If you are upgrading from any NSIC Helm chart version 1.x to 2.x, note down the metrics\_endpoint(analyticsConfig.endpoint.server) and transactions\_endpoint (analyticsConfig.endpoint.service) values, and then upgrade using the following set parameters in Helm upgrade: --set analyticsConfig.endpoint.metrics. service=<metrics\_endpoint>,analyticsConfig.endpoint.transactions .service=<transactions\_endpoint>.

## Creating a ConfigMap for analytics configuration

You can create a ConfigMap for analytics configuration during NSIC installation or upgrade using the Helm chart or the NetScaler Operator. For information about the parameters that need to be configured to create a ConfigMap, see Configuration.

Following is a sample values.yaml for a ConfigMap. For information on supported variables, see Supported environment variables for analytics configuration using ConfigMap.

```
1
    analyticsConfig:
2
    required: true
3
  distributedTracing:
4
     enable: true
5
      samplingrate: 100
6 endpoint:
7
     metrics:
         service: "1.1.1.1"
8
9
       transactions:
         service: "1.1.1.2"
10
   timeseries:
11
12
     port: 5563
13
       metrics:
         enable: true
14
15
         mode: 'prometheus'
16
         exportFrequency: 30
17
         schemaFile: schema.json
         enableNativeScrape: false
18
```

```
19 auditlogs:
20 enable: true
21 events:
22 enable: true
23 transactions:
24 enable: true
25 port: 5557
```

For more information on how to configure ConfigMap support on the NetScaler Ingress Controller, see Configuring ConfigMap support for the NetScaler Ingress Controller.

## Supported environment variables for analytics configuration using ConfigMap

You can configure the following parameters under NS\_ANALYTICS\_CONFIG using a ConfigMap:

- distributed\_tracing: This variable enables or disables OpenTracing in NetScaler and has the following attributes:
  - enable: Set this value to **true** to enable OpenTracing. The default value is **false**.
  - samplingrate: Specifies the OpenTracing sampling rate in percentage. The default value is 100.

For more information, see Distributed tracing.

- endpoint: Specifies the IP address or DNS address of the observability endpoint.
  - metrics:
    - \* service: Set this value as the IP address or DNS address of the observability endpoint.
  - transactions:
    - \* service: Set this value as the IP address or namespace/service of the NetScaler Observability Exporter service.
- timeseries: Enables exporting time series data from NetScaler. You can specify the following attributes for time series configuration.
  - port: Specifies the port number of time series end point of the analytics server. The default value is 5563.
  - metrics: Enables exporting metrics from NetScaler.
    - \* enable: Set this value to true to enable sending metrics. The default value is false.
    - \* mode: Specifies the mode of metric endpoint. The possible values are avro and prometheus.

- auditlogs: Enables exporting audit log data from NetScaler.
  - \* enable: Set this value to **true** to enable audit log data. The default value is **false**.
- events: Enables exporting events from the NetScaler.
  - enable: Set this value to true to enable exporting events. The default value is false.
- transactions: Enables exporting transactions from NetScaler.
  - enable: Set this value to true to enable sending transactions. The default value is false.
  - port: Specifies the port number of the transactional endpoint of analytics server. The default value is 5557.

You can change the ConfigMap settings at runtime while the NetScaler Ingress Controller is running.

The attributes of NS\_ANALYTICS\_CONFIG should follow a well-defined schema. If any value provided does not confirm with the schema, then the entire configuration is rejected. For reference, see the schema file ns\_analytics\_config\_schema.yaml.

## IP address management using the IPAM controller

## December 16, 2024

The IPAM controller is a container provided by NetScaler for IP address management and it runs in parallel to NetScaler Ingress Controller as a pod in the Kubernetes cluster. For services of type Load-Balancer, you can use the IPAM controller to automatically allocate IP addresses to services from a specified IP address range. You can specify this IP range in the YAML file while deploying the IPAM controller using YAML. NetScaler Ingress Controller configures the IP address allocated to the service as a virtual IP address (VIP) in NetScaler MPX or VPX.

Using this IP address, you can externally access the service.

## **Overview of services of type LoadBalancer**

In a Kubernetes environment, a microservice is deployed as a set of pods that are created and destroyed dynamically. Since the set of pods that refer to a microservice are constantly changing, Kubernetes provides a logical abstraction known as service to expose your microservice running on a set of pods. A service defines a logical set of pods, and policies to access them.

A service of type LoadBalancer is the simplest way to expose a microservice inside a Kubernetes cluster to the external world. Services of type LoadBalancer are natively supported in Kubernetes deployments on public clouds such as, AWS, GCP, or Azure. In cloud deployments, when you create a service of type LoadBalancer, a cloud managed load balancer is assigned to the service. The service is then exposed using the load balancer.

## NetScaler IPAM solution for services of type LoadBalancer

There may be several situations where you want to deploy your Kubernetes cluster on bare metal or on-premises rather than deploy it on public cloud. When you are running your applications on bare metal Kubernetes clusters, it is much easier to route TCP or UDP traffic using a service of type LoadBalancer than using ingress. Even for HTTP traffic, it is sometimes more convenient than ingress. However, there is no load balancer implementation natively available for bare metal Kubernetes clusters. NetScaler provides a way to load balance such services using the NetScaler Ingress Controller and NetScaler.

In the NetScaler solution for services of type LoadBalancer, NetScaler Ingress Controller deployed inside the Kubernetes cluster configures a NetScaler deployed outside the cluster to load balance the incoming traffic. Using the NetScaler solution, you can load balance the incoming traffic to the Kubernetes cluster regardless of whether the deployment is on bare metal, on-premises, or public cloud. Since the NetScaler Ingress Controller provides flexible IP address management that enables multitenancy for NetScalers, you can use a single NetScaler to load balance multiple services as well as to perform ingress functions. Hence, you can maximize the utilization of load balancer resources and significantly reduce your operational expenses.

## IP address management using the IPAM controller

NetScaler IPAM controller requires the VIP CRD provided by NetScaler. The VIP CRD contains fields for service-name, namespace, and IP address. The VIP CRD is used for internal communication between the NetScaler Ingress Controller and the IPAM controller.

The following diagram shows a deployment of service type load balancer where the IPAM controller is used to assign an IP address to a service.



When a new service of type Loadbalancer is created, the following events occur:

- 1. The NetScaler Ingress Controller creates a VIP CRD object for the service whenever the loadBalancerIP field in the service is empty.
- 2. The IPAM controller assigns an IP address for the VIP CRD object.
- 3. Once the VIP CRD object is updated with the IP address, the NetScaler Ingress Controller automatically configures the NetScaler.

## Note:

Custom resource definitions (CRDs) offered by NetScaler

also support services of type LoadBalancer. That means, you can specify a service of type LoadBalancer as a service name when you create a CRD object and apply the CRD to the service.

# Expose services of type LoadBalancer with IP addresses assigned by the IPAM controller

This topic provides information on how to expose services of type LoadBalancer with IP addresses assigned by the IPAM controller.

To expose a service of type load balancer with an IP address from the IPAM controller, perform the following steps:

- 1. Deploy NetScaler Ingress Controller.
- 2. Deploy NetScaler IPAM controller.
- 3. Deploy a sample application.
- 4. Expose the sample application using service of type LoadBalancer.
- 5. Access the service.

#### **Deploy NetScaler Ingress Controller**

Perform the following steps to deploy the NetScaler Ingress Controller with the IPAM controller argument.

1. Add the NetScaler Helm chart repository to your local registry using the following command.

```
helm repo add netscaler https://netscaler.github.io/netscaler-helm
        -charts/
```

2. Install NetScaler Ingress Controller using the following command.

```
1 helm install netscaler-ingress-controller netscaler/netscaler-
ingress-controller --set nsIP=<NS_IP>,license.accept=yes,
adcCredentialSecret=<>,ingressClass[0]=netscaler,serviceClass
[0]=netscaler,ipam=true -n netscaler
```

For detailed information about deploying and configuring NetScaler Ingress Controller using Helm charts, see the Helm chart repository.

#### **Deploy IPAM controller**

1. Add the NetScaler Helm chart repository to your local registry using the following command.

```
1 helm repo add netscaler https://netscaler.github.io/netscaler-helm
        -charts/
```

2. Install NetScaler IPAM controller using the following command.

```
helm install netscaler-ipam-controller netscaler/netscaler-ipam-
controller --set vipRange='[{
  "Prod": ["10.1.2.0 - 10.1.2.255"] }
] ' -n netscaler
```

#### **Deploy a sample application**

Perform the following steps to deploy an apache application in your Kubernetes cluster.

## Note:

In this example, an **apache** application is used. You can deploy a sample application of your choice.

1. Deploy a sample application using the following command:

```
yml kubectl apply -f - <<EOF apiVersion: apps/v1 kind: Deployment
metadata: name: apache namespace: netscaler labels: name: apache
spec: selector: matchLabels: app: apache replicas: 2 template:
metadata: labels: app: apache spec: containers: - name: apache
image: httpd:latest ports: - name: http containerPort: 80
imagePullPolicy: IfNotPresent EOF
```

2. Verify if the pods are running using the following command:

```
1 kubectl get pods
```

#### Expose the sample application using service of type LoadBalancer

Perform the following steps to create a service of type LoadBalancer:

1. Deploy a service to expose apache application, for which the IP address is allocated from the Prod VIP range specified during the IPAM installation.

```
kubectl apply -f - <<EOF</pre>
1
2
     apiVersion: v1
3
     kind: Service
4
     metadata:
5
      name: apache
      namespace: netscaler
6
7
       labels:
8
         name: apache
9
      annotations:
         service.citrix.com/class: 'netscaler'
11
         service.citrix.com/ipam-range: 'Prod'
     spec:
12
13
       externalTrafficPolicy: Local
14
       type: LoadBalancer
15
       ports:
16
       - name: http
        port: 80
17
18
         targetPort: http
19
       selector:
20
         app: apache
21
22
     EOF
```

When you create the service, the IPAM controller assigns an IP address to the apache service from the IP address range you had defined in the IPAM controller deployment. The IP address allocated by the IPAM controller is provided in the status.loadBalancer.ingress: field of the service definition. NetScaler Ingress Controller configures the IP address allocated to the service as a virtual IP address (VIP) in NetScaler.

2. View the service using the following command:

```
1 kubectl get service apache --output yaml
```

Output:

```
apiVersion: v1
kind: Service
metadata:
 annotations:
  NETSCALER_VPORT: "80"
 creationTimestamp: "2019-10-04T18:40:42Z"
 labels:
   name: apache
 name: apache
 namespace: default
 resourceVersion: "22245641"
 selfLink: /api/v1/namespaces/default/services/apache
 uid: 7811c632-e6d6-11e9-8359-527c8bde541f
spec:
 clusterIP: 10.108.104.133
 externalTrafficPolicy: Local
 healthCheckNodePort: 32646
 ports:
 – name: http
   nodePort: 31408
   port: 80
   protocol: TCP
   targetPort: http
 selector:
   app: apache
 sessionAffinity: None
 type: LoadBalancer
status:
 loadBalancer:
   ingress:
   - ip: 10.105.158.196
```
#### Access the service

You can access the apache service using the IP address assigned by the IPAM controller to the service. You can find the IP address in the status.loadBalancer.ingress: field of the service definition. Use the following curl command to access the service:

```
curl <IP_address>
1
```

The response should be:

```
<html><body><h1>It works!</h1></body></html>
```

# **Securing Ingress**

December 31, 2023

The topic covers the various ways to secure your Ingress using NetScaler and the annotations provided by the NetScaler Ingress Controller.

The following table lists the TLS use cases with sample annotations that you can use to secure your Ingress using the Ingress NetScaler and the NetScaler Ingress Controller:

U	se cases	Sample annotations
tle	n <b>iáble%sSv1n3tprotocos</b> perauthcont	<pre>ingress.citrix.com/frontend-sslprofile: '{ "tls13 ext":"1", "dhekeyexchangewithpsk":"yes" } '</pre>
Н	TTP strict"לדאז איז איז איז איז איז איז איז איז איז א	<pre>ingress.citrix.com/frontend-sslprofile: '{ "hsts" subdomain":"ves" }</pre>
		<pre>ingress.citrix.com/frontend-sslprofile: '{ "ocsps</pre>
Ŭ	CSP stapling	<pre>ingress.citrix.com/frontend-sslprofile: '{ "clien</pre>
S	et client authenticationtto mandatoryat	tory" } '
Т	LS session ticket extension ketlifeti	<pre>ingress.citrix.com/frontend-sslprofile: '{ "sessi ime : "300" } '</pre>
ç	SI session remeastimeout • "120"	<pre>ingress.citrix.com/frontend-sslprofile: '{ "sessr } '</pre>
5		<pre>ingress.citrix.com/frontend-sslprofile:'{ "sniena</pre>
phers" : [{ "c ] } ' C	ciphername": "secure", "ciph iphergroups	erpriority" :"1" } , { "ciphername": "secure", "c <sup>-</sup>
ers" : [{ "cipl recturl" } ' C	hername": "secure", "cipherp ipherredirect	<pre>ingress.citrix.com/frontend-sslprofile:'{ "sniena priority" :"1" } ], "cipherredirect":"enabled", "can priority" :"1" } ]</pre>

# Enable TLS v1.3 protocol

Using the annotations for SSL profiles, you can enable TLS 1.3 protocol support on the SSL profile and set the tls13SessionTicketsPerAuthContext and dheKeyExchangeWithPsk parameters in the SSL profile for the Ingress NetScaler.

The tls13SessionTicketsPerAuthContext parameter enables you to set the number of tickets the Ingress NetScaler issues anytime TLS 1.3 is negotiated, ticket-based resumption is enabled, and either a handshake completes or post-handhsake client authentication completes. The value can be increased to enable clients to open multiple parallel connections using a fresh ticket for each connection. The minimum value you can set is 1 and the maximum is 10. By default, the value is set to 1.

Note:

No tickets are sent if resumption is disabled.

The dheKeyExchangeWithPsk parameter allows you to specify whether the Ingress NetScaler requires a DHE key exchange to occur when a preshared key is accepted during a TLS 1.3 session resumption handshake. A DHE key exchange ensures forward secrecy, even if ticket keys are compromised, at the expense of extra resources required to carry out the DHE key exchange.

The following is a sample annotation for the HTTP profile to enable TLS 1.3 protocol support on SSL profile and set the tls13SessionTicketsPerAuthContext and dheKeyExchangeWithPsk parameters in the SSL profile.

# HTTP strict transport security (HSTS)

The Ingress NetScaler appliances support HTTP strict transport security (HSTS) as an inbuilt option in SSL profiles. Using HSTS, a server can enforce the use of an HTTPS connection for all communication with a client. That is, the site can be accessed only by using HTTPS. Support for HSTS is required for A+ certification from SSL Labs. For more information, see NetScaler support for HSTS.

Using the annotations for SSL profiles, you can enable HSTS in an SSL front-end profile on the Ingress NetScaler. The following is a sample ingress annotation:

```
ingress.citrix.com/frontend-sslprofile: '{
    "hsts":"enabled", "maxage" : "157680000", "includesubdomain":"yes" }
    '
```

Where:

- HSTS The state of HTTP Strict Transport Security (HSTS) on the SSL profile. Using HSTS, a server can enforce the use of an HTTPS connection for all communication with a client. The supported values are ENABLED and DISABLED. By default, the value is set to DISABLED.
- maxage Allows you to set the maximum time, in seconds, in the strict transport security (STS) header during which the client must send only HTTPS requests to the server. The minimum time you can set is 0 and the maximum is 4294967294. By default the value is to 0.
- IncludeSubdomains Allows you to enable HSTS for subdomains. If set to Yes, a client must send only HTTPS requests for subdomains. By default the value is set to No.

# **OCSP** stapling

The Ingress NetScaler can send the revocation status of a server certificate to a client, at the time of the SSL handshake, after validating the certificate status from an OCSP responder. The revocation status of a server certificate is "stapled" to the response the appliance sends to the client as part of the SSL handshake. For more information on NetScaler implementation of CRL and OCSP reports, see OCSP stapling.

To use the OCSP stapling feature, you can enable it using an SSL profile with the following ingress annotation:

```
ingress.citrix.com/frontend-sslprofile: '{
    "ocspstapling":"enabled" }
    '
```

#### Note:

To use OCSP stapling, you must add an OCSP responder on the NetScaler appliance.

#### Set Client authentication to mandatory

Using the annotations for SSL profiles, you can enable client authentication, the Ingress NetScaler appliance asks for the client certificate during the SSL handshake.

The appliance checks the certificate presented by the client for normal constraints, such as the issuer signature and expiration date.

Here are some use cases:

- Require a valid client certificate before website content is displayed. This restricts website content to only authorized machines and users.
- Request a valid client certificate. If a valid client certificate is not provided, then prompt the user for multifactor authentication.

Client authentication can be set to mandatory, or optional.

- When it is set as mandatory, if the SSL Client does not transmit a valid Client Certificate, then the connection is dropped. Valid means: signed/issued by a specific Certificate Authority, and not expired or revoked.
- When it is optional, then the NetScaler requests the client certificate, but proceeds with the SSL transaction even if the client presents an invalid certificate or no certificate. This configuration is useful for authentication scenarios (for example require two-factor authentication if a valid Client Certificate is not provided)

Using the annotations for SSL profiles, you can enable client authentication on an SSL virtual server and set client authentication as Mandatory.

The following is a sample annotation of the SSL profile:

```
1 ingress.citrix.com/frontend-sslprofile: '{
2 "clientauth":"enabled", "clientcert" : "mandatory" }
3 '
```

Note:

Make sure that you bind the client-certificate to the SSL virtual server on the Ingress NetScaler.

# **TLS session ticket extension**

An SSL handshake is a CPU-intensive operation. If session reuse is enabled, the server or client key exchange operation is skipped for existing clients. They are allowed to resume their sessions. This improves the response time and increases the number of SSL transactions per second that a server can support. However, the server must store details of each session state, which consumes memory and is difficult to share among multiple servers if requests are load balanced across servers.

The Ingress NetScaler appliances support the SessionTicket TLS extension. Use of this extension indicates that the session details are stored on the client instead of on the server. The client must indicate that it supports this mechanism by including the session ticket TLS extension in the client Hello message. For new clients, this extension is empty. The server sends a new session ticket in the NewSessionTicket handshake message. The session ticket is encrypted by using a key-pair known only to the server. If a server cannot issue a new ticket currently, it completes a regular handshake.

Using the annotations for SSL profiles, you can enable the use of session tickets, as per the RFC 5077. Also, you can set the life time of the session tickets issued by the Ingress NetScaler, using the sessionticketlifetime parameter.

The following is the sample ingress annotation:

```
2 "sessionticket" : "enabled", "sessionticketlifetime : "300" }
```

3

## **SSL** session reuse

You can reuse an existing SSL session on a NetScaler appliance. While the SSL renegotiation process consists of a full SSL handshake, the SSL reuse consists of a partial handshake because the client sends the SSL ID with the request.

Using the annotations for SSL profiles, you can enable session reuse and also set the session timeout value (in seconds) on the Ingress NetScaler.

The following is the sample ingress annotation:

```
1 ingress.citrix.com/frontend-sslprofile: '{
2 "sessreuse" : "enabled", "sesstimeout : "120" }
3 '
```

By default, the session reuse option is enabled on the appliance and the timeout value for the same is set to 120 seconds. Therefore, if a client sends a request on another TCP connection and the earlier SSL session ID within 120 seconds, then the appliance performs a partial handshake.

## Using cipher groups

The Ingress NetScaler ships with built-in cipher groups. To use ciphers that are not part of the DEFAULT cipher group, you have to explicitly bind them to an SSL profile. You can also create a user-defined cipher group to bind to the SSL virtual server on the Ingress NetScaler.

The built-in cipher groups can be used in Tier-1 and Tier-2 NetScaler, and the user-defined cipher group can be used only in Tier-1 NetScaler.

To use a user-defined cipher group, ensure that the NetScaler has a user-defined cipher group. Perform the following:

- 1. Create a user-defined cipher group. For example, testgroup.
- 2. Bind all the required ciphers to the user-defined cipher group.
- 3. Note down the user-defined cipher group name.

For detailed instructions, see Configure a user-defined cipher group.

Using the annotations for SSL profiles, you can bind the built-in cipher groups, a user-defined cipher group or both to the SSL profile.

The following is the syntax of the ingress annotation that you can use to bind the built-in cipher groups and a user-defined cipher group to an SSL profile:

```
1 ingress.citrix.com/frontend-sslprofile:'{
2 "snienable":"enabled", "ciphers" : [{
3 "ciphername": "secure", "cipherpriority" :"1" }
4 , {
5 "ciphername": "testgroup", "cipherpriority" :"2" }
6 ] }
7 '
```

The ingress annotation binds the built-in cipher group, SECURE, and the user-defined cipher group, testgroup, to the SSL profile.

# **Using cipher redirect**

During the SSL handshake, the SSL client (usually a web browser) announces the suite of ciphers that it supports, in the configured order of cipher preference. From that list, the SSL server then selects a cipher that matches its own list of configured ciphers.

If the ciphers announced by the client does not match those ciphers configured on the SSL server, the SSL handshake fails. The failure is announced by a cryptic error message displayed in the browser. These messages rarely mention the exact cause of the error.

With cipher redirection, you can configure an SSL virtual server to deliver accurate, meaningful error messages when an SSL handshake fails. When the SSL handshake fails, the NetScaler appliance redirects the user to a previously configured URL or, if no URL is configured, displays an internally generated error page.

The following is the syntax of the ingress annotation that you can use to bind cipher groups and enable cipher redirect to redirect the request to redirecturl.

```
ingress.citrix.com/frontend-sslprofile:'{
    "snienable": "enabled", "ciphers" : [{
    "ciphername": "secure", "cipherpriority" :"1" }
    ], "cipherredirect":"enabled", "cipherurl": "https://redirecturl" }
    '
```

# **TCP use cases**

July 22, 2024

In a Kubernetes environment, an ingress object allows access to the Kubernetes services from outside the Kubernetes cluster. Standard Kubernetes ingress resources assume that all the traffic is HTTPbased and do not cater to non-HTTP based protocols such as TCP, UDP, and SSL. Hence, any non-HTTP applications such as DNS, FTP or LDAP cannot be exposed using the standard Kubernetes ingress.

# How to load balance TCP ingress traffic

NetScaler provides a solution using ingress annotations to load balance TCP-based ingress traffic. When you specify these annotations in the ingress resource definition, NetScaler Ingress Controller configures NetScaler to load balance TCP ingress traffic.

You can use the following annotations in your Kubernetes ingress resource definition to load balance the TCP-based ingress traffic:

- ingress.citrix.com/insecure-service-type: This annotation enables L4 load balancing with TCP for NetScaler.
- ingress.citrix.com/insecure-port: This annotation configures the port for TCP traffic. It is helpful when micro service access is required on a non-standard port. By default, port 80 is configured.

For more information about ingress annotations, see Ingress annotations.

Sample: Ingress definition for TCP-based ingress.

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: networking.k8s.io/v1
3 kind: Ingress
4 metadata:
5
   annotations:
6
       ingress.citrix.com/insecure-port: '6379'
       ingress.citrix.com/insecure-service-type: tcp
7
   name: redis-master-ingress
8
9 spec:
10
    ingressClassName: guestbook
   defaultBackend:
11
     service:
13
        name: redis-master-pods
14
         port:
15
           number: 6379
16 ---
17
   apiVersion: networking.k8s.io/v1
18 kind: IngressClass
19 metadata:
20 name: guestbook
21 spec:
22
    controller: citrix.com/ingress-controller
23
    EOF
```

You can use the following service annotation in the service definition YAML to load balance the TCPbased ingress traffic: service.citrix.com/service-type-<index>. For more information about service annotations, see Service annotations.

Sample: Service type LoadBalancer YAML for load balancing TCP-based ingress traffic.

```
1 apiVersion: v1
```

```
2 kind: Service
3 metadata:
4 name: backend
5 annotations:
        service.citrix.com/class: 'netscaler'
6
7
        service.citrix.com/ipam-range: 'Dev'
        service.citrix.com/service-type-0: TCP
8
9
     labels:
10
         app: backend
11 spec:
12 ports:
13
     - name: port-6379
14
         port: 6379
15
        targetPort: 6379
   type: LoadBalancer
16
   selector:
17
18
     name: backend
```

# Load balance ingress traffic based on SSL over TCP

NetScaler Ingress Controller provides ingress.citrix.com/secure-service-type: ssl\_tcp annotation that you can use to load balance ingress traffic based on SSL over TCP.

Sample: Ingress definition for SSL over TCP based Ingress.

```
kubectl apply -f - <<EOF</pre>
1
2
     apiVersion: networking.k8s.io/v1
3
   kind: Ingress
4
    metadata:
       annotations:
5
         ingress.citrix.com/secure-service-type: "ssl_tcp"
6
         ingress.citrix.com/secure-backend: '{
7
    "frontendcolddrinks":"True" }
8
   1
9
       name: colddrinks-ingress
10
11
    spec:
       ingressClassName: colddrink
13
       defaultBackend:
14
         service:
15
           name: frontendcolddrinks
16
           port:
17
             number: 443
18
       tls:
       - secretName: "colddrink-secret"
19
20
21
     apiVersion: networking.k8s.io/v1
22
     kind: IngressClass
   metadata:
23
24
       name: colddrink
25
     spec:
26
    controller: citrix.com/ingress-controller
```

27 EOF

# Monitor and improve the performance of your TCP-based applications

Application developers can closely monitor the health of UDP-based applications through rich monitors (such as TCP-ECV) in NetScaler. The ECV (extended content validation) monitors help in checking whether the application returns expected content or not. NetScaler Ingress Controller provides ingress.citrix.com/monitor annotation that can be used to monitor the health of the backend service.

Also, the application performance can be improved by using persistence methods such as Source IP. You can use these NetScaler features through Smart Annotations in Kubernetes.

The following ingress resource example uses smart annotations:

```
1 kubectl apply -f - <<EOF</pre>
   apiVersion: networking.k8s.io/v1
2
3 kind: Ingress
4 metadata:
5
   annotations:
6
       ingress.citrix.com/frontend-ip: "192.168.1.1"
       ingress.citrix.com/insecure-port: "80"
7
       ingress.citrix.com/lbvserver: '{
8
    "mongodb-svc":{
9
    "lbmethod":"SRCIPDESTIPHASH" }
10
11
    }
12
    1
13
       ingress.citrix.com/monitor: '{
    "mongodbsvc":{
14
    "type":"TCP-ECV" }
15
    }
16
    1
17
18
    name: mongodb
19
   spec:
20
     rules:
21
     - host: mongodb.beverages.com
22
       http:
         paths:
23
24
         - backend:
25
              service:
26
                name: mongodb-svc
27
                port:
28
                  number: 80
29
            path: /
            pathType: Prefix
31
   EOF
```

# How to expose non-standard HTTP ports in the NetScaler CPX service

Sometimes you need to expose ports other than 80 and 443 in a NetScaler CPX service for allowing TCP traffic on other ports.

This section provides information on how to expose other non-standard HTTP ports on the NetScaler CPX service when you deploy it in the Kubernetes cluster.

## For Helm chart deployments

To expose non-standard HTTP ports while deploying NetScaler CPX with ingress controller using Helm charts, see the Helm chart installation guide.

#### For deployments using the OpenShift operator

For deployments using the OpenShift operator, you need to edit the YAML definition to create CPX with ingress controller as specified in the step 6 of Deploy the NetScaler Ingress Controller as a sidecar with NetScaler CPX using NetScaler Operator and specify the ports as shown in the following example:

1 servicePorts: - port: 80 3 protocol: TCP 4 name: http 5 - port: 443 6 protocol: TCP 7 name: https 8 - port: 6379 9 protocol: TCP 10 name: tcp

The following sample configuration is an example for deployment using the OpenShift Operator. The service port definitions are highlighted in green.

#### NetScaler ingress controller

Hancopeccust
nitroReadTimeout: 20
nodeSelector:
key: **
value: "
nsConfigDnsRec: false
nsCookieVersion: '0'
nsDnsNameserver: ''
nsGateway: 192.168.1.1
nsHTTP2ServerSide: 'OFF'
nsIP: 192.168.1.2
nsLbHashAlgo:
hashAlgorithm: DEFAULT
hashFingers: 256
required: false
nsProtocol: http
nsSvcLbDnsRec: false
openshift: true
optimizeEndpointBinding: false
podAnnotations: {}
profileHttpFrontend: {}
profileSslFrontend: {}
profileTcpFrontend: {}
pullPolicy: IfNotPresent
rbacRole: false
replicaCount: 1
resources: {}
routeLabels: ''
serviceAccount:
create: true
serviceAnnotations: {}
servicePorts:
- port: 80
protocol: TCP
name: http
- port: 443
protocol: TCP
name: https
- port: 6379
protocol: TCP
name: tcp
servicespec:
externalinatticPolicy: Cluster
Toducalancersourcekanges: []
Servicetype.
loadealancer:

# **TCP profile support**

This section covers various ways to configure TCP parameters on NetScaler using smart annotations for services of type LoadBalancer and ingress using the annotations in NetScaler Ingress Controller.

A TCP profile is a collection of TCP settings. Instead of configuring the settings on each entity, you can configure TCP settings in a profile and bind the profile to all the required entities. The front-end TCP profiles can be attached to the client-side content switching virtual server and the back-end TCP profiles can be configured for a service group.

# TCP profile support for services of type LoadBalancer

NetScaler Ingress Controller provides the following service annotations for TCP profile for services of type LoadBalancer. You can use these annotations to define the TCP settings for NetScaler.

Service annotation	Description
service.citrix.com/frontend- tcpprofile	Use this annotation to create the front-end TCP profile (client plane).
service.citrix.com/backend- tcpprofile	Use this annotation to create the back-end TCP profile (server plane).

**User-defined TCP profiles** Using service annotations for TCP, you can create custom profiles with the same name as that of content switching virtual server or service group and bind them to the corresponding virtual server (frontend-tcpprofile) and service group (backend-tcpprofile).

Service annotation	Sample	
ervice.citrix.com/front'emablepp'rofi'le	<pre>service.citrix.com/frontend-tcpprofile</pre>	
service.citrix.com/backendatlepp/rofile	service.citrix.com/backend-tcpprofi	

**Built-in TCP profiles** Built-in TCP profiles do not create any profiles but bind a given profile name in the annotation to the corresponding virtual server (frontend-tcpprofile) and service group (backend-tcpprofile).

Examples for built-in TCP profiles:

```
service.citrix.com/frontend-tcpprofile: "tcp_preconf_profile"
service.citrix.com/backend-tcpprofile: '{ "citrix-svc" : "tcp_preconf_profile
" }
```

**Example: Service of type LoadBalancer with TCP profile configuration** In this example, TCP profiles are configured for a sample application tea-beverage. This application is deployed and exposed using a service of type LoadBalancer using the YAML file.

Note:

For information about exposing services of type LoadBalancer, see service of type LoadBalancer.

Deploy a sample application (tea-beverage.yaml) using the following command:

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: apps/v1
3 kind: Deployment
4 metadata:
5
   name: tea-beverage
6 labels:
7
      name: tea-beverage
8 spec:
9
   selector:
    matchLabels:
10
11
       name: tea-beverage
  replicas: 2
12
13 template:
14
     metadata:
15
        labels:
16
          name: tea-beverage
17
     spec:
18
        containers:
19
         - name: tea-beverage
20
           image: quay.io/sample-apps/hotdrinks:latest
           ports:
21
22
           - name: tea-80
23
            containerPort: 80
24
           - name: tea-443
25
            containerPort: 443
26
           imagePullPolicy: Always
27
   ___
28 apiVersion: v1
29 kind: Service
30 metadata:
31 name: tea-beverage
32 annotations:
33
     service.citrix.com/frontend-ip: 10.105.158.194
34
       service.citrix.com/frontend-tcpprofile: '{
   "ws" : "enabled", "sack" : "enabled" }
     service.citrix.com/backend-tcpprofile: '{
37
   "ws" : "enabled", "sack" : "enabled" }
38
   1
39
40 spec:
   type: LoadBalancer
41
   ports:
42
43
    - name: tea-80
44
     port: 80
45
      targetPort: 80
46 selector:
47
      name: tea-beverage
48 EOF
```

After the application is deployed, the corresponding entities and profiles are created on NetScaler. Run the following commands on NetScaler to verify the same: show cs vserver k8s-teabeverage\_80\_default\_svcandshow servicegroup k8s-tea-beverage\_80\_sgp\_f4lezsannv •

1	<pre># show cs vserver k8s-tea-beverage_80_default_svc</pre>
2	k8s-tea-beverage_80_default_svc (10.105.158.194:80) - TCP
	Type: CONTENT
3	State: UP
4	Last state change was at Wed Apr 3 09:37:59 2024
5	Time since last state change: 0 days, 00:00:09.790
6	Client Idle Timeout: 9000 sec
7	Down state flush: ENABLED
8	Disable Primary Vserver On Down : DISABLED
9	Comment: uid=
	VIGQWRCYKCM6WFYX2GFKRVT3ZF6JSFISPW6XM24JADBXEYRLITOQ====
10	**TCP profile name: k8s-tea-beverage_80_default_svc**
11	Appflow logging: ENABLED
12	State Update: DISABLED
13	Default: k8s-tea-
	beverage_80_lbv_f4lezsannvu7tk2ftpjbhi4hza2tvdnk Content
	Precedence: RULE
14	L2Conn: OFF Case Sensitivity: ON
15	Authentication: OFF
16	401 Based Authentication: OFF
17	HTTP Redirect Port: 0 Dtls : OFF
18	Persistence: NONE
19	Listen Policy: NONE
20	IcmpResponse: PASSIVE
21	RHIstate: PASSIVE
22	Traffic Domain: 0
23	
24	1) Default Target LB: k8s-tea-
	beverage_80_lbv_f4lezsannvu7tk2ftpjbhi4hza2tvdnk Hits: 0
25	Done

1	# show servicegroup k8s-tea-
	beverage_80_sgp_f4lezsannvu7tk2ftpjbhi4hza2tvdnk
2	k8s-tea-beverage_80_sgp_f4lezsannvu7tk2ftpjbhi4hza2tvdnk -
	TCP
3	State: ENABLED Effective State: UP Monitor Threshold : 0
4	Max Conn: O Max Req: O Max Bandwidth: O kbits
5	Use Source IP: NO
6	Client Keepalive(CKA): NO
7	Monitoring Owner: 0
8	TCP Buffering(TCPB): NO
9	HTTP Compression(CMP): NO
10	Idle timeout: Client: 9000 sec Server: 9000 sec
11	Client IP: DISABLED
12	Cacheable: NO
13	SC: ???
14	SP: OFF
15	Down state flush: ENABLED
16	Monitor Connection Close : NONE
17	Appflow logging: ENABLED
18	TCP profile name: k8s-tea-

	beverage_80_sgp_f4lezsannvu7tk2ftpjbhi4hza2tvdnk
19	ContentInspection profile name: ???
20	Process Local: DISABLED
21	Traffic Domain: 0
22	Comment: "lbsvc:tea-beverage,svcport:80,ns:default"
23	
24	
25	1) 10.146.107.38:30524 State: UP Server Name: 10.146.107.38
	Server ID: None Weight: 1 Order: Default
26	Last state change was at Wed Apr 3 09:38:00 2024
27	Time since last state change: 0 days, 00:02:27.660
28	
29	Monitor Name: tcp- <b>default</b> State: UP Passive: 0
30	Probes: 30 Failed [Total: 0 Current: 0]
31	Last response: Success - TCP syn+ack received.
32	Response Time: 0.000 millisec
33	Done

# Note:

The TCP profile is supported for single-port services.

# Configure TCP profiles using Ingress annotations

The following table lists some of the TCP use cases with sample annotations:

Use case	Sample annotation	
Silently drop idle TCP connections	<pre>ingress.citrix.com/frontend- tcpprofile: '{ "</pre>	
Delayed TCP connection acknowledgments	<pre>drophalfclosedconnontimeout" : " ingress.citrix.com/frontend- enabled", "dropestconnontimeout" tcpprofile: '{,"delayedack" :</pre>	
Client side Multipath TCP session management	ingress:eitrix:eom/proktedda- ingress:eitrix:com/backenda- tepprofile: 'i "matchatsvendbackend" tepprofile: 'i "matchatsvendbackend" drophatceteesedentation:	
TCP Optimization	enabled", "dropestconnontimeout"	
Selective acknowledgment	ingressieftrixieom/backend- froptondlacpprofiterixisveack"::"	
Forward acknowledgment	<pre>Mptabled"emabled", " ingress.citrix.com/ mptablessiont*meemt": "7200" } } frontend_tcpprofile: '{{"tack": backend_tcpprofile: '{{"citrix- "enabled" } '</pre>	
Window Scaling	<pre>shere's is it is a ck "compensate of the second secon</pre>	
	<pre>backend_tcpprofile: '{ "citrix-</pre>	
© 1997–2025 Citrix Systems, Inc. All rights reserved. <sub>SVC</sub> ": { "ws" : "enabled", "wsva@		
	": "9" } } '	

Use case	Sample annotation
Maximum Segment Size	ingress.citrix.com/
	<pre>trontend_tcpprofile: '{ "mss" :</pre>
Keep-Alive	"1460", "maxpktpermss" : "512" } įngress.citrix.com/
	frontend_tcpprofile: '{ "ka" : " ingress.citrix.com/
bufferSize	enabled", " Pagkesg_tcpprofile; '{ "citrix- kaprobeupdatelastactivity" : " \$Yonteng t@pprofile460{,""
МРТСР	enabled", "kaconnidletime"; MaxekFeefees: %81962"; "Ageess"&81962"; "Ageess"&81962"; "Ageess"&9063"; "Ageoss", citrix, com/ "Ageoseednteebyacfile; "{ "mptcp"; backend_tcpprofile; "{ "citrix-
flavor	Inghabteditrix.com/ Bygreidelentersize": "8190" Herchdebdebbohpreestsfllcitrix- trentend_trepprofilgabltddlflqvor"
Dynamic receive buffering	kappesserver verse
Defending TCP against spoofing attacks	「日日の高いは、法律権が日本日日第二、「「3」、「"×」 日日の高いは、法律権が日本日日第二、「3」、「"×」 日日の日本日本日本日本日本日本日本日本日本日本日本日本日本日本日本日本日本日本
	<pre>MgressicitPfspcom/ enabled", "     satwindowattenuate", ''     satwindowattenuate", '' </pre>
	มลอดจรลงกฎแพลนู เม่นอองจากเรา เน่น เน่น เ

#### Note:

The above Ingress annotations can also be used with service annotations in the format described earlier.

#### "spootsynarop" : "enabled" } } '

**Silently drop idle TCP connections** In a network, when a large number of TCP connections become idle, NetScaler sends RST packets to close them. The packets sent over the channels activate those channels unnecessarily, causing a flood of messages that in turn causes NetScaler to generate a flood of service-reject messages.

Using the drophalfclosedconnontimeout and dropestconnontimeout parameters in TCP profiles, you can silently drop TCP half closed connections on idle timeout or drop TCP established connections on an idle timeout. By default, these parameters are disabled on NetScaler. If you enable both of them, neither a half closed connection nor an established connection causes an RST packet to be sent to the client when the connection times out. The NetScaler just drops the connection.

Using the annotations for TCP profiles, you can enable or disable the drophalfclosedconnontimeout and dropestconnontimeout on NetScaler. The following is a sample annotation of TCP profile to enable these parameters:

```
ingress.citrix.com/frontend-tcpprofile: '{ "drophalfclosedconnontimeout
" : "enable", "dropestconnontimeout" : "enable" } '
```

```
ingress.citrix.com/backend-tcpprofile: '{ "citrix-svc" : { "drophalfclosedconn
" : "enable", "dropestconnontimeout" : "enable" } '
```

**Delayed TCP connection acknowledgments** To avoid sending several ACK packets, NetScaler supports TCP delayed acknowledgment mechanism. It sends delayed ACK with a default timeout of 100 ms. NetScaler accumulates data packets and sends ACK only if it receives two data packets in continuation or if the timer expires. The minimum delay you can set for the TCP deployed ACK is 10 ms and the maximum is 300 ms. By default the delay is set to 100 ms.

Using the annotations for TCP profiles, you can manage the delayed ACK parameter. The following is a sample annotation of TCP profile to enable these parameters:

```
ingress.citrix.com/frontend-tcpprofile: '{ "delayedack" : "150" } '
ingress.citrix.com/backend-tcpprofile: '{ "citrix-svc" : { "delayedack
" : "150" } } '
```

**Client side Multipath TCP session management** You can perform TCP configuration on NetScaler for Multipath TCP (MPTCP) connections between the client and NetScaler. MPTCP connections are not supported between NetScaler and the back-end communication. Both the client and NetScaler appliance must support the same MPTCP version.

You can enable MPTCP and set the MPTCP session timeout (mptcpsessiontimeout) in seconds using TCP profiles in NetScaler. If the mptcpsessiontimeout value is not set then the MPTCP sessions are flushed after the client idle timeout. The minimum timeout value you can set is 0 and the maximum is 86400. By default, the timeout value is set to 0.

Using the annotations for TCP profiles, you can enable MPTCP and set the mptcpsessiontimeout parameter value on NetScaler. The following is a sample annotation of TCP profile to enable MPTCP and set the mptcpsessiontimeout parameter value to 7200 on NetScaler:

```
ingress.citrix.com/frontend-tcpprofile: '{ "mptcp" : "enabled", "
mptcpsessiontimeout" : "7200" } '
ingress.citrix.com/backend-tcpprofile: '{ "citrix-svc" : { "mptcp" :
"enabled", "mptcpsessiontimeout" : "7200" } } '
```

**TCP optimization** Most of the relevant TCP optimization capabilities of NetScaler are exposed through a corresponding TCP profile. Using the annotations for TCP profiles, you can enable the following TCP optimization capabilities on NetScaler:

• Selective acknowledgment (SACK): TCP SACK addresses the problem of multiple packet losses which reduces the overall throughput capacity. With selective acknowledgment the receiver can inform the sender about all the segments which are received successfully, enabling sender to only retransmit the segments which were lost. This technique helps T1 improve overall throughput and reduce the connection latency.

The following is a sample annotation of TCP profile to enable SACK on NetScaler:

```
ingress.citrix.com/frontend_tcpprofile: '{ "sack" : "enabled" } '
ingress.citrix.com/backend_tcpprofile: '{ "citrix-svc" : { "sack"
    : "enabled" } } '
```

• Forward acknowledgment (FACK): To avoid TCP congestion by explicitly measuring the total number of data bytes outstanding in the network, and helping the sender (either T1 or a client) control the amount of data injected into the network during retransmission timeouts.

The following is a sample annotation of TCP profile to enable FACK on NetScaler:

```
ingress.citrix.com/frontend_tcpprofile: '{ "fack" : "enabled" } '
ingress.citrix.com/backend_tcpprofile: '{ "citrix-svc" : { "fack"
    : "enabled" } } '
```

• Window Scaling (WS): TCP window scaling allows increasing the TCP receive window size beyond 65535 bytes. It helps improve TCP performance overall and specially in high bandwidth and long delay networks. It helps with reducing latency and improving response time over TCP.

The following is a sample annotation of TCP profile to enable WS on NetScaler:

```
ingress.citrix.com/frontend_tcpprofile: '{ "ws" : "enabled", "
wsval" : "9" } '
```

```
ingress.citrix.com/backend_tcpprofile: '{ "citrix-svc" : { "ws" :
    "enabled", "wsval" : "9" } } '
```

Where wsval is the factor used to calculate the new window size. The argument is mandatory only when window scaling is enabled. The minimum value you can set is 0 and the maximum is 14. By default, the value is set to 4.

• Maximum Segment Size (MSS): MSS of a single TCP segment. This value depends on the MTU setting on intermediate routers and end clients. A value of 1460 corresponds to an MTU of 1500.

The following is a sample annotation of TCP profile to enable MSS on NetScaler:

```
ingress.citrix.com/frontend_tcpprofile: '{ "mss" : "1460", "
maxpktpermss" : "512" } '
```

```
ingress.citrix.com/backend_tcpprofile: '{ "citrix-svc" : { "mss"
: "1460", "maxpktpermss" : "512" } }
```

Where:

- mss is the MSS to use for the TCP connection. Minimum value: 0; Maximum value: 9176.
- maxpktpermss is the maximum number of TCP packets allowed per maximum segment size (MSS). Minimum value: 0; Maximum value: 1460.
- Keep-Alive (KA): Send periodic TCP keep-alive (KA) probes to check if the peer is still up.

The following is a sample annotation of TCP profile to enable TCP keep-alive (KA) on NetScaler:

```
ingress.citrix.com/frontend_tcpprofile: '{ "ka" : "enabled", "
kaprobeupdatelastactivity" : "enabled", "kaconnidletime": "900",
"kamaxprobes" : "3", "kaprobeinterval" : "75" } '
ingress.citrix.com/backend_tcpprofile: '{ "citrix-svc" : { "
ka" : "enabled", "kaprobeupdatelastactivity" : "enabled", "
kaconnidletime": "900", "kamaxprobes" : "3", "kaprobeinterval"
    : "75" } '
```

Where:

- ka is used to enable sending periodic TCP keep-alive (KA) probes to check if the peer is still up. Possible values: ENABLED, DISABLED. Default value: DISABLED.
- kaprobeupdatelastactivity updates the last activity for the connection after receiving keep-alive (KA) probes. Possible values: ENABLED, DISABLED. Default value: ENABLED.
- kaconnidletime is the duration (in seconds) for the connection to be idle, before sending a keep-alive (KA) probe. The minimum value you can set is 1 and the maximum is 4095.
- kaprobeinterval is the time interval (in seconds) before the next keep-alive (KA) probe, if the peer does not respond. The minimum value you can set is 1 and the maximum is 4095.
- bufferSize: Specify the TCP buffer size, in bytes. The minimum value you can set is 8190 and the maximum is 20971520. By default the value is set to 8190.

The following is a sample annotation of TCP profile to specify the TCP buffer size:

```
ingress.citrix.com/frontend_tcpprofile: '{ "bufferSize" : "8190"
} '
ingress.citrix.com/backend_tcpprofile: '{ "citrix-svc" : { "
bufferSize" : "8190" } } '
```

• Multipath TCP (MPTCP): Enable MPTCP and set the optional MPTCP configuration. The following is a sample annotation of TCP profile to enable MPTCP and se the optional MPTCP configurations:

```
ingress.citrix.com/frontend_tcpprofile: '{ "mptcp" : "enabled", "
mptcpdropdataonpreestsf" : "enabled", "mptcpfastopen": "enabled",
    "mptcpsessiontimeout" : "7200" } '
```

```
ingress.citrix.com/backend_tcpprofile: '{ "citrix-svc" : { "
mptcp" : "enabled", "mptcpdropdataonpreestsf" : "enabled", "
mptcpfastopen": "enabled", "mptcpsessiontimeout" : "7200" } } '
Wherea
```

Where:

- mptcpdropdataonpreests f is used to silently dropping the data on Pre-Established subflow. When enabled, DSS data packets are dropped silently instead of dropping the connection when data is received on pre-established subflow. Possible values: ENABLED, DISABLED. Default value: DISABLED.
- mptcpfastopen can be enabled so that DSS data packets are accepted before receiving the third ack of SYN handshake. Possible values: ENABLED, DISABLED. Default value: DISABLED
- flavor: Set the TCP congestion control algorithm. Possible values: Default, BIC, CUBIC, Westwood, and Nile. Default value: Default. The following sample annotation of TCP profile sets the TCP congestion control algorithm:

```
ingress.citrix.com/frontend_tcpprofile: '{ "flavor" : "westwood"
} '
```

```
ingress.citrix.com/backend_tcpprofile: '{ "citrix-svc" : { "
flavor" : "westwood" } } '
```

• Dynamic receive buffering: Enable or disable dynamic receive buffering. When enabled, it allows the receive buffer to be adjusted dynamically based on memory and network conditions. Possible values: ENABLED, DISABLED, and the Default value: DISABLED.

Note:

The buffer size argument must be set for dynamic adjustments to take place.

```
ingress.citrix.com/frontend_tcpprofile: '{ "dynamicReceiveBuffering
" : "enabled" } '
```

```
ingress.citrix.com/backend_tcpprofile: '{ "citrix-svc" : { "
dynamicReceiveBuffering" : "enabled" } } '
```

**Defend TCP against spoofing attacks** You can enable NetScaler to defend TCP against spoof attacks using the rstWindowAttenuate parameter in TCP profiles. By default, the rstWindowAttenuate parameter is disabled. This parameter is enabled to protect NetScaler against spoofing. If you enable rstWindowAttenuate, it replies with corrective acknowledg-ment (ACK) for an invalid sequence number. Possible values: Enabled, Disabled. Additionally, spoofSynDrop parameter can be used to enable or disable drop of invalid SYN packets to protect against spoofing. When disabled, established connections will be reset when a SYN packet is received. The default value for this parameter is ENABLED.

The following is a sample annotation of TCP profile to enable rstWindowAttenuate on NetScaler:

```
ingress.citrix.com/frontend_tcpprofile: '{ "rstwindowattenuate" : "
enabled", "spoofsyndrop" : "enabled" } '
ingress.citrix.com/backend_tcpprofile: '{ "citrix-svc" : { "rstwindowattenuate"
" : "enabled", "spoofsyndrop" : "enabled" } '
```

**Example for applying TCP profile using Ingress annotation** This example shows how to apply TCP profiles.

1. Deploy the front-end ingress resource with the TCP profile. In this Ingress resource, backend and TLS are not defined.

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: networking.k8s.io/v1
3 kind: Ingress
4 metadata:
5 name: frontend-ingress
   annotations:
6
7
      ingress.citrix.com/insecure-termination: "allow"
      ingress.citrix.com/frontend-ip: "10.221.36.190"
8
9
      ingress.citrix.com/frontend-tcpprofile: '{
10 "ws" : "enabled", "sack" : "enabled" }
11 '
12 spec:
13
    tls:
14
     - hosts:
15
    rules:
    - host:
16
17 EOF
```

2. Deploy the secondary ingress resource with the same front-end IP address. Back end and TLS are defined, which creates the load balancing resource definition.

```
1 kubectl apply -f - <<EOF
2 apiVersion: networking.k8s.io/v1
3 kind: Ingress</pre>
```

```
4 metadata:
5 name: backend-ingress
6 annotations:
      ingress.citrix.com/insecure-termination: "allow"
7
8 ingress.citrix.com/frontend-ip: "10.221.36.190"
9 spec:
10 tls:
    - secretName: <hotdrink-secret>
11
12
   rules:
13 - host: hotdrink.beverages.com
14
     http:
        paths:
15
16
        - path: /
17
          pathType: Prefix
18
          backend:
19
           service:
             name: frontend-hotdrinks
20
             port:
21
22
              number: 80
23 EOF
```

3. After the ingress resources are deployed, the corresponding entities, profiles are created on NetScaler.

Run the following command on NetScaler: show cs vserver k8s-10.221.36.190 \_443\_ssl.

1	#	show cs vserver k8s-10.221.36.190_443_ssl
2		
3		k8s-10.221.36.190_443_ssl (10.221.36.190:443) - SSL Type: CONTENT
4		State: UP
5		Last state change was at Wed Apr 3 04:21:38 2024
6		Time since last state change: 0 days, 00:00:57.420
7		Client Idle Timeout: 180 sec
8		Down state flush: ENABLED
9		Disable Primary Vserver On Down : DISABLED
10		Comment: uid=
		XMX2KPYG2GUJIHGTLVCPA7QVXDUBDRMJFTAWNCPAA2TVXB33EL5A====
11		TCP profile name: k8s-10.221.36.190_443_ssl
12		Appflow logging: ENABLED
13		State Update: DISABLED
14		Default: Content Precedence: RULE
15		Vserver IP and Port insertion: OFF
16		L2Conn: OFF Case Sensitivity: ON
17		Authentication: OFF
18		401 Based Authentication: OFF
19		Push: DISABLED Push VServer:
20		Push Label Rule: none
21		HTTP Redirect Port: 0 Dtls : OFF
22		Persistence: NONE
23		Listen Policy: NONE
24		IcmpResponse: PASSIVE

```
25 RHIstate: PASSIVE
26 Traffic Domain: 0
27
28 1) Content-Switching Policy: k8s-
backend_80_csp_2k75kfjrr6ptgzwtncozwxdjqrpbvicz Rule: HTTP
.REQ.HOSTNAME.SERVER.EQ("hotdrink.beverages.com") && HTTP.
REQ.URL.PATH.SET_TEXT_MODE(IGNORECASE).STARTSWITH("/")
Priority: 200000008 Hits: 0
29 Done
```

#### Note:

For an exhaustive list of the various TCP parameters supported on NetScaler, refer to Supported TCP Parameters. The key and value that you pass in the JSON format must match the NetScaler NITRO format. For more information on the NetScaler NITRO API, see NetScaler 14.1 REST APIs - NITRO Documentation for TCP profiles.

# **HTTP use cases**

#### July 22, 2024

This topic covers various HTTP use cases that you can configure on NetScaler using smart annotations for services of type LoadBalancer and ingress.

HTTP configurations for NetScaler can be specified in an entity called HTTP profile, which is a collection of HTTP settings. The HTTP profile can then be associated with services or virtual servers that can use these HTTP configurations.

# HTTP profile support for services of type LoadBalancer

NetScaler Ingress Controller provides the following service annotations for HTTP profile for services of type LoadBalancer. You can use these annotations to define the HTTP settings for NetScaler.

Service annotation	Description
service.citrix.com/frontend- httpprofile	Use this annotation to create the front-end HTTP profile (client plane).
service.citrix.com/backend- httprofile	Use this annotation to create the back-end HTTP profile (server plane).

## **User-defined HTTP profiles**

Using service annotations for HTTP, you can create custom profiles with the same name as that of content switching virtual server or service group and bind these profiles to the corresponding content switching virtual server (frontend-httpprofile) or service group (backend-httpprofile).

Service annotation	Sample
service citrix com/fr'ont'end-shutt'orrd'fride	<pre>service.citrix.com/frontend-httpprofi ed" } '</pre>
service.ertrix.com/right@lasvecpprolenabit	service.citrix.com/backend-httpprofil
service.citrix.com/backæinds/hdttpprbaniabeled	}

#### **Built-in HTTP profiles**

Built-in HTTP profiles do not create any profiles but bind a given profile name in the annotation to the corresponding virtual server (frontend-httpprofile) and the service group (backend-httpprofile).

The following examples are for built-in HTTP profiles.

```
service.citrix.com/frontend-httpprofile: "http_preconf_profile"
service.citrix.com/backend-httpprofile: '{ "apache" : "http_preconf_profile
" } '
```

#### Example: Service of type LoadBalancer with HTTP profile configuration

In this example, HTTP profiles are configured for a sample application tea-beverage. This application is deployed and exposed using a service of type LoadBalancer using the following YAML file.

Note:

```
For information about exposing services of type LoadBalancer, see <a href="mailto:seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeservice.seeses
```

Deploy a sample application (tea-beverage.yaml) using the following command:

```
1 kubectl apply -f - <<EOF
2 apiVersion: apps/v1
3 kind: Deployment
4 metadata:
5 name: tea-beverage
6 labels:
7 name: tea-beverage</pre>
```

```
8 spec:
9
   selector:
10
    matchLabels:
11
      name: tea-beverage
   replicas: 2
12
   template:
13
14
       metadata:
15
         labels:
16
           name: tea-beverage
17
      spec:
18
        containers:
19
         - name: tea-beverage
20
           image: quay.io/sample-apps/hotdrinks:latest
           ports:
22
           - name: tea-80
23
             containerPort: 80
24
           - name: tea-443
25
             containerPort: 443
26
           imagePullPolicy: Always
27
28 apiVersion: v1
29 kind: Service
30 metadata:
31
   name: tea-beverage
   annotations:
       service.citrix.com/service-type-0: 'HTTP'
34
       service.citrix.com/frontend-ip: 10.105.158.194
35
       service.citrix.com/frontend-httpprofile: '{
   "http2direct" : "enabled", "altsvc" : "enabled" }
37
38
       service.citrix.com/backend-httpprofile: '{
39
    "http2direct" : "enabled", "altsvc" : "enabled" }
40
41 spec:
42
    type: LoadBalancer
43
    ports:
     - name: tea-80
44
45
       port: 80
46
       targetPort: 80
47
     selector:
48
       name: tea-beverage
49 EOF
```

After the application is deployed, the corresponding entities and profiles are created on NetScaler. Run the following commands on NetScaler to verify the same: show cs vserver k8s-teabeverage\_80\_default\_svcand show servicegroup k8s-tea-beverage\_80\_sgp\_f4lezsannv

```
# show cs vserver k8s-tea-beverage_80_default_svc
k8s-tea-beverage_80_default_svc (10.105.158.194:80) - HTTP Type:
CONTENT
```

1 2

3

4	State: UP
5	Last state change was at Wed May 8 06:44:11 2024
6	Time since last state change: 0 days, 00:04:24.740
7	Client Idle Timeout: 180 sec
8	Down state flush: ENABLED
9	Disable Primary Vserver On Down : DISABLED
10	Comment: uid=2TZGFGG3JKYS2D7KWJIPS70DTL5D4S7IMU5K6GETMXHXRDE2WIAQ
11	==== HTTP profile name: k8s-tea-beverage 80 default svc
12	Appflow logging: ENABLED
13	Port Rewrite : DISABLED
14	State Update: DISABLED
15	Default: k8s-tea-beverage_80_lbv_f4lezsannvu7tk2ftpjbhi4hza2tvdnk
	Content Precedence: RULE
16	Vserver IP and Port insertion: OFF
17	L2Conn: OFF Case Sensitivity: ON
18	Authentication: OFF
19	401 Based Authentication: OFF
20	Push: DISABLED Push VServer:
21	Push Label Rule: none
22	HTTP Redirect Port: 0 Dtls : OFF
23	Persistence: NONE
24	Listen Policy: NONE
25	IcmpResponse: PASSIVE
26	RHIstate: PASSIVE
27	Traffic Domain: 0
28	
29	1) Default Target LB: k8s-tea-
30	Done

1	# show servicegroup k8s-tea-
	beverage_80_sgp_f4lezsannvu7tk2ftpjbhi4hza2tvdnk
2	k8s-tea-beverage_80_sgp_f4lezsannvu7tk2ftpjbhi4hza2tvdnk – HTTP
3	State: ENABLED Effective State: UP Monitor Threshold : 0
4	Max Conn: 0 Max Req: 0 Max Bandwidth: 0 kbits
5	Use Source IP: NO
6	Client Keepalive(CKA): NO
7	Monitoring Owner: 0
8	TCP Buffering(TCPB): NO
9	HTTP Compression(CMP): NO
10	Idle timeout: Client: 180 sec Server: 360 sec
11	Client IP: DISABLED
12	Cacheable: NO
13	SC: ???
14	SP: OFF
15	Down state flush: ENABLED
16	Monitor Connection Close : NONE
17	Appflow logging: ENABLED
18	HTTP profile name: k8s-tea-
	beverage_80_sgp_t4lezsannvu7tk2ttpjbhi4hza2tvdnk
19	ContentInspection profile name: ???
20	Process Local: DISABLED

```
Traffic Domain: 0
21
22
           Comment: "lbsvc:tea-beverage,svcport:80,ns:default"
23
24
           1) 10.146.107.146:30092 State: UP Server Name:
25
              10.146.107.146 Server ID: None Weight: 1 Order: Default
             Last state change was at Wed May 8 06:44:25 2024
26
27
             Time since last state change: 0 days, 00:04:51.610
28
29
             Monitor Name: tcp-default State: UP Passive: 0
             Probes: 59 Failed [Total: 0 Current: 0]
31
             Last response: Success - TCP syn+ack received.
32
             Response Time: 0.000 millisec
33
           2) 10.146.107.147:30092 State: UP Server Name:
34
              10.146.107.147 Server ID: None Weight: 1 Order: Default
             Last state change was at Wed May 8 06:44:45 2024
             Time since last state change: 0 days, 00:04:31.690
37
             Monitor Name: tcp-default State: UP Passive: 0
38
39
             Probes: 56 Failed [Total: 2 Current: 0]
             Last response: Success - TCP syn+ack received.
40
41
             Response Time: 0.000 millisec
42
     Done
```

#### Note:

The HTTP profile is supported for single-port services.

# **Configure HTTP profiles using Ingress annotations**

The following table lists the HTTP use cases with sample annotations:

Use case	Sample annotation
Configuring HTTP/2	<pre>ingress.citrix.com/frontend- httpprofile: '{ "http2" : " enabled" } '</pre>
	ingress.citrix.com/frontend- http://tepstendctticom/tepstendctticom/tepstendctic
Handling HTTP session timeouts	hhgpess:c1@haple0m/frontend-         hhgppsofsitrix{com/bpsdendct": "         bhgppsofsitrix{com/bpsdendct": "         bhgppsofsitrix{com/bpsdendct": "         hhttppdofeile: '{"Englished: "
	<pre>ihgressreqtimeoutomybiekend-"drop" httpprofile: '{ "apache" : { " httpprofile: '{ "apache" : { " httppred\$reitrix.cem/biedtendaltsvc httpprefile.u.'{ "reatimeout" :</pre>
© 1997–2025 Citrix Systems, Inc. All rights reserved	<pre>int:ppenabled" } }reqtimeout . j"10", "adpttimeout" : "enabled"674 '</pre>
	<pre>ingress.citrix.com/frontend- httpprofile: '{ "reusepooltimeout</pre>

Use case	Sample annotation
	ingress.citrix.com/backend-
	httpprofile: '{ "apache" : { "
	reqtimeout" : "10", "
	<pre>reqtimeoutaction" : "drop" } } '</pre>
	ingress.citrix.com/backend-

## Note:

The above Ingress annotations can also be used with service annotations in the format described earlier.

	<pre>httpprofile: '{ "apache" : { "</pre>
Configuring HTTP/2	<pre>reusepooltimeout" : "20000" } } '</pre>

NetScaler configures HTTP/2 on the client side and on the server side. For more information, see HTTP/2 support on NetScaler. For an HTTP load balancing configuration, NetScaler uses one of the following methods to start communicating with the client/server using HTTP/2.

NetScaler provides configurable options in an HTTP profile for the HTTP/2 methods. These HTTP/2 options can be applied to the client side and the server side of an HTTPS or HTTP load balancing setup. NetScaler Ingress Controller provides annotations to configure the HTTP profile on NetScaler. You use these annotations to configure the various HTTP load balancing configuration on NetScaler to communicate with the client/server using HTTP/2.

Note:

- In the following method descriptions, client and server are generals terms for an HTTP/2 connection. For example, for a load balancing setup of a NetScaler using HTTP/2, NetScaler acts as a server on the client side and acts as a client to the server side.
- Ensure that the HTTP/2 service side global parameter (HTTP2Serverside) is enabled on NetScaler.

# HTTP/2 upgrade

In this method, a client sends an HTTP/1.1 request to a server. The request includes an upgrade header, which asks the server for upgrading the connection to HTTP/2. If the server supports HTTP/2, the server accepts the upgrade request and notifies it in its response. The client and the server start communicating using HTTP/2 after the client receives the upgrade confirmation response.

Using the annotations for HTTP profiles, you can configure the HTTP/2 upgrade method on NetScaler. The following sample annotation configures the HTTP/2 upgrade method on NetScaler:

ingress.citrix.com/frontend-httpprofile: '{ "http2" : "enabled" } '

```
ingress.citrix.com/backend-httpprofile: '{ "apache" : { "http2" : "
enabled" } } '
```

## Direct HTTP/2

In this method, a client directly starts communicating to a server in HTTP/2 instead of using the HTTP/2 upgrade method. If the server does not support HTTP/2 or is not configured to directly accept HTTP/2 requests, it drops the HTTP/2 packets from the client. This method is helpful if the admin of the client device already knows that the server supports HTTP/2.

Using the annotations for HTTP profiles, you can configure the direct HTTP/2 method on NetScaler. The following is a sample annotation of the HTTP profile to configure the direct HTTP/2 method on NetScaler:

```
ingress.citrix.com/backend-httpprofile: '{ "apache" : { "http2direct"
  : "enabled" } } '
```

#### **Direct HTTP/2 using Alternative Service (ALT-SVC)**

In this method, a server advertises that it supports HTTP/2 to a client by including an Alternative Service (ALT-SVC) field in its HTTP/1.1 response. If the client is configured to understand the ALT-SVC field, the client and the server start directly communicating using HTTP/2 after the client receives the response.

The following is a sample annotation for the HTTP profile to configure the direct HTTP/2 using alternative service (ALT-SVC) method on NetScaler:

```
ingress.citrix.com/frontend-httpprofile: '{ "http2direct" : "enabled
", "altsvc" : "enabled" } '
ingress.citrix.com/backend-httpprofile: '{ "apache" : { "http2direct"
  : "enabled", "altsvc" : "enabled" } } '
```

# **Handling HTTP session timeouts**

Various timeouts can be configured to handle different types of HTTP requests and also mitigate attacks such as Slowloris DDOS attack, where you want to restrict client-initiated connections. On NetScaler, you can configure the following timeouts for such scenarios:

reqTimeout and reqTimeoutAction

- adptTimeout
- reusePoolTimeout

#### reqTimeout and reqTimeoutAction

In NetScaler, you can configure the HTTP request timeout value and the request timeout action using the reqTimeout and reqTimeoutAction parameters in the HTTP profile. The reqTimeout value is set in seconds and the HTTP request must complete within the specified time in the reqTimeout parameter. If the HTTP request does not complete within the defined time, the specified request timeout action in the reqTimeoutAction is performed. The minimum timeout value that you can set is 0 and the maximum is 86400. By default, the timeout value is set to 0.

Using the reqTimeoutAction parameter, you can specify the type of action that must be taken in case the HTTP request timeout value (reqTimeout) elapses. You can specify the following actions:

- RESET
- DROP

Using the annotations for HTTP profiles, you can configure the HTTP request timeout and HTTP request timeout action. The following sample annotation of the HTTP profile is used to configure the HTTP request timeout and HTTP request timeout action on NetScaler:

```
ingress.citrix.com/frontend-httpprofile: '{ "reqtimeout" : "10", "
reqtimeoutaction" : "drop" } '
ingress.citrix.com/backend-httpprofile: '{ "apache" : { "reqtimeout"
: "10", "reqtimeoutaction" : "drop" } } '
```

#### adptTimeout

Instead of using a set timeout value for the requested sessions, you can also enable adptTimeout . The adptTimeout parameter adapts the request timeout as per the flow conditions. If enabled, the request timeout is increased or decreased internally and applied on the flow. By default, this parameter is set to DISABLED.

Using annotations for HTTP profiles, you can enable or disable the adpttimeout parameter as following:

```
ingress.citrix.com/frontend-httpprofile: '{ "reqtimeout" : "10", "
adpttimeout" : "enabled" } '
ingress.citrix.com/backend-httpprofile: '{ "apache" : { "reqtimeout"
    : "10", "adpttimeout" : "enabled" } '
```

#### reusePoolTimeout

You can configure a reuse pool timeout value to flush any idle server connections from the reuse pool. If the server is idle for the configured time, then the corresponding connections are flushed.

The minimum timeout value that you can set is 0 and the maximum is 31536000. By default, the timeout value is set to 0.

Using annotations for HTTP profiles, you can configure the required timeout value as follows:

```
ingress.citrix.com/frontend-httpprofile: '{ "reusepooltimeout" :
"20000" } '
ingress.citrix.com/backend-httpprofile: '{ "apache" : { "reusepooltimeout
" : "20000" } } '
```

## Example for applying HTTP profile using Ingress annotation

This example shows how to apply HTTP profiles.

1. Deploy the front-end ingress resource with the HTTP profile. In this Ingress resource, back end and TLS are not defined.

```
kubectl apply -f - <<EOF</pre>
1
2
     apiVersion: networking.k8s.io/v1
3
   kind: Ingress
    metadata:
4
   name: frontend-ingress
annotations:
5
6
7
         ingress.citrix.com/insecure-termination: "allow"
8
         ingress.citrix.com/frontend-ip: "10.221.36.190"
9
         ingress.citrix.com/frontend-httpprofile: '{
10 "reqtimeout" : "10", "adpttimeout" : "enabled" }
11 '
    spec:
13
      tls:
14
       - hosts:
15
     rules:
16
       - host:
17 EOF
```

2. Deploy the secondary ingress resource with the same front-end IP address. In this Ingress resource, back end and TLS are defined, which creates the load balancing resource definition.

```
1 kubectl apply -f - <<EOF
2 apiVersion: networking.k8s.io/v1
3 kind: Ingress
4 metadata:
5 name: backend-ingress
6 annotations:
```

```
ingress.citrix.com/insecure-termination: "allow"
7
8
         ingress.citrix.com/frontend-ip: "10.221.36.190"
9
   spec:
10
      tls:
       - secretName: <hotdrink-secret>
11
12
       rules:
13
       - host: hotdrink.beverages.com
14
         http:
           paths:
15
16
           - path: /
17
             pathType: Prefix
18
             backend:
19
               service:
                 name: frontend-hotdrinks
21
                 port:
22
                   number: 80
23 EOF
```

3. After the ingress resources are deployed, the corresponding entities and profiles are created on NetScaler.

Run the following command on NetScaler to verify the same: show cs vserver k8s -10.221.36.190\_443\_ssl.

1	# show cs vserver k8s-10.221.36.190_443_ssl
2	k8s-10.221.36.190_443_ssl (10.221.36.190:443) - SSL Type:
	CONTENT
3	State: UP
4	Last state change was at Wed Apr 3 04:50:56 2024
5	Time since last state change: 0 days, 00:00:31.590
6	Client Idle Timeout: 180 sec
7	Down state flush: ENABLED
8	Disable Primary Vserver On Down : DISABLED
9	Comment: uid=
	IIFGNN4JZ4LBQ7VCBTX6DN7TIBGG0U30D20BIC7MB7302B2FNL2Q====
10	**HTTP profile name: k8s-10.221.36.190_443_ssl**
11	Appflow logging: ENABLED
12	State Update: DISABLED
13	Default: Content Precedence: RULE
14	Vserver IP and Port insertion: OFF
15	L2Conn: OFF Case Sensitivity: ON
16	Authentication: OFF
17	401 Based Authentication: OFF
18	Push: DISABLED Push VServer:
19	Push Label Rule: none
20	HTTP Redirect Port: 0 Dtls : OFF
21	Persistence: NONE
22	Listen Policy: NONE
23	IcmpResponse: PASSIVE
24	RHIstate: PASSIVE
25	Trattic Domain: 0
26	
27	<ol> <li>Content-Switching Policy: k8s-</li> </ol>

backend_80_csp_2k75kfjrr6ptgzwtncozwxdjqrpbvicz Rule: HTTP
.REQ.HOSTNAME.SERVER.EQ("hotdrink.beverages.com") && HTTP.
REQ.URL.PATH.SET_TEXT_MODE(IGNORECASE).STARTSWITH("/")
Priority: 20000008 Hits: 0
Done

#### Note:

28

For various HTTP parameters supported on NetScaler, refer to NetScaler 14.1 REST APIs - NITRO Documentation for HTTP profiles. The key and value that you pass in the JSON format must match the values in the NITRO format.

# HTTP callout with the rewrite and responder policy

#### December 31, 2023

An HTTP callout allows NetScaler to generate and send an HTTP or HTTPS request to an external server (callout agent) as part of the policy evaluation. The information that is retrieved from the server (callout agent) can be analyzed by advanced policy expressions and an appropriate action can be performed. For more information about the HTTP callout, see the NetScaler documentation.

You can initiate the HTTP callout through the following expressions with the rewrite and responder CRD provided by NetScaler:

- sys.http\_callout(): This expression is used for blocking the call when the httpcallout agent response needs to be evaluated.
- sys.non\_blocking\_http\_callout(): This expression is used for non-blocking calls (for example: traffic mirroring)

These expressions accept the httpcallout\_policy name defined in the CRD as a parameter, where the name needs to be specified in double quotes.

For example: sys.http\_callout("callout\_name").

In this expression, callout\_name refers to the appropriate httpcallout\_policy defined in the rewrite and responder CRD YAML file.

The following table explains the attributes of the HTTP callout request in the rewrite and responder CRD.

Parameter	Description
name	Specifies the name of the callout, maximum is
	up to 32 characters.
server_ip	Specifies the IP Address of the server (callout
	agent) to which the callout is sent.
server_port	Specifies the Port of the server (callout agent) to
	which the callout is sent.
http_method	Specifies the method used in the HTTP request
	that this callout sends. The default value is GET.
host_expr	Specifies the text expression to configure the
	host header. This expression can be a literal
	value (for example, 192.101.10.11) or it can be an
	advanced expression (for example,
	http.req.header("Host")) that derives the value.
	The literal value can be an IP address or a fully
	qualified domain name. Mutually exclusive with
	the full HTTP request expression.
url_stem_expr	Specifies a string expression for generating the
	URL stem. The string expression can contain a
	literal string (for example, "/mysite/index.html")
	or an expression that derives the value (for
	example, http.req.url).
headers	Specifies one or more headers to insert into the
	HTTP request. Each header name and exp,
	where exp is an expression that is evaluated at
	runtime to provide the value for the named
	header.
parameters	Specifies one or more query parameters to insert
	into the HTTP request URL (for a GET request) or
	into the request body (for a POST request). Each
	parameter is represented by a name and an
	expr, where expr is an expression that is
	evaluated at run time to provide the value for the
	named parameter (name=value). The parameter
	values are URL encoded.
body_expr	An advanced string expression for generating the
	body of the request. The expression can contain
	a literal string or an expression that derives the
	value (for example, client.ip.src).

Parameter	Description
full_req_expr	Specifies the exact HTTP request, in the form of an expression, which the NetScaler sends to the callout agent. The request expression is constrained by the feature for which the callout
	is used. For example, an HTTP.RES expression cannot be used in a request-time policy bank or in a TCP content switching policy bank
scheme	Specifies the type of scheme for the callout server. Example: HTTP, HTTPS
return_type	Specifies the type of data that the target callout agent returns in response to the callout. The available settings function as follows: TEXT - Treat the returned value as a text string. NUM -
cache_for_secs	Treat the returned value as a boolean value. Specifies the duration, in seconds, for which the callout response is cached. The cached
	responses are stored in an integrated caching content group named calloutContentGroup. If the duration is not configured, the callout responses are not cached
	unless a normal caching configuration is used to cache them. This parameter takes precedence over any normal caching configuration that would otherwise apply to these responses.
result_expr	Specifies the expression that extracts the callout results from the response sent by the HTTP callout agent. This expression must be a response based expression, that is, it must begin with HTTP, RES. The operations in this
	expression must match the return type. For example, if you configure a return type of TEXT,
	the result expression must be a text based expression. If the return type is NUM, the result expression (result_expr) must return a numeric
	value, as in the following example: http.res.body(10000).length

Parameter	Description
comment	Specifies any comments to preserve the
	information about this HTTP callout.

# Using the rewrite and responder CRD to validate whether a client IP address is blocklisted

This section shows how to initiate an HTTP callout using the rewrite and responder CRD to validate whether a client IP address is blocklisted or not and take appropriate action.

The following diagram explains the workflow of a request where each number in the diagram denotes a step in the workflow:



- 1. Client request
- 2. HTTP callout request to check if the client is blocklisted (The client IP address is sent as a query parameter with the name Cip)
- 3. Response from the HTTP callout server
- 4. Request is forwarded to the service if the response in step 3 indicates a safe IP address (the client IP address is not matching with the blocklisted IP addresses on the callout server).
- 5. Respond to the client as Access denied, if the response in step 3 indicates a bad IP address (the client IP address is matching with the blocklisted IP addresses on the callout server).

The following is a sample YAML file (ip\_validate\_responder.yaml) for validating a blocklisted IP address:
## Note:

You must deploy the rewrite and responder CRD before deploying the ip\_validate\_responder YAML file.

```
1 apiVersion: citrix.com/v1
2
  kind: rewritepolicy
3 metadata:
4 name: validateip
5 spec:
6 responder-policies:
       - servicenames:
7
8
           - frontend
         responder-policy:
9
10
           respondwith:
             http-payload-string: '"HTTP/1.1 401 Access denied\r\n\r\n"'
11
           respond-criteria: 'sys.http_callout("blocklist_callout").
12
               CONTAINS("IP Matched")' #Callout name needs to be given in
               double quotes to pick httpcallout_policy
           comment: 'Invalid access'
13
14
     httpcallout_policy:
15
16
       - name: blocklist_callout
17
         server_ip: "192.2.156.160"
18
         server_port: 80
         http_method: GET
19
20
         host_expr: '"192.2.156.160"'
         url_stem_expr: '"/validateIP.pl"'
21
22
         headers:
         - name: X-Request
23
24
           expr: '"Callout Request"'
25
         parameters:
26
         - name: Cip
27
           expr: 'CLIENT.IP.SRC'
28
         return_type: TEXT
29
         result_expr: 'HTTP.RES.BODY(100)'
```

# Using the rewrite and responder CRD to update the URL with a valid path requested by the client

This section shows how to initiate an HTTP callout using the rewrite and responder CRD when a path exposed to the client is different from the actual path due to security reasons.

The work flow of a request is explained in the following diagram where each number in the diagram denotes a step in the workflow.



- 1. Client request
- 2. HTTP callout request to get the valid path (the path requested from the client is sent as a query parameter with the name path to the callout server)
- 3. Response from the HTTP callout server
- 4. The URL request is rewritten with a valid path and forwarded to the service (where the valid path is mentioned between the tags newpath in the callout response).

The following is a sample YAML (path\_rewrite) file.

Note:

You must deploy the rewrite and responder CRD before deploying the path\_rewrite YAML file.

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4 name: getvalidpath
5 spec:
   rewrite-policies:
6
7
       - servicenames:
           - frontend
8
         rewrite-policy:
9
           operation: replace
10
           target: http.req.url
           modify-expression: 'sys.http_callout("mapping_callout")' #
              Callout name needs to be given in double quotes to pick
              httpcallout_policy
           comment: 'Get the valid path'
13
14
           direction: REQUEST
```

```
15
            rewrite-criteria: 'TRUE'
16
17
     httpcallout_policy:
18
       - name: mapping_callout
         server_ip: "192.2.156.160"
19
20
         server_port: 80
21
         http_method: GET
         host_expr: '"192.2.156.160"'
22
23
         url_stem_expr: '"/getPath.pl"'
24
         headers:
25
         - name: X-Request
           expr: '"Callout Request"'
26
27
         parameters:
28
         - name: path
29
           expr: 'http.req.url'
         return_type: TEXT
         result_expr: '"HTTP.RES.BODY(500).AFTER_STR(\"<newpath>\").
31
             BEFORE_STR(\"</newpath>\")"'
```

# **Configure session affinity or persistence on the Ingress NetScaler**

December 31, 2023

Session affinity or persistence settings on the Ingress NetScaler allows you to direct client requests to the same selected server regardless of which virtual server in the group receives the client request. When the configured time for persistence expires, any virtual server in the group is selected for the incoming client requests.

If persistence is configured, it overrides the load balancing methods once the server has been selected. It maintains the states of connections on the servers represented by that virtual server. The NetScaler then uses the configured load balancing method for the initial selection of a server, but forwards to that same server all subsequent requests from the same client.

The most commonly used persistence type is persistence based on cookies.

#### **Configure persistence based on cookies**

When you enable persistence based on cookies, the NetScaler adds an HTTP cookie into the Set-Cookie header field of the HTTP response. The cookie contains information about the service to which the HTTP requests must be sent. The client stores the cookie and includes it in all subsequent requests, and the ADC uses it to select the service for those requests.

The NetScaler inserts the cookie <NSC\_XXXX>= <ServiceIP> <ServicePort>.

Where:

- <<NSC\_XXXX> is the virtual server ID that is derived from the virtual server name.
- <<ServiceIP> is the hexadecimal value of the IP address of the service.
- <<ServicePort> is the hexadecimal value of the port of the service.

The NetScaler encrypts ServiceIP and ServicePort when it inserts a cookie, and decrypts them when it receives a cookie.

```
For example, a.com=fffffff02091f1045525d5f4f58455e445a4a423660;expires= Fri, 23-Aug-2019 07:01:45.
```

You can configure persistence setting on the ingress NetScaler, using the following Ingress annotation provided by the NetScaler Ingress Controller:

Where:

- timeout specifies the duration of persistence. If session cookies are used with a timeout value of 0, no expiry time is specified by NetScaler regardless of the HTTP cookie version used. The session cookie expires when the Web browser is closed
- cookiename specifies the name of cookie with a maximum of 32 characters. If not specified, cookie name is internally generated.
- persistenceType here specifies the type of persistence to be used, COOKIEINSERT is used to cookie based persistence. Apart from cookie, other options can also be used along with appropriate arguments and other required parameters.

Possible values are SOURCEIP, SSLSESSION, DESTIP, SRCIPDESTIP, and so on.

#### **Source IP address persistence**

When source IP persistence is configured on the Ingress NetScaler, you can set persistence to an load balancing virtual server, that creating a stickiness for the subsequest requests from the same client.

The following is a sample Ingress annotation to configure source IP address persistence:

```
1 ingress.citrix.com/lbvserver: '{
2 "apache":{
3 "persistenceType":"SOURCEIP", "timeout":"10" }
4 }
5 '
```

### **SSL** session ID persistence

When SSL session ID persistence is configured, the NetScaler appliance uses the SSL session ID, which is part of the SSL handshake process, to create a persistence session before the initial request is directed to a service. The load balancing virtual server directs subsequent requests that have the same SSL session ID to the same service. This type of persistence is used for SSL bridge services.

The following is a sample Ingress annotation to configure SSL session ID persistence:

```
1 ingress.citrix.com/lbvserver: '{
2 "apache":{
3 "persistenceType":"SSLSESSION" }
4 }
5 '
```

#### **Destination IP address-based persistence**

In this type of persistence, when the Ingress NetScaler receives a request from a new client, it creates a persistence session based on the IP address of the service selected by the virtual server (the destination IP address). Subsequently, it directs requests to the same destination IP to the same service. This type of persistence is used with link load balancing.

The following is a sample Ingress annotation to configure destination IP address-based persistence:

```
1 ingress.citrix.com/lbvserver: '{
2 "apache":{
3 "persistenceType":"DESTIP" }
4 }
5 '
```

# Source and destination IP address-based persistence

In this type of persistence, when the NetScaler appliance receives a request, it creates a persistence session based on both the IP address of the client (the source IP address) and the IP address of the service selected by the virtual server (the destination IP address). Subsequently, it directs requests from the same source IP and to the same destination IP to the same service.

The following is a sample Ingress annotation to configure source and destination IP address-based persistence:

```
1 ingress.citrix.com/lbvserver: '{
2 "apache":{
3 "persistenceType":"SRCIPDESTIP" }
4 }
5 '
```

# **Allowlisting or blocklisting IP addresses**

## December 31, 2023

Allowlisting IP addresses allows you to create a list of trusted IP addresses or IP address ranges from which users can access your domains. It is a security feature that is often used to limit and control access only to trusted users.

Blocklisting IP addresses is a basic access control mechanism. It denies access to the users accessing your domain using the IP addresses that you have blocklisted.

The Rewrite and Responder CRD provided by NetScaler enables you to define extensive rewrite and responder policies using datasets, patsets, and string maps and also enable audit logs for statistics on the Ingress NetScaler.

Using the rewrite or responder policies you can allowlist or blocklist the IP addresses/CIDR using which users can access your domain.

The following sections cover various ways you can allowlist or blocklist the IP addresses/CIDR using the rewrite or responder policies.

## **Allowlist IP addresses**

Using a responder policy, you can allowlist IP addresses and silently drop the requests from the clients using IP addresses different from the allowlisted IP addresses.

Create a file named allowlist-ip.yaml with the following rewrite policy configuration:

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4 name: allowlistip
5 spec:
   responder-policies:
6
7
       - servicenames:
8
           - frontend
9
         responder-policy:
           drop:
           respond-criteria: '!client.ip.src.TYPECAST_text_t.equals_any("
              allowlistip")'
           comment: 'Allowlist certain IP addresses'
12
13
     patset:
14
       - name: allowlistip
15
         values:
16
           - '10.xxx.170.xx'
17
           - '10.xxx.16.xx'
```

You can also provide the IP addresses as a list:

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4
    name: allowlistip
5 spec:
6
    responder-policies:
7
       - servicenames:
           - frontend
8
9
         responder-policy:
10
           drop:
           respond-criteria: '!client.ip.src.TYPECAST_text_t.equals_any("
11
               allowlistip")'
           comment: 'Allowlist certain IP addresses'
12
     patset:
14
       - name: allowlistip
15
         values: [ '10.xxx.170.xx', '10.xxx.16.xx' ]
```

Then, deploy the YAML file (allowlist-ip.yaml) using the following command:

1 kubectl create -f allowlist-ip.yaml

# Allowlist IP addresses and send 403 response to the request from clients not in the allowlist

Using a responder policy, you can allowlist a list of IP addresses and send the HTTP/1.1 403 Forbidden response to the requests from the clients using IP addresses different from the allowlisted IP addresses.

Create a file named allowlist-ip-403.yaml with the following rewrite policy configuration:

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
   name: allowlistip
4
5
   spec:
   responder-policies:
6
       - servicenames:
7
8
           - frontend
9
         responder-policy:
10
           respondwith:
11
             http-payload-string: '"HTTP/1.1 403 Forbidden\r\n\r\n" + "
                 Client: " + CLIENT.IP.SRC + " is not authorized to access
                URL:" + HTTP.REQ.URL.HTTP_URL_SAFE +"\n"'
12
           respond-criteria: '!client.ip.src.TYPECAST_text_t.equals_any("
               allowlistip")'
13
           comment: 'Allowlist a list of IP addresses'
14
     patset:
15
       - name: allowlistip
16
         values: [ '10.xxx.170.xx', '10.xxx.16.xx' ]
```

Then, deploy the YAML file (allowlist-ip-403.yaml) using the following command:

```
1 kubectl create -f allowlist-ip-403.yaml
```

#### **Allowlist a CIDR**

You can allowlist a CIDR using a responder policy. The following is a sample responder policy configuration to allowlist a CIDR:

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4
   name: blocklistips1
5
  spec:
   responder-policies:
6
       - servicenames:
7
           - frontend
8
9
         responder-policy:
           respondwith:
10
             http-payload-string: '"HTTP/1.1 403 Forbidden\r\n\r\n" + "
11
                Client: " + CLIENT.IP.SRC + " is not authorized to access
                URL:" + HTTP.REQ.URL.HTTP_URL_SAFE +"\n"'
12
           respond-criteria: '!client.ip.src.IN_SUBNET(10.xxx.170.xx/24)'
           comment: 'Allowlist certain IPs'
13
```

#### **Blocklist IP addresses**

Using a responder policy, you can blocklist IP addresses and silently drop the requests from the clients using the blocklisted IP addresses.

Create a file named blocklist-ip.yaml with the following responder policy configuration:

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4
   name: blocklistips
5 spec:
   responder-policies:
6
       - servicenames:
7
8
           - frontend
9
         responder-policy:
10
           respondwith:
11
           drop:
           respond-criteria: 'client.ip.src.TYPECAST_text_t.equals_any("
              blocklistips")'
13
           comment: 'Blocklist certain IPS'
14
15
     patset:
16

    name: blocklistips
```

```
    17
    values:

    18
    - '10.xxx.170.xx'

    19
    - '10.xxx.16.xx'
```

Then, deploy the YAML file (blocklist-ip.yaml) using the following command:

```
1 kubectl create -f blocklist-ip.yaml
```

# **Blocklist a CIDR**

You can blocklist a CIDR using a responder policy. The following is a sample responder policy configuration to blocklist a CIDR:

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4
    name: blocklistips1
5 spec:
    responder-policies:
6
7
       - servicenames:
8
           - frontend
9
         responder-policy:
10
           respondwith:
             http-payload-string: '"HTTP/1.1 403 Forbidden\r\n\r\n" + "
11
                Client: " + CLIENT.IP.SRC + " is not authorized to access
                URL:" + HTTP.REQ.URL.HTTP_URL_SAFE +"\n"'
           respond-criteria: 'client.ip.src.IN_SUBNET(10.xxx.170.xx/24)'
12
           comment: 'Blocklist certain IPs'
13
```

# Allowlist a CIDR and blocklist IP addresses

You can allowlist a CIDR and also blocklist IP addresses using a responder policy. The following is a sample responder policy configuration:

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
   name: allowlistsub
4
5 spec:
6
    responder-policies:
7
       - servicenames:
           - frontend
8
9
         responder-policy:
10
           drop:
           respond-criteria: 'client.ip.src.TYPECAST_text_t.equals_any("
11
              blocklistips") || !client.ip.src.IN_SUBNET(10.xxx.170.xx/24)
           comment: 'Allowlist a subnet and blocklist few IP's'
12
13
```

```
14 patset:
15 - name: blocklistips
16 values:
17 - '10.xxx.170.xx'
```

# **Blocklist a CIDR and allowlist IP addresses**

You can blocklist a CIDR and also allowlist IP addresses using a responder policy. The following is a sample responder policy configuration:

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4
  name: blocklistips1
5 spec:
   responder-policies:
6
       - servicenames:
7
8
           - frontend
         responder-policy:
9
10
           drop:
           respond-criteria: 'client.ip.src.IN_SUBNET(10.xxx.170.xx/24) &&
11
                !client.ip.src.TYPECAST_text_t.equals_any("allowlistips")'
           comment: 'Blocklist a subnet and allowlist few IP's'
12
13
14
     patset:
15
       - name: allowlistips
         values:
17
           - '10.xxx.170.xx'
18
           - '10.xxx.16.xx'
```

# Interoperability with ExternalDNS

#### December 31, 2023

In a Kubernetes environment, you can expose your deployment using a service of type LoadBalancer . Also, an IP address can be assigned to the service using . The assigns IP address to the service from a defined pool of IP addresses. For more information, see Expose services of type LoadBalancer with IP addresses assigned by the IPAM controller.

The service can be accessed using the IP address assigned by the IPAM controller and for service discovery you need to manually register the IP address to a DNS provider. If the IP address assigned to the service changes, the associated DNS record must be manually updated and the entire process becomes cumbersome. In such cases, you can use a ExternalDNS to keep the DNS records synchronized with your external entry points. Also, ExternalDNS allows you to control DNS records dynamically through Kubernetes resources in a DNS provider-agnostic way. For the ExternalDNS integration to work, the external-dns.alpha.kubernetes.io/ hostname annotation must contain the host name.

#### Note:

For ExtenalDNS to work, ensure that you add the annotation external-dns.alpha. kubernetes.io/hostname in the service specification and specify a host name for the service using the annotation.

#### To integrate with ExternalDNS:

1. Install the ExternalDNS with Infoblox provider.

#### Note:

The interoperability solution has been tested with Infoblox provider and the solution might work for other providers as well.

- 2. Specify the domain name in the ExternalDNS configuration.
- 3. In the service of type LoadBalancer specification, add the following annotation and specify a host name for the service using the annotation:

1 external-dns.alpha.kubernetes.io/hostname

4. Deploy the service using the following command:

1 kubectl create -f <service-name>.yml

# Using NetScaler credentials stored in a Vault server for the NetScaler Ingress Controller

#### December 31, 2023

In most organizations, tier 1 NetScaler Ingress devices and Kubernetes clusters are managed by separate teams. Usually, network administrators manage tier 1 NetScaler Ingress devices, while developers manage Kubernetes clusters. The NetScaler Ingress Controller requires NetScaler credentials such as NetScaler user name and password to configure the NetScaler. You can specify NetScaler credentials as part of the NetScaler Ingress Controller specification and store the ADC credentials as Kubernetes secrets. However, you can also store NetScaler credentials in a Vault server and pass credentials to the NetScaler Ingress Controller to minimize any security risk. This topic provides information on how to use NetScaler credentials stored in a Vault server for the NetScaler Ingress Controller. The following diagram explains the steps for using NetScaler credentials which are stored in a Vault server with the NetScaler Ingress Controller.



# Prerequisites

Ensure that you have setup a Vault server and enabled key-value (KV) secret store. For more information, see Vault documentation.

# Using NetScaler credentials from a Vault server for the NetScaler Ingress Controller

Perform the following tasks to use NetScaler credentials from a Vault server for the NetScaler Ingress Controller.

- 1. Create a service account for Kubernetes authentication.
- 2. Create a Key Vault secret and setup Kubernetes authentication on Vault server.
- 3. Leverage Vault Auto-Auth functionality to fetch NetScaler credentials for the NetScaler Ingress Controller.

### Create a service account for Kubernetes authentication

Create a service account for Kubernetes authentication by using the following steps:

1. Create a service account cic-k8s-role and provide the service account necessary permissions to access the Kubernetes TokenReview API by using the following command.

Following is a part of the sample cic-k8s-role-service-account.yml file.

```
1 apiVersion: rbac.authorization.k8s.io/v1
2 kind: ClusterRoleBinding
3 metadata:
4 name: role-tokenreview-binding
5 namespace: default
6 roleRef:
7 apiGroup: rbac.authorization.k8s.io
8 kind: ClusterRole
9 name: system:auth-delegator
10 subjects:
11 - kind: ServiceAccount
12 name: cic-k8s-role
13 namespace: default
```

2. Set the VAULT\_SA\_NAME environment variable to the name of the service account you have already created.

3. Set the SA\_JWT\_TOKEN environment variable to the JWT of the service account that you used to access the TokenReview API.

4. Get a Kubernetes CA signed certificate to communicate with Kubernetes API.

#### Create a key vault secret and setup Kubernetes authentication on the Vault server

Log in to the Vault server and perform the following steps to create a Key Vault secret and setup Kubernetes authentication.

1. Review the sample vault policy file citrix-adc-kv-ro.hcl and create a read-only policy, citrixadc-kv-ro in Vault.

```
1 $ tee citrix-adc-kv-ro.hcl <<EOF</pre>
2 # If working with K/V v1
3 path "secret/citrix-adc/*"
4 {
5
       capabilities = ["read", "list"]
6
7
    }
8
9 # If working with K/V v2
  path "secret/data/citrix-adc/*"
10
11 {
12
13
       capabilities = ["read", "list"]
14
    }
15
16 EOF
17
18 # Create a policy named citrix-adc-kv-ro
19 $ vault policy write citrix-adc-kv-ro citrix-adc-kv-ro.hcl
```

2. Create a KV secret with NetScaler credentials at the secret/citrix-adc/path.

```
1 vault kv put secret/citrix-adc/credential username='<ADC
username>' \
2 password='<ADC password>' \
3 ttl='30m'
```

3. Enable Kubernetes authentication at the default path (auth/kubernetes).

1 # \$ vault auth enable kubernetes

4. Specify how to communicate with the Kubernetes cluster.

```
1 $ vault write auth/kubernetes/config \
2 token_reviewer_jwt="$SA_JWT_TOKEN" \
3 kubernetes_host="https://<K8S_CLUSTER_URL>:<API_SERVER_PORT>" \
4 kubernetes_ca_cert="$SA_CA_CRT"
```

5. Create a role to map the Kubernetes service account to Vault policies and the default token TTL. This role authorizes the cic-k8s-role service account in the default namespace and maps the service account to the citrix-adc-kv-ro policy.

1 \$ vault write auth/kubernetes/role/cic-vault-example\

```
2 bound_service_account_names=cic-k8s-role \
3 bound_service_account_namespaces=default \
4 policies=citrix-adc-kv-ro \
5 ttl=24h
```

#### Note:

Authorization with Kubernetes authentication back-end is role based. Before a token is used for login, it must be configured as part of a role.

#### Leverage Vault agent auto-authentication for the NetScaler Ingress Controller

Perform the following steps to leverage Vault auto-authentication.

1. Review the provided Vault Agent configuration file, vault-agent-config.hcl.

```
1
        exit_after_auth = true
        pid_file = "/home/vault/pidfile"
2
3
4
        auto_auth {
5
6
            method "kubernetes" {
7
8
                mount_path = "auth/kubernetes"
9
                config = {
10
11
                     role = "cic-vault-example"
                 }
12
13
             }
14
15
16
17
            sink "file" {
18
                config = {
19
20
21
                     path = "/home/vault/.vault-token"
22
                 }
23
             }
24
25
26
         }
```

#### Note:

The Vault agent Auto-Auth is configured to use the Kubernetes authentication method enabled at the auth/kubernetes path on the Vault server. The Vault Agent uses the cic-vault-example role to authenticate.

The sink block specifies the location on disk where to write tokens. Vault Agent Auto-

Auth sink can be configured multiple times if you want Vault Agent to place the token into multiple locations. In this example, the sink is set to /home/vault/.vault-token.

# 2. Review the Consul template consul-template-config.hcl

```
file.
```

1	vault {
2	
3	renew_token = <b>false</b>
4	<pre>vault_agent_token_file = "/home/vault/.vault-token"</pre>
5	retry {
6	
7	<pre>backoff = "1s"</pre>
8	}
9	
10	}
11	
12	
13	template {
14	
15	<pre>destination = "/etc/citrix/.env"</pre>
16	contents = < <eoh< td=""></eoh<>
17	NS_USER=
18	NS_PASSWORD=
19	
20	EOH
21	}

This template reads secrets at the secret/citrix-adc/credential path and sets the user name and password values.

If you are using KV store version 1, use the following template.

```
1
       template {
2
3
           destination = "/etc/citrix/.env"
4
           contents = <<EOH
5
           NS_USER=
           NS_PASSWORD=
6
7
8
           EOH
9
                }
```

3. Create a Kubernetes config-map from vault-agent-config.hcl and consul-template-config.hcl.

```
1 kubectl create configmap example-vault-agent-config --from-
file=./vault-agent-config.hcl --form-file=./consul-template
-config.hcl
```

4. Create a NetScaler Ingress Controller pod with Vault and consul template as init container citrixk8s-ingress-controller-vault.yaml. Vault fetches the token using the Kubernetes authentication method and pass it on to a consul template which creates the .env file on shared volume. This token is used by the NetScaler Ingress Controller for authentication with tier 1 NetScaler.

1 kubectl apply citrix-k8s-ingress-controller-vault.yaml

The citrix-k8s-ingress-controller-vault.yaml file is as follows:

1	apiVersion: v1
2	kind: Pod
3	metadata:
4	annotations:
5	name: cic-vault
6	namespace: <b>default</b>
7	spec:
8	containers:
9	- args:
10	ingress-classes tier-1-vpx
11	feature-node-watch <b>true</b>
12	env:
13	- name: NS_IP
14	value: <tier 1="" adc="" ip-address=""></tier>
15	- name: EULA
16	value: "yes"
17	<pre>image: in-docker-reg.eng.citrite.net/cpx-dev/kumar-cic</pre>
1.0	
10	ImagePullPolicy: Always
19	name: cic-k8s-ingress-controller
20	volumemounts:
21	- mountPath: /etc/citrix
22	name: snared-data
23	initContainers:
24	- args:
25	- agent
26	config=/etc/vault/vault-agent-config.ncl
21	log-level=debug
28	
29	
30	value: <vault< td=""></vault<>
31	image: Vault
32	name woult agent outh
33	name: Vault-agent-auth
34	volumemounts:
30	- mountPath: /etc/vault
30	name: contrig
20	- mountPath: /nome/vautt
20	
39	- digs.
40	hcl
41	log-level=debug
42	once
43	env:
44	- name: HOME
45	value: /home/vault
46	<pre>- name: VAULT_ADDR</pre>

47	value: <vault_url></vault_url>
48	<pre>image: hashicorp/consul-template:alpine</pre>
49	<pre>imagePullPolicy: Always</pre>
50	name: consul-template
51	volumeMounts:
52	<pre>- mountPath: /home/vault</pre>
53	name: vault-token
54	<pre>- mountPath: /etc/consul-template</pre>
55	name: config
56	- mountPath: /etc/citrix
57	name: shared-data
58	serviceAccountName: vault-auth
59	volumes:
60	- emptyDir:
61	medium: Memory
62	name: vault-token
63	- configMap:
64	defaultMode: 420
65	items:
66	- key: vault-agent-config.hcl
67	path: vault-agent-config.hcl
68	- key: consul-template-config.hcl
69	name: example-vault-agent-config
70	name: config
71	- emptyDir:
72	medium: Memory
73	name: shared-data

If the configuration is successful, the Vault server fetches a token and passes it on to a Consul template container. The Consul template uses the token to read NetScaler credentials and write it as an environment variable in the path /etc/citrix/.env. The NetScaler Ingress Controller uses these credentials for communicating with the tier 1 NetScaler.

Verify that the NetScaler Ingress Controller is running successfully using credentials fetched from the Vault server.

# How to use Kubernetes secrets for storing NetScaler credentials

#### December 31, 2023

In most organizations, Tier 1 NetScaler Ingress devices and Kubernetes clusters are managed by separate teams. The NetScaler Ingress Controller requires NetScaler credentials such as NetScaler user name and password to configure the NetScaler. Usually, NetScaler credentials are specified as environment variables in the NetScaler Ingress Controller pod specification. But, another secure option is to use Kubernetes secrets to store the NetScaler credentials. This topic describes how to use Kubernetes secrets to store the ADC credentials and various ways to provide the credentials stored as secret data for the NetScaler Ingress Controller.

## **Create a Kubernetes secret**

Perform the following steps to create a Kubernetes secret.

1. Create a file adc-credential-secret.yaml which defines a Kubernetes secret YAML with NetScaler user name and password in the data section as follows.

```
1 apiVersion: v1
2 kind: Secret
3 metadata:
4 name: adc-credential
5 data:
6 username: <ADC user name>
7 password: <ADC password>
```

2. Apply the adc-credential-secret.yaml file to create a secret.

1 kubectl apply -f adc-credential-secret.yaml

Alternatively, you can also create the Kubernetes secret using --from-literal option of the kubectl command as shown as follows:

```
1 kubectl create secret generic adc-credentials --from-literal=
    username=<username> --from-literal=password=<password>
```

Once you have created a Kubernetes secret, you can use one of the following options to use the secret data in the NetScaler Ingress Controller pod specification.

- Use secret data as environment variables in the NetScaler Ingress Controller pod specification
- Use a secret volume mount to pass credentials to the NetScaler Ingress Controller

# Use secret data as environment variables in the NetScaler Ingress Controller pod specification

You can use secret data from the Kubernetes secret as the values for the environment variables in the NetScaler Ingress Controller deployment specification.

A snippet of the YAML file is shown as follows.

```
    name: "NS_USER"
    valueFrom:
    secretKeyRef:
    name: adc-credentials
    key: username
```

```
6 # Set user password for Nitro
7 - name: "NS_PASSWORD"
8 valueFrom:
9 secretKeyRef:
10 name: adc-credentials
11 key: password
```

Here is an example of the NetScaler Ingress Controller deployment with value of environment variables sourced from the secret object.

```
1 apiVersion: apps/v1
2 kind: Deployment
3 metadata:
   name: cic-k8s-ingress-controller
4
5 spec:
    selector:
6
7
       matchLabels:
8
         app: cic-k8s-ingress-controller
9
   replicas: 1
10
     template:
11
       metadata:
12
         name: cic-k8s-ingress-controller
13
         labels:
14
           app: cic-k8s-ingress-controller
         annotations:
16
       spec:
         serviceAccountName: cic-k8s-role
17
18
         containers:
19
         - name: cic-k8s-ingress-controller
           image: <image location>
20
21
           env:
            # Set NetScaler NSIP/SNIP, SNIP in case of HA (mgmt has to be
22
                enabled)
23
            - name: "NS_IP"
              value: "x.x.x.x"
24
            # Set username for Nitro
25
            - name: "NS_USER"
26
27
              valueFrom:
28
               secretKeyRef:
29
                 name: adc-credentials
                 key: username
31
            # Set user password for Nitro
            - name: "NS_PASSWORD"
33
              valueFrom:
34
               secretKeyRef:
                 name: adc-credentials
                 key: password
37
            # Set log level
            - name: "EULA"
39
              value: "yes"
40
           imagePullPolicy: Always
```

# Use a secret volume mount to pass credentials to the NetScaler Ingress Controller

Alternatively, you can also use a volume mount using the secret object as a source for the NetScaler credentials. The NetScaler Ingress Controller expects the secret to be mounted at path /etc/ citrix and it looks for the credentials in files username and password.

You can create a volume from the secret object and then mount the volume using volumeMounts at /etc/citrix as shown in the following deployment example.

```
1 apiVersion: apps/v1
2 kind: Deployment
3 metadata:
4 name: cic-k8s-ingress-controller
5 spec:
6
    selector:
7
      matchLabels:
8
         app: cic-k8s-ingress-controller
9
   replicas: 1
   template:
11
       metadata:
12
         name: cic-k8s-ingress-controller
13
         labels:
14
           app: cic-k8s-ingress-controller
15
         annotations:
16
       spec:
         serviceAccountName: cic-k8s-role
17
18
         containers:
         - name: cic-k8s-ingress-controller
19
20
           image: <image location>
21
           env:
            # Set NetScaler NSIP/SNIP, SNIP in case of HA (mgmt has to be
22
                enabled)
            - name: "NS_IP"
23
              value: "x.x.x.x"
24
25
            # Set log level
            - name: "EULA"
26
              value: "yes"
27
28
           volumeMounts:
29
           # name must match the volume name below
             - name: secret-volume
31
               mountPath: /etc/citrix
           imagePullPolicy: Always
32
         # The secret data is exposed to Containers in the Pod through a
            Volume.
34
         volumes:
35
         - name: secret-volume
           secret:
             secretName: adc-credentials
37
```

# Use NetScaler credentials stored in a Hashicorp Vault server

You can also use the NetScaler credentials stored in a Hashicorp Vault server for the NetScaler Ingress Controller and push the credentials through a sidecar container.

For more information, see Using NetScaler credentials stored in a Vault server.

# How to load balance ingress traffic to UDP-based application

#### May 23, 2024

In a Kubernetes environment, an ingress object allows access to the Kubernetes services from outside the Kubernetes cluster. Standard Kubernetes ingress resources assume that all the traffic is HTTPbased and do not cater to non-HTTP based protocols such as TCP, UDP, and SSL. Hence, any non-HTTP applications such as DNS, FTP or LDAP cannot be exposed using the standard Kubernetes ingress.

NetScaler provides a solution using ingress annotations to load balance UDP-based ingress traffic. When you specify these annotations in the ingress resource definition, NetScaler Ingress Controller configures NetScaler to load balance UDP-based ingress traffic.

You can use the following annotations in your Kubernetes ingress resource definition to load balance the UDP-based ingress traffic:

- ingress.citrix.com/insecure-service-type: This annotation enables L4 load balancing with UDP or ANY as a protocol for NetScaler.
- ingress.citrix.com/insecure-port: This annotation configures the port for UDP traffic. It is helpful when micro service access is required on a non-standard port. By default, port 80 is configured.

For more information about annotations, see annotations.

You can also use the standard Kubernetes solution of creating a service of type LoadBalancer with NetScaler. You can find out more about Service Type LoadBalancer in NetScaler.

Sample: Ingress definition for UDP-based ingress.

```
1 kubectl apply -f - <<EOF
2 apiVersion: networking.k8s.io/v1
3 kind: Ingress
4 metadata:
5 annotations:
6 ingress.citrix.com/insecure-port: "5084"
7 ingress.citrix.com/insecure-service-type: "udp"
8 name: udp-ingress
9 spec:
10 defaultBackend:
```

Sample: Service definition where the service port name is defined as udp-53:

```
1 kubectl apply -f - <<EOF</pre>
2 apiVersion: v1
3 kind: Service
4 metadata:
5
   name: bind
6 labels:
7
     app: bind
8 spec:
9 ports:
10
    - name: udp-53
     port: 53
    targetPort: 53
protocol: UDP
13
14 selector:
15 name: bind
16 EOF
```

# Monitor and improve the performance of your UDP-based applications

Application developers can closely monitor the health of UDP-based applications through rich monitors (such as UDP-ECV) in NetScaler. The ECV (extended content validation) monitors help in checking whether the application returns expected content or not. NetScaler Ingress Controller provides ingress.citrix.com/monitor annotation that can be used to monitor the health of the backend service.

Also, the application performance can be improved by using persistence methods such as Source IP. You can use these NetScaler features through Smart Annotations in Kubernetes.

The following ingress resource example uses smart annotations:

```
1 kubectl apply -f - <<EOF
2 apiVersion: networking.k8s.io/v1
3 kind: Ingress
4 metadata:
5 annotations:
6 ingress.citrix.com/frontend-ip: "192.168.1.1"
7 ingress.citrix.com/insecure-port: "80"
8 ingress.citrix.com/lbvserver: '{
9 "mongodb-svc":{
10 "lbmethod":"SRCIPDESTIPHASH" }
```

```
NetScaler ingress controller
```

```
11 }
12
       ingress.citrix.com/monitor: '{
13
   "mongodbsvc":{
14
    "type":"UDP-ECV" }
15
16
    }
   1
17
   name: mongodb
18
19 spec:
    rules:
    - host: mongodb.beverages.com
22
      http:
23
         paths:
24
         - backend:
25
             service:
26
               name: mongodb-svc
               port:
27
28
                number: 80
29
           path: /
           pathType: Prefix
31 EOF
```

# How to expose non-standard HTTP ports in the NetScaler CPX service

Sometimes you need to expose ports other than 80 and 443 in a NetScaler CPX service for allowing UDP traffic on other ports.

This section provides information on how to expose other non-standard HTTP ports on the NetScaler CPX service when you deploy it in the Kubernetes cluster.

# For Helm chart deployments

To expose non-standard HTTP ports while deploying NetScaler CPX with ingress controller using Helm charts, see the Helm chart installation guide.

# For deployments using the OpenShift operator

For deployments using the OpenShift operator, you need to edit the YAML definition to create CPX with ingress controller as specified in the step 6 of Deploy the NetScaler Ingress Controller as a sidecar with NetScaler CPX using NetScaler Operator and specify the ports as shown in the following example:

```
1 servicePorts:
2 - port: 80
3 protocol: UDP
4 name: http
5 - port: 443
6 protocol: UDP
```

7	name: https
8	- port: 6379
9	protocol: UDF
10	name: UDP

# How to set up dual-tier deployment

December 31, 2023

In a dual-tier deployment, NetScaler VPX or MPX is deployed outside the Kubernetes cluster (Tier-1) and NetScaler CPXs are deployed inside the Kubernetes cluster (Tier-2).

NetScaler MPX or VPX devices in Tier-1 proxy the traffic (North-South) from the client to NetScaler CPXs in Tier-2. The Tier-2 NetScaler CPX then routes the traffic to the microservices in the Kubernetes cluster. The NetScaler Ingress Controller deployed as a standalone pod configures the Tier-1 NetScaler. And, the sidecar NetScaler Ingress Controller in one or more NetScaler CPX pods configures the associated NetScaler CPX in the same pod.

The Dual-Tier deployment can be set up on Kubernetes in bare metal environment or on public clouds such as, AWS, GCP, or Azure.

The following diagram shows a Dual-Tier deployment:



#### Setup process

The NetScaler Ingress Controller repo provides a sample Apache microservice and manifests for NetScaler CPX for Tier-2, ingress object for Tier-2 NetScaler CPX, NetScaler Ingress Controller, and an ingress object for Tier-1 NetScaler for demonstration purpose. These samples are used in the setup process to deploy a dual-tier topology.

Perform the following:

- 1. Create a Kubernetes cluster in cloud or on-premises. The Kubernetes cluster in cloud can be a managed Kubernetes (for example: GKE, EKS, or AKS) or a custom created Kubernetes deployment.
- 2. Deploy NetScaler MPX or VPX on a multi-NIC deployment mode outside the Kubernetes cluster.
  - For instructions to deploy NetScaler MPX, see NetScaler documentation.
  - For instructions to deploy NetScaler VPX, see Deploy a NetScaler VPX instance.

Perform the following after you deploy NetScaler VPX or MPX:

a) Configure an IP address from the subnet of the Kubernetes cluster as SNIP on the NetScaler. For information on configuring SNIPs in NetScaler, see Configuring Subnet IP Addresses

#### (SNIPs).

b) Enable management access for the SNIP that is the same subnet of the Kubernetes cluster. The SNIP should be used as NS\_IP variable in the NetScaler Ingress Controller YAML file to enable NetScaler Ingress Controller to configure the Tier-1 NetScaler.

#### Note:

It is not mandatory to use SNIP as NS\_IP. If the management IP address of the NetScaler is reachable from NetScaler Ingress Controller then you can use the management IP address as NS\_IP.

- c) In cloud deployments, enable MAC-Based Forwarding mode on the Tier-1 NetScaler VPX. As NetScaler VPX is deployed in multi-NIC mode, it would not have the return route to reach the POD CNI network or the Client network. Hence, you need to enable MAC-Based Forwarding mode on the Tier-1 NetScaler VPX to handle this scenario.
- d) Create a NetScaler system user account specific to NetScaler Ingress Controller. NetScaler Ingress Controller uses the system user account to automatically configure the Tier-1 NetScaler.
- e) Configure your on-premises firewall or security groups on your cloud to allow inbound traffic to the ports required for NetScaler. The Setup process uses port 80 and port 443, you can modify these ports based on your requirement.
- 3. Deploy a sample microservice. Use the following command:

```
1 kubectl create -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/deployment/dual-tier/manifest/
apache.yaml
```

4. Deploy NetScaler CPX as Tier-2 ingress. Use the following command:

```
1 kubectl create -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/deployment/dual-tier/manifest/
tier-2-cpx.yaml
```

5. Create an ingress object for the Tier-2 NetScaler CPX. Use the following command:

```
kubectl create -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/deployment/dual-tier/manifest/
ingress-tier-2-cpx.yaml
```

- 6. Deploy the NetScaler Ingress Controller for Tier-1 NetScaler. Perform the following:
  - a) Download the NetScaler Ingress Controller manifest file. Use the following command:

```
1 wget https://raw.githubusercontent.com/citrix/citrix-k8s-
ingress-controller/master/deployment/dual-tier/manifest/
tier-1-vpx-cic.yaml
```

Environment Variable	Mandatory or Optional	Description
NS_IP	Mandatory	The IP address of the NetScaler appliance. For more details, see Prerequisites.
NS_USER and NS_PASSWORD	Mandatory	The user name and password of the NetScaler VPX or MPX appliance used as the Ingress device. For more details, see Prerequisites.
EULA	Mandatory	The End User License Agreement. Specify the value as Yes.
LOGLEVEL	Optional	The log levels to control the logs generated by NetScaler Ingress Controller. By default, the value is set to DEBUG. The supported values are: CRITICAL, ERROR, WARNING, INFO, and DEBUG. For more information, see Log Levels
NS_PROTOCOL and NS_PORT	Optional	Defines the protocol and port that must be used by NetScaler Ingress Controller to communicate with NetScaler. By default, NetScaler Ingress Controller uses HTTPS on port 443. You can also use HTTP on port 80.

b) Edit the NetScaler Ingress Controller manifest file and enter the values for the following environmental variables:

Environment Variable	Mandatory or Optional	Description
ingress-classes	Optional	If multiple ingress load
		balancers are used to load
		balance different ingress
		resources. You can use this
		environment variable to specify
		NetScaler Ingress Controller to
		configure NetScaler associated
		with specific ingress class. For
		information on Ingress classes,
		see Ingress class support
NS_VIP	Optional	NetScaler Ingress Controller
		uses the IP address provided in
		this environment variable to
		configure a virtual IP address to
		the NetScaler that receives
		Ingress traffic.

c) Deploy the updated NetScaler Ingress Controller manifest file. Use the following command:

```
1 kubectl create -f tier-1-vpx-cic.yaml
```

7. Create an ingress object for the Tier-1 NetScaler. Use the following command:

```
1 kubectl create -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/deployment/dual-tier/manifest/
ingress-tier-1-vpx.yaml
```

8. Update DNS server details in the cloud or on-premises to point your website to the VIP of the Tier-1 NetScaler.

```
For example: citrix-ingress.com 192.250.9.1
```

Where 192.250.9.1 is the VIP of the Tier-1 NetScaler and citrix-ingress.com is the microservice running in your Kubernetes cluster.

9. Access the URL of the microservice to verify the deployment.

# Set up dual-tier deployment using one step deployment manifest file

For easy deployment, the NetScaler Ingress Controller repo includes an all-in-one deployment manifest. You can download the file and update it with values for the following environmental variables and deploy the manifest file.

Note:

Ensure that you have completed step 1–2 in the Setup process.

#### Perform the following:

1. Download the all-in-one deployment manifest file. Use the following command:

```
1 wget https://raw.githubusercontent.com/citrix/citrix-k8s-ingress-
controller/master/deployment/dual-tier/manifest/all-in-one-dual
-tier-demo.yaml
```

2. Edit the all-in-one deployment manifest file and enter the values for the following environmental variables:

Environment Variable	Mandatory or Optional	Description
NS_IP	Mandatory	The IP address of the NetScaler appliance. For more details, see Prerequisites.
NS_USER and NS_PASSWORD	Mandatory	The user name and password of the NetScaler VPX or MPX appliance used as the Ingress device. For more details, see Prerequisites.
EULA	Mandatory	The End User License Agreement. Specify the value as Yes.
LOGLEVEL	Optional	The log levels to control the logs generated by NetScaler Ingress Controller. By default, the value is set to DEBUG. The supported values are: CRITICAL, ERROR, WARNING, INFO, and DEBUG. For more information, see Log Levels

#### NetScaler ingress controller

Environment Variable	Mandatory or Optional	Description
NS_PROTOCOL and NS_PORT	Optional	Defines the protocol and port that must be used by NetScaler Ingress Controller to communicate with NetScaler. By default, NetScaler Ingress Controller uses HTTPS on port 443. You can also use HTTP on port 80.
ingress-classes	Optional	If multiple ingress load balancers are used to load balance different ingress resources. You can use this environment variable to specify NetScaler Ingress Controller to configure NetScaler associated with specific ingress class. For information on Ingress classes, see [Ingress class support](/en- us/netscaler-k8s-ingress- controller/configure/ingress-
NS_VIP	Optional	classes.html NetScaler Ingress Controller uses the IP address provided in this environment variable to configure a virtual IP address to the NetScaler that receives Ingress traffic. <b>Note:</b> NS_VIP acts as a fallback when the frontend-ip annotation is not provided in Ingress yaml. Not supported for Type Loadbalancer service.

# 3. Deploy the updated all-in-one deployment manifest file. Use the following command:

1 kubectl create -f all-in-one-dual-tier-demo.yaml

# Horizontal pod autoscaler for NetScaler CPX with custom metrics

## December 31, 2023

While deploying workloads in a Kubernetes cluster for the first time, it is difficult to exactly predict the resource requirements and how those requirements might change in a production environment. Using Horizontal pod autoscaler (HPA), you can automatically scale the number of pods in your workload based on different metrics like actual resource usage. HPA is a resource provided by Kubernetes which scales Kubernetes based resources like deployments, replicasets, and replication controllers.

Traditionally, HPA gets the required metrics from a metrics server. It then periodically adjusts the number of replicas in a deployment to match the observed average metrics to the target you specify.



NetScaler provides a custom-metric based HPA solution for NetScaler CPX.

By default, the metrics server only gives CPU and memory metrics for a pod.

NetScaler provides a rich set of in-built metrics for analyzing application performance and based on these metrics you can take a better autoscaling judgment. A custom metric based HPA is a better solution like autoscaling based on HTTP request rate, SSL transactions, or ADC bandwidth.

# **NetScaler CPX HPA solution**

NetScaler CPX HPA solution consists of the following components:

• NetScaler VPX: NetScaler VPX or MPX is deployed at Tier-1 and load balances the client requests among the NetScaler CPX pods inside the cluster.

- NetScaler CPX: NetScaler CPX deployed inside the cluster acts as a Tier-2 load balancer for the endpoint application pods. The NetScaler CPX pod is running along with the NetScaler Ingress Controller and NetScaler metric exporter as sidecars.
- NetScaler Ingress Controller: The NetScaler Ingress Controller is an ingress controller which is built around the Kubernetes Ingress and automatically configures NetScaler based on the Ingress resource configuration. The NetScaler Ingress Controller deployed as a stand-alone pod configures the NetScaler VPX and other instances configures NetScaler CPXs.
- NetScaler Metrics Exporter: The NetScaler Metrics Exporter exports the application performance metrics to the open-source monitoring system Prometheus. The NetScaler Metrics Exporter collects metrics from NetScaler CPX and exposes it in a format that Prometheus can understand.
- Prometheus: Prometheus is an open-source systems monitoring and alerting toolkit. Prometheus is used to collect metrics from NetScaler CPXs and expose them using a Prometheus adapter which is queried by the HPA controller to keep a check on metrics.
- Prometheus adapter: Prometheus adapter contains an implementation of the Kubernetes resource metrics API and custom metrics API. This adapter is suitable for use with the autoscaling/v2 HPA in Kubernetes version 1.6+. It can also replace the metrics server on clusters that already run Prometheus and collect the appropriate metrics.

The following diagram is a visual representation of how the NetScaler CPX HPA solution works.



The Tier-1 NetScaler VPX load balances the NetScaler CPXs at Tier-2. NetScaler CPXs load balance applications. Other components like Prometheus, Prometheus-adapter, and an HPA controller is also deployed.

The HPA controller keeps polling the Prometheus-adapter for custom metrics like HTTP requests rate or bandwidth. Whenever the limit defined by the user in the HPA is reached, the HPA controller scales the NetScaler CPX deployment and creates another NetScaler CPX pod to handle the load.

## **Deploy NetScaler CPX HPA solution**

Perform the following steps to deploy the NetScaler CPX HPA solution.

1. Clone the citrix-k8s-ingress-controller repository from GitHub using the following command.

1 git clone https://github.com/citrix/citrix-k8s-ingress-controller.
 git

After cloning, change your directory to the HPA folder with the following command.

1 cd citrix-k8s-ingress-controller/blob/master/docs/how-to/hpa

- 2. From the HPA directory, open and edit the values.sh file and set the following values for NetScaler VPX.
  - VPX\_IP: IP address of the NetScaler VPX
  - VPX\_PASSWORD: The password of the nsroot user on the NetScaler VPX
  - VIRTUAL\_IP\_VPX: The IP address on which the sample guesbook application is accessed.
- 3. Create all the required resources by running the create\_all.sh file.

1 ./create\_all.sh

This step creates the following resources:

- Prometheus and Grafana for monitoring
- NetScaler CPX with the NetScaler Ingress Controller and metrics exporter as sidecars
- NetScaler Ingress Controller as a stand-alone pod to configure NetScaler VPX
- A sample guestbook application
- HPA controller for monitoring the NetScaler CPX autoscale deployment
- Prometheus adapter for exposing the custom metrics
- Add an entry in the hosts file. The route must be added in the hosts file to route traffic for the guestbook application to the NetScaler VPX virtual IP address.
   For most Linux distros, the hosts file is present in the /etc folder.
- 5. Send some generated traffic and verify the NetScaler CPX autoscale deployment.

The NetScaler CPX deployment HPA has been configured in such a way that when the average HTTP requests rate of the NetScaler CPX goes above 20 requests per second, it autoscales. You can use the following scripts provided in the HPA folder for sending traffic:

- 16\_curl.sh Send 16 HTTP requests per second (lesser than the threshold)
- 30\_curl.sh Send 30 HTTP requests per second (greater than the threshold)

a. Run the 16\_curl.sh script to send 16 HTTP requests per second to the NetScaler CPX.

```
1 ./16_curl.sh
```

The following diagram a Grafana dashboard which displays HTTP requests per second.



The following output shows the HPA state with 16 HTTP RPS.						
NAME	REFERENCE	TARGETS	MINPODS	MAXPODS	REPLICAS	AGE
cpx-builtin	Deployment/cpx-builtin	16/20	1	3	1	4m9s

b. Run the 30\_curl.sh script to send 30 HTTP requests per second to NetScaler CPX.

1 ./30\_curl.sh

When you run this script, the threshold of 20 requests that was set has been crossed and the NetScaler CPX deployment autoscales from one pod to two pods. The average value of the metric HTTP request rate also goes down from 30 to 15 as there are two NetScaler CPX pods.

The following output shows the state of HPA when the target is crossed.

NAME	REFERENCE	TARGETS	MINPODS	MAXPODS	REPLICAS	AGE
cpx-builtin	Deployment/cpx-builtin	30/20	1	3	1	8m4s

The following output shows that the number of replicas of NetScaler CPX have gone up to 2 and the average value of HTTP RPS comes down to 15.

NAME	REFERENCE	TARGETS	MINPODS	MAXPODS	REPLICAS	AGE
cpx-builtin	Deployment/cpx-builtin	15/20	1	3	2	9m24s

The following diagram shows a Grafana dashboard with two NetScaler CPXs load balancing the traffic.



6. Clean up by executing the delete\_all.sh script.

1 ./delete\_all.sh

#### Note:

If the Tier-1 NetScaler VPX is not present, use NodePort to expose the NetScaler CPX service.

# **Deploy Direct Server Return**

December 31, 2023

In a typical load-balanced system, a load balancer acts as a mediator between web servers and clients. Incoming client requests are received by the load balancer and it passes the requests to the appropriate server with slight modifications to the data packets. The server responds to the load balancer with the required data and then the load balancer forwards the response to the client.

In a Direct Server Return (DSR) deployment, load balancer forwards the client request to the server, but the back-end server directly sends the response to the client. The use of different network paths for request and response helps to avoid extra hops and reduces the latency. Because the server directly responds to the client, DSR speeds up the response time between the client and the server and also removes some extra load from the load balancer. Using DSR is a transparent way to achieve increased network performance for your applications with little to no infrastructure changes. For more information on DSR using NetScaler, see the NetScaler documentation.

DSR solution is useful in the following situations:

- While handling applications that deliver video streaming where low latency (response time) matters.
- Where intelligent load balancing is not required
- When the output capacity of the load-balancer can be the bottleneck
However, when you use the DSR advanced layer 7 load balancing features are not supported.

#### DSR network topology for Kubernetes using NetScaler

In this topology, there is an external load-balancer (Tier-1 ADC) that distributes the traffic to the ingress ADC (Tier 2 ADC) deployed inside the Kubernetes cluster over an overlay (L3 DSR IPIP). Tier-2 ADC picks up the packet, decapsulate the packet, and performs load balancing among services. The Tier-2 ADC sends the return traffic from service to the client instead of sending it via Tier-1 ADC.



#### Deploying DSR for cloud native applications using NetScaler

Perform the steps in the following sections to deploy DSR for applications deployed on the Kubernetes cluster.

#### Deploy NetScaler CPX as Tier-2 ADC

This section contains steps to create configurations required on the ingress device for DSR topology.

1. Create a namespace for DSR using the following command:

```
1 kubectl apply -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/example/dsr/KubernetesConfig/
dsr_namespace.yaml
```

2. Create a ConfigMap using the following command.

1 kubectl apply -f https://raw.githubusercontent.com/citrix/citrixk8s-ingress-controller/master/example/dsr/KubernetesConfig/ cpx\_config.yaml -n dsr

#### Note:

In this example, the node controller network is configured as 192.168.1.0/24. Hence, the command to create IP tunnel is provided as add iptunnel dsr 192.168.1.254 255.255.255.0 \*. You need to specify the value according to your CNC configuration.

3. Deploy NetScaler CPX on the namespace dsr.

```
1 kubectl apply -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/example/dsr/KubernetesConfig/
citrix-k8s-cpx-ingress.yml -n dsr
```

#### Deploying a sample application on the Kubernetes cluster

Perform the steps in this section to deploy a sample application on Kubernetes cluster.

1. Deploy the guestbook application using the following command.

```
1 kubectl apply -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/example/dsr/KubernetesConfig/
guestbook-all-in-one.yaml -n dsr
```

- 2. Expose the guestbook application using Ingress.
  - a) Download the guestbook ingress YAML file using the following command.

```
1 wget https://raw.githubusercontent.com/citrix/citrix-k8s-
ingress-controller/master/example/dsr/KubernetesConfig/
guestbook-all-in-one.yaml
```

b) Edit and provide the DSR IP or public IP address through which you access your application using the ingress.citrix.com/frontend-ip: annotations.

1 ingress.citrix.com/frontend-ip: "<ip-address>"

c) Save the YAML file and deploy the Ingress resource using the following command.

1 kubectl apply -f guestbook-ingress.yaml -n dsr

#### Establish network connectivity between Tier-1 and Tier-2 ADCs

Perform the steps in this section to establish network connectivity between Tier-1 and Tier-2 ADCs.

1. Download the YAML to deploy node controller using the following command.

1 wget https://raw.githubusercontent.com/citrix/citrix-k8s-nodecontroller/master/deploy/citrix-k8s-node-controller.yaml

- 2. Edit the YAML file and provide the values for NS\_IP, NS\_USER, NS\_PASSWORD, and RE-MOTE\_VTEPIP arguments. For detailed information, see node controller.
- 3. Save the YAML file and deploy the node controller.

```
1 kubectl create -f citrix-k8s-node-controller.yaml -n dsr
```

#### Deploy the NetScaler Ingress Controller for Tier-1 ADC and expose NetScaler CPX as a service

Perform the following steps to deploy the NetScaler Ingress Controller as a stand-alone pod and create an Ingress resource for Tier-2 NetScaler CPX.

1. Download the NetScaler Ingress Controller YAML file using the following command.

```
1 wget https://raw.githubusercontent.com/citrix/citrix-k8s-
ingress-controller/master/example/dsr/KubernetesConfig/
citrix-k8s-ingress-controller.yaml
```

- 2. Edit the YAML file and update the following values for NetScaler Ingress Controller.
  - NS\_IP
  - NS\_USER
  - NS\_PASSWORD

For more information, see Deploy the NetScaler Ingress Controller using YAML.

3. Save the YAML file and deploy the NetScaler Ingress Controller.

```
1 kubectl create -f citrix-k8s-ingress-controller.yaml -n dsr
```

4. Create DSR configuration on Tier-1 ADC by creating an ingress resource for the Tier-2 NetScaler CPX.

```
wget https://raw.githubusercontent.com/citrix/citrix-k8s-ingress
-controller/master/example/dsr/KubernetesConfig/vpx-ingress.
yaml
```

5. Edit the YAML file and provide the DSR or public IP address through which user access your application using the ingress.citrix.com/frontend-ip: annotation. This IP address must be same as the IP address you have specified in step 2.

```
1 kubectl apply -f vpx-ingress.yaml -n dsr
```

#### Test the DSR deployment

To test the DSR deployment, access the application from a browser using the IP address specified for the ingress.citrix.com/frontend-ip: annotation. A guestbook page is populated.

A sample output is given as follows:

## Guestbook

Message	es		
Submit			

#### Troubleshooting

When you test the application, it might not populate any pages even though all the required configurations are created. This is because of rp\_filter rules on the host. If you experience such an issue, use the following commands on all the hosts to disable the rules.

```
sysctl -w net.ipv4.conf.all.rp_filter=0
sysctl -w net.ipv4.conf.cni0.rp_filter=0
sysctl -w net.ipv4.conf.eth0.rp_filter=0
sysctl -w net.ipv4.conf.cni0.rp_filter=0
sysctl -w net.ipv4.conf.default.rp_filter=0
```

## Support for admission controller webhooks

#### December 31, 2023

Admission controllers are powerful tools for intercepting requests to the Kubernetes API server prior to the persistence of the object. Using Kubernetes admission controllers, you can define and customize what is allowed to run on your cluster. Hence, they are useful tools for cluster administrators to deploy preventive security controls on your cluster. But you need to compile the admission controllers into the kube-apiserver binary and they offer limited flexibility.

To overcome this limitation, Kubernetes supports dynamic admission controllers that can be developed as extensions and run as webhooks configured at runtime.

Using the Admission controller webhooks Kubernetes cluster administrators can create additional plug-ins to the admission chain of API server without recompiling them. Admission controller webhooks can be executed whenever a resource is created, updated, or deleted.

You can define two types of admission controller webhooks:

- validating admission webhook
- mutating admission webhook

Mutating admission webhooks are invoked first, and they can modify objects sent to the API server to enforce custom defaults. Once all the object modifications are complete, and the incoming object is validated by the API server, validating admission webhooks are invoked. Validating admission hooks process requests and accept or reject requests to enforce custom policies.



The following diagram explains how the admission controller webhook works:

Here are some of the scenarios where admission webhooks are useful:

- To mandate a reasonable security baseline across an entire namespace or cluster mandating. For example, disallowing containers from running as root or making sure the container's root filesystem is always mounted as read-only.
- To enforce the adherence to certain standard and practices for labels, annotations, or resource limits. For example, enforce label validation on different objects to ensure proper labels are being used for various objects.
- To validate the configuration of the objects running in the cluster and prevent any obvious misconfigurations from hitting your cluster.

For example, to detect and fix images deployed without semantic tags.

## How to apply admission controllers

Writing an admission controller for each specific use case is not scalable and it helps to have a system that that supports multiple configurations covering different resource types and fields. You can

use Open policy agent (OPA) and Gatekeeper to implement a customizable admission webhook for Kubernetes.

OPA is an open source, general-purpose policy engine that unifies policy enforcement across the stack. Gatekeeper is a customizable validating webhook that enforces CRD-based policies executed by OPA.



#### credit)

Gatekeeper introduces the following functionalities

- An extensible, parameterized policy library
- Native Kubernetes CRDs for instantiating the policy library (constraints)
- Native Kubernetes CRDs for extending the policy library (constraint templates)
- Audit functionality

#### Writing and deploying an admission controller webhook

#### Prerequisites

- Kubernetes 1.14.0 or later with the admissionregistration.k8s.io/v1beta1 API enabled. You can verify whether the API is enabled by using the following command:
  - 1 kubectl api-versions | grep admissionregistration.k8s.io/v1beta1

The following output indicates that the API is enabled:

1 admissionregistration.k8s.io/v1beta1

• The mutating admission webhook and validate admission webhook admission controllers should be added and listed in the correct order in the admission-control flag of kube-apiserver.

With Minikube, you can perform this task by starting Minikube with the following command:

```
1 minikube start --extra-config=apiserver.enable-admission-plugins=
NamespaceLifecycle,LimitRanger,ServiceAccount,
DefaultStorageClass,DefaultTolerationSeconds,NodeRestriction,
MutatingAdmissionWebhook,ValidatingAdmissionWebhook`
```

• Ensure that you have cluster administrator permissions.

```
1 kubectl create clusterrolebinding cluster-admin-binding --
clusterrole cluster-admin --user <YOUR USER NAME>
```

#### Mutating admission webhook configuration

For more information on mutating admission webhook configuration, see ingress-admission-webhook.

The following use cases are covered in the mutating admission webhook example:

- Update port in an Ingress based on the Ingress name
- Enable secure back-end forcefully based on a namespace

#### Validating admission webhook configuration using Gatekeeper

Gatekeeper uses a CRD that allows you to create constraints as Kubernetes resources. This CRD is called a ConstraintTemplate in Gatekeeper. The schema of the constraint allows an administrator to fine-tune the behavior of a constraint, similar to arguments to a function. Constraints are used to inform Gatekeeper that the administrator wants a constraint template to be enforced, and how.

You can apply various policies using constraint templates. Various examples are listed at the Gatekeeper library.

#### Deploying a sample policy

Perform the following steps to deploy HttpsOnly as a sample policy using Gatekeeper. The HttpsOnly policy allows only an Ingress configuration with HTTPS.

1. Install Gatekeeper using the following command.

#### Note:

In this step, Gatekeeper is installed using a prebuilt image. You can install Gatekeeper using various methods mentioned in the Gatekeeper installation.

1 # kubectl apply -f https://raw.githubusercontent.com/open-policyagent/gatekeeper/master/deploy/gatekeeper.yaml

You can verify the installation using the following command.

```
1 kubectl get crd | grep -i constraintsonstrainttemplates.templates.
gatekeeper.sh
```

You can check all the constraint templates using the following command:

kubectl get constrainttemplates.templates.gatekeeper.sh

2. Apply the httpsonly constraint template.

```
1 kubectl apply -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/docs/how-to/webhook/httpsonly/
template.yaml
```

3. Apply a constraint to enforce the httpsonly policy.

```
1 kubectl apply -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/docs/how-to/webhook/httpsonly/
constraint.yaml
```

4. Deploy a sample Ingress which violates the policy to verify the policy. It should display an error while creating the Ingress.

```
1 kubectl apply -f https://raw.githubusercontent.com/citrix/citrix-
k8s-ingress-controller/master/docs/how-to/webhook/httpsonly/bad
-example-ingress.yaml
2
3 Error from server ([denied by ingress-https-only] Ingress must be
https. tls configuration is required for test-ingress): error
when creating "ingress.yaml": admission webhook "validation.
gatekeeper.sh" denied the request: [denied by ingress-https-
only] Ingress must be https. tls configuration is required for
test-ingress
```

5. Now, deploy an Ingress which has the required TLS section in Ingress.

```
1 # kubectl apply -f https://raw.githubusercontent.com/citrix/
citrix-k8s-ingress-controller/master/docs/how-to/webhook/
httpsonly/good-example-ingress.yaml
2
3 ingress.networking.k8s.io/test-ingress created
```

- 6. Clean up the installation using the following commands once you have finished the verification of Gatekeeper policies.
  - Uninstall all packages and template installed.
     kubectl delete -f https://raw.githubusercontent.com/citrix/citrixk8s-ingress-controller/master/docs/how-to/webhook/httpsonly/ good-example-ingress.yaml
     kubectl delete -f https://raw.githubusercontent.com/citrix/citrixk8s-ingress-controller/master/docs/how-to/webhook/httpsonly/ constraint.yaml
     kubectl delete -f https://raw.githubusercontent.com/citrix/citrixk8s-ingress-controller/master/docs/how-to/webhook/httpsonly/ constraint.yaml
     kubectl delete -f https://raw.githubusercontent.com/citrix/citrixk8s-ingress-controller/master/docs/how-to/webhook/httpsonly/ template.yaml
     kubectl delete -f https://raw.githubusercontent.com/open-policyagent/gatekeeper/master/deploy/gatekeeper.yaml

#### More sample use cases

There are multiple use cases listed under the webhook directory.

The steps are similar to what is specified in the example and can be summarized as follows:

- 1. Apply the template YAML file given in each use case directory.
- 2. Apply the constraint YAML file.
- 3. Verify by applying bad or good sample YAML files to validate the use case.

For further use cases, see the Gatekeeper library.

## Enable gRPC support using the NetScaler Ingress Controller

#### June 5, 2025

gRPC is a high performance, open-source universal RPC framework created by Google. In gRPC, a client application can directly call methods on a server application from a different server in the same way you call local methods.

You can easily create distributed applications and services using GRPC.

#### Enable gRPC support

Perform the following steps to enable GRPC support using HTTP2.

1. Create a YAML file cic-configmap.yaml and enable the global parameter for HTTP2 server side support using the following entry in the ConfigMap. For more information on using ConfigMap, see the ConfigMap documentation.

1 NS\_HTTP2\_SERVER\_SIDE: 'ON'

2. Apply the ConfigMap using the following command.

1 kubectl apply -f cic-configmap.yaml

3. Edit the cic.yaml file for deploying the NetScaler Ingress Controller to support ConfigMap.

```
1 args:

2 --ingress-classes

3 citrix

4 --configmap

5 default/cic-configmap
```

4. Deploy the NetScaler Ingress Controller as a stand-alone pod by applying the edited YAML file.

```
1 kubectl apply -f cic.yaml
```

5. To test the gRPC traffic, you may need to install grpcurl. Perform the following steps to install grpcurl on a Linux machine.

```
1 go get github.com/fullstorydev/grpcurl
2 go install github.com/fullstorydev/grpcurl/cmd/grpcurl
```

6. Apply the gRPC test service YAML file (grpc-service.yaml).

```
1 kubectl apply -f grpc-service.yaml
```

Following is a sample content for the grpc-service.yaml file.

```
1
     apiVersion: apps/v1
2
     kind: Deployment
3
    metadata:
4
     name: grpc-service
5
    spec:
6
       replicas: 1
7
       selector:
8
         matchLabels:
9
           app: grpc-service
      template:
11
         metadata:
12
           labels:
13
             app: grpc-service
14
         spec:
15
           containers:
           - image: registry.cn-hangzhou.aliyuncs.com/acs-sample/grpc
16
               -server:latest
17
             imagePullPolicy: Always
18
             name: grpc-service
19
             ports:
             - containerPort: 50051
20
               protocol: TCP
```

```
restartPolicy: Always
23
24 apiVersion: v1
25
   kind: Service
26
   metadata:
27
    name: grpc-service
   spec:
28
29
      ports:
       - port: 50051
31
        protocol: TCP
32
        targetPort: 50051
      selector:
34
        app: grpc-service
35
       sessionAffinity: None
       type: NodePort
```

7. Create a certificate for the gRPC Ingress configuration.

```
1 openssl req -x509 -nodes -days 365 -newkey rsa:2048 -keyout tls.
key -out tls.crt -subj "/CN=grpc.example.com/0=grpc.example.com
"
2
3 kubectl create secret tls grpc-secret --key tls.key --cert tls.crt
4
5 secret "grpc-secret" created
```

- 8. Enable HTTP2 using Ingress annotations. See HTTP/2 support for steps to enable HTTP2 using the NetScaler Ingress Controller.
  - Create a YAML file for the front-end Ingress configuration and apply it to enable HTTP/2 on the content switching virtual server.

kubectl apply -f frontend-ingress.yaml

The content of the frontend-ingress.yaml file is provided as follows:

```
1
     apiVersion: networking.k8s.io/v1
2
     kind: Ingress
3
    metadata:
4
       annotations:
5
         ingress.citrix.com/frontend-httpprofile: '{
6
    "http2":"enabled", "http2direct" : "enabled" }
7
8
         ingress.citrix.com/frontend-ip: 192.0.2.1
9
         ingress.citrix.com/secure-port: "443"
       name: frontend-ingress
11
     spec:
       ingressClassName: citrix
12
13
       rules:
14
       - {
     }
15
16
17
       tls:
```

18 - { 19 }

• Create a YAML file for the back-end Ingress configuration with the following content and apply it to enable HTTP2 on back-end (service group).

kubectl apply -f backend-ingress.yaml

The content of the backend-ingress.yaml file is provided as follows:

```
apiVersion: networking.k8s.io/v1
1
     kind: Ingress
2
3
     metadata:
       annotations:
4
5
         ingress.citrix.com/backend-httpprofile: '{
    "grpc-service":{
6
    "http2": "enabled", "http2direct" : "enabled" }
7
8
     }
    i.
9
10
         ingress.citrix.com/frontend-ip: 192.0.2.2
         ingress.citrix.com/secure-port: "443"
11
12
      name: grpc-ingress
13
    spec:
14
       ingressClassName: citrix
15
       rules:
       - host: grpc.example.com
16
17
         http:
           paths:
18
19
           - backend:
20
               service:
21
                  name: grpc-service
22
                  port:
23
                    number: 50051
24
              path: /
              pathType: Prefix
26
       tls:
27
       - hosts:
28
         - grpc.example.com
29
         secretName: grpc-secret
```

9. Test the gRPC traffic using the grpcurl command.

```
1 grpcurl -v -insecure -d '{
2 "name": "gRPC" }
3 ' grpc.example.com:443 helloworld.Greeter.SayHello
```

The output of the command is shown as follows:

```
1 Resolved method descriptor:
2 rpc SayHello ( .helloworld.HelloRequest ) returns ( .helloworld.
HelloReply );
3
4
```

```
5 Request metadata to send:
6 (empty)
7
8
9 Response headers received:
10 content-type: application/grpc
11
12
13 Response contents:
14 {
15
16
       "message": "Hello gRPC"
17
    }
18
19
20
21 Response trailers received:
22 (empty)
23 Sent 1 request and received 1 response
```

#### Validate the rate limit CRD

Perform the following steps to validate the rate limit CRD.

1. Apply the rate limit CRD using the ratelimit-crd.yaml file.

kubectl create -f ratelimit-crd.yaml

2. Create a YAML file (ratelimit-crd-object.yaml) with the following content for the rate limit policy.

```
apiVersion: citrix.com/v1beta1
1
2
    kind: ratelimit
3
   metadata:
4
    name: throttle-req-per-clientip
5
   spec:
   servicenames:
6
       - grpc-service
7
8 selector_keys:
9
     basic:
      path:
10
11
        - "/"
12
       per_client_ip: true
13
     req_threshold: 5
      timeslice: 60000
14
15
      throttle_action: "RESPOND"
```

3. Apply the YAML file using the following command.

1 kubectl create -f ratelimit-crd-object.yaml

4. Test gRPC traffic using the grpcurl command.

```
1 grpcurl -v -insecure -d '{
2 "name": "gRPC" }
3 ' grpc.example.com:443 helloworld.Greeter.SayHello
```

The command returns the following error in response after the rate limit is reached:

```
1 Error invoking method "helloworld.Greeter.SayHello": failed to
    query for service descriptor "helloworld.Greeter": rpc error:
    code = Unavailable desc =
2
3 Too Many Requests: HTTP status code 429; transport: missing
    content-type field
```

#### Validate the Rewrite and Responder CRD with gRPC

Perform the following steps to validate the Rewrite and Responder CRD.

1. Apply the Rewrite and Responder CRD using the rewrite-responder-policies-deployment.yaml file.

kubectl create -f rewrite-responder-policies-deployment.yaml

2. Create a YAML file (rewrite-crd-object.yaml) with the following content for the rewrite policy.

```
1 apiVersion: citrix.com/v1
2 kind: rewritepolicy
3 metadata:
4 name: addcustomheaders
5 spec:
   rewrite-policies:
6
7
      - servicenames:
8
           - grpc-service
         rewrite-policy:
9
          operation: insert_http_header
           target: 'sessionID'
11
           modify-expression: '"48592th42gl24456284536tgt2"'
12
13
           comment: 'insert SessionID in header'
14
           direction: RESPONSE
15
           rewrite-criteria: 'http.res.is_valid'
```

3. Apply the YAML file using the following command.

1 kubectl create -f rewrite-crd-object.yaml

4. Test the gRPC traffic using the grpcurl command.

```
1 grpcurl -v -insecure -d '{
2 "name": "gRPC" }
3 ' grpc.example.com:443 helloworld.Greeter.SayHello
```

This command adds a session id in the gRPC request response.

```
1 Resolved method descriptor:
2 rpc SayHello ( .helloworld.HelloRequest ) returns ( .helloworld.
      HelloReply );
3
4 Request metadata to send:
5 (empty)
6
7 Response headers received:
8 content-type: application/grpc
9 sessionid: 48592th42gl24456284536tgt2
10
11 Response contents:
12 {
13
14
     "message": "Hello gRPC"
15
    }
16
17
18 Response trailers received:
19 (empty)
20 Sent 1 request and received 1 response
```

## **Policy-based routing for multiple Kubernetes clusters**

#### May 6, 2024

When using a single NetScaler to load balance multiple Kubernetes clusters, NetScaler Ingress Controller adds pod CIDR networks in NetScaler through static routes. These routes establish network connectivity between Kubernetes pods and NetScaler. However, when the pod CIDRs overlap there may be route conflicts. NetScaler supports policy-based routing (PBR) to address the networking conflicts in such scenarios. In PBR, routing decisions are taken based on the criteria that you specify. Typically, a next hop is specified where NetScaler VPX or NetScaler MPX sends the selected packets. In a GSLB Kubernetes environment, PBR is implemented by reserving a subnet IP address (SNIP) for each Kubernetes cluster or the NetScaler Ingress Controller. Using net profile, the SNIP is bound to all service groups created by the same NetScaler Ingress Controller. For all the traffic generated from service groups belonging to the same cluster, the source IP address is the same SNIP.

In the following sample topology, PBR is configured for two Kubernetes clusters, which are load balanced using NetScaler MPX or NetScaler VPX.



#### **Configure PBR using NetScaler Ingress Controller**

To configure PBR, you need one SNIP or more per Kubernetes cluster. You can provide SNIP values using the nsSNIPS parameter in the NSIC Helm install command.

Use the following values.yml file in the helm install command to deploy NetScaler Ingress Controller and configure PBR.

```
1 nsIP: <NS_IP>
2 license:
3 accept: yes
4 adcCredentialSecret: <secret>
5 nodeWatch: True
6 clusterName: <name of cluster>
7 nsSNIPS: '["1.2.3.4", "5.6.7.8"]'
```

#### Validate PBR configuration on NetScaler after deploying NetScaler Ingress Controller

This validation example uses a two-node Kubernetes cluster with NetScaler Ingress Controller deployed along with the following ConfigMap with two SNIPs. NetScaler ingress controller



You can verify that NetScaler Ingress Controller adds the following configurations to NetScaler:

1. To verify whether the IPset of all NS\_SNIPs is added, run the following command: show ipset k8s-pbr\_ipset.

> show	ipset k8s-pbr_ipset
1)	Name: k8s-pbr_ipset
	IP:1.2.3.4
	IP:5.6.7.8
Done	

2. A net profile is added with the SrcIP set to the IPset.



3. To verify whether the service group added by NetScaler Ingress Controller contains the net profile set, run show servicegroup.



4. Run the following command to check whether the PBR is addded: shpbr.

```
shpbr
1)
        Name: k8s-1.2.3.4_10.106.175.86
        Action: ALLOW
                                                Hits: 0
        srcIP = (1.2.3.4
        destIP = 192.168.0.0-192.168.0.255
        srcMac:
        Protocol:
                                               Interface:
        Vlan:
        Active Status: ENABLED
                                                Applied Status: APPLIED
        Priority: 30
        NextHop: 10.106.175.86
2)
        Name: k8s-5.6.7.8_10.106.175.86
        Action: ALLOW
                                                Hits: 0
        srcIP = (5.6.7.8)
        destIP = 192.168.0.0-192.168.0.255
        srcMac:
        Protocol:
        Vlan:
                                               Interface:
        Active Status: ENABLED
                                                Applied Status: APPLIED
        Priority: 40
        NextHop: 10.106.175.86
3)
        Name: k8s-1.2.3.4_10.106.175.87
        Action: ALLOW
                                                Hits: 0
        srcIP = (1.2.3.4)
        destIP = 192.168.16.0-192.168.16.255
        srcMac:
        Protocol:
        Vlan:
                                               Interface:
        Active Status: ENABLED
                                                Applied Status: APPLIED
        Priority: 50
        NextHop: 10.106.175.87
4)
        Name: k8s-5.6.7.8_10.106.175.87
        Action: ALLOW
                                                Hits: 0
        srcIP = (5.6.7.8)
        destIP = 192.168.16.0-192.168.16.255
        srcMac:
        Protocol:
        Vlan:
                                               Interface:
        Active Status: ENABLED
                                                Applied Status: APPLIED
        Priority: 60
        NextHop: 10.106.175.87
```

Here:

- The number of PBRs is equivalent to (number of SNIPs) \*(number of Kubernetes nodes). In this case, it adds four (2\*2) PBRs.
- The srcIP of the PBR is the NS\_SNIPS provided to NetScaler Ingress Controller by ConfigMap. The destIP is the CNI overlay subnet range of the Kubernetes node.
- NextHop is the IP address of the Kubernetes node.
- 5. You can also use the logs of the NetScaler Ingress Controller to validate the configuration.



#### Configure PBR using the NetScaler node controller

You can configure PBR using the NetScaler node controller for multiple Kubernetes clusters. When you are using a single NetScaler to load balance multiple Kubernetes clusters with NetScaler node controller for networking, the static routes added by it to forward packets to the IP address of the VXLAN tunnel interface might cause route conflicts. To support PBR, NetScaler node controller needs to work in conjunction with the NetScaler Ingress Controller to bind the net profile to the service group.

Perform the following steps to configure PBR using the NetScaler node controller:

 While starting NetScaler node controller, provide the CLUSTER\_NAME as an environment variable. Specifying this variable indicates that it is a multi-cluster deployment and the NetScaler node controller configures PBR instead of static routes.

Example:

```
helm install nsnc netscaler/netscaler-node-controller --set
license.accept=yes,nsIP=<NSIP>,vtepIP=<NetScaler SNIP>,vxlan.
id=<VXLAN ID>,vxlan.port=<VXLAN PORT>,network=<IP-address-range
-for-VTEP-overlay>,adcCredentialSecret=<Secret-for-NetScaler-
credentials>,cniType=<CNI-overlay-name>,clusterName=<cluster-name
>
```

2. While deploying NetScaler Ingress Controller, provide the CLUSTER\_NAME as an environment variable. This value should be the same as the value provided in the node controller, and enable PBR using the argument nsncPbr=True to configure PBR on NetScaler.

#### Notes:

- The value provided for CLUSTER\_NAME in NetScaler node controller and NetScaler Ingress Controller deployment files should match when they are deployed in the same Kubernetes cluster.
- The CLUSTER\_NAME is used while creating the net profile entity and binding it to service groups on NetScaler VPX or NetScaler MPX.

#### Validate PBR configuration on NetScaler after deploying the NetScaler node controller

This validation example uses a two-node Kubernetes cluster with NetScaler node controller and NetScaler Ingress Controller deployed.

You can verify that the following configurations are added to the NetScaler by NetScaler node controller:

1. Run the following command: show netprofile.

A net profile is added with the value of srcIP set to the SNIP added by node controller while creating the VXLAN tunnel network between NetScaler and Kubernetes nodes.

> show	netprofile	
1)	Name: cnc-cluster1_netprof	
	SrcIP: 162.1.11.254	
	SrcIPPersistency: DISABLED	
	MBF: DISABLED	
	Proxy Protocol: DISABLED	
	Proxy Protocol version: V1	
Done		

2. Run the following command: shservicegroup.

NetScaler Ingress Controller binds the net profile to the service groups it creates.

```
shservicegroup
       k8s-frontend_80_sgp_5udfadwklqkfukj3rsuoao5b6wxauf5j - HTTP
       State: ENABLED Effective State: UP
                                               Monitor Threshold : 0
       Max Conn: 0
                       Max Req: 0
                                       Max Bandwidth: 0 kbits
       Use Source IP: NO
       Client Keepalive(CKA): NO
       Monitoring Owner: 0
       TCP Buffering(TCPB): NO
       HTTP Compression(CMP): NO
       Idle timeout: Client: 180 sec
                                       Server: 360 sec
       Client IP: DISABLED
       Cacheable: NO
       SC: OFF
       SP: OFF
      Down state flush: ENABLED
       Monitor Connection Close : NONE
       Appflow logging: ENABLED
       ContentInspection profile name: ???
      Network profile name: cnc-cluster1_netprof
       Process Local: DISABLED
       Traffic Domain: θ
       Comment: "ing:web-ingress,ingport:80,ns:default,svc:frontend,svcport:80"
Done
```

3. Run the following command: show pbr.

NetScaler node controller adds PBRs.



Here:

- The number of PBRs is equal to the number of Kubernetes nodes. In this case, it adds two PBRs.
- The srcIP of the PBR is the SNIP added by NetScaler node controller in the tunnel network . The destIP is the CNI overlay subnet range of the Kubernetes node. The NextHop is the IP address of Kubernetes node's VXLAN tunnel interface.

#### Note:

NetScaler node controller adds PBRs instead of static routes. The rest of the configuration of the VXLAN and bridge table remains the same. For more information, see the NetScaler node controller configuration.

## Single tier NetScaler Ingress solution for MongoDB

#### June 5, 2025

MongoDB is one of the most popular NoSQL databases which is designed to process and store massive amounts of unstructured data. Cloud-native applications widely use MongoDB as a NoSQL database in the Kubernetes platform.

To identify and troubleshoot performance issues are a challenge in a Kubernetes environment due to the massive scale of application deployments. For database deployments like MongoDB, monitoring

is a critical component of database administration to ensure that high availability and high performance requirements are met.

NetScaler provides an ingress solution for load balancing and monitoring MongoDB databases on a Kubernetes platform using the advanced load balancing and performance monitoring capabilities of NetScalers. NetScaler Ingress solution for MongoDB provides you deeper visibility into MongoDB transactions and helps you to quickly identify and address performance issues whenever they occur. Using NetScaler Observability Exporter, you can export the MongoDB transactions to Elasticsearch and visualize them using Kibana dashboards to get deeper insights.

The following diagram explains NetScaler Ingress solution for MongoDB using a single-tier deployment of NetScaler.



In this solution, a NetScaler VPX is deployed outside the Kubernetes cluster (Tier-1) and NetScaler Observability Exporter is deployed inside the Kubernetes cluster.

The Tier-1 NetScaler VPX routes the traffic (North-South) from MongoDB clients to Mongo DB query routers (Mongos) in the Kubernetes cluster. NetScaler Observability Exporter is deployed inside the Kubernetes cluster.

As part of this deployment, an Ingress resource is created for NetScaler VPX (Tier-1 Ingress). The Tier-1 Ingress resource defines rules to enable load balancing for MongoDB traffic on NetScaler VPX and specifies the port for Mongo. Whenever MongoDB traffic arrives on the specified port on a NetScaler VPX, it routes this traffic to one of the Mongo service instances mentioned in the Ingress rule. Mongo service is exposed by the MongoDB administrator, and the same service instance is specified in the Ingress.

The NetScaler Observability Exporter instance aggregates transactions from NetScaler VPX and uploads them to the Elasticsearch server. You can set up Kibana dashboards to visualize the required data (for example, query response time, most queried collection names) and analyze them to get meaningful insights. Only insert, update, delete, find, and reply operations are parsed and metrics are sent to the NetScaler Observability Exporter.

#### Prerequisites

You must complete the following steps before deploying the NetScaler Ingress solution for MongoDB.

- Set up a Kubernetes cluster in cloud or on-premises
- Deploy MongoDB in the Kubernetes cluster with deployment mode as sharded replica set. Other deployment modes for MongoDB are not supported.
- Ensure that you have Elasticsearch installed and configured. Use the elasticsearch.yaml file for deploying Elasticsearch.
- Ensure that you have installed Kibana to visualize your transaction data. Use the kibana.yaml file for deploying Kibana.
- Deploy a NetScaler VPX instance outside the Kubernetes cluster. For instructions on how to deploy NetScaler VPX, see Deploy a NetScaler VPX instance.

Perform the following after you deploy the NetScaler VPX:

- 1. Configure an IP address from the subnet of the Kubernetes cluster as SNIP on the NetScaler. For information on configuring SNIPs in NetScaler, see Configuring Subnet IP Addresses (SNIPs).
- 2. Enable management access for the SNIP that is the same subnet of the Kubernetes cluster. The SNIP should be used as NS\_IP variable in the NetScaler Ingress Controller YAML file to enable the NetScaler Ingress Controller to configure the Tier-1 NetScaler.

#### Note:

It is not mandatory to use SNIP as NS\_IP. If the management IP address of the NetScaler is reachable from the NetScaler Ingress Controller then you can use the management IP address as NS\_IP.

- 3. Create a NetScaler system user account specific to the NetScaler Ingress Controller. The NetScaler Ingress Controller uses the system user account to automatically configure the Tier-1 NetScaler.
- 4. Configure NetScaler VPX to forward DNS queries to CoreDNS pod IP addresses in the Kubernetes cluster.

<sup>1</sup> add dns nameServer <core-dns-pod-ip-address>

For example, if the pod IP addresses are 192.244.0.2 and 192.244.0.3, configure the name servers on NetScaler VPX as:

```
1 add dns nameServer 192.244.0.3
2 add dns nameServer 192.244.0.2
```

#### **Deploy the NetScaler Ingress solution for MongoDB**

When you deploy the NetScaler Ingress solution for MongoDB, you deploy the following components in the Kubernetes cluster:

- A stand-alone NetScaler Ingress Controller for NetScaler VPX
- An Ingress resource for NetScaler VPX
- NetScaler Observability Exporter

Perform the following steps to deploy the NetScaler Ingress solution for MongoDB.

1. Create a Kubernetes secret with the user name and password for NetScaler VPX.

kubectl create secret generic nslogin --from-literal=username='
username' --from-literal=password='mypassword'

2. Download the cic-configmap.yaml file and then deploy it using the following command.

1 kubectl create -f cic-configmap.yaml

- 3. Deploy the NetScaler Ingress Controller as a pod using the following steps.
  - a) Download the NetScaler Ingress Controller manifest file. Use the following command:

```
1 wget https://raw.githubusercontent.com/citrix/citrix-k8s-
ingress-controller/master/deployment/dual-tier/manifest/
tier-1-vpx-cic.yaml
```

b) Edit the NetScaler Ingress Controller manifest file and enter the values for the following environmental variables:

Environment Variable	Mandatory or Optional	Description
NS_IP	Mandatory	The IP address of the NetScaler appliance. For more details, see Prerequisites.
NS_USER and NS_PASSWORD	Mandatory	The user name and password of the NetScaler VPX or MPX appliance used as the Ingress device.

Environment Variable	Mandatory or Optional	Description
EULA	Mandatory	The End User License
		Agreement. Specify the value
		as Yes.
LOGLEVEL	Optional	The log levels to control the
		logs generated by NetScaler
		Ingress Controller. By default,
		the value is set to DEBUG. The
		supported values are:
		CRITICAL, ERROR, WARNING,
		INFO, and DEBUG.
NS_PROTOCOL and NS_PORT	Optional	Defines the protocol and port
		that must be used by the
		NetScaler Ingress Controller to
		communicate with NetScaler.
		By default, the NetScaler
		Ingress Controller uses HTTPS
		on port 443. You can also use
		HTTP on port 80.
ingress-classes	Optional	If multiple Ingress load
		balancers are used to load
		balance different Ingress
		resources. You can use this
		environment variable to specify
		the NetScaler Ingress
		Controller to configure
		NetScaler associated with a
		specific Ingress class. For
		information on Ingress classes,
		see Ingress class support
NS_VIP	Optional	NetScaler Ingress Controller
		uses the IP address provided in
		this environment variable to
		configure a virtual IP address to
		the NetScaler that receives
		Ingress traffic.

c) Specify or modify the following arguments in the NetScaler Ingress Controller YAML file.

```
1 args:
2 - --configmap
3 default/cic-configmap
4 - --ingress-classes
5 tier-1-vpx
```

d) Deploy the updated NetScaler Ingress Controller manifest file using the following command:

1 kubectl create -f tier-1-vpx-cic.yaml

4. Create an Ingress object for the Tier-1 NetScaler using the tier-1-vpx-ingress.yaml file.

```
1 kubectl apply -f tier-1-vpx-ingress.yaml
```

Following is the content for the tier-1-vpx-ingress.yaml file. As per the rules specified in this Ingress resource, NetScaler Ingress Controller configures the NetScaler VPX to listen for MongoDB traffic at port 27017. As shown in this example, you must specify the service that you have created for MongoDb query routers (for example:serviceName: mongodb-mongos) so that the NetScaler VPX can route traffic to it. Here, mongodb-mongos is the service for MongoDB query routers.

```
apiVersion: networking.k8s.io/v1
1
2
     kind: Ingress
3
     metadata:
4
       annotations:
5
         ingress.citrix.com/analyticsprofile: '{
6
    "tcpinsight": {
7
    "tcpBurstReporting":"DISABLED" }
8
     }
    ŧ.
9
         ingress.citrix.com/insecure-port: "27017"
10
         ingress.citrix.com/insecure-service-type: mongo
11
         ingress.citrix.com/insecure-termination: allow
12
13
       name: vpx-ingress
14
     spec:
       ingressClassName: tier-1-vpx
15
16
       defaultBackend:
17
           service:
18
             name: mongodb-mongos
19
              port:
20
                number: 27017
```

5. Deploy NetScaler Observability Exporter with Elasticsearch as the endpoint using the coe-esmongo.yaml file.

1 kubectl apply -f coe-es-mongo.yaml

#### Note:

You must set the Elasticsearch server details in the ELKServer environment variable either based on an IP address or the DNS name, along with the port information.

Following is a sample ConfigMap file.

```
1 apiVersion: v1
2 kind: ConfigMap
3 metadata:
4
    name: coe-config-es
5 data:
6
    lstreamd_default.conf:
7
       {
8
            "Endpoints": {
9
                "ES": {
11
12
                    "ServerUrl": "elasticsearch.default.svc.cluster.local
13
                        :9200",
                    "IndexPrefix":"adc_coe",
14
                    "IndexInterval": "daily",
15
16
                    "RecordType": {
17
                         "HTTP": "all",
18
                         "TCP": "all",
19
                         "SWG": "all",
20
                         "VPN": "all",
21
                        "NGS": "all",
                        "ICA": "all",
23
                        "APPFW": "none",
24
                        "BOT": "none",
25
                        "VIDEOOPT": "none".
26
                        "BURST_CQA": "none",
27
28
                         "SLA": "none",
                         "MONGO": "all"
29
                     }
31
    ,
32
                    "ProcessAlways": "no",
                    "ProcessYieldTimeOut": "500",
33
                    "MaxConnections": "512",
34
                    "ElkMaxSendBuffersPerSec": "64",
                    "JsonFileDump": "no"
37
                 }
             }
40
41
        }
```

### Verify the deployment of NetScaler Ingress solution

You can use the command as shown in the following example to verify that all the applications are deployed and list all services and ports.

oot@kubenode222:~# kubectl get pods,svcall-namespaces -o wide									
IAMESPACE	NAME	READY	STATUS	RESTARTS	AGE	IP	NODE	NOMINATED NODE	READINESS GATES
lefault	pod/cic-k8s-ingress-controller	1/1	Running		3h18m		kubenode223		
lefault	pod/coe-es-66bb7f795d-7hftr	1/1	Running		15d		kubenode223		
lefault	pod/elasticsearch-5775b46ccd-gnt97	1/1	Running		90d		kubenode223		
lefault	pod/kibana-5954647d85-6f4mj	1/1	Running		90d		kubenode223		
lefault	pod/mongo-786f4cb565-b49fm	1/1	Running		95d		kubenode223		
ube-system	pod/coredns-74ff55c5b-jtrnc	1/1	Running		96d		kubenode222		
ube-system	pod/coredns-74ff55c5b-rqvv4	1/1	Running		96d		kubenode222		
ube-system	pod/etcd-kubenode222	1/1	Running		96d		kubenode222		
ube-system	pod/kube-apiserver-kubenode222	1/1	Running		96d		kubenode222		
ube-system	pod/kube-controller-manager-kubenode222	1/1	Running		96d		kubenode222		
ube-system	pod/kube-flannel-ds-j4qrr	1/1	Running	θ	95d	10.102.217.222	kubenode222		
ube-system	pod/kube-flannel-ds-nm556	1/1	Running		95d		kubenode223		
ube-system	pod/kube-proxy-9q8n9	1/1	Running		95d	10.102.217.223	kubenode223		
ube-system	pod/kube-proxy-n755g	1/1	Running		96d		kubenode222		
ube-system	pod/kube-scheduler-kubenode222	1/1	Running		96d		kubenode222		

You can use the kubectl get ingress command as shown in the following example to get information about the Ingress objects deployed.

1	1 # kubectl get ingress					
2						
3	NAME	HOSTS	ADDRESS	PORTS	AGE	
4	vpx-ingress	*		80	22d	

#### Verify observability for MongoDB traffic

This topic provides information on how to get visibility into MongoDB transactions using the NetScaler Ingress solution and it uses Kibana dashboards to visualize the database performance statistics.

Before performing the steps in this topic ensure that:

- You have deployed MongoDB as a sharded replica set in the Kubernetes cluster
- Deployed the NetScaler Ingress solution for MongoDB
- A client application for MongoDB is installed to send traffic to the MongoDB.
- Kibana is installed for visualization

Perform the following steps to verify observability for MongoDB traffic.

1. Configure your client application for MongoDB to point to the virtual IP address of the Tier-1 NetScaler VPX.

For example:

```
1 mongodb://<vip-of-vpx>:27017/
```

- 2. Send multiple requests (for example insert, update, delete) to the MongoDB database using your MongoDB client application. The transactions are uploaded to the Elasticsearch server.
- 3. Set up a Kibana dashboard to visualize the MongoDB transactions. You can use the following sample Kibana dashboard.



In this dashboard, you can see performance statistics for your MongoDB deployment including the different type of queries and query response time. Analyzing this data helps you to find any anomalies like latency in a transaction and take immediate action.

## **Export telemetry data to Prometheus**

For your Kubernetes deployment, if you have your Prometheus server deployed in the same Kubernetes cluster, you can configure annotations to enable Prometheus to automatically add NetScaler Observability Exporter as a scrape target.

Following is a snippet of NetScaler Observability Exporter YAML file (coe-es-mongodb.yaml) with these annotations.

```
template:
1
2
    metadata:
3
       name: coe-es
4
       labels:
5
         app: coe-es
6
       annotations:
         prometheus.io/scrape: "true"
7
8
         prometheus.io/port: "5563"
```

Alternatively, you can manually add NetScaler Observability Exporter as the scrape target on your Prometheus server configuration file.

Also, ensure that metrics for Prometheus are enabled in the cic-configmap.yaml file as shown in the following YAML file.

```
1 apiVersion: v1
2 kind: ConfigMap
3 metadata:
4 name: cic-configmap
5 namespace: default
6 data:
7
   NS_ANALYTICS_CONFIG:
      distributed_tracing:
8
         enable: 'false'
9
10
         samplingrate: 0
11
       endpoint:
         server: 'coe-es.default.svc.cluster.local'
13
        timeseries:
14
         port: 5563
15
          metrics:
           enable: 'true'
16
17
           mode: 'prometheus'
18
         auditlogs:
           enable: 'false'
19
20
          events:
21
            enable: 'false'
        transactions:
22
23
         enable: 'true'
          port: 5557
24
```

In this YAML file, the following configuration enables metrics for Prometheus.

```
1 metrics:
2 enable: 'true'
3 mode: 'prometheus'
```

# Canary and blue-green deployment using NetScaler VPX and Azure pipelines for Kubernetes based applications

December 31, 2023

This topic provides information on how to achieve canary and blue-green deployment for Kubernetes applications using NetScaler VPX and Azure pipelines.

## Canary deployment using NetScaler VPX and Azure pipelines for Kubernetes based applications

Canary is a deployment strategy which involves deploying new versions of an application in small and phased incremental steps. The idea of canary is to first deploy the new changes to a small set of users to take a decision on whether to reject or promote the new deployments and then roll out the changes to the rest of the users. This strategy limits the risk involved in deploying a new version of the application in the production environment.

Azure pipelines are a cloud service provided by Azure DevOps which allows you to automatically run builds, perform tests, and deploy code to various development and production environments.

This section provides information on how to achieve canary deployment for Kubernetes based application using NetScaler VPX and NetScaler Ingress Controller with Azure pipelines.

#### Benefits of Canary deployment

- Canary version of application acts as an early warning for potential problems that might be present in the new code and the deployment issues.
- You can use the canary version for smoke tests and A/B testing.
- Canary offers easy rollback and zero-downtime upgrades.
- You can run multiple versions of applications together at the same time.

In this solution, NetScaler VPX is deployed on the Azure platform to enable load balancing of an application and achieve canary deployment using NetScaler VPX. For more information on how to deploy NetScaler on Microsoft Azure, see the NetScaler documentation link.

#### **Canary deployment using NetScaler**

You can achieve canary deployment using NetScaler with Ingress annotations which is a rule based canary deployment. In this approach, you need to define an additional Ingress object with specific annotations to indicate that the application request needs to be served based on the rule based canary deployment strategy. In the Citrix solution, Canary based traffic routing at the Ingress level can be achieved by defining various sets of rules as follows:

- Applying the canary rules based on weight
- Applying the canary rules based on the HTTP request header
- Applying the canary rules based on the HTTP header value

For more information, see simplified canary deployment using Ingress annotations

#### Canary deployment using NetScaler VPX with Azure pipelines

Citrix proposes a solution for canary deployment using NetScaler VPX and NetScaler Ingress Controller with Azure pipelines for Kubernetes based applications.



In this solution, there are three configuration directories:

- kubernetes\_configs
- deployment\_configs
- pipeline\_configs

**kubernetes\_configs** This directory includes the version based application specific deployment YAML files and the Helm based configuration files to deploy NetScaler Ingress Controller which is responsible to push NetScaler configuration to achieve canary deployment.



#### Note:

You can download the latest Helm charts from the NetScaler Ingress Controller Helm charts repository and place it under the cic\_helm directory.

**deployment\_configs** This directory includes the setup\_config and teardown\_config JSON files that specify the path of the YAML files available for the specific version of the application to be deployed or brought down during canary deployment.



**pipeline\_configs** This directory includes the Azure pipeline script and the python script which reads the user configurations and triggers the pipeline based on the user request to introduce a new version of the application or teardown a version of an application. The change in percentage of traffic weight in application ingress YAML would trigger the pipeline to switch the traffic between the available version of applications.



With all the three configuration files in place, any update to the files under deployment\_configs and kubernetes\_configs directories in GitHub, would trigger the pipeline in Azure.

The traffic split percentage can be adjusted using the ingress.citrix.com/canary-weight annotation in the ingress YAML of the application.

#### Deploy a sample application on Canary in Azure pipelines

This topic explains how to deploy a sample application on Canary mode using NetScaler and Azure pipelines.

#### Prerequisites Ensure that:

- NetScaler VPX is already deployed on the Azure platform and is ready to be used by our sample application.
- AKS cluster with Kubernetes service connection configured for the Azure pipeline.

#### Perform the following steps:

- 1. Clone the GitHub repository and go to the directory cd/canary-azure-devops.
- 2. Place the application deployment specific YAMLs (with the ingress file) under a versioned folder v1 in the kubernetes\_configs directory.
- 3. Create three Azure pipelines using the existing YAML files, deploy\_cic.yaml, deploy. yaml, and teardown.yaml, for deploying NetScaler Ingress Controller and deploying and tearing down the applications. See, Azure pipelines for creating a pipeline.
- 4. Update the subscription, agent pool, service connection and NetScaler details in the pipeline YAML.
- 5. Save the pipeline.

6. Update the path in deploy\_config.json with the path specifying the directory where the application YAMLs are placed.

```
1 {
2 
3 
4 "K8S_CONFIG_PATH" : "cd/canary-azure-devops/kubernetes_configs/v1
5 
6 }
```

- 7. Commit the deploy\_config.json file and v1 directory using Git to trigger the pipeline to deploy the v1 version of the application.
- 8. Access the application through NetScaler.
- Introduce the v2 version of the application by creating the v2 directory under kubernetes\_configs
   Make sure that the ingress under this version has the canary annotation specified with the right weight to be set for traffic split.
- 10. Deploy the version v2 of the application by updating deploy\_config.json with the path specifying the v2 directory. Now, the traffic is split between version v1 and v2 based on the canary weight set in the ingress annotation (for example, ingress.citrix.com/canary -weight: "40")
- 11. Continue progressively increasing the traffic weight in the ingress annotation until the new version is ready to serve all the traffic.

## Blue-green deployment using NetScaler VPX and Azure pipelines for Kubernetes based applications

Blue-green deployment is a technique that reduces downtime and risk by running two identical production environments called blue and green. At any time, only one of the environments is live that serves all the production traffic. The basis of the blue-green method is side-by-side deployments of two separate but identical environments. Deploying an application in both the environments can be fully automated by using jobs and tasks. This approach enforces duplication of every resource of an application. However, there are many different ways blue-green deployments can be carried out in various continuous deployment tools.

Using NetScaler VPX with Azure pipelines the same canary based solution can be used to achieve bluegreen deployment by adjusting the traffic weight to either zero or 100.
# Traffic management for external services

#### December 31, 2023

Sometimes, all the available services of an application may not be deployed completely on a single Kubernetes cluster. You may have applications that rely on the services outside of one cluster as well. In this case, micro services need to define an ExternalName service to resolve the domain name. However, in this approach, you would not be able to get features such as traffic management, policy enforcement, fail over management and so on. As an alternative, you can configure NetScaler to resolve the domain names and leverage the features of NetScaler.

# **Configure NetScaler to reach external services**

You can configure NetScaler as a domain name resolver using NetScaler Ingress Controller. When you configure NetScaler as domain name resolver, you need to resolve:

- Reachability of NetScaler from microservices
- Domain name resolution at NetScaler to reach external services

#### Configure a service for reachability from Kubernetes cluster to NetScaler

To reach NetScaler from microservices, you have to define a headless service which would be resolved to a NetScaler service and thus the connectivity between microservices and NetScaler establishes.

```
1 apiversion: v1
2 kind: Service
3 metadata:
4 name: external-svc
5 spec:
6 selector:
7 app: cpx
8 ports:
9 - protocol: TCP
10 port: 80
```

#### Configure NetScaler as a domain name resolver using NetScaler Ingress Controller

You can configure NetScaler through NetScaler Ingress Controller to create a domain based service group using the ingress annotation ingress.citrix.com/external-service. The value for ingress.citrix.com/external-service is a list of external name services with their corresponding domain names. For NetScaler VPX, name servers are configured on NetScaler using the ConfigMap.

# Note:

ConfigMaps are used to configure name servers on NetScaler only for NetScaler VPX. For NetScaler CPX, CoreDNS forwards the name resolution request to the upstream DNS server.

# Traffic management using NetScaler CPX

The following diagram explains NetScaler CPX deployment to reach external services. An Ingress is deployed where the external service annotation is specified to configure DNS on NetScaler CPX.



Note: A ConfigMap is used to configure name servers on NetScaler VPX.



In this deployment:

- 1. A microservice sends the DNS query for www.externalsvc.com which would get resolved to the NetScaler CPX service.
- 2. NetScaler CPX resolves www.externalsvc.com and reaches external service.

Following are the steps to configure NetScaler CPX to load balance external services:

- 1. Define a headless service to reach NetScaler.
  - 1 apiVersion: v1

```
2 kind: Service
3 metadata:
4 name: external-svc
5 spec:
6 selector:
7 app: cpx
8 ports:
9 - protocol: TCP
10 port: 80
```

2. Define an ingress and specify the external-service annotation as specified in the dbsingress.yaml file. When you specify this annotation, NetScaler Ingress Controller creates DNS servers on NetScaler and binds the servers to the corresponding service group.

```
1 annotations:
2 ingress.citrix.com/external-service: '{
3 "external-svc": {
4 "domain": "www.externalsvc.com" }
5 }
6 '
```

3. Add the IP address of the DNS server on NetScaler using ConfigMap.

#### Note:

This step is applicable only for NetScaler VPX.

```
1 apiVersion: v1
2 kind: ConfigMap
3 metadata:
4 name: nameserver-cmap
5 namespace: default
6 data:
7 NS_DNS_NAMESERVER: '[]'
```

# Support for external name service across namespaces

#### December 31, 2023

Namespaces are used to isolate resources within a Kubernetes cluster. Sometimes, services in a different namespace might have to access a service located in another namespace. In such scenarios, you can use the ExternalName service provided by Kubernetes. An ExternalName service is a special service that does not have selectors and instead uses DNS names.

In the service definition, the externalName field must point to the namespace and also to the service which we are trying to access on that namespace. Citrix ingress controller supports services of type ExternalName when you have to access services within the cluster.

When you create the ExternalName service, the following criteria must be met:

- The externalName field in the service definition must follow the format: svc: <name-of-the-service>.<namespace-of-the-service>.svc.cluster .local
- The port number in the ExternalName service must exactly match the port number of the targeted service.

#### Note:

When the service of an application is outside the Kubernetes cluster and you have created an ExternalName service, you can resolve the domain name using the Traffic management for external services feature.

#### Sample ExternalName service

In this example, a mysql service is running in the default namespace and a sample ExternalName service is created to access the mysql service from the namespace1 namespace.

The following is a sample service definition for a MySQL service running in the default namespace.

```
1 apiVersion: v1
2 kind: Service
3 metadata:
4 name: mysql
5 namespace: default
6 spec:
7
   clusterIP: None
8 ports:
  - port: 3306
9
    protocol: TCP
targetPort: 3306
10
11
12 selector:
13
     app: mysql
14 type: ClusterIP
```

The following is a sample ExternalName service definition to access the mysql service from the namespace1 namespace.

```
1 kind: Service
2 apiVersion: v1
3 metadata:
4 name: dbservice
5 namespace: namespace1
6 spec:
7 type: ExternalName
8 externalName: mysql.default.svc.cluster.local
9 ports:
10 - port: 3306
```

11 protocol: TCP
12 targetPort: 3306

In the example, the service points to the namespace where mysql is deployed as specified in the field externalName: mysql.default.svc.cluster.local. Here mysql is the service name and default is the namespace. You can see that the port name is also the same as the mysql service.

# **Troubleshooting NetScaler Ingress Controller**

December 20, 2024

You can debug NetScaler Ingress Controller using the following methods. First, use the event-based debugging method followed by the log-based debugging method. Use the NetScaler kubectl plug-in and NSIC diagnostic tool for advanced debugging.

- Event-based debugging
- Log-based debugging
- NetScaler Kubectl plug-in
- NSIC diagnostic tool

## **Event-based debugging**

Events in Kubernetes are entities that offer insights into the operational flow of other Kubernetes entities.

Event-based debugging for NetScaler Ingress Controller is enabled at the pod level. Use the following command to view the events for NetScaler Ingress Controller.

```
1 kubectl describe pods <citrix-k8s-ingress-controller pod name> -n <
    namespace of pod>
```

You can view the events under the Events section.

In the following example, NetScaler has been deliberately made unreachable and the same information can be seen under the Events section.

```
kubectl describe pods cic-vpx-functionaltest -n functionaltest
Name: cic-vpx-functionaltest
Namespace: functionaltest
Events:
```

```
Type Reason Age From
7
         Message
8
       ____ ____
                         ____
                               ____
         _____
9
      Normal Pulled
                         33m kubelet, rak-asp4-node2
         Container image "citrix-ingress-controller:latest" already
         present on machine
                               kubelet, rak-asp4-node2
10
      Normal Created 33m
         Created container cic-vpx-functionaltest
11
      Normal Started 33m
                               kubelet, rak-asp4-node2
          Started container cic-vpx-functionaltest
12
      Normal Scheduled 33m
                               default-scheduler
          Successfully assigned functionaltest/cic-vpx-functionaltest to
          rak-asp4-node2
13
14
      Normal Created
                        33m CIC ENGINE, cic-vpx-functionaltest
         CONNECTED: NetScaler:<NetScaler IP>:80
15
      Normal Created 33m CIC ENGINE, cic-vpx-functionaltest
         SUCCESS: Test LB Vserver Creation on NetScaler:
16
      Normal Created 33m CIC ENGINE, cic-vpx-functionaltest
         SUCCESS: ENABLING INIT features on NetScaler:
      Normal Created 33m CIC ENGINE, cic-vpx-functionaltest
17
          SUCCESS: GET Default VIP from NetScaler:
18
      Warning Created 17s CIC ENGINE, cic-vpx-functionaltest
         UNREACHABLE: NetScaler: Check Connectivity::<NetScaler IP>:80
```

For further debugging, check the logs of the NetScaler Ingress Controller pod.

# Log-based debugging

You can change the log level of NetScaler Ingress Controller at runtime using the ConfigMap feature. For changing the log level during runtime, see the ConfigMap documentation.

To check logs on NetScaler Ingress Controller, use the following command.

1 kubectl logs <citrix-k8s-ingress-controller> -n namespace

The following table describes some of the common issues and workarounds.

Problem	Log	Workaround
NetScaler instance is not reachable	2019-01-10 05:05:27,250 - ERROR - [nitrointer- face.py:login_logout:94] (MainThread) Exception: HTTPConnectionPool(host=' 10.106.76.200', port=80): Max retries exceeded with url: /nitro/v1/config/login (Caused by NewConnectionError(' <url- lib3.connection.HTTPConnect object at 0x7f4d45bd63d0&gt;: Failed to establish a new connection: [Errno 113] No route to host',))</url- 	Ensure that NetScaler is up and running, and you can ping the NSIP address.
Wrong user name or password	2019-01-10 05:03:05,958 - ERROR - [nitrointer- face.py:login_logout:90] (MainThread) Nitro Excep- tion::login_logout::errorcode: username or password	=354,message=Invalid
SNIP is not enabled with management access	2019-01-10 05:43:03,418 - ERROR - [nitrointer- face.py:login_logout:94] (MainThread) Exception: HTTPConnectionPool(host=' 10.106.76.242', port=80): Max retries exceeded with url: /nitro/v1/config/login (Caused by NewConnectionError(' <url- lib3.connection.HTTPConnect object at 0x7f302a8cfad0&gt;: Failed to establish a new connection: [Errno 110] Connection timed out',))</url- 	Ensure that you have enabled the management access in NetScaler (for NetScaler VPX high availability) and set the IP address, <b>NSIP</b> , with management access enabled.

Problem	Log	Workaround
Error while parsing annotations	2019-01-10 05:16:10,611 - ERROR - [kuber- netes.py:set_annotations_to_co (MainThread) set_annotations_to_csapp: Error message=No JSON object could be decodedInvalid Annotation \$service_weights please fix	sapp:1040]
	and apply \${"frontend":, "catalog":95}	
Wrong port for NITRO access	2019-01-10 05:18:53,964 - ERROR - [nitrointer- face.py:login_logout:94] (MainThread) Exception: HTTPConnectionPool(host=' 10.106.76.242', port=34438): Max retries exceeded with url: /nitro/v1/config/login (Caused by NewConnectionError(' <url- lib3.connection.HTTPConnecti object at 0x7fc592cb8b10&gt;: Failed to establish a new connection: [Errno 111] Connection refused',))</url- 	Verify if the correct port is specified for NITRO access. By default, NetScaler Ingress Controller uses the port <b>80</b> for communication.
Ingress class is wrong	2019-01-10 05:27:27,149 - INFO - [kuber- netes.py:get_all_ingresses:132 (MainThread) Unsupported Ingress class for ingress	Verify that the ingress file belongs to the ingress class P <b>g</b> hat NetScaler Ingress Controller monitors.

Problem	Log	Workaround			
Kubernetes API is not reachable	2019-01-10 05:32:09,729 -	Check if the kubernetes_url is			
	ERROR -	correct. Use the command,			
	[kubernetes.py:_get:222]	kubectl cluster-info <b>to</b>			
	(Thread-1) Error while calling	get the URL information.			
	/ser-	Ensure that the Kubernetes			
	vices:HTTPSConnectionPool(h	<b>ost</b> <del>, a</del> in node is running at			
	10.106.76.237', port=6443):	https://			
	Max retries exceeded with	kubernetes_master_addres			
	url: /api/v1/services (Caused	: 6443 and the Kubernetes API			
	by NewConnectionError( ' <url-< td=""><td>server pod is up and running.</td></url-<>	server pod is up and running.			
	lib3.connection.VerifiedHTTPS	Connection			
	object at 0x7fb3013e7dd0>:				
	Failed to establish a new				
	connection: [Errno 111]				
	Connection refused',))				
Incorrect service port specified	NA	Provide the correct port details			
in the YAML file		in the ingress YAML file and			
		reapply the ingress YAML to			
		solve the issue.			
Load balancing virtual server	NA	Check for the service name and			
and service group are created		port used in the YAML file. For			
but are down		NetScaler VPX, ensure that			
		feature-node-watchis			
		set to <b>true</b> when bringing up			
		NetScaler Ingress Controller.			
Content switching (CS) virtual	NA	Use the annotation,			
server is not getting created for		ingress.citrix.com/			
NetScaler VPX.		frontend-ip, in the ingress			
Incorrect secret provided in the	2019-01-10 09:30:50.673 -	YAML file for NetScaler VPX. Correct the values in the YAML			
TLS section in the ingress YAML	INFO -	file and reapply YAML to solve			
file	[kubernetes.nv: get:231]	the issue.			
	(MainThread) Resource not				
	found:				
	/secrets/default-secret12345				
	namespace default				

Problem	Log	Workaround
	2019-01-10 09:30:50,673 -	
	INFO - [kuber-	
	netes.py:get_secret:1712]	
	(MainThread) Failed to get	
	secret for the app	
	default-secret12345.default	
The feature-node-watch	ERROR - [nitrointer-	This error occurs when
argument is specified, but	face.py:add_ns_route:4495]	feature-node-watchis
static routes are not added in	(MainThread) Nitro Excep-	enabled and NetScaler VPX and
NetScaler VPX	tion::add_ns_route::errorcode	<b>=60.4),enestagel</b> e <b>The</b> are not in
	gateway is not directly	the same network. You must
	reachable	removethefeature-
		node-watch argument from
		NetScaler Ingress Controller
		YAML file. Static routes do not
		work when NetScaler VPX and
		Kubernetes cluster are in
		different networks. Use node
		controller to create tunnels
		between NetScaler VPX and
	EDDOD	cluster nodes.
CRD status not updated	Error -	Verify the permission to push
	(MainThroad) Excention	<b>TRB</b> status is provided in the
	during CBD status undate for	RBAC. The permission should
	negrwaddmuloccmod: 403	be similar to the following YAML
	Client Error: Forbidden for	• apiGroups: ["citrix.com"]
	url https://	resources: [
	10.96.0.1:443/anis/	"rewritepolicies/status",
	citrix.com/v1/	"canarycrds/status",
	namespaces/default/	"authpolicies/status",
	rewritepolicies/	"ratelimits/status",
	negrwaddmuloccmod/	"listeners/status",
	status	"httproutes/status",
		``wafs/status'']

Problem	Log	Workaround
NetScaler Ingress Controller	ERROR -	Verify that the permission to
event not updated	[clienthelper.py:post:94]	update the NetScaler Ingress
	(MainThread) Reuqest	Controller pod events is
	/events to api server is	provided in the RBAC rules.
	forbidden	<ul> <li>apiGroups: ['"'] resources:</li> </ul>
		["events"] verbs: [
		"create"]
Rewrite-responder policy not	ERROR - [con-	Such errors are due to incorrect
added	fig dispatcher.py: dispatch	confige backsi241
	(Dispatcher) Status: 104,	rewrite-responder CRDs. Fix
	ErrorCode: 3081, Reason:	the expression and reapply the
	Nitro Exception: Expression	CRD.
	syntax error [D(10,	
	20).^RE_SELECT(, Offset 15] <	
	ERROR - [con-	
	fig_dispatcher.py:dispatch_	config_pack:324]
	(Dispatcher) Status: 104,	
	ErrorCode: 3098, Reason:	
	Nitro Exception: Invalid	
	expression data type	
	[ent.ip.src^, Offset 13]	
Application of a CRD failed. The	2020-07-13 08:49:07.620 -	Log shows that the NITRO
NetScaler Ingress Controller	ERROR - [con-	command has failed. The same
converts a CRD into a set of	fig dispatcher.pv: dispatch	cologeappearszigevetScaler also.
configurations to configure	(Dispatcher) Failed to execute	Check the NetScaler ns.log
NetScaler to the desired state	config	and search for the error string
as per the specified CRD. If the	ADD sslprofile k8s crd k8ser	viceing the grep command to
configuration fails, then the	service_default_80 <i>tcp_backen</i>	_ d{fishie.ksst_୧ନଥିବଃବୋଦ୍ଧice_kuai
CRD instance may not get	service_default_80_tcp_backe	<b>nd</b> ommand that failed during
applied on NetScaler.	sslprofiletype:BackEnd	the application of CRD. Try to
	tls12:enabled } from	delete the CRD and add it
	ConfigPack	again.
	'default.k8service.kuard-	
	service.add_spec'	

Problem	Log	Workaround	
	2020-07-13 08:49	:07,620 -	
	ERROR - [con-		
	fig_dispatcher.py:dispatch_config_pack:257]		
	(Dispatcher) Status: 104,		
	ErrorCode: 1074,	ErrorCode: 1074, Reason:	
	Nitro Exception: Invalid value		
	[sslProfileType, value differs		
	from existing entity and it		
	cant be updated.	I	
	2020-07-13 08:49:07,620 -		
	INFO - [con-		
	fig_dispatcher.py	:dispatch_config_pack:263]	
	(Dispatcher) Proc	essing of	
	ConfigPack		
	'default.k8servic	e.kuard-	
	service.add_spec'failed		

## NetScaler Kubernetes kubectl plug-in

NetScaler provides a kubectl plug-in to inspect NetScaler Ingress Controller deployments and perform troubleshooting operations. You can perform troubleshooting operations using the subcommands available with this plug-in.

Note:

This plugin is supported from NSIC version 1.32.7 onwards.

#### Installation using curl

You can install the kubectl plug-in by downloading it from the NetScaler Modern Apps tool kit repository using curl as follows.

#### For Linux:

```
1 curl -L0 https://github.com/netscaler/modern-apps-toolkit/releases/
	download/v1.0.0-netscaler-plugin/netscaler-plugin_v1.0.0-netscaler-
	plugin_Linux_x86_64.tar.gz
2 gunzip netscaler-plugin_v1.0.0-netscaler-plugin_Linux_x86_64.tar.gz
3 tar -xvf netscaler-plugin_v1.0.0-netscaler-plugin_Linux_x86_64.tar
4 chmod +x kubectl-netscaler
```

```
5 sudo mv kubectl-netscaler /usr/local/bin/kubectl-netscaler
```

#### For Mac:

```
1 curl -s -L https://github.com/netscaler/modern-apps-toolkit/releases/
download/v1.0.0-netscaler-plugin/netscaler-plugin_v1.0.0-netscaler-
plugin_Darwin_x86_64.tar.gz | tar xvz -
2 chmod +x kubectl=netscaler
```

```
2 chmod +x kubectl-netscaler
```

```
3 sudo mv kubectl-netscaler /usr/local/bin/kubectl-netscaler
```

#### Note:

For Mac, you need to enable allow a developer app.

### For Windows:

```
1 curl.exe -L0 https://github.com/netscaler/modern-apps-toolkit/releases/
download/v1.0.0-netscaler-plugin/netscaler-plugin_v1.0.0-netscaler-
plugin_Windows_x86_64.zip | tar xvz
```

#### Note:

For Windows, you must set your \$PATH variable to the directory where the kubectlnetscaler.exe file is extracted.

#### **Installation using Krew**

Krew helps you discover and install kubectl plugins on your machine. Follow the Krew Quickstart Guide to install and set up Krew.

- 1. Install and set up Krew on your machine.
- 2. Download the plugin list:

1 kubectl krew update

3. Discover plugins available on Krew:

```
1 kubectl krew search netscaler
```

```
1NAMEDESCRIPTIONINSTALLED2netscalerInspect NetScaler Ingressesno
```

4. Install the plug-in:

1 kubectl krew install netscaler

#### Note:

For Mac, you need to enable allow a developer app.

#### Examples for usage of subcommands in Kubectl plug-in

The following subcommands are available with this plug-in:

Subcommand	Description
help	Provides information about the various options.
	You can also run this command after installation
	to check if the installation is successful and see
	which commands are available.
status	Displays the status (up, down, or active) of
	NetScaler entities for provided prefix input (the
	default value of the prefix is k8s).
conf	Displays NetScaler configuration (show run
	output).
support	Gets NetScaler (show techsupport) and
	NetScaler Ingress Controller support bundle.
	Support-related information is extracted as two
	tar.gz files. These two tar files are
	show tech support information from
	NetScaler ADC and Kubernetes related
	information for troubleshooting where the
	ingress controller is deployed.

#### Examples for usage of subcommands

**Help command** The help command is used to know about the available commands.

1 # kubectl netscaler --help

For more information about a subcommand use the help command as follows:

1 # kubectl netscaler <command> --help

**Status command** The status subcommand shows the status of various components of NetScaler that are created and managed by NetScaler Ingress Controller in the Kubernetes environment.

The components can be filtered based on either both application prefix (NS\_APPS\_NAME\_PREFIX environment variable for NetScaler Ingress Controller pods or the entity prefix value in the Helm chart) and ingress name or one of them. The default search prefix is k8s.

Flag	Short form	Description
deployment		Name of the ingress controller deployment.
-ingress	-i	Specify the option to retrieve the config status of a particular Kubernetes ingress resource.
–label	-l	Label of the ingress controller deployment.
–output		Output format. Supported formats are tabular (default) and JSON.
–pod		Name of the ingress controller pod.
–prefix	-р	Specify the name of the prefix provided while deploying NetScaler Ingress Controller.
–verbose	-V	If this option is set, additional information such as NetScaler configuration type or service
		port is displayed.

The following example shows the status of NetScaler components created by NetScaler Ingress Controller with the label app=cic-tier2-citrix-cpx-with-ingress-controller and the prefix plugin2 in the NetScaler namespace.

1	<pre># kubectl r     -contro</pre>	netscaler status ller -n netscale	-l app r -p p	p=cic-tier2-citrix lugin	-cpx-with-ingress
2					
3	Showing Net	tScaler componen <sup>.</sup>	ts <b>for</b>	prefix: plugin2	
4	NAMESPACE	INGRESS	PORT	RESOURCE	NAME
					STATUS
5				Listener	plugin
	-198.16	8.0.1_80_http			up
6	default			Traffic Policy	plugin-
	apache2	_80_csp_mqwmhc66	h3bkd5	i4hd224lve7hjfzvoi	active
7	default			Traffic Action	plugin-
	apache2	_80_csp_mqwmhc66	h3bkd5	i4hd224lve7hjfzvoi	attached
8	default	plugin-apache2	80	Load Balancer	plugin-
	apache2	_80_lbv_mqwmhc66	h3bkd5	i4hd224lve7hjfzvoi	up

9	default	plugin-apache2	80	Service		plugin-
	apache2	2_80_sgp_mqwmhc60	Sh3bkd5	i4hd224lv	/e7hjfzvoi	
10	default	plugin-apache2	80	Service	Endpoint	198.168.0.2
					up	
11	netscaler			Traffic	Policy	plugin-
	apache2	2_80_csp_lhmi6gp3	Baytmvn	ww3zczp2y	zlyoacebl	active
12	netscaler			Traffic	Action	plugin-
	apache2	2_80_csp_lhmi6gp3	Baytmvn	ww3zczp2y	zlyoacebl	attached
13	netscaler	plugin-apache2	80	Load Bal	ancer	plugin-
	apache2	2_80_lbv_lhmi6gp3	Baytmvn	ww3zczp2y	zlyoacebl	up
14	netscaler	plugin-apache2	80	Service		plugin-
	apache2	2_80_sgp_lhmi6gp3	Baytmvn	ww3zczp2y	zlyoacebl	
15	netscaler	plugin-apache2	80	Service	Endpoint	198.168.0.3
					an	

**Conf command** The conf subcommand shows the running configuration information on NetScaler (show run output). The l option is used for querying the label of NetScaler Ingress Controller pod.

Flag	Short form	Description
-deployment		Name of the ingress controller deployment.
–label	-l	Label of the ingress controller deployment.
–pod		Name of the ingress controller pod.

The sample output for the kubectl NetScaler conf subcommand is as follows:

```
1
       # kubectl netscaler conf -l app=cic-tier2-citrix-cpx-with-ingress-
          controller -n netscaler
2
3
       set ns config -IPAddress 198.168.0.4 -netmask 255.255.255.255
4
       set ns weblogparam -bufferSizeMB 3
       enable ns feature LB CS SSL REWRITE RESPONDER AppFlow CH
5
       enable ns mode L3 USNIP PMTUD
6
       set system user nsroot -encrypted
7
       set rsskeytype -rsstype ASYMMETRIC
8
       set lacp -sysPriority 32768 -mac 8a:e6:40:7c:7f:47
9
       set ns hostName cic-tier2-citrix-cpx-with-ingress-controller-7
10
          bf9c46cb9-xpwvm
11
       set interface 0/1 -haHeartbeat OFF -throughput 0 -bandwidthHigh 0 -
          bandwidthNormal 0 -intftype Linux -ifnum 0/1
       set interface 0/2 -speed 1000 -duplex FULL -throughput 0 -
12
          bandwidthHigh 0 -bandwidthNormal 0 -intftype Linux -ifnum 0/2
```

# **Support command** The support subcommand gets NetScaler (show techsupport) and NetScaler Ingress Controller support bundle.

#### Warning:

For NetScaler CPX, technical support bundle files are copied to the location you specify. For security reasons, if NetScaler Ingress Controller is managing a NetScaler VPX or NetScaler MPX, then the tech support bundle is extracted only and not copied. You need to get the technical support bundle files from NetScaler manually.

Flag	Short form	Description
-deployment		Name of the ingress controller
		deployment.
–label	-l	Label of the ingress controller
		deployment.
–pod		Name of the ingress controller
		pod.
–appns		List of space-separated
		namespaces (within quotes)
		from where Kubernetes
		resource details such as ingress,
		services, pods, and crds are
		extracted (For example, default
		"namespace1""namespace2")
		(default "default").
–dir	-d	Specify the absolute path of
		the directory to store support
		files. If not provided, the
		current directory is used.
–unhidelP		Set this flag to unhide IP
		addresses while collecting
		Kubernetes information. By
		default, this flag is set to
		false.
-skip-nsbundle		This option disables extraction
		of techsupport from NetScaler.
		By default, this flag is set to
		false.

## Flags for support subcommand:

The following sample output is for the kubectl netscaler support command.

# **NSIC diagnostic tool**

The NSIC diagnostic tool is a shell script that collects information about NetScaler Ingress Controller and applications deployed in the Kubernetes cluster. This tool takes namespace, CNI, and output directory path as an input to extract the necessary information and stores the output files in tar format. If there is any information that user considers sensitive and not to be shared, scan through the output\_ directory under the user provided output directory path and recreate the tar file to share.

Download the NSIC diagnostic tool script and run nsic\_diagnostics\_tool.sh.

# **Upgrade NetScaler Ingress Controller**

#### October 22, 2024

This section describes how to upgrade the NetScaler Ingress Controller instance for both standalone and sidecar (NetScaler CPX with NetScaler Ingress Controller) deployments. The upgrade procedure ensures that data traffic is not affected during the upgrade process.

## Prerequisites

- **Secrets**: If you want to have the upgraded NSIC version in a different namespace, place your secrets in that namespace.
- **Deployment names**: Different names for the YAML-based deployment and the Helm-based deployment should not impact functionality. However, you can retain the same name by using the fullnameOverride argument.
- For a Helm-based NSIC deployment, run the helm get values <release\_name> -a command to retrieve the values.yaml for the existing deployment.

# Upgrade NSIC that was deployed by using a YAML file

If you have an NSIC instance that was deployed by using a YAML file, follow this procedure to upgrade NSIC by using a Helm chart.

1. Back up the CRD instances of your existing deployment. To back up the CRD instances to a YMAL file, run the following command:

2. Set the correct arguments in values.yaml (for a standalone deployment) or values.yaml (for a sidecar deployment) based on the current deployment and ConfigMap YAML files.

Notes:

- We recommend that you specify a unique value for entityPrefix.
- We recommend that you specify the ingressClass as needed to associate your NSIC instance with specific ingress resources.
- We recommend that you specify the serviceClass as needed to associate your NSIC instance with specific services of type LoadBalancer.
- 3. Delete your existing NSIC deployment and other related resources, such as ClusterRole, Cluster-RoleBinding, and ServiceAccount, that were created as part of the existing NSIC deployment.
  - 1 kubectl delete deployment <NSIC\_DEPLOYMENT\_NAME>
    2 kubectl delete clusterrole <NSIC\_CLUSTERROLE\_NAME>
    3 kubectl delete clusterrolebinding <NSIC\_CLUSTERROLEBINDING\_NAME
    4 kubectl delete serviceaccount <NSIC\_SERVICEACCOUNT\_NAME>
- 4. Delete the Configmap.

1 kubectl delete configmap <CONFIGMAP\_NAME>

- 5. Add the NetScaler Helm chart repository to your local registry by using the following command.
  - 1 helm repo add netscaler https://netscaler.github.io/netscaler-helm -charts/

If the NetScaler Helm chart repository is already added to your local registry, update the repository by using the following command:

1 helm repo update netscaler

6. Deploy NetScaler Ingress Controller with the modified values.yaml.

helm install <existing NSIC release name deployed using YAML>
netscaler/netscaler-ingress-controller -f values.yaml

7. Deploy the CRD specification by using the following command.

```
1 kubectl apply -f https://raw.githubusercontent.com/netscaler/
netscaler-helm-charts/refs/heads/master/netscaler-ingress-
controller/crds/crds.yaml
```

8. If needed, deploy the CRD instances that were backed up in step 1.

#### Upgrade NSIC that was deployed by using a Helm chart

Follow this procedure to upgrade NSIC that was deployed using a Helm chart.

1. Back up the CRD instances of your existing deployment. To back up the CRD instances to a YMAL file, run the following command:

```
1 for crd in $(kubectl get crds -o jsonpath='{
2 .items[*].metadata.name }
3 ' | tr ' ' \n' | grep 'citrix.com'); do echo "Getting
            instances for CRD: $crd"; kubectl get $crd --all-namespaces
            -o yaml >> "all_crd_instances.yaml"; done
```

2. (Optional) Add the NetScaler Helm chart repository to your local registry by using the following command.

```
1 helm repo add netscaler https://netscaler.github.io/netscaler-
helm-charts/
```

3. If the NetScaler Helm chart repository is already added to your local registry, update the repository by using the following command:

```
1 helm repo update netscaler
```

- 4. Replace cic.values with nsic.values.
- 5. To upgrade NSIC:
  - a) (Optional) If the existing NSIC is installed using a chart available in the Citrix repository, do the following. Else, skip to the next step.
    - For a standalone NSIC, update the values.yaml with the following attributes.
      - Specify the imageRepository value as netscaler/netscaler-k8singress-controller.
      - (Optional) To upgrade NSIC version 1.35.6 and earlier to any later versions starting from 1.36.5, replace cncPbr with nsncPbr to match the parameter naming in the NetScaler Helm chart.

- Run the helm get values <release\_name> -a command and update the values.yaml file accordingly.
- For NSIC as a sidecar, update the values.yaml with the following attributes.
  - Specify the nsic.imageRepository value as netscaler/netscalerk8s-ingress-controller.
  - (Optional) To upgrade NSIC version 1.35.6 and earlier to any later versions starting from 1.36.5, replace cncPbr with nsncPbr to match the parameter naming in the NetScaler Helm chart.
  - Run the helm get values <release\_name> -a command and update the values.yaml file accordingly.

#### Notes:

- We recommend that you specify a unique value for entityPrefix.
- We recommend that you specify the ingressClass as needed to associate your NSIC instance with specific ingress resources.
- We recommend that you specify the serviceClass as needed to associate your NSIC instance with specific services of type LoadBalancer.

#### b) Run the following command.

```
1 helm upgrade netscaler-ingress-controller netscaler/netscaler-
ingress-controller -f values.yaml [--version=<supported nsic
version>]
```

#### Note:

For a list of supported NSIC versions, see this page.

#### 6. Deploy the CRD specification by using the following command.

```
1 kubectl apply -f https://raw.githubusercontent.com/netscaler/
netscaler-helm-charts/refs/heads/master/netscaler-ingress-
controller/crds/crds.yaml
```

#### 7. If needed, deploy the CRD instances that were backed up in step 1.

# **Entity name change**

December 31, 2023

While adding the NetScaler entities, the NetScaler Ingress Controller maintains unique names per Ingress, service or namespace. Sometimes, it results in NetScaler entities with large names even exceeding the name limits in NetScaler.

Now, the naming format in the NetScaler Ingress Controller is updated to shorten the entity names. In the updated naming format, a part of the entity name is hashed and all the necessary information is provided as part of the entity comments.

After this update, the comments available on lbvserver and servicegroup entity names provides all the necessary details like the ingress name, ingress port, service name, service port, and the namespace of the application.

# Format for comments

Ingress: ing:<ingress-name>,ingport:<ingress-port>,ns:<k8s-namespace>,
svc:<k8s-servicename>,svcport:<k8s-serviceport>

Service of type LoadBalancer:lbsvc:<k8s-servicename>,svcport:<k8s-serviceport
>,ns:<k8s-namespace>

The following table explains the entity name changes introduced with the NetScaler Ingress Controller version 1.12.

Entity	Old naming format	New naming format	Description/Comments
csvserver (ingress)	k8s	k8s	no changes
	-192.2.170.67	-192.2.170.67	
	_80_http	_80_http	
csvserver(type	k8s-	k8s-	Now, the port is
LoadBalancer)	apache_default_80	_ <b>spac</b> he_80_default	_followed by a
			namespace
lbvserver(type	k8s-	k8s-	The comment for
LoadBalancer)	apache_default_80	_ <b>apvac_kke</b> _s80_lbv_wli	k <b>egpæ</b> o5vunbthsoj4lxegk7cc
	-	Comment:	LoadBalancer is
	apache_default_80	_bbscvc:apache,	now different
		<pre>svcport:80,ns:</pre>	
		default	

Entity	Old naming format	New naming format	Description/Comments		
servicegroup	k8s-	k8s-	The suffix sgp is		
(type LoadBalancer)	apache_default_80 -	0_ <b>spac_ke</b> _80_sgp_wl	ik <b>ædpled</b> o5vunbthsoj4lxegk7cc		
	apache_default_80	0_svc			
cspolicy or	k8s-web-	k8s-	Moved service-name,		
csaction <b>or</b> responder	ingress_default_4 -	44 <b>ອີ<u>r</u>ໝ8ts</b> end_80_csp_	26 <b>sepvice-pbstto®he</b> ygvrqrzpm4k beginning,added		
policy	frontend_default_	_80_svc	suffix of cs, hashed ingress-name,		
			ingress-port, and		
			namespace		
lbvserver (ingress)	k8s-web-	k8s-	Suffix lbv and		
	ingress_default_44 <u>8r</u> @stend_80_lbv_26 <b>cqmmeentkaddedtoytge</b> /rqrzpm4k				
	-	Comment:	entity		
	frontend_default_80nguceb-ingress				
		,ingport:5080,			
		ns: <b>default</b> ,svc:			
		frontend,			
		svcport:80			
servicegroup	k8s-web-	k8s-	Suffix sgp is added.		
(ingress)	ingress_default_4 -	44 <b>ፄ<u>r</u>ໝ8ቴend_80_sgp_</b>	267pneiak5rw6hoygvrqrzpm4k		
	frontend_default_	_80_svc			
lbvserver(UDP)	k8s-web-	k8s-bind_53-	-udp is still appended		
	ingress_default_9 -udp_k8s- bind_default_53	_ ∂0 <b>5</b> đp_lbv_uyomblbl	ag <b>toxhevportxas eadiea</b> k6wkpfmw		
	-udp_svc				

When you upgrade from an older version of the NetScaler Ingress Controller to the latest version, the NetScaler Ingress Controller renames all the entities with the new naming format. However, the NetScaler Ingress Controller does not handle the downgrade from the latest version to an older version.

# Licensing

December 31, 2023

For licensing the NetScaler CPX, you need to provide the following information in the YAML for the NetScaler Application Delivery Management (ADM) to automatically pick the licensing information:

- LS\_IP (License server IP) Specify the NetScaler ADM IP address.
- LS\_PORT (License server Port) This is not a mandatory field. You must specify the ADM port only if you have changed it. The default port is 27000.
- PLATFORM Specify the Platform License. Platform is CP1000.

The following is a sample yaml file:

1	apiVersion: apps/v1
2	kind: Deployment
3	metadata:
4	labels:
5	name: cpx-ingress
6	name: cpx-ingress
7	spec:
8	replicas: 1
9	selector:
10	matchLabels:
11	name: cpx-ingress
12	template:
13	metadata:
14	annotations:
15	NETSCALER_AS_APP: "True"
16	labels:
17	name: cpx-ingress
18	spec:
19	serviceAccountName: cpx
20	containers:
21	- args:
22	ingress-classes citrix-ingress
23	env:
24	- name: EULA
25	value: "YES"
26	- name: NS_PROTOCOL
21	
28	- name: NS_PURI
29	value: "9080"
30	- name: LS_IP
31	value: <adm ip=""></adm>
32	
33	
34 25	
30	Value. CP1000

36	image: cpx-ingress:latest
37	imagePullPolicy: Always
38	name: cpx-ingress
39	ports:
40	- containerPort: 80
41	name: http
42	protocol: TCP
43	- containerPort: 443
44	name: https
45	protocol: TCP
46	- containerPort: 9080
47	name: nitro-http
48	protocol: TCP
49	- containerPort: 9443
50	name: nitro-https
51	protocol: TCP
52	securityContext:
53	privileged: true

# net>scaler

© 2025 Cloud Software Group, Inc. All rights reserved. This document is subject to U.S. and international copyright laws and treaties. No part of this document may be reproduced in any form without the written authorization of Cloud Software Group, Inc. This and other products of Cloud Software Group may be covered by registered patents. For details, please refer to the Virtual Patent Marking document located at https://www.cloud.com/legal. Citrix, the Citrix logo, NetScaler, and the NetScaler logo and other marks appearing herein are either registered trademarks or trademarks of Cloud Software Group, Inc. and/or its subsidiaries in the United States and/or other countries. Other marks are the property of their respective owner(s) and are mentioned for identification purposes only. Please refer to Cloud SG's Trademark Guidelines and Third Party Trademark Notices (https://www.cloud.com/legal) for more information.

© 1997–2025 Citrix Systems, Inc. All rights reserved.