

NetScaler SDXTM

Contents

NetScaler SDX hardware-software compatibility matrix	15
How to read the table	16
Related information	19
SDX hardware platforms	20
NetScaler SDX 8900	20
NetScaler SDX 9100	22
NetScaler SDX 9200	24
Key specifications	24
Prepare for installation	25
Hardware components	25
Infrastructure requirements	25
Network requirements	25
Power requirements	26
Detailed AC input voltage specifications	26
Maximum input current requirements	26
Management station requirements	26
Hardware platform	27
Front panel components	27
Rear panel components	28
Install the hardware	29
Rack-mount the appliance	29
Network and power connections	32
Connection sequence	32
Power supply status indicators	33
System access and configuration	34
Console connection setup by using the GUI	34

Common default IP address configuration	35
NetScaler SDX Management Service and DOM0 Configuration	37
System specifications	40
Hardware configurations	40
Power and environmental specifications	41
Physical specifications	42
System LED status	43
Front panel LED status	43
Power supply LED status	43
NVMe storage LED status	44
Additional information resources	44
Information required for technical support	45
NetScaler SDX 17000	45
Key specifications	45
Prepare for installation	46
Infrastructure requirements	47
Network requirements	47
Power requirements	47
Detailed AC input voltage specifications	48
Maximum input current requirements	48
Management station requirements	48
Hardware platform	48
Front panel components	48
Rear panel components	50
Install the hardware	51
Rack-mount the appliance	51
Network and power connections	55
Connection sequence	55

Power supply status indicators	56
System access and configuration	57
SVM GUI access method	57
Common default IP address configuration	59
Console connection setup by using the CLI	60
NetScaler SDX Management Service sample configuration	61
System specifications	65
Hardware configurations	65
Power and environmental specifications	66
Physical specifications	67
System LED status	68
Front panel LED status	68
Power supply LED status	68
NVMe SSD tray LED status	69
Additional information resources	69
Technical support access	70
Information required for technical support	70
NetScaler SDX 14000	70
NetScaler SDX 14000-40G	73
NetScaler SDX 14xxx-40S	75
NetScaler SDX 14030 FIPS, 14060 FIPS, 14080 FIPS	76
NetScaler SDX 15000	79
NetScaler SDX 15000-50G	81
NetScaler SDX 16000	84
NetScaler SDX 25100A, SDX 25160A	87
NetScaler SDX 25100-40G	89
NetScaler SDX 26000	91
NetScaler SDX 26000-50S	93

NetScaler SDX 26000-100G	96
Common hardware components	99
LCD display and LED status indicators	99
LED port-status indicators	100
LED power supply indicators	102
Ports	102
RS232 serial port	103
Copper Ethernet ports	103
Management ports	103
1G SFP, 10G SFP+, and 40G QSFP+ ports	104
Ports compatibility	104
25G, 40G, 50G, and 100G ports	104
1G pluggable media	107
Copper 1G SFP distance specifications	107
Short reach fiber 1G SFP distance specifications	107
Short reach fiber 1G SFP distance specifications	108
Long reach fiber 1G SFP distance specifications	108
Long reach fiber 1G SFP distance specifications	109
10 GE pluggable media	109
Short reach fiber 10G SFP+ distance specifications	109
Long reach fiber 10G SFP+ distance specifications	110
Field replaceable units	111
Power supply	111
Electrical Safety Precautions for Power Supply Replacement	112
Replace an AC or a DC power supply	112
To install or replace an AC or a DC power supply on a NetScaler SDX appliance	112
Solid-State drives	114
Direct attach cable	114

To replace a DAC	115
Transceivers for MPX 9200 and 17000 platforms	116
Transceiver, DACs, AOCs Master Matrix	116
Platform port topology	116
Supported Transceivers, DACs & AOCs' matrix	117
Pre-Deployment checklist	126
Glossary	127
Transceiver, DAC, AOC FAQs	129
Speed & Form Factor	129
Q1: Are 1G SFP transceivers supported?	129
Q2: Which 10G transceivers are supported?	129
Q3: Which 25G transceivers are supported?	129
Q4: Is 50G supported?	129
Q5: Are 40G QSFP+ transceivers supported?	129
Q6: Are 100G/40G QSFP28 transceivers supported?	129
Q7: Is 200G supported?	130
Q8: Can I use a QSFP28 adapter at 25G or 10G with regular SFP optics on the MPX/ SDX 17000 200G QSFP56 ports?	130
Auto-Negotiation, Speed Forcing & FEC	130
Q1: Do ports auto-negotiate speed with these transceivers?	130
Q2: What FEC mode should I configure on the peer switch?	131
Q3: Why does not my 40G link come up with auto-negotiation?	131
Mixing & Compatibility	132
Q1: Can I use third-party optics or DAC/AOC cables?	132
Q2: Can I mix SR and LR (or different vendor optics) on each end of the link?	132
Q3: How do I choose SR vs LR?	132
Q4: What MPO polarity do I need for SR4 modules (40G/100G/200G)?	132
Verification & Troubleshooting	132

Q1: How do I verify that the transceiver is detected and the link is up?	132
Q2: Can I use a Cisco Nexus N9K with non-Cisco optics?	133
Q3: My 9530413 (Citrix MPN) 100G LR4 link is not coming up over OS2 fiber to a Mellanox SN3700 or Cisco N9K switch. The transceivers are seated and powered. What should I check first?	133
Q4: What is the QSFP28-to-SFP28 adapter 9530416 (Citrix MPN) and when do I use it?	134
Q5: When would I use the QSFP28-to-SFP28 adapter 9530416 (Citrix MPN) instead of just using an SFP56 port directly?	134
Q6: Is the QSFP28-to-SFP28 adapter 9530416 (Citrix MPN) supported on the MPX/SDX 9200?	134
Q7: Does the QSFP28-to-SFP28 adapter 9530416 (Citrix MPN) add latency or signal loss?	135
Transceiver Interop Use Cases	135
Citrix MPN: 9530409 (200G QSFP56 SR4 MPO-12)	135
Use Case: Native 200G to Mellanox SN3700	135
Citrix MPN: 9530410 (100/40G QSFP28 SR4 MPO-12)	135
Use Case 1: 100G Native Link to Cisco Nexus N9K	135
Use Case 2: 40G Forced Link to Cisco Nexus N9K (Mixed-Speed Interop)	135
Use Case 3: 100G Native Link to Mellanox SN3700 (Spectrum-2)	136
MPN: 9530411 (40G QSFP+ BiDi SR LC)	136
Use Case 1: BiDi Link with 9530411 on Both Ends (Stable)	136
Use Case 2: Mixed BiDi Pairing (Link Flap Observed)	137
Citrix MPN: 9530413 (100G QSFP28 LR4 LC, 10 km)	137
Use Case 1: 100G LR4 to Cisco Nexus N9K	137
Use Case 2: 100G LR4 to Mellanox SN3700	137
Citrix MPN: 9530416 (QSFP28-to-SFP28 Adapter)	137
Use Case 1: 25G SFP28 SR Uplink via QSFP56 Port	137
Use Case 2: 10G SFP+ SR Uplink via QSFP56 Port	138
Citrix MPN: 9530513 (40G QSFP+ DAC, 5 m)	138

Use Case 1: 40G QSFP+ DAC to Cisco Nexus N9K	138
Use Case 2: 40G QSFP+ DAC to Mellanox SN3700	138
Citrix MPN: 9530517 (100G QSFP28 AOC, 20 m)	139
Use Case 1: 100G AOC to Cisco Nexus N9K	139
Use Case 2: 100G AOC to Mellanox SN3700	139
Managing RAID disk allocation on the SDX platform	140
View drive properties and operations	140
Storage repository	140
Logical drives	141
Create a new logical drive	142
Add a logical drive to the SDX Platform	142
Add a second logical drive on the SDX Platform	142
Physical drives	143
Raid support	144
Replace a defective RAID-supported SSD drive with a blank SSD drive	144
Replace a defective RAID logical Drive with two blank SSD Drives	145
Safety, cautions, warnings, and other information	146
Prepare for installation	147
Unpack the appliance	147
Prepare the site and rack	148
Site requirements	148
Rack requirements	148
Electrical safety precautions	149
Install the hardware	150
Rack-mount the appliance	150
Remove the inner rails from the rail assembly	151
Attach the inner rails to the appliance	151
Install the rack rails on the rack	151

Install the appliance in the rack	153
Install and remove 1G SFP transceivers	154
Install a 1G SFP transceiver	155
Remove a 1G SFP transceiver	156
Install and remove 10G SFP+ transceivers	156
Install a 10G SFP+ transceiver	156
Remove a 10G SFP+ transceiver	157
Connect the cables	157
Connect the Ethernet cables	157
Connect an Ethernet cable to a 10/100/1000BASE-T port or 1G SFP copper transceiver	158
Connect the Ethernet cable to a 1G SFP fiber, 10G SFP+ transceiver	158
Connect the Console cable	158
Connect the console cable to a computer or terminal	158
Connect the power cable	159
Connect the appliance to the power source	159
Switch on the appliance	160
To switch on the appliance	160
Initial configuration	161
Initial configuration through the Management Service user interface	162
To configure the NetScaler SDX appliance by using the Management Service user interface	162
Initial Configuration through the serial console	165
To configure the NetScaler SDX appliance by using the serial console	165
Change the password of the default user account	166
Change the administrative password by using the GUI	166
Lights out management port of the NetScaler SDX appliance	167
LOM support matrix	168
Upgrade LOM version to 3.11.0	169
Upgrade procedure	169

Configuration compatibility	169
Configure the network settings on the LOM port	170
To configure the NetScaler LOM port by using the GUI	171
To configure the NetScaler LOM port by using the shell	172
Obtain health monitoring information	173
To obtain health monitoring information	174
Install the MIB	174
Configure SNMP Alerts	174
To configure SNMP alerts on the LOM by using the GUI	175
Set up SNMP alerts on the LOM by using the NetScaler shell	175
Create and install an SSL certificate on SDX 9100 and SDX 16000	176
Prerequisites	176
Create an SSL certificate by using OpenSSL	176
Install an SSL certificate by using the LOM GUI	177
Example	177
Sample Private_Key.pem (2048 bits)	181
Sample request.csr	183
Sample lom_certificate.pem	184
Install a certificate and key on older platforms using the LOM GUI	185
To upload a trusted certificate and private key by using the LOM GUI	185
Obtaining the MAC address, serial number, and host properties of the appliance	190
To obtain the MAC address, serial number, and host properties of the appliance by using the LOM GUI	191
To obtain the MAC address and host properties of the BMC by using the appliance shell	191
Example	191
Perform power control operations by using the LOM port	193
Perform power control operations by using the GUI	193
Perform a power cycle of the BMC	194

Perform power control operations on the BMC by using the appliance shell	194
Perform a warm restart remotely from another computer on the network	194
Perform a cold restart of the LOM from the appliance	194
Perform a warm restart remotely from another computer on the network	194
Perform a core dump	194
Perform a core dump by using the GUI	194
Perform a core dump remotely from another computer on the network by using the shell	194
Restoring the BMC configuration to factory defaults	195
To reset the configuration to factory defaults by using the GUI	195
To reset the configuration to factory defaults by using the shell	195
Specifying the port for IPMI BMC failover	195
Using the BIOS POST code to detect errors	196
To read the BIOS Post Code by using the LOM GUI	196
To read the BIOS POST code by using the shell	196
LOM port management for SDX 9200 and SDX 17000	196
Prerequisites	196
Initial network access	197
Default LOM settings	197
Network connection methods	197
Configure LOM by using the GUI	197
Initial login and password change	197
Network configuration	198
System lockdown mode restriction	200
Symptoms of lockdown mode	200
Affected functions	200
Resolution	200
Configure LOM by using the CLI (IPMI tool)	200
Access requirements	200

System lockdown management	201
Static IP configuration	201
Verify configuration	202
Display current LOM network settings	202
Test network connectivity	202
User password management	202
LOM factory reset	204
Remote system power control	204
Power on procedure	204
Graceful shutdown procedure	205
Troubleshooting	207
Cannot access the default IP address	207
Login fails with default credentials	207
BMC not responding after configuration	208
Best practices for LOM configuration and management	208
Security best practices	208
Power control best practices	209
Hardware health attributes of NetScaler SDX appliances	211
Health attributes for the SDX 9100 platform	211
Health attributes for the SDX 8900 platform	213
Health attributes for the SDX 15000 and SDX 26000 platforms	213
Health attributes for the SDX 14000 platform	215
Health attributes for the SDX 22040 platform	216
Health attributes for SDX 24100 platform	220
Hardware health attributes for the SDX 16000 platform	225
Sample output	226
CLI	227
Shell	230

LOM GUI	232
Fortville NIC firmware upgrade on NetScaler SDX appliances	234
Supported platforms	234
Support for updated NIC driver and firmware on NetScaler SDX with 10G/25G/40G NICs	235
Converting a NetScaler MPX appliance to a NetScaler SDX appliance	236
Convert a NetScaler MPX 26000-50S appliance to a NetScaler SDX 26000-50S appliance	237
Verify proper operation of the MPX appliance's components	238
Upgrade the appliance	269
Reconfigure the converted appliance	272
Converting a NetScaler MPX 24100 and 24150 appliance to a NetScaler SDX 24100 and 24150 appliance	273
To verify proper operation of the MPX appliance's components	275
To upgrade the appliance	275
Convert a NetScaler MPX 14020/14030/14040/ 14060/14080/14100 appliance to a NetScaler SDX 14020/14030/14040/ 14060/14080/14100 appliance	277
To verify proper operation of the MPX appliance's components	278
To upgrade the appliance	284
To reconfigure the converted appliance	286
Convert a NetScaler MPX 14020 40G, MPX 14040 40G, MPX 14060 40G, MPX 14080 40G appliance to a NetScaler SDX 14020 40G, SDX 14040 40G, SDX 14060 40G, SDX 14080 40G appliance	287
Verify proper operation of the MPX appliance's components	288
Upgrade the appliance	315
Reconfigure the converted appliance	317
Convert a NetScaler MPX 8900 appliance to a NetScaler SDX 8900 appliance	318
To verify proper operation of the MPX appliance's components	319
Upgrade the appliance	330
Replace the SSD	330
Start the appliance	331

Convert a NetScaler MPX 15000 appliance to a NetScaler SDX 15000 appliance	332
Verify proper operation of the MPX appliance's components	334
Upgrade the appliance	358
Reconfigure the converted appliance	361
Migrate the configuration of a NetScaler instance across SDX appliances	363
Prerequisites	365
Supported migration scenarios	366
Case 1: No channel and same interface speed and numbering	366
Case 2: No channel and different interface speed or numbering	366
Case 3: Same channel numbering	366
Case 4: Different channel numbering	367
Phases in migration	367
Phase 1: Break HA (if applicable) and remove the secondary VPX	367
Phase 2: Form an HA pair between the old primary VPX and the new VPX	368
Phase 3: Migrate VPX configuration	368
Phase 4 (optional): Additional steps required for case 2 and case 4	368
Phase 5: Force failover to migrate the traffic to the new VPX on the new SDX appliance	369
Troubleshooting	370
NetScaler SDX hardware document history	371
fix	374

NetScaler SDX hardware-software compatibility matrix

Keep the following points in mind about NetScaler SDX hardware-software and feature-software compatibility:

- For each platform, only the starting build for each release is mentioned. All subsequent builds within a release are supported.
- In 10.5 and prior releases, the SDX appliance was shipped with a virtual bundle. The bundle included a specific combination of Management Service image, Citrix Hypervisor image, Citrix Hypervisor supplemental packs.
- From 11.0 and later releases, the single bundle image is shipped, which combines all the components.
- Any combinations of SDX image version and VPX image versions are supported as long as they meet minimum version requirements for a particular hardware platform.
- From 12.0 57.19, the single bundle image also includes the lights-out management (LOM) firmware.
- Direct upgrade from version 10.5 to version 13.1 is not supported. First upgrade from 10.5 to 11.1 or 12.0 or 12.1, and then upgrade to SDX 13.1.
- Usually, the SDX image and the VPX image are not interdependent, which means the SDX appliance and the VPX instances don't need the same image version.

Exception: If you provision a new VPX instance or upgrade an existing VPX instance to 12.0-58.15, you must also upgrade the SDX version to 12.0-58.15. See [<https://support.citrix.com/article/CTX236730>](https://support.citrix.com/article/CTX236730).

If you use NetScaler ADM to manage NetScaler appliances running 12.0–58.15, you must upgrade NetScaler ADM to 12.0-58.15.

The following tables list the compatibility matrix for all NetScaler SDX hardware platforms and the SDX and VPX images supported on these platforms.

How to read the table

The SDX image version and the VPX image version are listed for each platform. The SDX image and the VPX image are not interdependent, which means the SDX appliance and the VPX instances don't need the same image version. Check the note for exception to compatibility.

IMPORTANT

The first supported build for each hardware platform and software release are listed in the following table. **All subsequent builds are supported unless the word "only" follows the build number. When it does, only the specified build is supported on that platform.**

HARDWARE PLATFORM S/SDX SBI VERSION AND VPX IMAGE VERSION	11.1	12.1	13.0	13.1	14.1
SDX 8015	11.1–47.14	12.1–48.13	13.0–36.x	13.1–4.x	14.1–4.x
SDX 8400/8600	11.1–47.14	12.1–48.13	13.0–36.x	13.1–4.x	14.1–4.x
SDX 8920/8930	11.1–57.11 (See note after the table)	12.1–48.13	13.0–36.x	13.1–4.x	14.1–47.x
SDX 9100	X	X	X	13.1–21.50	14.1–4.x
SDX 9200	X	X	X	X	14.1–72.x
SDX 11500/13500/14500/16500/18500/20500	11.1–47.14	12.1–48.13	13.0–36.x	13.1–4.x	X
SDX 11515/1152	11.1–47.14	12.1–48.13	13.0–36.x	13.1–4.x	X

HARDWARE PLATFORM S/SDX SBI VERSION AND VPX IMAGE VERSION	11.1	12.1	13.0	13.1	14.1
0/11530/115 40/11542					
SDX 14020/1403 0/14040/140 60/14080/14 100	11.1–47.14	12.1–48.13	13.0–36.x	13.1–4.x	14.1–29.x
SDX 14020-40G/ 14040-40G/ 14060-40G/ 14080-40G/ 14100-40G	11.1–47.14	12.1–48.13	13.0–36.x	13.1–4.x	14.1–29.x
SDX 14040-40S/ 14060-40S/ 14080-40S/ 14100-40S	11.1–48.10	12.1–48.13	13.0–36.x	13.1–4.x	14.1–29.x
SDX 14030 FIPS/14060 FIPS/14080 FIPS	11.1–52.13	12.1–48.13	13.0–36.x	13.1–4.x	14.1–29.x
SDX 15020/1503 0/15040/150 60/15080/15 100/15120	11.1–63.x	12.1–56.x	13.0–52.x	13.1–4.x	14.1–4.x
SDX 15020-50G/ 15030-50G/	11.1–59.10	12.1–48.13	13.0–36.x	13.1–4.x	14.1–4.x

HARDWARE PLATFORM S/SDX SBI VERSION AND VPX IMAGE VERSION	11.1	12.1	13.0	13.1	14.1
15040-50G/ 15060-50G/ 15080-50G/ 15100-50G/ 15120-50G					
SDX 16000	X	X	X	13.1–37.x	14.1–4.x
SDX 17000	X	X	X	X	14.1–72.x
SDX 22040/2206 0/22080/221 00/22120	11.1–47.14	12.1–48.13	13.0–36.x	13.1–4.x	14.1–4.x
SDX 24100/2415 0	11.1–47.14	12.1–48.13	13.0–36.x	13.1–4.x	14.1–4.x
SDX 25100A/ 25160A/ 25200A	11.1–51.21	12.1–48.13	13.0–36.x	13.1–4.x	14.1–4.x
SDX 25100-40G/ 25160-40G/ 25200-40G	11.1–51.21	12.1–48.13	13.0–36.x	13.1–4.x	14.1–4.x
SDX 26100/2616 0/26200	11.1–63.9	12.1–56.x	13.0–52.x	13.1–4.x	14.1–4.x
SDX 26100-100G /	11.1–58.13	12.1–48.13	13.0–36.x	13.1–4.x	14.1–4.x

HARDWARE PLATFORM S/SDX SBI VERSION AND VPX IMAGE VERSION	11.1	12.1	13.0	13.1	14.1
26160-100G / 26200-100G					
SDX 26100-50S/ 26160-50S/ 26200-50S	11.1–63.9	12.1–56.x	13.0–52.x	13.1–4.x	14.1–4.x

NOTES:

- On the NetScaler SDX 8900 appliance, the 11.1 56.15 VPX image is supported. SDX 11.1 56.15 image is not supported.
- NetScaler Secure Web Gateway (SWG) is supported from release 12.0 56.20 and later.
- NetScaler cluster is not supported on NetScaler SDX SWG instances.

Related information

For information about how to upgrade NetScaler SDX and the VPX instance, see [Single bundle upgrade for 12.1 release](#).

SDX hardware platforms

The various NetScaler hardware platforms offer a wide range of features, communication ports, and processing capacities. All platforms have multicore processors.

The NetScaler SDX appliance is a multitenant platform on which you can provision and manage multiple virtual instances. All NetScaler appliances have front to rear airflow.

For the updated list of current platforms, click [Product matrix](#) and select **Citrix ADC**.

NetScaler SDX 8900

The NetScaler SDX 8900 appliance is a 1U appliance. This platform has a single 8-core processor and 32 GB of memory. The appliance provides a total of 10 network ports:

- Six 10/100/1000Base-T RJ45 copper Ethernet Ports.
- Four 10G SFP+ Ethernet Ports.


1G copper transceivers are supported in 10G slots. 1G fiber transceivers are not supported.

The appliance comes with two VPX instances and supports up to seven instances. You can separately buy extra five instances.

For information on the software releases supported on the NetScaler hardware platforms, see [NetScaler SDX Hardware-Software Compatibility Matrix](#).

The figures are for illustrative purposes only.


Figure 1. NetScaler SDX 8900 front panel



The NetScaler SDX 8900 appliances have the following ports:

- RS232 serial Console Port.
- One 10/100/1000Base-T RJ45 copper Ethernet LOM Port. Use this port to remotely monitor and manage the appliance independently of the NetScaler software.
- One 10/100/1000Base-T RJ45 copper Ethernet Management Port, numbered 0/1. This port is used to connect directly to the appliance for NetScaler administration functions.
- Six 10/100/1000Base-T RJ45 copper Ethernet Ports, numbered 1/1 to 1/6 from left to right.
- Four 10G SFP+ Ethernet Ports, numbered 10/1 to 10/4 from left to right.
- USB port (reserved for a future release).

Figure 2. NetScaler SDX 8900 back panel



The following components are visible on the back panel of the SDX 8900 appliances:

- One 240 GB or larger removable solid-state drive (SSD).

Note: Drive densities might increase as components become EOL but its size is never smaller than the original.

- Power switch, which turns power to the appliance on or off.
 - If the OS is functional, press the switch for less than two seconds to power down the system with a graceful shutdown.
 - If the OS is not responsive, press the power switch for more than 4 seconds to force the power off.
- One power supply, 100–240 VAC (second power supply for redundancy is a customer installable option). Each power supply has an LED indicating its status:

LED COLOR	LED INDICATES
OFF	No power to any power supply in the appliance.

LED COLOR	LED INDICATES
Flashing RED	No power to this power supply.
Flashing GREEN	Power supply is in standby mode.
GREEN	Power supply is functional.
RED	Power supply failure.

- **Disable alarm button**, which is functional only when the appliance has two power supplies. Press this button to silence the power alarm when one of two power supplies loses input power (second power supply optional) or when a power supply is malfunctioning.
- **Non-Maskable Interrupt (NMI) Button**, used at the request of Technical Support to initiate a core dump. To press this red button, which is recessed to prevent unintentional activation, use a pen, pencil, or other pointed object. The **NMI Button** is also available remotely over the network in the LOM GUI, in the **Remote Control** menu. For more information, see [Lights out management port of the NetScaler SDX appliance](#).


NetScaler SDX 9100

The NetScaler SDX 9100 appliance is a 1U appliance. This platform has a single 10-core processor and 64 GB of memory. The appliance provides a total of eight 25G SFP28 ports.

For information on the software releases supported on the ADC hardware platforms, see [Hardware-Software Release Matrix](#).

The following figure shows the front panel of the SDX 9100 appliance.

Figure 1. NetScaler SDX 9100, front panel




The NetScaler SDX 9100 series appliances have the following ports:

- RS232 serial Console port.
- One 10/100/1000Base-T RJ45 copper Ethernet LOM port. Use this port to remotely monitor and manage the appliance independently of the ADC software.
- One 10/100/1000Base-T RJ45 copper Ethernet management port, numbered 0/1. This port is used to connect directly to the appliance for ADC administration functions.
- Two USB ports (reserved for a future release).
- Eight 25G SFP28 ports, numbered 25/1 to 25/8. For information about supported transceivers per port, see [25G, 40G, 50G, and 100G ports](#).

The following figure shows the back panel of the SDX 9100 appliance.

Figure 2. NetScaler SDX 9100, back panel



The following components are visible on the back panel of the SDX 9100 appliance:

- One 480 GB removable solid-state drive (SSD).

Note: Drive densities might increase as components become EOL but its size is never smaller than the original.

- Power switch, which turns power to the appliance on or off.
 - If the OS is functional, press the switch for less than two seconds to power down the system with a graceful shutdown.
 - If the OS is not responsive, press the power switch for more than 4 seconds to force the power off.
- Two power supplies (for hardware shipped after 13th Dec, 2023), rated at 450 watts, 100–240 VAC. Maximum power consumption is 275 watts and typical power consumption is 225 watts. The following table indicates the LED status of each power supply:

LED COLOR	LED INDICATES
OFF	No power to any power supply in the appliance.
Flashing RED	No power to this power supply.
Flashing GREEN	Power supply is in standby mode.
GREEN	Power supply is functional.
RED	Power supply failure.

- **Disable alarm button**, which is functional only when the appliance has two power supplies. Press this button to silence the power alarm when one of two power supplies loses input power (second power supply optional) or when a power supply is malfunctioning.
- **Non-Maskable Interrupt (NMI) Button**, used at the request of Technical Support to initiate a core dump. To press this red button, which is recessed to prevent unintentional activation, use a pen, pencil, or other pointed object. The NMI Button is also available remotely over the network in the LOM GUI, in the **Remote Control** menu. For more information about the lights out management port of the appliance, see [Lights out management port of the NetScaler SDX appliance](#).

NetScaler SDX 9200

Welcome to the next-generation NetScaler SDX 9200 platform. This section provides comprehensive instructions on how to deploy your appliance efficiently and configure it for optimal performance.

The SDX 9200 1U platform represents the latest advancement in NetScaler technology, featuring enhanced processing power, improved connectivity options, and robust management capabilities.

Key specifications

- Intel Xeon 5th-generation processor architecture
- 128 GB DDR5 ECC memory for enhanced performance
- Dual 4 x 50G high-speed ports with PCIe 4.0 support
- Redundant CRPS power supplies for high availability
- Dedicated BMC for comprehensive out-of-band management

Prepare for installation

Before installing your new appliance, carefully unpack the unit and verify that all the components are present. The only tools required for installing the appliance are a Phillips screwdriver, a flathead screwdriver, and a cage nut insertion tool (if installing into square-hole racks). Your NetScaler 1U package should include the following items:

Hardware components

- 1x NetScaler SDX 9200 appliance
- 1x standard 4-post rail kit (19-inch rack compatible), which includes:
 - Left rail (x1)
 - Right rail (x1)
 - Mounting screw kit (includes M4x4L, M4x6L, M3x6L, M5x20L screws)
 - Rackmount Ear Kit (x1)

The included accessory kit contains essential connectivity hardware, including two power cords to support redundant power supply operation.

NOTES:

- Rail kit compatibility varies by rack manufacturer and configuration. If the provided rail kit does not fit your specific rack infrastructure, contact your NetScaler sales representative immediately to obtain an appropriate alternative mounting solution.
- Before proceeding, examine all components for shipping damage. Report any damaged items to your carrier and NetScaler support before beginning the installation process.

Infrastructure requirements

- Compatible 19-inch rack with adequate depth clearance (minimum 550 mm)
- Proper airflow management with front-to-rear cooling path
- Available rack units (1U height requirement)
- Adequate weight support (system weight: 25.2 lbs)

Network requirements

- Administrative LAN with available switch ports

- Dedicated out-of-band management network
- Appropriate network cables (Cat5e/Cat6 for management, fiber/DAC for data)

Power requirements

- Dual power sources (recommended for redundancy)
- Frequency operating range: 47 Hz – 63 Hz
- Adequate power capacity (full load 110 V: 406 W/240 V: 394.57 W)
- PSU input connector: IEC320-C14 (Male)
- Required power cord: IEC320-C13 to NEMA 5-15P (or regional equivalent)
- Rating: Minimum 10 A / 250 V

Detailed AC input voltage specifications

VOLTAGE RANGE	MINIMUM	RATED OPERATING RANGE	MAXIMUM	BROWN OUT	BROWN IN
115 V Range	90 Vrms	100-127 Vrms	140 Vrms	75 Vac+/-4 V	85 Vac+/-4 V
230 V Range	180 Vrms	200-240 Vrms	264 Vrms	N/A	N/A

Maximum input current requirements

INPUT VOLTAGE RANGE	MAXIMUM CURRENT
100 Vac – 127 Vac	10 A
200 Vac – 240 Vac	5 A

Management station requirements


- Workstation or laptop (Windows or macOS) with network connectivity
- Web browser with JavaScript and HTML5 support (Chrome, Firefox, Safari, or Edge)
- Terminal emulator software for CLI access
- Administrative privileges for network configuration

Hardware platform

Front panel components

The 1U front panel integrates comprehensive connectivity and management interfaces designed for enterprise data center environments. Each component serves a specific function in system operation and management.

The following figure shows the front panel of the SDX 9200 appliance.



LABEL	COMPONENT TYPE	TECHNICAL SPECIFICATION	PRIMARY FUNCTION
CONSOLE	RJ-45 Port	RS232 compatible	Primary console interface for direct system access
MGMT 0/1 & 0/2	Dual Ethernet Ports (1000 Base-T)	1 GbE Ethernet	NetScaler management network connectivity
LOM	Dedicated IPMI port	Intelligent Platform Management	Out-of-band monitoring and control system hardware
NMI	Pinhole button	Non-maskable interrupt	Kernel core dump generation
PWR	Soft-Push button 4-second holds for on/off	Power management	System power control with graceful shutdown

LABEL	COMPONENT TYPE	TECHNICAL SPECIFICATION	PRIMARY FUNCTION
TYPE-C	Optional Port	USB Type-C	Alternate console and debug interface
UID	Locator LED	Blue LED indicator	Remote identification capability
GPIO1/GPIO2	User-defined GPIOs	General purpose I/O	Currently disabled in 1U platform design
50/1–50/8	Data Ports	50GbE network interface card NIC	Quad-port SFP56 with PCIe 4.0 x16

Console port configuration: The primary console port eliminates the need for legacy RS232 adapters by providing direct RJ-45 connectivity, which simplifies cable management and improves connection reliability.


Management port operation: Dual management ports support link aggregation and provide necessary redundancy for secure administrative access.

Data Port Capabilities: High-performance data ports deliver 50GbE per port and support various transceiver types, including SFP56 modules.

Rear panel components

The rear panel houses critical infrastructure components, including dual power supplies, the cooling system, a UID indicator, a PSU alarm, and a secondary power control button.

The following figure shows the back panel of the SDX 9200 appliance.



COMPONENT	TECHNICAL DETAILS	OPERATIONAL CHARACTERISTICS	SERVICE INFORMATION
PSU1 / PSU2	CRPS PSU	Redundant N+1 configuration	Hot-swappable during operation
2nd UID	Locator LED	Blue LED indicator	Remote identification capability
Rear Power Button	Optional control	Mirrors front button function	Dual power control access
Fan Modules	5x system fans	Smart fan control by using the BMC	No field replaceable Fan FRU

Power supply architecture: The Common Redundant Power Supply (CRPS) design ensures continuous operation in the event of a single power supply failure. Each unit is hot-swappable, enabling maintenance without system downtime.

Cooling system design: The intelligent fan control system automatically adjusts cooling based on internal temperatures and current load conditions. Five fan modules provide N+1 redundancy, ensuring simplified serviceability.

Identification system: The Unit Identifier (UID) LED facilitates rapid system identification in dense rack environments. It supports both local activation and remote management through the system interface.


Install the hardware

Rack-mount the appliance

Perform the following procedure to ensure proper appliance mounting and system performance.

Remove the inner rails from the rail assembly

1. Pull the white disconnect tab to release and detach the inner rail member from each slide rail.
2. With the stamped text facing outward, align the inner rail with the chassis mount points, using the black reference markers on both sides of the chassis.
3. Slide the rail toward the rear until it clicks firmly into the locked position.
4. Fasten each inner rail using the provided M4x4L screws, as shown in the diagram.
5. Repeat this process for both the left and right sides of the chassis.




Chassis mounting ear installation

1. Place the rack-mount ears against the left and right sides of the chassis.
2. Align each ear with the designated screw holes on the front of the chassis.
3. Fasten each mounting ear using three M3x6L screws per side.
4. Ensure that the ears are flush against the chassis and properly aligned.
5. Confirm that the mounting ears do not obstruct front-panel ports or LED indicators.

NOTE:

The left and right-side rails are symmetrical. Repeat the installation steps for both sides as illustrated in the diagrams.




Outer slide rail assembly installation

1. Hold the outer rail assembly against the rack frame.
2. Align the mounting brackets with the rack post holes at the selected U-height.
3. Insert the rail mounting mechanism into the rack posts until it engages.
4. Listen for an audible "click" to confirm that the rail is properly seated and locked.

5. Ensure both the front and rear mounting points are securely attached to the posts.
6. Verify that the rails are level and parallel using the numbered rack post markings.
7. Repeat this process for both the left and right outer rail assemblies.

Mount the chassis into the rack

1. Carefully align the chassis-mounted inner rails with the rack-mounted outer rails.
2. Smoothly slide the system into the rack until the locking mechanism engages.
3. If necessary, apply light, even pressure to ensure the unit locks securely into place.



Safety requirements

For safety and equipment durability, ensure that the following requirements are met:

1. Maintain a minimum of 2 inches of clearance at the front and rear of the chassis to ensure proper ventilation.
2. Verify the rack's weight capacity before installation and ensure that the rack is stable and properly secured to the floor.
3. Use proper lifting techniques. The system weight is approximately 25.2 lbs.

4. Strictly follow electrical safety procedures when connecting the appliance to a power source.

Airflow management

The 1U platform is designed for front-to-rear airflow. Unobstructed passage is required for optimal cooling performance. Restricted airflow results in thermal overheating and the automatic activation of system protection mechanisms (throttling or shutdown).

Network and power connections

Connection sequence

Perform the following sequence to ensure proper system initialization and to prevent configuration conflicts.

1. Management network connectivity

- a. Connect the 0/1 MGMT port to your administrative LAN by using an appropriate Ethernet cable.
- b. Ensure that the network switch port is configured to support auto-negotiation.
- c. Verify network connectivity and ensure that the correct VLAN configuration is applied, if applicable.
- d. Document the assigned IP address and network parameters for future reference.

2. Out-of-Band (OOB) management setup

- a. Connect the LOM port to your dedicated OOB management network.
- b. Use network infrastructure that is physically or logically separate from production traffic.
- c. Configure the port for dedicated mode operation (default).
- d. Ensure management network isolation for security.

3. Data port configuration

- a. Insert the appropriate SFP+, SFP28, or SFP56 transceiver modules into the data ports.
- b. Ensure that transceiver specifications align with your existing network infrastructure requirements.
- c. Connect the fiber optic or direct cables to the data ports.
- d. Confirm that the cable specifications support the required port speed (example: 50GbE).

4. Power system connection

- a. Plug the first power cord into the PSU1 input receptacle.
- b. Plug the second power cord into the PSU2 input receptacle.
- c. Ensure that the power source meets requirements (100–240 VAC, 50–60 Hz).
- d. To ensure power redundancy, connect each PSU to a separate power circuit.

5. System startup

- a. Press the front-panel power button to begin the boot sequence.
- b. Observe the front-panel LED indicators during the startup process.
- c. Verify that both power supply LEDs show a solid green state, indicating normal operation.
- d. Allow the system to complete the full boot cycle before attempting to access the management interface.

Power supply status indicators

The CRPS power supplies provide comprehensive status monitoring through integrated LED indicators located on each unit's faceplate.

PSU LED indicator status:

LED INDICATOR	PSU STATE	SYSTEM IMPACT	ACTION REQUIRED
Solid GREEN	AC present, all DC rails active under load	Normal operation	None - normal
1 Hz Blink GREEN	AC present, 12Vsb active including 5Vsb for BMC power, main rails idle	Redundant standby (N+1)	None - normal
Solid AMBER	AC input lost or disconnected	Lost redundancy (if N+1) or system down	Check AC power connection
1 Hz Blink AMBER	OCP/OVP/OTP warning threshold reached	Degraded PSU operation/Potential shutdown	Verify PSU cooling and airflow
Solid AMBER	Critical Fault (OCP/OVP/OTP tripped)	PSU shutdown for protection	Replace the PSU immediately

Monitoring recommendations:

- Regularly inspect PSU LED indicators during system operation.
- Document all occurrences of amber LED states to assist in long-term reliability and trending analysis.
- Replace modules with blinking amber status indicators during the next scheduled maintenance cycle.
- Contact customer support for any persistent warning conditions.

System access and configuration

Console connection setup by using the GUI

Use this method for initial deployment and ongoing administration.

Platform-specific access information

- Management through virtualized service interfaces
- Management Service VM: 192.168.100.1 with 255.255.0.0 (/16) subnet mask
- XenServer hypervisor: 192.168.100.2 with 255.255.0.0 (/16) subnet mask
- NSIP Pool Range: 192.168.100.10 to 192.168.100.100 for NetScaler instances

Perform the following steps to initialize the appliance

1. Network preparation
 - a. Connect your management workstation to the MGMT 0/1 port or a port on the same management VLAN.
 - b. Configure your workstation network settings (if required).
 - c. Verify the network path by performing a ping test to the default NSIP address: 192.168.100.1.
 - d. Ensure that any local firewalls allow HTTPS (port 443) traffic to the appliance.
2. Browser initialization
 - a. Open a web browser (Chrome, Firefox, Safari, or Edge)
 - b. Enter the URL `https://192.168.100.1` (Management Service VM).

NOTE:

Accept security certificate warnings (self-signed certificate)

3. Initial authentication

- a. Enter the default administrative user: `nsroot`.
- b. Enter the system serial number exactly as it appears on the physical chassis label (case-sensitive).
- c. Click **Login** or press **Enter**.

On successful authentication, the system redirects you to the **Initial Configuration Wizard**.

4. Setup wizard configuration

The **Setup Wizard** guides you through the following essential parameters:

- a. NSIP address configuration: Set a NetScaler IP address for management access.
- b. Network parameters: Configure the subnet mask and default gateway.
- c. SNIP configuration: Add subnet IP addresses for client communication.
- d. DNS settings: Configure DNS server addresses for name resolution.
- e. Hostname setup: Assign a unique, descriptive hostname to identify the appliance on the network.
- f. Administrative password: Change the default password to a strong, secure administrative password.

Common default IP address configuration

INTERFACE OR SERVICE	DEFAULT USER LOGIN CREDENTIALS	DEFAULT IP ADDRESS	SUBNET MASK OR GATEWAY	DESCRIPTION
Management Service (SVM) (0/1)	nsroot / product SN	192.168.100.1	255.255.0.0/16	Management Service is the management interface for NetScaler SDX like provisioning

INTERFACE OR SERVICE	DEFAULT USER LOGIN CREDENTIALS	DEFAULT IP ADDRESS	SUBNET MASK OR GATEWAY	DESCRIPTION
				and monitoring VPX instances
XenServer (DOM0)	root / product SN	192.168.100.2	255.255.0.0 /16 GW: 192.168.0.1	DOM0 is the hypervisor control domain (based on Linux) that interfaces with system hardware. Note: The product S/ N as the default password for DOM0 is only valid after Management Service initialization.
XenServer (Management Service virtual port 0/2)	nsroot / product SN	169.254.0.10	255.255.0.0 /16	Management Service internal default IP for 1U platform
LOM or IPMI Port	nsroot / product SN	192.168.1.3	255.255.255.0/ 24	Out-of-band management IPMI


Navigation best practices

- Use tab-based navigation to move quickly between input fields.
- Hover your cursor over field labels or information icons to view additional help information.
- Ensure that you save your changes before navigating to a different configuration section.
- Avoid using the browser's back button, as it might result in the loss of unsaved configuration data or session timeouts.

NetScaler SDX Management Service and DOM0 Configuration

1. Management Service shell access

a. SDX Management Service initial network configuration



b. SSH connection

- i. From your management laptop, initiate an SSH session to the Management Service IP address : `ssh nsroot@192.168.100.1`
- ii. Authenticate using the username as nsroot and password as the product serial number.
- iii. Log in to the NetScaler CLI prompt, run the `shell` command.

You can verify that you have successfully gained access when the prompt changes to `bash-3.2#`.

2. Configuration file verification

- a. **Access the configuration file:** From the `bash-3.2#` prompt, run the `cat /flash/mpsconfig/svm.conf` command to display the contents of the Management Service configuration file.
- b. **Review default contents:** Ensure the output matches the standard default settings, which include the following parameters:

Default svm.conf output:

- `arp -d -a`
- `route flush`
- `ifconfig 0/1 192.168.100.1 netmask 255.255.0.0`
- `ifconfig 0/2 169.254.0.10 netmask 255.255.0.0`
- `/mps/changenameserver.sh 127.0.0.2`

- c. **Verify XenServer Connectivity:** Confirm the XenServer IP address is listed at the end of the file as: `#XenServer : 192.168.100.2`.

3. Network configuration modification

- a. **Launch the configuration utility:** From the bash shell environment, initiate the interactive network setup tool by running the command: `/mps/networkconfig`.

b. **Interactive configuration menu:** After launching the utility, the following menu appears:

1. Management Service Host Name [ns]
2. Management Service IPv4 address [192.168.100.1]
3. XenServer IPv4 address [192.168.100.2]
4. Netmask [255.255.0.0]
5. Gateway IPv4 address []
6. DNS IPv4 Address [127.0.0.2]
7. Cancel and quit
8. Save and quit

NOTE:

The current values are shown in square brackets. Press the corresponding number (1–8) to modify each parameter.

c. **Navigate the interactive menu:** Use the menu options (1–6) to modify the following parameters as required for your network environment.

- Update Management Service IP address for SVM access
- Configure XenServer IP address for DOM0 management
- Set the subnet mask and default gateway for external connectivity.
- Configure the primary DNS server address for name resolution.


- d. **Save changes:** Once all parameters are updated, select option 8 (Save and quit) to commit your changes.

NOTE:

The Service VM automatically applies the new network configuration during the next startup process.

4. Management Service GUI access verification

- Test connectivity: Open a web browser and enter the new Management Service IP address in the address bar.
- Login: Authenticate using the username `nsroot` and the product Serial Number as the password.



1. DOM0 Console access

- Power on the NetScaler SDX appliance and connect your management laptop to the RJ45 console port.
- Set your terminal emulator (PuTTY) to 9600 baud, 8 data bits, None for parity, 1 stop bit, and None for flow control.
- Log in to the DOM0 using the username root and the product serial number as the password.

- d. Confirm that you have reached the DOM0 command line when the prompt displays as `[root@netscaler-sdx ~]#`.

System specifications

Hardware configurations

The SDX 9200 platform integrates enterprise-grade hardware built to deliver high-performance applications with consistent reliability.

Core system architecture

COMPONENT CATEGORY	SPECIFICATION DETAILS	PERFORMANCE CHARACTERISTICS
Regulatory Model	1U 3P3A	FCC and international compliance designation
Processor Architecture	Intel Xeon 5th Generation	24 cores per socket
Platform Chipset	Intel C741 PCH	Advanced platform controller hub
BMC Controller	Baseboard Management Controller with IPMI 2.0	Remote KVM, and hardware monitoring
System Memory	128 GB DDR5 5600 ECC RDIMM	2x 64 GB modules with error correction
Primary Storage	800 GB NVMe U.2/U.3	2.5" form factor, high-performance SSD
Management Connectivity	Two 1GbE RJ45 (1000Base-T)	Dedicated management port access
Data Connectivity	Two 4×50G SFP56 network interfaces	PCIe 4.0 x16 interface

Memory architecture

- **DDR5 technology:** Provides enhanced bandwidth and improved power efficiency.
- **ECC (Error Correcting Code):** Maintains data integrity and enhances overall system stability.
- **RDIMM design:** Uses registered DIMMs to support high-density memory configurations.

Storage system

- **NVMe interface:** Enables high-throughput storage with low latency.
- **U.2/U.3 compatibility:** Supports current and next-generation storage standards.
- **Optimized capacity:** Features 800 GB of storage, specifically tuned for system configurations and log retention.

Network interfaces

- **SFP56 connectivity:** Flexible ports that support a wide range of transceiver types and cabling options.
- **PCIe 4.0 integration:** Provides optimized bandwidth for 4x 50GbE high-speed connectivity.
- **Packet offloading:** Hardware-level acceleration for optimized network performance.

Thermal management

- Acoustic optimization for data center environments.
- Redundant cooling with fan modules.
- Temperature monitoring through the BMC interface.

Power and environmental specifications

Power system specifications

PARAMETER	SPECIFICATION	OPERATING RANGE	NOTES
Power Supply Configuration	2x CRPS PSUs	Redundant N+1 operation	Hot-swappable units
Input Voltage Range	100–240 VAC	Universal input compatibility	Auto-sensing
Input Frequency	50–60 Hz	Standard utility frequency	±3 Hz tolerance
Input Current	100–127 V: 10.0 A	Variable by input voltage	Maximum ratings
	200–240 V: 5.0 A		
Typical System Power Consumption	300 W	Normal operational load	Steady-state operation

PARAMETER	SPECIFICATION	OPERATING RANGE	NOTES
Maximum System Power Consumption	406 W	Peak system utilization	Full load + redundancy

Environmental specifications

ENVIRONMENTAL FACTOR	OPERATING RANGE	STORAGE RANGE	CRITICAL LIMITS
Ambient Temperature	0°C to 45°C (32°F to 113°F)	-40°C to 70°C	50°C Absolute maximum
Relative Humidity	10% to 90% RH	10% to 95% RH	Non-condensing
Operating Altitude	0 to 5,000 m	0 to 15,200 m	Derated above 8,000 m
Typical System Airflow	80.3 CFM	N/A	100% load at 25°C ambient
Maximum System Airflow	104 CFM	N/A	100% load at elevated temp

Physical specifications

DIMENSION	MEASUREMENT	STANDARD COMPLIANCE
Width	443 mm (17.4")	EIA 310-D 19-inch rack
Depth	550 mm (21.7")	Standard server depth
Height	44 mm (1.73")	1U Rack unit
System Weight	11.4 kg (25.2 lbs)	Approximate operational weight
Package Weight	18.6 kg (41.0 lbs)	Shipping configuration

System LED status

The system provides real-time status monitoring by using the LED indicators located on the front and rear panels.

Front panel LED status

LED INDICATOR	NORMAL OPERATION	FAULT CONDITION	DESCRIPTION
PWR (Power)	Solid Green	Off	Indicates system power and initialization status
GPIO1	Off	N/A	User-defined function (Disabled by default)
GPIO2	Off	N/A	User-defined function (Disabled by default)

Power supply LED status

LED INDICATOR	PSU STATE	SYSTEM IMPACT	ACTION REQUIRED
Solid GREEN	AC present, all DC rails active under load	Normal operation	None - normal
1 Hz Blink GREEN	AC present, 12Vsb active including 5Vsb for BMC power, main rails idle	Redundant standby (N+1)	None - normal
Solid AMBER	AC input lost or disconnected	Lost redundancy (if N+1) or system down	Check AC power connection
1 Hz Blink AMBER	OCP/OVP/OTP warning threshold reached	Degraded PSU operation/Potential shutdown	Verify PSU cooling and airflow

LED INDICATOR	PSU STATE	SYSTEM IMPACT	ACTION REQUIRED
Solid AMBER	Critical Fault (OCP/ OVP/OTP tripped)	PSU shutdown for protection	Replace the PSU immediately

NVMe storage LED status

The NVMe SSD activity indicators provide real-time storage operation status.



Left LED (Power/Present) indicator

CONDITION	LED BEHAVIOR
No SSD present	Off
SSD installed and powered	Solid Green

Right LED (Activity) indicator

NVME BRAND	CONDITION	LED BEHAVIOR
Micron	Idle	Very slow blinking
	Active (I/O in flight)	Fast Blinking (4 Hz)
Samsung	Idle	LED Off
	Active (I/O in flight)	Fast Blinking (4 Hz)

Additional information resources

- For information about step-by-step procedures for performing system upgrades and downgrades, see [Upgrade and downgrade a NetScaler appliance](#).

- For information about deploying redundant NetScaler configurations, see [Configuring high availability](#).

Information required for technical support

Before contacting NetScaler Support, gather the following information to ensure an efficient resolution process:

- **Hardware serial number:** Locate this on the physical chassis label.
- **Software version:** Provide the current release version and specific build number (Example: Release 14.1 Build 12.35).
- **Problem description:** Document the specific issue, including exact error codes, timestamps of the occurrence, and any recent changes to the environment.
- **System Configuration:** Provide the network environment and relevant configuration files (Example: ns.conf).
- **Log files:** Collect the system logs and diagnostic information.

NetScaler SDX 17000

Welcome to the next-generation NetScaler SDX 17000 platform. This section provides comprehensive instructions on how to deploy your appliance efficiently and configure it for optimal performance.

The SDX 17000 2U platform represents the latest advancement in NetScaler technology, featuring enhanced processing power, improved connectivity options, and robust management capabilities.

Key specifications

COMPONENT	SPECIFICATION
Processor	Intel Xeon 5th generation processor, 2x sockets, 24 cores / 48 threads per socket
Total Logical CPUs	96 (2 sockets x 24 cores x 2 threads per core)
System Memory	256 GB DDR5 ECC RDIMM
Boot Storage (RAID-1)	2x 960 GB NVMe SSD (slots 1 and 2)
VM Storage (RAID-1)	4x 800 GB NVMe SSD (slots 3 to 6)

COMPONENT	SPECIFICATION
Data Connectivity	2x Quad-port 50G SFP56 NIC (PCIe 4.0 x16) and 2x Dual-port 200G QSFP56 NIC (PCIe 5.0 x16)
Management Ports	2x 1 GbE RJ45 (MGMT 0/1 and 0/2) and 1x dedicated LOM port
Power Supplies	2x CRPS PSUs, redundant N+1, hot-swappable, up to 1600 W each
BMC / IPMI	Dedicated BMC for out-of-band management
Form Factor	2U, EIA 310-D 19-inch rack compatible, 88 mm height

Prepare for installation

Before installing your new appliance, carefully unpack the unit and verify that all the components are present. The only tools required for installing the appliance are a Phillips screwdriver, a flathead screwdriver, and a cage nut insertion tool (if installing into square-hole racks). Your NetScaler 2U package should include the following items:

- 1x NetScaler SDX 17000 appliance
- 1x standard 4-post rail kit (19-inch rack compatible), which includes:
 - Left rail (x1)
 - Right rail (x1)
 - Mounting screw kit (includes M4x4L, M4x6L, M3x6L, M5x20L screws)
 - Rackmount ear kit (x1)

The included accessory kit contains essential connectivity hardware, including two power cords to support redundant power supply operation.

NOTES:

- Rail kit compatibility varies by rack manufacturer and configuration. If the provided rail kit does not fit your specific rack infrastructure, contact your NetScaler sales representative immediately to obtain an appropriate alternative mounting solution.
- Before proceeding, examine all components for shipping damage. Report any damaged items to your carrier and NetScaler support before beginning the installation process.

Infrastructure requirements

- Ensure that the 19-inch rack has a minimum of 550 mm depth clearance and 2U available height (88 mm).
- Proper airflow management with front-to-rear cooling path.
- Adequate weight support (system weight: 38 lbs).

Network requirements

- Administrative LAN switch port available for MGMT 0/1 (1 GbE RJ-45).
- Dedicated out-of-band management network or LOM/IPMI access.
- Appropriate network cables (Cat5e or Cat6 for management, fiber, or DAC for data).
- 50 GbE capable switch for SFP56 data ports; 200 GbE capable switch for QSFP56 data ports.

Power requirements

- Dual power sources (recommended for N+1 redundancy).
- Frequency operating range: 47 Hz – 63 Hz.
- Adequate power capacity (full load 100 V: 806 W / 240 V: 765 W).
- PSU input connector: IEC320-C14 (Male).
- Required power cord: IEC320-C13 to NEMA 5-15P (or regional equivalent).
- Rating: Minimum 10 A / 250 V.

Detailed AC input voltage specifications

VOLTAGE RANGE	MINIMUM	RATED OPERATING RANGE	MAXIMUM	BROWN OUT	BROWN IN
115 V Range	90 Vrms	100-127 Vrms	140 Vrms	75 Vac+/-4 V	85 Vac+/-4 V
230 V Range	180 Vrms	200-240 Vrms	264 Vrms	N/A	N/A

Maximum input current requirements

INPUT VOLTAGE RANGE	MAXIMUM CURRENT
100 Vac – 127 Vac	12 A
200 Vac – 240 Vac	10 A

Management station requirements

- Workstation or laptop (Windows or macOS) with network connectivity.
- Web browser with JavaScript and HTML5 support (Chrome, Firefox, Safari, or Edge).
- Terminal emulator software for CLI access.
- Administrative privileges for network configuration.


Hardware platform

Front panel components

The SDX 17000 front panel provides management interfaces, NVMe drive bays, data ports, and power controls. Six hot-swappable NVMe bays are located at the front left of the panel.

NVMe drive bay assignments:

- Slots 1–2: Two 960 GB NVMe SSDs configured as RAID-1 for the SDX OS/system volume.
- Slots 3–6: Four 800 GB NVMe SSDs allocated as two independent RAID-1 pairs for the VM datastore.




Front panel component description

LABEL	COMPONENT TYPE	TECHNICAL SPECIFICATION	PRIMARY FUNCTION
CONSOLE	RJ-45 Console Port	RS-232 compatible, 9600-8N1	Primary serial console for direct system access
MGMT 0/1 & 0/2	Dual Ethernet Ports (1000Base-T)	1000Base-T, autonegotiate	Primary/secondary in-band management
LOM	Dedicated IPMI Port	Intelligent Platform Management Interface, IPMI 2.0, Redfish 1.15.1	Out-of-band hardware monitoring and control
NMI	Pinhole Button	Non-maskable interrupt	Force kernel core dump for crash analysis
PWR	Soft-push Button (4 sec on/off)	Power management	Power on / graceful shutdown control
TYPE-C	USB Type-C console port (optional)	Alternate console (not used in standard deployment)	Secondary console and debug interface

LABEL	COMPONENT TYPE	TECHNICAL SPECIFICATION	PRIMARY FUNCTION
UID	Locator LED	Blue LED indicator	Visual system identification in rack
GPIO1/GPIO2	User-defined GPIOs	General-purpose I/O	Not utilized; disabled by default on 2U platform
NVMe Slots 1-2	U.2/U.3 NVMe SSDs	960 GB per drive (2x), RAID-1	SDX OS boot volume
NVMe Slots 3-6	U.2/U.3 NVMe SSDs	800 GB per drive (4x), RAID-1	VM datastore (two pairs)
50/1-50/8	SFP56 Data Ports	50 GbE network interface card (NIC)	Two 4-port SFP56 with PCIe 4.0 x16
200/1-200/4	QSFP56 Data Ports	200 GbE network interface card (NIC)	Two 2-port QSFP56 with PCIe 5.0 x16

Rear panel components

The rear panel houses the dual Common Redundant Power Supply (CRPS), cooling fans, a secondary Unit Identification (UID) LED, a PSU alarm mute button, a VGA output port, and a rear power button.



Panel component description


COMPONENT	TECHNICAL DETAILS	OPERATIONAL CHARACTERISTICS	SERVICE INFORMATION
PSU1/PSU2	CRPS PSU, up to 1600 W each	Redundant N+1 configuration	Hot-swappable during operation
Rear UID LED	Blue LED locator	Blue LED indicator	Remote chassis identification
PSU Alarm Mute	Red push-button	Silences audible alarm on PSU fault or loss of AC input	Press once to acknowledge PSU alarm
VGA Output	DB-15 (female D-sub)	System video output	Local KVM connection; also accessible via LOM
Rear Power Button	Soft-push, optional	Mirrors front panel power button function	Dual power control access
Fan Modules	5 system fans	Smart fan intelligent control by BMC	Note: No field-replaceable Fan FRU

Install the hardware**Rack-mount the appliance**

Perform the following procedure to ensure proper appliance mounting and system performance.


Remove the inner rails from the rail assembly

1. Pull the white disconnect tab to release and detach the inner rail member from each slide rail.
2. With the stamped text facing outward, align the inner rail with the chassis mount points, using the black reference markers on both sides of the chassis.
3. Slide the rail toward the rear until it clicks firmly into the locked position.
4. Fasten each inner rail using the provided M4x4L screws. Use five screws per side for a total of ten.
5. Repeat this process for both the left and right sides of the chassis.




Chassis mounting ear installation

1. Attach rack mount ears to both the left and right sides of the chassis front.
2. Align the ears with the designated screw holes on the chassis.
3. Secure each ear with M4x6L screws using four screws per side.
4. Verify that the ears are flush and do not obstruct front panel access.





Outer slide rail assembly installation

1. Align the mounting brackets with the rack post holes at the selected 2 U-height.
2. Insert the rail mounting mechanism into the rack posts until it engages.
3. Listen for an audible "click" to confirm that the rail is properly seated and locked.
4. Ensure both the front and rear mounting points are securely attached to the posts.
5. Verify that the rails are level and parallel using the numbered rack post markings.
6. Repeat this process for both the left and right outer rail assemblies.



Mount the chassis into the rack

1. Carefully align the chassis-mounted inner rails with the rack-mounted outer rails.
2. Smoothly slide the system into the rack until the locking mechanism engages.
3. If necessary, apply light, even pressure to ensure the unit locks securely into place.
4. Secure each front ear to the rack post using one M5x20 screw per side.



Safety requirements

For safety and equipment durability, ensure that the following requirements are met:

1. Maintain a minimum of 2 inches of clearance at the front and rear of the chassis to ensure proper ventilation.
2. Verify the rack's weight capacity before installation and ensure that the rack is stable and properly secured to the floor.
3. Use proper lifting techniques. The system weight is approximately 38 lbs.
4. Strictly follow electrical safety procedures when connecting the appliance to a power source.

Airflow management

The 2U platform is designed for front-to-rear airflow. Unobstructed passage is required for optimal cooling performance. Restricted airflow results in thermal overheating and the automatic activation of system protection mechanisms (throttling or shutdown).

Network and power connections

Connection sequence

Perform the following procedure to ensure proper system initialization and to prevent configuration conflicts.

1. Management network connectivity

- a. Connect the 0/1 MGMT port to your administrative LAN by using a Cat5e/Cat6 Ethernet cable.
- b. Ensure that the network switch port is configured to support auto-negotiation.
- c. Verify network connectivity and ensure that the correct VLAN configuration is applied, if applicable.
- d. Document the assigned IP address and network parameters for future reference.

2. Out-of-Band (OOB) management setup

- a. Connect the Lights-Out Management (LOM) port to a dedicated out-of-band (OOB) management network.
- b. Use separate network infrastructure from production traffic.
- c. Leave the LOM port in dedicated mode operation (factory default).

3. Data port configuration

- a. Install appropriate SFP56 transceivers (50 GbE ports: 50/1-50/8) or QSFP56 transceivers (200 GbE ports: 200/1-200/4).
- b. Match the transceiver type to switch port capabilities. Use Citrix-qualified transceiver MPNs for guaranteed interoperability.
- c. Connect fiber or DAC cables to data ports as required.
- d. Verify that the cable specifications match the port capabilities (50 GbE and 200 GbE).

4. Power system connection

- a. Plug the first power cord into the PSU1 input receptacle.
- b. Plug the second power cord into the PSU2 input receptacle.
- c. Connect the unit to independent power feeds where possible to ensure N+1 redundancy.
- d. Ensure that the power source meets requirements (100–240 VAC, 50–60 Hz).

5. System startup

- a. Press the front panel power button to initiate the boot sequence.
- b. Observe the front panel LEDs; the power LED should illuminate solid green.
- c. Verify both PSU LEDs are solid green within ten seconds of power-on.
- d. Allow the full boot cycle to complete (approximately 4–5 minutes) before accessing the system.

Power supply status indicators

The CRPS power supplies provide comprehensive status monitoring through integrated LED indicators located on each unit's faceplate.

PSU LED indicator status:

LED INDICATOR	PSU STATE	SYSTEM IMPACT	ACTION REQUIRED
Solid GREEN	AC present, all DC rails active under load	Normal operation	None - normal
1 Hz Blink GREEN	AC present, 12Vsb active including 5Vsb for BMC power, main rails idle	Redundant standby (N+1)	None - normal

LED INDICATOR	PSU STATE	SYSTEM IMPACT	ACTION REQUIRED
Solid AMBER	AC input lost or disconnected	Lost redundancy (if N+1) or system down	Check AC power connection
1 Hz Blink AMBER	OCP/OVP/OTP warning threshold reached	Degraded PSU operation/Potential shutdown	Verify PSU cooling and airflow
Solid AMBER	Critical Fault (OCP/OVP/OTP tripped)	PSU shutdown for protection	Replace the PSU immediately

System access and configuration

SVM GUI access method

The web-based Service VM (SVM) management interface provides a centralized graphical user interface (GUI) for platform configuration and monitoring. Through the SVM GUI, administrators can manage key system functions and access platform resources associated with the SVM and the XenServer Dom0 hypervisor.

Platform-specific access information

- Management through virtualized service interfaces
- Management Service VM: 192.168.100.1 with 255.255.0.0 (/16) subnet mask
- XenServer hypervisor (DOM0): 192.168.100.2 with 255.255.0.0 (/16) subnet mask
- NSIP pool range: 192.168.100.10 to 192.168.100.100 for NetScaler instances

Perform the following steps to initialize the appliance

1. Network preparation
 - a. Connect your management workstation to the MGMT 0/1 port or a port on the same management VLAN.
 - b. Set a static IP address on the workstation in the 192.168.100.X/16 range (for example, 192.168.100.50).
 - c. Verify the network path by performing a ping test to the default SVM IP address: 192.168.100.1.
 - d. Ensure that any local firewalls allow HTTPS (port 443) traffic to the appliance.

2. Browser initialization

- a. Open a Web browser (Chrome, Firefox, Safari, or Edge).
- b. Enter the URL `https://192.168.100.1` (Management Service VM).

NOTE:

Accept a security certificate warnings (self-signed certificate) to proceed to the login page.

3. Initial authentication

- a. Username: `nsroot`.
- b. Password: Enter the system serial number exactly as it appears on the physical chassis label (case-sensitive).
- c. Click **Login** or press **Enter**.

On successful authentication, the system redirects you to the **Initial Configuration Wizard**.

4. Setup wizard configuration

The **Setup Wizard** guides you through the following essential parameters:

- a. Management service IP address (NSIP): Assign a static IP address for ongoing SVM access.
- b. XenServer IP address: Assign the DOM0 management address.
- c. Subnet mask and default gateway: Configure these for your network segment.
- d. DNS server: Enter the primary DNS IP address.
- e. Hostname: Assign a unique, descriptive hostname to identify the appliance on the network.
- f. Administrative password: Change the default password to a strong, secure administrative password.

Common default IP address configuration

INTERFACE OR SERVICE	DEFAULT USER LOGIN CREDENTIALS	DEFAULT IP ADDRESS	SUBNET MASK OR GATEWAY	DESCRIPTION
Management Service (SVM) (0/1)	nsroot / product SN	192.168.100.1	255.255.0.0/16	Management Service is the management interface of NetScaler SDX for provisioning and monitoring VPX instances
XenServer (DOM0)	root / product SN	192.168.100.2	255.255.0.0 /16 GW: 192.168.0.1	DOM0 is the hypervisor control domain (based on Linux) that interfaces with system hardware. Note: The product serial number as the default password for DOM0 is only valid after Management Service initialization.
XenServer (Management Service virtual port 0/2)	nsroot / product SN	169.254.0.10	255.255.0.0 /16	Management Service internal default IP address for 2U platform. Used for internal DOM0 and Management Service communication

INTERFACE OR SERVICE	DEFAULT USER LOGIN CREDENTIALS	DEFAULT IP ADDRESS	SUBNET MASK OR GATEWAY	DESCRIPTION
LOM or IPMI Port	nsroot / product SN	192.168.1.3	255.255.255.0/24	Out-of-band management IPMI, Redfish, remote iKVM, power control

Console connection setup by using the CLI

1. Physical connection

- Connect an RJ-45 serial cable to the console port on the front panel.
- Connect the other end of the cable to your workstation's serial port or to a USB-to-serial adapter.

Ensure that the cable is seated firmly at both ends to prevent data corruption or character "scrambling" during the session.

2. Terminal emulator configuration

Open your terminal emulation software (such as PuTTY, Tera Term, or SecureCRT) and configure the following settings:

SETTING	VALUE
Baud Rate	9600 bps
Data Bits	8
Parity	None
Stop Bits	1
Flow Control	None (Disabled)
Terminal Type	ANSI
Putty Keypad Mode	VT 100

3. Session establishment

- a. Open the terminal emulator application.
- b. Establish a connection by using the configured settings.
- c. Press Enter to trigger the login prompt if it does not appear automatically.

Wait for the login prompt to display.

4. Authentication process

- a. At the login prompt, type `nsroot`.
- b. At the password prompt, type the system serial number (found on the chassis label).
- c. On successful authentication, the system displays the NetScaler command prompt.


NOTE:

All login credentials are case-sensitive.

NetScaler SDX Management Service sample configuration

1. Management Service shell access

a. SDX Management Service initial network configuration



b. SSH connection

- i. From your management laptop, initiate an SSH session to the Management Service IP address : `ssh nsroot@192.168.100.1`
- ii. Authenticate using the username as nsroot and the password as the product serial number.
- iii. Log in to the NetScaler CLI prompt, run the `shell` command.

You can verify that you have successfully gained access when the prompt changes to `bash-3.2#`.

2. Configuration file verification

- a. **Access the configuration file:** From the `bash-3.2#` prompt, run the `cat /flash/mpsconfig/svm.conf` command to display the contents of the Management Service configuration file.

- b. **Review default contents:** Ensure the output matches the standard default settings, which include the following parameters:

Default svm.conf output:

- ifconfig 0/2 down
- ifconfig 0/2 delete
- arp -d -a
- route flush
- ifconfig 0/1 192.168.100.1 netmask 255.255.0.0
- ifconfig 0/2 169.254.0.10 netmask 255.255.0.0
- /mps/changenameserver.sh 127.0.0.2

- c. **Verify XenServer Connectivity:** Confirm the XenServer IP address is listed at the end of the file as: `#XenServer : 192.168.100.2`.

NOTE:

The SDX 17000 uses virtual port 0/3 for the SVM internal IP address (169.254.0.10). The port number differs from the SDX 1U platform, which uses port 0/2. Ensure that the svm.conf file reflects port 0/3 for the 2U platform.

3. Network configuration modification

- a. **Launch the configuration utility:** From the bash shell environment, initiate the interactive network setup tool by running the command: `/mps/networkconfig`.
- b. **Interactive configuration menu:** After launching the utility, the following menu appears:

```
-----
-----

NetScaler SDX initial network configuration

This menu allows you to set and modify the initial IPv4 network
addresses.
```

The current value is displayed in brackets ([]).

Selecting the listed number allows the address to be changed.

1. Management Service Host Name [ns]
2. Management Service IPv4 address [192.168.100.1]
3. XenServer IPv4 address [192.168.100.2]
4. Netmask [255.255.0.0]
5. Gateway IPv4 address []
6. DNS IPv4 Address [127.0.0.2]
7. Cancel and quit
8. Save and quit

NOTE:

The current values are shown in square brackets. Press the corresponding number (1–6) to modify each parameter.


- c. **Navigate the interactive menu:** Use the menu options (1–6) to modify the following parameters as required for your network environment.
- d. **Save changes:** Once all parameters are updated, select option 8 (Save and quit) to commit your changes.

NOTE:

The Service VM automatically applies the new network configuration during the next startup process.

4. Management Service GUI access verification

- a. **Test connectivity:** Open a web browser and enter the new Management Service IP address in the address bar.
- b. **Login:** Authenticate using the username `nsroot` and the product serial number as the password.



net.scaler | NetScaler SDX (NSSDX-17000) | Mar 11 2026 18:42:41 UTC | nsroot

Dashboard | Configuration | Documentation | Downloads

Search Menu

NetScaler SDX

Management Service Information

Platform	NSSDX-17000
Product	NetScaler SDX
Build	14.1: Build 43.137, Date: Jan 27 2026, 04:00:02
IP Address	10.217.116.86
Host ID	02357131ef3a
System ID	450103
Serial Number	MAN4ADR345
System Time	Wed Mar 11 18:42:37 UTC 2026
Uptime	5 days 12 hours 23 minutes


Hypervisor Information

Version	XenServer 8.4
Kernel Version	4.19.0-1
Platform Version	14.1.0-451
Applied Hotfixes	View Applied Hotfixes
Uptime	5 days 12 hours 28 minutes

System Information

BMC Firmware Version	Version: 4.01
BIOS Version	5.1a

- c. **Navigate to **Configuration > Interfaces**** to view all the management and data interfaces on SDX 17000 from the SVM GUI.



net.scaler | NetScaler SDX (NSSDX-17000) | Mar 11 2026 18:43:58 UTC | nsroot

Dashboard | Configuration | Documentation | Downloads

Search Menu

System > Interfaces

Interfaces

Edit | Reset | VLAN Filter | VLAN Stripping | Select Action

	INTERFACE	LINK STATE	SPEED	DUPLEX	MTU	BASE MAC ADDRESS	ADMIN STATE
☐	0/1	UP	1000Mb/s	full	1500	-NA-	Enabled
☐	0/2	UP	1000Mb/s	full	1500	-NA-	Enabled
☐	200/1	UP	20000Mb/s	full	1500	-NA-	Enabled
☐	200/2	UP	20000Mb/s	full	1500	-NA-	Enabled
☐	200/3	UP	20000Mb/s	full	1500	-NA-	Enabled
☐	200/4	UP	20000Mb/s	full	1500	-NA-	Enabled
☐	50/1	UP	5000Mb/s	full	1500	-NA-	Enabled
☐	50/2	UP	5000Mb/s	full	1500	-NA-	Enabled
☐	50/3	UP	5000Mb/s	full	1500	-NA-	Enabled
☐	50/4	UP	5000Mb/s	full	1500	-NA-	Enabled
☐	50/5	UP	5000Mb/s	full	1500	-NA-	Enabled
☐	50/6	UP	5000Mb/s	full	1500	-NA-	Enabled
☐	50/7	UP	5000Mb/s	full	1500	-NA-	Enabled
☐	50/8	UP	5000Mb/s	full	1500	-NA-	Enabled

5. DOM0 console access

- a. Power on the NetScaler SDX appliance and connect your management laptop to the RJ45 console port.
- b. Set your terminal emulator (PuTTY) to 9600 baud, 8 data bits, None for parity, 1 stop bit, and None for flow control.
- c. Log in to DOM0 using the username root and the product serial number as the password.
- d. Confirm that you have reached the DOM0 command line when the prompt displays as `[root@netscaler-sdx ~]#`.

System specifications

Hardware configurations

The SDX 17000 platform integrates enterprise-grade hardware built to deliver high-performance applications with consistent reliability.

Core system architecture

COMPONENT CATEGORY	SPECIFICATION DETAILS	PERFORMANCE CHARACTERISTICS
Regulatory Model	2U3P3A	FCC and international compliance designation
Processor Architecture	Intel Xeon 5th generation processor	2 sockets, 24 physical cores / 48 threads per socket; 96 logical CPUs total at 2.2 GHz
Platform Chipset	Intel C741 PCH	Advanced platform controller hub
BMC Controller	Baseboard Management Controller with IPMI 2.0	Remote KVM, hardware monitoring, power control
System Memory	256 GB DDR5 5600 ECC RDIMM	8x 32 GB modules, error correction enabled
Primary Storage	2x 960 GB NVMe U.2/U.3 (slot 1 and 2)	Configured as RAID-1 for SDX OS and system volume

COMPONENT CATEGORY	SPECIFICATION DETAILS	PERFORMANCE CHARACTERISTICS
VM Storage	4x 800 GB NVMe U.2/U.3 (slot 3 to 6)	Two independent RAID-1 pairs for VM datastore
Management Connectivity	Two 1 GbE RJ45 (1000Base-T)	MGMT 0/1 (primary), MGMT 0/2 (secondary)
50G Data Connectivity	Two 4-port SFP56 NICs	PCIe 4.0 x16 - ports 50/1-50/8, 8 total 50G ports
200G Data Connectivity	Two 2-port QSFP56 NICs	PCIe 5.0 x16 - ports 200/1-200/4, 4 total 200G ports
NUMA Topology	2 NUMA nodes	Each NUMA node corresponding to one CPU socket

Power and environmental specifications

Power system specifications

PARAMETER	SPECIFICATION	OPERATING RANGE	NOTES
Power Supply Configuration	2x CRPS PSUs	Redundant N+1 operation	Hot-swappable, up to 1600 W each
Input Voltage Range	100-240VAC	Universal input compatibility	Auto-sensing
Input Frequency	50-60 Hz	Standard utility frequency	+/-3Hz tolerance
Input Current	100-127Vac: 12A max 200-240Vac: 10A max	Variable by input voltage	Maximum ratings
Typical System Power Consumption	465 W	Normal operational load	Steady-state operation

PARAMETER	SPECIFICATION	OPERATING RANGE	NOTES
Maximum System Power Consumption	806 W at 100 V / 765 W at 240 V	Peak system utilization	Full load with both PSUs active

Environmental specifications

ENVIRONMENT FACTOR	OPERATING RANGE	STORAGE RANGE	CRITICAL LIMITS
Ambient Temperature	0-45 deg C (32-113 deg F)	-40 to 70 deg C	50 deg C absolute maximum
Relative Humidity	10-90% RH	10-95% RH	Non-condensing
Operating Altitude	0-5,000 m	0-15,200 m	Derated above 8,000 m
Typical System Airflow (front-to-back)	~126.54 CFM	N/A	Full load at nominal 25 deg C ambient
Maximum System Airflow (front-to-back)	~214.83 CFM	N/A	Full load at elevated temperature

Physical specifications

DIMENSION	MEASUREMENT	STANDARD COMPLIANCE
Width	443 mm (17.4 in)	EIA 310-D 19-inch rack
Depth	550 mm (21.7 in)	Standard server depth
Height	88 mm (3.5 in)	2U rack unit
System Weight	38 lbs	Approximate operational weight
Package Weight	52 lbs	Shipping configuration

System LED status

The system provides real-time status monitoring by using the LED indicators located on the front and rear panels.

Front panel LED status

LED INDICATOR	NORMAL OPERATION	FAULT CONDITION	DESCRIPTION
PWR (Power)	Solid Green	Off	System power and initialization status
GPIO1	Off (default)	N/A	User-defined function (disabled)
GPIO2	Off (default)	N/A	User-defined function (disabled)

Power supply LED status

LED INDICATOR	PSU STATE	SYSTEM IMPACT	ACTION REQUIRED
Solid GREEN	AC present, all DC rails active under load	Normal operation	None - normal
1 Hz Blink GREEN	AC present, 12Vsb active including 5Vsb for BMC power, main rails idle	Redundant standby (N+1)	None - normal
Solid AMBER	AC input lost or disconnected	Lost redundancy (if N+1) or system down	Check AC power connection
1 Hz Blink AMBER	OCP/OVP/OTP warning threshold reached	Degraded PSU operation/Potential shutdown	Verify PSU cooling and airflow
Solid AMBER	Critical Fault (OCP/OVP/OTP tripped)	PSU shutdown for protection	Replace the PSU immediately

NVMe SSD tray LED status

Each NVMe tray includes two LEDs. The right LED indicates drive presence and activity. The left LED indicates drive status and is controlled by the NVMe backplane controller through the BMC.



Left LED status indicators:

LEFT LED (LOCATION / STATUS)	LED STATE
Locate/Identify Drive	Yellow Blink at 4 Hz

Use the LOM web UI (Storage section) to toggle the location LED for a specific NVMe slot (0-5) to identify a drive in a populated rack. Select the target drive and toggle it from OFF (default) to ON (locate active).

Right LED status indicators:

RIGHT LED (ACTIVE / PRESENT)	LED STATE
SSD present and idle	Solid Green
Storage data transfer/active	Green, Blinking at 4 Hz
No SSD Present or Removed	LED Off

Additional information resources

- For information about detailed specifications, configuration guides, and troubleshooting procedures, see the [NetScaler Hardware Platforms Documentation](#).
- For information about step-by-step procedures for performing system upgrades and downgrades, see [Upgrade and downgrade a NetScaler appliance](#).

- For information about deploying redundant NetScaler configurations, see [Configuring high availability](#).

Technical support access

- Primary support portal: Customer portal at <https://www.netscaler.com/>
- Online documentation: Complete technical library at www.docs.netscaler.com

Information required for technical support

Before contacting NetScaler Support, gather the following information to ensure an efficient resolution process:

- **Hardware serial number:** Locate this on the physical chassis label.
- **Software version:** Provide the current release version and specific build number (Example: Release 14.1 Build 12.35).
- **Problem description:** Document the specific issue, including exact error codes, timestamps of the occurrence, and any recent changes to the environment.
- **System configuration:** Provide the network environment and relevant configuration files (Example: ns.conf).
- **Log files:** Collect the system logs and diagnostic information.

NetScaler SDX 14000


The NetScaler SDX 14020/14030/14040/14060/14080/14100 are 2U appliances. Each model has two 6-core processors and 64 GB of memory and sixteen 10G SFP+ ports (16x10G SFP+).

NOTE

For information about NetScaler SDX hardware and component compatibility matrix, see [NetScaler SDX Hardware-Software Compatibility Matrix](#).

The following figure shows the front panel of the SDX 14020/14030/14040/14060/14080/ 14100 (16x10G SFP+) appliance.

Figure 1. NetScaler SDX 14020/14030/14040/14060/14080/14100 (16x10G SFP+), front panel



The NetScaler SDX 14020/14030/14040/14060/14080/14100 appliances have the following ports:


- RS232 serial Console Port.
- 10/100Base-T copper Ethernet Port (RJ45), also called the LOM port. You can use this port to remotely monitor and manage the appliance independently of the NetScaler software.
- Two 10/100/1000Base-T copper Ethernet management ports (RJ45), numbered 0/1 and 0/2 from left to right. These ports are used to connect directly to the appliance for system administration functions.
- Network ports, sixteen 10G SFP+ ports (16x10G SFP+).
- USB ports (reserved for a future release).

NOTE

The 10G SFP+ ports on these appliances support copper 1G SFP transceivers.

The following figure shows the back panel of the SDX 14020/14030/14040/14060/14080/ 14100 appliance.

Figure 2. NetScaler SDX 14020/14030/14040/14060/14080/14100, back panel



The following components are visible on the back panel of the SDX 14020/14030/14040/14060/14080/14100 appliance:

- Two 240 GB or larger removable solid-state drives in a redundant array of independent disks (RAID) configuration.

Note: These drives store the NetScaler® software, and cannot be increased by users.

- Two to six 300 GB or larger solid-state drives.

Note: These drives are used to provision VPX instances. Contact NetScaler sales Support to purchase extra drives.

Note: Drive densities might increase as components become EOL but its size is never smaller than the original.

For more information about managing RAID allocation, see [Managing RAID disk allocation on SDX 22XXX](#).

- **Power switch** which turns power to the appliance on or off. Press the switch for less than two seconds to shut off the power.
- **Two power supplies.** each rated at 1000 watts, 100–240 volts. Each power supply has an LED that indicates the status of the power supply, as described in [Common hardware components](#).
- **Disable alarm button.** which is functional only when the appliance has two power supplies. Press this button to stop the power alarm from sounding when either of the following conditions is true:
 - You have plugged the appliance into only one power outlet

- One power supply is malfunctioning and you want to continue operating the appliance until it is repaired.
- **Non-maskable interrupt (NMI) Button.** used at the request of Technical Support to initiate a core dump. To press this red button, which is recessed to prevent unintentional activation, use a pen, pencil, or other pointed object. The NMI Button is also available remotely over the network in the LOM GUI, in the Remote Control menu. For more information about the lights out management port of the appliance, see [Lights out management port of the NetScaler SDX appliance](#).

Note: If you buy an instance license pack for SDX 14000 that includes extra SSDs, read [CTX215581](#). Complete the steps for adding a logical drive before trying to use the additional instance licenses.

NetScaler SDX 14000-40G

The NetScaler SDX 14020-40G, SDX 14040-40G, SDX 14060-40G, SDX 14080-40G, and SDX 14100-40G are 2U appliances. Each model has two 6-core processors and 64 GB of memory. The SDX 14020-40G/14040-40G/14060-40G/14080-40G/14100-40G appliances are available in the four-40G QSFP+ ports and sixteen 10G SFP+ ports (4x40G QSFP+ + 16x10G SFP+) configuration.

For information on the software releases supported on the NetScaler hardware platforms, see [NetScaler SDX hardware-software compatibility matrix](#).

The following figure shows the front panel of the SDX 14000-40G appliance.

Figure 1. NetScaler SDX™ 14020-40G, SDX 14040-40G, SDX 14060-40G, SDX 14080-40G, SDX 14100-40G (4x40G QSFP+ + 16x10G SFP+), front panel



The NetScaler SDX 14000-40G appliances have the following ports:

- RS232 serial Console Port.
- 10/100Base-T copper Ethernet Port (RJ45), also called the LOM port. You can use this port to remotely monitor and manage the appliance independently of the NetScaler software.
- Two 10/100/1000Base-T copper Ethernet Management Ports (RJ45), numbered 0/1 and 0/2 from left to right. These ports are used to connect directly to the appliance for system administration functions.

- Network Ports: Four 40G QSFP+ ports and sixteen 10G SFP+ ports (4x40G QSFP+ + 16x10G SFP+). For information about supported transceivers per port, see [25G, 40G, 50G, and 100G ports](#).
- USB ports (reserved for a future release).

The following figure shows the back panel of the SDX 14000-40G appliance.

Figure 2. NetScaler SDX 14020/14040/14060/14080/14100-40G (4x40G QSFP+, 16x10G SFP+), back panel



The following components are visible on the back panel of the SDX 14000-40G appliance:

- Two 240 GB or larger removable solid-state drives (SSDs). The two leftmost solid-state drives store the NetScaler® software. The next two solid-state drives store user data. The remaining four solid-state drives are reserved for future use. These appliances are a redundant array of independent disks (RAID) devices. In a RAID configuration, the same data is stored on multiple drives to improve performance, increase storage capacity, lower the risk of data loss, and provide fault tolerance.

Note: Drive densities might increase as components become EOL but its size is never smaller than the original.

- Power switch, which turns the power to the appliance on or off. Press the switch for less than two seconds to shut off the power.
- Two power supplies, each power supply is rated at 1000 watts, 100–240 volts. Each power supply has an LED that indicates the status of the power supply.
- **Disable alarm button.** This button is functional only when the appliance has two power supplies. Press this button to stop the power alarm from sounding when either of the following conditions is true:
 - You have plugged the appliance into only one power outlet.
 - One power supply is malfunctioning and you want to continue operating the appliance until it is repaired.
- Non-maskable interrupt (NMI) button that is used at the request of Technical Support to initiate a core dump. To press this red button, which is recessed to prevent unintentional activation, use a pen, pencil, or other pointed object. The NMI Button is also available remotely over the network in the LOM GUI, in the Remote Control menu. For more information about the lights

out management port of the appliance, see [Lights out management port of the NetScaler SDX appliance](#).

NetScaler SDX 14xxx-40S

The NetScaler SDX 14040-40S, SDX 14060-40S, SDX 14080-40S, and SDX 14100-40S are 2U appliances. Each model has two 6-core processors and 64 GB of memory. The SDX 14040-40S/14060-40S/14080-40S/SDX 14100-40S appliances are available in the four-40S QSFP+ ports and sixteen 10G SFP+ ports (4x40G QSFP+ + 8x10G SFP+) configuration.

For information on the software releases supported on the NetScaler hardware platforms, see [NetScaler SDX hardware-software compatibility matrix](#).

The following figure shows the front panel of the SDX 14xxx-40S appliance.

Figure 1. NetScaler SDX™ 14040-40S, SDX 14060-40S, SDX 14080-40S, SDX 14100-40S (4x40G QSFP+ + 8x10G SFP+), front panel



The NetScaler SDX 14xxx-40S appliances have the following ports:

- RS232 serial Console Port.
- 10/100Base-T copper Ethernet Port (RJ45), also called the LOM port. You can use this port to remotely monitor and manage the appliance independently of the NetScaler software.
- Two 10/100/1000Base-T copper Ethernet Management Ports (RJ45), numbered 0/1 and 0/2 from left to right. These ports are used to connect directly to the appliance for system administration functions.
- Network Ports: Four 40S QSFP+ ports and eight 10G SFP+ ports (4x40G QSFP+ + 8x10G SFP+). For information about supported transceivers per port, see [25G, 40G, 50G, and 100G ports](#).

The following figure shows the back panel of the SDX 14xxx-40S appliance.

Figure 2. NetScaler SDX 14040/14060/14080/14100-40S (4x40G QSFP+, 8x10G SFP+), back panel



The following components are visible on the back panel of the SDX 14xxx-40S appliance:

- Two 240 GB or larger removable solid-state drives (SSDs). The two leftmost solid-state drives store the NetScaler® software. The next two solid-state drives store user data. The remaining

four solid-state drives are reserved for future use. These appliances are a redundant array of independent disks (RAID) devices. In a RAID configuration, the same data is stored on multiple drives to improve performance, increase storage capacity, lower the risk of data loss, and provide fault tolerance.

Note: Drive densities might increase as components become EOL but its size is never smaller than the original.

- Power switch, which turns the power to the appliance on or off. Press the switch for less than two seconds to shut off the power.
- Two power supplies, each power supply is rated at 1000 watts, 100–240 volts. Each power supply has an LED that indicates the status of the power supply.
- **Disable alarm button.** This button is functional only when the appliance has two power supplies. Press this button to stop the power alarm from sounding when either of the following conditions is true:
 - You have plugged the appliance into only one power outlet.
 - One power supply is malfunctioning and you want to continue operating the appliance until it is repaired.
- Non-maskable interrupt (NMI) button that is used at the request of Technical Support to initiate a core dump. To press this red button, which is recessed to prevent unintentional activation, use a pen, pencil, or other pointed object. The NMI Button is also available remotely over the network in the LOM GUI, in the Remote Control menu. For more information about the lights out management port of the appliance, see [Lights out management port of the NetScaler SDX appliance](#).


NetScaler SDX 14030 FIPS, 14060 FIPS, 14080 FIPS

IMPORTANT:

LAS support for the NetScaler SDX 14000 FIPS is available only in NetScaler SVM version 13.1, starting from SVM build 13.1.60.x onwards. Customers must upgrade NetScaler SDX 14000 FIPS to SVM version 13.1 build 60.x or later for LAS transition. LAS is not supported in SVM version 14.1 for the NetScaler SDX 14000 FIPS. Therefore, customers must use only the latest SVM 13.1 build.

The NetScaler models SDX 14030 FIPS, SDX 14060 FIPS, and SDX 14080 FIPS are 2U appliances. Each model has two 6-core processors, 64 GB memory, 16 10G SFP+ ports (16x10G SFP+).

Figure 1. NetScaler SDX 14030/14060/14080 FIPS front panel




The NetScaler SDX 14030/14060/14080 FIPS appliances have the following ports:

- RS232 serial Console Port.
- 10/100Base-T copper Ethernet Port (RJ45), also called the LOM port. You can use this port to remotely monitor and manage the appliance independently of the NetScaler software.
- Two 10/100/1000Base-T copper Ethernet Management Ports (RJ45), numbered 0/1 and 0/2 from left to right. These ports are used to connect directly to the appliance for system administration functions.
- Sixteen 10G SFP+ ports.
- USB ports (reserved for a future release).

NOTE:

- 10G ports do not support 1G copper or 1G fiber transceivers.
- 40G ports do not support 10G and 1G transceivers.

Figure 2. NetScaler SDX 14030/14060/14080 FIPS appliance back panel



The following components are visible on the back panel of the SDX 14030 FIPS, SDX 14060 FIPS, and SDX 14080 FIPS appliances:

- Two 240 GB or larger removable solid-state drives in a redundant array of independent disks (RAID) configuration.

Note: These drives store the NetScaler® software, and cannot be increased by users.

- Two to six 300 GB or larger solid-state drives.

Note: These drives are used to provision VPX instances. Contact NetScaler sales Support to purchase extra drives.

Note: Drive densities might increase as components become EOL but its size is never smaller than the original.

- **Power switch.** Which turns power to the appliance on or off. Press the switch for less than two seconds to shut off the power.
- Two power supplies, each rated at 1000 watts, 100–240 volts. Each power supply has an LED that indicates the status of the power supply, see [Common hardware components](#).
- **Disable alarm button.** Which is functional only when the appliance has dual redundant, hot-swappable power supplies. Press this button to silence the power alarm when one of the two power supplies loses input power or when a power supply is malfunctioning.
- **Non-maskable interrupt (NMI) button.** Used at the request of Technical Support to initiate a core dump. To press this red button, which is recessed to prevent unintentional activation, use a pen, pencil, or other pointed object. The NMI Button is also available remotely over the network in the lights out management GUI, in the Remote Controlment. For more information, see [Lights out management port of the NetScaler SDX appliance](#).

See [Configuring an SDX 14000 FIPS appliance](#) for more information about how to configure an SDX 14030/14060/14080 FIPS appliance.


NetScaler SDX 15000

The NetScaler SDX 15020/15030/15040/15060/15080/15100/15120 appliances are 2U appliances. This platform has two 8-core processors and 128 GB of memory. The appliance provides a total of 16 Ethernet 25G ports.

For information on the software releases supported on the NetScaler hardware platforms, see [NetScaler SDX hardware-software compatibility matrix](#).

The following figure shows the front panel of the SDX 15000 series appliances.

Figure 1. NetScaler SDX 15000, front panel




The NetScaler SDX 15000 series appliances have the following ports:

- RS232 serial console port.
- One 10/100/1000Base-T RJ45 copper Ethernet LOM Port. Use this port to remotely monitor and manage the appliance independently of the NetScaler software.
- One 10/100/1000Base-T RJ45 copper Ethernet Management Port, numbered 0/1. This port is used to connect directly to the appliance for NetScaler administration functions.
- Sixteen 25G Ethernet ports, numbered 25/1 to 25/16. For information about supported transceivers per port, see [25G, 40G, 50G, and 100G ports](#).

The following figure shows the back panel of the SDX 15000 appliance.

Figure 2. NetScaler SDX 15000, back panel



The following components are visible on the back panel of the SDX 15000 appliance:

- Four 240 GB and four 480 GB removable solid-state drives (SSD). Two 240 GB are used by the system while the remaining are used for virtual instances and in RAID 1 pairs.

Note: Drive densities might increase as components become EOL but its size is never smaller than the original.
- Power switch, which turns power to the appliance on or off.
 - If the OS is functional, press the switch for less than two seconds to shut down the system gracefully.
 - If the OS is not responsive, press the power switch for more than 4 seconds to force the power off.
- Two hot-swappable 100–240 VAC input power-supply modules. Each power supply has an LED indicating its status:

LED COLOR	LED INDICATES
OFF	No power to any power supply in the appliance.
Flashing RED	No power to this power supply.

LED COLOR	LED INDICATES
Flashing GREEN	Power supply is in standby mode.
GREEN	Power supply is functional.
RED	Power supply failure.
Flashing RED and GREEN	Warning (OVP/UVP/OCP/OTP/Fan); OVP = Over Voltage Protection; UVP = Under Voltage Protection; OCP = Over Current Protection; OTP = Over Temperature Protection

- **Disable alarm button:** functional only when the appliance has two power supplies. Press this button to silence the power alarm when one of two power supplies loses input power (second power supply optional) or when a power supply is malfunctioning.
- **Non-Maskable Interrupt (NMI) button:** recessed to prevent unintentional activation and used at the request of Technical Support to initiate a core dump. To press this red button, use a pen, pencil, or other pointed object. The NMI button is also available remotely over the network in the LOM GUI, in the Remote Control menu. For more information, see [Lights out management port of the NetScaler SDX appliance](#) topic.

NetScaler SDX 15000-50G


The NetScaler SDX 15020/15030/15040/15060/15080/15100/15120 50G models are 2U appliances. This platform has a dual 8-core processor and 128 GB memory. The SDX 15000-50G appliance provides a total of 12 network ports:

- Eight 10G SFP+ Ethernet Ports
- Four 50G Ethernet Ports

For information about the software releases supported on the NetScaler hardware platforms, see [NetScaler SDX Hardware and Component Compatibility Matrix](#)

The illustrations in this guide are only indicative and might not represent the actual NetScaler appliance.


Figure 1. NetScaler SDX 15000-50G, front panel



The NetScaler SDX 15000-50G appliance has the following ports:

- One RS232 Serial Console Port.
- One 10/100/1000 Base-T RJ45 copper Ethernet lights out management port. Use this port to remotely monitor and manage the appliance independently of the NetScaler software.
- Two 10/100/1000Base-T RJ45 copper Ethernet Management Ports, numbered 0/1 and 0/2. These ports are used to connect directly to the appliance for NetScaler administration functions.
- Eight 10G SFP+ Ethernet ports, numbered 10/1 to 10/8. Each port has its own LED.
- Four 50G ports, numbered 50/1 to 50/4. Each port has its own LED. For information about supported transceivers per port, see [25G, 40G, 50G, and 100G ports](#).

Figure 2. NetScaler SDX 15000-50G, rear panel



The back panel of the SDX 15000-50G appliance has the following components:

- Two 240 GB RAID-supported removable boot solid-state drives (SSDs), (slots 1 and 2) as shown in figure 2. Two 240 GB RAID-supported removable storage repositories (slots 3 and 4 paired) SSDs, and four 480 GB storage repositories (slots 5–6 paired and 7–8 paired) SSDs.

Note: Drive densities might increase as components become EOL but its size is never smaller than the original.

- A power switch, which turns power to the appliance on or off. If the OS is functional, press the switch for less than two seconds to power down the system with a graceful shutdown. If the OS is not responsive, press the power switch for more than 4 seconds to force the power off.
- Two hot-swappable 100–240 VAC input power supply modules. Each power supply has an LED indicating its status, as shown in the following table.

LED COLOR	LED INDICATES
OFF	No power to any power supply in the appliance.
Flashing RED	No power to this power supply module.
Flashing GREEN	Power supply is in standby mode.
GREEN	Power supply is functional.
RED	Power supply failure.

LED COLOR	LED INDICATES
Flashing RED and GREEN	Warning (OVP/UVP/OCP/OTP/Fan); OVP = Over Voltage Protection; UVP = Under Voltage Protection; OTP = Over Temperature Protection.

- A **Disable alarm button**, which is functional only when the appliance has dual redundant, hot-swappable power supplies. Press this button to silence the power alarm when one of the two power supplies loses input power or when a power supply is malfunctioning.
- A Non-Maskable Interrupt (NMI) button, used at the request of Technical Support to initiate a core dump. To press this red button, which is recessed to prevent unintentional activation, use a pen, pencil, or other pointed object. The NMI Button is also available remotely over the network in the LOM GUI, in the Remote Control menu. For more information about the lights out management port of the appliance, see [Lights out management port of the NetScaler SDX appliance](#).


NetScaler SDX 16000

The NetScaler SDX 16000 appliance is a 2U appliance. This platform has two 16-core processors and 256 GB (16 x 16 GB DIMM) of memory. The appliance provides a total of eight 25G SFP28 ports and four 100G QSFP28 Ethernet ports.

For information on the software releases supported on the ADC hardware platforms, see Hardware-Software Release Matrix.

The following figure shows the front panel of the SDX 16000 appliance.

Figure 1. NetScaler SDX 16000, front panel




The NetScaler SDX 16000 series appliances have the following ports:

- RS232 serial Console port.
- One 10/100/1000Base-T RJ45 copper Ethernet LOM port. Use this port to remotely monitor and manage the appliance independently of the ADC software.
- Two 10/100/1000Base-T RJ45 copper Ethernet management ports, numbered 0/1 and 0/2. These ports are used to connect directly to the appliance for ADC administration functions.
- Two USB ports (reserved for a future release).
- Eight 25G SFP28 ports, numbered 25/1 to 25/8. Four 100G QSFP28 ports, numbered 100/1 to 100/4. For information about supported transceivers per port, see [25G](#), [40G](#), [50G](#), and [100G ports](#).

The following figure shows the back panel of the SDX 16000 appliance.

Figure 2. NetScaler SDX 16000, back panel



The following components are visible on the back panel of the SDX 16000 appliance:

- Two 480 GB removable solid-state drives (boot drives).
- Six 480 GB removable solid-state drives for storing data.

Note: Drive densities might increase as components become EOL but its size is never smaller than the original.

- Power switch, which turns power to the appliance on or off.
 - If the OS is functional, press the switch for less than two seconds to power down the system with a graceful shutdown.
 - If the OS is not responsive, press the power switch for more than 4 seconds to force the power off.
- Two power supplies, each rated at 850 watts, 100–240 VAC. The following table indicates the LED status of each power supply:

LED COLOR	LED INDICATES
OFF	No power to any power supply in the appliance.
Flashing RED	No power to this power supply.
Flashing GREEN	Power supply is in standby mode.


LED COLOR	LED INDICATES
GREEN	Power supply is functional.
RED	Power supply failure.

- **Disable alarm button**, which is functional only when the appliance has two power supplies. Press this button to silence the power alarm when one of two power supplies loses input power (second power supply optional) or when a power supply is malfunctioning.
- **Non-Maskable Interrupt (NMI) Button**, used at the request of Technical Support to initiate a core dump. To press this red button, which is recessed to prevent unintentional activation, use a pen, pencil, or other pointed object. The NMI Button is also available remotely over the network in the LOM GUI, in the **Remote Control** menu. For more information about the lights out management port of the appliance, see [Lights out management port of the NetScaler SDX appliance](#).

NetScaler SDX 25100A, SDX 25160A

The NetScaler SDX 25100A and SDX 25160A are 2U appliances. Each model has two 10-core processors, 256 GB of memory, eight 40GE QSFP+ ports (8x40GE QSFP+).

Figure 1. NetScaler SDX 25100A and SDX 25160A front panel




The NetScaler SDX 25160A and SDX 25100A appliances have the following ports:

- RS232 serial Console Port.
- 10/100Base-T copper Ethernet Port (RJ45), also called the LOM port. You can use this port to remotely monitor and manage the appliance independently of the NetScaler software.
- Network ports, eight 40G QSFP+ ports. For information about supported transceivers per port, see [25G, 40G, 50G, and 100G ports](#).

NOTE

40G ports do not support 10G and 1G transceivers.

Figure 2. NetScaler SDX 25100A and SDX 25160A back panel



The following components are visible on the back panel of the SDX 25100A and SDX 25160A appliance:


- **Two 300 GB removable solid-state drives (SSDs).** The two leftmost solid-state drives store the NetScaler software. The next two solid-state drives store user data. The remaining four solid-state drives are reserved for future use. These appliances are a redundant array of independent disks (RAID) devices.
- **Note:** Drive densities might increase as components become EOL but its size is never smaller than the original.
- **Power switch.** Which turns power to the appliance on or off. Press the switch for less than two seconds to shut off the power.
- **Two power supplies.** each rated at 1000 watts, 100–240 volts. Each power supply has an LED that indicates the status of the power supply.
- **Disable alarm button.** Which is functional only when the appliance has dual redundant, hot-swappable power supplies. Press this button to silence the power alarm when one of the two power supplies loses input power or when a power supply is malfunctioning.
- **Non-maskable interrupt (NMI) button.** Used at the request of Technical Support to initiate a core dump. To press this red button, which is recessed to prevent unintentional activation, use a pen, pencil, or other pointed object. The NMI Button is also available remotely over the network in the LOM GUI, in the Remote Control menu. For more information about the lights

out management port of the appliance, see [Lights out management port of the NetScaler SDX appliance](#).

NetScaler SDX 25100-40G

The NetScaler SDX 25100-40G and SDX 25160-40G are 2U appliances. Each model has two 10-core processors, 256 GB of memory, four 40G QSFP+ ports, and sixteen 10G SFP+ ports (4x40G QSFP+ + 16x10G SFP+).

Figure 1. NetScaler SDX 25100 40G, SDX 25160 40G (4x40G QSFP+ + 16x10G SFP+) front panel




The NetScaler SDX 25100 40G, SDX 25160 40G appliances have the following ports:

- RS232 serial Console Port.
- 10/100Base-T copper Ethernet Port (RJ45), also called the LOM port. You can use this port to remotely monitor and manage the appliance independently of the NetScaler software.
- Two 10/100/1000Base-T copper Ethernet Management Ports (RJ45), numbered 0/1 and 0/2 from left to right. These ports are used to connect directly to the appliance for system administration functions.
- Network Ports: Four 40G QSFP+ ports and sixteen 10G SFP+ ports (4x40G QSFP+ + 16x10G SFP+). For information about supported transceivers per port, see [25G, 40G, 50G, and 100G ports](#).
- USB ports (reserved for a future release).

NOTE:

- 10G ports do not support 1G copper or 1G fiber transceivers.
- 40G ports do not support 10G and 1G transceivers.

Figure 2. NetScaler SDX 25100-40G and SDX 25160-40G back panel

The following components are visible on the back panel of the SDX 25100-40G, SDX 25160-40G appliance:

- **Two 300 GB removable solid-state drives (SSDs).** The two leftmost solid-state drives store the NetScaler software. The next two solid-state drives, of 300 GB each, store user data. The remaining four solid-state drives are reserved for future use.

Note: Drive densities might increase as components become EOL but its size is never smaller than the original.

- **Power switch.** Which turns power to the appliance on or off? Press the switch for less than two seconds to shut off the power.
- **Two power supplies.** Each rated at 1000 watts, 100–240 volts. Each power supply has an LED that indicates the status of the power supply, as described in [Common Hardware Components](https://en-us.netScaler-hardware-platforms/sdx/common-hardware-components.html)[(/en-us/netScaler-hardware-platforms/sdx/common-hardware-components.html)].
- **Disable alarm button.** Which is functional only when the appliance has dual redundant, hot-swappable power supplies. Press this button to silence the power alarm when one of the two power supplies loses input power or when a power supply is malfunctioning.

- **Non-maskable interrupt (NMI) button.** Used at the request of Technical Support to initiate a core dump. To press this red button, which is recessed to prevent unintentional activation, use a pen, pencil, or other pointed object. The NMI Button is also available remotely over the network in the LOM GUI, in the Remote Control menu. For more information about the lights out management port of the appliance, see [Lights out management port of the NetScaler SDX appliance](#).


NetScaler SDX 26000

The NetScaler SDX 26000 appliances are 2U appliances. These appliances have two 14-core processors and 256 GB of memory. The appliance has eight 50G and sixteen 25G network ports.

For information about the software releases supported on the NetScaler hardware platforms, see [NetScaler SDX hardware-software compatibility matrix](#).

The following figure shows the front panel of the NetScaler SDX 26000 appliances.

Figure 1. NetScaler SDX 26000, front panel




The NetScaler SDX 26000 appliances have the following ports:

- RS232 serial console port.
- One 10/100/1000Base-T RJ45 copper Ethernet LOM Port. Use this port to remotely monitor and manage the appliance independently of the NetScaler software.
- Two 10/100/1000Base-T RJ45 copper Ethernet Management Ports, numbered 0/1 and 0/2. These ports are used to connect directly to the appliance for NetScaler administration functions.

- Eight 50G ports, numbered 50/1 to 50/8.
- Sixteen 25G ports, numbered 25/1 through 25/16. For information about supported transceivers per port, see [25G, 40G, 50G, and 100G ports](#).

The following figure shows the back panel of the SDX 26000 appliances.

Figure 2. NetScaler SDX 26000, back panel



The following components are visible on the back panel of the NetScaler SDX 26000 appliances:

- Two 240 GB and six 960 GB removable solid-state drive (SSD).

Note: Drive densities might increase as components become EOL but its size is never smaller than the original.

- Power switch, which turns power to the appliance on or off.
 - If the OS is functional, press the switch for less than two seconds to power down the system with a graceful shutdown.
 - If the OS is not responsive, press the power switch for more than 4 seconds to force the power off.
- Two hot-swappable 100–240 VAC, 1000 W power supply modules. Max power consumption is 672 W. Typical power consumption is 540 W. Each power supply has an LED indicating its status as follows:

LED COLOR	LED INDICATES
OFF	No power to any power supply on the appliance.
Flashing RED	No power to this power supply.
Flashing GREEN	Power supply is in standby mode.
GREEN	Power supply is functional.
RED	Power supply failure.
Flashing RED and GREEN	Warning (OVP/UVP/OCP/OTP/Fan); OVP = Over Voltage Protection; UVP = Under Voltage Protection; OCP = Over Current Protection; OTP = Over Temperature Protection

- **Disable alarm button.** Press this button to silence the power alarm when one of two power supplies loses input power or when a power supply is malfunctioning.
- **Non-Maskable Interrupt (NMI) Button,** used at the request of Technical Support to initiate a core dump. To press this red button, which is recessed to prevent unintentional activation, use a pen, pencil, or other pointed object. The NMI Button is also available remotely over the network in the LOM GUI, in the Remote Control menu. For more information, see [Lights out management port of the NetScaler SDX appliance](#).


NetScaler SDX 26000-50S

The NetScaler SDX 26xxx-50S appliances are 2U appliances. These appliances have two 14-core processors and 256 GB of memory. The appliances have four 50G and sixteen 25G network ports.

For information about software releases supported on the NetScaler hardware platforms, see [NetScaler SDX hardware-software compatibility matrix](#).

The following figure shows the front panel of the NetScaler SDX 26xxx-50S appliances.

Figure 1. NetScaler SDX 26xxx-50S, front panel




The NetScaler SDX 26xxx-50S appliances have the following ports:

- RS232 serial console port.
- One 10/100/1000Base-T RJ45 copper Ethernet LOM Port. Use this port to remotely monitor and manage the appliance independently of the NetScaler software.
- Two 10/100/1000Base-T RJ45 copper Ethernet Management Ports, numbered 0/1 and 0/2. These ports are used to connect directly to the appliance for NetScaler administration functions.
- Four 50G ports, numbered 50/1 to 50/4.
- Sixteen 25G ports, numbered 25/1 through 25/16. For information about supported transceivers per port, see [25G, 40G, 50G, and 100G ports](#).

The following figure shows the back panel of the SDX 26xxx-50S appliances.

Figure 2. NetScaler SDX 26xxx-50S, back panel



The following components are visible on the back panel of the NetScaler SDX 26xxx-50S appliances:

- Two 240GB and six 960GB removable solid-state drive (SSD).

Note: Drive densities might increase as components become EOL but its size is never smaller than the original.

- Power switch, which turns power to the appliance on or off.
 - If the OS is functional, press the switch for less than two seconds to shut down gracefully.
 - If the OS is not responsive, press the power switch for more than 4 seconds, for a hard shutdown.
- Two hot-swappable 100–240 VAC, 1200 W power supply modules. Max power consumption is 764 W. Typical power consumption is 628 W. Each power supply has an LED indicating its status:

LED COLOR	LED INDICATES
OFF	No power to any power supply on the appliance.
Flashing RED	No power to this power supply.
Flashing GREEN	Power supply is in standby mode.
GREEN	Power supply is functional.
RED	Power supply failure.
Flashing RED and GREEN	Warning;(OVP/UVP/OCP/OTP/Fan); OVP = Over Voltage Protection; UVP = Under Voltage Protection; OCP = Over Current Protection; OTP = Over Temperature Protection

- **Disable alarm button.** Press this button to silence the power alarm when the power supply loses input power or when a power supply is malfunctioning.
- **Non-Maskable Interrupt (NMI) button:** is recessed to prevent unintentional activation and used at the request of Technical Support to start a core dump. To press this red button, use a pen, pencil, or other pointed object. The NMI Button is also available remotely over the network in

the LOM GUI, in the Remote Control menu. For more information, see [Lights out management port of the NetScaler SDX appliance](#).


NetScaler SDX 26000-100G

The NetScaler SDX 26000-100G appliance is a 2U appliance. This platform has two 14-core processors and 256 GB memory. The appliance provides a total of 8 network ports.

For information about the software releases supported on the NetScaler hardware platforms, see [NetScaler SDX Hardware-Software Compatibility Matrix](#)[(/en-us/netscaler-hardware-platforms/sdx/sdx-hardware-software-compatibility-matrix.html)].

The illustrations in this guide are only indicative and might not represent the actual NetScaler appliance.


Figure 1. NetScaler SDX 26000-100G front panel



The NetScaler SDX 26000-100G appliances have the following ports:

- RS232 serial Console Port.
- One 10/100/1000Base-T RJ45 copper Ethernet LOM Port. Use this port to remotely monitor and manage the appliance independently of the NetScaler software.
- Two 10/100/1000Base-T RJ45 copper Ethernet Management Ports, numbered 0/1 and 0/2. These ports are used to connect directly to the appliance for NetScaler administration functions.
- Eight 100G Ethernet Ports, numbered 100/1 to 100/8 from left to right. For information about supported transceivers per port, see [25G, 40G, 50G, and 100G ports](#).

Figure 2. NetScaler SDX 26000-100G back panel



The following components are visible on the back panel of the SDX 26000-100G appliances:

- Two 240 GB removable boot solid-state drives (SSDs).
- Six 960 GB removable SSDs.

Note: Drive densities might increase as components become EOL but its size is never smaller than the original.

- Power switch, which turns power to the appliance on or off.
 - If the OS is functional, press the switch for less than two seconds to power down the system with a graceful shutdown.
 - If the OS is not responsive, press the power switch for more than 4 seconds to force the power off.
- Two power supplies, 100–240 VAC. Each power supply has an LED indicating its status, as described in the following table:

LED COLOR	LED INDICATES
OFF	No power to any power supply on the appliance.
Flashing RED	No power to this power supply.
Flashing GREEN	Power supply is in standby mode.
GREEN	Power supply is functional.

LED COLOR	LED INDICATES
RED	Power supply failure.
Flashing RED and GREEN	Warning (OVP/UVP/OCP/OTP/Fan); OVP = Over Voltage Protection; UVP = Under Voltage Protection; OCP = Over Current Protection; OTP = Over Temperature Protection.

- **Disable alarm button**, which is functional only when the appliance has dual redundant, hot-swappable power supplies. Press this button to silence the power alarm when one of the two power supplies loses input power or when a power supply is malfunctioning.
- **Non-Maskable Interrupt (NMI) Button**, used at the request of Technical Support to initiate a core dump. To press this red button, which is recessed to prevent unintentional activation, use a pen, pencil, or other pointed object. The NMI Button is also available remotely over the network in the LOM GUI, in the Remote Control menu. For more information about the lights out management port of the appliance, see [Lights out management port of the NetScaler SDX appliance](#).

Common hardware components

NOTE:

This section does not apply to the SDX 9200 platform. For information specific to the SDX 9200, see [NetScaler SDX 9200 platform](#).

Each platform has front panel and back panel hardware components. The front panel has an LCD display and an RS232 serial console port. The number, type, and location of ports—copper Ethernet, copper and fiber 1G SFP, 10GSFP+, and 40G QSFP+. The back panel provides access to the fan and the field-replaceable units (power supplies, CompactFlash cards, and solid-state and hard-disk drives).

LCD display and LED status indicators

On some NetScaler SDX appliances, the LCD on the front panel displays the appliance's model number, but the number shown might not be the licensed model number. To view the licensed model number of any SDX appliance, log on to the Management Service and check the licensed model number in the top left corner of the screen. For example, if you have purchased an SDX 11515 license, the LCD screen displays CITRIX NSSDX-11500, and the Management Service screen displays NetScaler SDX (11515).

The LCD backlight on the NetScaler appliance is always on.

On the appliance's back panel, system status LEDs indicate the overall status of the appliance. The following table describes the indicators of the system status LED.

On the appliance's back panel, power status LEDs indicate the status of each power supply.

NOTE

System status LEDs are available on only some NetScaler SDX appliances.

LED COLOR	LED INDICATES
OFF	No power

LED COLOR	LED INDICATES
Green	Appliance is receiving power
Red	Appliance has detected an error

The port LEDs show whether a link is established and traffic is flowing through the port. The following table describes the LED indicators for each port. There are two LED indicators for each port type.

NOTE

This section applies to all the appliances.

LED port-status indicators

PORT TYPE	LED COLOR	LED INDICATES
50 Gbps	Off	A link has not been established
	Blinking amber	Indicates a problem with the link
	Solid green	Indicates a valid link with no active traffic
	Blinking green	Indicates a valid link with active traffic

PORT TYPE	LED	LED	LED COLOR	LED INDICATES
10 Gbps	Top	Speed	Off	No connection.
			Solid blue	Traffic rate of 10 gigabits per second
	Bottom	Link/Activity	Off	No link

PORT TYPE	LED	LED	LED COLOR	LED INDICATES
			Solid green	A link is established but no traffic is passing through the port
			Blinking green	Traffic is passing through the port
1G SFP (1 Gbps)	Left	Link/Activity	Off	No link
			Solid green	A link is established but no traffic is passing through the port
			Blinking green	Traffic is passing through the port
	Right	Speed	Off	No connection
			Yellow	Traffic rate of 1 gigabit per second
Management (RJ45)	Left	Speed	Off	No connection, or a traffic rate of 10 megabits per second (Mbps)
			Green	Traffic rate of 100 Mbps

PORT TYPE	LED	LED	LED COLOR	LED INDICATES
			Amber	Traffic rate of 1 gigabit per second
	Right	Link/Activity	Off	No link
			Solid yellow	A link is established but no traffic is passing through the port

On each power supply, a bicolor LED indicator shows the condition of the power supply.

LED power supply indicators

POWER SUPPLY TYPE	LED COLOR	LED INDICATES
AC	OFF	No power to any power supply
	Flashing RED	No power to this power supply
	Flashing GREEN	The power supply is in standby mode
	GREEN	The power supply is functional
	RED	Power supply failure

Ports

Ports are used to connect the appliance to external devices. NetScaler appliances support RS232 serial ports, 10/100/1000Base-T copper Ethernet ports, 1-gigabit copper and fiber 1G SFP ports, 10-gigabit fiber SFP+, 25 GB, 40 GB QSFP+, 50 GB, and 100 GB ports. All NetScaler appliances

have a combination of some or all of these ports. For details on the type and number of ports available on your appliance, see the section describing that platform.

NOTE

NetScaler appliances leverage the Receive Side Scaling (RSS) feature to optimally distribute and process traffic. The RSS implementation is based on Toeplitz hashing which distributes the packet among its queues based on the 5-tuple: TCP/IP protocol, source IP address, destination IP address, source port, and destination port.

RS232 serial port

The RS232 serial console port provides a connection between the appliance and a computer, allowing direct access to the appliance for initial configuration or troubleshooting.

All hardware platforms ship with an appropriate serial cable used to connect your computer to the appliance. For instructions on connecting your computer to the appliance, see [Installing the Hardware](#).

Copper Ethernet ports

The copper Ethernet ports installed on many models of the appliance are standard RJ45 ports.

There are two types of copper Ethernet ports that are installed on your appliance:

- 10/100BASE-T port

The 10/100BASE-T port has a maximum transmission speed of 100 megabits per second (Mbps). Most platforms have at least one 10/100BASE-T port.

- 10/100/1000BASE-T port

The 10/100/1000BASE-T port has a maximum transmission speed of 1 gigabit per second, 10 times faster than the other type of copper Ethernet port. Most platforms have at least one 10/100/1000Base-T port.

To connect any of these ports to your network, you plug one end of a standard Ethernet cable into the port and plug the other end into the appropriate network connector.

Management ports

Management ports are standard copper Ethernet ports (RJ45), which are used for direct access to the appliance for system administration functions.

1G SFP, 10G SFP+, and 40G QSFP+ ports

A 1G SFP port can operate at a speed of 1 Gbps. It accepts either a copper 1G SFP transceiver for operation as a copper Ethernet port, or a fiber 1G SFP transceiver for operation as a fiber optic port.

The 10G SFP+ and Base-T 10G, 40G QSFP+ are high-speed ports that can operate at speeds of up to 40 Gbps. You need a fiber optic cable to connect to a port. If the other end of the fiber optic cable is attached to a 1G SFP port, the 10G SFP+ port negotiates to match the speed of the 1G SFP port.

Ports compatibility

On some appliances, the 10G slot supports **copper** 1G transceivers, which can operate at up to 1 Gbps in a 10 Gbps slot.

NOTE

- Certain platforms have 10G slots that do not support copper transceivers. Check with your account representative for support details.
- You cannot insert a fiber 1G transceiver into a 10G slot.
- You cannot insert a 10G transceiver into a 1G slot.


The 10G SFP+ modules are dual-speed capable and support both 1 Gbps and 10 Gbps, depending on the peer switch that the model connects to.


25G, 40G, 50G, and 100G ports

A 100G port can support five speeds: 10G, 25G, 40G, 50G, and 100G. 1G speed is not supported on the 100G port. 50G and 100G ports use the same transceiver. 40G QSFP+ are high-speed ports that can operate at speeds of up to 40 Gbps. The appliance determines the speed, and not the port.

Only 50G/100G (QSFP28) transceivers can be directly used on a QSFP28 interface. Use a QSA28 adapter on a QSFP28 interface to use 10G (SFP+) and 25G (SFP28) transceivers. Use a 40G QSFP+ transceiver in the 50G/100G port if the peer is 40G.

The following diagram shows the transceiver compatibility.



**NOTES:**

- Fortville firmware 8.70 does not support 1G connectivity. Dual speed 10G/1G SFP+ (both SR LC and LR LC) don't operate at 1G, if Fortville firmware 8.70 is used.
- Use `i40e_firmware_view.sh` to check the FW version.
- Ensure that the FW is at version 7.0. To downgrade the firmware version to 7.0, type:

```
[root@netscaler-sdx ~]# echo "version 7.00" > /opt/xensource/packages/
files/sdx-firmware/i40efirmware-target
```

```
[root@netscaler-sdx ~]# echo "version 7.00" > /opt/xensource/packages/
files/sdx-firmware/i40efirmware-direction
```

```
[root@netscaler-sdx ~]#reboot
```

Most tables have the following columns:

- Description: The price list description of the part.
- Transmit Wavelength: The nominal transmit wavelength.
- Cable/Fiber Type: Fiber characteristics affect the maximum transmit distance achievable, especially with 10G on multi-mode fiber (MMF), where various dispersion components become dominant. For more information, see [<http://www.thefoa.org/tech/ref/basic/fiber.html>](http://www.thefoa.org/tech/ref/basic/fiber.html).
- Typical Reach: Maximum transmit distance.
- Products: Some chassis are available with different media options. Use the appropriate data sheet to confirm that your particular chassis type supports the media.

1G pluggable media

The following table lists the maximum distance specifications for 1G transceivers.

Copper 1G SFP distance specifications

Description: NetScaler 1G SFP Ethernet copper (100 m) - 4 pack

Transmitter wavelength (nm): Not applicable

Cable type: Category 5 (Cat-5) copper cable

Typical reach (m): 100 m

Applicable platforms:

- SDX 14000
- SDX 16000
- SDX 22040/22060/22080/22100/22120
- SDX 24100/24150

Short reach fiber 1G SFP distance specifications

Description: NetScaler 1G SFP Ethernet SX (300 m) - 4 pack

Transmitter wavelength (nm): 850 nm (nominal)

Fiber type: 50/125um MMF, 2000MHz-km (OM3) Typical Reach (m): 550 m

Fiber type: 50/125um MMF, 500MHz-km (OM2) Typical Reach (m): 550 m

Fiber type: 50/125um MMF, 400MHz-km Typical Reach (m): 550 m

Fiber type: 62.5/125um MMF, 200MHz-km (OM1) Typical Reach (m): 300 m

Fiber type: 62.5/125um MMF, 160MHz-km Typical Reach (m): 300 m

Applicable platforms: SDX 22040/22060/22080/22100/22120 SDX 24100/24150

Short reach fiber 1G SFP distance specifications

Description: NetScaler 1G SFP Ethernet short range (300 m) - Single

Transmitter wavelength (nm): 850 nm (nominal)

Fiber type: 50/125um MMF, 2000MHz-km (OM3) Typical Reach (m): 550 m

Fiber type: 50/125um MMF, 500MHz-km (OM2) Typical Reach (m): 550 m

Fiber type: 50/125um MMF, 400MHz-km Typical Reach (m): 550 m

Fiber type: 62.5/125um MMF, 200MHz-km (OM1) Typical Reach (m): 275 m

Fiber type: 62.5/125um MMF, 160MHz-km Typical Reach (m): 220 m

Applicable platforms:

- SDX 22040/22060/22080/22100/22120
- SDX 24100/24150

Long reach fiber 1G SFP distance specifications

Description: NetScaler 1G SFP Ethernet LX - Single

Transmitter wavelength (nm): 1310 nm (nominal)

Fiber type: 9/125um SMF

Typical reach (m): 10 km

Applicable platforms:

- SDX 22040/22060/22080/22100/22120

- SDX 24100/24150

Long reach fiber 1G SFP distance specifications

Description: NetScaler 1G SFP Ethernet long range (10 km) - Single

Transmitter wavelength (nm): 1310 nm (nominal)

Fiber type: 9/125um SMF

Typical reach (m): 10 km

Applicable platforms:

- SDX 22040/22060/22080/22100/22120
- SDX 24100/24150

10 GE pluggable media

The following table lists the maximum distance specifications for 10G transceivers.

Short reach fiber 10G SFP+ distance specifications

Description: NetScaler 10G SFP+, Ethernet Short Range (300 m) - Single

Transmitter wavelength (nm): 850 nm (nominal)

Fiber types 50/125um MMF, 2000MHz-km (OM3), typical reach (m): 300 m 50/125um MMF, 500MHz-km (OM2), typical reach (m): 82 m 50/125um MMF, 400MHz-km, typical reach 66 m 62.5/125um MMF, 200MHz-km (OM1), typical reach: 33 m 62.5/125um MMF, 160MHz-km, typical reach: 26 m

Applicable platforms:

- SDX 9100
- SDX 14000
- SDX 14000-40G
- SDX 14000-40S
- SDX 14000 FIPS
- SDX 15000
- SDX 15000-50G (adapter required to connect 50G ports to 10G transceiver)
- SDX 16000 (adapter required to connect 100G ports to 10G transceiver)

- SDX 22040/22060/22080/22100/22120
- SDX 24100/24150
- SDX 25000-40G
- SDX 26000 (adapter required to connect 50G ports to 10G transceiver)
- SDX 26000-50S (adapter required to connect 50G ports to 10G transceiver)
- SDX 26000-100G (adapter required to connect 100G ports to 10G transceiver)

Long reach fiber 10G SFP+ distance specifications

Description: NetScaler 10G SFP+ Ethernet long range (10 km) - Single

Transmitter wavelength (nm): 1310 nm (nominal)

Fiber type: 9/125um SMF

Typical reach (m): 10 km

Applicable platforms:

- SDX 9100
- SDX 14000
- SDX 14000-40G
- SDX 14000-40S
- SDX 14000 FIPS
- SDX 15000
- SDX 15000-50G (adapter required to connect 50G ports to 10G transceiver)
- SDX 16000 (adapter required to connect 100G ports to 10G transceiver)
- SDX 22040/22060/22080/22100/22120
- SDX 24100/24150
- SDX 25000-40G
- SDX 26000 (adapter required to connect 50G ports to 10G transceiver)
- SDX 26000-50S (adapter required to connect 50G ports to 10G transceiver)
- SDX 26000-100G (adapter required to connect 100G ports to 10G transceiver)

Field replaceable units

NetScaler field replaceable units (FRUs) are components that can be removed from the appliance and replaced by the user or a technician at the user's site. The FRUs in a NetScaler SDX appliance can include DC or AC power supplies, solid-state drives (SSDs), and a direct-attach cable (DAC).

All NetScaler FRUs must be purchased from Citrix. Components not provided by NetScaler are not supported on NetScaler appliances. Contact your NetScaler sales representative to buy FRUs for your appliance.

NOTE

The figures in this section are illustrative. The actual product might vary depending on the model number.

This topic includes the following sections:

- Power Supply
- Solid-State Drive
- Direct Attach Cable

Power supply

For appliances containing two power supplies, the second power supply acts as a backup. Some appliances can accommodate four power supplies, and require two power supplies for proper operation. The third and fourth power supplies act as backup. All power supplies must be of the same type (AC or DC).

The SDX appliance ships with a country-specific power cord and a plug.

For power-supply specifications, see the [SDX hardware platforms](#) topic, which describes the various platforms.

NOTE

If you suspect that a power-supply fan is not working, check the system health through the SDX GUI or lights out management. For more information, see [System Health Monitoring](#) and [Lights out management port of the NetScaler SDX appliance](#).

On each power supply, a bicolor LED indicator shows the condition of the power supply. Green means that the appliance receives power. Red means that the power supply has detected an error. For more information about the LED power supply indicator, see "LCD Display and LED Status Indicators" in [Common hardware components](#).

Electrical Safety Precautions for Power Supply Replacement

- Ensure that the appliance has a direct physical connection to earth ground during normal use. When installing or repairing an appliance, always connect the ground circuit first and disconnect it last.
- NetScaler SDX appliances do not contain user serviceable components, except as noted in the (FRU) section. Do not attempt to open the chassis.
- Always unplug the power cord from the power source before removing the power supply.
- Never touch a power supply when the power cord is plugged in. As long as the power cord is plugged in, line voltages are present in the power supply even if the power switch is turned off.

For the complete list of safety precautions, see [Safety, cautions, warnings, and other information](#).

Replace an AC or a DC power supply

If the appliance has only one power supply, you have to shut down the appliance before replacing the power supply. If the appliance has two power supplies, you can replace one power supply without shutting down the appliance, provided the other power supply is working.

To install or replace an AC or a DC power supply on a NetScaler SDX appliance

1. Align the semicircular handle perpendicular to the power supply. Loosen the thumbscrew and press the lever toward the handle and pull out the existing power supply, as shown in the following figure.

Figure 1. Remove an Existing AC Power Supply



Figure 2. Remove an Existing DC Power Supply



2. Carefully remove the new power supply from its box.
3. On the back of the appliance, align the power supply with the power supply slot.
4. Insert the power supply into the slot and press against the semicircular handle until you hear the power supply snap into place.

Figure 3. Insert the Replacement AC Power Supply



Figure 4. Insert the Replacement DC Power Supply



5. Connect the power supply to a power source. If connecting all power supplies, plug separate power cords into the power supplies and connect them to separate wall sockets.

NOTE

The NetScaler SDX appliance emits a high-pitched alert if the following conditions are met:

- One power supply fails.
- If you connect only one power cable to the appliance in which two power supplies are installed.

To silence the alarm, press the button labeled Alarm Disable on the back panel of the appliance.

The **Disable alarm button** is functional only when the appliance has two power supplies.

Solid-State drives

In a NetScaler SDX appliance, the solid-state drives (SSD) store system bootup files and VDI of Management Service instances. Some support RAID and some do not.

SSDs with RAID support are hot swappable.

NOTE:

For more information about managing RAID allocation, see [Managing RAID disk allocation on the SDX platform](#).

Direct attach cable

A direct attach cable (DAC) assembly is a high performance integrated duplex data link for bi-directional communication. The cable is compliant with the IPF MSA (SFF-8432) for mechanical form factor and SFP+ MSA for DACs. The cable, which can be up to 5 meters long, is data-rate agnostic. Supporting speeds more than 10 Gbps, it is a cost-effective alternative to optical links (SFP+ transceivers and fiber optic cables.) The transceiver with DAC is hot-swappable. You can insert and remove the transceiver with the attached cable without shutting down the appliance. The NetScaler appliance supports only the passive DAC.

Points to note:

- DAC is supported only on 10G ports. Do not insert a DAC into a 1G port.

- Do not attempt to unplug the integrated copper cable from the transceiver and insert a fiber cable into the transceiver.


NOTE

Autonegotiation is not supported on an interface to which a DAC is connected.

To replace a DAC


1. To remove the DAC, pull the tab on the top of the DAC, and then pull the DAC out of the port.

Figure 9. Remove a DAC from the 10G port



2. To insert the DAC, slide it into the 10G port on the appliance. You hear a click when the DAC properly fits into the port.

Figure 10. Insert a DAC into the 10G port



Transceivers for MPX 9200 and 17000 platforms

This section describes the transceivers, direct-attach copper (DAC) cables, and active optical cables (AOCs) that are supported on the NetScaler MPX 9200 and MPX 17000 platforms. It also provides answers to frequently asked questions and real-world interoperability use cases to help you plan and validate your deployments.

The MPX 9200 provides 8 SFP56 ports (10G/25G/50G) and has no QSFP-class ports. The MPX 17000 provides the same 8 SFP56 ports along with 4 QSFP56 ports (40G/100G/200G). As a result, all 40G, 100G, and 200G modules are supported only on the MPX 17000.

Use the following topics to find the information that you need:

- [Supported transceivers, DACs, and AOCs](#): The platform port topology and the complete matrix of supported transceivers, DACs, and AOCs, including speed, form factor, reach, and per-platform support.
- [FAQs](#): Answers to common questions about supported speeds, form factors, and module types for the MPX 9200 and MPX 17000 platforms.
- [Transceiver Interop use cases](#): Validated real-world deployment topologies and negotiation details for connecting the MPX 17000 to industry-standard data center switches.

Transceiver, DACs, AOCs Master Matrix

Platform port topology

PLATFORM	FORM FACTOR	NIC CONFIGURATION	SFP56 PORTS (10G/25G/50G)	QSFP56 PORTS (40G/100G/200G)
NetScaler MPX/SDX 9200	1U	2x ConnectX-7 quad-port 50G	8	0
NetScaler MPX/SDX 17000	2U	2x ConnectX-7 quad-port 50G AND 2x	8	4

PLATFORM	FORM FACTOR	NIC CONFIGURATION	SFP56 PORTS (10G/25G/50G)	QSFP56 PORTS (40G/100G/200G)
		ConnectX-7 dual-port 200G		

Implication: The 9200 has no QSFP-class ports. Any 40G, 100G, or 200G transceiver, DAC, or AOC is supported only on the MPX/SDX 17000 via the two 200G NICs. The 8x SFP56 ports on both platforms accept 10G SFP+, 25G SFP28, and 50G SFP56 modules, DACs, and AOCs.

Supported Transceivers, DACs & AOCs' matrix

NOTE:

When customers order a transceiver using a Citrix SKU, the received package or box may show only the Citrix MPN.

CITRIX SKU	CITRIX MPN	DESCRIPTION	SPEED	FORM FACTOR	ENCODING	REACH / MEDIA	FORWARD ERROR CORRECTION	AUTO-NEG	MPX / SDX 9200	MPX / SDX 17000
6003085	9530401	NetScaler 50GbE SFP56 SR LC, OM3 70m ; NS 920	50G	SFP56	PAM4	OM3 70m / OM4 100m	RS-FEC	Yes	Supported	Supported

CIT RIX SKU	CIT RIX MPN	DES CRI PTION	SPE ED	FOR M FAC TOR	ENC ODI NG	REA CH / MED IA	FOR WA RD ERR OR COR REC TION	AUT O- NEG	MPX / SDX 9200	MPX / SDX 1700 0
		0/17 000								
600 308 8	953 040 2	Net Scal er 25G bE SFP 28 SR LC, OM 3 70m ; NS 920 0/17 000	25G	SFP 28	NRZ	OM 3 70m / OM 4 100 m	FC- FEC	Yes	Sup port ed	Sup port ed
600 309 0	953 040 3	Net Scal er 10G bE SFP + SR LC, OM 3 300 m; NS	10G	SFP +	NRZ	OM 3 300 m / OM 4 400 m	Non e	Yes	Sup port ed	Sup port ed

CIT RIX SKU	CIT RIX MPN	DES CRI PTION	SPE ED	FOR M FAC TOR	ENC ODI NG	REA CH / MED IA	FOR WA RD ERR OR COR REC TION	AUT O- NEG	MPX / SDX 9200	MPX / SDX 1700 0
		920 0/17 000								
600 314 2	953 040 5	Net Scal er 50G bE SFP 56 LR LC, OS2 10k m; NS 920 0/17 000	50G	SFP 56	PAM 4	SMF (OS 2) 10 km	RS- FEC	Yes	Sup port ed	Sup port ed
600 312 7	953 040 6	Net Scal er 25G bE SFP 28 LR LC, OS2 10k m; NS	25G	SFP 28	NRZ	SMF (OS 2) 10 km	FC- FEC	Yes	Sup port ed	Sup port ed

CIT RIX SKU	CIT RIX MPN	DES CRI PTION	SPE ED	FOR M FAC TOR	ENC ODI NG	REA CH / MED IA	FOR WA RD ERR OR COR REC TION	AUT O- NEG	MPX / SDX 9200	MPX / SDX 1700 0
		920 0/17 000								
600 312 8	953 040 7	Net Scal er 10G bE SFP + LR LC, OS2 10k m; NS 920 0/17 000	10G	SFP +	NRZ	SMF (OS 2) 10 km	Non e	Yes	Sup port ed	Sup port ed
600 309 3	953 050 1	Net Scal er 50G BAS E- CU SFP 56- to- SFP 56 DAC , 3m	50G	SFP 56	PAM 4	Pas sive cop per	RS- FEC	Yes	Sup port ed	Sup port ed

CIT RIX SKU	CIT RIX MPN	DES CRI PTION	SPE ED	FOR M FAC TOR	ENC ODI NG	REA CH / MED IA	FOR WA RD ERR OR COR REC TION	AUT O- NEG	MPX / SDX 9200	MPX / SDX 1700 0
		(Mel lano x/ Juni per); NS 920 0/17 000								
600 308 6	953 050 7	Net Scal er 50G BAS E- CU SFP 56- to- SFP 56 DAC , 2.5 m; NS 920 0/17 000	50G	SFP 56	PAM 4	Pas sive cop per	RS- FEC	Yes	Sup port ed	Sup port ed
600 313 1	953 040 9	Net Scal er	200 G	QSF P56	PAM 4	OM 3 70m	RS- FEC	Yes	Not pres ent	Sup port ed

CIT RIX SKU	CIT RIX MPN	DES CRI PTION	SPE ED	FOR M FAC TOR	ENC ODI NG	REA CH / MED IA	FOR WA RD ERR OR COR REC TION	AUT O- NEG	MPX / SDX 9200	MPX / SDX 1700 0
		200 GbE QSF P56 SR4 MP O-1 2, OM 3 70m ; NS 170 00			(4x5 0G)	/ OM 4 100 m				
600 313 2	953 041 0	Net Scal er 100/ 40G bE Dual - Rate QSF P28 SR4 MP O-1 2, OM 3 70m ; NS	100/ 40G	QSF P28	NRZ (4x2 5G at 100 G; 4x1 0G at 40G)	OM 3 70m / OM 4 100 m	RS- FEC at 100 G; Non e at 40G	100 G: Yes; 40G : Forc ed spe ed both end s	Not pres ent	Sup port ed

CIT RIX SKU	CIT RIX MPN	DES CRI PTION	SPE ED	FOR M FAC TOR	ENC ODI NG	REA CH / MED IA	FOR WA RD ERR OR COR REC TION	AUT O- NEG	MPX / SDX 9200	MPX / SDX 1700 0
		170 00								
600 313 3	953 0411	Net Scal er 40G bE QSF P+ BiDi SR LC, OM 3 100 m; NS 170 00	40G	QSF P+	NRZ (4x1 0G)	OM 3 100 m / OM 4 150 m	Non e	Forc ed spe ed to 40G ; Citri x MP N requ ired on both end s	Not pres ent	Sup port ed
600 313 5	953 041 3	Net Scal er 100 GbE QSF P28 LR4 LC, OS2 10k m; NS	100 G	QSF P28	NRZ (4x2 5G CW DM)	SMF (OS 2) 10 km	Non e	Yes	Not pres ent	Sup port ed

CIT RIX SKU	CIT RIX MPN	DES CRI PTION	SPE ED	FOR M FAC TOR	ENC ODI NG	REA CH / MED IA	FOR WA RD ERR OR COR REC TION	AUT O- NEG	MPX / SDX 9200	MPX / SDX 1700 0
		170 00								
600 314 3	953 041 6	Net Scal er QSF P28- to- SFP 28 cag e ada pter; NS 170 00	Ada pter for 25G or 10G	QSF P28- to- SFP 28	Pas sthr oug h (1 lane)	N/A - acce pts SFP 28 (25 G) or SFP + (10 G) optic	Inhe rits from inse rted optic	Inhe rits from inse rted optic	Not pres ent	Sup port ed
600 314 4	953 051 3	Net Scal er 40G bE QSF P+ Pas sive DAC , 5m; NS 170 00	40G	QSF P+ DAC	NRZ (4x1 0G)	Pas sive cop per	Non e	Yes	Not pres ent	Sup port ed

CIT RIX SKU	CIT RIX MPN	DES CRI PTION	SPE ED	FOR M FAC TOR	ENC ODI NG	REA CH / MED IA	FOR WA RD ERR OR COR REC TION	AUT O- NEG	MPX / SDX 9200	MPX / SDX 1700 0
600 314 5	953 051 4	Net Scal er 200 G QSF P56- to- QSF P56 Pas sive Twin ax 2.5 m; NS 170 00	200 G	QSF P56 DAC	PAM 4 (4x5 0G)	Pas sive cop per	RS- FEC	Yes	Not pres ent	Sup port ed
600 314 6	953 051 7 *	Net Scal er 100 GbE QSF P28 AO C, 20m ; NS 170 00	100 G	QSF P28 AO C	NRZ (4x2 5G)	Acti ve opti cal (inte grat ed)	RS- FEC	Yes	Not pres ent	Sup port ed

CIT RIX SKU	CIT RIX MPN	DES CRI PTION	SPE ED	FOR M FAC TOR	ENC ODI NG	REA CH / MED IA	FOR WA RD ERR OR COR REC TION	AUT O- NEG	MPX / SDX 9200	MPX / SDX 17000
600 314 7	953 052 2	Net Scal er 200 GbE QSF P56 AOC, 20m (Mel lano x); NS 170 00	200 G	QSF P56 AOC	PAM 4 (4x5 0G)	Acti ve opti cal (inte grat ed)	RS- FEC	Yes	Not pres ent	Sup port ed

*Note:

Qualified as 100G only. Although the AOC may appear to partially train at lower per-lane rates in some lab environments or automation reports, it does not advertise lower-speed compliance in its EEPROM. Lower-speed operation requires manual force on both ends and is not supported. For 10G / 25G / 40G / 50G links, use the dedicated Citrix MPN qualified for that speed.

Pre-Deployment checklist

1. Confirm platform and port count. MPX/SDX 9200 is SFP56-only (8 ports). MPX/SDX 17000 has 8x SFP56 + 4x QSFP56 ports.
2. Match the form factor to MPN. SFP+ / SFP28 / SFP56 go in SFP56 ports; QSFP+ / QSFP28 / QSFP56 go in QSFP56 ports on the MPX/SDX 17000.

3. Verify reach and fiber type:
 - SR / SR4 -> OM3 or OM4 multimode
 - LR / LR4 -> OS2 single-mode
 - MPO-12 SR4 -> Type B polarity
4. Configure the FEC on the peer switch to match the speed and modulation (see the FEC table in the FAQ KB).
5. For Cisco Nexus 9000 deployments, enable "service unsupported-transceiver" on the Cisco switch side to allow non-Cisco-branded optics.
6. For 40G QSFP+ links and 40G via 9530410 dual-rate, force 40G on both host and switch; auto-negotiation does not converge.
7. For 40G BiDi via 9530411, install the same Citrix MPN on both ends.
8. For LR4 links (9530413), verify OS2 fiber cleanliness and polarity. If a link does not come up, switch-side RX power at the floor (-40 dBm) on all lanes indicates a fiber-path issue, not a transceiver issue.
9. Verify MPO polarity and LC duplex polarity before powering up the link. Most "link-down with clean optics" tickets resolve to a polarity issue.
10. For 9530416 adapter deployments, install a qualified Citrix SFP28 or SFP+ optic into the adapter and configure host and peer switch for the lower link speed (25G or 10G).

Glossary

TERM	DEFINITION
AOC	Active Optical Cable => fixed-length cable with integrated optical electronics at each end
BiDi	Bi-Directional => uses two wavelengths over a single fiber pair, enabling 40G over duplex LC (instead of 8-fiber MPO)
DAC	Direct Attach Copper => fixed-length passive copper cable with integrated SFP/QSFP connectors
FC-FEC	Firecode FEC (IEEE 802.3 Clause 74 / BASE-R) => used at 25G NRZ

TERM	DEFINITION
LR / LR4	Long Reach - single-mode-fiber optic, typically up to 10 km
MPO-12	Multi-fiber Push-On 12-fiber connector used by 40G/100G/200G SR4 modules
NRZ	Non-Return-to-Zero binary signaling (2-level) used at 10G/25G/100G NRZ links
OS2	Single-mode fiber for LR4 and other long-reach optics
OM3 / OM4	Multi-mode fiber grades - OM3 (300 m at 10G, 70 m at SR4); OM4 (extended reach)
PAM4	4-Level Pulse-Amplitude Modulation => used at 50G/200G/400G to double bit density per symbol
QSFP+	4-lane 40G form factor - superseded by QSFP28
QSFP28	4-lane form factor for 100G NRZ; backwards-compatible with 40G in dual-rate modules
QSFP56	4-lane form factor for 200G PAM4
RS-FEC	Reed-Solomon FEC (IEEE 802.3 Clause 91 / 108) => used at 50G PAM4, 100G SR4 NRZ, 200G PAM4
SFP+	1-lane 10G form factor
SFP28	1-lane 25G form factor
SFP56	1-lane 50G PAM4 form factor
SR / SR4	Short Reach - multimode-fiber optic (single-fiber-pair for SR; 8-fiber MPO for SR4)

Transceiver, DAC, AOC FAQs

Speed & Form Factor

Q1: Are 1G SFP transceivers supported?

No. The SFP56 and QSFP56 ports on the 50G and 200G NICs of NetScaler MPX/SDX 9200 and 17000 appliances do not negotiate to 1G SGMII or 1000BASE-X (Ethernet over fiber). The minimum supported port speed is 10G.

Q2: Which 10G transceivers are supported?

NetScaler-qualified 10GbE SFP+ optics are supported on both the MPX/SDX 9200 and MPX/SDX 17000 platforms. This includes SR (multimode, OM3, up to 300m) and LR (single-mode, up to 10km) transceivers. For more information, see the [Transceiver, DACs, AOCs' matrix](#).

Q3: Which 25G transceivers are supported?

NetScaler-qualified 25GbE SFP28 optics are supported on both the MPX/SDX 9200 and MPX/SDX 17000 platforms. Supported options include SR (multimode, OM3 up to 70m / OM4 up to 100m) and LR (single-mode, up to 10km) transceivers. For more information, see the [Transceiver, DACs, AOCs' matrix](#)[(/en-us/netscaler-hardware-platforms/mpx/mpx-9200-17000-master-matrix)].

Q4: Is 50G supported?

Yes. NetScaler-qualified 50GbE SFP56 SR and LR optics, as well as 50G SFP56 Direct-attach copper (DAC) cables, are supported on both the MPX/SDX 9200 and MPX/SDX 17000 platforms. For more information, see the [Transceiver, DACs, AOCs' matrix](#)[(/en-us/netscaler-hardware-platforms/mpx/mpx-9200-17000-master-matrix)].

Q5: Are 40G QSFP+ transceivers supported?

Yes, only on the MPX and SDX 17000 platform. The Citrix-qualified options are:

- **9530411 (Citrix MPN)** (40GbE QSFP+ BiDi SR LC optical): Ensure that the same NetScaler-branded module is used on both ends of the connection. Mixing this transceiver with a Cisco-branded 40G-BiDi module on the switch side can cause link flapping.
- **9530513 (Citrix MPN)** (40GbE QSFP+ DAC, 5 m passive copper): A DAC alternative to BiDi optics, ideal for short-reach 40G links.

Q6: Are 100G/40G QSFP28 transceivers supported?

Yes, only on the MPX and SDX 17000 platform. The Citrix-qualified options are:

- **9530410 (Citrix MPN):** 100/40GbE Dual-Rate QSFP28 SR4 (MPO-12 MMF)
- **9530413 (Citrix MPN):** 100GbE QSFP28 LR4 (LC, SMF, 10 km)
- **9530517 (Citrix MPN):** 100GbE QSFP28 AOC (20 m)

Q7: Is 200G supported?

Yes, only on the QSFP56 ports of the dual-port 200G NICs on the MPX/SDX 17000 platform. The Citrix-qualified options are:

- **9530409 (Citrix MPN):** 200GbE QSFP56 SR4 (MPO-12 MMF)
- **9530514 (Citrix MPN):** 200GbE QSFP56 to QSFP56 Passive Twinax DAC (2.5 m)
- **9530522 (Citrix MPN):** 200GbE QSFP56 AOC (20 m)

Q8: Can I use a QSFP28 adapter at 25G or 10G with regular SFP optics on the MPX/SDX 17000 200G QSFP56 ports?

Yes, by using the Citrix adapter (Citrix MPN: 9530416). This is a QSFP28-to-SFP28 mechanical adapter based on the Mellanox/Nvidia MAM1Q00A-QSA28 design.

The adapter slots into a 200G QSFP56 port on the MPX/SDX 17000 platform and accepts any qualified Citrix-branded SFP28 (25G) or SFP+ (10G) optic. This is particularly useful when the appliance's native SFP56 ports are fully utilized, allowing customers to provision an additional 25G or 10G uplink from an available QSFP56 port.

Auto-Negotiation, Speed Forcing & FEC

Q1: Do ports auto-negotiate speed with these transceivers?

It depends on the link speed and the transceiver type:

SPEED	ENCODING	AUTO-NEGOTIATION BEHAVIOR
10G SFP+	NRZ	Yes
25G SFP28	NRZ	Yes (with FC-FEC)
50G SFP56 (PAM4)	PAM4	Yes (with RS-FEC)
40G QSFP+ BiDi	NRZ	No - forced speed may be required on the Nexus N9K/ N3K peer switch

SPEED	ENCODING	AUTO-NEGOTIATION BEHAVIOR
100G QSFP28 SR4	NRZ (4x25G)	Yes (with RS-FEC)
100G QSFP28 LR4	NRZ (4x25G CWDM)	Yes (No FEC - clean SMF link)
200G QSFP56 PAM4	PAM4 (4x50G)	Yes (with RS-FEC)

Q2: What FEC mode should I configure on the peer switch?

Match FEC to the link speed and media. FEC mismatch is the #1 cause of link-up-but-no-traffic or high-CRC-error symptoms.

LINK SPEED	MODULATION	REQUIRED FEC
10GbE	NRZ	None
25GbE NRZ (SFP28)	NRZ	FC-FEC
40GbE QSFP+	NRZ (4x10G)	None
50GbE SFP56	PAM4	RS-FEC
100GbE SR4	NRZ (4x25G)	RS-FEC
100GbE LR4	NRZ (4x25G CWDM)	None
200GbE SR4 / AOC	PAM4 (4x50G)	RS-FEC

Q3: Why does not my 40G link come up with auto-negotiation?

Whether a 40G link comes up on its own depends on the type of cable or transceiver used, not on the NIC card itself. The behavior differs by cable type:

- **40G copper cable (passive DAC):** The link comes up at 40G automatically. No configuration is needed on either end.
- **40G fiber transceiver (BiDi):** These run at a fixed speed and do not support auto-negotiation. If a port is left set to auto-negotiate, it does not find a matching partner and the link does not come up. To fix this, set the port to a fixed 40G speed on whichever device defaults to auto-negotiation, typically the Cisco peer switch (set "speed 40000" on the Nexus peer switch).

- **100G transceiver (QSFP28) used at 40G:** With auto-negotiation enabled, the link comes up at its native 100G speed, not 40G. To run it at 40G, manually set a fixed 40G speed on both ends.

Mixing & Compatibility

Q1: Can I use third-party optics or DAC/AOC cables?

Only Citrix-qualified, Citrix-branded MPNs (953xxxx series) are supported. Using non-qualified modules may cause link bring-up issues, FEC mismatch, or DDM read failures, and may not be covered under standard support terms.

Q2: Can I mix SR and LR (or different vendor optics) on each end of the link?

Both ends of a link must use compatible optics, which means same wavelength, same compliance code (for example, SR4 to SR4, LR4 to LR4), and matching connector or fiber type. Mixing SR with LR will not establish a link. For 40G BiDi optical, both ends must be the same Citrix MPN 9530411.

Q3: How do I choose SR vs LR?

- **SR (Short Reach):** Multimode fiber (OM3 70m / OM4 100m typical for SR4 at 100G+; OM3 300m / OM4 400m for 10G SR). Use within a single rack or row.
- **LR (Long Reach):** Single-mode fiber (OS2), typically up to 10 km. Use for inter-row, inter-DC, or campus links.

Always confirm the qualified reach in the Transceiver, DACs, AOCs' matrix](/en-us/netscaler-hardware-platforms/mpx/mpx-9200-17000-master-matrix). Actual reach depends on the optic and the fiber grade.

Q4: What MPO polarity do I need for SR4 modules (40G/100G/200G)?

All Citrix-qualified MPO-12 SR4 modules (Citrix MPN 9530409, 9530410) require Type B (key-up to key-down) MPO-12 trunk cables for correct fiber polarity. Type A or Type C trunks result in no-link or a one-way link.

Verification & Troubleshooting

Q1: How do I verify that the transceiver is detected and the link is up?

- Check the NetScaler host interface state: link, speed, FEC.
- Check the peer switch interface: link, speed, FEC, error counters.
- If the link is down: reseal the module, verify cable polarity (LC duplex or MPO Type B), confirm both sides match in speed and FEC.

Q2: Can I use a Cisco Nexus N9K with non-Cisco optics?

Yes, with the `service unsupported-transceiver` command enabled in the global configuration of the Nexus switch. Without this command, the Nexus N9K may refuse to bring up the port even if the optic is electrically compatible. Per-lane DDM may be reported as N/A on the N9K for non-Cisco QSFP modules. This is cosmetic, not a fault.

Q3: My 9530413 (Citrix MPN) 100G LR4 link is not coming up over OS2 fiber to a Mellanox SN3700 or Cisco N9K switch. The transceivers are seated and powered. What should I check first?

The 9530413 (Citrix MPN) samples should auto-negotiate to 100G LR4 over OS2 single-mode fiber without any manual speed forcing. If the link is not coming up, the most common root cause is a fiber-side issue (dirty or damaged LC connector, or a bad/broken patch cable), not the transceiver or NIC.

Diagnostic indicator: Check switch-side per-lane RX power. If all four lanes report approximately -40 dBm (the optical floor / "no light"), this is not a transceiver fault. It means that no optical signal is reaching the switch receiver. The link never reaches PHY training under these conditions, which is why auto-negotiation appears to fail.

Example symptom on a Mellanox SN3700:

```
Rx Power (per channel):
```

```
Ch1 Rx Power : 0.00010 mW / -40.00000 dBm  ← floor / no light
```

```
Ch2 Rx Power : 0.00010 mW / -40.00000 dBm  ← floor / no light
```

```
Ch3 Rx Power : 0.00010 mW / -40.00000 dBm  ← floor / no light
```

```
Ch4 Rx Power : 0.00010 mW / -40.00000 dBm  ← floor / no light
```

```
Ch1-Ch4 Rx Power Alarm : Rx Pwr Too Low
```

Meanwhile, TX power on both host and switch reads normal (~1.5–1.7 mW / ~2 dBm per channel) with no Tx alarms and the lasers are healthy. The optical path between them is broken.

Resolution steps:

1. Replace the OS2 LC-to-LC duplex single-mode patch cable with a known-good fiber.
2. Reseat both transceivers (host side and switch side) after fiber replacement.

3. Verify LC polarity and confirm that the duplex connectors are not crossed or mis-keyed.
4. Re-check switch-side RX power. After fiber replacement, RX power on all four channels should return to a normal range (typically +1 to +3 dBm), all "Rx Pwr Too Low" alarms should clear, and the link should come up at 100Gx4, auto-negotiation, without manual configuration.

Example after fiber replacement (Mellanox SN3700):

```
Eth1/31 Operational state: Up / Speed: 100Gx4 / Negotiation: Auto
```

```
Rx Power (per channel) → alarms now NONE
```

```
TX Power (per channel) → alarms NONE
```

Q4: What is the QSFP28-to-SFP28 adapter 9530416 (Citrix MPN) and when do I use it?

This QSFP28-to-SFP28 mechanical adapter (based on the Mellanox MAM1Q00A-QSA28) fits a 1-lane SFP-class optic into a 4-lane QSFP56 port on the MPX/SDX 17000. With this adapter installed, the 200G QSFP56 port can operate at lower link speeds, specifically 25G with an SFP28 optic (for example, 9530402 SR), or 10G with an SFP+ optic (for example, 9530403 SR), when the host NIC and peer switch are configured to match.

Q5: When would I use the QSFP28-to-SFP28 adapter 9530416 (Citrix MPN) instead of just using an SFP56 port directly?

Use the adapter when:

- All 8x SFP56 ports on the MPX/SDX 17000 are already in use and you need an additional 25G or 10G uplink.
- A deployment design specifically requires a 25G or 10G link from a QSFP56 port (for example, to match a particular switch panel layout).

For a fresh deployment, prefer the native SFP56 ports on the MPX/SDX 9200 and 17000 for 10G / 25G / 50G optics as they require no adapter and have the simplest interop path.

Q6: Is the QSFP28-to-SFP28 adapter 9530416 (Citrix MPN) supported on the MPX/SDX 9200?

No. The MPX/SDX 9200 has no QSFP56 ports. The 9530416 (Citrix MPN) is only applicable to the MPX/SDX 17000.

Q7: Does the QSFP28-to-SFP28 adapter 9530416 (Citrix MPN) add latency or signal loss?

No. The 9530416 (Citrix MPN) is a passive mechanical adapter. It physically routes a single lane between the SFP28/SFP+ optic and one lane of the QSFP56 port with no active components or retiming. Signal integrity and latency are the same as a native SFP socket.

Transceiver Interop Use Cases

The following use cases reflect real-world deployment topologies. These setups were validated on the NetScaler MPX and SDX 17000 platform using Mellanox ConnectX-7 200G NICs (FW 28.43.3608) with industry-standard data center peer switches. All QSFP-class use cases apply exclusively to the 17000 series. The MPX and SDX 9200 platform do not feature QSFP56 ports.

Citrix MPN: 9530409 (200G QSFP56 SR4 MPO-12)

Use Case: Native 200G to Mellanox SN3700

MPX/SDX 17000 (200G NIC) >> 9530409 <--- MPO-12 (Type B) ---> 9530409 << Mellanox SN3700 (QSFP56 port)

- **Negotiation:** Auto-negotiation at 200G PAM4 (4 x 50G), RS-FEC (544,514).

Citrix MPN: 9530410 (100/40G QSFP28 SR4 MPO-12)

Use Case 1: 100G Native Link to Cisco Nexus N9K

MPX/SDX 17000 (200G NIC) >> 9530410 <--- MPO-12 (Type B) ---> 9530410 << Cisco N9K-C93180YC-FX (QSFP28 port)

- **Customer scenario:** Data-center deployment uplinking the NetScaler appliance into a Cisco Nexus 100G port using Citrix-branded optics on both host and switch ends.
- **Negotiation:** Auto-negotiation establishes the link at 100G (4 x 25G NRZ) with RS-FEC (528,514). No manual configuration required.

Use Case 2: 40G Forced Link to Cisco Nexus N9K (Mixed-Speed Interop)

MPX/SDX 17000 (200G NIC) >> 9530410 <--- MPO-12 (Type B) ---> Cisco QSFP-40G-SR4 or Intel FTL410QE2C-IT << Cisco N9K-C93180YC-FX (QSFP+ port)

- **Customer scenario:** Migration or mixed-generation environments where the appliance must connect to an existing 40G Nexus port using a legacy QSFP-40G-SR4 transceiver on the switch side.
- **Negotiation:** Auto-negotiation is **not possible**. Speed must be set manually to 40G on both the host and the switch.
- **IEEE 802.3 architectural constraint:** The host QSFP28 SerDes initializes at 4 x 25G. The switch QSFP-40G-SR4 runs 4 x 10G. These are incompatible lane rates that require a common forced rate first.
- On the Nexus 9000, enable `service unsupported-transceiver` to allow the non-Cisco-branded QSFP-40G-SR4 on the switch side.
- Per-lane DDM may show **N/A** on the N9K side at 40G. This is expected and has no impact to link health.

Use Case 3: 100G Native Link to Mellanox SN3700 (Spectrum-2)

MPX/SDX 17000 (200G NIC) >> 9530410 <--- MPO-12 (Type B) ---> 9530410 << Mellanox SN3700 (QSFP28 port)

- **Customer scenario:** Data-center deployment uplinking the NetScaler appliance into a Mellanox/NVIDIA SN3700 100G switch.
- **Negotiation:** Auto-negotiation establishes 100G with RS-FEC. No manual configuration required.

Key takeaway for 9530410: Native 100G works automatically on both N9K and SN3700. The 40G path requires forced speed on both ends and applies only when the peer switch port carries a QSFP-40G-SR4 module.

MPN: 9530411 (40G QSFP+ BiDi SR LC)

Use Case 1: BiDi Link with 9530411 on Both Ends (Stable)

MPX/SDX 17000 (200G NIC) >> 9530411 <--- LC duplex MMF ---> 9530411 << Cisco N9K or Mellanox SN3700 (QSFP+ port)

- **Result:** Link is stable, **no flap observed** when the same Citrix MPN 9530411 is installed at both host and switch ends. Forced 40G speed required on peer switch sides (no auto-negotiation at 40G).

Use Case 2: Mixed BiDi Pairing (Link Flap Observed)

9530411 on host, Cisco-branded 40G-BiDi (for example, AFBR-79EBPZ-CS2) on switch side.

- **Symptom:** Intermittent link flaps when the switch side uses a Cisco-branded 40G-BiDi module while the host uses 9530411.
- **Status: Not recommended.** For supported BiDi deployments, use the Citrix MPN 9530411 on **both** ends.

Key takeaway for 9530411: Use the same Citrix-branded BiDi MPN on host **and** switch ends. Cross-vendor BiDi pairings (Citrix host + Cisco-branded switch) are known to flap and are not supported.

Citrix MPN: 9530413 (100G QSFP28 LR4 LC, 10 km)

Use Case 1: 100G LR4 to Cisco Nexus N9K

MPX/SDX 17000 (200G NIC) >> **9530413** <--- OS2 LC-LC duplex SMF ---> **9530413** << **Cisco N9K-C93180YC-FX** (QSFP28 port)

- **Negotiation:** Auto-negotiation at 100G (4 x 25G NRZ), no FEC (LR4 SMF link has sufficient SNR margin to operate without FEC).

Use Case 2: 100G LR4 to Mellanox SN3700

MPX/SDX 17000 (200G NIC) >> **9530413** <--- OS2 LC-LC duplex SMF ---> **9530413** << **Mellanox SN3700** (QSFP28 port)

- **Negotiation:** Auto-negotiation at 100G (4 x 25G NRZ), no FEC.

Key takeaway for 9530413: Native 100G LR4 links auto-negotiate cleanly with no FEC between host and both validated interop switches over OS2 SMF.

Citrix MPN: 9530416 (QSFP28-to-SFP28 Adapter)

Use Case 1: 25G SFP28 SR Uplink via QSFP56 Port

MPX/SDX 17000 (200G NIC) >> **9530416** adapter >> **9530402** (25G SFP28 SR) <--- LC duplex MMF ---> **Cisco SFP-25G-SR-S** << Cisco Nexus (25G port)

- **Customer scenario:** The MPX/SDX 17000's eight native SFP56 ports are already fully utilized. The customer needs an additional 25G uplink and uses a QSFP56 port through the 9530416

adapter + a 9530402 SFP28 SR optic on the host side, matched against a Cisco SFP-25G-SR-S on a Cisco Nexus 25G switch port.

- **Negotiation:** Standard 25G auto-negotiation behavior. Link comes up at 25G with FC-FEC.

Use Case 2: 10G SFP+ SR Uplink via QSFP56 Port

MPX/SDX 17000 (200G NIC) >> **9530416** adapter >> **9530403** (10G SFP+ SR) <--- LC duplex MMF ---> **Cisco SFP-10G-SR** << Cisco Nexus (10G port)

- **Customer scenario:** The customer needs an additional 10G uplink from a 17000 QSFP56 port and uses a 9530416 adapter + 9530403 SFP+ SR optic against a Cisco SFP-10G-SR on a Nexus 10G port.
- **Negotiation:** Standard 10G operation. The link comes up at 10G with no FEC.

Key takeaway for 9530416: The adapter enables a 200G QSFP56 port to operate as a single-lane 25G or 10G port using a standard SFP28 / SFP+ optic. Configure the host and switch for the matching lower speed. For 10G or 25G deployments, prefer the native SFP56 cages on the MPX/SDX 17000 (or any port on the MPX/SDX 9200) as they require no adapter.

Citrix MPN: 9530513 (40G QSFP+ DAC, 5 m)

Use Case 1: 40G QSFP+ DAC to Cisco Nexus N9K

MPX/SDX 17000 (**200G NIC**, host end labeled **CITRIX/9530513**) >> **9530513** <--- 5 m passive copper ---> **9530513** << (switch end labeled **CISCO-CITRIX/C-QCIQMX-PDAC5M**) << **Cisco N9K-C93180YC-FX** (QSFP+/QSFP28 port)

- **Customer scenario:** Short-reach 40G uplink (within a rack or adjacent rack) from the NetScaler appliance to a Cisco Nexus 40G port using a passive copper DAC.
- **Cable end labeling (by design):** The host end is labeled "CITRIX / 9530513". The switch end is labeled "CISCO-CITRIX / C-QCIQMX-PDAC5M". The asymmetric labeling is intentional so each end is recognized natively by the peer it faces.
- **Negotiation:** Auto-negotiation at 40G (4 x 10G NRZ), no FEC.

Use Case 2: 40G QSFP+ DAC to Mellanox SN3700

MPX/SDX 17000 (**200G NIC**, host end labeled **CITRIX/9530513**) >> **9530513** <--- 5 m passive copper ---> **9530513** << (switch end labeled **CISCO-CITRIX/C-QCIQMX-PDAC5M**) << **Mellanox SN3700** (QSFP+/QSFP28 port)

- **Customer scenario:** Short-reach 40G uplink (within a rack or adjacent rack) from the NetScaler appliance to a Mellanox SN3700 switch port using a passive copper DAC.
- **Negotiation:** Auto-negotiation at 40G (4 x 10G NRZ), no FEC.

Key takeaway for 9530513: The 40G QSFP+ DAC is a passive copper alternative to the 9530411 BiDi optic for short (5 m or less) 40G links. Unlike the 9530411 BiDi case, mixed-end labeling (CITRIX on host / CISCO-CITRIX on switch) is the expected design for the 9530513, not a constraint.

Citrix MPN: 9530517 (100G QSFP28 AOC, 20 m)

Use Case 1: 100G AOC to Cisco Nexus N9K

MPX/SDX 17000 (200G NIC, CITRIX end) >> 9530517 <--- 20 m AOC ---> 9530517 (CISCO-CITRIX end / Vendor PN Q100GBAOC20MCMXC) << Cisco N9K (QSFP28 port)

- **EEPROM programming asymmetry:** Host end is labeled "CITRIX / 9530517" (Vendor OUI: 00:00:00); switch end is labeled "CISCO-CITRIX / Q100GBAOC20MCMXC" (Vendor OUI: 00:00:00). Vendor Rev: End A = **A1**, End B = **D1**. The asymmetric labeling is intentional so each end is recognized natively by the peer it faces. No anomalies observed against intended host/switch role mapping.

Use Case 2: 100G AOC to Mellanox SN3700

MPX/SDX 17000 (200G NIC, CITRIX end) >> 9530517 <--- 20 m AOC ---> 9530517 (CISCO-CITRIX end / Vendor PN Q100GBAOC20MCMXC) << Mellanox SN3700 (QSFP28 port)

Key constraint for 9530517: This module also advertises 40G capability in EEPROM, but does not link in a QSFP+ peer-switch port (the AOC end-electronics are QSFP28-only). For 40G interop, use 9530411 (BiDi) or 9530513 (40G QSFP+ DAC, 5 m) instead.

Managing RAID disk allocation on the SDX platform

NetScaler SDX 14000/15000/16000/25000/26000 series of platforms now include a Redundant Array of Independent Disks (RAID) controller, which can support up to eight physical disks. Multiple disks provide not only performance gains, but also enhanced reliability. Reliability is especially important for an SDX appliance, because the appliance hosts many virtual machines, and a disk failure affects multiple virtual machines. The RAID controller on the Management Service supports the RAID 1 configuration, which implements disk mirroring. That is, two disks maintain the same data. If a disk in the RAID 1 array fails, its mirror immediately supplies all needed data.

RAID 1 disk mirroring combines two physical drives in one logical drive. The usable capacity of a logical drive is equivalent to the capacity of one of its physical drives. Combining two 1-terabyte drives, for example, creates a single logical drive with a total usable capacity of 1 terabyte. This combination of drives appears to the appliance as a single logical drive.

The SDX appliance is shipped with a configuration that includes logical drive 0 and logical drive 1. Logical drive 0 is allocated for the Management Service and the Citrix Hypervisor and logical drive 1 is allocated for the NetScaler instances that you provision. To use more physical drives, you have to create new logical drives.

View drive properties and operations

An SDX appliance supports a maximum of eight physical-drive slots, that is, a pair of four slots on each side of the appliance. You can insert physical drives into the slots. Before you can use a physical drive, you must make it part of a logical drive.

In the Management Service, the **Configuration > System > RAID** screen includes tabs for logical drives, physical drives, and storage repositories.

Storage repository

On the **Configuration > System > RAID > Storage Repository** tab, you can view the status of storage repositories on the SDX appliance. You can also view information about a storage-repository drive that is not attached, and you can remove such a drive by selecting it and then

clicking **Remove**. The **Storage Repository** tab displays the following information about each storage repository:

- **Name:** Name of the storage repository drive.
- **Is Drive Attached:** Whether the storage repository is attached or not. If the drive is not attached, you can click **Remove** to delete.
- **Size:** Size of the storage repository.
- **Utilized:** Amount of storage-repository space in use.

Logical drives

On the **Configuration > System > RAID > Logical Drives** tab, you can view the name, state, size, of each logical drive, and information about its component physical drives. The following table describes the states of the logical drive.

STATE	DESCRIPTION
Optimal	The logical drive operating condition is good. All configured drives are online.
Degraded	The logical drive operating condition is not optimal. One of the configured drives has failed or is offline.
Failed	The logical drive has failed.
Offline	The logical drive is not available to the RAID controller.
Active	The logical drive is fully active, disk I/O and/or resync can be happening.
Clean	The logical drive is active, but does not have any pending write operations.
Recovering	All data on the logical drive is being written from the running array to fill up a new disk with all the relevant data.
Resync	All data on the array is being synchronized from the running array. However, most, if not all, of the data is OK.

You can also view the details the physical drives associated with the logical drive by selecting the logical drive and clicking **Show Physical Drive**.

Create a new logical drive

1. Navigate to **Configuration > System > RAID**, and select the **Logical Drives** tab.
2. Click **Add**.
3. In the **Create Logical Disk** dialog box, select two slots that contain operational physical drives, and then click **Create**.

Add a logical drive to the SDX Platform

1. On the back of the appliance, insert the two blank SSDs in slot numbers 5 and 6. You can add the SSDs in a running system.

NOTE:

Make sure that the SSDs are NetScaler® certified.

2. In the Management Service, navigate to **Configuration > System > RAID** and the **Physical Drives** tab. You can see the SSDs that you added.
3. Navigate to the **Logical Drive** tab and click **Add**.
4. In the **Create Logical Disk** page:
 - a. In the **First Slot** drop-down list, select 5.
 - b. In the **Second Slot** drop-down list, select 6.
 - c. Click **Create**.

Note: In Management Service, the slot number begins with zero. So the slot numbering in Management Service differs from the slot numbering on the physical appliance.

 - d. The logical drive is created and is listed under the **Logical Drive tab**. Click the refresh icon to update the order of the logical drives.

Add a second logical drive on the SDX Platform

1. To add another logical drive, insert the SSDs in slot numbers 7 and 8.
2. In the **Create Logical Disk** page:
 - a. In the **First Slot** drop-down list, select 7.

- b. In the **Second Slot** drop-down list, select 8.
 - c. Click **Create**
3. The logical drive is created and is listed under the **Logical Drive tab**. Click the refresh icon to update the order of the logical drives.

Physical drives

An SDX appliance supports a maximum of eight physical slots, that is, a pair of four slots on each side of the appliance. On the **Configuration > System > RAID > Physical Drives** tab, you can view the following information:

- **Slot:** Physical slot associated with the physical drive.
- **State:** Indicates if a physical drive is present on the appliance.
- **Present** Physical drive is present on the appliance.
- **Size:** Size of the physical drive.
- **Not Present** It means that either the physical drive is absent on the appliance or it has been deactivated so that it can be removed.
- **Firmware State:** State of the firmware. Possible Values:
 - **Online, spun up:** Physical drive is up and is controlled by RAID.
 - **Unconfigured (good):** Physical drive is in good condition and can be added as a part of the logical drive pair.
 - **Unconfigured (bad):** Physical drive is not in good condition and cannot be added as part of a logical drive.
- **Foreign State:** Indicates if the disk is empty.
- **Logical Drive:** Associated logical drive.
- **RAID State** It defines the raid state of the physical drives.
 - **active, sync:** Physical drive in good condition in sync with pending writes.
 - **clean, sync:** Physical drive in good condition in sync with no pending writes.
 - **spare, rebuilding:** A new physical drive is being filled up with relevant data.
 - **faulty:** Physical drive in a faulty state.

In the **Physical Drives** pane, you can perform the following actions on the physical drives:

- **Initialize:** Initialize the disk. You can initialize the physical drive if it is not in good state and must be added as a part of the logical drive pair.
- **Rebuild:** Initiate a rebuild of the drive. When a drive in a drive group fails, you can rebuild the drive by re-creating the data that was stored on the drive before it failed. The RAID controller re-creates the data stored on the other drives in the drive group.
 - **Locate:** Locate the drive on the appliance, indicated by causing the Drive Activity LED associated with the drive to blink.
 - **Stop Locate:** Stop locating the drive on the appliance.
 - **Prepare to Remove:** Deactivate the selected physical drive so that it can be removed.

Raid support

Replace a defective RAID-supported SSD drive with a blank SSD drive

1. Navigate to **Configuration > System > RAID**.
2. On the **Physical Drives** tab, select the defective drive that you want to replace.
3. Click **Prepare to Remove** to remove the drive.
4. Click **Physical Drives**. The state of the drive must now appear as **Not Present**.
5. Locate the SSD on the back panel of the appliance. Push the safety latch of the drive cover down while pulling out on the drive handle to disengage. Pull out the faulty drive.
6. Ensure that the new SSD is the same size as the old one. For example, if the old drive is 480 GB, the new one must be 480 GB.
7. Open the drive handle fully up, and insert the new Citrix-certified SSD into the slot as far as possible. To seat the drive, close the handle flush with the rear of the appliance so that the drive locks securely into the slot.

When inserting the drive, orient the drive handle the same way as the installed drives.

8. Click the refresh icon for the list of physical drives. The state of the drive must appear as **Present**.
9. Select the drive first and then click the **Rebuild** option. Click **OK** when the **Rebuild Physical Drive** dialog box is displayed. After you replace one of the SSDs, the configuration on the other SSD in the mirrored SSD is copied to the replacement SSD.

NOTE:

A Force Clean check option is available in **Rebuild Physical Drive** dialog box to erase any existing RAID metadata on the drive. The drive is considered as a foreign drive if it came from another RAID setup or has old RAID metadata that must be erased.

10. Click the refresh icon to check the status of the rebuild process. When the rebuild process is complete, you can see **Online, Spun Up** status in the **Firmware State** column.

Replace a defective RAID logical Drive with two blank SSD Drives

If VPX-SR failure happens, where both the drives that form a RAID1 pair have failed, replace both the SSD drives.

To replace the defective RAID logical Drive do the following:

1. In the Management Service **NetScaler** tab, select and delete each of the VPX instances residing on the failed VPX-SR.
2. Navigate to **Configuration > System > RAID**, and select the **Logical Drives** tab.
3. Select the failed logical drives and click **Delete**.
4. Replace the failed drives with new Citrix-certified SSD drives in the corresponding slot pair.
5. Select the logical drives, and click **Create**.
6. The logical drive is re-created and ready for use. The VPX instances that were deleted earlier can now be created or restored as required.

Safety, cautions, warnings, and other information

Read the safety and other compliance information before unpacking and installing the product. The safety and other compliance information remains the same for both SDX and MPX appliances. See the following documents:

- [Safety, cautions, warnings, and other information.](#)
- [Taiwan BSMI RoHS statement](#)

Note: For the list of safety certifications, standards, and ROHS compliance for each model, see the data sheet. The data sheet is available on www.netscaler.com. Navigate to **Platform > Form Factors > Hardware**.

- [FCC compliance statement](#)

Prepare for installation

NOTE:

This section does not apply to the SDX 9200 platform. For information specific to the SDX 9200, see [NetScaler SDX 9200 platform](#).

Before you install your new appliance, carefully unpack your appliance and make sure that all parts were delivered. Verify that the location where the appliance is to be installed meets temperature and power requirements. Ensure that the server cabinet or floor-to-ceiling cabinet is securely bolted to the floor and has sufficient airflow.

Only trained and qualified personnel must install, maintain, or replace the appliance, and efforts must be taken to ensure that all cautions and warnings are followed.

This document includes the following details:

- Unpacking the appliance
- Preparing the site and rack
- Electrical safety precautions

Unpack the appliance

The hardware accessories for your particular appliance, such as cables, adapters, and rail kit, vary depending on the hardware platform you ordered. Unpack the box that contains your new appliance on a sturdy table with plenty of space and inspect the contents.

Verify that you received the cables, adapter, and rail kits specified for your appliance.

NOTE:

Make sure that a power outlet is available for each cable. If the kit that you received does not fit your rack, contact your NetScaler® sales representative to order the appropriate kit. Transceiver modules are sold separately. Contact your NetScaler sales representative to order transceiver modules for your appliance. Only transceivers supplied by Citrix® are supported on the appliance.

For Brazilian customers, NetScaler does not ship a power cable. Use a cable that conforms to the **ABNT NBR 14136:2002** standard.

In addition to the items included in the box with your new appliance, you need the following items to complete the installation and initial configuration process.

- Ethernet cables for each additional Ethernet port that you connect to your network
- One available Ethernet port on your network switch or hub for each NetScaler Ethernet port you want to connect to your network Note: Transceiver modules are sold separately. Contact your NetScaler sales representative to order transceiver modules for your appliance. Only transceivers supplied by Citrix are supported on the appliance.
- A computer to serve as a management workstation

Prepare the site and rack

There are specific site and rack requirements for the NetScaler appliance. Ensure that adequate environmental control and power density are available. Racks must be bolted to the ground, have sufficient airflow, and have adequate power and network connections. Preparing the site and rack are important steps in the installation process and help ensure a smooth installation.

Site requirements

The appliance must be installed in a server room or server cabinet with the following features:

- **Environment control**

An air conditioner, preferably a dedicated computer room air conditioner (CRAC), capable of maintaining the cabinet or server room at a temperature of no more than 27 degrees C/80.6 degrees F at altitudes of up to 2100 m/7000 ft, or 18 degrees C/64.4 degrees F at higher altitudes, a humidity level no greater than 45 percent, and a dust-free environment .

- **Power density**

Wiring capable of handling at least 4000 watts per rack unit in addition to power needs for the CRAC.

Rack requirements

The rack on which you install your appliance must meet the following criteria:

- Rack characteristics

Racks must be either integrated into a purpose-designed server cabinet or be the floor-to-ceiling type, bolted down at both top and bottom to ensure stability. If you have a cabinet, it must be installed perpendicular to a load-bearing wall for stability and sufficient airflow. If you have a server room, your racks must be installed in rows spaced at least 1 meter/3 ft apart for sufficient airflow. Your rack must allow your IT personnel unfettered access to the front and back of each server and to all power and network connections.

- Power connections

At minimum, two standard power outlets per unit.

- Network connections

At minimum, four Ethernet connections per rack unit.

- Space requirements

One empty rack unit for the NetScaler SDX 8015/8400/8600, and two consecutive empty rack units for all other appliance models.

You can order the following rail kits separately.

- Compact 4-post rail kit, which fits racks of 23–33 inches.
- 2-post rail kit, which fits 2-post racks.

Electrical safety precautions

Read the caution and danger information you need to know, before unpacking and installing the product. For more information, see [Safety, cautions, warnings, and other information](#).

Install the hardware

Determining that the location where you plan to install your appliance meets the environmental standards and the server rack is in place according to the instructions. After you mount the appliance, you are ready to connect it to the network, to a power source, and to the console terminal that you will use for initial configuration. To complete the installation, you turn on the appliance. Be sure to observe the cautions and warnings listed with the installation instructions.

Note: Keep the serial number handy before you rack-mount the appliance. The serial number is the password for the first time logon to the appliance and can be found at the back of the appliance.

Rack-mount the appliance

Most appliances can be installed in standard server racks that conform to the EIA-310-D specification. The appliances ship with a set of rails, which you must install before you mount the appliance. The only tools that you need for installing an appliance are a Phillips screwdriver and a flathead screwdriver.

CAUTION:

If you are installing the appliance as the only unit in the rack, mount it at the bottom. If the rack contains other units, make sure that the heaviest unit is at the bottom. If the rack has stabilizing devices available, install them before mounting the appliance.

To check the different hardware platforms and the rack units required for each platform, see the details given for each model under [SDX hardware platforms](#).

Each appliance ships with a mounting rail kit that contains two rail assemblies, one each for the left and the right side of the appliance, and screws to attach the rails. An assembly consists of an inner rail and a rack rail. The supplied rail kit is 28 inches long (38 inches extended). Contact your NetScaler® sales representative to order a 23 inches (33 inches extended) rail kit.

Note: The same rail kit is used for both square-hole and round-hole racks. See [To install the rack rails on the rack](#) for specific instructions for threaded, round-hole racks.

To mount the appliance, you must first install the rails and then install the appliance in the rack.

Perform the following tasks to mount the appliance:

- Remove the inner rails from the rail assembly.
- Attach the inner rails to the appliance.
- Install the rack rails on the rack.
- Install the appliance in the rack.


The appliance is shipped with rack-rail hardware. This hardware consists of two inner rails that you attach to the appliance, one on each side, and a rack-rail assembly that you attach to the rack. The following figure illustrates the steps involved in mounting the NetScaler SDX appliance to a rack.

Remove the inner rails from the rail assembly

1. Place the rail assembly on a flat surface.
2. Slide out the inner rail toward the front of the assembly.
3. Press the latch until the inner rail comes all the way out of the rail assembly.
4. Repeat steps 1 through 3 to remove the second inner rail.

Attach the inner rails to the appliance

1. Position the right inner rail behind the handle on the right side of the appliance.
2. Align the holes on the rail with the corresponding holes on the side of the appliance.
3. Attach the rail to the appliance with the provided screws: 4 per side for a 1U appliance and 5 per side for a 2U appliance, as shown in the following figure. Figure 1. Attaching inner rails



4. Repeat steps 1 through 3 to install the left inner rail on the other side of the appliance.

Install the rack rails on the rack

1. If you have a round-hole, threaded rack, skip to step 3.
2. Install square nut retainers into the front post and back post of the rack as shown in the following figures. Before inserting a screw, be sure to align the square nut with the correct hole

for your 1U or 2U appliance. The three holes are not evenly spaced. Figure 2. Installing Retainers into the Front Rack Posts






Figure 3. Installing Retainers into the Rear Rack Posts




3. Install the adjustable rail assembly into the rack as shown in the following figures. Use a screw to lock the rear rail flange into the rack. With the screw securing the rail in place, you can optionally remove the latching spring. Figure 4. Installing the Rail Assembly to the Rack



Install the appliance in the rack

1. Align the inner rails, attached to the appliance, with the rack rails.
2. Slide the appliance into the rack rails, keeping the pressure even on both sides.
3. Verify that the appliance is locked in place by pulling it all the way out from the rack. Figure 5.

Rack Mounting the Appliance



Watch this quick video about [how to rack mount a NetScaler hardware appliance](#).

Install and remove 1G SFP transceivers

A Small Form-Factor Pluggable (SFP) is a compact transceiver that can operate at speeds of up to 1 gigabit per second and is available in both copper and fiber types. Inserting a 1G SFP copper transceiver converts the 1G SFP port to a 1000BASE-T port. Inserting a 1G SFP fiber transceiver converts the 1G SFP port to a 1000BASE-X port. Auto-negotiation is enabled by default on the 1G SFP port into which you insert your 1G SFP transceiver. When a link between the port and the network is established, the speed and mode are matched on both ends of the cable.

CAUTION:

NetScaler appliances do not support 1G SFP transceivers from vendors other than Citrix Systems. Attempting to install third-party 1G SFP transceivers on your NetScaler appliance voids the warranty.

Insert 1G SFP transceivers into the 1G SFP ports on the front panel of the appliance. Frequent installation and removal of transceivers shortens their life span. Follow the removal procedure carefully to avoid damaging the 1G SFP transceiver or the appliance.

CAUTION:

Do not install the transceivers with the cables attached. Doing so can damage the cable, the connector, or the optical interface of the transceiver.


Install a 1G SFP transceiver

1. Remove the 1G SFP transceiver carefully from its box. Danger: Do not look directly into fiber optic transceivers or cables. They emit laser beams that can damage your eyes.
2. Align the 1G SFP transceiver to the front of the 1G SFP transceiver port on the front panel of the appliance, as shown in the following figure.

NOTE

The illustration in the following figures might not represent your actual appliance.

Figure 6. Installing a 1G SFP transceiver



3. Hold the 1G SFP transceiver between your thumb and index finger and insert it into the 1G SFP transceiver port. Press it in until you hear the transceiver snap into place.
4. Lock the transceiver.
5. Verify that the LED is green and blinks twice, which indicates that the transceiver is functioning correctly.
6. When using a fiber 1G SFP transceiver, do not remove the dust caps attached to the transceiver and the cable until you are ready to insert the cable.

Remove a 1G SFP transceiver

1. Disconnect the cable from the 1G SFP transceiver. If you are using a fiber optic cable, replace the dust cap on the cable before putting it away. Danger: Do not look directly into fiber optic transceivers or cables. They emit laser beams that can damage your eyes.
2. Unlock the 1G SFP transceiver.
3. Hold the 1G SFP transceiver between your thumb and index finger and slowly pull it out of the port.
4. If you are removing a fiber 1G SFP transceiver, replace the dust cap before putting it away.
5. Put the 1G SFP transceiver into its original box or another appropriate container.

Install and remove 10G SFP+ transceivers

A 10-Gigabit Small Form-Factor Pluggable (SFP+) is a compact optical transceiver that can operate at speeds of up to 10 gigabits per second. Autonegotiation is enabled by default on the 10G SFP+ ports into which you insert your 10G SFP+ transceiver. When a link between the port and the network is established, the mode is matched on both ends of the cable and for 10G SFP+ transceivers, the speed is also autonegotiated.

CAUTION:

NetScaler appliances do not support 10G SFP+ transceivers provided by vendors other than Citrix Systems. Attempting to install third-party 10G SFP+ transceivers on your NetScaler appliance voids the warranty.

Insert the 10G SFP+ transceivers into the 10G SFP+ ports on the front panel of the appliance. Frequent installation and removal of transceivers shortens their life span. Follow the removal procedure carefully to avoid damaging the transceiver or the appliance.

CAUTION:

Do not install the transceivers with the cables attached. Doing so can damage the cable, the connector, or the optical interface of the transceiver.

Install a 10G SFP+ transceiver

1. Remove the 10G SFP+ transceiver carefully from its box.

WARNING:

Do not look directly into fiber optic transceivers and cables. They emit laser beams that can damage your eyes.

2. Align the 10G SFP+ transceiver to the front of the 10G SFP+ transceiver port on the front panel of the appliance.
3. Hold the 10G SFP+ transceiver between your thumb and index finger and insert it into the 10G SFP+ transceiver port, pressing it in until you hear the transceiver snap into place.
4. Lock the transceiver.
5. Verify that the LED is green and blinks twice, which indicates that the transceiver is functioning correctly.
6. Do not remove the dust caps attached to the transceiver and cable until you are ready to insert the cable.

Remove a 10G SFP+ transceiver

1. Disconnect the cable from the 10G SFP+ transceiver. Replace the dust cap on the cable before putting it away. Danger: Do not look directly into fiber optic transceivers or cables. They emit laser beams that can damage your eyes.
2. Unlock the 10G SFP+ transceiver.
3. Hold the 10G SFP+ transceiver between your thumb and index finger and slowly pull it out of the port.
4. Replace the dust cap on the transceiver before putting it away.
5. Put the 10G SFP+ transceiver into its original box or another appropriate container.

Connect the cables

When the appliance is securely mounted on the rack, you are ready to connect the cables.

Ethernet cables and the optional console cable are connected first. Connect the power cable last.

Warning: Before installing or repairing the appliance, remove all jewelry and other metal objects that might come in contact with power sources or wires. When you touch both a live power source or wire and ground, any metal objects can heat up rapidly and cause burns, set clothing on fire, or fuse the metal object to an exposed terminal.


Connect the Ethernet cables

Ethernet cables connect your appliance to the network. The type of cable you need depends on the type of port used to connect to the network. Use a category 5e or category 6 Ethernet cable with a standard RJ-45 connector on a 10/100/1000BASE-T port or 1G SFP copper transceiver. Use a fiber optic cable with an LC duplex connector with a 1G SFP fiber transceiver, 10G SFP+

transceiver. The type of connector at the other end of the fiber optic cable depends on the port of the device that you are connecting to.

Connect an Ethernet cable to a 10/100/1000BASE-T port or 1G SFP copper transceiver

1. Insert the RJ-45 connector on one end of your Ethernet cable into an appropriate port on the front panel of the appliance, as shown in the following figure. Figure 7. Inserting an Ethernet cable



2. Insert the RJ-45 connector on the other end into the target device, such as a router or switch.
3. Verify that the LED glows amber when the connection is established.

Connect the Ethernet cable to a 1G SFP fiber, 10G SFP+ transceiver


1. Remove the dust caps from the transceiver and cable.
2. Insert the LC connector on one end of the fiber optic cable into the appropriate port on the front panel of the appliance.
3. Insert the connector on the other end into the target device, such as a router or switch.
4. Verify that the LED glows amber when the connection is established.

Connect the Console cable

You can use the console cable to connect your appliance to a computer or terminal, from which you can configure the appliance. Alternatively, you can use a computer connected to the network. Before connecting the console cable, configure the computer or terminal to support VT100 terminal emulation, 9600 baud, 8 data bits, 1 stop bit, parity, and flow control set to NONE. Then connect one end of the console cable to the RS232 serial port on the appliance and the other end to the computer or terminal.

Connect the console cable to a computer or terminal

1. Insert the DB-9 connector at the end of the cable into the console port on the front panel of the appliance, as shown in the following figure. Figure 8. Inserting a console cable

**NOTE**

To use a cable with an RJ-45 converter, insert the optional converter provided into the console port and attach the cable to it.


2. Insert the RJ-45 connector at the other end of the cable into the serial port of the computer or terminal.

Connect the power cable

An SDX 8015/8400/8600 appliance has one power cable. All the other appliances come with two power cables, but they can also operate if only one power cable is connected. A separate ground cable is not required, because the three-prong plug provides grounding.

Connect the appliance to the power source

1. Connect one end of the power cable to the power outlet on the back panel of the appliance, next to the power supply, as shown in the following figure. Figure 9. Inserting a power cable



2. Connect the other end of the power cable to a standard 110V/220V power outlet.
3. If a second power supply is provided, repeat steps 1 and 2 to connect the second power supply.

NOTE

Some appliances emit a high-pitched alert if one power supply fails or if you connect only one power cable to the appliance. To silence the alarm, you can press the small red button on the back panel of the appliance.

Switch on the appliance

After you have installed the appliance in a rack and connected the cables, verify that the power cable is properly connected. If you have installed a second power supply, make sure that the second cable is connected to an outlet for a different circuit than the first. After verifying the connections, you are ready to switch on the appliance.

To switch on the appliance

1. Verify that the appliance is connected through a console or Ethernet port. This check ensures that you can configure the appliance after it is switched on.
2. Press the ON/OFF toggle power switch on the back panel of the appliance.

CAUTION:

Be aware of the location of the emergency power off (EPO) switch, so that if an electrical accident occurs you can quickly remove power from the appliance.

Initial configuration

NOTE:

This section does not apply to the SDX 9200 platform. For information specific to the SDX 9200, see [NetScaler SDX 9200 platform](#).

After you have installed your appliance in a rack, you are ready to perform the initial configuration. To perform the initial configuration, you can use the Management Service user interface or the serial console. You can access the Management Service user interface from any computer that is on the same network as the new SDX appliance. If you do not have a computer on the same network, use the serial console to perform the initial configuration of the SDX appliance. Citrix recommends that, when you complete the initial configuration, you change the root-user password. For information about changing the root-user password, see [Configuring Authentication and Authorization Settings](#).

Determine the following information for performing the initial configuration.

- NetScaler SDX IP address and subnet mask: The management IP address and the mask used to define the subnet in which the SDX appliance is located. This IP address is used to access the NetScaler SDX Management Service user interface.
- Citrix Hypervisor™ IP address: The IP address of the Citrix Hypervisor.
- Default gateway: The IP address of the router that forwards traffic out of the appliance's subnet. The default gateway must be in the same subnet as the NSIP address.
- Root password: The root user has full administrative privileges on the appliance. The root password is used to authenticate the root user. Change this password during the initial configuration of the appliance.

This topic includes the following sections:

- Initial Configuration through the Management Service User Interface
- Initial Configuration through the Serial Console
- Changing the Password of the Default User Account

Initial configuration through the Management Service user interface

To set up the appliance by using the Management Service user interface, connect a workstation or laptop to the same network as the appliance.

To configure the NetScaler SDX appliance by using the Management Service user interface

1. Connect the NetScaler SDX appliance to a management workstation or network by using interface 0/1.
2. Open a browser and type: `<http://192.168.100.1>`

Note: The NetScaler SDX Management Service is preconfigured with the IP address 192.168.100.1 and the Citrix Hypervisor is preconfigured with the IP address 192.168.100.2.

1. In **User Name**, type `nsroot`. In **Password**, if the earlier default password does not work, try typing the serial number of the appliance. The serial number bar code is available at the back of the appliance. Citrix recommends that you change the password after the first logon. For information about changing the password, see [Change the administrative password](#).
2. In the navigation pane, click **System**.
3. In the details pane, under **Setup Appliance**, click **Network Configuration**.
4. In the **Network Configuration** page, do the following:

a. In the **Interface** field, select the management interface that connects the appliance to a management workstation or network. Possible values: 0/1, 0/2. Default: 0/1.

Note: If you create an LA channel of the management ports, you can assign only one IP address to the channel. In the **Interface** field, only the LA channel option appears instead of 0/1 and 0/2. Also, even if both ports are connected, the Management Service is active only on one port.

b. In the **Appliance supportability IP** field, enter the IP address of the Citrix Hypervisor.

c. In the **Gateway** field, enter the IP address of the router that forwards traffic out of the appliance's subnet.

d. In the **DNS** field, enter the IPv4 address of the primary DNS server.

NOTE:

IPv6 addresses are not support for the primary DNS server.

e. Select the IPv4 checkbox if you want to use the IPv4 address for the Management Service and enter the details for the following parameters:

i. **Appliance Management IP** — The IPv4 address that is used to access the Management Service by using a Web browser.

ii. **Netmask** —The mask used to define the subnet in which the SDX appliance is located.

f. Select the **IPv6** checkbox if you want to use the IPv6 address for the Management Service and enter the details for the following parameters:

i. **Management Service IP Address** —The IPv6 address that is used to access the Management Service by using a Web browser.

NOTE

The Citrix Hypervisor IP address and Management Service IP address must be in the same subnet.

ii. **Gateway IPv6 Address** — The IPv4 address of the router that forwards traffic out of the appliance's subnet.

g. Select the **Additional DNS checkbox** to add DNS server IP addresses as an extra DNS server apart from the primary DNS server. The IP addresses can be either IPv4 or IPv6.

NOTE:

Make sure that:

- You add a DNS server IP address or two DNS server IP addresses as an extra DNS server.
- You do not use the same DNS server IP address for the primary DNS server and extra DNS servers.
- Cascading of DNS servers for authentication, authorization, and auditing is not supported. For authentication, authorization, and auditing of LDAP referrals, RADIUS, and TACACS, the primary DNS server configured for address resolution is always considered for authentication, authorization, and auditing.

8. Click **OK**, and then click **Close**.

To confirm that the appliance is configured correctly, ping the new Management Service IP address or use the new IP address to open the user interface in a browser.

NOTE:

Log on to the NetScaler SDX appliance using the CLI, make sure that the file `/etc/resolv.conf` file included the added extra DNS server IP addresses. Also, the `/mpsconfig/svm.conf` file reflects the added extra DNS server IP addresses.

For example:

```
/mps/changenameserver.sh 127.0.0.1
```

```
/mps/addnameserver.sh 1.2.3.4
```

Initial Configuration through the serial console

To perform the initial configuration of the SDX appliance from outside the L2 domain, connect to the console port of the appliance and follow the instructions carefully.

NOTE

The `networkconfig` utility is available from build 72.5 and later.

To configure the NetScaler SDX appliance by using the serial console

1. Connect the console cable into your appliance.
2. Connect the other end of the cable to your computer and run the vt100 terminal emulation program of your choice.
 - For Microsoft Windows, you can use HyperTerminal.
 - For Apple Macintosh OSX, you can use the GUI-based Terminal program or the shell-based telnet client. Note: OSX is based on the FreeBSD UNIX platform. Most standard UNIX shell programs are available from the OSX command line.
 - To For UNIX-based workstations, you can use any supported terminal emulation program.
3. Press ENTER. The terminal screen displays the Logon prompt. Note: You might have to press ENTER two or three times, depending on which terminal program you are using.
4. At the prompt, type: `ssh nsrecover@169.254.0.10` to log on to the Management Service. In **Password**, if the earlier default password does not work, try typing the serial number of the

appliance. The serial number bar code is available at the back of the appliance. Citrix recommends that you change the password after the first login. For information about changing the password, see [Change the administrative password](#).

5. Type `shell` to switch to the shell prompt and type `networkconfig`

You can now use the new IP address to log on to the Management Service user interface.

Change the password of the default user account

The default user account provides complete access to all features of the NetScaler SDX appliance. For security reasons, the administrator account must be used only when necessary, and only individuals whose duties require full access must know the password for the administrator account. Citrix recommends changing the administrator password frequently. If you lose the password, you can reset the password to the default by reverting the appliance settings to factory defaults, and you can then change the password.

You can change the password of the default user account in the **Users** pane. In the **Users** pane, you can view the following details:

- Name

Lists the user accounts configured on the SDX appliance.

- Permission

Displays the permission level assigned to the user account.

Change the administrative password by using the GUI

1. Log on to the appliance by using the administrative credentials.
2. Navigate to **System > User Administration > Users**.
3. In the **Users** pane, click the default user account, and then click **Modify**.
4. In the **Modify System User** dialog box, in **Password** and **Confirm Password**, enter the password of your choice.
5. Click **OK**.

For more information about more NetScaler SDX software configuration, see [NetScaler SDX software documentation](#).

Lights out management port of the NetScaler SDX appliance

The SDX ADC appliances have an Intelligent Platform Management Interface (IPMI), also known as the lights out management (LOM) port, on the front panel of the appliance. You can use the LOM port to remotely monitor and manage the appliance, independently of the NetScaler software.

By connecting the LOM port to a dedicated channel that is separate from the data channel, you can make sure that connectivity to the appliance is maintained even if the data network is down. Further you can reduce the data cable and data network as a single point of failure.

You can access the LOM port through a browser and use the GUI for most tasks. All tasks can be performed through the NetScaler shell.

You can use either the GUI or a shell for the following tasks:

- Configuring the network settings
- Health monitoring
- Power control operations
- Factory reset

Different NetScaler® ADC appliances support different shells:

- For Citrix Hypervisor based NetScaler SDX appliances, use the dom0 Linux root shell. To access the dom0 shell, log on to the Citrix Hypervisor management IP address instead of the SDX Management Service IP address, using the `root` account, not the `nsroot` account.
- For Linux based appliances, use the Linux bash root shell.

NOTE

The terms LOM and Baseboard Management Controller (BMC) are used interchangeably.

Caution: LOM firmware versions are platform specific. Upgrading to a LOM firmware version other than the one shown for your platform in the LOM support matrix table, might result in the LOM becoming unusable.

LOM support matrix

The LOM support matrix shows the recommended LOM firmware version for different platforms.

Table. LOM support matrix

HARDWARE	RECOMMENDED VERSION
SDX 5900	4.61
SDX 8000	3.56
SDX 8900	4.61
SDX 9100	2.13.12
	3.11.0
	Supported NetScaler builds
	• 14.1 - 43.x • 13.1 - 58.x
SDX 9200	4.1.1
SDX 14000	4.14
SDX 14000-40S	4.14
SDX 14000 FIPS	4.14
SDX 14000-40G	4.14
SDX 15000	5.56
SDX 15000-50G	5.56
SDX 16000	2.13.12
	3.11.0
	Supported NetScaler builds

HARDWARE	RECOMMENDED VERSION	
	• 14.1 - 43.x • 13.1 - 58.x	
SDX 25000A	4.14	
SDX 25000TA	4.14	
SDX 25000-40G	4.14	
SDX 26000	5.56	
SDX 26000-50S	5.56	
SDX 26000-100G	5.56	

Upgrade LOM version to 3.11.0

Starting from release 13.1 build 58.x and release 14.1 build 43.x, LOM version 3.11.0 is available for the following platforms. This version addresses multiple functional issues.

- SDX 9100
- SDX 16000

Upgrade procedure

1. Login to dom0 Linux root shell.
2. Run the following command:

```
echo "3.11" > /opt/xensource/packages/files/sdx-lom/platforms/matua/version
```

3. In the Management Service user interface, navigate to the **Configuration > System**.
4. Click **Reboot Appliance**.

Configuration compatibility

The LOM configuration format is not compatible between the 2.12.12 and 2.13.12 LOM versions and the new 3.11.00 LOM version.

You might lose some LOM configurations during the upgrade process from the 2.12.12 or 2.13.12 version to the new LOM 3.11.00 version. Once the upgrade is complete, you cannot use any manually created LOM 2.12.12 or 2.13.12 version configuration backups to restore the configuration.

NOTE

We recommend saving the current LOM configuration before the upgrade. By saving the LOM configuration you can manually reapply any missing settings afterwards.

The upgrade process does not retain all the LOM configuration settings; manual reapplication might be necessary in the following areas:

- Manually created local LOM user accounts (specifically, all accounts except the default ns root account).
- RADIUS server shared a secret.
- Password for LDAP authentication.
- SMTP password in the **Notifications > SMTP** section.
- Authentication key and encryption key for SNMPv3 alerts in the **Notifications > Alerts** section.
- The SNMPv1 alert Trap Community string in the **Notifications > Alerts** section.
- Any SNMP settings in the **Notifications > SNMP** section.

Configure the network settings on the LOM port

The default IP address for initial access to the LOM port is 192.168.1.3. Change the default credentials and IP address the first time you log on. All LOM GUI operations require you to connect to the appliance by typing the LOM IP address in a web browser and then entering the administrator credentials. Alternatively, you can access LOM functionality through the command line by using the `ipmitool` utility. Using the `ipmitool` utility remotely, you can determine the LOM firmware version number, perform warm and cold restarts, configure LOM network settings, monitor the health of the appliance, and perform power control operations. The utility is available for download at <http://ipmitool.sourceforge.net/>. The `ipmitool` utility is also included in NetScaler MPX and CloudBridge/SDX (dom0) appliances for initial LOM port network configuration. When using the shell, you can choose to use DHCP or static IP settings for initial network configuration. After configuring the network settings, you can use the `ipmitool`

commands over the network. For example, the BMC firmware revision command would need the same user name, password, and IP address that is used to access the BMC/LOM GUI port.

For initial configuration, connect the network port on your laptop or workstation directly to the LOM port with a crossover cable, or to a switch in the same local subnet(192.168.1.x) as the LOM port. Assign a network-reachable IP address and change the default credentials. After saving the new settings, the LOM restarts and the changes take effect. After the restart, you must use the new address to access the LOM.

If you make a mistake that results in losing network connectivity at both the old and new IP addresses, you must use the local shell method to recover.

See the [Secure Deployment Guide](#) for best practices for managing administrative credentials and configuring your network for a secure LOM deployment.

NOTE

On all SDX platforms, except SDX 22040/22060/22080/22100/22120 and SDX 24100/24150, the LEDs on the LOM port are nonoperational by design.

Tip: For first-time setup in a network, to facilitate troubleshooting, make sure that a laptop/PC is connected directly to the LOM port. If you can ping and access the LOM GUI at the default IP address (192.168.1.3) by using static addressing on the computer, but remote access does not work:

- Take a closer look at the network firewall settings and access control list (ACL) policies of all network devices along the network path.

Tip: If some LOM GUI features work but others do not (for example, normal NetScaler console output is visible in the NetScaler console window in the LOM GUI, but typing in the console does not work), try the preceding method to isolate the cause to the specific BMC protocol being blocked by the network.

Tip: Use the iKVM (HTML5) feature to access the LOM GUI. Alternately, use Java. Ensure that the latest Java updates are installed on your computer.

To configure the NetScaler LOM port by using the GUI

1. In a web browser, type `<http://192.168.1.3>` and enter the default user credentials. In Password, if the earlier default password does not work, try typing the serial number of the appliance. The serial number bar code is available at the back of the appliance.

NOTE

The NetScaler LOM port is preconfigured with IP address 192.168.1.3 and subnet mask 255.255.255.0. See [Troubleshooting](#) if some of the options are grayed out.

2. On the **Configuration** tab, click **Network** and type new values for the following parameters:
 - IP Address—IP address of the LOM port.
 - Subnet Mask—Subnet mask used to define the subnet of the LOM port.
 - Default Gateway—IP address of the router that connects the LOM port to the network.
3. Click **Save**.
4. If you want to change the user credentials, navigate to **Configuration > Users**, select the user, click **Modify User**, and change the credentials.

To configure the NetScaler LOM port by using the shell

Note: You need superuser (admin) credentials to access the shell.

1. Configure the IP addressing mode:
 - To use DHCP, at the shell prompt, type:

```
ipmitool lan set 1 ipsrc dhcp
```

No further IP-level configuration is required.

- To use static addressing, at the shell prompt type:

```
>ipmitool lan set 1 ipsrc static  
  
>ipmitool lan set 1 ipaddr (LOM IP address)  
  
>ipmitool lan set 1 netmask (netmask IP address)  
  
>ipmitool lan set 1 defgw ipaddr <default gateway IP address>
```

The BMC reboots to apply the changes. Pings to the BMC succeed after approximately 60 seconds.

2. Optionally, to configure Ethernet VLAN ID and priority, at the NetScaler shell prompt type:


```
>ipmitool lan set 1 vlan id <off|ID>  
  
>ipmitool lan set 1 vlan priority <priority>
```

You can either disable or enable the VLAN. Set the VLAN ID to a value from 1 to 4094, and the VLAN priority to a value from 0 to 7. After the network settings have been correctly applied, you can access the `ipmitool` remotely from a physically separate machine over the network. For remote access, enter the BMC user name, BMC password, and the BMC IP address. For example, to run the `ipmitool mc info` command, at the shell prompt on a remote machine, type:

```
ipmitool -U <username> -P <password> -H <bmc IP address> mc info
```

Obtain health monitoring information

There are two NetScaler MIBs: the NetScaler software management MIB and the NetScaler IPMI LOM hardware management MIB. The software management MIB is primarily used for monitoring the application software and the application software's utilization of hardware resources, such as CPU % and memory %. It provides a high level view of the appliance and is therefore suitable for the application monitoring function carried out by an application group within an organization. The LOM MIB is used for monitoring the hardware health and therefore provides a lower level view of the appliance. This type of monitoring is more applicable to the network monitoring function carried out by a network monitoring group.

The LOM SNMP traps in the LOM MIB report hardware failures. The NetScaler SNMP traps in the NetScaler MIB report software failures and hardware load issues.

The NetScaler MIB has a small subset of hardware sensors. It does not cover any BIOS level failures, because the BIOS checks the hardware primarily during boot time, before the NetScaler software starts. If the BIOS detects a failure, it does not load the boot loader. If the boot loader does not load, the operating system does not load, and therefore the NetScaler SNMP software service responsible for sending the traps does not load.

The NetScaler software Management MIB issues a warning under the following conditions only:

1. If the failure is gradual enough for the main CPU to issue an SNMP alert. An electrical failure close to the CPU, such as a failed electrical capacitor, occurs too quickly for the CPU to issue an alert.
2. If the failure happens after the BIOS, Operating System, and SNMP service have started and normal boot-up has been successful.

3. If the failure happens while the operating system and other system software are in a stable enough state for the SNMP software service to run.

Whenever the NetScaler MIB is unable to report these warnings, because of hardware or software failure, the LOM MIB monitors and reports the warnings. The LOM microcontroller operates independently of the NetScaler software. To monitor the hardware and software of the NetScaler appliance, you must use both the NetScaler MIB and the LOM MIB.

The NetScaler IPMI LOM hardware management MIB SNMP firmware runs on the BMC microcontroller chip. The BMC chip CPU sends a warning in the case of a hardware failure, regardless of whether any of the preceding conditions occur. For example, if the BIOS halts the system during boot-up because of a memory DIMM failure, the BMC chip uses the BIOS POST code snooping mechanism to detect the failure, and sends a bad DIMM SNMP alert.

You can log on to the LOM port to view the health information about the appliance. All system sensor information, such as system temperature, CPU temperature, and status of fans and power supplies, appears on the sensor readings page. The Event Log records time stamps of routine events such as a power cycle, in addition to recording hardware-failure events. If SNMP traps are enabled, these events can be sent to your SNMP Network Monitoring software. For more information about how to set up an SNMP alert, see [Configuring SNMP Alerts](#).

To obtain health monitoring information

1. In the **Menu** bar, click **System Health**.
2. Under **Options**, click **Sensor Readings**.

Install the MIB

Download the IPMI SNMP management information base (MIB) for your LOM firmware version, and import it into the SNMP monitoring software.

For a sample configuration, see `<http://www.net-snmp.org/tutorial/tutorial-5/commands/snmptrap.html>`. For the exact steps of this procedure specific to your environment, contact your SNMP network monitoring software provider.

Configure SNMP Alerts

You can configure SNMP alerts on the LOM. Optionally, you can configure an alert to send emails.

To configure the alerts, you can use the LOM GUI or the NetScaler shell.

To configure SNMP alerts on the LOM by using the GUI

1. Download the IPMI View utility from `<ftp://ftp.supermicro.com/utility/IPMIView/>` and install it on your computer. Use this utility to test the configuration. For more information, see the section about configuring the alert settings in the IPMI View User Guide at `<http://supermicro.com>`.
2. Open the IPMI View utility.
3. In the LOM GUI, navigate to **Configuration > Alerts**, click **Alert No 1**, and then click **Modify**.
4. Select the severity level of the events for which to generate alerts.
5. Set Destination IP to the IP address at which you installed the IPMI View utility.
6. Optionally, to receive alerts by email, specify an email address. To avoid receiving email for routine alerts, specify a severity higher than Informational.
7. Click **Save**.
8. The LOM starts sending alerts to the IPMI View utility within a minute or two. After the IPMI View utility starts receiving alerts from the LOM, reconfigure the destination IP address to point to your SNMP Network Management Software, such as HP OpenView.

Set up SNMP alerts on the LOM by using the NetScaler shell

To customize your filter and policy settings, see the IPMI Specification 2.0 rev. 1.1 documentation.

The latest IPMI specifications are available from the IPMI section of the Intel website:

```
<http://www.intel.com/content/www/us/en/servers/ipmi/ipmi-specifications.html>
```

Usually, customization in the SNMP Network Management Software is the preferred method, because it can be done one time at a central location. Therefore, the following settings send all events for all sensors to the SNMP network management software. These events are low traffic events and therefore do not result in any significant network usage.

To set up SNMP filters to allow all events:

```
ipmitool raw 4 0x12 0x6 0x10 0x80 1 1 0 0xff 0xff 0xff 0xff 0xff 0xff 0  
0xff 0 0 0xff 0 0 0xff 0
```

To set up a policy list for all sensors and events:

```
ipmitool raw 4 0x12 9 0x10 0x18 0x11 0x81
```

To set up a destination IP address for an SNMP event:

```
ipmitool lan alert set 1 1 ipaddr <x.x.x.x>
```

Where, <x.x.x.x> is the IP address to which the SNMP event must be sent.

To specify an SNMP community string name:

```
ipmitool lan set 1 snmp <community string>
```

Create and install an SSL certificate on SDX 9100 and SDX 16000

Use OpenSSL to create a key, generate a certificate signing request (CSR), and generate a self-signed SSL certificate. Use the LOM GUI to install the SSL certificate. After successful installation, the LOM GUI uses the new SSL certificate for secure connections.

Prerequisites

- Access to an SDX platform running FreeBSD or Linux machine with OpenSSL installed.
- Administrative access credentials (nsroot/password) to the LOM web interface.

Create an SSL certificate by using OpenSSL

Perform the following steps from OpenSSL:

1. Generate a 2048-bit RSA private key. Type: `openssl genrsa -out private_key.pem 2048`

A file named private_key.pem is created in the current directory containing the RSA private key.

2. Generate a CSR using this private key. Type: `openssl req -new -key private_key.pem -out request.csr` Enter values for the following parameters:

- Country Name (2 letter code): For example, US
- State or Province Name (full name): For example, California
- Locality Name: For example, San Francisco
- Organization Name: For example, My Company
- Organizational Unit Name: For example, IT Department
- Common Name For example, yourname (or your server's hostname)
- Email Address: For example, <mailto:admin@mycompany.com>
- Challenge password: Leave blank
- Optional company name: Leave blank

A file named request.csr is created in the current directory containing the certificate signing request.

3. Generate a self-signed SSL certificate using the CSR and the private key. The certificate is valid for 365 days. Type: `openssl x509 -req -in request.csr -signkey private_key.pem -out lom_certificate.pem -days 365`

A file named lom_certificate.pem is created in the current directory containing the self-signed SSL certificate.

Install an SSL certificate by using the LOM GUI

1. Log in to the LOM web interface using the default administrator (`nsroot`) account.
2. Navigate to **Configuration > Network > SSL Certificates**.
3. In **New SSL Certificate**, click **Select File**.
4. Browse to the location of the certificate file (lom_certificate.pem).
5. In **New Private Key**, click **Select File**.
6. Browse to the location of the key file (private_key.pem).
7. Click **Upload**.

The BMC will reboot after uploading the SSL certificate and private Key to the LOM UI and the system is locked by default. Check the SSL certificate details in the browser to confirm it matches the lom_certificate.pem you generated.

Example

```
[root@netscaler-sdx SSL_Cert]# openssl version

OpenSSL 1.0.2k-fips  26 Jan 2017

[root@netscaler-sdx SSL_Cert]# uname -a

Linux netscaler-sdx 4.4.0+2 #1 SMP Fri Apr 30 02:46:33 UTC 2021 x86_64
x86_64 x86_64 GNU/Linux

[root@netscaler-sdx SSL_Cert]# cat /etc/centos-release

XenServer release 7.1.2 (xenenterprise)

[root@netscaler-sdx SSL_Cert]# openssl genrsa -out private_key.pem 2048
```

```
Generating RSA private key, 2048 bit long modulus
```

```
.....+++
```

```
.....+++
```

```
e is 65537 (0x10001)
```

```
[root@netscaler-sdx SSL_Cert]# ls
```

```
private_key.pem
```

```
[root@netscaler-sdx SSL_Cert]# openssl req -new -key private_key.pem -out  
request.csr
```

You are about to be asked to enter information that will be incorporated
into your certificate request.

What you are about to enter is what is called a Distinguished Name or a
DN.

There are quite a few fields but you can leave some blank

For some fields there will be a default value,

If you enter '.', the field will be left blank.

```
-----
```

```
Country Name (2 letter code) [XX]:US
```

```
State or Province Name (full name) []:California
```

```
Locality Name (eg, city) [Default City]:Santa_Clara
```

```
Organization Name (eg, company) [Default Company Ltd]:CSG
```

```
Organizational Unit Name (eg, section) []:Engineering
```

```
Common Name (eg, your name or your server's hostname) []: cloud.com
```

```
Email Address []: none@cloud.com
```

Please enter the following 'extra' attributes

to be sent with your certificate request

A challenge password []:

An optional company name []:

```
[root@netscaler-sdx SSL_Cert]# ls
```

```
private_key.pem  request.csr
```

```
[root@netscaler-sdx SSL_Cert]# openssl x509 -req -in request.csr -signkey  
private_key.pem -out lom_certificate.pem -days 365
```

Signature ok

```
subject=/C=US/ST=California/L=Santa_Clara/O=CSG/OU=Engineering/  
CN=cloud.com/emailAddress=none@cloud.com
```

Getting Private key

```
[root@netscaler-sdx SSL_Cert]# ls
```

```
lom_certificate.pem  private_key.pem  request.csr
```

```
[root@netscaler-sdx SSL_Cert]# openssl version
```

```
OpenSSL 1.0.2k-fips  26 Jan 2017
```

```
[root@netscaler-sdx SSL_Cert]# uname -a
```

```
Linux netscaler-sdx 4.4.0+2 #1 SMP Fri Apr 30 02:46:33 UTC 2021 x86_64  
x86_64 x86_64 GNU/Linux
```

```
[root@netscaler-sdx SSL_Cert]# cat /etc/centos-release
```

```
XenServer release 7.1.2 (xenenterprise)
```

```
[root@netscaler-sdx SSL_Cert]# openssl genrsa -out private_key.pem 2048
```

Generating RSA private key, 2048 bit long modulus

.....+++

.....+++

e is 65537 (0x10001)

```
[root@netscaler-sdx SSL_Cert]# ls
```

private_key.pem

```
[root@netscaler-sdx SSL_Cert]# openssl req -new -key private_key.pem -out  
request.csr
```

You are about to be asked to enter information that will be incorporated
into your certificate request.

What you are about to enter is what is called a Distinguished Name or a
DN.

There are quite a few fields but you can leave some blank

For some fields there will be a default value,

If you enter '.', the field will be left blank.

Country Name (2 letter code) [XX]:US

State or Province Name (full name) []:California

Locality Name (eg, city) [Default City]:Santa_Clara

Organization Name (eg, company) [Default Company Ltd]:CSG

Organizational Unit Name (eg, section) []:Engineering

Common Name (eg, your name or your server's hostname) []: cloud.com


```
Email Address []: none@cloud.com
```

```
Please enter the following 'extra' attributes
```

```
to be sent with your certificate request
```

```
A challenge password []:
```

```
An optional company name []:
```

```
[root@netscaler-sdx SSL_Cert]# ls
```

```
private_key.pem  request.csr
```

```
[root@netscaler-sdx SSL_Cert]# openssl x509 -req -in request.csr -signkey  
private_key.pem -out lom_certificate.pem -days 365
```

```
Signature ok
```

```
subject=/C=US/ST=California/L=Santa_Clara/O=CSG/OU=Engineering/  
CN=cloud.com/emailAddress=none@cloud.com
```

```
Getting Private key
```

```
[root@netscaler-sdx SSL_Cert]# ls
```

```
lom_certificate.pem  private_key.pem  request.csr
```

Sample Private_Key.pem (2048 bits)

```
-----BEGIN RSA PRIVATE KEY-----
```

```
MIIEogIBAAKCAQEAlW0N9EghsirHkPsgAiwXHuDnW6D0stWC9V7ZLu38g36pqYk
```

```
X1yh0N4yF61nCqS+Mfiu+p1PEDbXc5dwtd7Le4JEW80FHko5mBz3Uj9MSenJTBVT
```

```
V1d+hyu82nR0VWdAxe72kPiiTG7PuM34ZZqeHuy+tLU9FkAXDtxyj1Pf/TDJPQ1/
```

```
7h1k1nTW5t9s2rSXznj3jsX0g88xF8S/ksGP106F507XoZYtpafUB/SqfZuX95qv
```

12EbbusifBwrmKVd/Z5/FguSGCeszBmrCol0dfUT44lMRPP2zW0m60234j4Hm0uh
Z/qL34seA1SmIMHkr9K57VHWqWoc3NmTDK2n0wIDAQABAOIBAEaI7AVb50wLFZlC
caGU3nUeatYKa7fdQE7VvyGV2wbPwnnYA2Ngf1oxyMOUfwdvLLi6Ax8YL1IgESPn
YLT5LmRm4Y3PhY7+yqEbS0PvzTs2uGbJyKV4b6mu477bKe8GhnTokEtukm4yAKvj
z1Vn3ywJ5Ru4UdWm+ztwPEK7JQiGvkoMSn7+qDhEAw5sw9cdCRxqLITm0xVhda5L
x0576NCAvSMCGlkwu9F6nnscP2TIT6Do44AiERAQRUDwVrNWNAbQQwQ+9Q0LRLIn
lg+h8UIzBdUo4+HPm6ewSKTM1VBtQWTrLzKSEatUoqxXBznx1FpPbjo0G5xcTnXS
pzoztIECgYEAXxV1c706Hcy56hNv1q4G1UYDhRvBMIyeCnZWrr5lFD0Q5v7I/azH
i2EN4B/XZEQjstGxjGzCUXwngZoMykth67+bfynGkuAzdDQFgTXZCT1LSGwbJJQw
XhMTkL9Qpc9F1pByeX3G/o5pBc4xtjInMqf1wPEFD14ujeE/WMC0eCsCgYEAwCU9
azqLK1LMPgjcmeP4Zl51xTVbaCzo5mreQAMoljxc4QDEW+HIK7/hhD57/MyA7Zya
6xIhIQ0q26xGV+MrJ9tWIhyUfoC07kV9jc0TcszaMAzNk0+vKVEwmvQC7833csC9
Tx/hS4jSiYkyqeIVYV0IZWjZvSlFqfmsZCl1UvkCgYAGSPo1RkMQb8cqyaRNPb8/
em9gA2lM9B0FNwkziRU9110yRzbnM041a0H54md1ZeV0nuohCcn0spu4rrlapZiY
zbrF//ah6mwVbhMCakXI0e0PjI2iUP9Z6PZEKpYMNgf5ZyQGC999Z9e09FyDaP/t
zHmotrM3Mfz8Q7ExRxwrcwKBgCwRnneUcszZ7Zwgvi/S9hBMg+haJ3/KJRSL2D0I
RevHJmo4mdCVBSr7lEaXaip940gebe2SN+Tztyuw5GVN5dz7U1L6iRhGfrjTWUuH
iHl1AH3awd+SfBMW2vX/Fx1W/PTlvcWDA1ImrFr9C5CpGMw+4SkZF19rYt3sNjy4
YX9JAoGABAKC332U9iTPXajrygPkAtDjJQXnwp0BAXHU9NxNHF00X2gkU+RQSotC
8muqyDgA0SLuagd0bAZS59wQCjTtHXB8WR5R5fri7a1Mdmw+/0nhGG09CwV4Vw/n
ZEFexXmC1uU01jffj0HtXrKgmka1j/1e60tjr/r6UxV/vM900gtw=

```
-----END RSA PRIVATE KEY-----
```

Sample request.csr

```
-----BEGIN CERTIFICATE REQUEST-----
```

```
MIIC3DCCAcQCAQAwwgZYxCzAJBgNVBAYTA1VTMRMwEQYDVQQIDApDYWxpZm9ybmlh
MRQwEgYDVQQHDAtTYW50YV9DbGFyYTEMMAoGA1UECgwDQ1NHMRQwEgYDVQQLDAtF
bmdpbmVlcm1uZzESMBAGA1UEAwwJY2xvdWQuY29tMSQwIgYJKoZIhvcNAQkBFhVr
aG9hLm5ndXl1bkBjbG91ZC5jb20wgwEiMA0GCSqGSIb3DQEBAQUAA4IBDwAwggEK
AoIBAQCvBq30SCGyKseQ+yACLBce4903DoM6y3AL1Xtku7fyDfqmpiRfXKE43jIX
rWcKpL4x+K76mU8QNtdz13C13st7gkRbw4UeSjmYHPdSP0xJ6c1MFVNxV36HK7za
dHRVZ0DF7vaQ+KJMbs+4zfhlmp4e7L60tT0WQBc03HK0U9/9MMk9CX/uHWSwdNbm
32zatJf0ePe0xc6DzzEXxL+SwY/U7oXk7teh1i2lp9QH9Kp9m5f3mq/XYRtu6yJ8
HCuYpV39nn8WC5IYJ6zMGasKiXR19RPjiUxE8/bNbSbo7bfiPgeY66Fn+ovfix4D
VKYgweSv0rntUdapahzc2ZMMrafTAgMBAAGgADANBgkqhkiG9w0BAQsFAA0CAQEA
I9rsE/tsUgGNHzgsJ6Bq45a1AFN1hDS2IPuvEVl+3eGECR6ZRMxcHRn3ZGS0ISmD
mYeTfzwxT/IcCRXbVqFvAPkdiRwZwDvcjDg3BN1petrSURbyYbd9BTRkthSn+N10
9Szyik0EKbuHnQ9TbeCB0vAKpSFiz+jImfNYXYilvFTsJRJaDFDN0Xlv6QgDAX2J
qxjVai0isJVz5QXjzD/NpWtFFQJYb1LphS4uXmM1lg3sWaAI9LhfAd1wJC9uSr8R
CZyqnV6UjC0AXMsBbchr2SNa5lWNBH6F9Kjv2StlrpglmJfq63MqDQLq3HFfKrfA
aQi90lrjSMWkdIQX22pZxw==
```

```
-----END CERTIFICATE REQUEST-----
```

Sample lom_certificate.pem

```
-----BEGIN CERTIFICATE-----

MIIDqjCCApICCDaiXN+twSKSjANBgkqhkiG9w0BAQsFADCB1jELMAkGA1UEBhMC

VVMxEzARBgNVBAgMCkNhbg1mb3JuaWExFDASBgNVBACMC1NhbnRhX0NsYXJhMQww

CgYDVQQKDANDU0cxFDASBgNVBASMC0VuZ2luZWVyaW5nMRIwEAYDVQQDDA1jbG91

ZC5jb20xJDAiBgkqhkiG9w0BCQEFWtob2Eubmd1ewVuQGNsb3VkLmNvbTAeFw0y

NDA2MTEwMjAxMDRaFw0yNTA2MTEwMjAxMDRaMIGWMQswCQYDVQQGEwJVUzETMBEG

A1UECAwKQ2FsaWZvcn5pYTEUMBIGA1UEBwwLU2FudGFfQ2xhcmExDDAKBgNVBAoM

A0NTRzEUMBIGA1UECwwLRW5naW5lZXJpbmcxEjAQBgNVBAMMCWNSb3VkLmNvbTEk

MCIGCSqGSib3DQEJARYVa2hvYS5uZ3V5ZW5AY2xvdWQuY29tMIIBIjANBgkqhkiG

9w0BAQEFAAOCAQ8AMIIBCgKCAQEAlW0N9EghsirHkPsgAiwXHuPdNw6D0stwc9V7

ZLu38g36pqYkX1yh0N4yF61nCqS+Mfiu+p1PEDbXc5dwt7Le4JEW80FHko5mBz3

Uj9MSenJTBVTV1d+hyu82nR0VWdAxe72kPiiTG7PuM34ZZqeHuy+tLU9FkAXDtxy

j1Pf/TDJPQ1/7h1k1nTW5t9s2rSXznj3jsX0g88xF8S/ksGP106F507XoZYtpafU

B/SqfZuX95qv12EbbusifBwrnKVd/Z5/FguSGCeszBmrCol0dfUT441MRPP2zw0m

60234j4Hm0uhZ/qL34seA1SmIMHkr9K57VHWqWoc3NmTDK2n0wIDAQABMA0GCSqG

Sib3DQEBCwUAA4IBAQBLojhxNsw24N0uTIQ3dJh6WZATiiBjw8kQyEJqyiB8oCm0

oAVleDAjI44C2eaR1vj321yNQI6bmBGPffwUoIX6YMAf1l6nJq0f19+rJd1FYCCd

FIqt76sC9YTu8WL3j7X1LE2lhQj7RZUt321QcG30qxQoXlQIM5oP7q17WkmPY0tW

JQZ4LjQRGHtc9rDiSlkzeMeBgtG3HqdNSorn2S15JJf/4sm5JXQXd7GByicv9aNm

Aagjqw1kziJUpL02r2bRX+3Qn0NE5WlxaYYisIPe9py3TsnLXHcrnTqrHbh6e4wc

+yF9+4nouCHPj0s2i0QV7koFHZ8lnEiUYaxYT8w1
```

```
-----END CERTIFICATE-----
```

Install a certificate and key on older platforms using the LOM GUI

Citrix® recommends using HTTPS to access the LOM GUI. To use HTTPS, you must replace the default SSL certificate with one from a trusted certificate authority and upload a private key to the LOM GUI.

To encrypt SNMP alerts, set up an SSL certificate and private key. In the GUI, navigate to **Configuration > SSL Certification** and apply the SSL certificate and private key. See the NetScaler Secure Deployment Guide for more information about how to securely deploy the LOM in your network. To enable encryption and learn the security measures for LOM, see [<http://support.citrix.com/article/CTX200084>](http://support.citrix.com/article/CTX200084).


If you make a mistake, you must restore the [BMC to the factory defaults](#) to erase the certificate and key.

NOTE


The certificate file must contain only the certificate. The certificate and key must not be in the same file. Make sure that the certificate contains only the certificate and that the key file contains only the key.

To upload a trusted certificate and private key by using the LOM GUI


1. Navigate to **Configuration > SSL Certification**.




2. In the right pane, click the **Choose File** buttons to select a new SSL certificate and a new private key.




3. To verify that you have selected the correct certificate and private key, check the file names of the certificate and key. The file names appear next to the **Choose File** buttons.




4. Click **Upload**. A message informs you that uploading a new SSL certificate replaces the existing (default) certificate.
5. Click **OK**.



6. When a message informs you that the certificate and key have been uploaded successfully, click **OK** to reset the device.



The reset takes approximately 60 seconds. You are then redirected to the login page.




7. Log on to the LOM GUI by using your default credentials.


NOTE

If the certificate or key is invalid, the BMC reboots, tries the new settings, and reverts to using the previous settings.

8. In the address bar, click the lock icon to display the connection tab, as shown in the following screen.



9. Click **Certificate** information to display details about the certificate that you uploaded.



Note: For the best practices for LOM and NetScaler security, see <http://support.citrix.com/article/CTX200084>.

Obtaining the MAC address, serial number, and host properties of the appliance

A Media Access Control address (MAC address) is a unique identifier assigned to network interfaces for communication on the physical network segment. The serial number is on the back

panel of the appliance. If you do not have easy access to the back panel, you can get the appliance's serial number by logging on to the lights out management (LOM) port. You can also retrieve the parameter settings assigned to the IP addresses configured on the appliance, such as the state of ARP, ICMP, telnet, secure shell access, and dynamic routing.

To obtain the MAC address, serial number, and host properties of the appliance by using the LOM GUI

1. Log on to the LOM GUI.
2. Navigate to **System > Component Info**.
3. In the **Overview** tab, you can view the MAC address, serial number, and host properties of the appliance.

To obtain the MAC address and host properties of the BMC by using the appliance shell

At the shell prompt, type:

```
ipmitool lan print
```

Example

```
Set in Progress      : Set Complete

Auth Type Support    : MD2 MD5 OEM

Auth Type Enable     : Callback : MD2 MD5 OEM

                     : User      : MD2 MD5 OEM

                     : Operator  : MD2 MD5 OEM

                     : Admin     : MD2 MD5 OEM

                     : OEM       :

IP Address Source    : Static Address

IP Address           : 192.168.1.3
```

```
Subnet Mask           : 255.255.255.0

MAC Address           : 00:25:90:3f:5e:d0

SNMP Community String : public

IP Header             : TTL=0x00 Flags=0x00 Precedence=0x00 TOS=0x00

BMC ARP Control       : ARP Responses Enabled, Gratuitous ARP Disabled

Gratituous ARP Intrvl : 0.0 seconds

Default Gateway IP    : 0.0.0.0

Default Gateway MAC    : 00:00:00:00:00:00

Backup Gateway IP     : 0.0.0.0

Backup Gateway MAC    : 00:00:00:00:00:00

802.1q VLAN ID       : Disabled

802.1q VLAN Priority  : 0

RMCP+ Cipher Suites   : 1,2,3,6,7,8,11,12,0

Cipher Suite Priv Max : aaaaXXaaaXXaaXX

                        :      X=Cipher Suite Unused

                        :      c=CALLBACK

                        :      u=USER

                        :      o=OPERATOR

                        :      a=ADMIN

                        :      O=OEM
```

Perform power control operations by using the LOM port

Through the LOM port, you can remotely perform power control operations, such as graceful shutdown and restart, power cycling the appliance, and restarting the BMC microcontroller. A cold restart takes longer than a warm restart. In a cold restart, you switch off power to the appliance and then switch it back on.

Perform power control operations by using the GUI

1. In the **Menu** bar, click **Remote Control**.
2. Under **Options**, click **Power Control**, and then select one of the following options:
 - **Reset System**—Gracefully restart the appliance. All operations on the appliance are stopped, no new connections to the client or server are accepted, and all existing connections are closed before the appliance restarts. This option is similar to a warm restart, such as by entering the reboot command. The BMC does not reboot itself during this operation.
 - **Power Off System – Immediate**—Disconnect power to the appliance immediately, without gracefully shutting down the appliance. The BMC continues to operate normally in this mode to allow the user to remotely power on the appliance. This option is the same as pushing the power button until the unit powers off.
 - **Power Off System – Orderly Shutdown**—Gracefully shut down the appliance, and then disconnect power to the appliance. Has the same effect as pressing the power button on the back panel of the appliance for less than four seconds. All operations on the appliance are stopped, no new connections to the client or server are accepted, and all existing connections are closed before the appliance shuts down. The BMC continues to operate normally in this mode to allow the user to remotely power on the appliance. This option is the same as entering the shutdown command in the appliance shell.
 - **Power On System**—Switch on the appliance. The BMC does not reboot itself during this operation. This option is the same as pushing the power button.
 - **Power Cycle System**—Switch off the appliance, and then turn it back on. The BMC does not reboot itself during this operation. This option is the same as pushing the power button until the unit powers off, and then pushing the power button to power on the unit.
3. Click **Perform Action**.

Perform a power cycle of the BMC

A warm restart, cold restart, or a power cycle of the appliance, using the power button, does not include power cycling the BMC. The BMC runs on standby power directly from the power supply. Therefore, the BMC is not affected by any state of the power button on the appliance. The only way to power cycle the BMC is to remove all power cords from the appliance for 60 seconds.

Perform power control operations on the BMC by using the appliance shell

When performing either a warm or cold restart of the BMC microcontroller, you cannot communicate with the LOM port. Both actions restart the BMC but not the main CPU. To perform a warm restart of LOM from the appliance, type:

```
ipmitool mc reset warm
```

Perform a warm restart remotely from another computer on the network

```
ipmitool -U <bmc_gui_username> -P <bmc_gui_password> -H <bmc IP address> mc  
reset warm
```

Perform a cold restart of the LOM from the appliance

```
ipmitool mc reset cold
```

Perform a warm restart remotely from another computer on the network

```
ipmitool -U <bmc_gui_username> -P <bmc_gui_password> -H <bmc IP address> mc  
reset cold
```

Perform a core dump

If the appliance fails or becomes unresponsive, you can remotely perform a core dump. This procedure has the same effect as pressing the **NMI** button on the back panel of the appliance.

Perform a core dump by using the GUI

1. In the **Menu** bar, click **Remote Control**.
2. Under **Options**, click **NMI**, and then click **Initiate NMI**.

Perform a core dump remotely from another computer on the network by using the shell

At the shell prompt, type:

```
ipmitool -U <bmc_gui_username> -P <bmc_gui_password> -H <bmc IP address> chassis
power diag
```

Restoring the BMC configuration to factory defaults

You can restore the BMC to its factory-default settings, including deleting the SSL Certificate and SSL key.

To reset the configuration to factory defaults by using the GUI

1. Navigate to **Maintenance > Factory Default**.
2. Click **Restore**.

To reset the configuration to factory defaults by using the shell


At the shell prompt, type:

```
ipmitool raw 0x30 0x40
```

Specifying the port for IPMI BMC failover

With LOM firmware version 3.x or later, the default mode for failover between the dedicated LOM port and the shared LOM/management port is to fail over to the active port. By default, no user configuration is needed other than selecting the port to which to connect the cable. The motherboard has an Ethernet switch between the management MAC and the management port, and between the LOM MAC and the LOM port. The following figure shows the Ethernet switch.

Figure 1. Ethernet Switch



You can set this switch to direct LOM traffic through the dedicated LOM port or through the shared management port. A dedicated LOM port removes the management port as a single point of failure, while a shared LOM/management port reduces the cabling costs.

Using the BIOS POST code to detect errors

You can read the BIOS POST code by using the LOM GUI or the shell. To interpret the BIOS Beep codes, see `<https://www.supermicro.com/manuals/other/AMI\_AptioV\_BIOS\_POST\_Codes\_for\_SM\_Motherboards.pdf>`.

To read the BIOS Post Code by using the LOM GUI

Navigate to **Miscellaneous > BIOS Post Snooping**.

To read the BIOS POST code by using the shell

At the prompt, type:

```
ipmitool raw 0x30 0x2a
```

LOM port management for SDX 9200 and SDX 17000

This section provides step-by-step instructions for the initial configuration of the Lights Out Management (LOM) interface on NetScaler SDX 9200 and 17000 series platforms. The LOM enables essential out-of-band management, allowing you to manage the appliance independently of the main operating system.

Prerequisites

Before beginning the LOM configuration, ensure that the following requirements are met:

- Physical or console access to NetScaler.
- Network connectivity to the LOM port (direct or switched).
- Device serial number (required for initial authentication).
- Planned IP addressing scheme for the LOM network segment.
- Administrative access credentials.
- Web browser with HTML5 support (Chrome, Firefox, Edge, or Safari).
- Standard HTTPS (port 443) network access to the LOM IP address.

NOTE:

Java or third-party plug-ins are not required.

Initial network access

The LOM interface is configured with default settings to facilitate initial configuration.

Default LOM settings

PARAMETER	DEFAULT VALUE
IP Address	192.168.1.3
Subnet Mask	255.255.255.0
Default Gateway	192.168.1.1
Username	nsroot
Password	Device serial number (case-sensitive)

Network connection methods

Select one of the following connection methods for initial setup:

Direct connection:

- Connect a laptop Ethernet port to the LOM interface using a crossover cable.
- Set the laptop static IP address to 192.168.1.10 with a subnet mask of 255.255.255.0.

Switched connection:

- Connect the LOM interface and a management workstation to the same network switch.
- Ensure that both devices are in the 192.168.1.0/24 broadcast domain.

Configure LOM by using the GUI

Initial login and password change

Access the LOM web interface using the default IP address.

1. Open a web browser

- a. Type https://192.168.1.3.
- b. Accept the self-signed SSL certificate warning that appears in the browser.

2. Authenticate with default credentials

- a. Enter **nsroot** as the Username.
- b. Enter the device serial number as the password (note: this is case-sensitive).
- c. Click **Login**.

3. Mandatory password change

On the first login, the system enforces a password change for security.

- **Current Password:** Enter the device serial number.
- **New Password:** Enter a password that meets the complexity requirements shown on the page.
- **Confirm Password:** Re-enter the new password.
- Click **Submit** to save.

NOTE:

The system logs you out automatically once the password is saved. Log in again using the **nsroot** account and your newly created password.

Password should be changed for default user.

- Password should have a minimum of 8 characters.
- Password is not allowed combination of <> characters, ex: "123<4567>"
- Password is not allowed reverse of username.
- Password should have combination of at least two classes from a-z, A-Z, 0-9, special characters.
- Special characters !, \$, %, &, (,), *, +, -, ., /, <, =, >, @, [\] ^ _ { | } ~, ' are allowed.

New Password

Confirm Password

Submit

net>scaler

Network configuration

To integrate LOM into your production management network:

1. Navigate to **Settings > Network IP Settings > LAN Interface (eth0)**.

2. Enter the following details:

a. **IPv4 Address**

b. **IPv4 Subnet**

c. **IPv4 Gateway**

3. Click **Save**. The Baseboard Management Controller (BMC) takes approximately 30 seconds to apply changes and then disconnects your current session.

4. Verify connectivity by pinging the new LOM IP address by using the following command from your management workstation:

```
ping <New_LOM_IP_Address>
```

5. Access the LOM interface using the new IP address:

```
https://<new_LOM_IP_address>
```

6. Log using `nsroot` and updated password.

The screenshot shows the NetScaler SDX web interface. On the left is a dark blue sidebar with the 'net>scaler' logo at the top. Below the logo is a 'Firmware Information' section showing version 4.01.01, a timestamp 'Oct 29 2025 09:48:35 UTC', and a 'Host Online' status. A 'Quick Links..' dropdown is below that. The sidebar menu includes: Dashboard, Sensor, System Inventory, FRU Information, Logs & Reports, Settings, Remote Control, Image Redirection, Power Control, Maintenance, and Sign Out. The main content area has a blue header with a hamburger menu icon. The title 'Network IP Settings' is displayed. A green question mark icon is in the top right of the settings box. The settings are as follows: 'Enable LAN' is checked; 'LAN Interface' is set to 'eth0'; 'MAC Address' is '08:35:71:31:D8:6E'; 'Enable IPv4' is checked; 'Enable IPv4 DHCP' is unchecked; 'IPv4 Address' is '10.217.116.102'; 'IPv4 Subnet' is '255.255.255.0'; and 'IPv4 Gateway' is '10.217.116.1'.

Setting	Value
Enable LAN	☒
LAN Interface	eth0
MAC Address	08:35:71:31:D8:6E
Enable IPv4	☒
Enable IPv4 DHCP	☐
IPv4 Address	10.217.116.102
IPv4 Subnet	255.255.255.0
IPv4 Gateway	10.217.116.1

System lockdown mode restriction

When **System Lockdown** mode is enabled, all out-of-band management interfaces (LOM Web UI, IPMI, and Redfish) are restricted to read-only access. As a result, users cannot modify LOM/BMC settings such as **Network Settings** and **User Management** under the **Settings** menu.

Symptoms of lockdown mode

If the system is locked, you encounter the following issues within the GUI:

- **Error Messages:** Insufficient privilege (code 8000).
- **Notifications:** You do not have permission to view this content.
- **GUI Behavior:** Settings menu items appear grayed out or are completely inaccessible.

Affected functions

The following operations are restricted during lockdown:

- Network configuration changes.
- User management and password operations.
- LOM firmware updates.
- Field Replaceable Unit (FRU) data modifications.

Resolution

System Lockdown must be disabled from the NetScaler host OS using `ipmitool` commands. For information about how to disable lockdown, see `.(#system-lockdown-management)` section in the CLI section for lockdown management procedures.

Configure LOM by using the CLI (IPMI tool)

If the GUI is inaccessible or **System Lockdown** is active, use the NetScaler host shell to run the `ipmitool` commands.

Access requirements

To run command-line instructions, connect to the appliance by using one of the following methods:

- Serial console connection
- SSH session (if a network is configured)
- Direct access or monitor access

NOTE:

All commands must run from the NetScaler shell environment. If you are at the NetScaler CLI prompt, type `shell` to enter the FreeBSD environment.

System lockdown management

System Lockdown mode enforces read-only access across all out-of-band management interfaces. This security feature prevents unauthorized configuration changes via LOM, IPMI, and Redfish.

- **Check Status:** `ipmitool raw 0x34 0x81 (00 =Disabled (Configuration changes allowed), 01 =Enabled (Read-only mode active)).`
- **Disable Lockdown:** `ipmitool raw 0x34 0x81 0x0 .`

This command enables full read-write access to all BMC configuration interfaces.

- **Enable Lockdown:** `ipmitool raw 0x34 0x81 0x1 .`

This command restricts all BMC interfaces to read-only access for enhanced security.

Static IP configuration

Perform the following procedure to manually configure the interface:

1. **Set the IP address source to static:** Change the interface mode to disable DHCP and prepare the interface for static IP assignment using the following command:

```
ipmitool lan set 1 ipsrc static
```

2. **Assign the IP address:** Specify the static IP address for the LOM interface using the following command:

```
ipmitool lan set 1 ipaddr <LOM_IP_Address>
```

Example: `ipmitool lan set 1 ipaddr 192.0.2.255`

3. **Assign the Subnet mask:** Define the subnet mask for the LOM network segment using the following command:

```
ipmitool lan set 1 netmask <Subnet_Mask>
```

Example: `ipmitool lan set 1 netmask 255.255.255.0`

4. **Configure the default gateway:** Specify the gateway address for routing traffic outside the local segment using the following command:

```
ipmitool lan set 1 defgw ipaddr <Gateway_IP>
```

Example: `ipmitool lan set 1 defgw ipaddr 192.0.2.255`

5. Reset or reboot the BMC to apply changes:

```
ipmitool bmc reset cold
```

The BMC requires approximately 120 seconds to fully initialize after a reset.

Verify configuration

After the 120-second initialization period, verify the new network configuration.

Display current LOM network settings

Run the following command to display all current network parameters, including the IP address, subnet mask, gateway, and MAC address.

```
ipmitool lan print 1
```

Test network connectivity

From an external management workstation, test connectivity to the new LOM IP address:

```
ping <new_LOM_IP_address>
```

Access the LOM web interface to confirm full functionality:

```
https://<new_LOM_IP_address>
```

User password management

Change the default LOM password from the NetScaler host environment for enhanced security.

1. List current users

Run the following command to display all configured LOM user accounts with their user IDs:

```
ipmitool user list 1
```

Example output:

ID	NAME	CALLIN	LINK AUTH	IPMI MSG	CHANNEL PRIV LIMIT
1	nsroot	false	false	false	NO ACCESS
2		false	false	true	ADMINISTRATOR
3		true	false	false	NO ACCESS

NOTE:

The default `nsroot` account has a User ID of `2`.

2. Change user password

Run the following command to change the password for the `nsroot` account (User ID `2`):

```
ipmitool user set password <user_ID> <new_LOM_password>
```

Where:

- `<user_ID>` is the appropriate user ID (typically 2 for nsroot account).
- `<new_LOM_password>` is the new password.

Example:

```
ipmitool user set password 2
```

You are prompted to enter the new password interactively.

Alternative (Non-Interactive):

```
ipmitool user set password 2 NewSecurePassword123
```

NOTE:

The system prompts for the new password interactively. Enter the password when prompted.

LOM factory reset

Restore the BMC to factory default settings. If the **Preserve Configuration (IPMI/network)** option is not enabled in the **Maintenance** → **Preserve configuration** section of the LOM GUI, this operation resets all settings - including LOM network configuration, user accounts, and IPMI web settings—to LOM factory defaults.

Run the following command to reset the LOM to factory default.

```
ipmitool raw 0x32 0x66
```

Reset effect:

- All settings revert to defaults (192.168.1.3/24).
- All passwords reset to the device serial number.
- System Lockdown mode resets to enable, by default.
- SSL certificates reset to self-signed defaults.
- All user configurations and custom settings are erased and restored back to factory defaults.

Post-reset actions:

The BMC automatically restarts after a factory reset. Wait 120 seconds for full initialization, then reconfigure the LOM interface using either GUI or command-line methods.

Remote system power control

The LOM interface provides comprehensive remote power management capabilities for system startup, shutdown, and reset operations.

Power on procedure

Run the following steps to remotely start the system.

1. Pre-startup verification

Before initiating power-on, verify the following conditions:

- Power supplies are connected and operational.
- Network infrastructure is available and stable.
- Environmental conditions are within an acceptable range (temperature, airflow).
- No hardware alerts or critical warnings are active.

2. Access power control

Navigate to the power control interface, **LOM Web UI > Dashboard > Power Control**.

3. Run power on

- a. Select **Power On** from the available power control actions.
- b. Review the confirmation dialog.
- c. Click **Perform Action** to initiate the startup sequence.

4. Monitor startup process

Observe the following indicators during system boot:

- The power LED changes from off to solid green.
- The system completes POST (Power-On Self-Test).
- Network interfaces initialize and link up.
- The operating system loads and becomes accessible.

NOTE:

Use the iKVM console feature to monitor boot progress in real time, if needed.

Graceful shutdown procedure

Perform a graceful operating system shutdown before removing power.

1. Pre-Shutdown verification

Confirm the following conditions before initiating a shutdown:

- No critical operations are in progress.
- All active user sessions are canceled.
- Configuration changes are saved.
- Backup operations are complete.

2. Access power control

Navigate to the power control interface, **LOM Web UI > Dashboard > Power Control**.

3. Initiate ACPI shutdown

- a. Select **ACPI Shutdown** from the power control actions.
- b. Review the confirmation dialog.
- c. Click **Perform Action** to initiate the graceful shutdown.

The **ACPI Shutdown** option sends a graceful shutdown signal to the operating system, allowing all services to exit.

4. Monitor shutdown process

Observe the following indicators during system shutdown:

- The operating system begins the shutdown sequence.
- Network interfaces go offline.
- System services exit gracefully.
- The power LED changes from green to off.

NOTE:

Allow adequate time for a complete shutdown before initiating any power-on operations.

The screenshot displays the NetScaler SDX web interface. On the left, a dark blue sidebar contains the 'net>scaler' logo and a navigation menu with items: Dashboard, Sensor, System Inventory, FRU Information, Logs & Reports, Settings, Remote Control, Image Redirection, and Power Control (highlighted). The main content area has a blue header with a hamburger menu icon. Below the header, the title 'Power Control on Host Server' is shown. A white box titled 'Power Actions' contains a list of actions: Power Off, Power On, Power Cycle, Hard Reset (selected with a blue checkmark), Power Diag (Initiate NMI), and ACPI Shutdown. A blue button labeled 'Perform Action' is located at the bottom right of the 'Power Actions' box. Above the 'Power Actions' box, the status 'Host is currently on' is displayed. A 'Quick Links..' dropdown menu is visible in the sidebar.

Troubleshooting

Common issues and resolution procedures for LOM initial configuration.

Cannot access the default IP address

Issue: Browser cannot reach `https://192.168.1.3`.

Possible causes:

- The workstation is not configured with the correct IP address.
- Network cable is not connected properly.
- Layer 2 switching issue or VLAN mismatch.
- BMC is not fully initialized after system boot.

Resolution steps:

1. **Verify workstation IP configuration:** Use `ipconfig` (Windows) or `ifconfig/ip addr` (Linux/macOS).
2. **Verify physical cable connection:** Ensure the cable is plugged into the dedicated LOM port.
3. **Check link status:** Verify the LED on the LOM port is active.
4. **Wait for initialization:** Allow up to 5 minutes after system boot for the BMC to initialize.
5. **Direct Connection:** Try a direct crossover cable connection if a switched network issue is suspected.

Login fails with default credentials

Issue: Authentication fails when using `nsroot` and the serial number.

Possible causes:

- Incorrect serial number entry (passwords are case-sensitive).
- The password was previously changed from the default.
- The user account is locked due to multiple failed attempts.

Resolution steps:

1. **Verify serial number:** Check the physical label or device documentation.
2. **Check case-sensitivity:** Ensure the serial number is entered exactly as displayed.
3. **Clear browser data:** Clear the browser cache and cookies, then retry.

4. **Account Lockout:** Wait 15 minutes if the account is locked, then retry.
5. **Factory Reset:** Perform a factory reset if the password is lost and cannot be recovered.

BMC not responding after configuration

Issue: LOM interface becomes inaccessible after applying network changes.

Possible causes:

- An invalid IP address or subnet mask is entered.
- The network gateway is unreachable from the new IP address.
- IP address conflicts with another device on the network.
- BMC requires additional time for initialization.

Resolution steps:

1. **Wait for initialization:** Allow a full 120 seconds for BMC network initialization.
2. **Check for conflicts:** Verify that no other device on the network segment is using the same IP address.
3. **Test connectivity:** Attempt to `ping` the new IP address from a management workstation.
4. **Revert to default:** Try connecting to the original default IP address (`192.168.1.3`) if the new IP address fails.
5. **Factory Reset:** Perform a factory reset if the BMC remains completely inaccessible.

Best practices for LOM configuration and management

Security best practices

Implement the following security measures immediately after the initial LOM configuration to ensure the integrity of the management environment.

Password policy

- **Change default passwords:** Update the default credentials immediately after the first login.
- **Complexity requirements:** Use strong passwords with a minimum of 8 characters.
- **Character variety:** Include a mix of uppercase letters, lowercase letters, numbers, and special characters.
- **Rotation:** Rotate passwords regularly according to your organizational security policy.

- **Uniqueness:** Avoid using common words or device-specific information (such as serial numbers) in passwords.

System lockdown

- **Production readiness:** Enable **System Lockdown** mode in all production environments.
- **Maintenance window:** Disable lockdown only when configuration changes are explicitly required.
- **Re-enablement:** Re-enable lockdown immediately after completing maintenance tasks.

Access logging

- **Audit logging:** Enable audit logging for all administrative actions to maintain a clear history of changes.
- **Regular reviews:** Regularly review access logs for any signs of unauthorized attempts.
- **Alerting:** Implement alerting mechanisms for failed authentication attempts.

SSL/TLS configuration

- **Certificate replacement:** Replace the default LOM self-signed SSL certificate with certificates from a trusted **multi-tier CA root chain**.
- **Chain validation:** Implement proper certificate chain validation.
- **Protocol version:** Use minimum TLS 1.2 protocol.
- **Monitoring:** Configure appropriate monitoring for certificate expiration.

Power control best practices

Follow these guidelines for safe and effective power management:

- **Prioritize graceful shutdowns:** Use graceful shutdown procedures whenever possible to prevent data corruption and service disruption.
- **Observe power cycle delays:** Wait a minimum of 30 seconds between power-off and power-on operations.
- **Maintain documentation:** Document all power control actions in change management or maintenance logs.
- **Schedule operations:** Coordinate all power operations with established maintenance Windows to minimize impact.
- **Ensure redundant access:** Maintain alternative access methods (such as serial console or iKVM) while power cycling the system.

- **Verify system state:** Always verify the system's operational state both before and after performing power operations.

Hardware health attributes of NetScaler SDX appliances

Operating ranges for NetScaler hardware platforms vary for different attributes.

Note

Use the `stat system -detail` command to display the current values of the attributes.

Health attributes for the SDX 9100 platform

The following table lists the health attributes for the SDX 9100 platform.

HEALTH ATTRIBUTE	RECOMMENDED RANGE
CPU Temp (Celsius)	10–95
PCH Temp (Celsius)	10–85
System Temp (Celsius)	10–80
Peripheral Temp (Celsius)	10–80
CPU_VRMIN Temp (Celsius)	10–95
VRMABCD Temp (Celsius)	10–95
VRMEFGH Temp (Celsius)	10–95
DIMMABCD Temp (Celsius)	10–80
DIMMEFGH Temp (Celsius)	10–80
System Fan 1 Speed (RPM)	1300–18000
System Fan 2 Speed (RPM)	1300–18000
System Fan 3 Speed (RPM)	1300–18000

HEALTH ATTRIBUTE	RECOMMENDED RANGE
System Fan 4 Speed (RPM)	1300–18000
System Fan 5 Speed (RPM)	1300–18000
System Fan 6 Speed (RPM)	1300–18000
Standby 3.3 V Supply (Volts)	2.885–3.69
+5.0 V Supply (Volts)	4.466–5.737
+12.0 V Supply (Volts)	10.704–13.476
VBAT (Volts)	2.591–3.464
3.3VSB (Volts)	2.954–3.552
5VSB (Volts)	4.496–5.378
1.8 V PCH (Volts)	1.618–1.943
PVNN PCH (Volts)	0.892–1.072
1.05 V PCH (Volts)	0.915–1.176
1.0 V BMC (Volts)	0.892–1.072
1.2 V BMC (Volts)	1.074–1.29
1.8 V BMC (Volts)	1.62–1.944
2.5 V BMC (Volts)	2.177–2.804
VcpuVRM (Volts)	1.616–1.944
vCPU (Volts)	0.931–1.195
VDimmABCD (Volts)	1.096–1.344
VDimmEFGH (Volts)	1.096–1.344

Health attributes for the SDX 8900 platform

The following table lists the health attributes for the SDX 8900 platform.

HEALTH ATTRIBUTE	RECOMMENDED RANGE
CPU 0 core (Volts)	1.77–1.81
Main 3.3 V Supply (Volts)	3.25–3.32
Standby 3.3 V Supply (Volts)	3.25–3.32
+5.0 V Supply (Volts)	4.84–4.95
+12.0 V Supply (Volts)	11.75–12.00
Battery (Volts)	3.02–3.13
5 V Standby (Volts)	5.05–5.16
CPU Fan 0 Speed (RPM)	6600–12600
CPU Fan 1 Speed (RPM)	6300–12200
System Fan Speed (RPM)	6700–8600
System Fan 1 Speed (RPM)	6600–8500
System Fan 2 Speed (RPM)	6700–8500
CPU 0 Temperature (Celsius)	40–57
CPU 1 Temperature (Celsius)	35–45
Internal Temperature (Celsius)	30–49
Power Supply 1 Status	NORMAL
Power Supply 2 Status	NOT PRESENT–NORMAL

Health attributes for the SDX 15000 and SDX 26000 platforms

The following table lists the health attributes for the SDX 15000 and SDX 26000 platforms.

HEALTH ATTRIBUTE	UNIT	MINIMUM VALUE	NOMINAL VALUE	MAXIMUM VALUE
CPU 0 Core Voltage	Volt	1.47	1.82	1.835
3_3VCC	Volt	3.12	3.30	3.48
3_3VSB	Volt	3.12	3.30	3.48
5VCC	Volt	4.72	5.00	5.28
p12V	Volt	11.33	12.00	12.67
VBAT	Volt	2.75	3.00	3.200
Vtt	-NA-	-NA-	-NA-	-NA-
5VSB	Volt	4.72	5.00	5.28
Voltage Sensor2	-NA-	-NA-	-NA-	-NA-
CPU Fan 0 Speed	RPM	1500	-NA-	7500
CPU Fan 1 Speed	RPM	1500	-NA-	7500
System Fan Speed	RPM	1500	-NA-	7500
System Fan 1 Speed	RPM	1500	-NA-	7500
System Fan 2 Speed	RPM	1500	-NA-	7500
Fan 0 Speed	RPM	1500	-NA-	7500
CPU 0 Temp	Celsius	0	-NA-	85
CPU 1 Temp	Celsius	0	-NA-	85

Internal Temp	Celsius	0	-NA-	80
---------------	---------	---	------	----

Health attributes for the SDX 14000 platform

The following table lists the health attributes for the SDX 14000 platform.

HEALTH ATTRIBUTE	RECOMMENDED RANGE
CPU 0 core (Volts)	0.85–0.91
CPU 1 core Volt	0.85–0.93
Main 3.3 V Supply (Volts)	3.31–3.36
Standby 3.3 V Supply (Volts)	3.26–3.31
+5.0 V Supply (Volts)	4.99–5.06
+12.0 V Supply (Volts)	11.98–12.08
Battery Voltage	3.02–3.12
Internal CPU Vtt Power	0.99–1.01
CPU Fan 0 Speed (RPM)	1875–4350
CPU Fan 1 Speed (RPM)	1875–4200
System Fan Speed (RPM)	1875–4350
System Fan 1 Speed (RPM)	1875–4275
System Fan 2 Speed (RPM)	1875–4200
CPU 0 Temperature (Celsius)	33–43
CPU 1 Temperature (Celsius)	35–45
Internal Temperature (Celsius)	28–38
Power Supply 1 Status	NORMAL
Power Supply 2 Status	NORMAL

Health attributes for the SDX 22040 platform

The following table lists the health attributes for the SDX 22040 platform.

HEALTH ATTRIBUTE	UNIT	LOWER NON-RECOVERABLE	LOWER CRITICAL	LOWER NON-CRITICAL	UPPER NON-CRITICAL	UPPER CRITICAL	UPPER NON-RECOVERABLE
CPU1 Temp	degrees C	0.000	0.000	0.000	90.000	93.000	95.000
CPU2 Temp	degrees C	0.000	0.000	0.000	90.000	93.000	95.000
System Temp	degrees C	-9.000	-7.000	-5.000	80.000	85.000	90.000
Peripheral Temp	degrees C	-9.000	-7.000	-5.000	80.000	85.000	90.000
PCH Temp	degrees C	-11.000	-8.000	-5.000	90.000	95.000	100.000
FPC_Temp 1	degrees C	• NA -	• NA -	• NA -	66.000	70.000	75.000
FPC_Temp 2	degrees C	• NA -	• NA -	• NA -	72.000	76.000	82.000
FPC_Temp 3	degrees C	• NA -	• NA -	• NA -	72.000	76.000	82.000
HDDBP_Temp 1	degrees C	• NA -	• NA -	• NA -	72.000	76.000	82.000

HEALTH ATTRIBUTE	UNIT	LOWER NON-RECOVERABLE	LOWER CRITICAL	LOWER NON-CRITICAL	UPPER NON-CRITICAL	UPPER CRITICAL	UPPER NON-RECOVERABLE
HDDBP_Temp 2	degrees C	• NA -	• NA -	• NA -	72.000	76.000	82.000
FAN 1	RPM	• NA -	1980.000	• NA -	• NA -	• NA -	• NA -
FAN 2	RPM	• NA -	1980.000	• NA -	• NA -	• NA -	• NA -
FAN 3	RPM	• NA -	1980.000	• NA -	• NA -	• NA -	• NA -
FAN 4	RPM	• NA -	1980.000	• NA -	• NA -	• NA -	• NA -
FAN 5	RPM	• NA -	1980.000	• NA -	• NA -	• NA -	• NA -
FAN 6	RPM	• NA -	1980.000	• NA -	• NA -	• NA -	• NA -
FAN 7	RPM	• NA -	1980.000	• NA -	• NA -	• NA -	• NA -
FAN 8	RPM	• NA -	1980.000	• NA -	• NA -	• NA -	• NA -

HEALTH ATTRIBUTE	UNIT	LOWER NON-RECOVERABLE	LOWER CRITICAL	LOWER NON-CRITICAL	UPPER NON-CRITICAL	UPPER CRITICAL	UPPER NON-RECOVERABLE
PS_1 Status	discrete	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
PS_1 FAN	RPM	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
PS_1 Temp	degrees C	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
PS_2 Status	discrete	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
PS_2 FAN	RPM	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
PS_2 Temp	degrees C	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
PS_3 Status	discrete	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
PS_3 FAN	RPM	• NA -	1980.00 0	• NA -	• NA -	• NA -	• NA -
PS_3 Temp	degrees C	• NA -	• NA -	• NA -	72.000	76.000	82.000

HEALTH ATTRIBUTE	UNIT	LOWER NON-RECOVERABLE	LOWER CRITICAL	LOWER NON-CRITICAL	UPPER NON-CRITICAL	UPPER CRITICAL	UPPER NON-RECOVERABLE
PS_4 Status	discrete	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
PS_4 FAN	RPM	• NA -	1980.00 0	• NA -	• NA -	• NA -	• NA -
PS_4 Temp	degrees C	• NA -	• NA -	• NA -	72.000	76.000	82.000
FPC Status	discrete	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
VTT	Volts	0.816	0.864	0.912	1.344	1.392	1.440
CPU1 Vcore	Volts	0.480	0.512	0.544	1.488	1.520	1.552
CPU2 Vcore	Volts	0.480	0.512	0.544	1.488	1.520	1.552
VDIMM AB	Volts	1.104	1.152	1.200	1.648	1.696	1.744
VDIMM CD	Volts	1.104	1.152	1.200	1.648	1.696	1.744
VDIMM EF	Volts	1.104	1.152	1.200	1.648	1.696	1.744
VDIMM GH	Volts	1.104	1.152	1.200	1.648	1.696	1.744
+1.5 V	Volts	1.248	1.296	1.344	1.648	1.696	1.744

HEALTH ATTRIBUTE	UNIT	LOWER NON-RECOVERABLE	LOWER CRITICAL	LOWER NON-CRITICAL	UPPER NON-CRITICAL	UPPER CRITICAL	UPPER NON-RECOVERABLE
3.3 V	Volts	2.640	2.784	2.928	3.648	3.792	3.936
+3.3VSB	Volts	2.640	2.784	2.928	3.648	3.792	3.936
5 V	Volts	4.096	4.288	4.480	5.504	5.696	6.912
12 V	Volts	10.176	10.494	10.812	13.250	13.568	13.886
VBAT	Volts	2.400	2.544	2.688	3.312	3.456	3.600

Health attributes for SDX 24100 platform

The following table lists the health attributes for SDX 24100/24150.

HEALTH ATTRIBUTE	UNIT	LOWER NON-RECOVERABLE	LOWER CRITICAL	LOWER NON-CRITICAL	UPPER NON-CRITICAL	UPPER CRITICAL	UPPER NON-RECOVERABLE
CPU1 Temp	degrees C	0.000	0.000	0.000	90.000	93.000	95.000
CPU2 Temp	degrees C	0.000	0.000	0.000	90.000	93.000	95.000
System Temp	degrees C	-9.000	-7.000	-5.000	80.000	85.000	90.000
Peripheral Temp	degrees C	-9.000	-7.000	-5.000	80.000	85.000	90.000
PCH Temp	degrees C	-11.000	-8.000	-5.000	90.000	95.000	100.000

HEALTH ATTRIBUTE	UNIT	LOWER NON-RECOVERABLE	LOWER CRITICAL	LOWER NON-CRITICAL	UPPER NON-CRITICAL	UPPER CRITICAL	UPPER NON-RECOVERABLE
FPC_Temp 1	degrees C	• NA -	• NA -	• NA -	66.000	70.000	75.000
FPC_Temp 2	degrees C	• NA -	• NA -	• NA -	72.000	76.000	82.000
FPC_Temp 3	degrees C	• NA -	• NA -	• NA -	72.000	76.000	82.000
HDDBP_Temp 1	degrees C	• NA -	• NA -	• NA -	72.000	76.000	82.000
HDDBP_Temp 2	degrees C	• NA -	• NA -	• NA -	72.000	76.000	82.000
FAN 1	RPM	• NA -	1980.000	• NA -	• NA -	• NA -	• NA -
FAN 2	RPM	• NA -	1980.000	• NA -	• NA -	• NA -	• NA -
FAN 3	RPM	• NA -	1980.000	• NA -	• NA -	• NA -	• NA -
FAN 4	RPM	• NA -	1980.000	• NA -	• NA -	• NA -	• NA -

HEALTH ATTRIBUTE	UNIT	LOWER NON- RECOVERABLE	LOWER CRITICAL	LOWER NON- CRITICAL	UPPER NON- CRITICAL	UPPER CRITICAL	UPPER NON- RECOVERABLE
FAN 5	RPM	• NA -	1980.00 0	• NA -	• NA -	• NA -	• NA -
FAN 6	RPM	• NA -	1980.00 0	• NA -	• NA -	• NA -	• NA -
FAN 7	RPM	• NA -	1980.00 0	• NA -	• NA -	• NA -	• NA -
FAN 8	RPM	• NA -	1980.00 0	• NA -	• NA -	• NA -	• NA -
PS_1 Status	discrete	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
PS_1 FAN	RPM	• NA -	1980.00 0	• NA -	• NA -	• NA -	• NA -
PS_1 Temp	degrees C	• NA -	• NA -	• NA -	72.000	76.000	82.000
PS_2 Status	discrete	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
PS_2 FAN	RPM	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -

HEALTH ATTRIBUTE	UNIT	LOWER NON-RECOVERABLE	LOWER CRITICAL	LOWER NON-CRITICAL	UPPER NON-CRITICAL	UPPER CRITICAL	UPPER NON-RECOVERABLE
PS_2 Temp	degrees C	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
PS_3 Status	discrete	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
PS_3 FAN	RPM	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
PS_3 Temp	degrees C	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
PS_4 Status	discrete	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
PS_4 FAN	RPM	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
PS_4 Temp	degrees C	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
FPC Status	discrete	• NA -	• NA -	• NA -	• NA -	• NA -	• NA -
VTT	Volts	0.816	0.864	0.912	1.344	1.392	1.440
CPU1 Vcore	Volts	0.480	0.512	0.544	1.488	1.520	1.552

HEALTH ATTRIBUTE	UNIT	LOWER NON- RECOVERABLE	LOWER CRITICAL	LOWER NON- CRITICAL	UPPER NON- CRITICAL	UPPER CRITICAL	UPPER NON- RECOVERABLE
CPU2 Vcore	Volts	0.480	0.512	0.544	1.488	1.520	1.552
VDIMM AB	Volts	1.104	1.152	1.200	1.648	1.696	1.744
VDIMM CD	Volts	1.104	1.152	1.200	1.648	1.696	1.744
VDIMM EF	Volts	1.104	1.152	1.200	1.648	1.696	1.744
VDIMM GH	Volts	1.104	1.152	1.200	1.648	1.696	1.744
+1.5 V	Volts	1.248	1.296	1.344	1.648	1.696	1.744
3.3 V	Volts	2.640	2.784	2.928	3.648	3.792	3.936
+3.3VSB	Volts	2.640	2.784	2.928	3.648	3.792	3.936
5 V	Volts	4.096	4.288	4.480	5.504	5.696	6.912
12 V	Volts	10.176	10.494	10.812	13.250	13.568	13.886
VBAT	Volts	2.400	2.544	2.688	3.312	3.456	3.600

Hardware health attributes for the SDX 16000 platform

The following table lists the recommended ranges for the health attributes on the SDX 16000 platform.

Run the `ipmitool sdr list` command from the shell. This command lists the sensory parameters data in detail.

Run the `stat system -detail` command from the CLI. This command lists a subset of the parameters.

HEALTH ATTRIBUTE (SHELL)	HEALTH ATTRIBUTE (CLI)	RECOMMENDED RANGE
CPU1 Temp	CPU 0 Temperature (Celsius)	10–95
CPU2 Temp	CPU 1 Temperature (Celsius)	10–95
PCH Temp	Temperature 1 (Celsius)	10–85
System Temp	Internal Temperature (Celsius)	10–80
Peripheral Temp	Temperature 0 (Celsius)	10–80
Fan 1	CPU Fan 0 Speed (RPM)	1260–8260
Fan 2	CPU Fan 1 Speed (RPM)	1260–8260
Fan 3	System Fan Speed (RPM)	1260–8260
Fan 4	System Fan 1 Speed (RPM)	1260–8260
Fan 5	System Fan 2 Speed (RPM)	1260–8260
Fan 6	Fan 0 Speed (RPM)	1260–8260
12V	+12.0 V Supply Voltage	11.330–12.520

HEALTH ATTRIBUTE (SHELL)	HEALTH ATTRIBUTE (CLI)	RECOMMENDED RANGE
5VCC	+5.0 V Supply Voltage	4.713–5.217
3.3VCC	Main 3.3 V Supply Voltage	3.134–3.464
VBAT	Battery Voltage (Volts)	2.591–3.464
Vcpu1VCCIN	CPU 0 Core Voltage (Volts)	1.616–2.016
Vcpu2VCCIN	CPU 1 Core Voltage (Volts)	1.616–2.016
Vcpu1VDDQABCD	Voltage 0 (Volts)	1.096–1.344
Vcpu1VDDQEFGH	Voltage 1 (Volts)	1.096–1.344
Vcpu2VDDQABCD	Voltage 2 (Volts)	1.096–1.344
Vcpu2VDDQEFGH	Voltage 3 (Volts)	1.096–1.344
5VSB	5V Standby Voltage (Volts)	4.745–5.249
3.3VSB	Standby 3.3 V Supply Voltage	3.117–3.447
1.8V PCH	Voltage 4 (Volts)	1.698–1.883
PS1 Status	Power supply 1 status	not applicable -
PS2 Status	Power supply 2 status	not applicable -

Sample output

The following sections list the different ways to view the system sensor health attributes. The data is a sample representation only. The output from your appliance might be different. For the acceptable ranges, see the previous table.

- From the CLI
- From the shell
- From the LOM GUI

CLI

At the command line, type: `stat system -detail`

```
>stat system -detail
```

NetScaler Executive View

System Information:

Up since Tue Nov 15 13:50:59 2022

Up since(Local) Tue Nov 15 13:50:59 2022

Memory usage (MB) 3562

InUse Memory (%) 6.24

Number of CPUs 9

System Health Statistics (Standard):

CPU 0 Core Voltage (Volts) 1.77

CPU 1 Core Voltage (Volts) 1.77

Main 3.3 V Supply Voltage 3.49

Standby 3.3 V Supply Voltage 3.45

+5.0 V Supply Voltage 5.26

+12.0 V Supply Voltage 12.44

Battery Voltage (Volts) 3.11

Intel CPU Vtt Power(Volts) 0.00

5V Standby Voltage(Volts)	5.04
Voltage Sensor2(Volts)	0.00
CPU Fan 0 Speed (RPM)	3360
CPU Fan 1 Speed (RPM)	3360
System Fan Speed (RPM)	3360
System Fan 1 Speed (RPM)	3360
System Fan 2 Speed (RPM)	3360
CPU 0 Temperature (Celsius)	54
CPU 1 Temperature (Celsius)	53
Internal Temperature (Celsius)	33
Power supply 1 status	NORMAL
Power supply 2 status	NORMAL
Power supply 3 status	NOT SUPPORTED
Power supply 4 status	NOT SUPPORTED
System Disk Statistics:	
/flash Size (MB)	29748
/flash Used (MB)	3927
/flash Available (MB)	23441
/flash Used (%)	14
/var Size (MB)	800856
/var Used (MB)	20127

/var Available (MB)	716660
---------------------	--------

/var Used (%)	2
---------------	---

System Health Statistics(Auxiliary):

Voltage 0 (Volts)	1.20
-------------------	------

Voltage 1 (Volts)	1.20
-------------------	------

Voltage 2 (Volts)	1.20
-------------------	------

Voltage 3 (Volts)	1.20
-------------------	------

Voltage 4 (Volts)	1.88
-------------------	------

Voltage 5 (Volts)	0.00
-------------------	------

Voltage 6 (Volts)	0.00
-------------------	------

Voltage 7 (Volts)	0.00
-------------------	------

Fan 0 Speed (RPM)	3360
-------------------	------

Fan 1 Speed (RPM)	0
-------------------	---

Fan 2 Speed (RPM)	0
-------------------	---

Fan 3 Speed (RPM)	0
-------------------	---

Temperature 0 (Celsius)	28
-------------------------	----

Temperature 1 (Celsius)	34
-------------------------	----

Temperature 2 (Celsius)	0
-------------------------	---

Temperature 3 (Celsius)	0
-------------------------	---

Done

>

Shell

At the command line, type: `shell`

```
> shell
```

```
root@ns# ipmitool sdr list
```

CPU1 Temp	54 degrees C	ok
CPU2 Temp	53 degrees C	ok
PCH Temp	34 degrees C	ok
System Temp	33 degrees C	ok
Peripheral Temp	28 degrees C	ok
CPU1_VRMIN Temp	65 degrees C	ok
CPU1_VRMIO Temp	44 degrees C	ok
CPU2_VRMIN Temp	46 degrees C	ok
CPU2_VRMIO Temp	47 degrees C	ok
P1_VRMABCD Temp	38 degrees C	ok
P1_VRMEFGH Temp	43 degrees C	ok
P2_VRMABCD Temp	42 degrees C	ok
P2_VRMEFGH Temp	45 degrees C	ok
FAN1	3360 RPM	ok
FAN2	3360 RPM	ok

FAN3	3360 RPM	ok
FAN4	3360 RPM	ok
FAN5	3360 RPM	ok
FAN6	3360 RPM	ok
P1_DIMMA~D Temp	38 degrees C	ok
P1_DIMME~H Temp	33 degrees C	ok
P2_DIMMA~D Temp	30 degrees C	ok
P2_DIMME~H Temp	30 degrees C	ok
12V	12.44 Volts	ok
5VCC	5.26 Volts	ok
3.3VCC	3.49 Volts	ok
VBAT	3.11 Volts	ok
Vcpu1VCCIN	1.77 Volts	ok
Vcpu1VCCIO	1.00 Volts	ok
Vcpu2VCCIN	1.77 Volts	ok
Vcpu2VCCIO	1.00 Volts	ok
Vcpu1VDDQABCD	1.20 Volts	ok
Vcpu1VDDQEFGH	1.20 Volts	ok
Vcpu2VDDQABCD	1.20 Volts	ok
Vcpu2VDDQEFGH	1.20 Volts	ok
5VSB	5.04 Volts	ok
3.3VSB	3.45 Volts	ok

```
2.5V BMC      | 2.56 Volts    | ok
1.8V BMC      | 1.83 Volts    | ok
1.2V BMC      | 1.20 Volts    | ok
1.0V BMC      | 1.02 Volts    | ok
1.8V PCH      | 1.88 Volts    | ok
PVNN PCH      | 1.06 Volts    | ok
1.05V PCH     | 1.06 Volts    | ok
PS1 Status    | 0x01          | ok
PS2 Status    | 0x01          | ok
AOC_NIC2 Temp | 51 degrees C  | ok
AOC_NIC4 Temp | 51 degrees C  | ok

root@ns#
```

LOM GUI

1. Log on to the LOM GUI.
2. On the Dashboard, click **Sensor** for temperature and voltage values. Click **Fan** for fan values.

Dashboard
System
Configuration
Remote Control
Maintenance

Overview
CPU
Memory
PSU
Power
Sensor
Fan

Filter
Show Thresholds
Select a sensor type category
X Voltage X Temperature

Sensor Readings
Export to Excel

Severity	Name	Reading	Type	Low NR	Low CT	High CT	High NR
OK	CPU2 Temp	55	Temperature	5	5	100	100
OK	CPU1 Temp	55	Temperature	5	5	100	100
OK	PCH Temp	33	Temperature	5	5	90	105
OK	System Temp	33	Temperature	5	5	85	90
OK	Peripheral Temp	29	Temperature	5	5	85	90
OK	CPU1_VRMIN Temp	66	Temperature	5	5	100	105
OK	CPU1_VRMI0 Temp	44	Temperature	5	5	100	105
OK	CPU2_VRMIN Temp	46	Temperature	5	5	100	105
OK	CPU2_VRMI0 Temp	47	Temperature	5	5	100	105
OK	P1_VRMABCD Temp	38	Temperature	5	5	100	105
OK	P1_VRMCFGH Temp	43	Temperature	5	5	100	105
OK	P2_VRMABCD Temp	42	Temperature	5	5	100	105
OK	P2_VRMCFGH Temp	45	Temperature	5	5	100	105
OK	P1_DIMMA-D Temp	38	Temperature	5	5	85	90
OK	P1_DIMME-H Temp	33	Temperature	5	5	85	90
OK	P2_DIMMA-D Temp	30	Temperature	5	5	85	90
OK	P2_DIMME-H Temp	30	Temperature	5	5	85	90
OK	AOC_NIC2 Temp	51	Temperature	5	5	100	105
OK	AOC_NIC4 Temp	51	Temperature	5	5	100	105
OK	12V	12.435	Voltage	10.735	10.735	13.2	13.2
OK	5VCC	5.259	Voltage	4.481	4.461	5.469	5.469
OK	3.3VCC	3.464	Voltage	2.968	2.968	3.629	3.629

Dashboard
System
Configuration
Remote Control
Maintenance

Overview
CPU
Memory
PSU
Power
Sensor
Fan

Fan

Operating State	Health Status	Name	RPM	Low NR	Low CT	High CT	High NR
OK	OK	FAN1	3360	980	1120	8680	8960
OK	OK	FAN2	3360	980	1120	8680	8960
OK	OK	FAN3	3360	980	1120	8680	8960
OK	OK	FAN4	3360	980	1120	8680	8960
OK	OK	FAN5	3360	980	1120	8680	8960
OK	OK	FAN6	3360	980	1120	8680	8960

Fortville NIC firmware upgrade on NetScaler SDX appliances

Upgrade the firmware on NetScaler SDX appliances containing Fortville NICs to remediate [<https://support.citrix.com/article/CTX263807>](https://support.citrix.com/article/CTX263807) .

Supported platforms

Important: Before upgrading the single bundle image (SBI), first upgrade the VPX instance to one of the following releases:

- Release 11.1 build 65.x and later
- Release 12.1 build 57.x and later
- Release 13.0 build 58.x and later

The firmware and NIC driver are automatically upgraded to version 7.0 when the SDX SBI is upgraded to the following minimum versions. The time taken to complete the upgrade varies by platform. For example, on an SDX 14000 platform with 6 Fortville NICs, the total time to upgrade is around 30 min.

SOFTWARE VERSION	SBI VERSION
11.1 build 65.x	11.1 build 65.x
12.1 build 57.x	12.1 build 57.x
13.0 build 58.x	13.0 build 58.x

The following SDX appliances and corresponding zero-capacity appliances contain Fortville NICs:

- SDX 8900
- SDX 14000-40S
- SDX 14000-40G
- SDX 15000
- SDX 15000-50G

- SDX 25000-40G
- SDX 26000
- SDX 26000-50S

Support for updated NIC driver and firmware on NetScaler SDX with 10G/25G/40G NICs

When you upgrade to Single Bundle Image (SBI) version 14.1-8.x and later or 13.1-50.x and later, the 10G/25G/40G NIC driver and firmware are automatically upgraded to version 8.70 on the following platforms. NIC firmware version 8.70 remediates CVE-2020-8690, CVE-2020-8691, CVE-2020-8692, and CVE-2020-8693.

- SDX 8900
- SDX 14000-40G
- SDX 15000
- SDX 15000-50G
- SDX 25000-40G
- SDX 16000
- SDX 9100
- SDX 26000
- SDX 26000-50S

Converting a NetScaler MPX appliance to a NetScaler SDX appliance

You can convert a NetScaler MPX appliance to an SDX appliance to deploy multiple virtualized ADC instances on a single, purpose-built physical appliance with full multiservice and multitenant support.

Note: Conversion of an SDX appliance to an MPX appliance is not supported.

You can convert the NetScaler MPX 11515/11520/11530/11540/11542 appliances to NetScaler SDX 11515/11520/11530/11540/11542 appliances. To convert, upgrade the software through a new Solid State Drive (SSD) and a new Hard Disk Drive (HDD).

The NetScaler models SDX 11515/11520/11530/11540/11542 are 2U appliances. Each model has two 6-core processors for a total of 12 physical cores (24 cores with hyper-threading), and 48 GB of memory.

The SDX 11515/11520/11530/11540/11542 appliances have the following ports:

- RS232 serial console port.
- 10/100Base-T copper Ethernet Port (RJ45), also called the LOM port. You can use this port to remotely monitor and manage the appliance independently of the NetScaler software.

NOTE:

On the SDX 11515/11520/11530/11540/11542 appliances, the LEDs on the LOM port are not operational, by design.

- Two 10/100/1000Base-T copper Ethernet management ports (RJ45), numbered 0/1 and 0/2 from left to right. These ports are used to connect directly to the appliance for system administration functions. Eight 10G SFP+ ports and four copper or fiber 1G SFP ports.

You can convert the NetScaler MPX 8005/8010/8015/8200/8400/8600/8800 appliances to NetScaler SDX 8010/8015/8400/8600 appliances by upgrading the software through a new Solid State Drive (SSD).

The NetScaler models SDX 8010/8015/8400/8600 are 1U appliances. Each model has one quad-core processor (8 cores with hyper-threading) and 32 GB of memory. The SDX 8010/8015/8400/8600 appliances are available in two port configurations:

- Six 10/100/1000Base-T copper Ethernet ports and six 1G SFP ports (6x10/100/1000Base-T copper Ethernet ports + 6x1G SFP)
- Six 10/100/1000Base-T copper Ethernet ports and two 10G SFP+ ports (6x10/100/1000Base-T copper Ethernet ports + 2x10G SFP+)

Convert a NetScaler MPX 26000-50S appliance to a NetScaler SDX 26000-50S appliance

You can convert a NetScaler MPX appliance to an NetScaler SDX appliance by upgrading the software through a new solid-state drive (SSD). NetScaler supplies a field conversion kit to migrate a NetScaler MPX appliance to an SDX appliance.

The conversion requires all eight SSDs.

NOTE:

Citrix recommends that you configure the lights out management (LOM) port of the appliance before starting the conversion process. For more information on the LOM port of the ADC appliance, see [Lights out management port of the NetScaler SDX appliance](#).

To convert an MPX appliance to an SDX appliance, you must access the appliance through a console cable attached to a computer or terminal. Before connecting the console cable, configure the computer or terminal to support the following configuration:

- VT100 terminal emulation
- 9600 baud
- 8 data bits
- 1 stop bit
- Parity and flow control set to NONE

Connect one end of the console cable to the RS232 serial port on the appliance, and the other end to the computer or terminal.

NOTE:

To use a cable with an RJ-45 converter, insert the optional converter into the console port and attach the cable to it.

With the cable attached, verify that the MPX appliance's components are functioning correctly. You are then ready to begin the conversion. The conversion process modifies the BIOS, installs a Citrix Hypervisor and a Service Virtual Machine image, and copies the NetScaler VPX image to the Solid State Drive.

The conversion process also sets up a redundant array of independent disks (RAID) controller for local storage and NetScaler VPX storage. SSD slots #1 and #2 are used for local storage and SSD slots #3, #4, #5, #6, #7, and #8 are used for NetScaler VPX storage.

After the conversion process, modify the appliance's configuration and apply a new license. You can then provision the VPX instances through the Management Service on what is now an NetScaler SDX appliance.

Verify proper operation of the MPX appliance's components

1. Access the console port and enter the administrator credentials.
2. Run the following command from the command line interface of the appliance to display the serial number: `show hardware`. You might need the serial number to log on to the appliance after the conversion.

Example

```
> show hardware
```

```
Platform: NSMPX-26000-50S 28*CPU+256GB+4*MLX(50)+16*F2X+2*E1K+4*4-CHIP  
C0 520310
```

```
Manufactured on: 9/15/2020
```

```
CPU: 2400MHZ
```

```
Host Id: 1862303878
```

```
Serial no: 83WRVD62EM
```

```
Encoded serial no: 83WRVD62EM
```

```
Netscaler UUID: 258f9971-f7db-11ea-b021-00e0ed9041b6
```

```
BMC Revision: 5.03
```

```
Done
```

The serial number might be helpful when you want to contact Citrix® technical support.

3. Run the following command to display the status of the active interfaces:

```
show interface
```

Example

```
> show interface
```

```
1) Interface 0/1 (Gig Ethernet 10/100/1000 Mbits) #8
```

```
flags=0xc020 <ENABLED, UP, UP, autoneg, HAMON, HEARTBEAT, 802.1q>
```

```
MTU=1500, native vlan=1, MAC=0c:c4:7a:56:64:7e, uptime 120h29m33s
```

```
Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,
```

```
throughput 0
```

```
Actual: media UTP, speed 1000, duplex FULL, fctl OFF, throughput 1000
```

```
LLDP Mode: NONE, LR Priority: 1024
```

```
RX: Pkts(1614945) Bytes(104805024) Errs(0) Drops(1601312) Stalls(0)
```

```
TX: Pkts(9166) Bytes(2556380) Errs(0) Drops(0) Stalls(0)
```

```
NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)
```

```
Bandwidth thresholds are not set.
```

2) Interface 0/2 (Gig Ethernet 10/100/1000 MBits) #9

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=0c:c4:7a:56:64:7f, downtime 120h29m53s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE, LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

3) Interface 25/1 (25G Ethernet, unknown device) #10

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:90:41:bc, downtime 120h29m53s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE, LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic

4) Interface 25/2 (25G Ethernet, unknown device) #11

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:90:41:bd, downtime 120h29m53s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE, LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic

5) Interface 25/3 (25G Ethernet, unknown device) #12

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:90:41:be, downtime 120h29m53s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE, LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic

6) Interface 25/4 (25G Ethernet, unknown device) #13

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:90:41:bf, downtime 120h29m53s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE, LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic

7) Interface 25/5 (25G Ethernet, unknown device) #14

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:90:41:f2, downtime 120h29m53s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE, LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic

8) Interface 25/6 (25G Ethernet, unknown device) #15

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:90:41:f3, downtime 120h29m53s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE, LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic

9) Interface 25/7 (25G Ethernet, unknown device) #16

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:90:41:f4, downtime 120h29m53s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE, LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic

10) Interface 25/8 (25G Ethernet, unknown device) #17

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:90:41:f5, downtime 120h30m00s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE, LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic

11) Interface 25/9 (25G Ethernet, unknown device) #2

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:90:41:b4, downtime 120h30m00s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE, LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic

12) Interface 25/10 (25G Ethernet, unknown device) #3

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:90:41:b5, downtime 120h30m00s

```
Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,  
  
          throughput 0
```

```
LLDP Mode: NONE,          LR Priority: 1024
```

```
RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
```

```
TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
```

```
NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)
```

```
Bandwidth thresholds are not set.
```

```
Rx Ring: Configured size=2048, Actual size=512, Type: Elastic
```

```
13) Interface 25/11 (25G Ethernet, unknown device) #0
```

```
flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>
```

```
MTU=1500, native vlan=1, MAC=00:e0:ed:90:41:b6, downtime 120h30m00s
```

```
Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,  
  
          throughput 0
```

```
LLDP Mode: NONE,          LR Priority: 1024
```

```
RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
```

```
TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
```

```
NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)
```

```
Bandwidth thresholds are not set.
```

```
Rx Ring: Configured size=2048, Actual size=512, Type: Elastic
```

```
14) Interface 25/12 (25G Ethernet, unknown device) #1

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:90:41:b7, downtime 120h30m00s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

        throughput 0

LLDP Mode: NONE,                LR Priority: 1024


RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic


15) Interface 25/13 (25G Ethernet, unknown device) #6

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:90:41:ee, downtime 120h30m00s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

        throughput 0

LLDP Mode: NONE,                LR Priority: 1024


RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
```

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic

16) Interface 25/14 (25G Ethernet, unknown device) #7

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:90:41:ef, downtime 120h30m00s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE, LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic

17) Interface 25/15 (25G Ethernet, unknown device) #4

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:90:41:f0, downtime 120h30m00s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE, LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic

18) Interface 25/16 (25G Ethernet, unknown device) #5

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:90:41:f1, downtime 120h30m00s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE, LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic

19) Interface 50/1 (mce Ethernet, unknown device) #20

flags=0xc020 <ENABLED, UP, UP, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=9216, native vlan=1, MAC=24:8a:07:a4:91:0c, uptime 99h59m54s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

Actual: media UTP, speed 50000, duplex FULL, fctl OFF, throughput 50000

LLDP Mode: NONE, LR Priority: 1024

RX: Pkts(215681) Bytes(11431096) Errs(0) Drops(215680) Stalls(0)

TX: Pkts(131) Bytes(5516) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

20) Interface 50/2 (mce Ethernet, unknown device) #21

flags=0xc020 <ENABLED, UP, UP, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=24:8a:07:a4:91:0d, uptime 100h00m22s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

Actual: media UTP, speed 50000, duplex FULL, fctl OFF, throughput 50000

LLDP Mode: NONE, LR Priority: 1024

RX: Pkts(313720) Bytes(17365302) Errs(0) Drops(312232) Stalls(0)

TX: Pkts(114) Bytes(4788) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

21) Interface 50/3 (100G Ethernet) #18

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=24:8a:07:a4:91:00, downtime 120h30m07s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE, LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

22) Interface 50/4 (100G Ethernet) #19

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=24:8a:07:a4:91:01, downtime 120h30m07s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE, LR Priority: 1024

```
RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.


23)      Interface L0/1 (Netscaler Loopback interface) #22

        flags=0x20008020 <ENABLED, UP, UP>

MTU=9216, native vlan=1, MAC=0c:c4:7a:56:64:7e, uptime 120h30m01s

LLDP Mode: NONE,                LR Priority: 1024


RX: Pkts(611167959) Bytes(89143315781) Errs(0) Drops(0) Stalls(0)

TX: Pkts(1207140213) Bytes(120806672260) Errs(0) Drops(0) Stalls(0)

Bandwidth thresholds are not set.


Done

>
```

4. In the `show interface` command's output, verify that all the interfaces are enabled and the status of every interface is shown as UP/UP.

NOTES:

- The interface status is displayed as UP/UP only if the cables are connected to the interfaces.
- If you do not have an SFP+ transceiver for every port, verify the interfaces in stages. After checking the first set of interfaces, unplug the SFP+ transceivers and plug them in to the next set of ports.

5. Run the following command for each of the interfaces that are not in the UP/UP state: `enable interface 50/x` Where x is the new interface number.

```
> enable interface 50/1

Done

> enable interface 50/2

Done

> enable interface 50/3

Done

> enable interface 50/4

Done
```

6. Run the following command to verify that the status of the power supplies is normal:

```
stat system -detail
```

Example

```
> stat system -detail
```

Citrix ADC Executive View

System Information:

Up since Wed Jan 20 21:35:18 2021

Up since(Local) Wed Jan 20 21:35:18 2021

Memory usage (MB)	5802
-------------------	------

InUse Memory (%)	2.52
------------------	------

Number of CPUs	15
----------------	----

System Health Statistics (Standard):

CPU 0 Core Voltage (Volts)	1.80
----------------------------	------

CPU 1 Core Voltage (Volts)	1.80
----------------------------	------

Main 3.3 V Supply Voltage	3.35
---------------------------	------

Standby 3.3 V Supply Voltage	3.25
------------------------------	------

+5.0 V Supply Voltage	5.03
-----------------------	------

+12.0 V Supply Voltage	12.13
------------------------	-------

Battery Voltage (Volts)	2.98
-------------------------	------

Intel CPU Vtt Power(Volts)	0.00
----------------------------	------

5V Standby Voltage(Volts)	4.97
---------------------------	------

Voltage Sensor2(Volts)	0.00
------------------------	------

CPU Fan 0 Speed (RPM)	3100
-----------------------	------

CPU Fan 1 Speed (RPM)	3100
-----------------------	------

System Fan Speed (RPM)	3100
------------------------	------

System Fan 1 Speed (RPM)	3100
--------------------------	------

System Fan 2 Speed (RPM)	3100
CPU 0 Temperature (Celsius)	47
CPU 1 Temperature (Celsius)	51
Internal Temperature (Celsius)	35
Power supply 1 status	NORMAL
Power supply 2 status	NORMAL
Power supply 3 status	NOT SUPPORTED
Power supply 4 status	NOT SUPPORTED

System Disk Statistics:

/flash Size (MB)	23801
/flash Used (MB)	787
/flash Available (MB)	21110
/flash Used (%)	3
/var Size (MB)	341167
/var Used (MB)	112391
/var Available (MB)	201482
/var Used (%)	35

System Health Statistics(Auxiliary):

Voltage 0 (Volts)	1.20
Voltage 1 (Volts)	1.20

```
Voltage 2 (Volts)                1.21
Voltage 3 (Volts)                1.20
Voltage 4 (Volts)                1.54
Voltage 5 (Volts)                0.00
Voltage 6 (Volts)                0.00
Voltage 7 (Volts)                0.00
Fan 0 Speed (RPM)               3200
Fan 1 Speed (RPM)                0
Fan 2 Speed (RPM)                0
Fan 3 Speed (RPM)                0
Temperature 0 (Celsius)          27
Temperature 1 (Celsius)          38
Temperature 2 (Celsius)          0
Temperature 3 (Celsius)          0

Done

>
```

1. Run the following command to generate a tar of system configuration data and statistics:

```
show techsupport
```

Example

```
> show techsupport
```

```
showtechsupport data collector tool - $Revision$!

NetScaler version 13.0

The NS IP of this box is 10.217.206.43

This is not HA configuration

Copying selected configuration files ....

Running shell commands ....

Running CLI show commands ....

Collecting ns running configuration....

Collecting running gslb configuration....

Running CLI stat commands ....

Running vtysh commands ....

Copying newnslog files ....

Copying core files from /var/core ....

Copying core files from /var/crash ....

Copying GSLB location database files ....

Copying GSLB auto sync log files ....

Copying Safenet Gateway log files ....

Copying messages, ns.log, dmesg and other log files ....

Creating archive ....

/var/tmp/support/support.tgz ---- points to ---> /var/tmp/support/
collector_P_10.217.206.43_5Dec2020_05_32.tar.gz
```

```
showtechsupport script took 1 minute(s) and 17 second(s) to execute.
```

```
Done
```

NOTE

The output of the command is available in the `/var/tmp/support/collector_<IP_address>_P_<date>.tar.gz` file. Copy this file to another computer for future reference. The output of the command might be helpful when you want to contact Citrix technical support.

2. At the command line interface, switch to the shell prompt. Type: `shell`

Example

```
> shell
```

```
Copyright (c) 1992-2013 The FreeBSD Project.
```

```
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993,  
1994
```

```
The Regents of the University of California. All rights reserved.
```

```
root@ns#
```

3. Run the following command to verify the memory on your appliance:

```
root@ns# grep "memory" /var/nslog/dmesg.boot
```

Example

```
root@ns# grep "memory" /var/nslog/dmesg.boot
```

```
real memory = 274877906944 (262144 MB)
```

```
avail memory = 264024535040 (251793 MB)
```

```
root@ns#
```

4. Run the following command to verify the number of CPU cores depending upon your appliance:

```
root@ns# grep "cpu" /var/nslog/dmesg.boot
```

Example

```
root@ns# grep "cpu" /var/nslog/dmesg.boot
```

```
cpu0 (BSP): APIC ID: 0
```

```
cpu1 (AP): APIC ID: 2
```

```
cpu2 (AP): APIC ID: 4
```

```
cpu3 (AP): APIC ID: 6
```

```
cpu4 (AP): APIC ID: 8
```

```
cpu5 (AP): APIC ID: 10
```

```
cpu6 (AP): APIC ID: 12
```

```
cpu7 (AP): APIC ID: 16
```

```
cpu8 (AP): APIC ID: 18
```

```
cpu9 (AP): APIC ID: 20
```

```
cpu10 (AP): APIC ID: 22
```

```
cpu11 (AP): APIC ID: 24
```

```
cpu12 (AP): APIC ID: 26
```

```
cpu13 (AP): APIC ID: 28
```

```
cpu14 (AP): APIC ID: 32
```

```
cpu15 (AP): APIC ID: 34
```

```
cpu16 (AP): APIC ID: 36

cpu17 (AP): APIC ID: 38

cpu18 (AP): APIC ID: 40

cpu19 (AP): APIC ID: 42

cpu20 (AP): APIC ID: 44

cpu21 (AP): APIC ID: 48

cpu22 (AP): APIC ID: 50

cpu23 (AP): APIC ID: 52

cpu24 (AP): APIC ID: 54

cpu25 (AP): APIC ID: 56

cpu26 (AP): APIC ID: 58

cpu27 (AP): APIC ID: 60

cpu0: <ACPI CPU> on acpi0

cpu1: <ACPI CPU> on acpi0

cpu2: <ACPI CPU> on acpi0

cpu3: <ACPI CPU> on acpi0

cpu4: <ACPI CPU> on acpi0

cpu5: <ACPI CPU> on acpi0

cpu6: <ACPI CPU> on acpi0

cpu7: <ACPI CPU> on acpi0

cpu8: <ACPI CPU> on acpi0

cpu9: <ACPI CPU> on acpi0
```



```
cpu10: <ACPI CPU> on acpi0
cpu11: <ACPI CPU> on acpi0
cpu12: <ACPI CPU> on acpi0
cpu13: <ACPI CPU> on acpi0
cpu14: <ACPI CPU> on acpi0
cpu15: <ACPI CPU> on acpi0
cpu16: <ACPI CPU> on acpi0
cpu17: <ACPI CPU> on acpi0
cpu18: <ACPI CPU> on acpi0
cpu19: <ACPI CPU> on acpi0
cpu20: <ACPI CPU> on acpi0
cpu21: <ACPI CPU> on acpi0
cpu22: <ACPI CPU> on acpi0
cpu23: <ACPI CPU> on acpi0
cpu24: <ACPI CPU> on acpi0
cpu25: <ACPI CPU> on acpi0
cpu26: <ACPI CPU> on acpi0
cpu27: <ACPI CPU> on acpi0

est0: <Enhanced SpeedStep Frequency Control> on cpu0
p4tcc0: <CPU Frequency Thermal Control> on cpu0

est1: <Enhanced SpeedStep Frequency Control> on cpu1
p4tcc1: <CPU Frequency Thermal Control> on cpu1
```

est2: <Enhanced SpeedStep Frequency Control> on cpu2

p4tcc2: <CPU Frequency Thermal Control> on cpu2

est3: <Enhanced SpeedStep Frequency Control> on cpu3

p4tcc3: <CPU Frequency Thermal Control> on cpu3

est4: <Enhanced SpeedStep Frequency Control> on cpu4

p4tcc4: <CPU Frequency Thermal Control> on cpu4

est5: <Enhanced SpeedStep Frequency Control> on cpu5

p4tcc5: <CPU Frequency Thermal Control> on cpu5

est6: <Enhanced SpeedStep Frequency Control> on cpu6

p4tcc6: <CPU Frequency Thermal Control> on cpu6

est7: <Enhanced SpeedStep Frequency Control> on cpu7

p4tcc7: <CPU Frequency Thermal Control> on cpu7

est8: <Enhanced SpeedStep Frequency Control> on cpu8

p4tcc8: <CPU Frequency Thermal Control> on cpu8

est9: <Enhanced SpeedStep Frequency Control> on cpu9

p4tcc9: <CPU Frequency Thermal Control> on cpu9

est10: <Enhanced SpeedStep Frequency Control> on cpu10

p4tcc10: <CPU Frequency Thermal Control> on cpu10

est11: <Enhanced SpeedStep Frequency Control> on cpu11

p4tcc11: <CPU Frequency Thermal Control> on cpu11

est12: <Enhanced SpeedStep Frequency Control> on cpu12

p4tcc12: <CPU Frequency Thermal Control> on cpu12

est13: <Enhanced SpeedStep Frequency Control> on cpu13

p4tcc13: <CPU Frequency Thermal Control> on cpu13

est14: <Enhanced SpeedStep Frequency Control> on cpu14

p4tcc14: <CPU Frequency Thermal Control> on cpu14

est15: <Enhanced SpeedStep Frequency Control> on cpu15

p4tcc15: <CPU Frequency Thermal Control> on cpu15

est16: <Enhanced SpeedStep Frequency Control> on cpu16

p4tcc16: <CPU Frequency Thermal Control> on cpu16

est17: <Enhanced SpeedStep Frequency Control> on cpu17

p4tcc17: <CPU Frequency Thermal Control> on cpu17

est18: <Enhanced SpeedStep Frequency Control> on cpu18

p4tcc18: <CPU Frequency Thermal Control> on cpu18

est19: <Enhanced SpeedStep Frequency Control> on cpu19

p4tcc19: <CPU Frequency Thermal Control> on cpu19

est20: <Enhanced SpeedStep Frequency Control> on cpu20

p4tcc20: <CPU Frequency Thermal Control> on cpu20

est21: <Enhanced SpeedStep Frequency Control> on cpu21

p4tcc21: <CPU Frequency Thermal Control> on cpu21

est22: <Enhanced SpeedStep Frequency Control> on cpu22

p4tcc22: <CPU Frequency Thermal Control> on cpu22

est23: <Enhanced SpeedStep Frequency Control> on cpu23

p4tcc23: <CPU Frequency Thermal Control> on cpu23

```

est24: <Enhanced SpeedStep Frequency Control> on cpu24

p4tcc24: <CPU Frequency Thermal Control> on cpu24

est25: <Enhanced SpeedStep Frequency Control> on cpu25

p4tcc25: <CPU Frequency Thermal Control> on cpu25

est26: <Enhanced SpeedStep Frequency Control> on cpu26

p4tcc26: <CPU Frequency Thermal Control> on cpu26

est27: <Enhanced SpeedStep Frequency Control> on cpu27

p4tcc27: <CPU Frequency Thermal Control> on cpu27

root@ns#

```

- Run the following command to verify that the /var drive is mounted as `/dev/ ar0s1a:`

```
root@ns# df -h
```

Example

```

root@ns# df -h

Filesystem      Size  Used Avail Capacity  Mounted on
/dev/md0        422M  404M   9.1M    98%    /
devfs           1.0k   1.0k     0B   100%   /dev
procfs          4.0k   4.0k     0B   100%  /proc
/dev/ar0s1a     23G   6.9G   14G    32%   /flash
/dev/ar0s1e    333G   32G  274G    10%   /var

root@ns#

```

- Type the following command to run the `ns_hw_err.bash` script, which checks for latent hardware errors:

```
root@ns# ns_hw_err.bash
```

Example

```
root@ns# ns_hw_err.bash

NetScaler NS13.0: Build 73.17.nc, Date: Nov 11 2020, 07:56:08 (64-bit)

platform: serial 83WRVD62EM

platform: sysid 520310 - NSMPX-26000-50S
28*CPU+256GB+4*MLX(50)+16*F2X+2*E1K+4*4-CHIP COL 8955

HDD MODEL: ar0: 434992MB <Intel MatrixRAID RAID1> status: READY

Generating the list of newslog files to be processed...

Generating the events from newslog files...

Checking for HDD errors...

Checking for HDD SMART errors...

Checking for Flash errors...

Checking for Mega Raid Controller errors...

Checking for SSL errors...

Checking for BIOS errors...

Checking for SMB errors...

Checking for MotherBoard errors...

Checking for CMOS errors...

License year: 2021: OK

Checking for SFP/NIC errors...
```

```
Checking for Firmware errors...

Checking for License errors...

Checking for Undetected CPUs...

Checking for DIMM flaps...

Checking for Memory Channel errors...

Checking for LOM errors...

Checking the Power Supply Errors...

Checking for Hardware Clock errors...

Script Done.

root@ns#
```

Important: Physically disconnect all ports, except the LOM port, including the management port from the network.

7. At the shell prompt, switch to the ADC command line. Type: exit

Example

```
root@ns# exit

logout

Done
```

8. Run the following command to shut down the appliance. You are asked if you want to completely stop the ADC. Type: `shutdown -p now`

Example

```
> shutdown -p now

Are you sure you want to completely stop NetScaler (Y/N)? [N]:y

Done

> Dec  5 06:09:11 <auth.notice> ns shutdown: power-down by root:

Dec  5 06:09:13 <auth.emerg> ns init: Rebooting via init mechanism

Dec  5 06:09:13 <syslog.err> ns syslogd: exiting on signal 15

Dec  5 06:09:13 aslearn[1662]: before pthread_join(), task name:
Aslearn_Packet_Loop_Task

Dec  5 06:09:15 aslearn[1662]: Exiting function ns_do_logging

Dec  5 06:09:15 aslearn[1662]: before pthread_join(), task name:
Aslearn_WAL_Cleanup_Task

Dec  5 06:09:15 aslearn[1662]: before pthread_join(), task name:
Aslearn_HA_Primary_Task

Dec  5 06:09:15 aslearn[1662]: 1662 exiting gracefully

Dec  5 06:09:18 [1672]: nsnet_tcpipconnect: connect() failed; returned
-1 errno=61

qat0: qat_dev0 stopped 12 acceleration engines

pci4: Resetting device

qat1: qat_dev1 stopped 12 acceleration engines

pci6: Resetting device

qat2: qat_dev2 stopped 12 acceleration engines

pci132: Resetting device

qat3: qat_dev3 stopped 12 acceleration engines

pci134: Resetting device
```

```
Dec  5 06:09:33 init: some processes would not die; ps axl advised

reboot initiated by init with parent kernel

Waiting (max 60 seconds) for system process `vnlru' to stop...done

Waiting (max 60 seconds) for system process `bufdaemon' to stop...done

Waiting (max 60 seconds) for system process `syncer' to stop...

Syncing disks, vnodes remaining...0 0 0 done

All buffers synced.

Uptime: 1h53m18s

ixl_shutdown: lldp start 0

ixl_shutdown: lldp start 0

ixl_shutdown: lldp start 0

ixl_shutdown: lldp start 0

usb0: Controller shutdown

uhub0: at usb0, port 1, addr 1 (disconnected)

usb0: Controller shutdown complete

usb1: Controller shutdown

uhub1: at usb1, port 1, addr 1 (disconnected)

ugen1.2: <vendor 0x8087> at usb1 (disconnected)

uhub3: at uhub1, port 1, addr 2 (disconnected)

ugen1.3: <FTDI> at usb1 (disconnected)

uftdi0: at uhub3, port 1, addr 3 (disconnected)

ugen1.4: <vendor 0x1005> at usb1 (disconnected)
```



```
umass0: at uhub3, port 3, addr 4 (disconnected)

(da0:umass-sim0:0:0:0): lost device - 0 outstanding, 0 refs

(da0:umass-sim0:0:0:0): removing device entry

usb1: Controller shutdown complete

usb2: Controller shutdown

uhub2: at usb2, port 1, addr 1 (disconnected)

  ugen2.2: <vendor 0x8087> at usb2 (disconnected)

uhub4: at uhub2, port 1, addr 2 (disconnected)

  ugen2.3: <vendor 0x0557> at usb2 (disconnected)

uhub5: at uhub4, port 7, addr 3 (disconnected)

  ugen2.4: <vendor 0x0557> at usb2 (disconnected)

ukbd0: at uhub5, port 1, addr 4 (disconnected)

ums0: at uhub5, port 1, addr 4 (disconnected)

usb2: Controller shutdown complete

ixl_shutdown: lldp start 0

ixl_shutdown: lldp start 0

ixl_shutdown: lldp start 0


ixl_shutdown: lldp start 0

acpi0: Powering system off
```

Upgrade the appliance

To upgrade the appliance, follow these steps:

1. Power off the ADC appliance.
2. Locate the two solid-state drives (SSDs) on the back of the appliance in slot #1 and slot #2, as shown in the following figure:



3. Verify that the replacement solid-state drives (SSDs) are the ones required for your ADC model. The conversion requires a minimum of four SSDs. The NetScaler label is on the top of one of the SSDs. The SSD is pre-populated with a new version of the BIOS and a recent build of the required NetScaler SDX Management Service. This SSD must be installed in slot # 1.
4. Remove the SSDs by pushing the safety latch of the drive cover down while pulling the drive handle.
5. On the new Citrix certified SSD drive, open the drive handle completely to the left. Then insert the new drive into the slot #1 as far as possible.
6. To seat the drive, close the handle flush with the rear side of the appliance so that the drive locks securely into the slot.

Important: The orientation of the SSD is important. When you insert the drive, make sure that the NetScaler® product label is at the side, similar to the already inserted SSDs.

7. Insert a second Citrix certified SSD, which matches the capacity of the SSD in slot #1, in slot # 2.

IMPORTANT

Mixing and matching of old and new SSDs is not supported. SSDs in slot #1 and slot #2, which constitute the first RAID pair (local storage), must be of the same size and type. Similarly, SSDs in slot #3 and slot #4, which constitute the second RAID pair (VPX storage), must be of the same size and type. Use only drives that are part of the provided conversion kit.

8. Insert the 960 GB blank SSD drives in slots 3, 4, 5, 6, 7, and 8.
9. Disconnect all the network cables from the data ports and the management ports.
10. Start the ADC appliance. For instructions, see [Switch on the appliance](#).

The conversion process can take up to 90 minutes, during which you must not power cycle the appliance. The entire conversion process might not be visible on the console and might appear to be unresponsive.

The conversion process updates the BIOS, installs the Citrix Hypervisor™, and the Management Service Operating system.

11. Keep the console cable attached during the conversion process. Allow the process to complete, at which point the SDX login prompt appears.
12. During the conversion process the LOM port connection might be lost as it resets the IP address to the default value of 192.168.1.3.
13. The default credentials for the Citrix Hypervisor are changed to root/nsroot after the appliance is converted from an MPX to SDX. If this password does not work, try typing nsroot/the serial number of the appliance. The serial number bar code is available at the back of the appliance, and is also available in the output of the `show hardware` command.

NOTE:

The serial number of the appliance remains the same.

14. To make sure that the conversion is successful, verify that the FVT result indicates success. Run the following command: `tail /var/log/fvt/fvt.log`

Reconfigure the converted appliance

After the conversion process, the appliance no longer has its previous working configuration. Therefore, you can access the appliance through a web browser only by using the default IP address: 192.168.100.1/16. Configure a computer on network 192.168.0.0 and connect it directly to the appliance's management port (0/1) with a cross-over Ethernet cable. Alternately, access the NetScaler SDX appliance through a network hub by using a straight-through Ethernet cable. Use the default credentials to log on, and then do the following:

1. Select the **Dashboard** tab.
2. Verify that the **System Resource** section displays the accurate number of CPU cores, SSL cores, and the total memory for your NetScaler SDX appliance.
3. Navigate to **Configuration > System** to modify the network information of the Management Service.
4. In the **Modify Network Configuration** dialog box, specify the following details:
 - **Interface:** Management interface that connects the appliance to a management workstation or network. Possible values: 0/1, 0/2. Default: 0/1.
 - **Gateway:** IP address of the router that forwards traffic out of the appliance's subnet.
 - Select the IPv4 checkbox if you want to use the IPv4 address for the Management Service and enter the details for the following parameters:
 - **Appliance Management IP:** The IPv4 address that is used to access the Management Service by using a web browser.
 - **Netmask:** The subnet mask in which the SDX appliance is located.
 - **DNS:** IPv4 address of the primary DNS server. IPv6 addresses are not supported for the primary DNS server.
 - Select the IPv6 checkbox if you want to use the IPv6 address for the Management Service and enter the details for the following parameters:
 - **Management Service IP Address:** The IPv6 address that is used to access the Management Service by using a Web browser.
 - **Gateway IPv6 Address:** The IPv4 address of the router that forwards traffic out of the appliance's subnet.
 - Select **Additional DNS** to add DNS server IP addresses as an extra DNS server apart from the primary DNS server. The IP addresses can be either IPv4 or IPv6.

Network Configuration

Management Service

Interface*
0/1

Gateway*
10 . 217 . 29 . 1

☒ IPv4

Appliance Management IP*
10 . 217 . 29 . 27

Netmask*
255 . 255 . 255 . 0

DNS
127 . 0 . 0 . 2

☐ IPv6

☐ Additional DNS

Appliance Supportability

☒ Configure Appliance supportability

Appliance supportability IP*
10 . 217 . 29 . 26

Gateway*
10 . 217 . 29 . 1

Netmask*
255 . 255 . 255 . 0

DNS
127 . 0 . 0 . 2

☐ Auto Disable Hypervisor Access

OK Close

5. Click **OK**. Connection to the Management Service is lost as the network information was changed.
6. Connect the NetScaler SDX appliance's management port 0/1 to a switch to access it through the network. Browse to the IP address used earlier and log in with the default credentials.
7. Apply the new licenses. For instructions, see [SDX Licensing Overview](#).
8. Navigate to **Configuration > System** and, in the **System Administration** group, click **Reboot Appliance**. Click **Yes** to confirm. You are now ready to provision the VPX instances on the NetScaler SDX appliance. For instructions, see [Provisioning NetScaler instances](#).

Converting a NetScaler MPX 24100 and 24150 appliance to a NetScaler SDX 24100 and 24150 appliance

You can convert a NetScaler MPX appliance to a NetScaler SDX appliance by upgrading the software through a new Solid State Drive (SSD). NetScaler supplies a field conversion kit to migrate a NetScaler MPX appliance to a NetScaler SDX appliance.

The conversion requires minimum of four SSDs.

NOTE

Citrix recommends that you configure the lights out management (LOM) Port of the NetScaler appliance before starting the conversion process. For more information on the LOM port of the NetScaler appliance, see [Lights out Management Port of the NetScaler appliance](#).

To convert a NetScaler MPX appliance to a NetScaler SDX appliance, you must access the appliance through a console cable attached to a computer or terminal. Before connecting the console cable, configure the computer or terminal to support the following configuration:

- VT100 terminal emulation
- 9600 baud
- 8 data bits
- 1 stop bit
- Parity and flow control set to NONE

Connect one end of the console cable to the RS232 serial port on the appliance, and the other end to the computer or terminal.

NOTE

To use a cable with an RJ-45 converter, insert the optional converter into the console port and attach the cable to it.

Citrix® recommends you to connect a VGA monitor to the appliance to monitor the conversion process, because the LOM connection is lost during the conversion process.

With the cable attached, verify that the MPX appliance's components are functioning correctly. You are then ready to begin the conversion. The conversion process modifies the BIOS, installs Citrix Hypervisor and a Management Service image, and copies the NetScaler VPX image to the Solid State Drive.

The conversion process also sets up a Redundant Array of Independent Disks (RAID) controller for local storage (SSD slot # 1 and SSD slot # 2) and NetScaler VPX storage (SSD slot # 3 and SSD slot # 4).

After the conversion process, modify the appliance's configuration and apply a new license. You can then provision the VPX instances through the Management Service on what is now a NetScaler SDX appliance.

To verify proper operation of the MPX appliance's components

1. Access the console port and enter the administrator credentials.
2. Run the following command from the command line interface of the appliance to display the serial number: `show hardware` The serial number might be helpful if you want to contact Citrix technical support.
3. Run the following command to display the status of the active 10G interfaces: `show interface`
4. In the `show interface` command's output, verify that all the interfaces are enabled and the status of every interface is shown as UP/UP.


NOTE

If you do not have an SFP+ transceiver for every port, verify the interfaces in stages. After checking the first set of interfaces, unplug the SFP+ transceivers and plug them in to the next set of ports.

5. Run the following command for each of the interfaces that are not in the UP/UP state:

To upgrade the appliance

1. Power off the NetScaler appliance.
2. Locate two solid-state drives (SSDs) on the back of the appliance in slot #1 and slot #2, as shown in the following figure:



3. Verify that the replacement solid-state drives (SSDs) are the ones required for your NetScaler model. The conversion requires minimum of four SSDs. The NetScaler label is on the top of one of the solid-state drives, which is pre-populated with a new version of BIOS and a recent build of the required NetScaler SDX Management Service. This SSD must be installed in slot # 1.
4. Remove the SSDs by pushing the safety latch of the drive cover down while pulling the drive handle.
5. On the new NetScaler® certified SSD drive, open the drive handle completely to the left, and then insert the new drive into the slot #1 as far as possible.
6. To seat the drive, close the handle flush with the rear side of the appliance so that the drive locks securely into the slot.

IMPORTANT

The orientation of the SSD is important. When you insert the drive, make sure that the NetScaler product label is at the top.

7. Insert a second NetScaler certified SSD, which matches the capacity of the SSD in slot #1, in slot # 2. Insert more blank NetScaler certified SSDs in slots #3 and #4.

IMPORTANT

Mixing and matching of old and new SSDs is not supported. SSDs in slot #1 and slot # 2, which constitute the first RAID pair (local storage), must be of the same size and type. Similarly, SSDs in slot # 3 and slot #4, which constitute the second RAID pair (VPX storage), must be of the same size and type. Only use drives that are part of the provided conversion kit.

8. Disconnect all network cables from the data ports and the management ports.
9. Start the NetScaler appliance. For instructions, see “Switching on the Appliance” in Installing the Hardware. The conversion process can run for approximately 30 minutes, during which you must not power cycle the appliance. The entire conversion process might not be visible on the console and might appear to be unresponsive. The conversion process updates the BIOS, installs the Citrix Hypervisor and the Management Service. It also copies the NetScaler VPX image to the SSD for instance provisioning, and forms the Raid1 pair.

NOTE

The serial number of the appliance remains the same.

10. Keep the console cable attached during the conversion process. Allow the process to complete, at which point the netscaler-sdx login: prompt appears.
11. During the conversion process the LOM port connection might be lost as it resets the IP address to the default value of 192.168.1.3. The conversion status output is available on the VGA monitor.

Convert a NetScaler MPX 14020/14030/14040/14060/14080/14100 appliance to a NetScaler SDX 14020/14030/14040/ 14060/14080/14100 appliance

You can convert a NetScaler MPX appliance to a NetScaler SDX appliance by upgrading the software through a new solid-state drive (SSD). NetScaler supplies a field conversion kit to migrate a NetScaler MPX appliance to an SDX appliance.

The conversion requires minimum of four SSDs.

Note: Citrix recommends that you configure the lights out management (LOM) port of the NetScaler appliance before starting the conversion process. For more information on the LOM port of the NetScaler appliance, see [Lights out management port of the NetScaler SDX appliance](#).

To convert a NetScaler MPX appliance to an SDX appliance, you must access the appliance through a console cable attached to a computer or terminal. Before connecting the console cable, configure the computer or terminal to support the following configuration:

- VT100 terminal emulation
- 9600 baud
- 8 data bits
- 1 stop bit
- Parity and flow control set to NONE

Connect one end of the console cable to the RS232 serial port on the appliance, and the other end to the computer or terminal.

Note: To use a cable with an RJ-45 converter, insert the optional converter into the console port and attach the cable to it.

Citrix® recommends that you connect a VGA monitor to the appliance to monitor the conversion process, because the LOM connection might be lost during the conversion process.

With the cable attached, verify that the MPX appliance's components are functioning correctly. You are then ready to begin the conversion. The conversion process modifies the BIOS, installs the Citrix Hypervisor™ and a Management Service image, and copies the VPX image to the solid-state drive.

The conversion process sets up a redundant array of independent disks (RAID) controller for local storage:

- (SSD slot # 1 and SSD slot # 2)
- NetScaler VPX storage (SSD slot # 3 and SSD slot # 4)

After the conversion process, you modify the configuration of the appliance and apply a new license. You can then provision the VPX instances through the Management Service on what is now an SDX appliance.

To verify proper operation of the MPX appliance's components

1. Access the console port and enter the administrator credentials.

2. Run the following command from the command line interface of the appliance to display the serial number: `show hardware`

The serial number might be helpful in case you want to contact Citrix technical support.

Example

```
show hardware

Platform: NSMPX-14000 12*CPU+16*IX+2*E1K+2*CVM N3 250101

Manufactured on: 10/2/2015

CPU: 2600MHZ

Host Id: 234913926

Serial no: JSW4UCKKM5

Encoded serial no: JSW4UCKKM5

Done
```

3. Run the following command to display the status of the active 10G interfaces: `show interface`
4. In the `show interface` command's output, verify that all the interfaces are enabled and the status of every interface is shown as UP/UP.

Note: If you do not have an SFP+ transceiver for every port, verify the interfaces in stages. After checking the first set of interfaces, unplug the SFP+ transceivers and plug them in to the next set of ports.

5. Run the following command for each of the interfaces that are not in the UP/UP state: `enable interface 10/x` where x is the new interface number.
6. Run the following command to verify that the status of the power supplies is normal: `stat system -detail`
7. Run the following command to generate a tar of system configuration data and statistics: `show techsupport`

Note: The output of the command is available in the `/var/tmp/support/collector_<IP_address>_P_<date>.tar.gz` file. Copy this file to another computer for

future reference. The output of the command might be helpful in case you want to contact Citrix technical support.

8. At the command line interface, switch to the shell prompt. Type: `shell`
9. Run the following command to verify the number of Cavium cards available depending upon your appliance: `root@ns# grep "cavium" /var/nslog/dmesg.boot`

Example

```
root@ns# grep "cavium"; /var/nslog/dmesg.boot

Cavium cavium_probe : found card 0x177d,device=0x11

cavium0 mem 0xdd600000-0xdd6fffff irq 32 at device 0.0 on pci3

Cavium cavium_probe : found card 0x177d,device=0x11

cavium1 mem 0xfaa00000-0xfaafffff irq 64 at device 0.0 on pci136
```

10. Run the following command to verify the RAM memory reserved for shared memory depending upon your appliance: `root@ns# grep "memory" /var/nslog/dmesg.boot`

Example

```
root@ns# grep "memory"; /var/nslog/dmesg.boot

real memory = 70866960384 (67584 MB)

avail memory = 66267971584 (63198 MB)
```

11. Run the following command to verify the number of CPU cores depending upon your appliance: `root@ns# grep "cpu" /var/nslog/dmesg.boot`

Example

```
root@ns# grep "cpu"; /var/nslog/dmesg.boot

cpu0 (BSP): APIC ID: 0

cpu1 (AP): APIC ID: 2
```

```
cpu2 (AP): APIC ID: 4

cpu3 (AP): APIC ID: 6

cpu4 (AP): APIC ID: 8

cpu5 (AP): APIC ID: 10

cpu6 (AP): APIC ID: 32

cpu7 (AP): APIC ID: 34

cpu8 (AP): APIC ID: 36

cpu9 (AP): APIC ID: 38

cpu10 (AP): APIC ID: 40

cpu11 (AP): APIC ID: 42

cpu0: <ACPI CPU> on acpi0

acpi_throttle0: <ACPI CPU Throttling> on cpu0

cpu1: <ACPI CPU> on acpi0

acpi_throttle1: <ACPI CPU Throttling> on cpu1

cpu2: <ACPI CPU> on acpi0

acpi_throttle2: <ACPI CPU Throttling> on cpu2

cpu3: <ACPI CPU> on acpi0

acpi_throttle3: <ACPI CPU Throttling> on cpu3

cpu4: <ACPI CPU> on acpi0

acpi_throttle4: <ACPI CPU Throttling> on cpu4

cpu5: <ACPI CPU> on acpi0

acpi_throttle5: <ACPI CPU Throttling> on cpu5
```

```

cpu6: <ACPI CPU> on acpi0

acpi_throttle6: <ACPI CPU Throttling> on cpu6

cpu7: <ACPI CPU> on acpi0

acpi_throttle7: <ACPI CPU Throttling> on cpu7

cpu8: <ACPI CPU> on acpi0

acpi_throttle8: <ACPI CPU Throttling> on cpu8

cpu9: <ACPI CPU> on acpi0

acpi_throttle9: <ACPI CPU Throttling> on cpu9

cpu10: <ACPI CPU> on acpi0

acpi_throttle10: <ACPI CPU Throttling> on cpu10

cpu11: <ACPI CPU> on acpi0

acpi_throttle11: <ACPI CPU Throttling> on cpu11

root@ns#

```

12. Run the following command to verify that the /var drive is mounted as /dev/ad8s1e: root@ns#
df -h
13. Run the `ns_hw_err.bash` script, which checks for latent hardware errors: root@ns#
`ns_hw_err.bash`

Example

```

root@ns# ns_hw_err.bash

NetScaler NS10.1: Build 133.11.nc, Date: Sep 21 2015, 17:59:51

platform: serial JSW4UCKKM5

platform: sysid 250101 - NSMPX-14000 12*CPU+16*IX+2*E1K+2*CVM N3

```

```
HDD MODEL: ar0: 227328MB &lt;Intel MatrixRAID RAID1&gt; status: READY
```

```
Generating the list of newslog files to be processed...
```

```
Generating the events from newslog files...
```

```
Checking for HDD errors...
```

```
Checking for HDD SMART errors...
```

```
Checking for Flash errors...
```

```
/var/nslog/dmesg.prev:* DEVELOPER mode - run NetScaler manually!
```

```
*****
```

```
FOUND          1 Flash errors: DEVELOPER mode - run NetScaler manually
```

```
*****
```

```
Checking for SSL errors...
```

```
Checking for BIOS errors...
```

```
Checking for SMB errors...
```

```
Checking for MotherBoard errors...
```

```
Checking for CMOS errors...
```

```
License year: 2015: OK
```

```
License server failed at startup. Check /var/log/license.log
```

```
Vendor daemon failed at startup. Check /var/log/license.log
```

```
Checking for SFP/NIC errors...
```

```
Checking for Firmware errors...
```

```
Checking for License errors...
```

```
Checking for Undetected CPUs...
```

```

Checking for DIMM flaps...

Checking for LOM errors...

Checking the Power Supply Errors...


root@ns#

```

14. Important: Physically disconnect all ports except the LOM port, including the management port, from the network.
15. At the shell prompt, switch to the NetScaler command line. Type: `exit`
16. Run the following command to shut down the appliance: `shutdown -p now`

To upgrade the appliance

1. Power off the NetScaler appliance.
2. Locate two solid-state drives (SSDs) on the back of the appliance in slot #1 and slot #2, as shown in the following figure:



3. Verify that the replacement SSDs are the ones required for your NetScaler model. The conversion requires minimum of four SSDs. The NetScaler label is on the top of one of the SSDs. The SSD is pre-populated with a new version of the BIOS and a recent build of the Management Service. This SSD must be installed in slot # 1.
4. Remove the SSDs by pushing the safety latch of the drive cover down while pulling the drive handle.

5. On the new NetScaler® certified SSD drive, open the drive handle completely to the left. Then insert the new drive into the slot #1 as far as possible.
6. To seat the drive, close the handle flush with the rear side of the appliance so that the drive locks securely into the slot.

Important: The orientation of the SSD is important. When you insert the drive, make sure that the NetScaler product label is at the top.

7. Insert a second NetScaler certified SSD, which matches the capacity of the SSD in slot #1, in slot # 2. Insert other blank NetScaler certified SSDs in slots #3 and #4.

Note: If the license of your appliance is 14040, insert other blank NetScaler certified SSDs in slots #3, #4, #5, and #6. If the license of your appliance is 14060/14080/14100, insert other blank NetScaler certified SSDs in slots #3, #4, #5, #6, #7, and #8.

Important: Mixing and matching of old and new SSDs is not supported. SSDs in slot #1 and slot # 2, which constitute the first RAID pair (local storage), must be of the same size and type. Similarly, SSDs in slot # 3 and slot # 4, which constitute the second RAID pair (VPX storage), must be of the same size and type. Only use drives that are part of the provided conversion kit.

8. Store the old SSDs for future handling.
9. Disconnect all network cables from the data ports and the management ports.
10. Start the NetScaler appliance. For instructions, see “Switching on the Appliance” in [Installing the hardware](#). The conversion process can run for approximately 30 minutes, during which you must not power cycle the appliance. The entire conversion process might not be visible on the console and might appear to be unresponsive. The conversion process updates the BIOS, installs the Citrix Hypervisor and the Management Service. It also copies the VPX image to the SSD for instance provisioning, and forms the Raid1 pair. **Note:** The serial number of the appliance remains the same.
11. Keep the console cable attached during the conversion process. Allow the process to complete, at which point the netscaler-sdx login: prompt appears.
12. During the conversion process the LOM port connection might be lost as it resets the IP address to the default value of 192.168.1.3. The conversion status output is available on the VGA monitor.
13. To make sure that the conversion is successful, verify that the FVT result indicates success. Run the following command: `tail /var/log/fvt/fvt.log`

Example

```
[root@netscaler-sdx ~]# tail /var/log/fvt/fvt.log

Wed, 28 Oct 2015 04:40:47 /opt/xensource/packages/files/fvt/workers/
check_vf_count --pf_device="0000:89:00.1" --
vf_count="40";

Wed, 28 Oct 2015 04:40:47 => PASS

Wed, 28 Oct 2015 04:40:47 /opt/xensource/packages/files/fvt/workers/
check_vf_count --pf_device="0000:03:00.0" --
vf_count="8";

Wed, 28 Oct 2015 04:40:47 => PASS

Wed, 28 Oct 2015 04:40:47 /opt/xensource/packages/files/fvt/workers/
check_vf_count --pf_device="0000:88:00.0" --
vf_count="8";

Wed, 28 Oct 2015 04:40:47 => PASS

Wed, 28 Oct 2015 04:40:47 FVT RESULT: SUCCESS!<br />

[root@netscaler-sdx ~]#
```

To reconfigure the converted appliance

After the conversion process, the appliance no longer has its previous working configuration. Therefore, you can access the appliance through a web browser only by using the default IP address: 192.168.100.1/16. Configure a computer on network 192.168.0.0 and connect it directly to the appliance's management port (0/1) with a cross-over Ethernet cable. Alternately, access the SDX appliance through a network hub by using a straight through Ethernet cable. Use the default credentials to log on, and then do the following:

1. Select the **Configuration** tab.
2. Verify that the System Resource section displays the accurate number of CPU cores, SSL cores, and the total memory for your SDX appliance.
3. Select the **System** node and, under **Set Up Appliance**, click **Network Configuration** to modify the IP address of the Management Service.
4. In the **Configure Network Configuration** dialog box, specify the following details:

- **Interface***—The interface through which clients connect to the Management Service. Possible values: 0/1, 0/2. Default: 0/1.
- **Citrix Hypervisor IP Address***—The IP address of Citrix Hypervisor.
- **Management Service IP Address***—The IP address of the Management Service.
- **Netmask***—The subnet mask for the subnet in which the SDX appliance is located.
- **Gateway***—The default gateway for the network.
- **DNS Server**—The IP address of the DNS server.

*A mandatory parameter

5. Click OK. Connection to the Management Service is lost as the network information was changed.
6. Connect the SDX appliance's management port 0/1 to a switch to access it through the network. Browse to the IP address used earlier and log on with the default credentials.
7. Apply the new licenses. For instructions, see [SDX Licensing Overview](#).
8. Navigate to **Configuration > System** and, in the **System Administration** group, click **Reboot Appliance**. Click **Yes** to confirm. You are now ready to provision the VPX instances on the NetScaler SDX appliance. For instructions, see [Provisioning NetScaler Instances](#).

Convert a NetScaler MPX 14020 40G, MPX 14040 40G, MPX 14060 40G, MPX 14080 40G appliance to a NetScaler SDX 14020 40G, SDX 14040 40G, SDX 14060 40G, SDX 14080 40G appliance

You can convert a NetScaler MPX appliance to a NetScaler SDX appliance by upgrading the software through a new solid-state drive (SSD). NetScaler supplies a field conversion kit to migrate a NetScaler MPX appliance to an SDX appliance.

The conversion requires minimum of four SSDs.

Note: Citrix recommends that you configure the lights out management (LOM) port of the appliance before starting the conversion process. For more information on the LOM port of the NetScaler appliance, see [Lights out management port of the NetScaler SDX appliance](#).

To convert an MPX appliance to an SDX appliance, you must access the appliance through a console cable attached to a computer or terminal. Before connecting the console cable, configure the computer or terminal to support the following configuration:

- VT100 terminal emulation
- 9600 baud
- 8 data bits
- 1 stop bit
- Parity and flow control set to NONE

Connect one end of the console cable to the RS232 serial port on the appliance, and the other end to the computer or terminal.

Note: To use a cable with an RJ-45 converter, insert the optional converter into the console port and attach the cable to it.

Citrix® recommends you to connect a VGA monitor to the appliance to monitor the conversion process, because the LOM connection is lost during the conversion process.

With the cable attached, verify that the MPX appliance's components are functioning correctly. You are then ready to begin the conversion. The conversion process modifies the BIOS, installs a Citrix Hypervisor™ and a Service Virtual Machine image, and copies the NetScaler VPX image to the Solid State Drive.

The conversion process also sets up a redundant array of independent disks (RAID) controller for local storage (SSD slot # 1 and SSD slot # 2) and NetScaler VPX™ storage (SSD slot # 3 and SSD slot # 4).

After the conversion process, make modifications to the appliance's configuration and apply a new license. You can then provision the VPX instances through the Management Service on what is now a NetScaler SDX™ appliance.

Verify proper operation of the MPX appliance's components

1. Access the console port and enter the administrator credentials.
2. Run the following command from the command line interface of the appliance to display the serial number: `show hardware`

Example

```
show hardware
```

```
Platform: NSMPX-14000-40G 12*CPU+16*F1X+4*F4X+2*E1K+2*CVM N3
250140
```

```
Manufactured on: 3/21/2016
```

```
CPU: 2600MHZ
```

```
Host Id: 234913926
```

```
Serial no: KZCHGCN810
```

```
Encoded serial no: KZCHGCN810
```

```
Done
```

The serial number might be helpful when you want to contact Citrix technical support.

3\ Run the following command to display the status of the active interfaces:

```
show interface
```

Example

```
show interface
```

```
1)      Interface 0/1 (Gig Ethernet 10/100/1000 Mbits) #10

        flags=0xc020 <ENABLED, UP, UP, autoneg, HAMON, 802.1q>

        MTU=1500, native vlan=1, MAC=0c:c4:7a:41:ea:94, uptime 43h59m55s

        Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

                throughput 0

        Actual: media UTP, speed 1000, duplex FULL, fctl OFF, throughput
1000

        LLDP Mode: NONE

        RX: Pkts(44337150) Bytes(2767558041) Errs(0) Drops(41260552)
Stalls(0)

        TX: Pkts(128116) Bytes(9004699) Errs(0) Drops(0) Stalls(0)
```

```
NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)
```

```
Bandwidth thresholds are not set.
```

2) Interface 0/2 (Gig Ethernet 10/100/1000 MBits) #11

```
flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>
```

```
MTU=1500, native vlan=1, MAC=0c:c4:7a:41:ea:95, downtime 44h00m08s
```

```
Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,
```

```
throughput 0
```

```
LLDP Mode: NONE
```

```
RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
```

```
TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
```

```
NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)
```

```
Bandwidth thresholds are not set.
```

3) Interface 10/1 (10G Ethernet) #17

```
flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>
```

```
MTU=1500, native vlan=1, MAC=68:05:ca:30:e1:40, downtime 44h00m08s
```

```
Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,
```

```
throughput 0
```

```
LLDP Mode: NONE
```

```
RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
```

```
TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
```

```
NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)
```

```
Bandwidth thresholds are not set.
```

4) Interface 10/2 (10G Ethernet) #18

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>

MTU=1500, native vlan=1, MAC=68:05:ca:30:e1:41, downtime 44h00m08s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

5) Interface 10/3 (10G Ethernet) #19

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>

MTU=1500, native vlan=1, MAC=68:05:ca:30:e1:42, downtime 44h00m08s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

6) Interface 10/4 (10G Ethernet) #20

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>

MTU=1500, native vlan=1, MAC=68:05:ca:30:e1:43, downtime 44h00m08s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

7) Interface 10/5 (10G Ethernet) #13

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>

MTU=1500, native vlan=1, MAC=68:05:ca:2e:6b:a0, downtime 44h00m08s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

8) Interface 10/6 (10G Ethernet) #14

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>

MTU=1500, native vlan=1, MAC=68:05:ca:2e:6b:a1, downtime 44h00m08s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

9) Interface 10/7 (10G Ethernet) #15

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>;

MTU=1500, native vlan=1, MAC=68:05:ca:2e:6b:a2, downtime 44h00m08s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

10) Interface 10/8 (10G Ethernet) #16

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>;

MTU=1500, native vlan=1, MAC=68:05:ca:2e:6b:a3, downtime 44h00m08s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE

```

    RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

    TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

    Bandwidth thresholds are not set.

11)    Interface 10/9 (10G Ethernet) #0

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>

MTU=1500, native vlan=1, MAC=68:05:ca:30:e1:a0, downtime 44h00m08s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

        throughput 0

LLDP Mode: NONE

    RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

    TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

    Bandwidth thresholds are not set.

12)    Interface 10/10 (10G Ethernet) #1

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>

MTU=1500, native vlan=1, MAC=68:05:ca:30:e1:a1, downtime 44h00m08s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

        throughput 0

LLDP Mode: NONE

    RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

    TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

```

```

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

13)   Interface 10/11 (10G Ethernet) #2

      flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>;

      MTU=1500, native vlan=1, MAC=68:05:ca:30:e1:a2, downtime 44h00m08s

      Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

               throughput 0

      LLDP Mode: NONE

      RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

      TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

      NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

      Bandwidth thresholds are not set.

14)   Interface 10/12 (10G Ethernet) #3

      flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>;

      MTU=1500, native vlan=1, MAC=68:05:ca:30:e1:a3, downtime 44h00m08s

      Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

               throughput 0

      LLDP Mode: NONE

      RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

      TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

      NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

```

Bandwidth thresholds are not set.

15) Interface 10/13 (10G Ethernet) #4

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>;

MTU=1500, native vlan=1, MAC=68:05:ca:30:e2:70, downtime
44h00m08s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,
throughput 0

LLDP Mode: NONE

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

16) Interface 10/14 (10G Ethernet) #5

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>;

MTU=1500, native vlan=1, MAC=68:05:ca:30:e2:71, downtime
44h00m15s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,
throughput 0

LLDP Mode: NONE

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

17) Interface 10/15 (10G Ethernet) #6

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>

MTU=1500, native vlan=1, MAC=68:05:ca:30:e2:72, downtime
44h00m15s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

18) Interface 10/16 (10G Ethernet) #7

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>

MTU=1500, native vlan=1, MAC=68:05:ca:30:e2:73, downtime
44h00m15s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

19) Interface 40/1 (40G Ethernet) #9

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>

MTU=1500, native vlan=1, MAC=3c:fd:fe:9c:f9:98, downtime
44h00m15s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,
throughput 0

LLDP Mode: NONE

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

20) Interface 40/2 (40G Ethernet) #21

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>

MTU=1500, native vlan=1, MAC=3c:fd:fe:9c:fa:68, downtime
44h00m15s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,
throughput 0

LLDP Mode: NONE

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

21) Interface 40/3 (40G Ethernet) #12

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>

MTU=1500, native vlan=1, MAC=3c:fd:fe:9c:f9:f8, downtime
44h00m15s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

22) Interface 40/4 (40G Ethernet) #8

flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, 802.1q>

MTU=1500, native vlan=1, MAC=3c:fd:fe:9c:f9:e0, downtime
44h00m15s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

23) Interface L0/1 (Netscaler Loopback interface) #22

```
flags=0x20008020 <ENABLED, UP, UP>

MTU=1500, native vlan=1, MAC=0c:c4:7a:41:ea:94, uptime 44h00m15s

LLDP Mode: NONE

RX: Pkts(103847611) Bytes(14441166499) Errs(0) Drops(0) Stalls(0)

TX: Pkts(206723664) Bytes(20488047637) Errs(0) Drops(0) Stalls(0)

Bandwidth thresholds are not set.
```

Done

4). In the show interface command's output, verify that all of the interfaces are enabled and the status of every interface is shown as UP/UP.

Notes:

- The interface status is displayed as UP/UP only if the cables are connected to the interfaces.
- If you do not have an SFP+ transceiver for every port, verify the interfaces in stages. After checking the first set of interfaces, unplug the SFP+ transceivers and plug them in to the next set of ports.

5). Run the following command for each of the interfaces that are not in the UP/UP state:

```
enable interface 10/x

enable interface 40/x
```

Where x is the new interface number.

Example:

```
enable interface 10/1

Done

> enable interface 40/1
```


Done

6). Run the following command to verify that the status of the power supplies is normal: `stat system -detail`

Example

```
stat system -detail
```

NetScaler Executive View

System Information:

Up since Mon Mar 21 13:53:22 2016

Memory usage (MB)	1177
-------------------	------

InUse Memory (%)	5.37
------------------	------

Number of CPUs	7
----------------	---

System Health Statistics (Standard):

```
CPU 0 Core Voltage (Volts) 0.99
```

```
CPU 1 Core Voltage (Volts) 0.98
```

Main 3.3 V Supply Voltage	3.36
---------------------------	------

Standby 3.3 V Supply Voltage	3.31
------------------------------	------

+5.0 V Supply Voltage	5.06
-----------------------	------

+12.0 V Supply Voltage	12.08
------------------------	-------

Battery Voltage (Volts)	3.17
-------------------------	------

```
Intel CPU Vtt Power(Volts) 1.01
```

5V Standby Voltage(Volts) 0.00

Voltage Sensor2(Volts)	0.00
------------------------	------

CPU Fan 0 Speed (RPM)	4650
CPU Fan 1 Speed (RPM)	4575
System Fan Speed (RPM)	4650
System Fan 1 Speed (RPM)	4650
System Fan 2 Speed (RPM)	4650
CPU 0 Temperature (Celsius)	42
CPU 1 Temperature (Celsius)	49
Internal Temperature (Celsius)	36
Power supply 1 status	NORMAL
Power supply 2 status	NORMAL
Power supply 3 status	NOT SUPPORTED
Power supply 4 status	NOT SUPPORTED
System Disk Statistics:	
/flash Size (MB)	16858
/flash Used (MB)	186
/flash Available (MB)	15323
/flash Used (%)	1
/var Size (MB)	143802
/var Used (MB)	781
/var Available (MB)	131517
/var Used (%)	0

System Health Statistics(Auxiliary):

Voltage 0 (Volts)	1.49
Voltage 1 (Volts)	1.49
Voltage 2 (Volts)	1.50
Voltage 3 (Volts)	1.49
Voltage 4 (Volts)	1.49
Voltage 5 (Volts)	0.00
Voltage 6 (Volts)	0.00
Voltage 7 (Volts)	0.00
Fan 0 Speed (RPM)	4650
Fan 1 Speed (RPM)	0
Fan 2 Speed (RPM)	0
Fan 3 Speed (RPM)	0
Temperature 0 (Celsius)	28
Temperature 1 (Celsius)	43
Temperature 2 (Celsius)	0
Temperature 3 (Celsius)	0
Done	

7). Run the following command to generate a tar of system configuration data and statistics: `show techsupport`

Example

```
show techsupport

showtechsupport data collector tool - $Revision: #2 $!
```

NetScaler version 10.5

Creating /var/tmp/support

The NS IP of this box is 192.168.100.1

Current HA state: Primary (or this is not part of HA pair!)

All the data will be collected under

/var/tmp/support/collector_P_192.168.100.1_23Mar2016_10_00

Copying selected configuration files from nsconfig

Copying WebInterface configuration files (if WI is installed)

.... WI is not installed on this system. Nothing to copy.

Running shell commands

Running CLI show commands

Running CLI stat commands

Running vtysh commands

Determining newslog files to archive

Last newslog file index=1

Warning! Missing newslog.0 and newslog.0.tar and newslog.0.tar.gz file!

Warning! Missing newslog.199 and newslog.199.tar and newslog.199.tar.gz file!

Warning! Missing newslog.198 and newslog.198.tar and newslog.198.tar.gz file!

Warning! Missing newslog.197 and newslog.197.tar and newslog.197.tar.gz file!

Warning! Missing newslog.196 and newslog.196.tar and

newslog.196.tar.gz file!

Five newslog files missing

.... copied 2 files from this directory.

Copying core files from /var/core (last 5 files created within the last week)

.... Nothing to copy No files created within the last one week

Copying core files from /var/crash (last 5 files created within the last week)

.... Nothing to copy No files created within the last one week

Copying imported files and mapping files

Copying GSLB location database files

Copying files from /var/log/db

Copying messages,ns.log,dmesg and other log files

Warning : cluster showtech node failure log file: /var/tmp/support/collector_P_192.168.100.1_23Mar2016_10_00/cluster_showtech_node_failure.log is not present

Archiving all the data into "/var/tmp/support/collector_P_192.168.100.1_23Mar2016_10_00.tar.gz"

Created a symbolic link for the archive with /var/tmp/support/support.tgz

/var/tmp/support/support.tgz ---- points to ---> /var/tmp/support/collector_P_192.168.100.1_23Mar2016_10_00.tar.gz

If this node is part of HA pair, please run it on the other node also!!

Done

Note: The output of the command is available in the `/var/tmp/support/collector\_IP\_address\_P\_date>.tar.gz` file. Copy this file to another computer for future reference. The output of the command might be helpful when you want to contact Citrix technical support.

8\). At the command line interface, switch to the shell prompt. Type: `shell`

Example

```
shell

Copyright (c) 1992-2013 The FreeBSD Project.

Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994
    The Regents of the University of California. All rights reserved.
```

9\). Run the following command to verify the number of Cavium cards available depending upon your appliance:

```
root@ns# grep "cavium" /var/nslog/dmesg.boot
```

Example

```
/var/nslog/dmesg.boot

Cavium cavium_probe : found card 0x177d,device=0x11

cavium0 mem 0xde200000-0xde2ffffff irq 32 at device 0.0 on pci3

cavium0: [ITHREAD]

Cavium cavium_probe : found card 0x177d,device=0x11

cavium1 mem 0xfb400000-0xfb4ffffff irq 64 at device 0.0 on pci136

cavium1: [ITHREAD]

Cavium cavium_probe : found card 0x177d,device=0x11
```

```
cavium0 mem 0xde200000-0xde2ffffff irq 32 at device 0.0 on pci3

cavium0: [ITHREAD]

Cavium cavium_probe : found card 0x177d,device=0x11

cavium1 mem 0xfb400000-0xfb4ffffff irq 64 at device 0.0 on pci136

cavium1: [ITHREAD]

root@ns#
```

10\). Run the following command to verify the RAM memory reserved for shared memory depending upon your appliance:

```
root@ns# grep "memory" /var/nslog/dmesg.boot
```

Example

```
/var/nslog/dmesg.boot

real memory  = 70866960384 (67584 MB)

avail memory = 66206515200 (63139 MB)

NS-KERN map_shared_mem_ioctl (cpu 11, NSPPE-00): Reserving 596 MB for
shared memory type 0

real memory  = 70866960384 (67584 MB)

avail memory = 66206515200 (63139 MB)

root@ns#
```

11\). Run the following command to verify the number of CPU cores depending upon your appliance:

```
root@ns# grep "cpu" /var/nslog/dmesg.boot
```

Example

```
root@ns# grep "cpu" /var/nslog/dmesg.boot
```

```
cpu0 (BSP): APIC ID: 0
```

```
cpu1 (AP): APIC ID: 2
```

```
cpu2 (AP): APIC ID: 4
```

```
cpu3 (AP): APIC ID: 6
```

```
cpu4 (AP): APIC ID: 8
```

```
cpu5 (AP): APIC ID: 10
```

```
cpu6 (AP): APIC ID: 32
```

```
cpu7 (AP): APIC ID: 34
```

```
cpu8 (AP): APIC ID: 36
```

```
cpu9 (AP): APIC ID: 38
```

```
cpu10 (AP): APIC ID: 40
```

```
cpu11 (AP): APIC ID: 42
```

```
cpu0: <ACPI CPU> on acpi0
```

```
cpu1: <ACPI CPU> on acpi0
```

```
cpu2: <ACPI CPU> on acpi0
```

```
cpu3: <ACPI CPU> on acpi0
```

```
cpu4: <ACPI CPU> on acpi0
```

```
cpu5: <ACPI CPU> on acpi0
```

```
cpu6: <ACPI CPU> on acpi0
```

```
cpu7: <ACPI CPU> on acpi0
```

```
cpu8: <ACPI CPU> on acpi0
```


cpu9: <ACPI CPU> on acpi0

cpu10: <ACPI CPU> on acpi0

cpu11: <ACPI CPU> on acpi0

est0: <Enhanced SpeedStep Frequency Control> on cpu0

p4tcc0: <CPU Frequency Thermal Control> on cpu0

est1: <Enhanced SpeedStep Frequency Control> on cpu1

p4tcc1: <CPU Frequency Thermal Control> on cpu1

est2: <Enhanced SpeedStep Frequency Control> on cpu2

p4tcc2: <CPU Frequency Thermal Control> on cpu2

est3: <Enhanced SpeedStep Frequency Control> on cpu3

p4tcc3: <CPU Frequency Thermal Control> on cpu3

est4: <Enhanced SpeedStep Frequency Control> on cpu4

p4tcc4: <CPU Frequency Thermal Control> on cpu4

est5: <Enhanced SpeedStep Frequency Control> on cpu5

p4tcc5: <CPU Frequency Thermal Control> on cpu5

est6: <Enhanced SpeedStep Frequency Control> on cpu6

p4tcc6: <CPU Frequency Thermal Control> on cpu6

est7: <Enhanced SpeedStep Frequency Control> on cpu7

p4tcc7: <CPU Frequency Thermal Control> on cpu7

est8: <Enhanced SpeedStep Frequency Control> on cpu8

p4tcc8: <CPU Frequency Thermal Control> on cpu8

est9: <Enhanced SpeedStep Frequency Control> on cpu9

p4tcc9: <CPU Frequency Thermal Control> on cpu9

est10: <Enhanced SpeedStep Frequency Control> on cpu10

p4tcc10: <CPU Frequency Thermal Control> on cpu10

est11: <Enhanced SpeedStep Frequency Control> on cpu11

p4tcc11: <CPU Frequency Thermal Control> on cpu11

NS-KERN map_shared_mem_ioctl (cpu 11, NSPPE-00): Reserving 596 MB for shared memory type 0

cpu0 (BSP): APIC ID: 0

cpu1 (AP): APIC ID: 2

cpu2 (AP): APIC ID: 4

cpu3 (AP): APIC ID: 6

cpu4 (AP): APIC ID: 8

cpu5 (AP): APIC ID: 10

cpu6 (AP): APIC ID: 32

cpu7 (AP): APIC ID: 34

cpu8 (AP): APIC ID: 36

cpu9 (AP): APIC ID: 38

cpu10 (AP): APIC ID: 40

cpu11 (AP): APIC ID: 42

cpu0: <ACPI CPU> on acpi0

cpu1: <ACPI CPU> on acpi0

cpu2: <ACPI CPU> on acpi0

```
cpu3: <ACPI CPU> on acpi0
cpu4: <ACPI CPU> on acpi0
cpu5: <ACPI CPU> on acpi0
cpu6: <ACPI CPU> on acpi0
cpu7: <ACPI CPU> on acpi0
cpu8: <ACPI CPU> on acpi0
cpu9: <ACPI CPU> on acpi0
cpu10: <ACPI CPU> on acpi0
cpu11: <ACPI CPU> on acpi0

est0: <Enhanced SpeedStep Frequency Control> on cpu0
p4tcc0: <CPU Frequency Thermal Control> on cpu0
est1: <Enhanced SpeedStep Frequency Control> on cpu1
p4tcc1: <CPU Frequency Thermal Control> on cpu1
est2: <Enhanced SpeedStep Frequency Control> on cpu2
p4tcc2: <CPU Frequency Thermal Control> on cpu2
est3: <Enhanced SpeedStep Frequency Control> on cpu3
p4tcc3: <CPU Frequency Thermal Control> on cpu3
est4: <Enhanced SpeedStep Frequency Control> on cpu4
p4tcc4: <CPU Frequency Thermal Control> on cpu4
est5: <Enhanced SpeedStep Frequency Control> on cpu5
p4tcc5: <CPU Frequency Thermal Control> on cpu5
est6: <Enhanced SpeedStep Frequency Control> on cpu6
```

```

p4tcc6: <CPU Frequency Thermal Control> on cpu6

est7: <Enhanced SpeedStep Frequency Control> on cpu7

p4tcc7: <CPU Frequency Thermal Control> on cpu7

est8: <Enhanced SpeedStep Frequency Control> on cpu8

p4tcc8: <CPU Frequency Thermal Control> on cpu8

est9: <Enhanced SpeedStep Frequency Control> on cpu9

p4tcc9: <CPU Frequency Thermal Control> on cpu9

est10: <Enhanced SpeedStep Frequency Control> on cpu10

p4tcc10: <CPU Frequency Thermal Control> on cpu10

est11: <Enhanced SpeedStep Frequency Control> on cpu11

p4tcc11: <CPU Frequency Thermal Control> on cpu11

root@ns#

```

12\ Run the following command to verify that the /var drive is mounted as /dev/ ar0s1a: root@ns#
df -h

Example

```

root@ns# df -h

```

Filesystem	Size	Used	Avail	Capacity	Mounted on
/dev/md0	327M	317M	3.1M	99%	/
devfs	1.0k	1.0k	0B	100%	/dev
procfs	4.0k	4.0k	0B	100%	/proc
/dev/ar0s1a	16G	186M	15G	1%	/flash
/dev/ar0s1e	140G	857M	128G	1%	/var

```
root@ns#
```

13\ Type the following command to run the ns_hw_err.bash script, which checks for latent hardware errors: root@ns# ns_hw_err.bash

Example

```
root@ns# ns_hw_err.bash

NetScaler NS10.5: Build 59.71.nc, Date: Mar  2 2016, 05:22:19

platform: serial KZCHGCN810

platform: sysid 250140 - NSMPX-14000-40G 12*CPU+16*F1X+4*F4X+2*E1K+2*CVM
N3

HDD MODEL: ar0: 227328MB &lt;Intel MatrixRAID RAID1&gt; status: READY

Generating the list of newslog files to be processed...

Generating the events from newslog files...

Checking for HDD errors...

Checking for HDD SMART errors...

Checking for Flash errors...

/var/nslog/dmesg.boot:* DEVELOPER mode - run NetScaler manually! *

/var/tmp/ns_hw_tmp_10831//dmesg-a.out:* DEVELOPER mode - run NetScaler
manually! *

*****

FOUND          2 Flash errors: DEVELOPER mode - run NetScaler manually

*****

Checking for Mega Raid Controller errors...
```

```
Checking for SSL errors...

Checking for BIOS errors...

Checking for SMB errors...

Checking for MotherBoard errors...

Checking for CMOS errors...

    License year: 2016: OK

License server failed at startup. Check /var/log/license.log

Vendor daemon failed at startup. Check /var/log/license.log

Checking for SFP/NIC errors...

Checking for Firmware errors...

Checking for License errors...

Checking for Undetected CPUs...

Checking for DIMM flaps...

Checking for LOM errors...

Checking the Power Supply Errors...

root@ns#
```

14\ **Important:** Physically disconnect all ports except the LOM port, including the management port, from the network.

15\ At the shell prompt, switch to the NetScaler command line. Type: exit

Example

```
root@ns# exit

logout
```

Done

16\ Run the following command to shut down the appliance. You are asked if you want to completely stop the NetScaler. Type: shutdown -p now

Example

```
root@ns#

root@ns# exit


logout

Done
```

Upgrade the appliance

To upgrade the appliance, follow these steps:

1. Power off the NetScaler® appliance.
2. Locate two solid-state drives (SSDs) on the back of the appliance in slot #1 and slot #2, as shown in the following figure:



3\ Verify that the replacement solid-state drives (SSDs) are the ones required for your NetScaler model. The conversion requires minimum of four SSDs. The NetScaler label is on the top of one of

the solid-state drives, which is pre-populated with a new version of BIOS and a recent build of the required NetScaler SDX Management Service. This SSD must be installed in slot # 1.

4\.. Remove the SSDs by pushing the safety latch of the drive cover down while pulling the drive handle.

5\.. On the new NetScaler certified SSD drive, open the drive handle completely to the left, and then insert the new drive into the slot #1 as far as possible.

6\.. To seat the drive, close the handle flush with the rear side of the appliance so that the drive locks securely into the slot.

Important: The orientation of the SSD is important. When you insert the drive, make sure that the NetScaler product label is at the side.

7\.. Insert a second NetScaler certified SSD, which matches the capacity of the SSD in slot #1, in slot # 2.

Note: If the license of your appliance is 14040 40G, 14060 40G, 14080 40G, insert more blank NetScaler certified SSDs in slots #3, #4, #5, and #6.

NetScaler SDX Model	Included Virtual Instances	Platform Maximum	SSDs included on base model	Extra SSDs for max instances
SDX 14020 40G	5	25	Two 240 GB; Two 300 GB	Two 300 GB
SDX 14040 40G, SDX 14060 40G, SDX 14080 40G	25	25	Two 240 GB; Four 300 GB	N/A

Important: Mixing and matching of old and new SSDs is not supported. SSDs in slot #1 and slot #2, which constitute the first RAID pair (local storage), must be of the same size and type. Similarly, SSDs in slot #3 and slot #4, which constitute the second RAID pair (VPX storage), must be of the same size and type. Use only drives that are part of the provided conversion kit.

8\.. Disconnect all network cables from the data ports and the management ports.

9\.. Start the NetScaler appliance. For instructions, see “Switching on the Appliance” in Installing the Hardware.

The conversion process can run for approximately 30 minutes, during which you must not power cycle the appliance. The entire conversion process might not be visible on the console and might appear to be unresponsive.

The conversion process updates the BIOS, installs the Citrix Hypervisor and the Management Service Operating system, and copies the NetScaler VPX image to the SSD for instance provisioning, and forms the Raid1 pair.

Note: The serial number of the appliance remains the same.

10\.. Keep the console cable attached during the conversion process. Allow the process to complete, at which point the netscaler-sdx login: prompt appears.

11\.. During the conversion process the LOM port connection might be lost as it resets the IP address to the default value of 192.168.1.3. The conversion status output is available on the VGA monitor.

12\.. The default credentials are now changed from nsroot/nsroot to root/nsroot once the appliance is converted from an MPX to SDX.

13\.. To make sure that the conversion is successful, verify that the FVT result indicates success. Run the following command: `tail /var/log/fvt/fvt.log`

Reconfigure the converted appliance

After the conversion process, the appliance no longer has its previous working configuration. Therefore, you can access the appliance through a web browser only by using the default IP address: 192.168.100.1/16. Configure a computer on network 192.168.0.0 and connect it directly to the appliance's management port (0/1) with a cross-over Ethernet cable, or access the NetScaler SDX appliance through a network hub by using a straight through Ethernet cable. Use the default credentials to log on, and then do the following:

1. Select the Configuration tab.
2. Verify that the System Resource section displays the accurate number of CPU cores, SSL cores, and the total memory for your NetScaler SDX appliance.
3. Select the **System** node and, under **Set Up Appliance**, click **Network Configuration** to modify network information of the Management Service.
4. In the **Modify Network Configuration** dialog box, specify the following details:
 - **Interface***—The interface through which clients connect to the Management Service. Possible values: 0/1, 0/2. Default: 0/1.

- Citrix Hypervisor IP Address*—The IP address of Citrix Hypervisor.
- Management Service IP Address*—The IP address of the Management Service.
- Netmask*—The subnet mask for the subnet in which the SDX appliance is located.
- Gateway*—The default gateway for the network.
- DNS Server—The IP address of the DNS server.

*A mandatory parameter

1. Click OK. Connection to the Management Service is lost as the network information was changed.
2. Connect the NetScaler SDX appliance's management port 0/1 to a switch to access it through the network. Browse to the IP address used earlier and log on with the default credentials.
3. Apply the new licenses. For instructions, see [SDX Licensing Overview](#).
4. Navigate to **Configuration > System** and, in the **System Administration** group, click **Reboot Appliance**. Click **Yes** to confirm. You are now ready to provision the VPX instances on the NetScaler SDX appliance. For instructions, see [Provisioning NetScaler instances](#).

Convert a NetScaler MPX 8900 appliance to a NetScaler SDX 8900 appliance

You can convert a NetScaler MPX appliance to a NetScaler SDX appliance by upgrading the software through a new solid-state drive (SSD). NetScaler supplies a field conversion kit to migrate a NetScaler MPX appliance to a NetScaler SDX appliance.

NOTE

Citrix recommends that you configure the lights out management (LOM) port of the NetScaler appliance before starting the conversion process. For more information on the LOM port of the NetScaler appliance, see [Lights out management port of the NetScaler appliance](#).

To convert a NetScaler MPX appliance to a NetScaler SDX appliance, you must access the appliance through a console cable attached to a computer or terminal. Before connecting the console cable, configure the computer or terminal to support the following configuration:

- VT100 terminal emulation
- 9600 baud
- 8 data bits

- 1 stop bit
- Parity and flow control set to NONE

Connect one end of the console cable to the RS232 serial port on the appliance, and the other end to the computer or terminal.

NOTE


To use a cable with an RJ-45 converter, insert the optional converter into the console port and attach the cable to it.

With the cable attached, verify that the MPX appliance's components are functioning correctly. You are then ready to begin the conversion. The conversion process modifies the BIOS, installs a Citrix Hypervisor and a Service Virtual Machine image, and copies the NetScaler VPX image to the SSD.

After the conversion process, you modify the appliance's configuration and apply a new license. You can then provision the VPX instances through the Management Service on what is now a NetScaler SDX appliance.

The following figure shows the front panel of the MPX 89xx.

Figure 1. NetScaler MPX 89xx front panel



To verify proper operation of the MPX appliance's components

1. Access the console port and enter the administrator credentials.
2. Run the following command from the command line interface of the appliance to display the serial number: **show hardware**

The serial number might be helpful in case you want to contact Citrix® technical support.

Example

```
> show hardware
```

```
Platform: NSMPX-8900 8*CPU+4*F1X+6*E1K+1*E1K+1*COL 8955 30010
```

```
Manufactured on: 12/3/2018
```

```
CPU: 2100MHZ
```

```
Host Id: 1862303878
```

```
Serial no: JVFUJCZT1E
```

```
Encoded serial no: JVFUJCZT1E
```

```
BMC Revision: 4.51
```

```
Done
```

3. Run the following command to display the status of the active 1G and 10G interfaces: **show interface**
4. In the show interface command's output, verify that all the interfaces are enabled and the status of every interface is shown as UP/UP.

NOTE

If you do not have an SFP+ transceiver for every port, verify the interfaces in stages. After checking the first set of interfaces, unplug the SFP+ transceivers and plug them in to the next set of ports. The SFP+ transceivers are not hot-swappable. Therefore, restart the MPX appliance after you connect the transceivers.

5. Run the following commands for each of the interfaces that are not in the UP/UP state:
 - **enable interface 1/x**
 - **enable interface 10/x**

where x is the new interface number.

6. Run the following command to verify that the status of the power supplies is normal: **stat system -detail**

Example

```
> stat system -detail
```

NetScaler Executive View

System Information:

Up since Tue Dec 4 14:01:49 2018

Memory usage (MB) 859

InUse Memory (%) 4.81

Number of CPUs 5

System Health Statistics (Standard):

CPU 0 Core Voltage (Volts) 1.78

CPU 1 Core Voltage (Volts) 0.00

Main 3.3 V Supply Voltage 3.28

Standby 3.3 V Supply Voltage 3.28

+5.0 V Supply Voltage 4.90

+12.0 V Supply Voltage 11.81

Battery Voltage (Volts) 3.02

Intel CPU Vtt Power(Volts) 0.00

5V Standby Voltage(Volts) 5.05

Voltage Sensor2(Volts) 0.00

CPU Fan 0 Speed (RPM) 6900

CPU Fan 1 Speed (RPM) 6700

System Fan Speed (RPM) 6800

System Fan 1 Speed (RPM) 6800

System Fan 2 Speed (RPM)	6900
CPU 0 Temperature (Celsius)	44
CPU 1 Temperature (Celsius)	0
Internal Temperature (Celsius)	37
Power supply 1 status	NORMAL
Power supply 2 status	NORMAL
Power supply 3 status	NOT SUPPORTED
Power supply 4 status	NOT SUPPORTED
System Disk Statistics:	
/flash Size (MB)	16858
/flash Used (MB)	323
/flash Available (MB)	15186
/flash Used (%)	2
/var Size (MB)	143802
/var Used (MB)	1880
/var Available (MB)	130418
/var Used (%)	1
System Health Statistics(Auxiliary):	
Voltage 0 (Volts)	1.20
Voltage 1 (Volts)	1.20
Voltage 2 (Volts)	0.00
Voltage 3 (Volts)	0.00

Voltage 4 (Volts)	0.00
Voltage 5 (Volts)	0.00
Voltage 6 (Volts)	0.00
Voltage 7 (Volts)	0.00
Fan 0 Speed (RPM)	7000
Fan 1 Speed (RPM)	0
Fan 2 Speed (RPM)	0
Fan 3 Speed (RPM)	0
Temperature 0 (Celsius)	28
Temperature 1 (Celsius)	34
Temperature 2 (Celsius)	0
Temperature 3 (Celsius)	0
Done	

7. Run the following command to generate a tar of system configuration data and statistics: **show techsupport**

Example

```
> show techsupport

showtechsupport data collector tool - $Revision: #13 $!

NetScaler version 11.1

Creating /var/tmp/support ....

The NS IP of this box is 10.221.44.30
```

```
This is not HA configuration

Copying selected configuration files ....

Running shell commands ....

Running CLI show commands ....

Collecting ns running configuration....

Collecting running gslb configuration....

Running CLI stat commands ....

Running vtysh commands ....

Copying newnslog files ....

Copying core files from /var/core ....

Copying core files from /var/crash ....

Copying GSLB location database files ....

Copying GSLB auto sync log files ....

Copying Safenet Gateway log files ....

Copying messages, ns.log, dmesg and other log files ....

Creating archive ....

/var/tmp/support/support.tgz ---- points to ---> /var/tmp/support/
collector_P_10.221.44.30_4Dec2018_14_14.tar.gz

Done
```


NOTE

The output of the command is available in the `/var/tmp/support/collector_<IP_address>_P_<date>.tar.gz` file. Copy this file to another computer for future reference. The output of the command might be helpful in case you want to contact Citrix technical support.

8. At the NetScaler command line interface, switch to the shell prompt. Type `shell`
9. Run the following command to verify that 270 MB of RAM is reserved for shared memory:

```
root@ns# dmesg | grep memory
```

Example

```
root@ns# dmesg | grep memory

real memory  = 36507222016 (34816 MB)

avail memory = 32728735744 (31212 MB)

NS-KERN nsppe_rendezvous: NSPPE-02 on CPU3NS-KERN map_shared_mem_ioctl
(cpu 2, NSPPE-01): Reserving 270 MB for shared memory type 0

root@ns#
```

10. Run the following command to verify that the appliance has 12 CPU cores: `root@ns# dmesg | grep cpu`

Example

```
root@ns# dmesg | grep cpu

cpu0 (BSP): APIC ID:  0

cpu1 (AP): APIC ID:  2

cpu2 (AP): APIC ID:  4

cpu3 (AP): APIC ID:  6
```

```
cpu4 (AP): APIC ID: 8

cpu5 (AP): APIC ID: 10

cpu6 (AP): APIC ID: 12

cpu7 (AP): APIC ID: 14

cpu0: <ACPI CPU> on acpi0

cpu1: <ACPI CPU> on acpi0

cpu2: <ACPI CPU> on acpi0

cpu3: <ACPI CPU> on acpi0

cpu4: <ACPI CPU> on acpi0

cpu5: <ACPI CPU> on acpi0

cpu6: <ACPI CPU> on acpi0

cpu7: <ACPI CPU> on acpi0

est0: <Enhanced SpeedStep Frequency Control> on cpu0

p4tcc0: <CPU Frequency Thermal Control> on cpu0

est1: <Enhanced SpeedStep Frequency Control> on cpu1

p4tcc1: <CPU Frequency Thermal Control> on cpu1

est2: <Enhanced SpeedStep Frequency Control> on cpu2

p4tcc2: <CPU Frequency Thermal Control> on cpu2

est3: <Enhanced SpeedStep Frequency Control> on cpu3

p4tcc3: <CPU Frequency Thermal Control> on cpu3

est4: <Enhanced SpeedStep Frequency Control> on cpu4

p4tcc4: <CPU Frequency Thermal Control> on cpu4
```

```

est5: <Enhanced SpeedStep Frequency Control> on cpu5

p4tcc5: <CPU Frequency Thermal Control> on cpu5

est6: <Enhanced SpeedStep Frequency Control> on cpu6

p4tcc6: <CPU Frequency Thermal Control> on cpu6

est7: <Enhanced SpeedStep Frequency Control> on cpu7

p4tcc7: <CPU Frequency Thermal Control> on cpu7

NS-KERN nsppe_rendezvous: NSPPE-02 on CPU3NS-KERN map_shared_mem_ioctl
(cpu 2, NSPPE-01): Reserving 270 MB for shared memory type 0

root@ns#

```

11. Run the following command to verify that the /var drive is mounted as /dev/ad0s1e: `root@ns# df -h`

Example

```

root@ns# df -h

```

Filesystem	Size	Used	Avail	Capacity	Mounted on
/dev/md0	354M	342M	5M	99%	/
devfs	1.0k	1.0k	0B	100%	/dev
procfs	4.0k	4.0k	0B	100%	/proc
/dev/ad0s1a	16G	323M	14G	2%	/flash
/dev/ad0s1e	140G	1.9G	127G	1%	/var

```

root@ns#

```

12. Run the `ns_hw_err.bash` script, which checks for latent hardware errors: `root@ns# /netscaler/ns_hw_err.bash`

Example

```
root@ns# ns_hw_err.bash

NetScaler NS11.1: Build 60.5.nc, Date: Oct  3 2018, 10:58:21

platform: serial JVFUJCZT1E

platform: sysid 30010 - NSMPX-8900 8*CPU+4*F1X+6*E1K+1*E1K+1*COL 8955

HDD MODEL: ad0: 228936MB <MICRON M510DC MTFDDAK240MBP 0013> at ata0-
master UDMA33


Generating the list of newslog files to be processed...

Generating the events from newslog files...

Checking for HDD errors...

Checking for HDD SMART errors...

Checking for Flash errors...

    skipping flash check because HDD and flash are same physical
device: ad0.

    Please refer to HDD error output.

Checking for Mega Raid Controller errors...

Checking for SSL errors...

Checking for BIOS errors...

Checking for SMB errors...

Checking for MotherBoard errors...
```

```
Checking for CMOS errors...

    License year: 2018: OK

License server failed at startup. Check /var/log/license.log

Vendor daemon failed at startup. Check /var/log/license.log

Checking for SFP/NIC errors...

Checking for Firmware errors...

Checking for License errors...

Checking for Undetected CPUs...

Checking for DIMM flaps...

Checking for LOM errors...

Checking the Power Supply Errors...

Checking for Hardware Clock errors...

*****

NS hardware check: Found 2 errors

*****

Script Done.

root@ns#
```

13. **Important:** Physically disconnect all ports except the LOM port, including the management port, from the network.
14. At the shell prompt, switch to the NetScaler command line. Type: **exit**
15. Run the following command to shut down the appliance: `shutdown -p now`

Example

```
> shutdown -p now
```

```
Are you sure you want to completely stop NetScaler (Y/N)? [N]:y
```


Upgrade the appliance

The upgrade process involves the following two steps:


- Replace the SSD
- Start the appliance

Replace the SSD

1. Locate the solid-state drive on the back panel of the appliance, as shown in the following figure:



2. Verify that the replacement solid-state drive (SSD) is the one required for your NetScaler model. The NetScaler label is on the top of the SSD. The SSD is pre-populated with a new version of the BIOS and a recent build of the required Service VM software.
3. Remove the currently installed SSD drive by pushing the safety latch of the drive cover to the right and pulling out the drive by its handle.
4. Open the drive handle on the new drive completely to the left, and insert the drive into the slot. The following figure shows the drive partially inserted. Push the drive all the way into the slot.



1. Close the handle flush with the rear side of the appliance so that the drive locks securely into the slot.

IMPORTANT

The orientation of the solid-state drive is important. When you insert the drive, make sure that the NetScaler® product label is at the top.

Start the appliance

1. Start the NetScaler appliance. For instructions, see the Switching on the Appliance section in [Installing the hardware](#).

The conversion process takes approximately 30 minutes to complete. The conversion process performs the following actions:

- Updates the BIOS.
- Installs the Citrix Hypervisor™ and the Service VM operating system.
- Copies the NetScaler VPX image to the solid state drive for instance provisioning.

When the conversion begins, the LCD screen on the front bezel indicates NSMPX-8900, as shown in the following figure.

```
![SDX 8900 before conversion](/content/dam/netscaler/us/en/netscaler-hardware-platforms/sdx/media/89xx-before-conversion.png)
```

When the conversion is successful, the LCD indicates Citrix NSSDX - 8900, as shown in the following figure.

```
![SDX 8900 after conversion](/content/dam/netscaler/us/en/netscaler-hardware-platforms/sdx/media/89xx-after-conversion.png)
```

> ****Note****

>

> The serial number of the appliance remains the same.

1. Keep the console cable attached during the conversion process. Allow the process to complete, at which point the netscaler-sdx login: prompt appears.
2. When the appliance finishes the conversion process, it no longer has the previously working configuration. Therefore, you can access the appliance through a Web browser only. Use the default IP address: 192.168.100.1/16. Configure a computer on network 192.168.0.0 and connect it directly to the management port 0/1 of the appliance by using a cross-over Ethernet cable. Alternately, access the appliance through a network hub by using a straight-through Ethernet cable. Use the default credentials.
3. Select the **Configuration** tab.
4. Verify that the System Resource section displays CPU cores, SSL cores, and total memory for the NetScaler SDX appliance.
5. Select System node and click the Network Configuration link on the System page to modify the IP address of the Service VM.
6. In the Modify Network Configuration dialog box, specify the following details:
 - Interface—The interface through which clients connect to the Management Service. Possible values: 0/1, 0/2. Default: 0/1.
 - Citrix Hypervisor IP Address—The IP address of the Citrix Hypervisor.
 - Management Service IP Address—The IP address of the Management Service.
 - Netmask—The subnet mask for the subnet in which the SDX appliance is located.
 - Gateway—The default gateway for the network.
 - DNS Server—The IP address of the DNS server. *An optional parameter
7. Click **OK**.
8. Connect the NetScaler SDX appliance to a switch to access it through the network. Browse to the Management Service IP and log on with the default credentials.
9. For instructions for applying the licenses, see [NetScaler SDX Licensing Overview](#).

Convert a NetScaler MPX 15000 appliance to a NetScaler SDX 15000 appliance

You can convert a NetScaler MPX appliance to an NetScaler SDX appliance by upgrading the software through a new solid-state drive (SSD). NetScaler supplies a field conversion kit to migrate a NetScaler MPX appliance to an SDX appliance.

The conversion requires all eight SSDs.

NOTE

Citrix recommends that you configure the lights out management (LOM) port of the appliance before starting the conversion process. For more information on the LOM port of the ADC appliance, see [Lights out management port of the NetScaler SDX appliance](#).

To convert an MPX appliance to an SDX appliance, you must access the appliance through a console cable attached to a computer or terminal. Before connecting the console cable, configure the computer or terminal to support the following configuration:

- VT100 terminal emulation
- 9600 baud
- 8 data bits
- 1 stop bit
- Parity and flow control set to NONE

Connect one end of the console cable to the RS232 serial port on the appliance, and the other end to the computer or terminal.

NOTE

To use a cable with an RJ-45 converter, insert the optional converter into the console port and attach the cable to it.

Citrix® recommends you to connect a VGA monitor to the appliance to monitor the conversion process, because the LOM connection is lost during the conversion process.

With the cable attached, verify that the MPX appliance's components are functioning correctly. You are then ready to begin the conversion. The conversion process modifies the BIOS, installs a Citrix Hypervisor and a Service Virtual Machine image, and copies the NetScaler VPX image to the Solid State Drive.

The conversion process also sets up a redundant array of independent disks (RAID) controller for local storage and NetScaler VPX storage. SSD slots #1 and #2 are used for local storage and SSD slots #3 and #4 are used for NetScaler VPX storage.

After the conversion process, modify the appliance's configuration and apply a new license. You can then provision the VPX instances through the Management Service on what is now an NetScaler SDX appliance.

Verify proper operation of the MPX appliance's components

1. Access the console port and enter the administrator credentials.
2. Run the following command from the command line interface of the appliance to display the serial number: `show hardware`. You might need the serial number to log on to the appliance after the conversion.

Example

```
> show hardware

Platform: NSMPX-15000-50G 16*CPU+128GB+4*MLX(50)+8*F1X+2*E1K+2*2-CHIP
COL 520400

Manufactured on: 9/13/2017

CPU: 2100MHZ

Host Id: 1862303878

Serial no: 4VCX9CUFN6

Encoded serial no: 4VCX9CUFN6

Netscaler UUID: d9de2de3-dc89-11e7-ab53-00e0ed5de5aa

BMC Revision: 5.56

Done
```

The serial number might be helpful when you want to contact Citrix technical support.

3. Run the following command to display the status of the active interfaces:

```
show interface
```

Example

```
> show interface
```

```
1) Interface 0/1 (Gig Ethernet 10/100/1000 Mbits) #4
```

```
flags=0xc020 <ENABLED, UP, UP, autoneg, HAMON, HEARTBEAT, 802.1q>
```

```
MTU=1500, native vlan=1, MAC=0c:c4:7a:e5:3c:50, uptime 1h08m02s
```

```
Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,
```

```
throughput 0
```

```
Actual: media UTP, speed 1000, duplex FULL, fctl OFF, throughput  
1000
```

```
LLDP Mode: NONE,
```

```
LR Priority: 1024
```

```
RX: Pkts(19446) Bytes(1797757) Errs(0) Drops(19096) Stalls(0)
```

```
TX: Pkts(368) Bytes(75619) Errs(0) Drops(0) Stalls(0)
```

```
NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)
```

```
Bandwidth thresholds are not set.
```

```
2) Interface 0/2 (Gig Ethernet 10/100/1000 Mbits) #5
```

```
flags=0x4000 <ENABLED, DOWN, down, autoneg, HEARTBEAT, 802.1q>
```

```
MTU=1500, native vlan=1, MAC=0c:c4:7a:e5:3c:51, downtime 1h08m15s
```

```
Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,
```

```
throughput 0
```

```
LLDP Mode: NONE,
```

```
LR Priority: 1024
```

```
RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
```

```
TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
```

```
NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)
```

```
Bandwidth thresholds are not set.
```

3) Interface 10/1 (10G Ethernet) #6

```
flags=0x4000 <ENABLED, DOWN, down, autoneg, HAMON, HEARTBEAT,
802.1q>
```

```
MTU=1500, native vlan=200, MAC=00:e0:ed:5d:e5:76, downtime 1h08m15s
```

```
Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,
```

```
throughput 0
```

```
LLDP Mode: NONE,
```

```
LR Priority: 1024
```

```
RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
```

```
TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
```

```
NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)
```

```
Bandwidth thresholds are not set.
```

```
Rx Ring: Configured size=2048, Actual size=512, Type: Elastic
```

4) Interface 10/2 (10G Ethernet) #7

```
flags=0x4000 <ENABLED, DOWN, down, autoneg, HEARTBEAT, 802.1q>
```

```
MTU=1500, native vlan=200, MAC=00:e0:ed:5d:e5:77, downtime 1h08m15s
```

```
Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,  
throughput 0
```

```
LLDP Mode: NONE, LR Priority: 1024
```

```
RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
```

```
TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
```

```
NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)
```

```
Bandwidth thresholds are not set.
```

```
Rx Ring: Configured size=2048, Actual size=512, Type: Elastic
```

5) Interface 10/3 (10G Ethernet) #8

```
flags=0x4000 <ENABLED, DOWN, down, autoneg, HEARTBEAT, 802.1q>
```

```
MTU=1500, native vlan=1, MAC=00:e0:ed:5d:e5:78, downtime 1h08m15s
```

```
Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,  
throughput 0
```

```
LLDP Mode: NONE, LR Priority: 1024
```

```
RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
```

```
TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
```

```
NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)
```

```
Bandwidth thresholds are not set.
```

```
Rx Ring: Configured size=2048, Actual size=512, Type: Elastic
```

6) Interface 10/4 (10G Ethernet) #9

flags=0x4000 <ENABLED, DOWN, down, autoneg, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:5d:e5:79, downtime 1h08m15s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE,

LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic

7) Interface 10/5 (10G Ethernet) #0

flags=0x4000 <ENABLED, DOWN, down, autoneg, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:5d:e5:aa, downtime 1h08m15s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE,

LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic

8) Interface 10/6 (10G Ethernet) #1

flags=0x4000 <ENABLED, DOWN, down, autoneg, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:5d:e5:ab, downtime 1h08m15s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE,

LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic

9) Interface 10/7 (10G Ethernet) #2

flags=0x4000 <ENABLED, DOWN, down, autoneg, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:5d:e5:ac, downtime 1h08m15s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE,

LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic

10) Interface 10/8 (10G Ethernet) #3

flags=0x4000 <ENABLED, DOWN, down, autoneg, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=00:e0:ed:5d:e5:ad, downtime 1h08m15s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE,

LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

Rx Ring: Configured size=2048, Actual size=512, Type: Elastic

11) Interface 50/1 (50G Ethernet) #13

flags=0x4000 <ENABLED, DOWN, down, autoneg, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=24:8a:07:a3:1f:84, downtime 1h08m22s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE,

LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

12) Interface 50/2 (50G Ethernet) #12

flags=0x4000 <ENABLED, DOWN, down, autoneg, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=24:8a:07:a3:1f:6c, downtime 1h08m22s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE,

LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

13) Interface 50/3 (50G Ethernet) #11

flags=0x4000 <ENABLED, DOWN, down, autoneg, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=24:8a:07:a3:1f:98, downtime 1h08m22s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE,

LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)

Bandwidth thresholds are not set.

14) Interface 50/4 (50G Ethernet) #10

flags=0x4000 <ENABLED, DOWN, down, autoneg, HEARTBEAT, 802.1q>

MTU=1500, native vlan=1, MAC=24:8a:07:94:b9:b6, downtime 1h08m22s

Requested: media AUTO, speed AUTO, duplex AUTO, fctl OFF,

throughput 0

LLDP Mode: NONE,

LR Priority: 1024

RX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)

```
TX: Pkts(0) Bytes(0) Errs(0) Drops(0) Stalls(0)
```

```
NIC: InDisc(0) OutDisc(0) Fctls(0) Stalls(0) Hangs(0) Muted(0)
```

```
Bandwidth thresholds are not set.
```

15) Interface L0/1 (Netscaler Loopback interface) #14

```
flags=0x20008020 <ENABLED, UP, UP>
```

```
MTU=1500, native vlan=1, MAC=0c:c4:7a:e5:3c:50, uptime 1h08m18s
```

```
LLDP Mode: NONE, LR Priority: 1024
```

```
RX: Pkts(5073645) Bytes(848299459) Errs(0) Drops(0) Stalls(0)
```

```
TX: Pkts(9923625) Bytes(968741778) Errs(0) Drops(0) Stalls(0)
```

```
Bandwidth thresholds are not set.
```

Done

4. In the show interface command's output, verify that all the interfaces are enabled and the status of every interface is shown as UP/UP.

NOTES:

5. Run the following command for each of the interfaces that are not in the UP/UP state:

```
> enable interface 50/1
```

Done

```
> enable interface 50/2
```

Done

```
> enable interface 50/3
```

```
Done
```

```
> enable interface 50/4
```

```
Done
```

Where x is the new interface number

6. Run the following command to verify that the status of the power supplies is normal:

```
stat system -detail
```

Example

```
> stat system -detail
```

```
NetScaler Executive View
```

```
System Information:
```

```
Up since          Sat Dec  5 04:17:29 2020
```

```
Up since(Local)  Sat Dec  5 04:17:29 2020
```

```
Memory usage (MB)                4836
```

```
InUse Memory (%)                 4.08
```

```
Number of CPUs                   13
```

```
System Health Statistics (Standard):
```

```
CPU 0 Core Voltage (Volts)       1.80
```

```
CPU 1 Core Voltage (Volts)       1.80
```

Main 3.3 V Supply Voltage	3.35
Standby 3.3 V Supply Voltage	3.23
+5.0 V Supply Voltage	5.00
+12.0 V Supply Voltage	12.06
Battery Voltage (Volts)	3.02
Intel CPU Vtt Power(Volts)	0.00
5V Standby Voltage(Volts)	4.95
Voltage Sensor2(Volts)	0.00
CPU Fan 0 Speed (RPM)	3500
CPU Fan 1 Speed (RPM)	3600
System Fan Speed (RPM)	3600
System Fan 1 Speed (RPM)	3600
System Fan 2 Speed (RPM)	3500
CPU 0 Temperature (Celsius)	37
CPU 1 Temperature (Celsius)	47
Internal Temperature (Celsius)	26
Power supply 1 status	NORMAL
Power supply 2 status	NORMAL
Power supply 3 status	NOT SUPPORTED
Power supply 4 status	NOT SUPPORTED

System Disk Statistics:

/flash Size (MB)	23801
/flash Used (MB)	7009
/flash Available (MB)	14887
/flash Used (%)	32
/var Size (MB)	341167
/var Used (MB)	56502
/var Available (MB)	257371
/var Used (%)	18

System Health Statistics(Auxiliary):

Voltage 0 (Volts)	1.20
Voltage 1 (Volts)	1.20
Voltage 2 (Volts)	1.20
Voltage 3 (Volts)	1.20
Voltage 4 (Volts)	1.54
Voltage 5 (Volts)	0.00
Voltage 6 (Volts)	0.00
Voltage 7 (Volts)	0.00
Fan 0 Speed (RPM)	3600
Fan 1 Speed (RPM)	0
Fan 2 Speed (RPM)	0
Fan 3 Speed (RPM)	0

```
Temperature 0 (Celsius)           24
Temperature 1 (Celsius)           30
Temperature 2 (Celsius)           0
Temperature 3 (Celsius)           0
Done
```

7. Run the following command to generate a tar of system configuration data and statistics:

```
show techsupport
```

Example

```
> show techsupport

showtechsupport data collector tool - $Revision$!

NetScaler version 13.0

The NS IP of this box is 10.217.206.43

This is not HA configuration

Copying selected configuration files ....

Running shell commands ....

Running CLI show commands ....

Collecting ns running configuration....

Collecting running gslb configuration....

Running CLI stat commands ....

Running vtysh commands ....
```

```
Copying newnslog files ....

Copying core files from /var/core ....

Copying core files from /var/crash ....

Copying GSLB location database files ....

Copying GSLB auto sync log files ....

Copying Safenet Gateway log files ....

Copying messages, ns.log, dmesg and other log files ....

Creating archive ....

/var/tmp/support/support.tgz ---- points to ---> /var/tmp/support/
collector_P_10.217.206.43_5Dec2020_05_32.tar.gz

showtechsupport script took 1 minute(s) and 17 second(s) to execute.

Done
```

NOTE

The output of the command is available in the `/var/tmp/support/collector\_IP\_address\_P\_date.tar.gz` file. Copy this file to another computer for future reference. The output of the command might be helpful when you want to contact Citrix technical support.

- At the command line interface, switch to the shell prompt. Type: `shell`

Example

```
> shell

Copyright (c) 1992-2013 The FreeBSD Project.
```



```
Copyright (c) 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993,  
1994
```

```
The Regents of the University of California. All rights reserved.
```

```
root@ns#
```

9. Run the following command to verify the number of Cavium cards available depending upon your appliance:

```
root@ns# grep "memory" /var/nslog/dmesg.boot
```

Example

```
root@ns# grep "memory" /var/nslog/dmesg.boot
```

```
real memory = 139586437120 (133120 MB)
```

```
avail memory = 132710871040 (126562 MB)
```

```
root@ns#
```

10. Run the following command to verify the number of CPU cores depending upon your appliance:

```
root@ns# grep "cpu" /var/nslog/dmesg.boot
```

Example

```
root@ns# grep "cpu" /var/nslog/dmesg.boot
```

```
cpu0 (BSP): APIC ID: 0
```

```
cpu1 (AP): APIC ID: 2
```

```
cpu2 (AP): APIC ID: 4
```

```
cpu3 (AP): APIC ID: 6
```

```
cpu4 (AP): APIC ID: 8
```

```
cpu5 (AP): APIC ID: 10

cpu6 (AP): APIC ID: 12

cpu7 (AP): APIC ID: 14

cpu8 (AP): APIC ID: 16

cpu9 (AP): APIC ID: 18

cpu10 (AP): APIC ID: 20

cpu11 (AP): APIC ID: 22

cpu12 (AP): APIC ID: 24

cpu13 (AP): APIC ID: 26

cpu14 (AP): APIC ID: 28

cpu15 (AP): APIC ID: 30

cpu0: <ACPI CPU> on acpi0

cpu1: <ACPI CPU> on acpi0

cpu2: <ACPI CPU> on acpi0

cpu3: <ACPI CPU> on acpi0

cpu4: <ACPI CPU> on acpi0

cpu5: <ACPI CPU> on acpi0

cpu6: <ACPI CPU> on acpi0

cpu7: <ACPI CPU> on acpi0

cpu8: <ACPI CPU> on acpi0

cpu9: <ACPI CPU> on acpi0

cpu10: <ACPI CPU> on acpi0
```

```
cpu11: <ACPI CPU> on acpi0
cpu12: <ACPI CPU> on acpi0
cpu13: <ACPI CPU> on acpi0
cpu14: <ACPI CPU> on acpi0
cpu15: <ACPI CPU> on acpi0

est0: <Enhanced SpeedStep Frequency Control> on cpu0
p4tcc0: <CPU Frequency Thermal Control> on cpu0

est1: <Enhanced SpeedStep Frequency Control> on cpu1
p4tcc1: <CPU Frequency Thermal Control> on cpu1

est2: <Enhanced SpeedStep Frequency Control> on cpu2
p4tcc2: <CPU Frequency Thermal Control> on cpu2

est3: <Enhanced SpeedStep Frequency Control> on cpu3
p4tcc3: <CPU Frequency Thermal Control> on cpu3

est4: <Enhanced SpeedStep Frequency Control> on cpu4
p4tcc4: <CPU Frequency Thermal Control> on cpu4

est5: <Enhanced SpeedStep Frequency Control> on cpu5
p4tcc5: <CPU Frequency Thermal Control> on cpu5

est6: <Enhanced SpeedStep Frequency Control> on cpu6
p4tcc6: <CPU Frequency Thermal Control> on cpu6

est7: <Enhanced SpeedStep Frequency Control> on cpu7
p4tcc7: <CPU Frequency Thermal Control> on cpu7

est8: <Enhanced SpeedStep Frequency Control> on cpu8
```

```

p4tcc8: <CPU Frequency Thermal Control> on cpu8

est9: <Enhanced SpeedStep Frequency Control> on cpu9

p4tcc9: <CPU Frequency Thermal Control> on cpu9

est10: <Enhanced SpeedStep Frequency Control> on cpu10

p4tcc10: <CPU Frequency Thermal Control> on cpu10

est11: <Enhanced SpeedStep Frequency Control> on cpu11

p4tcc11: <CPU Frequency Thermal Control> on cpu11

est12: <Enhanced SpeedStep Frequency Control> on cpu12

p4tcc12: <CPU Frequency Thermal Control> on cpu12

est13: <Enhanced SpeedStep Frequency Control> on cpu13

p4tcc13: <CPU Frequency Thermal Control> on cpu13

est14: <Enhanced SpeedStep Frequency Control> on cpu14

p4tcc14: <CPU Frequency Thermal Control> on cpu14

est15: <Enhanced SpeedStep Frequency Control> on cpu15

p4tcc15: <CPU Frequency Thermal Control> on cpu15

root@ns#

```

11. Run the following command to verify that the /var drive is mounted as `/dev/ ar0s1a:`

```
root@ns# df -h
```

Example

```

root@ns# df -h

Filesystem      Size      Used    Avail Capacity  Mounted on
/dev/md0        422M      404M      9.1M      98%        /

```

```

devfs          1.0k    1.0k      0B   100%    /dev
procfs         4.0k    4.0k      0B   100%    /proc
/dev/ar0s1a     23G     6.9G     14G   32%     /flash
/dev/ar0s1e    333G     32G     274G   10%     /var

root@ns#

```

12. Type the following command to run the `ns_hw_err.bash` script, which checks for latent hardware errors: `root@ns`

```
# ns_hw_err.bash
```

Example

```

root@ns# ns_hw_err.bash

NetScaler NS13.0: Build 71.3602.nc, Date: Nov 12 2020, 07:26:41   (64-bit)

platform: serial 4VCX9CUFN6

platform: sysid 520400 - NSMPX-15000-50G
16*CPU+128GB+4*MLX(50)+8*F1X+2*E1K+2*2-CHIP COL 8955

HDD MODEL: ar0: 434992MB <Intel MatrixRAID RAID1> status: READY

Generating the list of newslog files to be processed...

Generating the events from newslog files...

Checking for HDD errors...

Checking for HDD SMART errors...

Checking for Flash errors...

Checking for Mega Raid Controller errors...

```

```
Checking for SSL errors...

Dec  5 06:00:31 <daemon.err> ns monit[996]: 'safenet_gw' process is not
running

Checking for BIOS errors...

Checking for SMB errors...

Checking for MotherBoard errors...

Checking for CMOS errors...

License year: 2020: OK

Checking for SFP/NIC errors...

Dec  5 06:02:32 <daemon.err> ns monit[996]: 'safenet_gw' process is not
running

Checking for Firmware errors...

Checking for License errors...

Checking for Undetected CPUs...

Checking for DIMM flaps...

Checking for Memory Channel errors...

Checking for LOM errors...

Checking the Power Supply Errors...

Checking for Hardware Clock errors...

Script Done.

root@ns#
```

13. **Important:** Physically disconnect all ports except the LOM port, including the management port, from the network.

14. At the shell prompt, switch to the ADC command line. Type: `exit`

Example

```
root@ns# exit

logout

Done
```

15. Run the following command to shut down the appliance. You are asked if you want to completely stop the ADC. Type: `shutdown -p now`

Example

```
> shutdown -p now

Are you sure you want to completely stop NetScaler (Y/N)? [N]:y

Done

> Dec  5 06:09:11 <auth.notice> ns shutdown: power-down by root:

Dec  5 06:09:13 <auth.emerg> ns init: Rebooting via init mechanism

Dec  5 06:09:13 <syslog.err> ns syslogd: exiting on signal 15

Dec  5 06:09:13 aslearn[1662]: before pthread_join(), task name:
Aslearn_Packet_Loop_Task

Dec  5 06:09:15 aslearn[1662]: Exiting function ns_do_logging

Dec  5 06:09:15 aslearn[1662]: before pthread_join(), task name:
Aslearn_WAL_Cleanup_Task

Dec  5 06:09:15 aslearn[1662]: before pthread_join(), task name:
Aslearn_HA_Primary_Task

Dec  5 06:09:15 aslearn[1662]: 1662 exiting gracefully

Dec  5 06:09:18 [1672]: nsnet_tcpipconnect: connect() failed; returned
-1 errno=61
```

```
qat0: qat_dev0 stopped 12 acceleration engines

pci4: Resetting device

qat1: qat_dev1 stopped 12 acceleration engines

pci6: Resetting device

qat2: qat_dev2 stopped 12 acceleration engines

pci132: Resetting device

qat3: qat_dev3 stopped 12 acceleration engines

pci134: Resetting device

Dec  5 06:09:33 init: some processes would not die; ps axl advised

reboot initiated by init with parent kernel

Waiting (max 60 seconds) for system process `vnlrn' to stop...done

Waiting (max 60 seconds) for system process `bufdaemon' to stop...done

Waiting (max 60 seconds) for system process `syncer' to stop...

Syncing disks, vnodes remaining...0 0 0 done

All buffers synced.

Uptime: 1h53m18s

ixl_shutdown: lldp start 0

ixl_shutdown: lldp start 0

ixl_shutdown: lldp start 0

ixl_shutdown: lldp start 0

usb0: Controller shutdown

uhub0: at usb0, port 1, addr 1 (disconnected)
```



```
usb0: Controller shutdown complete

usb1: Controller shutdown

uhub1: at usb1, port 1, addr 1 (disconnected)

ugen1.2: <vendor 0x8087> at usb1 (disconnected)

uhub3: at uhub1, port 1, addr 2 (disconnected)

ugen1.3: <FTDI> at usb1 (disconnected)

uftdi0: at uhub3, port 1, addr 3 (disconnected)

ugen1.4: <vendor 0x1005> at usb1 (disconnected)

umass0: at uhub3, port 3, addr 4 (disconnected)

(da0:umass-sim0:0:0:0): lost device - 0 outstanding, 0 refs

(da0:umass-sim0:0:0:0): removing device entry

usb1: Controller shutdown complete

usb2: Controller shutdown

uhub2: at usb2, port 1, addr 1 (disconnected)

    ugen2.2: <vendor 0x8087> at usb2 (disconnected)

uhub4: at uhub2, port 1, addr 2 (disconnected)

ugen2.3: <vendor 0x0557> at usb2 (disconnected)

uhub5: at uhub4, port 7, addr 3 (disconnected)

ugen2.4: <vendor 0x0557> at usb2 (disconnected)

ukbd0: at uhub5, port 1, addr 4 (disconnected)

ums0: at uhub5, port 1, addr 4 (disconnected)

usb2: Controller shutdown complete
```

```

ixl_shutdown: lldp start 0

ixl_shutdown: lldp start 0

ixl_shutdown: lldp start 0

ixl_shutdown: lldp start 0


acpi0: Powering system off

```

Upgrade the appliance

To upgrade the appliance, follow these steps:

1. Power off the ADC appliance.
2. Locate two solid-state drives (SSDs) on the back of the appliance in slot #1 and slot #2, as shown in the following figure:



3. Verify that the replacement solid-state drives (SSDs) are the ones required for your ADC model. The conversion requires minimum of four SSDs. The NetScaler label is on the top of one of the SSDs. The SSD is pre-populated with a new version of the BIOS and a recent build of the required NetScaler SDX Management Service. This SSD must be installed in slot # 1.
4. Remove the SSDs by pushing the safety latch of the drive cover down while pulling the drive handle.
5. On the new NetScaler® certified SSD drive, open the drive handle completely to the left. Then insert the new drive into the slot #1 as far as possible.

6. To seat the drive, close the handle flush with the rear side of the appliance so that the drive locks securely into the slot.

Important: The orientation of the SSD is important. When you insert the drive, make sure that the NetScaler product label is at the side.

7. Insert a second NetScaler certified SSD, which matches the capacity of the SSD in slot #1, in slot # 2.

Note: If the license of your appliance is 14040 40G, 14060 40G, 14080 40G, insert more blank NetScaler certified SSDs in slots #3, #4, #5, and #6.

NetScaler SDX Model	Included Virtual Instances	Platform Maximum	SSDs included on base model	Extra SSDs for max instances
SDX 15020/ SDX 15020-50G	5	55	Two 240 GB RAID-supported removable boot solid-state drives (SSDs), (slots 1 and 2). Two 240 GB RAID-supported removable storage repositories (slots 3 and 4 paired) SSDs, and four 480 GB storage repositories (slots 5–6 paired and 7–8 paired) SSDs.	NA
SDX 15030/ SDX 15030-50G	20	55	Two 240 GB RAID-supported removable	NA

			boot solid-state drives (SSDs), (slots 1 and 2). Two 240 GB RAID-supported removable storage repositories (slots 3 and 4 paired) SSDs, and four 480 GB storage repositories (slots 5–6 paired and 7–8 paired) SSDs.	
--	--	--	---	--

IMPORTANT

Mixing and matching of old and new SSDs is not supported. SSDs in slot #1 and slot #2, which constitute the first RAID pair (local storage), must be of the same size and type. Similarly, SSDs in slot #3 and slot #4, which constitute the second RAID pair (VPX storage), must be of the same size and type. Use only drives that are part of the provided conversion kit.

8. Disconnect all network cables from the data ports and the management ports.
9. Start the ADC appliance. For instructions, see “Switching on the Appliance” in Installing the Hardware.

The conversion process can run for approximately 30 minutes, during which you must not power cycle the appliance. The entire conversion process might not be visible on the console and might appear to be unresponsive.

The conversion process updates the BIOS, installs the Citrix Hypervisor and the Management Service Operating system. It also copies the NetScaler VPX image to the SSD for instance provisioning, and forms the Raid1 pair.

Note: The serial number of the appliance remains the same.

10. Keep the console cable attached during the conversion process. Allow the process to complete, at which point the SDX login: prompt appears.
11. During the conversion process the LOM port connection might be lost as it resets the IP address to the default value of 192.168.1.3. The conversion status output is available on the VGA monitor.
12. The default credentials for the Citrix Hypervisor are changed to root/nsroot after the appliance is converted from an MPX to SDX. If this password does not work, try typing nsroot/the serial number of the appliance. The serial number bar code is available at the back of the appliance, and is also available in the output of the `show hardware` command.
13. To make sure that the conversion is successful, verify that the FVT result indicates success. Run the following command: `tail /var/log/fvt/fvt.log`

Reconfigure the converted appliance

After the conversion process, the appliance no longer has its previous working configuration. Therefore, you can access the appliance through a web browser only by using the default IP address: 192.168.100.1/16. Configure a computer on network 192.168.0.0 and connect it directly to the appliance's management port (0/1) with a cross-over Ethernet cable. Alternately, access the NetScaler SDX appliance through a network hub by using a straight through Ethernet cable. Use the default credentials to log on, and then do the following:

1. Select the Configuration tab.
2. Verify that the System Resource section displays the accurate number of CPU cores, SSL cores, and the total memory for your NetScaler SDX appliance.
3. Select the **System** node and, under **Set Up Appliance**, click **Network Configuration** to modify the network information of the Management Service.
4. In the **Modify Network Configuration** dialog box, specify the following details:
 - **Interface***—The interface through which clients connect to the Management Service. Possible values: 0/1, 0/2. Default: 0/1.
 - **Citrix Hypervisor™ IP Address***—The IP address of Citrix Hypervisor.
 - **Management Service IP Address***—The IP address of the Management Service.
 - **Netmask***—The subnet mask for the subnet in which the SDX appliance is located.
 - **Gateway***—The default gateway for the network.
 - **DNS Server**—The IP address of the DNS server.

*A mandatory parameter

5. Click **OK**. Connection to the Management Service is lost as the network information was changed.
6. Connect the NetScaler SDX appliance's management port 0/1 to a switch to access it through the network. Browse to the IP address used earlier and log on with the default credentials.
7. Apply the new licenses. For instructions, see [SDX Licensing Overview](#).
8. Navigate to **Configuration > System** and, in the **System Administration** group, click **Reboot Appliance**. Click **Yes** to confirm. You are now ready to provision the VPX instances on the NetScaler SDX appliance. For instructions, see [Provisioning NetScaler instances](#).

Migrate the configuration of a NetScaler instance across SDX appliances

Before you migrate the NetScaler instance configuration to a new SDX appliance, you must first set up the new SDX appliance.

The following table lists the device names involved in the migration process:

TERM	DESCRIPTION
SDX-OLD-A1, SDX-OLD-A2	Name of the old SDX platform
SDX-NEW-A1, SDX-NEW-A2	Name of the new SDX platform
VPX-OLD-A1, VPX-OLD-A2	Name of the VPX instance hosted on the old SDX platform
VPX-NEW-A1, VPX-NEW-A2	Name of the VPX instance hosted on the new SDX platform

The following images explain the migration process, showing the current, interim, and final states for the High Availability (HA) configurations.

Figure 1. Current state



Figure 2. Interim state





Figure 3. Desired state



The following table explains the color legends used in the diagram:

COLOR LEGENDS	DESCRIPTION
Green	Represents the platforms and instances in their current or desired state.
Orange	Represents the platforms and instances undergoing transition or pending migration.

Prerequisites

- Upgrade the new SDX appliance (SDX-NEW-A2) to the same build version as the older one. If the build on the old SDX appliance is not supported on the new SDX, first upgrade the older SDX appliance to a build that is supported on the new SDX.
- In the new SDX appliance (SDX-NEW-A2), complete the [initial configuration](#), including setting up usernames, passwords, IP addresses, and gateway configurations.
- If any specific configurations exist on the old SDX appliances that need to be replicated, use the Management Service to configure settings such as channel creation, mtu settings, and admin profiles on the new SDX appliance.

- Provision a NetScaler instance (VPX-NEW-A2) on the new SDX appliance (SDX-NEW-A2) with the same build and configuration as the older one. For more information, see [Provision NetScaler instance on SDX appliance](#).

Supported migration scenarios

This migration applies to the following scenarios:

- High Availability (HA) configuration
- Standalone deployments
- Migrating to an identical SDX platform
- Migrating to a different SDX platform

NOTES:

- This migration process doesn't support migrating from SDX to third-party hypervisor hosted VPX instances.
- These outlined steps are validated for migrating SDX platforms from SDX 14000 to 15000, SDX 14000 to 16000, and SDX 22000 to 15000. The same steps are also applicable to other SDX platforms.

The different scenarios are explained below:

Case 1: No channel and same interface speed and numbering

If the old VPX instance (VPX-OLD-A2) is configured with interfaces 10/1, 10/2, ..., 10/n, the new VPX instance (VPX-NEW-A2) must replicate this configuration precisely.

Case 2: No channel and different interface speed or numbering

If the old VPX instance (VPX-OLD-A2) is configured with interfaces 10/1 and 10/2, and the new VPX instance (VPX-NEW-A2) is configured with interfaces 10/3 and 10/4. Alternatively, the new VPX instance (VPX-NEW-A2) might use interfaces 25/1 and 25/2.

Case 3: Same channel numbering

If the old VPX instance (VPX-OLD-A2) is configured with channels LA/1, LA/2, ..., LA/n, then the new VPX instance (VPX-NEW-A2) must replicate this configuration precisely. The underlying member interface numbering or speed does not affect this requirement.

Case 4: Different channel numbering

If the old VPX instance (VPX-OLD-A2) is configured with channels LA/1 and LA/2, while the new VPX instance (VPX-NEW-A2) is configured with channels LA/3 and LA/4.

Phases in migration

Phase 1: Break HA (if applicable) and remove the secondary VPX

If HA is enabled, you must break the HA configuration and remove the secondary VPX (VPX-OLD-A2) from the HA setup.

1. Break the HA configuration

- a. Log on to the primary VPX instance (VPX-OLD-A1) CLI.
- b. To set the HA status on VPX-OLD-A1 to stay as the primary node, type:

```
set ha node -hastatus STAYPRIMARY
```

- c. To disable the HA sync status on the primary node (VPX-OLD-A1), type:

```
set ha node -haSync DISABLED
```

2. Remove the secondary VPX from the HA pair

- a. Log on to the secondary VPX instance (VPX-OLD-A2) CLI.
- b. To set the HA status on VPX-OLD-A2 to stay as the secondary node, type:

```
set ha node -hastatus STAYSECONDARY
```

- c. Remove VPX-OLD-A2 from the HA pair to break the ha configuration. Identify the node ID for VPX-OLD-A2 and run the following command:

```
rm ha node <node_id>
```

Phase 2: Form an HA pair between the old primary VPX and the new VPX

To establish an HA pair between the old primary VPX (VPX-OLD-A1) and the new VPX (VPX-NEW-A2):

1. On the CLI of the old primary VPX (VPX-OLD-A1), type:

```
add ha node 1 <IP address of VPX-NEW-A2>
```

2. On the CLI of the new VPX (VPX-NEW-A2), type:

```
add ha node 1 <IP address of VPX-OLD-A1>
```

Phase 3: Migrate VPX configuration

On the primary VPX (VPX-OLD-A1), enable HA sync using the following command:

```
set ha node -haSync ENABLED
```

Phase 4 (optional): Additional steps required for case 2 and case 4

Perform the following steps to address specific scenarios in case 2 (no channel and different interface speed or numbering) and case 4 (different channel numbering) migration.

On the new secondary VPX (VPX-NEW-A2):

1. Check the log file at `/var/nssynclog/sync_batch_status.log` for errors.
2. The interface-related settings might generate errors. Manually fix the `bind vlan` and `set interface` configurations.
3. If any other errors are found, manually address them.
4. Save the configuration using either the `save config` or `save ns config -all` command.

NOTES:

Before you initiate the traffic failover to the new VPX on the new SDX appliance, ensure the following:

- Verify that HA sync is enabled.

- Validate the configuration of the current secondary VPX (VPX-NEW-A2) using the command `show ns runningconfig`.
- Check for errors in the log file located at `/var/nssynclog/sync_batch_status.log` on the secondary node.

Phase 5: Force failover to migrate the traffic to the new VPX on the new SDX appliance

To initiate a forced failover to shift traffic to the new VPX (VPX-NEW-A2) hosted on the new SDX appliance, perform the following steps:

1. On VPX-OLD-A1, to enable the HA pair, type:

```
set ha node -haStatus ENABLED
```

Here, VPX-OLD-A1 becomes active in the HA pair.

2. To force the failover of traffic to the new VPX (VPX-NEW-A2), type:

```
force ha failover
```

VPX-NEW-A2 becomes the primary node handling the incoming traffic.

In the transition from the previous SDX platform, SDX-OLD-A2, to the new SDX platform, SDX-NEW-A2, the VPX-NEW-A2 on SDX-NEW-A2 is now the primary node and manages traffic. These steps are also applicable to replace SDX-OLD-A1 with the new SDX platform, SDX-NEW-A1.

Troubleshooting

Following is the list of troubleshooting articles for NetScaler SDX hardware issues:

1. Troubleshooting NetScaler Power Supply Issues: <https://support.citrix.com/article/CTX202340>
2. Troubleshooting NetScaler Small Form Pluggable (SFP) Issues: <https://support.citrix.com/article/CTX236074>
3. Troubleshooting NetScaler SSL Card Issues: <https://support.citrix.com/article/CTX236094>
4. Troubleshooting NetScaler LOM GUI issues: https://support.citrix.com/s/article/CTX477557-mpxsdx-series-910016000-gui-lom-access-to-configuration-settings-are-readonly?language=en_US

NetScaler SDX hardware document history

The following table describes important changes to the NetScaler hardware documentation. To receive notifications about these updates, subscribe to the RSS feed. To subscribe to the NetScaler SDX hardware document history, see [NetScaler SDX document history](#).

DATE	CHANGE	DESCRIPTION
19 Feb 2025	Updated the LOM GUI steps	See Obtaining the MAC address, serial number, and host properties of the appliance
23 Aug 2024	Added a link to an article about troubleshooting LOM issues	See Troubleshooting
02 Aug 2024	Updates to the procedure for replacing a RAID-supported SSD	See To replace a RAID-supported SSD
22 Jul 2024	New topic for creating and installing an SSL certificate on SDX 9100 and SDX 16000	See Create and install an SSL certificate on SDX 9100 and SDX 16000
20 Jun 2024	Removed topics for old platforms.	
30 May 2024	Updated the LOM version for some platforms and removed platforms that have reached EOL.	See LOM support matrix
12 Jan 2024	Two power supplies are shipped with SDX 9100 for	See NetScaler SDX 9100

DATE	CHANGE	DESCRIPTION
	hardware shipped after 13th Dec, 2023	
26 Sep 2023	Added support for updated NIC driver and firmware on NetScaler SDX with 10G/25G/40G NICs	See Fortville NIC firmware upgrade on NetScaler SDX appliances .
13 Sep 2023	Replaced nsroot with nsrecover in the initial configuration using the serial console section	See Initial Configuration through the serial console .
12 Jun 2023	Added a link to transceivers in the landing page	See Hardware form factors
29 May 2023	Updated LOM support matrix for SDX platforms	See LOM support matrix
20 Sep 2022	Support for SDX 16000 platform	See NetScaler SDX 16000
15 Jun 2022	Updated the hardware-software compatability matrix for SDX 9100	See NetScaler SDX hardware-software compatibility matrix
11 Oct 2021	Updated the hardware-software compatability matrix to include 13.1	See NetScaler SDX hardware-software compatibility matrix
15 Jul 2021	New topic about hardware health attributes	See Hardware health attributes of NetScaler SDX appliances
07 Dec 2020	New topic for converting MPX 15000 to SDX 15000	See Convert a NetScaler MPX 15000 appliance to a NetScaler SDX 15000 appliance
01 Aug 2020	Updates for CA password	See Initial configuration

DATE	CHANGE	DESCRIPTION
29 Jul 2020	Updated for LA channels on management interfaces	See Initial configuration
22 Jul 2020	Updated for 11.1-65.x support	See Fortville NIC firmware upgrade on NetScaler SDX appliances
24 Jun 2020	Updated for VPX instance to be upgraded before SDX	See Fortville NIC firmware upgrade on NetScaler SDX appliances
16 Jun 2020	Updated for USB port reserved for a future release	Updates to multiple topics
10 Jun 2020	New topic for Fortville NIC firmware upgrade	Fortville NIC firmware upgrade on NetScaler SDX appliances
28 May 2020	Added support for iKVM (HTML5) to access the LOM GUI	See Configure the network settings on the LOM port
09 Apr 2020	Added a note that SDX to MPX conversion is not supported	See Converting a NetScaler MPX appliance to a NetScaler SDX appliance
21 Aug 2019	Added new topics for SDX 15000 and SDX 26000	See SDX 15000 and SDX 26000
23 May 2019	Updated support for release 13.0 on all hardware platforms	See NetScaler SDX hardware-software compatibility matrix

fix

This is the content for the sample heading topic. It is hidden from the main table of contents but remains accessible via links.