

System

Oct 13, 2015

System

The following topics provide information of the NetScaler system.

Administration	Manages and monitors the NetScaler appliance by using built-in features such as authentication and authorization, SNMP management, audit logging, web server logging, NTP management, and Reporting tool.
AppFlow	Provides information about the reporting capabilities of AppFlow feature of the NetScaler.
AutoScale	Describes how users of Citrix CloudPlatform can use the AutoScale feature on the NetScaler appliance to enable automatic scaling of their applications.
CloudBridge Connector	Provides help in reducing the cost of moving your applications to the cloud, reduce the risk of application failure, and increase network efficiency in your cloud environment.
Clustering	A setup of multiple nCore NetScaler appliances that ensure high availability, high throughput, and scalability of a deployment of NetScaler appliances.
EdgeSight Monitoring for NetScaler	Monitors the end-user experience with web applications that are served in a NetScaler environment.
Flex Tenancy	A methodology that allows you to tune a group of NetScaler virtual appliance instances to the unique characteristics and needs of individual applications in a complex Web 2.0 setup.
High Availability	A setup of two NetScaler appliances that ensure the high availability of NetScaler appliances.
Web Interface	Provides access to Citrix XenApp and Citrix XenDesktop applications.

Basic Operations

Any changes you make to the configuration of a NetScaler appliance are not saved automatically. You have to save the settings manually. When an appliance is restarted, it loads the latest saved configuration.

This document includes the following details:

- [Viewing and Saving Configurations](#)
- [Clearing the NetScaler Configuration](#)

Viewing and Saving Configurations

Configurations are stored in the /nsconfig/ns.conf directory. For configurations to be available across sessions, you must save the configuration after every configuration change.

To view the running configuration by using the command line interface

At the command prompt, type:

```
show ns runningConfig
```

To view the running configuration by using the configuration utility

- Navigate to System > Diagnostics and, in the View Configuration group, click Running Configuration.

To find the difference between two configuration files by using the command line interface

At the command prompt, type:

```
diff ns config <configfile1> <configfile2>
```

To find the difference between two configuration files by using the configuration utility

- Navigate to System > Diagnostics and, in the View Configuration group, click Configuration difference.

To save configurations by using the command line interface

At the command prompt, type:

```
save ns config
```

To save configurations by using the configuration utility

On the Configuration tab, in the top-right corner, click the Save icon.

To view the saved configurations by using the command line interface

At the command prompt, type:

```
show ns ns.conf
```

To view the saved configurations by using the configuration utility

Navigate to System > Diagnostics and, in the View Configuration group, click Saved Configuration.

Clearing the NetScaler Configuration

You have the following three options for clearing the NetScaler configuration.

Basic level. Clearing your configuration at the basic level clears all settings except the following:

- NSIP, MIP(s), and SNIP(s)
- Network settings (Default Gateway, VLAN, RHI, NTP, and DNS settings)
- HA node definitions
- Feature and mode settings
- Default administrator password (nsroot)

Extended level. Clearing your configuration at the extended level clears all settings except the following:

- NSIP, MIP(s), and SNIP(s)
- Network settings (Default Gateway, VLAN, RHI, NTP, and DNS settings)
- HA node definitions

Feature and mode settings revert to their default values.

Full level. Clearing your configuration at the full level returns all settings to their factory default values. However, the NSIP and default gateway are not changed, because changing them could cause the appliance to lose network connectivity.

To clear the configuration by using the command line interface

At the command prompt, type:

```
clear ns config -force <level>
```

Example: To forcefully clear the basic configurations on an appliance.

```
clear ns config -force basic
```

To clear the configuration by using the configuration utility

- Navigate to System > Diagnostics and, in the Maintenance group, click Clear Configuration and select the configuration level to be cleared from the appliance.

Configuring Clock Synchronization

You can configure your NetScaler appliance to synchronize its local clock with a Network Time Protocol (NTP) server. This ensures that its clock has the same date and time settings as the other servers on your network.

You can configure clock synchronization on your appliance by adding NTP server entries to the `ntp.conf` file from either the configuration utility or the command line interface, or by manually modifying the `ntp.conf` file and then starting the NTP daemon (NTPD). The clock synchronization configuration does not change if the appliance is restarted, upgraded, or downgraded. However, the configuration does not get propagated to the secondary NetScaler in a high availability setup.

Note: If you do not have a local NTP server, you can find a list of public, open access, NTP servers at the official NTP site, <http://www.ntp.org>, under Public Time Servers List. Before configuring your NetScaler to use a public NTP server, be sure to read the Rules of Engagement page (link included on all Public Time Servers pages).

This document includes the following details:

- [Setting Up Clock Synchronization](#)
- [Starting the NTP Daemon](#)
- [Configuring Clock Synchronization Manually](#)

Setting Up Clock Synchronization

Updated: 2014-08-18

To configure clock synchronization, you must add NTP servers and then enable NTP synchronization.

To add an NTP server by using the command line interface

At the command prompt, type the following commands to add an NTP server and verify the configuration:

- `add ntp server (<serverIP> | <serverName>) [-minpoll <positive_integer>] [-maxpoll <positive_integer>]`
- `show ntp server`

Example

```
> add ntp server 10.102.29.30 -minpoll 6 -maxpoll 11
```

To configure an NTP server by using the configuration utility

Navigate to System > NTP Servers, and create the NTP server.

Starting the NTP Daemon

Updated: 2014-08-08

When you enable NTP synchronization, the NetScaler starts the NTP daemon and uses the NTP server entries in the `ntp.conf` file to synchronize its local time setting. If you do not want to synchronize the appliance time with the other servers in the network, you can disable NTP synchronization, which stops the NTP daemon (NTPD).

To enable NTP synchronization by using the command line interface

At the command prompt, type one of the following commands:

```
enable ntp sync
```

To enable NTP synchronization by using the configuration utility

Navigate to System > NTP Servers, click Action and select NTP Synchronization.

Configuring Clock Synchronization Manually

Updated: 2013-08-23

You can configure clock synchronization manually by logging on to the NetScaler and editing the `ntp.conf` file.

To enable clock synchronization on your NetScaler by modifying the `ntp.conf` file

1. Log on to the command line interface.
2. Switch to the shell prompt.
3. Copy the `/etc/ntp.conf` file to `/nsconfig/ntp.conf`, unless the `/nsconfig` directory already contains an `ntp.conf` file.

4. Check the /nsconfig/ntp.conf file for the following entries and, if they are present, remove them:

restrict localhost

restrict 127.0.0.2

5. Add the IP address for the desired NTP server to the /nsconfig/ntp.conf file, beneath the file's server and restrict entries.

Note: For security reasons, there should be a corresponding restrict entry for each server entry.

6. If the /nsconfig directory does not contain a file named rc.netscaler, create the file.

7. Add the following entry to /nsconfig/rc.netscaler: /usr/sbin/ntpd -c /nsconfig/ntp.conf -l /var/log/ntpd.log &

This entry starts the ntpd service, checks the ntp.conf file, and logs messages in the /var/log directory.

This process runs every time the NetScaler is restarted.

8. Reboot the NetScaler to enable clock synchronization.

Note:

If you want to start the time synchronization process without restarting the NetScaler, run the following command from the shell prompt:

```
/usr/sbin/ntpd -c /nsconfig/ntp.conf -l /var/log/ntpd.log &
```

Configuring System Session Timeout

A session timeout interval is provided to restrict the time duration for which a session (GUI, CLI, or API) remains active when not in use. For the NetScaler, the system session timeout can be configured at the following levels:

- **User level timeout.** Applicable to the specific user.

GUI	Navigate to System > User Administration > Users, select a user, and edit the user's timeout setting.
CLI	At the command prompt, enter the following command: set system user <name> -timeout <secs>

- **User group level timeout.** Applicable to all users in the group.

GUI	Navigate to System > User Administration > Groups, select a group, and edit the group's timeout setting.
CLI	At the command prompt, enter the following command: set system group <groupName> -timeout <secs>

- **Global system timeout.** Applicable to all users and users from groups who do not have a timeout configured.

GUI	Navigate to System > Settings, click Change global system settings, and update the timeout value as required.
CLI	At the command prompt, enter the following command: set system parameter -timeout <secs>

The timeout value specified for a user has the highest priority. If timeout is not configured for the user, the timeout configured for a member group is considered. If timeout is not specified for a group (or the user does not belong to a group), the globally configured timeout value is considered. If timeout is not configured at any level, the default value of 900 seconds is set as the system session timeout.

Additionally, you can specify timeout durations for each of the interfaces you are accessing. However, the timeout value specified for a specific interface is restricted to the timeout value configured for the user that is accessing the interface. For example, let us consider an user "publicadmin" who has a timeout value of 20 minutes. Now, when accessing an interface, the user must specify a timeout value that is within 20 minutes.

Note: You can choose to keep a check on the minimum and maximum timeout values by specifying the timeout as restricted (in CLI by specifying the restrictedTimeout parameter). This parameter is provided to account for previous NetScaler versions where the timeout value was not restricted.

- When enabled, the minimum configurable timeout value is 5 minutes (300 secs) and the maximum value is 1 day (86400 secs). If the timeout value is already configured to a value larger than 1 day, when this parameter is enabled, you are prompted to change it. If you do not change the value, the timeout value will automatically be reconfigured to the default timeout duration of 15 minutes (900 secs) on next reboot. The same will happen if the configured timeout value is less than 5 minutes.
- When disabled, the configured timeout durations are considered.

To configure the timeout duration at each interface:

NetScaler user interface	Timeout configuration
CLI	Specify the timeout value on the command prompt by using the following command: set cli mode -timeout <secs>
API	Specify the timeout value in the login payload.

Viewing the System Date and Time

To change the system date and time, you must use the shell interface to the underlying FreeBSD OS. However, to view the system date and time, you can use the command line interface or the configuration utility.

To view the system date and time by using the command line interface

At the command prompt, type:

```
show ns config
```

To view the system date and time by using the configuration utility

Navigate to System and select the System Information tab to view the system date.

Backing up and Restoring the NetScaler Appliance

You can back up the current state of a NetScaler appliance, and later use the backed up files to restore the appliance to the same state. You must use this feature before performing an upgrade or for precautionary reasons. A backup of a stable system enables you to restore the system to a stable point in the event that it becomes unstable.

Points to remember

- You cannot use the backup file taken from one appliance to restore a different appliance.
- You can back up and restore appliances in an HA setup, but make sure that you restore to the same appliance from which the backup file was created. For example, if the backup was taken from the primary appliance of the HA pair, when restoring make sure that the appliance you are restoring is the same appliance, even if it is no longer the primary appliance.
- You cannot perform the backup and restore operation on a NetScaler cluster.

This document includes the following details:

- [Backing up a NetScaler Appliance](#)
- [Restoring the NetScaler Appliance](#)

Backing up a NetScaler Appliance

Updated: 2015-03-11

Depending on the type of data to be backed up and the frequency at which you will create a backup, you can take a basic backup or a full backup.

- **Basic backup.** Backs up only configuration files. You might want to perform this type of backup frequently, because the files it backs up change constantly. The files that are backed up are:

Directory	Sub-Directory or Files
/nsconfig/	<ul style="list-style-type: none">• ns.conf• ZebOS.conf• rc.netscaler• snmpd.conf• nsbefore.sh• nsafter.sh• inetd.conf• ntp.conf• syslog.conf• newsyslog.conf• crontab• host.conf• hosts• ttys• sshd_config• httpd.conf• monitrc• rc.conf• ssh_config• localtime• issue• issue.net
/var/	<ul style="list-style-type: none">• download/*• log/wicmd.log• wi/tomcat/webapps/*• wi/tomcat/logs/*• wi/tomcat/conf/catalina/localhost/*• nslw.bin/etc/krb.conf• nslw.bin/etc/krb.keytab• netscaler/locdb/*• lib/likewise/db/*

	<ul style="list-style-type: none"> • vpn/bookmark/* • netscaler/crl • nstemplates/* • learnt_data/*
/netscaler/	<ul style="list-style-type: none"> • custom.html • vsr.htm

- o **Full backup.** In addition to the files that are backed up by a basic backup, a full backup backs up some less frequently updated files. The files that are backed up when using the full backup option are:

Directory	Sub-Directory or Files
/nsconfig/	<ul style="list-style-type: none"> • ssl/* • license/* • fips/*
/var/	<ul style="list-style-type: none"> • netscaler/ssl/* • wi/java_home/jre/lib/security/cacerts/* • wi/java_home/lib/security/cacerts/*

The backup is stored as a compressed TAR file in the /var/ns_sys_backup/ directory. To avoid issues due to non-availability of disk space, you can store a maximum of 50 backup files in this directory. You can use the rm system backup command to delete existing backup files so that you can create more backups.

Note:

- o While the backup operation is in progress, do not execute commands that affect the configuration.
- o If a file that is required to be backed up is not available, the operation skips that file.

To backup the NetScaler by using the NetScaler command line interface

At the command prompt, do the following:

1. Save the NetScaler configurations.

```
save ns config
```

2. Create the backup file.

```
create system backup [<fileName>] -level <basic | full> -comment <string>
```

Note: If the file name is not specified, the appliance creates a TAR file with the following naming convention: backup_<level>_<nsip_address>_<date-timestamp>.tgz.

Example: To backup the full appliance using the default naming convention for the backup file.

```
> create system backup -level full
```

3. Verify that the backup file was created.

```
show system backup
```

You can view properties of a specific backup file by using the fileName parameter.

To backup the NetScaler by using the configuration utility

Navigate to System > Backup and Restore, click Backup and then specify the details of the backup.

Restoring the NetScaler Appliance

Updated: 2014-09-11

When you restore the appliance from a backup file, the restore operation untars the backup file into the /var/ns_sys_backup/ directory. Once the untar operation is complete, the files are copied to their respective directories.

Attention: The restore operation does not succeed if the backup file is renamed or if the contents of the file are modified.

To restore the NetScaler by using the command line interface

At the command prompt, do the following:

1. Obtain a list of the backup files available on the appliance.

```
show system backup
```

2. Restore the appliance by specifying one of the backup files.

```
restore system backup -fileName <filename>
```

Example: To restore by using a full backup of an appliance.

```
> restore system backup -fileName backup_full_<nsip_address>_<date-timestamp>.tgz
```

3. Reboot the appliance.

```
reboot
```

To restore the NetScaler by using the configuration utility

Navigate to System > Backup and Restore, right-click the backup file to be restored and click Restore.

Restarting or Shutting down the Appliance

The NetScaler appliance can be remotely restarted or shut down from the available user interfaces. When a standalone NetScaler appliance is restarted or shut down, the unsaved configurations (configurations performed since the last `save ns config` command was issued) are lost.

In a high availability setup, when the primary appliance is rebooted/shut down, the secondary appliance takes over and becomes the primary. The unsaved configurations from the old primary are available on the new primary appliance.

You can also restart the appliance by only rebooting the NetScaler software and not rebooting the underlying operating system. This is called a warm reboot. For example, when you add a new license or change the NetScaler IP address, you can warm reboot the NetScaler appliance for these changes to take place.

Note: Warm reboot can be performed only on nCore appliances.

To restart the NetScaler by using the command line interface

At the command prompt, type:

```
reboot [-warm]
```

To restart the NetScaler by using the configuration utility

1. In the configuration utility, click Reboot on the home page of the Configuration tab.
2. When prompted to reboot, select Save configuration to make sure that you do not lose any configurations.

Note: You can perform a warm reboot by selecting Warm reboot.

To shut down the NetScaler by using the command line interface

At the command prompt, type:

- `shutdown -p now`: Shuts down the software and switches off the NetScaler. To restart NetScaler MPX, press the AC power switch. To Restart NetScaler VPX, restart the VPX instance.
- `shutdown -h now`: Shuts down the software and leaves the NetScaler switched on. Press any key to restart the NetScaler. This command does not switch off the NetScaler. Therefore, do not switch off the AC power or remove the AC power cables.

Note: The appliance cannot be shut down from the configuration utility.

Admin Partitioning

Note: The admin partitions feature is provided as a separate enhancement release that is based on NetScaler 10.5 Build 52.11. This feature is not available in the main or enhancement releases of NetScaler 10.5.

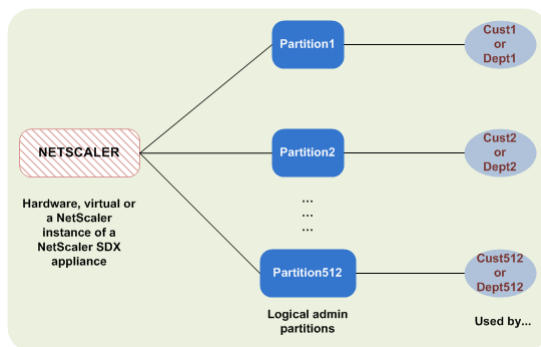
The NetScaler ADC provides an infrastructure called admin partitions that can be used to logically partition a NetScaler ADC.

Each admin partition:

- Has its own NetScaler configurations.
- Has its own administrators and users. Only users associated with a partition or system superuser can access and update the configurations.
- Uses a subset of NetScaler system resources such as bandwidth, connection pools, and memory.
- Handles traffic that is specific for that partition.

This means that each admin partition can function as a logical NetScaler ADC.

The following graphical representation shows a NetScaler ADC as a multi-tenant platform that can be used to service multiple customers, departments, or applications.



About the Admin Partitions Release

Review the release notes for the NetScaler 10.5 Enhancement (Admin Partitions) Build 52.1155.e release. This release is based on main release NetScaler 10.5 Build 52.11.

Release history:

- Build 52.1155.e (2014-11-18) (Current build)

This release notes includes information of all enhancement (Admin Partitions) builds of NetScaler 10.5 .

Note:

- This release notes does not document security fixes. To check if this build fixed any security issues, please refer to the security bulletin that is provided by the NetScaler security team.
- You can find the build in which the issue was provided by viewing the **[From Build xxxx]** label that is provided below each issue.

Points to Note

Some important aspects to keep in mind while using Build 52.1155.e.

- For large configurations, Citrix advises that you change the maximum memory limit from the default value of 10 MB.

[From Build 52.1155.e] [#504366]

- After adding an admin partition, make sure you save the configurations on the default partition. Otherwise, the partition setup configurations will be lost on system reboot.

[From Build 52.1155.e] [#493668]

- While creating an admin partition make sure that you first bind the user to the partition and then bind the required command policy to that user.

[From Build 52.1155.e] [#500821]

Fixed Issues

The issues addressed in Build 52.1155.e.

- In admin partitions, details of interfaces are are not displayed correctly.

[From Build 52.1155.e] [#511015, #512639]

- For all users, besides nsroot, failure in the execution of the "import responder htmlpage" command results in the "Done" message being displayed. This behavior is exhibited in default and non-default partitions.

[From Build 52.1155.e] [#506340]

- The NetScaler configuration utility incorrectly displays VLAN1, which cannot be bound to a partition. If you try to bind this VLAN, the following error is displayed: "Operation not permitted".

[From Build 52.1155.e] [#507394]

- Traffic domains can not be added by using the configuration utility.

[From Build 52.1155.e] [#512627]

- An invalid error message is being displayed when trying to add a VLAN.

[From Build 52.1155.e] [#513222, #512642, #512655]

- o The configuration utility throws an error when the values for minimum bandwidth, maximum bandwidth, and maximum connections are either set greater than the maximum supported value or lesser than the minimum supported value. Instead, the default value should have been set.

[From Build 52.1155.e] [#510293]

- o When deleting an admin partition, the NetScaler appliance can crash in some rarely occurring scenario.

[From Build 52.1155.e] [#504911, #505396]

- o Logging in to a NetScaler VPX too quickly after it has boot-up can cause the appliance to become unavailable.

[From Build 52.1155.e] [#510294]

- o When using the audit log feature, when deleting a partition, assert in hit when the NetScaler tries to transmit an NSB which is already freed. This issue is observed in the following scenarios:
 - When force failover is done on the primary node of HA setup.
 - When "clear config -f extended+" command was issued in default partition with 6 interfaces in a VPX and 2 user partitions configured with basic configs in it.

[From Build 52.1155.e] [#504101, #505766, #507193]

Known Issues

The known issues present in Build 52.1155.e.

- o In a NetScaler VPX with a large number of partitions, there can be lack of available disk space as newnslog files are created frequently due to rollover. This may result in the GUI and CLI becoming inaccessible.

[From Build 52.1155.e] [#515419]

- o RPCSVR services cannot be configured in admin partitions.

[From Build 52.1155.e] [#498477]

- o The GSLB configurations applied in the default partition can be viewed in admin partitions. This is not expected as user must not be able to view configurations that are defined in other partitions.

[From Build 52.1155.e] [#489512]

- o In a high availability (HA) deployment of NetScaler that has a high number of partitions (approximately 500) configured, the configuration synchronization might take more than 5 minutes.

[From Build 52.1155.e] [#515322]

- o Even though a partition operator cannot perform the add/rm/set/unset /bind/unbind/show operation for lb vserver, service, servicegroup, app flow collector, and cs policy, the configuration utility displays these operations.

[From Build 52.1155.e] [#507995]

- o The Surge protection feature cannot be configured in an admin partition. Since, surge protection parameters are part of the Change Global System Settings (System > Settings) dialog, when you try to update the global settings, the "Operation not supported" message is displayed.

[From Build 52.1155.e] [#498004]

- o When configuring an admin partition, in the list of channels (Network > Channels), the state of the LACP channel is incorrectly displayed. This issue is not present in the default partition.

[From Build 52.1155.e] [#517606, #518444]

- When a partition admin tries to perform the Download, Create, or Create Directory operation on the "Manage Certificate" screen, the operation not permitted error is displayed. The expected behavior is that the buttons must be disabled.

[From Build 52.1155.e] [#491353]

- In a high availability (HA) deployment of the NetScaler, if there have been more than two failovers, the load balancing persistency of previous sessions in admin partitions is not honored.

[From Build 52.1155.e] [#517023]

- When displaying the results of the "show lb monitor" command, the numbering of the user-defined monitors restarts from 1 instead of continuing the numbering from the list of built-in monitors.

[From Build 52.1155.e] [#511222]

- In both, default or admin partitions, when trying to import a password-protected key file, you get an error indicating that the key file is invalid. This error occurs because the NetScaler cannot import such key files.

[From Build 52.1155.e] [#512334]

- The maximum memory that can be configured for an admin partition is 2048 MB. Setting a value greater than this means that the value is automatically truncated to 2048 MB. This memory limit is per packet engine of the NetScaler.

[From Build 52.1155.e] [#504426]

- The Path MTU (PMTU) details are maintained at the global level instead of being maintained for each admin partition. This may cause NetScaler to use smaller segment size (MSS) to some hosts from all partitions even when some partitions may have routes which allow larger MSS.

[From Build 52.1155.e] [#502352]

- For a weblog client that is deployed on Linux, if there are more than 100 admin partitions, the partition names are not displayed properly.

[From Build 52.1155.e] [#518461]

- NetScaler might bypass TCP buffering for connections in non-default partitions since the memory pool is not initialized properly.

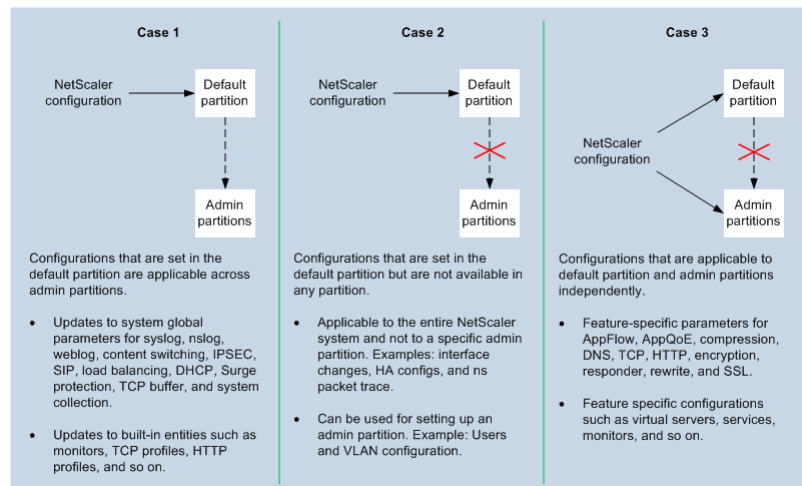
[From Build 52.1155.e] [#496522]

Types of Partitions

The NetScaler ADC ships with a single partition that is called the default partition. When the NetScaler is partitioned (when required to isolate the configurations or traffic), then, in addition to the default partition, the NetScaler can have one or more (maximum of 512) admin partitions. For details on default and admin partitions, see [table](#) that is provided below.

In a partitioned NetScaler, NetScaler configurations can be defined as follows:

- Global NetScaler configurations.** NetScaler configurations that are globally available (not restricted to specific partitions) must be defined in the default partition. Check Case 1 and Case 2 in the graphical illustration provided below.
- Partition-specific configurations.** NetScaler configurations that must be available on specific admin partitions can be defined on that admin partition. Check Case 3 in the graphical illustration provided below.



The following table provides details of default and admin partitions.

Â	Default partition	Admin partitions
Users	<p>By default, the following users have access to the default partition:</p> <ul style="list-style-type: none"> NetScaler superusers (including nsroot) NetScaler users who are not associated with a command policy that starts with <code>partition-</code>. 	<p>NetScaler superuser and other users who have been associated with that partition.</p> <p>Note: Partition users do not have shell access.</p>
File structure	<p>Uses the default file structure.</p> <p>For example, the NetScaler configuration files for default partition are stored in the <code>/nsconfig</code> directory. Similarly, log files are stored in the <code>/var/log/</code> directory.</p> <p>The same logic applies for SSL certificates for the default partition.</p>	<p>NetScaler configurations for an admin partition are stored in the <code>/nsconfig/partitions/<partitionName>/ns.conf</code> file.</p> <p>Other partition-specific files are stored in the <code>/var/partitions/<partitionName></code> directory.</p> <p>For example:</p> <ul style="list-style-type: none"> <code>/var/partitions/<partitionName>/download/</code> - To store downloaded files <code>/var/partitions/<partitionName>/log</code> - To store partition specific log files. Currently, however, logging is not supported at

		<p>partition-level. Therefore, this directory is empty and all logs are stored in the /var/log/ directory.</p> <ul style="list-style-type: none"> o /var/partitions/<partitionName>/netscaler/ssl - To store SSL CRL certificate related files
Resources available	All NetScaler resources.	NetScaler resources that are explicitly assigned to the partition.

Benefits and Uses of Admin Partitions

Admin partitions provide the following benefits:

- Allows delegation of administrative ownership of an application to the customer.
- Reduces the cost of ADC ownership without compromising on performance and ease-of-use.
- Safeguards from unwarranted configuration changes. In a non-partitioned NetScaler, authorized users of other application could intentionally or unintentionally change configurations that are required for your application. This could lead to undesirable behavior. This possibility is reduced in a partitioned NetScaler.
- Accelerates and allows to scale application deployments.
- Allows application-level or localized management and reporting.

Let us analyze a couple of scenarios and see how admin partitioning provides some of the benefits listed above.

Case 1 : To understand how an enterprise can benefit from this feature, consider a scenario faced by a company named Foo.com:

- Foo.com has a single NetScaler ADC.
- There are five departments and each department has one application that requires to be deployed with the NetScaler.
- Each application must be managed independently by a different set of users or administrators. Other users must be restricted from accessing the configurations.
- The application or back-end must be able to share resources like IP addresses.
- The global IT department must be able to control NetScaler-level settings which must be common to all partitions.
- Applications must be independent of one another. An error in configuration of one application must not affect the other.

A non-partitioned NetScaler would not be able to satisfy these requirements. However, you can achieve all these requirements by partitioning a NetScaler.

Simply create a partition for each of the applications, assign the required users to the partitions, specify the VLAN for each partition, and define global settings on the default partition.

Case 2 : To understand how a service provider can benefit from this feature, consider a scenario faced by a service provider named BigProvider:

- BigProvider has 5 customers: 3 small enterprises and 2 large enterprises.
- SmallBiz, SmallerBiz, and StartupBiz need only the most basic NetScaler functionality.
- BigBiz and LargeBiz are larger enterprises and have applications that attract a lot of traffic. They would like to use some of the more complex NetScaler functionality.

In a non-partitioned approach, the NetScaler administrator would typically use a NetScaler SDX appliance and provision a NetScaler instance for each customer. This solution would suit BigBiz and LargeBiz, because their applications need the undiminished power of the entire non-partitioned NetScaler. However, this solution might not be cost effective for servicing SmallBiz, SmallerBiz, and StartupBiz.

Therefore, BigProvider decides that admin partitioning is the way forward. They use NetScaler SDX appliance to bring up dedicated NetScaler instances for BigBiz and LargeBiz. They use a single NetScaler which is partitioned into three partitions, one each for SmallBiz, SmallerBiz, and StartupBiz.

The NetScaler administrator (superuser) then creates an admin partition for each of the small customers, specifies the users for the partitions, specifies the NetScaler resources for the partitions, and specifies the VLAN to be used by the traffic that is destined for each of the partitions.

NetScaler Feature-level Support in Admin Partitions

As described in [types of partitions](#), NetScaler configurations that are defined for individual admin partitions are called partition-specific configurations. There are some NetScaler features that cannot be defined in admin partitions, while being configurable in the default partition. The following table lists the configurations that are currently not partition-specific, that is, cannot be configured in admin partitions.

Note:

- Admin partitions cannot be set up in a NetScaler cluster.
- Subsequent release of the NetScaler appliance will provide support for these features.

Feature	NetScaler 10.5.e (special enhancement release)	NetScaler 11
Dynamic routing	No	Yes
Shared VLAN	No	No
Cluster	No	No
VXLAN	No	No
GRE tunnel	No	No
Load balancing for FTP, RTSP, TFTP, Diameter, SIP, RADIUS, RDP	No	No
SSL - FIPS	No	No
DNS - DNSSEC	No	No
Scriptable monitor	No	Yes
GSLB	No	No
Connection mirroring	No	No
FEO	No	No
nstrace	No	Yes
Sure connect	No	No
Priority queuing	No	No
HDOSP	No	No
Cache redirection	No	No
Integrated caching	No	Yes
Application firewall	No	No
AutoScale service group	No	No
AAA-TM	No	No
NetScaler Gateway	No	No

Partitioning a NetScaler

Note: Only superusers are authorized to partition a NetScaler.

Setting up an admin partition requires the following operations:

1. Create an admin partition and allocate the required NetScaler resources to that partition.
2. Specify the users that can access this partition.
3. Specify the level of authorization (using command policies) each user has in the partition.
4. Specify the configurations to isolate traffic across partitions.

After partitioning the NetScaler, make sure the users are made aware that their NetScaler configurations are now isolated from users who are not members of the partition.

Note:

- All configurations to set up an admin partition must be done from the default partition.
- Make sure the relevant users, command policies, VLANs, and bridgegroups are available on the NetScaler appliance.

To partition a NetScaler by using the command line interface

On the command prompt, do the following:

1. Create a partition and configure the NetScaler resources for that partition.

```
add ns partition <partitionName> [-maxBandwidth <positive_integer>] [-minBandwidth <positive_integer>] [-maxConn <positive_integer>] [-maxMemLimit <positive_integer>]
```

Note:

- For deployments that have large size of NetScaler configuration and large quantum of traffic, Citrix advises that you change the maximum memory limit from the default value of 10 MB.
- Use the set ns partition command if you want to modify any configuration settings.

2. Associate the appropriate users with the partition.

```
bind system user <name> -partitionName <string>
```

3. Associate one of the following command policies with the user: `partition-operator`, `partition-read-only`, `partition-network`, and `partition-admin`.

```
bind system user <name> <policyName> <priority>
```

4. Configure the VLAN through which traffic for this partition must be routed. You can use bridgegroups instead of VLANs to route the traffic.

- Add the VLAN and bind the required interfaces to it.

```
add vlan <id>
```

```
bind vlan <id> -ifnum <interface>
```

OR

- Add the bridgegroup and bind the required VLANs to it.

```
add bridgegroup <id>
```

```
bind bridgegroup <id> -vlan <id>
```

5. Bind the VLAN or bridgegroup to the partition.

```
bind ns partition <partitionName> -vlan <positive_integer>
```

OR

```
bind ns partition <partitionName> -bridgegroup <positive_integer>
```

Note: Use the show vlan or the show bridgegroup command to view the partitions associated with that VLAN or bridgegroup.

6. Verify the configurations of the partition.

```
show ns partition <partitionName>
```

Note: You can also use the `stat ns partition` command to view partition configurations.

7. Save the configuration.

`save ns config`

To partition a NetScaler by using the configuration utility

On the Configuration tab of the graphical user interface:

1. Navigate to System > Partition Administration, click Add and do the following:
 - a. Create and configure the resources for the admin partition.
 - b. Specify the VLANs or bridgegroups to be associated with the partition.
 - c. Associate user(s) with the partition.
Note: Make sure you bind users who are not yet associated with partition type command policies.
2. Navigate to System > User Administration, and to the partition user, bind the appropriate command policy. The command policy must be one of the `partition-` entries. The choice depends on the level of authorization you intend the user to have.
3. Save the configuration.

Configuring in a NetScaler Partition

Accessing a partitioned NetScaler is the same as accessing a non-partitioned NetScaler: through the NetScaler IP (NSIP) address or any other management IP address. As a user, after you provide your valid logon credentials, you are taken to the partition to which you are bound. Any configurations that you create are saved to that partition. If you are associated with more than one partition, you are taken to the first partition with which you were associated. If you want to configure entities on one of your other partitions, you must explicitly switch to that partition.

Note:

- NetScaler superusers and other non-partition users are taken to the default partition.
- Users of all the 512 partitions can log in simultaneously.

After accessing the appropriate partition, configurations that you perform are saved to that partition and are specific to that partition.

To configure in a NetScaler partition by using the command line interface

1. Log on to the NetScaler.
2. Check if you are in the correct partition. The command prompt displays the name of the currently selected partition.
 - If yes, skip to the next step.
 - If no, get a list of the partitions with which you are associated and switch over to the appropriate partition.
 - a. `show system user <username>`
 - b. `switch ns partition <partitionName>`
3. Now, you can perform the required configurations on the NetScaler just as you would do on a non-partitioned NetScaler.

To configure in a NetScaler partition by using the configuration utility


1. Log on to the NetScaler.
2. Check if you are in the correct partition. The top bar of the graphical user interface displays the name of the currently selected partition.
 - If yes, skip to the next step.
 - If no, navigate to System > Administrative Partitions > Partitions, right-click the partition to which you want to switch, and select Switch.
3. Now, you can perform the required configurations on the NetScaler just as you would do on a non-partitioned NetScaler.

Use Case: Resource Sharing Across Admin Partitions

By using admin partitions, resources such as IP addresses and entity names can be shared across different partitions. This means that you can now use an IP address (for example, 3.3.3.7) as the virtual server IP address across partitions or use the same name (for example, lbvserver1) for multiple virtual servers across partitions.

This is possible as each partition is associated with a different VLAN (or bridgegroup) and therefore traffic destined for different applications is segregated.

As shown in the following image, the virtual server IP address 3.3.3.7 is used in Partition1, Partition2, and Partition3.



Let us understand how the configurations must be performed.

Creating and configuring partitions that share the same IP address among virtual servers

1. **On default partition:** Log on to the NetScaler appliance as a super user and configure the three partitions as follows:

```
shell> ssh 10.102.29.60 -l nsroot
password: *****

> add ns partition Partition1
Done

> add ns partition Partition2
Done

> add ns partition Partition3
Done

> bind system user user1 -partitionName Partition1
Done

> bind system user user2 -partitionName Partition2
Done

> bind system user user3 -partitionName Partition3
Done

> bind system user user1 partition-admin 10
Done

> bind system user user2 partition-admin 20
Done

> bind system user user3 partition-admin 20
Done

> add vlan 20
Done

> bind vlan 20 -ifnum 2/1
Done

> add vlan 21
Done

> bind vlan 21 -ifnum 3/1
Done

> add vlan 30
```

Done

```
> bind vlan 30 -ifnum 4/1
Done
```

```
> bind ns partition Partition1 -vlan 20
Done
```

```
> bind ns partition Partition2 -vlan 21
Done
```

```
> bind ns partition Partition3 -vlan 30
Done
```

2. On Partition1: Log on to the NetScaler appliance as user1 and configure on Partition1.

```
shell> ssh 10.102.29.60 -l user1
password: *****
```

```
Partition1> add ns ip 3.3.3.2 255.255.255.0 -vServer DISABLED -type SNIP
Done
```

```
Partition1> add service s1 3.3.3.5 HTTP 80
Done
```

```
Partition1> add lb vserver lbvserver1 HTTP 3.3.3.7 80 -persistenceType NONE
Done
```

```
Partition1> bind lb vserver lbvserver1 s1
Done
```

```
Partition1> bind vlan 20 -IPAddress 3.3.3.2 255.255.255.0
Done
```

3. On Partition2: Log on to the NetScaler appliance as user2 and configure on Partition2.

```
shell> ssh 10.102.29.60 -l user2
password: *****
```

```
Partition2> add ns ip 5.5.5.3 255.255.255.0 -vServer DISABLED -type SNIP
Done
```

```
Partition2> add service s1 5.5.5.5 HTTP 80
Done
```

```
Partition2> add lb vserver lbvserver1 HTTP 3.3.3.7 80 -persistenceType NONE
Done
```

```
Partition2> bind lb vserver lbvserver1 s1
Done
```

```
Partition2> bind vlan 21 -IPAddress 5.5.5.3 255.255.255.0
Done
```

4. On Partition3: Log on to the NetScaler appliance as user3 and configure on Partition3.

```
shell> ssh 10.102.29.60 -l user3
password: *****
```

```
Partition3> add ns ip 6.6.6.3 255.255.255.0 -vServer DISABLED -type SNIP
Done
```

```
Partition3> add service s1 6.6.6.6 HTTP 80
Done
```

```
Partition3> add lb vserver lbvserver1 HTTP 3.3.3.7 80 -persistenceType NONE
Done
```

```
Partition3> bind lb vserver lbvserver1 s1
Done
```

```
Partition3> bind vlan 30 -IPAddress 6.6.6.3 255.255.255.0
Done
```

Admin Partitions FAQs

How can I get auditlogs for an admin partition?

In a partitioned NetScaler, you cannot have specific log servers for a specific partition. The servers that are defined at the default partition are applicable across all admin partitions. Therefore, to view the audit logs for a specific partition, you will have to use the "show audit messages" command.

Note: The users of an admin partition do not have access to the shell and therefore are not able to access the log files

How can I get web logs for an admin partition?

You can get the web logs for an admin partition as follows:

- For NetScaler 11 and later versions

The web logging feature must be enabled on each of the partitions that require web logging. Using the NetScaler Web Logging (NSWL) client, the NetScaler retrieves the web logs for all the partitions with which the user is associated.

- For versions prior to NetScaler 11

Web logs can be obtained only by nsroot and other superusers. Also, even though web logging is enabled on the default partition, the NetScaler Web Logging (NSWL) client fetches web logs for all the partitions.

To view the partition for each log entry, customize the log format to include the %P option. You can then filter the logs to view the logs for a specific partition.

How can I get the trace for an admin partition?

You can get the trace for an admin partition as follows:

- For NetScaler 11 and later versions

In a partitioned NetScaler appliance, the nstrace operation can be performed on individual admin partitions. The trace files are stored in the /var/partitions/<partitionName>/nstrace/directory.

- For versions prior to NetScaler 11

The nstrace operation can only be performed on the default partition. Therefore, packet captures are available for the entire NetScaler system. To get partition-specific packet captures, use VLAN-ID based filters.

How can I configure integrated caching in a partitioned NetScaler appliance?

Note: Integrated caching in admin partitions is supported from NetScaler 11 onwards.

To configure integrated caching (IC) on a partitioned NetScaler, after defining the IC memory on the default partition, the superuser can configure the IC memory on each admin partition such that the total IC memory allocated to all admin partitions does not exceed the IC memory defined on the default partition. The memory that is not configured for the admin partitions remains available for the default partition.

For example, if a NetScaler appliance with two admin partitions has 10 GB of IC memory allocated to the default partition, and IC memory allocation for the two admin partitions is as follows:

- Partition1: 4 GB
- Partition2: 3 GB

Then, the default partition has $10 - (4 + 3) = 3$ GB of IC memory available for use.

Note: If all IC memory is used by the admin partitions, no IC memory is available for the default partition.

Where can I find the logs for a partition?

NetScaler logs are not partition-specific. Log entries for all partitions must be stored in the /var/log/ directory.

Where can I get the NetScaler configuration file for a partition?

The configuration file (ns.conf) for the default partition is available in the /nsconfig directory. For admin partitions, the file is available in the /nsconfig/partitions/<partitionName> directory.

How can I get the technical support bundle specific to an admin partition?

To get the tech support bundle for a specific partition, you must execute the `show techsupport -scope partition -partitionname <string>` command from the default partition.

Note: This command also gives system-specific information.

What is the scope for L2 and L3 parameters in admin partitions?

Note: Applicable from NetScaler 11 onwards.

On a partitioned NetScaler appliance, the scope of updating the L2 and L3 parameters is as follows:

- For L2 parameters that are set by using the "set L2Param" command, the following parameters can be updated only from the default partition, and their values are applicable to all the admin partitions:

maxBridgeCollision, bdgSetting, garpOnVridIntf, garpReply, proxyArp, resetInterfaceOnHAfailover, and skip_proxying_bsd_traffic

The other L2 parameters can be updated in specific admin partitions, and their values are local to those partitions.

- For L3 parameters that are set by using the "set L3Param" command, all parameters can be updated in specific admin partitions, and their values are local to those partitions. Similarly, the values that are updated in the default partition are applicable only to the default partition.

How to enable dynamic routing in an admin partition?

Note: Dynamic routing in admin partitions is supported from NetScaler 11 onwards.

While dynamic routing (OSPF, RIP, BGP, ISIS, BGP+) is by default enabled on the default partition, in an admin partition, it must be enabled by using the following command:

```
set L3Param -dynamicRouting ENABLED
```

Note: A maximum of 63 partitions can run dynamic routing (62 admin partitions and 1 default partition).

On enabling dynamic routing on an admin partition, a virtual router (VR) is created.

- Each VR maintains its own vlan0 which will be displayed as vlan0_<partition-name>.
- All unbound IP addresses that are exposed to ZebOS are bound to vlan0.
- The default VR (of the default partition) shows all the VRs that are configured.
- The default VR shows the VLANs that are bound to these VRs (except default vlans).

Authentication and Authorization

To configure NetScaler authentication and authorization, you must first define the users who have access to the NetScaler appliance, and then you can organize these users into groups. After configuring users and groups, you need to configure command policies to define types of access, and assign the policies to users and/or groups.

You must log on as an administrator to configure users, groups, and command policies. The default NetScaler administrator user name is *nsroot*. After logging on as the default administrator, you should change the password for the *nsroot* account. Once you have changed the password, no user can access the NetScaler appliance until you create an account for that user. If you forget the administrator password after changing it from the default, you can reset it to *nsroot*.

Note: Local users can authenticate to the NetScaler even if external authentication servers are configured. You can restrict this by disabling the *localAuth* parameter of the *set system* parameter command.

Configuring Users, User Groups, and Command Policies

You must define your users by configuring accounts for them. To simplify the management of user accounts, you can organize them into groups.

You can also customize the command-line prompt for a user. Prompts can be defined in a user's configuration, in a user-group configuration, and in the global configuration. The prompt displayed for a given user is determined by the following order of precedence:

1. Display the prompt as defined in the user's configuration.
2. Display the prompt as defined in the group configuration for the user's group.
3. Display the prompt as defined in the system global configuration.

You can now specify a time-out value for inactive CLI sessions for a system user. If a user's CLI session is idle for a time that exceeds the time-out value, the NetScaler appliance terminates the connection. The timeout can be defined in a user's configuration, in a user-group configuration, and in the global configuration. The time-out for inactive CLI sessions for a user is determined by the following order of precedence:

1. Time-out value as defined in the user's configuration.
2. Time-out value as defined in the group configuration for the user's group.
3. Time-out value as defined in the system global configuration.

This document includes the following details:

- [Configuring User Accounts](#)
- [Configuring User Groups](#)
- [Configuring Command Policies](#)

Configuring User Accounts

Updated: 2014-08-07

To configure user accounts, you simply specify user names and passwords. You can change passwords and remove user accounts at any time.

To create a user account by using the command line interface

At the command prompt, type the following commands to create a user account and verify the configuration:

- `add system user <userName> [-promptString <string>] [-timeout <secs>]`
- `show system user <userName>`

Example

```
> add system user johnd -promptString user-%u-at-%T
```

```
Enter password:
```

```
Confirm password:
```

To configure a user account by using the configuration utility

Navigate to System > User Administration > Users, and create the user.

Configuring User Groups

Updated: 2014-08-07

After configuring a user group, you can easily grant the same access rights to everyone in the group. To configure a group, you create the group and bind users to the group. You can bind each user account to more than one group. Binding user accounts to multiple groups may allow more flexibility when applying command policies.

To create a user group by using the command line interface

At the command prompt, type the following commands to create a user group and verify the configuration:

- `add system group <groupName> [-promptString <string>] [-timeout <secs>]`
- `show system group <groupName>`

Example

```
> add system group Managers -promptString Group-Managers-at-%h
```

To bind a user to a group by using the command line interface

At the command prompt, type the following commands to bind a user account to a group and verify the configuration:

- bind system group <groupName> -userName <userName>
- show system group <groupName>

Example

```
> bind system group Managers -userName user1
```

To configure a user group by using the configuration utility

Navigate to System > User Administration > Groups, and create the user group.

Note: To add members to the group, in the Members section, click Add. Select users from the Available list and add them to the Configured list.

Configuring Command Policies

Command policies regulate which commands, command groups, vservers, and other entities that users and user groups are permitted to use.

The appliance provides a set of built-in command policies, and you can configure custom policies. To apply the policies, you bind them to users and/or groups.

Here are the key points to keep in mind when defining and applying command policies.

- You cannot create global command policies. Command policies must be bound directly to the users and groups on the appliance.
- Users or groups with no associated command policies are subject to the default (DENY-ALL) command policy, and are therefore unable to execute any configuration commands until the proper command policies are bound to their accounts.
- All users inherit the policies of the groups to which they belong.
- You must assign a priority to a command policy when you bind it to a user account or group account. This enables the appliance to determine which policy has priority when two or more conflicting policies apply to the same user or group.
- The following commands are available by default to any user and are unaffected by any command you specify:

help, show cli attribute, set cli prompt, clear cli prompt, show cli prompt, alias, unalias, history, quit, exit, whoami, config, set cli mode, unset cli mode, and show cli mode.

Built-in Command Policies

Updated: 2015-06-15

The following table describes the built-in policies.

Table 1. Built-in Command Policies

Policy name	Allows
read-only	Read-only access to all show commands except show ns runningConfig, show ns ns.conf, and the show commands for the NetScaler command group.
operator	Read-only access and access to commands to enable and disable services and servers.
network	Full access, except to the set and unset SSL commands, show ns ns.conf, show ns runningConfig, and show gslb runningConfig commands.
sysadmin	[From NetScaler 11 onwards] A sysadmin is lower than a superuser in terms of access allowed on the appliance. A sysadmin user can perform all NetScaler operations with the following exceptions: no access to the NetScaler shell, cannot perform user configurations, cannot perform partition configurations, and some other configurations as stated in the sysadmin command policy.
superuser	Full access. Same privileges as the nsroot user.

Creating Custom Command Policies

Updated: 2015-06-15

Regular expression support is offered for users with the resources to maintain more customized expressions, and for those deployments that require the flexibility that regular expressions offer. For most users, the built-in command policies are sufficient. Users who need additional levels of control but are unfamiliar with regular expressions may want to use only simple expressions, such as those in the examples provided in this section, to maintain policy readability.

When you use a regular expression to create a command policy, keep the following in mind.

- When you use regular expressions to define commands that will be affected by a command policy, you must enclose the commands in double quotation marks. For example, to create a command policy that includes all commands that begin with show, type the following:

```
"^show .*$"
```

To create a command policy that includes all commands that begin with rm, type the following:

```
"^rm .*$"
```

- Regular expressions used in command policies are not case sensitive.

The following table lists examples of regular expressions:

Table 2. Examples of Regular Expressions for Command Policies

Command specification	Matches these commands
"^rm\s+.*\$"	All remove actions, because all remove actions begin with the rm string, followed by a space and additional parameters and flags.
"^show\s+.*\$"	All show commands, because all show actions begin with the show string, followed by a space and additional parameters and flags.
"^shell\$"	The shell command alone, but not combined with any other parameters or flags.
"^add\s+vserver\s+.*\$"	All create vserver actions, which consist of the add vserver command followed by a space and additional parameters and flags.
"^add\s+(lb\s+vserver)\s+.*"	All create lb vserver actions, which consist of the add lb vserver command followed by a space and additional parameters and flags.

The following table shows the command specifications for each of the built-in command policies.

Table 3. Expressions Used in the Built-in Command Policies

Policy name	Command specification regular expression
read-only	(^man.*)(^show\s+(?!system)(?!configstatus)(?!ns ns\.conf)(?!ns savedconfig)(?!ns runningConfig)(?!gslib runnin
operator	(^man.*)(^show\s+(?!system)(?!configstatus)(?!ns ns\.conf)(?!ns savedconfig)(?!ns runningConfig)(?!gslib runnin
network	^(?!clear ns config.*)(?!scp.*)(?!set ssl fips)(?!reset ssl fips)(?!diff ns config)(?!shell)(?!reboot)(?!batch)\S+\s+(?!s
sysadmin	[From NetScaler 11 onwards] ^(?!(?!shell)(?!sftp)(?!scp)(?!batch)(?!source)(?!.*superuser)(?!.*nsroot)(?!show\s+sys
superuser	.*

To create a command policy by using the command line interface

At the command prompt, type the following commands to create a command policy and verify the configuration:

- add system cmdPolicy <policyname> <action> <cmds spec>
- show system cmdPolicy <policyName>

Example

```
> add system cmdPolicy read_all ALLOW (^show\s+(!system)(!ns ns.conf)(!ns runningConfig).*)|
```

To configure a command policy by using the configuration utility

Navigate to System > User Administration > Command Policies, and create the command policy.

Binding Command Policies to Users and Groups

Updated: 2014-08-07

Once you have defined your command policies, you must bind them to the appropriate user accounts and groups. When you bind a policy, you must assign it a priority so that the appliance can determine which command policy to follow when two or more applicable command policies are in conflict.

Command policies are evaluated in the following order:

- Command policies bound directly to users and the corresponding groups are evaluated according to priority number. A command policy with a lower priority number is evaluated before one with a higher priority number. Therefore, any privileges the lower-numbered command policy explicitly grants or denies are not overridden by a higher-numbered command policy.
- When two command policies, one bound to a user account and other to a group, have the same priority number, the command policy bound directly to the user account is evaluated first.

To bind command policies to a user by using the command line interface

At the command prompt, type the following commands to bind a command policy to a user and verify the configuration:

- bind system user <userName> -policyName <policyName> <priority>
- show system user <userName>

Example

```
> bind system user user1 -policyName read_all 1
```

To bind command policies to a user by using the configuration utility

Navigate to System > User Administration > Users, select the user and bind command policies.

Optionally, you can modify the default priority to ensure that the policy is evaluated in the proper order.

To bind command policies to a group by using the command line interface

At the command prompt, type the following commands to bind a command policy to a user group and verify the configuration:

- bind system group <groupName> -policyName <policyName> <priority>
- show system group <groupName>

Example

```
> bind system group Managers -policyName read_all 1
```

To bind command policies to a group by using the configuration utility

Navigate to System > User Administration > Groups, select the group and bind command policies.

Optionally, you can modify the default priority to ensure that the policy is evaluated in the proper order.

Resetting the Default Administrator (nsroot) Password

The nsroot account provides complete access to all features of the appliance. Therefore, to preserve security, the nsroot account should be used only when necessary, and only individuals whose duties require full access should know the password for the nsroot account. Frequently changing the nsroot password is advisable. If you lose the password, you can reset it to the default and then change it.

To reset the nsroot password, you must boot the appliance into single user mode, mount the file systems in read/write mode, and remove the set NetScaler user nsroot entry from the ns.conf file. You can then reboot, log on with the default password, and choose a new password.

To reset the nsroot password

1. Connect a computer to the console port of the NetScaler ADC and log on.
Note: You cannot log on by using SSH to perform this procedure; you must connect directly to the appliance.
2. Reboot the NetScaler ADC.
3. Press CTRL+C when the following message appears:

```
Press [Ctrl-C] for command prompt, or any other key to boot immediately.
```

```
Booting [kernel] in # seconds.
```

4. Run the following command to start the NetScaler in a single user mode:

```
boot -s
```

Note: If boot -s does not work, then try reboot -- -s and appliance will reboot in single user mode.

After the appliance boots, it displays the following message:

```
Enter full path name of shell or RETURN for /bin/sh:
```

5. Press ENTER key to display the # prompt, and type the following commands to mount the file systems:
 - a. Run the following command to check the disk consistency:

```
fsck /dev/ad0s1a
```

Note: Your flash drive will have a specific device name depending on your NetScaler; hence, you have to replace ad0s1a in the preceding command with the appropriate device name.

- b. Run the following command to display the mounted partitions:

```
df
```

If the flash partition is not listed, you need to mount it manually.

- c. Run the following command to mount the flash drive:

```
mount /dev/ad0s1a /flash
```

6. Run the following command to change to the nsconfig directory:

```
cd /flash/nsconfig
```

7. Run the following commands to rewrite the ns.conf file and remove the set of system commands defaulting to the nsroot user:

- a. Run the following command to create a new configuration file that does not have commands defaulting to the nsroot user:

```
grep -v "set system user nsroot" ns.conf > new.conf
```

- b. Run the following command to make a backup of the existing configuration file:

```
mv ns.conf old.ns.conf
```

- c. Run the following command to rename the new.conf file to ns.conf:

```
mv new.conf ns.conf
```

8. Run the following command to reboot the NetScaler:

```
reboot
```

9. Log on using the default nsroot user credentials.

10. Run the following command to reset the nsroot user password:

```
set system user nsroot <New_Password>
```

Example of a User Scenario

The following example shows how to create a complete set of user accounts, groups, and command policies and bind each policy to the appropriate groups and users. The company, Example Manufacturing, Inc., has three users who can access the NetScaler appliance:

- **John Doe.** The IT manager. John needs to be able to see all parts of the NetScaler configuration but does not need to modify anything.
- **Maria Ramiez.** The lead IT administrator. Maria needs to be able to see and modify all parts of the NetScaler configuration except for NetScaler commands (which local policy dictates must be performed while logged on as nsroot).
- **Michael Baldrock.** The IT administrator in charge of load balancing. Michael needs to be able to see all parts of the NetScaler configuration, but needs to modify only the load balancing functions.

The following table shows the breakdown of network information, user account names, group names, and command policies for the sample company.

Table 1. Sample Values for Creating Entities

Field	Value	Note
NetScaler host name	ns01.example.net	N/A
User accounts	johnd, mariar, and michaelb	John Doe, IT manager, Maria Ramirez, IT administrator and Michael Baldrock, IT administrator.
Groups	Managers and SysOps	All managers and all IT administrators.
Command Policies	read_all, modify_lb, and modify_all	Allow complete read-only access, Allow modify access to load balancing, and Allow complete modify access.

The following description walks you through the process of creating a complete set of user accounts, groups, and command policies on the NetScaler appliance named ns01.example.net.

The description includes procedures for binding the appropriate user accounts and groups to one another, and binding appropriate command policies to the user accounts and groups.

This example illustrates how you can use prioritization to grant precise access and privileges to each user in the IT department.

The example assumes that initial installation and configuration have already been performed on the NetScaler.

Configuration steps

1. Use the procedure described in "Configuring User Accounts" to create user accounts **johnd**, **mariar**, and **michaelb**.
2. Use the procedure described in "Configuring User Groups" to create user groups **Managers** and **SysOps**, and then bind the users **mariar** and **michaelb** to the **SysOps** group and the user **johnd** to the **Managers** group.
3. Use the procedure described in "Creating Custom Command Policies" to create the following command policies:
 - **read_all** with action **Allow** and command spec `(^show\s+(?!system)(?!ns ns.conf)(?!ns runningConfig).*)|(^stat.*)`
 - **modify_lb** with action as **Allow** and the command spec `^set\s+lb\s+.*$`
 - **modify_all** with action as **Allow** and the command spec `^S\s+(?!system).*`
4. Use the procedure described in "Binding Command Policies to Users and Groups" to bind the **read_all** command policy to the **SysOps** group, with priority value **1**.
5. Use the procedure described in "Binding Command Policies to Users and Groups" to bind the **modify_lb** command policy to user **michaelb**, with priority value **5**.

The configuration you just created results in the following:

- John Doe, the IT manager, has read-only access to the entire NetScaler configuration, but he cannot make modifications.
- Maria Ramirez, the IT lead, has near-complete access to all areas of the NetScaler configuration, having to log on only to perform NetScaler-level commands.
- Michael Baldrock, the IT administrator responsible for load balancing, has read-only access to the NetScaler configuration, and can modify the configuration options for load balancing.

The set of command policies that applies to a specific user is a combination of command policies applied directly to the user's account and command policies applied to the group(s) of which the user is a member.

Each time a user enters a command, the operating system searches the command policies for that user until it finds a policy with an ALLOW or DENY action that matches the command. When it finds a match, the operating system stops its command policy search and allows or denies access to the command.

If the operating system finds no matching command policy, it denies the user access to the command, in accordance with the NetScaler appliance's default deny policy.

Note: When placing a user into multiple groups, take care not to cause unintended user command restrictions or privileges. To avoid these conflicts, when organizing your users in groups, bear in mind the NetScaler command policy search procedure and policy ordering rules.

Configuring External User Authentication

External user authentication is the process of authenticating the users of the Citrix NetScaler appliance by using an external authentication server. The NetScaler supports LDAP, RADIUS, and TACACS+ authentication servers. To configure external user authentication, you must create authentication policies. You can configure one or many authentication policies, depending on your authentication needs. An authentication policy consists of an expression and an action. Authentication policies use NetScaler classic expressions, which are described in detail in "Policy Configuration and Reference."

After creating an authentication policy, you bind it to the system global entity and assign a priority to it. You can create simple server configurations by binding a single authentication policy to the system global entity. Or, you can configure a cascade of authentication servers by binding multiple policies to the system global entity. If no authentication policies are bound to the system, users are authenticated by the onboard system.

This document includes the following details:

- [Configuring LDAP Authentication](#)
- [Configuring RADIUS Authentication](#)
- [Configuring TACACS+ Authentication](#)
- [Binding the Authentication Policies to the System Global Entity](#)

Configuring LDAP Authentication

Updated: 2014-12-29

You can configure the NetScaler appliance to authenticate user access with one or more LDAP servers. LDAP authorization requires identical group names in Active Directory, on the LDAP server, and on the appliance. The characters and case must also be the same.

By default, LDAP authentication is secured by using SSL/TLS protocol. There are two types of secure LDAP connections. In the first type, the LDAP server accepts the SSL/TLS connection on a port separate from the port used to accept clear LDAP connections. After users establish the SSL/TLS connection, LDAP traffic can be sent over the connection. The second type allows both unsecure and secure LDAP connections and is handled by a single port on the server. In this scenario, to create a secure connection, the client first establishes a clear LDAP connection. Then the LDAP command StartTLS is sent to the server over the connection. If the LDAP server supports StartTLS, the connection is converted to a secure LDAP connection by using TLS.

The port numbers for LDAP connections are:

- 389 for unsecured LDAP connections
- 636 for secure LDAP connections
- 3268 for Microsoft unsecure LDAP connections
- 3269 for Microsoft secure LDAP connections

LDAP connections that use the StartTLS command use port number 389. If port numbers 389 or 3268 are configured on the appliance, it tries to use StartTLS to make the connection. If any other port number is used, connection attempts use SSL/TLS. If StartTLS or SSL/TLS cannot be used, the connection fails.

When configuring the LDAP server, the case of the alphabetic characters must match that on the server and on the appliance. If the root directory of the LDAP server is specified, all of the subdirectories are also searched to find the user attribute. In large directories, this can affect performance. For this reason, Citrix recommends that you use a specific organizational unit (OU).

The following table lists examples of user attribute fields for LDAP servers.

Table 1. User Attribute Fields for LDAP Servers

LDAP server	User attribute	Case sensitive?
Microsoft Active Directory	Server sAMAccountName	No
Novell eDirectory	cn	Yes
IBM Directory Server	uid	Yes

Lotus Domino	CN	Yes
Sun ONE directory (formerly iPlanet)	uid or cn	Yes

The following table lists examples of the base distinguished name (DN).

Table 2. Examples of Base Distinguished Name

LDAP server	Base DN
Microsoft Active Directory	DC=citrix, DC=local
Novell eDirectory	dc=citrix, dc=net
IBM Directory Server	cn=users
Lotus Domino	OU=City, O=Citrix, C=US
Sun ONE directory (formerly iPlanet)	ou=People, dc=citrix, dc=com

The following table lists examples of the bind distinguished name (DN).

Table 3. Examples of Bind Distinguished Name

LDAP server	Bind DN
Microsoft Active Directory	CN=Administrator, CN=Users, DC=citrix, DC=local
Novell eDirectory	cn=admin, dc=citrix, dc=net
IBM Directory Server	LDAP_dn
Lotus Domino	CN=Notes Administrator, O=Citrix, C=US
Sun ONE directory (formerly iPlanet)	uid=admin, ou=Administrators, ou=TopologyManagement, o=NetscapeRoot

To configure LDAP authentication by using the command line interface

At the command prompt, do the following:

1. Create an LDAP action.

```
add authentication ldapAction <name> {-serverIP <ip_addr|ipv6_addr|*> | {-serverName <string>}} >] [-authTimeout <positive_integer>] [-ldapBase <string>] [-ldapBindDn <string>] {-ldapBindDnPassword } [-ldapLoginName <string>] [-groupAttrName <string>] [-subAttributeName <string>]
```

Example:

```
add authentication ldapAction ldap70 -serverIP <IP> -authTimeout 30 -ldapBase "CN=xxxxx"
```

2. Create an LDAP policy.

```
add authentication ldapPolicy <name> <rule> [<reqAction>]
```

Example:

```
add authentication ldappolicy ldap_pol ns_true ldap70
```

3. Bind the LDAP policy to the following bind points at which the policy will be evaluated.
 - o **System Global:** bind system global <policyName> [-priority <positive_integer>]

- o **VPN Global:** bind vpn global <policyName> [-priority <positive_integer>]
- o **Authentication Server:** bind authentication vserver <name> [-policy <string> [-priority <positive_integer>]]
- o **VPN Server:** bind vpn vserver <name> [-policy <string> [-priority <positive_integer>]]

To configure LDAP authentication by using the configuration utility

Navigate to System > Authentication > LDAP, and create the LDAP authentication policy.

Determining attributes in the LDAP directory

If you need help determining your LDAP directory attributes, you can easily look them up with the free LDAP browser from Softerra.

You can download the LDAP browser from the Softerra LDAP Administrator Web site at <http://www.ldapbrowser.com>. After the browser is installed, set the following attributes:

- o The host name or IP address of your LDAP server.
- o The port of your LDAP server. The default is 389.
- o The base DN field can be left blank.
- o The information provided by the LDAP browser can help you determine the base DN needed for the Authentication tab.
- o The Anonymous Bind check determines whether the LDAP server requires user credentials for the browser to connect to it. If the LDAP server requires credentials, leave the check box cleared.

After completing the settings, the LDAP browser displays the profile name in the left pane and connects to the LDAP server.

Configuring RADIUS Authentication

Updated: 2014-08-08

You can configure the NetScaler appliance to authenticate user access with one or more RADIUS servers. If you are using RSA SecurID, SafeWord, or Gemalto Protiva products, use a RADIUS server.

Your configuration might require using a network access server IP address (NAS IP) or a network access server identifier (NAS ID). When configuring the appliance to use a RADIUS authentication server, use the following guidelines:

- o If you enable use of the NAS IP, the appliance sends its configured IP address to the RADIUS server, rather than the source IP address used in establishing the RADIUS connection.
- o If you configure the NAS ID, the appliance sends the identifier to the RADIUS server. If you do not configure the NAS ID, the appliance sends its host name to the RADIUS server.
- o When the NAS IP is enabled, the appliance ignores any NAS ID that was configured by using the NAS IP to communicate with the RADIUS server.

To configure RADIUS authentication by using the configuration utility

Navigate to System > Authentication > Radius, and create the RADIUS authentication policy.

Choosing RADIUS authentication protocols

The NetScaler appliance supports implementations of RADIUS that are configured to use any of several protocols for user authentication, including:

- o Password Authentication Protocol
- o Challenge-Handshake Authentication Protocol (CHAP)
- o Microsoft Challenge-Handshake Authentication Protocol (MS-CHAP Version 1 and Version 2)

If your deployment of the appliance is configured to use RADIUS authentication and your RADIUS server is configured to use Password Authentication Protocol, you can strengthen user authentication by assigning a strong shared secret to the RADIUS server. Strong RADIUS shared secrets consist of random sequences of uppercase and lowercase letters, numbers, and punctuation, and are at least 22 characters long. If possible, use a random character generation program to determine RADIUS shared secrets.

To further protect RADIUS traffic, assign a different shared secret to each appliance or virtual server. When you define clients on the RADIUS server, you can also assign a separate shared secret to each client. If you do this, you must configure separately each policy that uses RADIUS authentication.

Shared secrets are configured on the appliance when a RADIUS policy is created.

Configuring IP address extraction

You can configure the appliance to extract the IP address from a RADIUS server. When a user authenticates with the RADIUS server, the server returns a framed IP address that is assigned to the user. The following are attributes for IP address extraction:

- Allows a remote RADIUS server to supply an IP address from the internal network for a user logged on to the appliance.
- Allows configuration for any RADIUS attribute using the type `ipaddress`, including those that are vendor encoded.

When configuring the RADIUS server for IP address extraction, you configure the vendor identifier and the attribute type.

The vendor identifier enables the RADIUS server to assign an IP address to the client from a pool of IP addresses that are configured on the RADIUS server. The vendor ID and attributes are used to make the association between the RADIUS client and the RADIUS server. The vendor ID is the attribute in the RADIUS response that provides the IP address of the internal network. A value of zero indicates that the attribute is not vendor encoded. The attribute type is the remote IP address attribute in a RADIUS response. The minimum value is one and the maximum value is 255.

A common configuration is to extract the RADIUS attribute *framed IP address*. The vendor ID is set to zero or is not specified. The attribute type is set to eight.

To configure IP address extraction by using the configuration utility

1. Navigate to System > Authentication > Radius, and select a policy.
2. Modify the server parameters and set relevant values in Group Vendor Identifier and Group Attribute Type fields.

Configuring TACACS+ Authentication

Updated: 2014-08-07

You can configure a TACACS+ server for authentication. Similar to RADIUS authentication, TACACS+ uses a secret key, an IP address, and the port number. The default port number is 49. To configure the appliance to use a TACACS+ server, provide the server IP address and the TACACS+ secret. The port needs to be specified only when the server port number in use is something other than the default port number of 49.

To configure TACACS+ authentication by using the configuration utility

Navigate to System > Authentication > TACACS, and create the TACACS authentication policy.

After the TACACS+ server settings are configured on the appliance, bind the policy to the system global entity. For more information about binding authentication policies globally, see ["Binding the Authentication Policies to the System Global Entity."](#)

Binding the Authentication Policies to the System Global Entity

Updated: 2014-12-30

When the authentication policies are configured, bind the policies to the system global entity.

To bind an authentication policy to system global using the command line interface

At the command line prompt, do the following:

```
bind system global <policyName> [-priority <positive_integer>]
```

Example:

```
bind system global ldappoll -priority 10
```

To bind an authentication policy to system global using the configuration utility

1. Navigate to System > Authentication, and select the authentication type.
2. On the Policies tab, click Global Bindings and bind the authentication policies.

TCP Configurations

TCP configurations for a NetScaler appliance can be specified in an entity called a TCP profile, which is a collection of TCP settings. The TCP profile can then be associated with services or virtual servers that want to use these TCP configurations.

A default TCP profile can be configured to set the TCP configurations that will be applied by default, globally to all services and virtual servers.

Note: When a TCP parameter has different values for service, virtual server, and globally, the value of the most-specific entity (the service) is given the highest precedence.

The NetScaler appliance also provides other approaches for configuring TCP. Read on for more information.

The NetScaler appliance supports the following TCP capabilities:

- Defending TCP against spoofing attacks. The NetScaler implementation of window attenuation is RFC 4953 compliant.
- Explicit Congestion Notification (ECN), which sends notification of the network congestion status to the sender of the data and takes corrective measures for data congestion or data corruption. The NetScaler implementation of ECN is RFC 3168 compliant.
- Round Trip Time Measurement (RTTM) using the TimeStamp option. For the TimeStamp option to work, at least one side of the connection (client or server) must support it. The NetScaler implementation of TimeStamp option is RFC 1323 compliant.
- Detection of spurious retransmissions can be done using TCP duplicate selective acknowledgement (D-SACK) and forward RTO-Recovery (F-RTO). In case of spurious retransmissions, the congestion control configurations are reverted to their original state. The NetScaler implementation of D-SACK is RFC 2883 compliant, and F-RTO is RFC 5682 compliant.
- Congestion control using New-Reno, BIC, CUBIC,, and TCP Westwood algorithms.
- Window scaling to increase the TCP receive window size beyond its maximum value of 65,535 bytes.

Note: Before configuring window scaling, make sure that:

You do not set a high value for the scale factor, because this could have adverse effects on the appliance and the network.

You do not configure window scaling unless you clearly know why you want to change the window size. Both hosts in the TCP connection send a window scale option during connection establishment. If only one side of a connection sets this option, window scaling is not used for the connection.

Each connection for same session is an independent window scaling session. For example, when a client's request and the server's response flow through the appliance, it is possible to have window scaling between the client and the appliance without window scaling between the appliance and the server.

- TCP maximum congestion window size that is user configurable. The default value is 8190 bytes.
- Selective acknowledgment (SACK), using which the data receiver (either a NetScaler appliance or a client) notifies the sender about all the segments that have been received successfully.
- Forward acknowledgment (FACK) avoids TCP congestion by explicitly measuring the total number of data bytes outstanding in the network, and helping the sender (either a NetScaler ADC or a client) control the amount of data injected into the network during retransmission timeouts.
- TCP connection multiplexing enables reuse of existing TCP connections. The NetScaler appliance stores established TCP connections to the reuse pool. Whenever a client request is received, appliance checks for an available connection in the reuse pool and serves the new client if the connection is available. If it is unavailable, the appliance creates a new connection for the client request and stores the connection to the reuse pool.

NetScaler supports connection multiplexing for HTTP, SSL, and DataStream connection types.

- Dynamic receive buffering allows the receive buffer to be adjusted dynamically based on memory and network conditions.
- MPTCP connections between client and NetScaler. MPTCP connections are not supported between NetScaler and the backend server.

The NetScaler implementation of MPTCP is RFC 6824 compliant.

Note:

For MPTCP to work, both sides of the connection (client and server) must support it. If you use the NetScaler appliance as an MPTCP gateway for your servers, the servers do not have to support MPTCP

The NetScaler appliance does not initiate subflows (MP_JOIN's). The appliance expects the client to initiate subflows.

- TCP keep-alive to monitor the TCP connections to verify if the peers are up.

Additionally, NetScaler provides configuration support for the following:

- TCP segmentation offload.
- Synchronizing cookie for TCP handshake with clients. Disabling this capability prevents SYN attack protection on the NetScaler appliance.
- Learning MSS to enable MSS learning for all the virtual servers configured on the appliance.

This document includes the following details:

- [Setting Global TCP Parameters](#)
- [Setting Service or Virtual Server Specific TCP Parameters](#)
- [Built-in TCP Profiles](#)
- [Sample TCP Configurations](#)

Setting Global TCP Parameters

Updated: 2014-08-08

The NetScaler appliance allows you to specify values for TCP parameters that are applicable to all NetScaler services and virtual servers. This can be done using:

- Default TCP profile
- [Global TCP command](#)
- [TCP buffering feature](#)

Note: The `recvBuffSize` parameter of the `set ns tcpParam` command is deprecated from release 9.2 onwards. In later releases, set the buffer size by using the `bufferSize` parameter of the `set ns tcpProfile` command. If you upgrade to a release where the `recvBuffSize` parameter is deprecated, the `bufferSize` parameter is set to its default value.

Default TCP Profile

A TCP profile, named as `nstcp_default_profile`, is used to specify TCP configurations that will be used if no TCP configurations are provided at the service or virtual server level.

Note:

- Not all TCP parameters can be configured through the default TCP profile. Some settings have to be performed by using the global TCP command (see section below).
- The default profile does not have to be explicitly bound to a service or virtual server.

To configure the default TCP profile

- Using the command line interface, at the command prompt enter:

```
set ns tcpProfile nstcp_default_profile
```
- On the configuration utility, navigate to **System > Profiles**, click **TCP Profiles** and update `nstcp_default_profile`.

Global TCP command

Another approach you can use to configure global TCP parameters is the global TCP command. In addition to some unique parameters, this command duplicates some parameters that can be set by using a TCP profile. Any update made to these duplicate parameters is reflected in the corresponding parameter in the default TCP profile.

For example, if the SACK parameter is updated using this approach, the value is reflected in the SACK parameter of the default TCP profile (`nstcp_default_profile`).

Note: Citrix recommends that you use this approach only for TCP parameters that are not available in the default TCP profile.

To configure the global TCP command

- Using the command line interface, at the command prompt enter:

```
set ns tcpParam {
```

- On the configuration utility, navigate to System > Settings, click Change TCP parameters and update the required TCP parameters.

TCP buffering feature

NetScaler provides a feature called TCP buffering that you can use to specify the TCP buffer size. The feature can be enabled globally or at service level.

Note: The buffer size can also be configured in the default TCP profile. If the buffer size has different values in the TCP buffering feature and the default TCP profile, the greater value is applied.

To configure the TCP buffering feature globally

- At the command prompt enter:

```
enable ns mode TCPB
```

```
set ns tcpbufParam -size <positiveInteger> -memLimit <positiveInteger>
```

- On the configuration utility, navigate to System > Settings, click Configure Modes and select TCP Buffering.

And, navigate to System > Settings, click Change TCP parameters and specify the values for Buffer size and Memory usage limit.

Setting Service or Virtual Server Specific TCP Parameters

Updated: 2014-08-08

Using TCP profiles, you can specify TCP parameters for services and virtual servers. You must define a TCP profile (or use a built-in TCP profile) and associate the profile with the appropriate service and virtual server.

Note:

- You can also modify the TCP parameters of default profiles as per your requirements. For more information on built-in TCP profiles, see [Built-in TCP Profiles](#).
- You can specify the TCP buffer size at service level using the parameters specified by the TCP buffering feature. For more information, see [TCP buffering feature](#).

To specify service or virtual server level TCP configurations by using the command line interface

At the command prompt, perform the following:

1. Configure the TCP profile.

```
set ns tcpProfile <profile-name>...
```

2. Bind the TCP profile to the service or virtual server.

To bind the TCP profile to the service:

```
set service <name> ....
```

Example:

```
> set service servicel -tcpProfileName profile1
```

To bind the TCP profile to the virtual server:

```
set lb vserver <name> ....
```

Example:

```
> set lb vserver lbvserver1 -tcpProfileName profile1
```

To specify service or virtual server level TCP configurations by using the configuration utility

At the configuration utility, perform the following:

1. Configure the TCP profile.

Navigate to System > Profiles > TCP Profiles, and create the TCP profile.

2. Bind the TCP profile to the service or virtual server.

Navigate to Traffic Management > Load Balancing > Services/Virtual Servers, and create the TCP profile, which should be bound to the service or virtual server.

Built-in TCP Profiles

Updated: 2014-04-07

For convenience of configuration, the NetScaler provides some built-in TCP profiles. Review the built-in profiles listed below and select a profile and use it as it is or modify it to meet your requirements. You can bind these profiles to your required services or virtual servers.

Table 1. Built-in TCP Profiles

Built-in profile	Description
nstcp_default_profile	Represents the default global TCP settings on the appliance.
nstcp_default_tcp_lan	Useful for back-end server connections, where these servers reside on the same LAN as the appliance.
nstcp_default_tcp_lan_thin_stream	Similar to the nstcp_default_tcp_lan profile; however, the settings are tuned to small size packet flows.
nstcp_default_tcp_interactive_stream	Similar to the nstcp_default_tcp_lan profile; however, it has a reduced delayed ACK timer and ACK on PUSH packet settings.
nstcp_default_tcp_lfp	Useful for long fat pipe networks (WAN) on the client side. Long fat pipe networks have long delay, high bandwidth lines with minimal packet drops.
nstcp_default_tcp_lfp_thin_stream	Similar to the nstcp_default_tcp_lfp profile; however, the settings are tuned for small size packet flows.
nstcp_default_tcp_lnp	Useful for long narrow pipe networks (WAN) on the client side. Long narrow pipe networks have considerable packet loss once in a while.
nstcp_default_tcp_lnp_thin_stream	Similar to the nstcp_default_tcp_lnp profile; however, the settings are tuned for small size packet flows.
nstcp_internal_apps	Useful for internal applications on the appliance (for example, GSLB sitesyncing). This contains tuned window scaling and SACK options for the desired applications. This profile should not be bound to applications other than internal applications.
nstcp_default_Mobile_profile	Useful for mobile devices.
nstcp_default_XA_XD_profile	Useful for a XenApp or XenDesktop deployment.

Sample TCP Configurations

Updated: 2015-04-28

Sample command line interface examples for configuring the following:

- [Defending TCP against spoofing attacks](#)

- o Explicit Congestion Notification (ECN)
- o Selective ACKnowledgment (SACK)
- o Window Scaling (WS)
- o Maximum Segment Size (MSS)
- o NetScaler to learn the MSS of a virtual server
- o TCP keep-alive
- o Buffer size - using TCP profile
- o Buffer size - using TCP buffering feature
- o MPTCP
- o Congestion control
- o Dynamic receive buffering

Defending TCP against spoofing attacks

Enable the NetScaler to defend TCP against spoof attacks.

```
> set ns tcpProfile profile1 -rstWindowAttenuate ENABLED -spoofSynDrop ENABLED
Done
> set lb vserver lbvserver1 -tcpProfileName profile1
Done
```

Explicit Congestion Notification (ECN)

Enable ECN on the required TCP profile.

```
> set ns tcpProfile profile1 -ECN ENABLED
Done
> set lb vserver lbvserver1 -tcpProfileName profile1
Done
```

Selective ACKnowledgment (SACK)

Enable SACK on the required TCP profile.

```
> set ns tcpProfile profile1 -SACK ENABLED
Done
> set lb vserver lbvserver1 -tcpProfileName profile1
Done
```

Window Scaling (WS)

Enable window scaling and set the window scaling factor on the required TCP profile.

```
> set ns tcpProfile profile1 -WS ENABLED -WSVal 9
Done
> set lb vserver lbvserver1 -tcpProfileName profile1
Done
```

Maximum Segment Size (MSS)

Update the MSS related configurations.

```
> set ns tcpProfile profile1 -mss 1460 - maxPktPerMss 512
Done
> set lb vserver lbvserver1 -tcpProfileName profile1
Done
```

NetScaler to learn the MSS of a virtual server

Enable the NetScaler to learn the VSS and update other related configurations.

```
> set ns tcpParam -learnVsvrMSS ENABLED -mssLearnInterval 180 -mssLearnDelay 3600
Done
```

TCP keep-alive

Enable TCP keep-alive and update other related configurations.


```
> set ns tcpProfile profile1 â€"KA ENABLED â€"KaprobeUpdateLastactivity ENABLED -KAconnIdleT
Done
> set lb vserver lbvserver1 -tcpProfileName profile1
Done
```

Buffer size - using TCP profile

Specify the buffer size.

```
> set ns tcpProfile profile1 â€"bufferSize 8190
Done
> set lb vserver lbvserver1 -tcpProfileName profile1
Done
```

Buffer size - using TCP buffering feature

Enable the TCP buffering feature (globally or for a service) and then specify the buffer size and the memory limit.

```
> enable ns feature TCPB
Done
> set ns tcpbufParam -size 64 -memLimit 64
Done
```

MPTCP

Enable MPTCP and then set the optional MPTCP configurations.

```
> set ns tcpProfile profile1 -mptcp ENABLED
Done
> set ns tcpProfile profile1 -mptcpDropDataOnPreEstSF ENABLED -mptcpFastOpen ENABLED -mptcpS
Done
> set ns tcpParam -mptcpConCloseOnPassiveSF ENABLED -mptcpChecksum ENABLED -mptcpSFtimeout 0
-mptcpMaxSF 4 -mptcpMaxPendingSF 4 -mptcpPendingJoinThreshold 0 -mptcpRTOsToSwitchSF 2 -mptc
Done
```

Congestion control

Set the required TCP congestion control algorithm.

```
> set ns tcpProfile profile1 -flavor Westwood
Done
> set lb vserver lbvserver1 -tcpProfileName profile1
Done
```

Dynamic receive buffering

Enable dynamic receive buffering on the required TCP profile.

```
> set ns tcpProfile profile1 -dynamicReceiveBuffering ENABLED
Done
> set lb vserver lbvserver1 -tcpProfileName profile1
Done
```

HTTP Configurations

HTTP configurations for a NetScaler appliance can be specified in an entity called an HTTP profile, which is a collection of HTTP settings. The HTTP profile can then be associated with services or virtual servers that want to use these HTTP configurations.

A default HTTP profile can be configured to set the HTTP configurations that will be applied by default, globally to all services and virtual servers.

Note: When a HTTP parameter has different values for service, virtual server, and globally, the value of the most-specific entity (the service) is given the highest precedence.

The NetScaler appliance also provides other approaches for configuring HTTP. Read on for more information.

The NetScaler supports the following HTTP capabilities:

- WebSocket protocol which allows browsers and other clients to create a bi-directional, full duplex TCP connection to the servers. The NetScaler implementation of WebSocket is RFC 6455 compliant.
- SPDY (Speedy). For more information, see [SPDY](#).

This document includes the following details:

- [Setting Global HTTP Parameters](#)
- [Setting Service or Virtual Server Specific HTTP Parameters](#)
- [Built-in HTTP Profiles](#)
- [Sample HTTP Configurations](#)

Setting Global HTTP Parameters

Updated: 2014-08-06

The NetScaler appliance allows you to specify values for HTTP parameters that are applicable to all NetScaler services and virtual servers. This can be done using:

- [Default HTTP profile](#)
- [Global HTTP command](#)

Default HTTP profile

A HTTP profile, named as `nshttp_default_profile`, is used to specify HTTP configurations that will be used if no HTTP configurations are provided at the service or virtual server level.

Note:

- Not all HTTP parameters can be configured through the default HTTP profile. Some settings have to be performed by using the global HTTP command (see section below).
- The default profile does not have to be explicitly bound to a service or virtual server.

To configure the default HTTP profile

- Using the command line interface, at the command prompt enter:

```
set ns httpProfile nshttp_default_profile
```

- On the configuration utility, navigate to System > Profiles, click HTTP Profiles and update `nshttp_default_profile`.

Global HTTP command

Another approach you can use to configure global HTTP parameters is the global HTTP command. In addition to some unique parameters, this command duplicates some parameters that can be set by using a HTTP profile. Any update made to these duplicate parameters is reflected in the corresponding parameter in the default HTTP profile.

For example, if the `maxReusePool` parameter is updated using this approach, the value is reflected in the `maxReusePool` parameter of the default HTTP profile (`nshttp_default_profile`).

Note: Citrix recommends that you use this approach only for HTTP parameters that are not available in the default HTTP profile.

To configure the global HTTP command

- Using the command line interface, at the command prompt enter:

```
set ns httpParam &#x2013;
```

- On the configuration utility, navigate to System > Settings, click Change HTTP parameters and update the required HTTP parameters.

Setting Service or Virtual Server Specific HTTP Parameters

Updated: 2014-08-07

Using HTTP profiles, you can specify HTTP parameters for services and virtual servers. You must define a HTTP profile (or use a built-in HTTP profile) and associate the profile with the appropriate service and virtual server.

Note: You can also modify the HTTP parameters of default profiles as per your requirements. For more information on built-in HTTP profiles, see [Built-in HTTP Profiles](#).

To specify service or virtual server level HTTP configurations by using the command line interface

At the command prompt, perform the following:

- Configure the HTTP profile.

```
set ns httpProfile <profile-name>...
```

- Bind the HTTP profile to the service or virtual server.

To bind the HTTP profile to the service:

```
set service <name> .....
```

Example:

```
> set service service1 -httpProfileName profile1
```

To bind the HTTP profile to the virtual server:

```
set lb vserver <name> .....
```

Example:

```
> set lb vserver lbvserver1 -httpProfileName profile1
```

To specify service or virtual server level HTTP configurations by using the configuration utility

At the configuration utility, perform the following:

- Configure the HTTP profile.

Navigate to System > Profiles > HTTP Profiles, and create the HTTP profile.

- Bind the HTTP profile to the service or virtual server.

Navigate to Traffic Management > Load Balancing > Services/Virtual Servers, and create the HTTP profile, which should be bound to the service/virtual server.

Built-in HTTP Profiles

Updated: 2014-04-25

For convenience of configuration, the NetScaler provides some built-in HTTP profiles. Review the profiles listed below and use it as it is or modify it to meet your requirements. You can bind these profiles to the required services or virtual servers.

Table 1. Built-in HTTP Profiles

Built-in profile	Description
nshttp_default_profile	Represents the default global HTTP settings on the appliance.

nshttp_default_strict_validation	Settings for deployments that require strict validation of HTTP requests and responses.
----------------------------------	---

Sample HTTP Configurations

Updated: 2014-04-25

Sample command line interface examples to configure the following:

- [HTTP band statistics](#)
- [WebSocket connections](#)

HTTP band statistics

Specify the band size for HTTP requests and responses.

```
> set protocol httpBand reqBandSize 300 respBandSize 2048
Done
> show protocol httpband -type REQUEST
```

WebSocket connections

Enable webSocket on the required HTTP profile.

```
> set ns httpProfile http_profile1 -webSocket ENABLED
Done
> set lb vserver lbvserver1 -httpProfileName profile1
Done
```

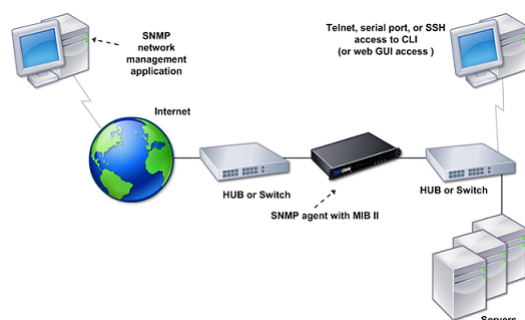
SNMP

You can use Simple Network Management Protocol (SNMP) to configure the SNMP agent on the Citrix NetScaler appliance to generate asynchronous events, which are called *traps*. The traps are generated whenever there are abnormal conditions on the NetScaler. The traps are then sent to a remote device called a *trap listener*, which signals the abnormal condition on the NetScaler appliance. Or, you can query the SNMP agent for System-specific information from a remote device called an *SNMP manager*. The agent then searches the management information base (MIB) for the data requested and sends the data to the SNMP manager.

The SNMP agent on the NetScaler can generate traps compliant with SNMPv1, SNMPv2, and SNMPv3. For querying, the SNMP agent supports SNMP version 1 (SNMPv1), SNMP version 2 (SNMPv2), and SNMP version 3 (SNMPv3).

The following figure illustrates a network with a NetScaler that has SNMP enabled and configured. In the figure, each SNMP network management application uses SNMP to communicate with the SNMP agent on the NetScaler. The SNMP agent searches its management information base (MIB) to collect the data requested by the SNMP Manager and provides the information to the application.

Figure 1. *NetScaler Supporting SNMP*



Importing MIB Files to the SNMP Manager and Trap Listener

To monitor a NetScaler appliance, you must download the MIB object definition files. The MIB files include the following:

- MIB-2 groups SYSTEM, IF, ICMP, UDP, and SNMP.
- NetScaler-specific configuration and statistics.

You can obtain the MIB object definition files from the `/netscaler/snmp` directory or from the Downloads tab of the NetScaler GUI.

If the SNMP management application is other than WhatsUpGold, download the following files to the SNMP management application:

- NS-MIB-smiv1.mib. Used by SNMPv1 managers and trap listeners.
- NS-MIB-smiv2.mib. Used by SNMPv2 and SNMPv3 managers and SNMPv2 trap listeners.

If the SNMP management application is WhatsUpGold, download the following files to the SNMP management application:

- mib.txt
- traps.txt

Configuring the NetScaler to Generate SNMP Traps

You can configure the NetScaler appliance to generate asynchronous events, which are called *traps*. The traps are generated whenever there are abnormal conditions on the appliance. The traps are sent to a remote device called a *trap listener*. This helps administrators monitor the appliance and respond promptly to any issues.

The NetScaler appliance provides a set of condition entities called *SNMP alarms*. When the condition in any SNMP alarm is met, the appliance generates SNMP trap messages that are sent to the configured trap listeners. For example, when the LOGIN-FAILURE alarm is enabled, a trap message is generated and sent to the trap listener whenever there is a login failure on the appliance.

To configure the NetScaler appliance to generate traps, you need to enable and configure alarms. Then, you specify trap listeners to which the appliance will send the generated trap messages.

This document includes the following details:

- [Enabling an SNMP Alarm](#)
- [Configuring Alarms](#)
- [Configuring SNMPv1 or SNMPv2 Traps](#)
- [Configuring SNMPv3 Traps](#)
- [Enabling Unconditional SNMP Trap Logging](#)

Enabling an SNMP Alarm

Updated: 2014-08-08

The NetScaler appliance generates traps only for SNMP alarms that are enabled. Some alarms are enabled by default, but you can disable them.

When you enable an SNMP alarm, the appliance generates corresponding trap messages when some events occur. Some alarms are enabled by default.

To enable an SNMP alarm by using the command line interface

At the command prompt, type the following commands to set the parameters and verify the configuration:

- `enable snmp alarm <trapName>`
- `show snmp alarm <trapName>`

To enable an SNMP alarm by using the configuration utility

1. Navigate to System > SNMP > Alarms, and select the alarm.
2. Click Actions and select Enable.

Configuring Alarms

Updated: 2014-08-06

The NetScaler appliance provides a set of condition entities called *SNMP alarms*. When the condition set for an SNMP alarm is met, the appliance generates SNMP traps messages that are sent to the configured trap listeners. For example, when the LOGIN-FAILURE alarm is enabled, a trap message is generated and sent to the trap listener whenever there is a login failure on the appliance.

You can assign an SNMP alarm with a severity level. When you do this, the corresponding trap messages are assigned that severity level.

The following are the severity levels, defined on the appliance, in decreasing order of severity.

- Critical
- Major
- Minor
- Warning
- Informational

For example, if you set a warning severity level for the SNMP alarm named LOGIN-FAILURE, the trap messages generated when there is a login failure will be assigned with the warning severity level.

You can also configure an SNMP alarm to log the corresponding trap messages generated whenever the condition on that alarm is met.

To configure an SNMP alarm by using the command line interface

At the command prompt, type the following commands to configure an SNMP alarm and verify the configuration:

- `set snmp alarm <trapName> [-thresholdValue <positive_integer> [-normalValue <positive_integer>]] [-time <secs>] [-state (ENABLED | DISABLED)] [-severity <severity>] [-logging (ENABLED | DISABLED)]`
- `show snmp alarm <trapName>`

To configure SNMP alarms by using the configuration utility

Navigate to System > SNMP > Alarms, select an alarm and configure the alarm parameters.

Configuring SNMPv1 or SNMPv2 Traps

Updated: 2014-08-06

After configuring the alarms, you need to specify the trap listener to which the appliance sends the trap messages. Apart from specifying parameters such as IP or IPv6 address and the destination port of the trap listener, you can specify the type of trap (either generic or specific) and the SNMP version.

You can configure a maximum of 20 trap listeners for receiving either generic or specific traps.

You can also configure the appliance to send SNMP trap messages with a source IP address other than the NetScaler IP (NSIP or NSIP6) address to a particular trap listener. For a trap listener that has an IPv4 address, you can set the source IP to either a mapped IP (MIP) address or a subnet IP (SNIP) address configured on the appliance. For a trap listener that has an IPv6 address, you can set the source IP to subnet IPv6 (SNIP6) address configured on the appliance.

You can also configure the appliance to send trap messages to a trap listener on the basis of a severity level. For example, if you set the severity level as Minor for a trap listener, all trap messages of the severity level equal to or greater than Minor (Minor, Major, and Critical) are sent to the trap listener.

If you have defined a community string for the trap listener, you must also specify a community string for each trap that is to be sent to the listener. A trap listener for which a community string has been defined accepts only trap messages that include a community string matching the community string defined in the trap listener. Other trap messages are dropped.

To add an SNMP trap by using the command line interface

At the command prompt, type the following commands to set the parameters and verify the configuration:

- `add snmp trap <trapClass> <trapDestination> -version (V1 | V2) -destPort <port> -communityName <string> -srcIP <ip_addr> -severity <severity>`
- `show snmp trap`

Example

```
> add snmp trap specific 10.102.29.3 -version V2 -destPort 80 -communityName com1 -severity
```

To configure SNMP Traps by using the configuration utility

Navigate to System > SNMP > Traps, and create the SNMP trap.

Configuring SNMPv3 Traps

Updated: 2014-08-08

SNMPv3 provides security capabilities such as authentication and encryption by using the credentials of SNMP users. An SNMP manager can receive SNMPv3 trap messages only if its configuration includes the password assigned to the SNMP user.

The trap destination can now receive SNMPv1, SNMPv2, and SNMPv3 trap messages.

To configure an SNMPv3 trap by using the command line interface

At the command prompt, do the following:

1. Add an SNMPv3 trap.

```
add snmp trap <trapClass> <trapDestination> -version ( V1 | V2 | V3 ) -destPort <port> -communityName <string> -srcIP <ip_addr> -severity <severity>
```

Note: Once set, the SNMP trap version cannot not be modified.

Example

```
> add snmp trap specific 10.102.29.3 -version V3 -destPort 80 -communityName com1 -seve
```

2. Add an SNMP user.

```
add snmp user <name> -group <string> [ -authType ( MD5 | SHA ) { -authPasswd } [-privType ( DES | AES ) { -privPasswd }]]
```

Example

```
> add snmp user edocs_user -group edocs_group
```

3. Bind the SNMPv3 trap to the SNMP user.

```
bind snmp trap <trapClass> <trapDestination> [-version <version>] (-userName <string> [-securityLevel <securityLevel>])
```

Example

```
> bind snmp trap specific 10.102.29.3 -version V3 -userName edocs_user -securityLevel a
```

To configure an SNMPv3 trap by using the configuration utility

1. Add an SNMPv3 trap.

Navigate to System > SNMP > Traps, and create the SNMP trap by selecting V3 as the SNMP version.

2. Add an SNMP user.

Navigate to System > SNMP > Users, and create the SNMP user.

3. Bind the SNMPv3 trap to the SNMP user.

- o Navigate to System > SNMP > Traps, and select the SNMP version 3 trap.
- o Select the user to which the trap should be bound and define the appropriate Security Level.

Enabling Unconditional SNMP Trap Logging

Updated: 2014-08-08

By default, the NetScaler appliance logs any SNMP trap messages (for SNMP alarms in which logging is enabled) when at least one trap listener is specified on the appliance. However, you can specify that SNMP trap messages be logged even when no trap listeners are configured.

To enable unconditional SNMP trap logging by using the command line interface

At the command prompt, type the following commands to configure unconditional SNMP trap logging and verify the configuration:

- o set snmp option -snmpTrapLogging (ENABLED | DISABLED)
- o show snmp option

To enable unconditional SNMP trap logging by using the configuration utility

Navigate to System > SNMP, click Change SNMP Options and select SNMP Trap Logging.

Configuring the NetScaler for SNMP v1 and v2 Queries

You can query the NetScaler SNMP agent for system-specific information from a remote device called *SNMP managers*. The agent then searches the management information base (MIB) for the data requested and sends the data to the SNMP manager.

The following types of SNMP v1 and v2 queries are supported by the SNMP agent:

- GET
- GET NEXT
- ALL
- GET BULK

You can create strings called *community strings* and associate each of these to query types. You can associate one or more community strings to each query type. Community strings are passwords and used to authenticate SNMP queries from SNMP managers.

For example, if you associate two community strings, such as **abc** and **bcd**, to the query type GET NEXT, the SNMP agent on the NetScaler appliance considers only those GET NEXT SNMP query packets that contain **abc** or **bcd** as the community string.

This document includes the following details:

- [Specifying an SNMP Manager](#)
- [Specifying an SNMP Community](#)

Specifying an SNMP Manager

Updated: 2014-08-06

You must configure the NetScaler appliance to allow the appropriate SNMP managers to query it. You must also provide the SNMP manager with the required NetScaler-specific information. You can add up to a maximum of 100 SNMP managers or networks.

For an IPv4 SNMP manager you can specify a host name instead of the manager's IP address. If you do so, you must add a DNS name server that resolves the host name of the SNMP manager to its IP address. You can add up to a maximum of five host-name based SNMP managers.

Note: The appliance does not support use of host names for SNMP managers that have IPv6 addresses. You must specify the IPv6 address.

If you do not configure at least one SNMP manager, the appliance accepts and responds to SNMP queries from all IP addresses on the network. If you configure one or more SNMP managers, the appliance accepts and responds only to SNMP queries from those specific IP addresses.

If you remove an SNMP manager from the configuration, that manager can no longer query the appliance.

To add SNMP managers by specifying IP addresses by using the command line interface

At the command prompt, type the following commands to set the parameters and verify the configuration:

- `add snmp manager <IPAddress> ... [-netmask <netmask>]`
- `show snmp manager`

Example

```
> add snmp manager 10.102.29.10 10.102.29.15 10.102.29.30
```

To add an SNMP manager by specifying its host name by using the command line interface

Important: If you specify the SNMP manager's host name instead of its IP address, you must configure a DNS name server to resolve the host name to the SNMP manager's IP address. For more information, see ["Adding a Name Server."](#)

At the command prompt, type the following commands to set the parameters and verify the configuration:

- `add snmp manager <IPAddress> [-domainResolveRetry <integer>]`
- `show snmp manager`

Example

```
> add nameserver 10.103.128.15  
  
> add snmp manager engwiki.eng.example.net â€"domainResolveRetry 10
```

To add an SNMP manager by using the configuration utility

1. Navigate to System > SNMP > Managers, and create the SNMP manager.
Important: If you specify the SNMP manager's host name instead of its IPv4 address, you must configure a DNS name server to resolve the host name to the SNMP manager's IP address. For more information, see ["Adding a Name Server."](#)
Note: The appliance does not support host names for SNMP managers that have IPv6 addresses.

Specifying an SNMP Community

Updated: 2014-08-06

You can create strings called *community strings* and associate them with the following SNMP query types on the appliance:

- GET
- GET NEXT
- ALL
- GET BULK

You can associate one or more community strings to each query types. For example, when you associate two community strings, such as **abc** and **bcd**, to the query type GET NEXT, the SNMP agent on the appliance considers only those GET NEXT SNMP query packets that contain **abc** or **bcd** as the community string.

If you do not associate any community string to a query type then the SNMP agent responds to all SNMP queries of that type.

To specify an SNMP community by using the command line interface

At the command prompt, type the following commands to set the parameters and verify the configuration:

- add snmp community <communityName> <permissions>
- show snmp community

Example

```
> add snmp community com all
```

To configure an SNMP community string by using the configuration utility

Navigate to System > SNMP > Community, and create the SNMP community.

Configuring the NetScaler for SNMPv3 Queries

Simple Network Management Protocol Version 3 (SNMPv3) is based on the basic structure and architecture of SNMPv1 and SNMPv2. However, SNMPv3 enhances the basic architecture to incorporate administration and security capabilities, such as authentication, access control, data integrity check, data origin verification, message timeliness check, and data confidentiality.

To implement message level security and access control, SNMPv3 introduces the user-based security model (USM) and the view-based access control model (VACM).

- **User-Based Security Model.** The user-based security model (USM) provides message-level security. It enables you to configure users and security parameters for the SNMP agent and the SNMP manager. USM offers the following features:
 - Data integrity:** To protect messages from being modified during transmission through the network.
 - Data origin verification:** To authenticate the user who sent the message request.
 - Message timeliness:** To protect against message delays or replays.
 - Data confidentiality:** To protect the content of messages from being disclosed to unauthorized entities or individuals.
- **View-Based Access Control Model.** The view-based access control model (VACM) enables you to configure access rights to a specific subtree of the MIB based on various parameters, such as security level, security model, user name, and view type. It enables you to configure agents to provide different levels of access to the MIB to different managers.

The Citrix NetScaler supports the following entities that enable you to implement the security features of SNMPv3:

- SNMP Engines
- SNMP Views
- SNMP Groups
- SNMP Users

These entities function together to implement the SNMPv3 security features. Views are created to allow access to subtrees of the MIB. Then, groups are created with the required security level and access to the defined views. Finally, users are created and assigned to the groups.

Note: The view, group, and user configuration are synchronized and propagated to the secondary node in a high availability (HA) pair. However, the engine ID is neither propagated nor synchronized as it is unique to each NetScaler appliance. To implement message authentication and access control, you need to:

- Set the Engine ID
- Configure Views
- Configure Groups
- Configure Users

This document includes the following details:

- [Setting the Engine ID](#)
- [Configuring a View](#)
- [Configuring a Group](#)
- [Configuring a User](#)

Setting the Engine ID

Updated: 2014-08-06

SNMP engines are service providers that reside in the SNMP agent. They provide services such as sending, receiving, and authenticating messages. SNMP engines are uniquely identified using engine IDs.

The NetScaler appliance has a unique engineID based on the MAC address of one of its interfaces. It is not necessary to override the engineID. However, if you want to change the engine ID, you can reset it.

To set the engine ID by using the command line interface

At the command prompt, type the following commands to set the parameters and verify the configuration:

- `set snmp engineid <engineID>`

- o show snmp engineId

Example

```
> set snmp engineId 8000173f0300c095f80c68
```

To set the engine ID by using configuration utility

Navigate to System > SNMP > Users, click Configure Engine ID and type an engine ID.

Configuring a View

Updated: 2014-08-06

SNMP views restrict user access to specific portions of the MIB. SNMP views are used to implement access control.

To add an SNMP view by using the command line interface

At the command prompt, type the following commands to set the parameters and verify the configuration:

- o add snmp view <name> <subtree> -type (included | excluded)
- o show snmp view <name>

Example

```
> add snmp view View1 -type included
```

To configure an SNMP view by using the configuration utility

Navigate to System > SNMP > Views, and create the SNMP view.

Configuring a Group

Updated: 2014-08-06

SNMP groups are logical aggregations of SNMP users. They are used to implement access control and to define the security levels. You can configure an SNMP group to set access rights for users assigned to that group, thereby restricting the users to specific views.

You need to configure an SNMP group to set access rights for users assigned to that group.

To add an SNMP group by using the command line interface

At the command prompt, type the following commands to set the parameters and verify the configuration:

- o add snmp group <name> <securityLevel> -readViewName <string>
- o show snmp group <name> <securityLevel>

Example

```
> add snmp group edocs_group2 authPriv -readViewName edocs_read_view
```

To configure an SNMP group by using the configuration utility

Navigate to System > SNMP > Groups, and create the SNMP group.

Configuring a User

Updated: 2014-08-06

SNMP users are the SNMP managers that the agents allow to access the MIBs. Each SNMP user is assigned to an SNMP group.

You need to configure users at the agent and assign each user to a group.

To configure a user by using the command line interface

At the command prompt, type the following commands to set the parameters and verify the configuration:

- add snmp user <name> -group <string> [-authType (MD5 | SHA) {-authPasswd } [-privType (DES | AES) {-privPasswd }]]
- show snmp user <name>

Example

```
> add snmp user edocs_user -group edocs_group
```

To configure an SNMP user by using the configuration utility

Navigate to System > SNMP > Users, and create the SNMP user.

Configuring SNMP Alarms for Rate Limiting

Citrix NetScaler appliances such as the NetScaler MPX 10500, 12500, and 15500 are rate limited. The maximum throughput (Mbps) and packets per second (PPS) are determined by the license purchased for the appliance. For rate-limited platforms, you can configure SNMP traps to send notifications when throughput and PPS approach their limits and when they return to normal.

Throughput and PPS are monitored every seven seconds. You can configure traps with high-threshold and normal-threshold values, which are expressed as a percentage of the licensed limits. The appliance then generates a trap when throughput or PPS exceeds the high threshold, and a second trap when the monitored parameter falls to the normal threshold. In addition to sending the traps to the configured destination device, the NetScaler logs the events associated with the traps in the /var/log/ns.log file as EVENT ALERTSTARTED and EVENT ALERTENDED.

Exceeding the throughput limit can result in packet loss. You can configure SNMP alarms to report packet loss.

For more information about SNMP alarms and traps, see ["Configuring the NetScaler to generate SNMP v1 and v2 Traps"](#).

This document includes the following details:

- [Configuring an SNMP Alarm for Throughput or PPS](#)
- [Configuring SNMP Alarm for Dropped Packets](#)

Configuring an SNMP Alarm for Throughput or PPS

Updated: 2014-08-12

To monitor both throughput and PPS, you must configure separate alarms.

To configure an SNMP alarm for the throughput rate by using the command line interface

At the command prompt, type the following commands to configure the SNMP alarm and verify the configuration:

- `set snmp alarm PF-RL-RATE-THRESHOLD [-thresholdValue <positive_integer> [-normalValue <positive_integer>]] [-state (ENABLED | DISABLED)] [-severity <severity>] [-logging (ENABLED | DISABLED)]`
- `show snmp alarm PF-RL-RATE-THRESHOLD`

Example

```
> set snmp alarm PF-RL-RATE-THRESHOLD -thresholdValue 70 -normalValue 50
```

To configure an SNMP alarm for PPS by using the command line interface

At the command prompt, type the following commands to configure the SNMP alarm for PPS and verify the configuration:

- `set snmp alarm PF-RL-PPS-THRESHOLD [-thresholdValue <positive_integer> [-normalValue <positive_integer>]] [-state (ENABLED | DISABLED)] [-severity <severity>] [-logging (ENABLED | DISABLED)]`
- `show snmp alarm PF-RL-PPS-THRESHOLD`

Example

```
> set snmp alarm PF-RL-PPS-THRESHOLD -thresholdValue 70 -normalValue 50
```

To configure an SNMP alarm for throughput or PPS by using the configuration utility

1. Navigate to System > SNMP > Alarms, and select PF-RL-RATE-THRESHOLD (for throughput rate) or PF-RL-PPS-THRESHOLD (for packets per second).
2. Set the alarm parameters and enable the selected SNMP alarm.

Configuring SNMP Alarm for Dropped Packets

Updated: 2014-08-12

You can configure an alarm for packets dropped as a result of exceeding the throughput limit and an alarm for packets dropped as a result of exceeding the PPS limit.

To configure an SNMP alarm for packets dropped because of excessive throughput, by using the command line interface

At the command prompt, type:

```
set snmp alarm PF-RL-RATE-PKTS-DROPPED [-state (ENABLED | DISABLED)] [-severity <severity>] [-logging (ENABLED | DISABLED )]
```

To configure an SNMP alarm for packets dropped because of excessive PPS, by using the command line interface

At the command prompt, type:

```
set snmp alarm PF-RL-PPS-PKTS-DROPPED [-state (ENABLED | DISABLED)] [-severity <severity>] [-logging (ENABLED | DISABLED )]
```

To configure an SNMP alarm for dropped packets by using the configuration utility

1. Navigate to System > SNMP > Alarms, and select PF-RL-RATE-PKTS-DROPPED (for packets dropped because of excessive throughput) or PF-RL-PPS-PKTS-DROPPED (for packets dropped because of excessive PPS).
2. Set the alarm parameters and enable the selected SNMP alarm.

Audit Logging

Auditing is a methodical examination or review of a condition or situation. The Audit Logging feature enables you to log the NetScaler states and status information collected by various modules in the kernel and in the user-level daemons. For audit logging, you can use the SYSLOG protocol, the native NSLOG protocol, or both.

SYSLOG is a standard protocol for logging. It has two components: the SYSLOG auditing module, which runs on the NetScaler appliance, and the SYSLOG server, which can run on the underlying FreeBSD operating system (OS) of the NetScaler appliance or on a remote system. SYSLOG uses user data protocol (UDP) for data transfer.

Similarly, the native NSLOG protocol has two components—the NSLOG auditing module, which runs on the NetScaler appliance, and the NSLOG server, which can run on the underlying FreeBSD OS of the NetScaler appliance or on a remote system. NSLOG uses transmission control protocol (TCP) for data transfer.

When you run a SYSLOG or NSLOG server, it connects to the NetScaler appliance. The NetScaler appliance then starts sending all the log information to the SYSLOG or NSLOG server, and the server can filter the log entries before storing them in a log file. An NSLOG or SYSLOG server can receive log information from more than one NetScaler appliance, and a NetScaler appliance can send log information to more than one SYSLOG server or NSLOG server.

The log information that a SYSLOG or NSLOG server collects from a NetScaler appliance is stored in a log file in the form of messages. These messages typically contain the following information:

- The IP address of a NetScaler appliance that generated the log message
- A time stamp
- The message type
- The predefined log levels (Critical, Error, Notice, Warning, Informational, Debug, Alert, and Emergency)
- The message information

To configure audit logging, you first configure the audit modules on the NetScaler appliance. That involves creating audit policies and specifying the NSLOG server or SYSLOG server information. You then install and configure the SYSLOG or the NSLOG server on the underlying FreeBSD OS of the NetScaler appliance or on a remote system.

Note: Because SYSLOG is an industry standard for logging program messages, and various vendors provide support, this documentation does not include SYSLOG server configuration information.

The NSLOG server has its own configuration file (auditlog.conf). You can customize logging on the NSLOG server system by making additional modifications to the configuration file (auditlog.conf).

Configuring the NetScaler Appliance for Audit Logging

On the NetScaler appliance, you configure SYSLOG and/or NSLOG policies. Each policy includes a rule, which is an expression identifying the messages to be logged, and a SYSLOG or NSLOG (depending on the type of policy) action. The action specifies the server to which to send the log message, the level of the messages to be logged, and the data format of the logged messages. You can bind the policies globally or to individual virtual servers.

The appliance logs the following information related to TCP connections:

- Source port
- Destination port
- Source IP
- Destination IP
- Number of bytes transmitted and received
- Time period for which the connection is open

Note:

- You can enable TCP logging on individual load balancing virtual servers. You must bind the audit log policy to a specific load balancing virtual server that you want to log.
- When using the NetScaler as the audit log server, by default, the ns.log file is rotated (new file is created) when the file size reaches 100K and the last 25 copies of the ns.log are archived and compressed with gzip. To accommodate more archived files after 25 files, the oldest archive is deleted. You can modify the 100K limit or the 25 file limit by updating the following entry in the /etc/newsyslog.conf file:

```
/var/log/ns.log 600 25 100 * z
```

where, 25 is the number of archived files to be maintained and 100K is the size of the ns.log file after which the file will be archived.

This document includes the following details:

- [Configuring SYSLOG and NSLOG Actions](#)
- [Configuring Audit Policies](#)
- [Binding the Audit Policies Globally](#)
- [Configuring Policy-Based Logging](#)

Configuring SYSLOG and NSLOG Actions

Updated: 2015-06-03

You can configure audit server actions for different servers and for different log levels.

To configure a SYSLOG server action by using the command line interface

At the command prompt, type the following commands to set the parameters and verify the configuration:

- `add audit syslogAction <name> <serverIP> [-serverPort <port>] -logLevel <logLevel> [-dateFormat (MMDDYYYY | DDMMYYYY)]`
- `show audit syslogAction [<name>]`

Example

```
> add audit syslogaction audit-action1 10.102.1.1 -loglevel INFORMATIONAL -dateformat MMDDYY
```

To configure an NSLOG server action by using the command line interface

At the command prompt, type the following commands to set the parameters and verify the configuration:

- `add audit nslogAction <name> <serverIP> [-serverPort <port>] -logLevel <logLevel> [-dateFormat (MMDDYYYY | DDMMYYYY)]`
- `show audit nslogAction [<name>]`

Example

```
> add audit nslogAction nslog-action1 10.102.1.3 -serverport 520 -loglevel INFORMATIONAL -da
```

To configure an auditing server action by using the configuration utility

Navigate to System > Auditing > Syslog or Nslog, click Servers tab and create the auditing server.

Configuring Audit Policies

Updated: 2015-04-29

Configure SYSLOG policies to log messages to a SYSLOG server, and/or NSLOG policy to log messages to an NSLOG server. Each policy includes a rule identifying the messages to be logged, and a SYSLOG or NS LOG (depending on the type of policy) action.

To configure a SYSLOG policy by using the command line interface

At the command prompt, type the following commands to set the parameters and verify the configuration:

- o add audit syslogPolicy <name> <rule> <action>
- o show audit syslogPolicy [<name>]

Example

```
> add audit syslogpolicy syslog-poll ns_true audit-action1
```

To configure an NSLOG policy by using the command line interface

At the command prompt, type the following commands to set the parameters and verify the configuration:

- o add audit nslogPolicy <name> <rule> <action>
- o show audit nslogPolicy [<name>]

Example

```
> add audit nslogPolicy nslog-poll ns_true nslog-action1
```

To configure an audit server policy by using the configuration utility

Navigate to System > Auditing > Syslog or Nslog, click Policies tab and create the auditing policy.

Binding the Audit Policies Globally

Updated: 2015-06-02

You must globally bind the audit log policies to enable logging of all NetScaler system events. By defining the priority level, you can set the evaluation order of the audit server logging. Priority 0 is the highest and is evaluated first. The higher the priority number, the lower is the priority of evaluation.

To configure a SYSLOG policy by using the command line interface

At the command prompt, type:

- o bind system global [<policyName> [-priority <positive_integer>]]
- o show system global

Example

```
> bind system global nslog-poll -priority 20
```

To globally bind the audit policy by using the configuration utility

1. Navigate to System > Auditing > Syslog or Nslog.
2. On Policies tab, click Action, and select Global Bindings to bind the audit global policies.

Configuring Policy-Based Logging

Updated: 2014-08-07

You can configure policy-based logging for rewrite and responder policies. Audit messages are then logged in a defined format when the rule in a policy evaluates to TRUE. To configure policy-based logging, you configure an audit-message action that uses default syntax expressions to specify the format of the audit messages, and associate the action with a policy. The policy can be bound either globally or to a load balancing or content switching virtual server. You can use audit-message actions to log messages at various log levels, either in syslog format only or in both syslog and newnslog formats.

Pre Requisites

- User Configurable Log Messages (userDefinedAuditlog) option is enabled for when configuring the audit action server to which you want to send the logs in a defined format. For more information about enabling policy-based logging on an audit action server, see ["Binding the Audit Policies Globally."](#)
- The related audit policy is bound to system global. For more information about binding audit policies to system global, see ["Binding the Audit Policies Globally."](#)

Configuring an Audit Message Action

You can configure audit message actions to log messages at various log levels, either in syslog format only or in both syslog and newnslog formats. Audit-message actions use expressions to specify the format of the audit messages.

To create an audit message action by using the command line interface

At the command prompt, type:

```
add audit messageaction <name> <logLevel> <stringBuilderExpr> [-logtoNewnslog (YES|NO)] [-bypassSafetyCheck (YES|NO)]
```

Example

```
> add audit messageaction log-act1 CRITICAL '"Client:"+CLIENT.IP.SRC+" accessed "+HTTP.REQ.U
```

To configure an audit message action by using the configuration utility

Navigate to System > Auditing > Message Actions, and create the audit message action.

Binding Audit Message Action to a Policy

After you have created an audit message action, you must bind it to a rewrite or responder policy. For more information about binding log message actions to a rewrite or responder policy, see ["Rewrite"](#) or ["Responder"](#).

Installing and Configuring the NSLOG Server

During installation, the NSLOG server executable file (auditserver) is installed along with other files. The auditserver executable file includes options for performing several actions on the NSLOG server, including running and stopping the NSLOG server. In addition, you use the auditserver executable to configure the NSLOG server with the IP addresses of the NetScaler appliances from which the NSLOG server will start collecting logs. Configuration settings are applied in the NSLOG server configuration file (auditlog.conf).

Then, you start the NSLOG server by executing the auditserver executable. The NSLOG server configuration is based on the settings in the configuration file. You can further customize logging on the NSLOG server system by making additional modifications to the NSLOG server configuration file (auditlog.conf).

Attention: The version of the NSLOG server package must be the same as that of the NetScaler. For example, if the version of the NetScaler is 10.1 Build 125.9, the NSLOG server must also be of the same version.

The following table lists the operating systems on which the NSLOG server is supported.

Table 1. Supported Platforms for the NSLOG Server

Operating system	Software requirements	Remarks
Windows	<ul style="list-style-type: none">Windows XP ProfessionalWindows Server 2003Windows 2000/NTWindows Server 2008Windows Server 2008 R2	
Linux	<ul style="list-style-type: none">RedHat Linux 4 or laterSUSE Linux Enterprise 9.3 or later	
FreeBSD	FreeBSD 6.3 or later	For NetScaler 10.5, use only FreeBSD 8.4.
Mac OS	Mac OS 8.6 or later	Not supported on NetScaler 10.1 and later releases.

The minimum hardware specifications for the platform running the NSLOG server are as follows:

- Processor- Intel x86 ~501 megahertz (MHz)
- RAM - 512 megabytes (MB)
- Controller - SCSI

This document includes the following details:

- Installing NSLOG Server on the Linux Operating System
- Installing NSLOG Server on the FreeBSD Operating System
- Installing NSLOG Server Files on the Windows Operating System
- NSLOG Server Command Options
- Adding the NetScaler Appliance IP Addresses on the NSLOG Server
- Verifying the NSLOG Server Configuration File

Installing NSLOG Server on the Linux Operating System

Updated: 2013-06-20

Log on to the Linux system as an administrator. Use the following procedure to install the NSLOG server executable files on the system.

To install the NSLOG server package on a Linux operating system

1. At a Linux command prompt, type the following command to copy the NSauditserver.rpm file to a temporary directory:

```
cp <path_to_cd>/Utilities/auditserver/Linux/NSauditserver.rpm /tmp
```

2. Type the following command to install the NSauditserver.rpm file:

```
rpm -i NSauditserver.rpm
```

This command extracts the files and installs them in the following directories:

- o /usr/local/netscaler/etc
- o /usr/local/netscaler/bin
- o /usr/local/netscaler/samples

To uninstall the NSLOG server package on a Linux operating system

1. At a command prompt, type the following command to uninstall the audit server logging feature:

```
rpm -e NSauditserver
```

2. For more information about the NSauditserver RPM file, use the following command:

```
rpm -qpi *.rpm
```

3. To view the installed audit server files use the following command:

```
rpm -qpl *.rpm
```

*.rpm: Specifies the file name.

Installing NSLOG Server on the FreeBSD Operating System

Before you can install the NSLOG server, you have to copy the NSLOG package from the NetScaler product CD or download it from www.citrix.com. The NSLOG package has the following name format:

AuditServer_<release number>-<build number>.zip

For example: AuditServer_10.5-58.11.zip

This package contains files for all supported platforms: Linux, Windows, and FreeBSD. On a FreeBSD operating system, install the NSLOG package that has the following name format:

audserver_bsd-<release number>-<build number>.tgz

For example: audserver_bsd-10.5-58.11.tgz

To download NSLOG package from www.citrix.com

1. In a web browser, go to www.citrix.com.
2. In the menu bar, click Log In.
3. Enter your login credentials, and then click Log In.
4. In the menu bar, click Downloads.
5. From the **Select a product** list, select **NetScaler ADC**.
6. On the **NetScaler ADC** page, select the release for which you want to download the NSLOG package (for example, Release 10.5), and then select **Firmware**.
7. Under **Firmware**, select the NetScaler firmware for the build number for which you want to download the NSLOG package.
8. On the page that appears, scroll down, select **Audit Servers**, and click **Download File** next to the package that you want to download.

To install the NSLOG server package on a FreeBSD operating system

1. On the system to which you have downloaded the NSLOG package AuditServer_<release number>-<build number>.zip (for example, AuditServer_9.3-51.5.zip), extract the FreeBSD NSLOG server package audserver_bsd-<release number>-<build number>.tgz (for example, audserver_bsd-9.3-51.5.tgz) from the package.
2. Copy the FreeBSD NSLOG server package audserver_bsd-<release number>-<build number>.tgz (for example, audserver_bsd-9.3-51.5.tgz) to a directory on a system running FreeBSD OS.
3. At a command prompt for the directory into which the FreeBSD NSLOG server package was copied, run the following command to install the package:

```
pkg_add audserver_bsd-<release number>-<build number>.tgz
```

Example

```
pkg_add audserver_bsd-9.3-51.5.tgz
```

The following directories are extracted:

- o <root directory extracted from the FreeBSD NSLOG server package tgz file>\netscaler\bin (for example, /var/auditserver/netscaler/bin)
 - o <root directory extracted from the FreeBSD NSLOG server package tgz file>\netscaler\etc (for example, /var/auditserver/netscaler/etc)
 - o <root directory extracted from the FreeBSD NSLOG server package tgz file>\netscaler\samples (for example, /var/auditserver/samples)
4. At a command prompt, type the following command to verify that the package is installed:

```
pkg_info | grep NSaudserver
```

To uninstall the NSLOG server package on a FreeBSD operating system

At a command prompt, type:

```
pkg_delete NSaudserver
```

Installing NSLOG Server Files on the Windows Operating System

Updated: 2013-06-20

Before you can install the NSLOG server, you have to copy the NSLOG package from the NetScaler product CD or download it from www.citrix.com. The NSLOG package has the following name format AuditServer _<release number>-<build number>.zip (for example, AuditServer_9.3-51.5.zip). This package contains NSLOG installation packages for all supported platforms.

To download NSLOG package from www.Citrix.com

1. In a web browser, go to www.citrix.com.
2. In the menu bar, click Log In.
3. Enter your login credentials, and then click Log In.
4. In the menu bar, click Downloads.
5. Search to find the page that provides the appropriate release number and build.
6. On that page, under Audit Servers, click Download to download the NSLOG package, having the format AuditServer_<release number>-<build number>.zip, to your local system (for example, AuditServer_9.3-51.5.zip).

To install NSLOG server on a Windows operating system

1. On the system, where you have downloaded the NSLOG package AuditServer_<release number>-<build number>.zip (for example, AuditServer_9.3-51.5.zip), extract audserver_win-<release number>-<build number>.zip (for example, audserver_win-9.3-51.5.zip) from the package.
2. Copy the extracted file audserver_<release number>-<build number>.zip (for example, audserver_win-9.3-51.5.zip) to a Windows system on which you want to install the NSLOG server.
3. Unzip the audserver_<release number>-<build number>.zip file (for example, audserver_win-9.3-51.5.zip).
4. The following directories are extracted:
 - a. <root directory extracted from the Windows NSLOG server package zip file>\bin (for example, C:\audserver_win-9.3-51.5\bin)
 - b. <root directory extracted from the Windows NSLOG server package zip file>\etc (for example, C:\audserver_win-9.3-51.5\etc)
 - c. <root directory extracted from the Windows NSLOG server package zip file>\samples (for example, C:\audserver_win-9.3-51.5\samples)
5. At a command prompt, run the following command from the <root directory extracted from the Windows NSLOG server package zip file>\bin path:

```
audserver -install -f <directorypath>\auditlog.conf
```

<directorypath>: Specifies the path to the configuration file (auditlog.conf). By default, log.conf is under <root directory extracted from Windows NSLOG server package zip file>\samples directory. But you can copy auditlog.conf to your desired directory.

To uninstall the NSLOG server on a Windows operating system

At a command prompt, run the following from the <root directory extracted from Windows NSLOG server package zip file>\bin path:

```
audserver -remove
```

NSLOG Server Command Options

Updated: 2013-06-20

The following table describes the commands that you can use to configure audit server options.

Table 2. Audit Server Options

Audit server commands	Specifies
<code>audserver -help</code>	The available Audit Server options.
<code>audserver -addns -f <path to configuration file></code>	<p>The system that gathers the log transaction data.</p> <p>You are prompted to enter the IP address of the NetScaler appliance.</p> <p>Enter the valid user name and password.</p>
<code>audserver -verify -f <path to configuration file></code>	Check for syntax or semantic errors in the configuration file (for example, auditlog.conf).
<code>audserver -start -f <path to configuration file></code>	<p>Start audit server logging based on the settings in the configuration file (auditlog.conf).</p> <p>Linux only: To start the audit server as a background process, type the ampersand sign (&) at the end of the command.</p>
<code>audserver -stop</code> (Linux only)	<p>Stops audit server logging when audit server is started as a background process.</p> <p>Alternatively, use the Ctrl+C key to stop audit server logging.</p>
<code>audserver -install -f <path to configuration file></code> (Windows only)	Installs the audit server logging client as a service on Windows.
<code>audserver -startservice</code> (Windows Only)	<p>Start the audit server logging service, when you enter this command at a command prompt.</p> <p>You can also start audit server logging from Start > Control Panel > Services.</p> <p>Note: Audit server logging starts by using the configuration settings in the configuration file, for example, auditlog.conf file specified in the audit server install option.</p>
<code>audserver -stopservice</code> (Windows Only)	Stop audit server logging.
<code>audserver -remove</code>	Removes the audit server logging service from the registry.

Run the `audserver` command from the directory in which the audit server executable is present:

- o On Windows: \ns\bin
- o On Solaris and Linux: \usr\local\netscaler\bin

The audit server configuration files are present in the following directories:

- o On Windows: \ns\etc
- o On Linux: \usr\local\netscaler\etc

The audit server executable is started as ./auditserver in Linux and FreeBSD.

Adding the NetScaler Appliance IP Addresses on the NSLOG Server

In the configuration file (auditlog.conf), add the IP addresses of the NetScaler appliances whose events must be logged.

To add the IP addresses of the NetScaler appliance

At a command prompt, type the following command:

```
audserver -addns -f <directorypath>\auditlog.conf
```

<directorypath>: Specifies the path to the configuration file (auditlog.conf).

You are prompted to enter the information for the following parameters:

NSIP: Specifies the IP address of the NetScaler appliance, for example, 10.102.29.1.

Userid: Specifies the user name, for example, nsroot.

Password: Specifies the password, for example, nsroot.

If you add multiple NetScaler IP addresses (NSIP), and later you do not want to log all of the NetScaler appliance event details, you can delete the NSIPs manually by removing the NSIP statement at the end of the auditlog.conf file. For a high availability (HA) setup, you must add both primary and secondary NetScaler IP addresses to auditlog.conf by using the audserver command. Before adding the IP address, make sure the user name and password exist on the system.

Verifying the NSLOG Server Configuration File

Check the configuration file (audit log.conf) for syntax correctness to enable logging to start and function correctly.

To verify configuration, at a command prompt, type the following command:

```
audserver -verify -f <directorypath>\auditlog.conf
```

<directorypath>: Specifies the path to the configuration file (audit log.conf).

Running the NSLOG Server

To start audit server logging

Type the following command at a command prompt:

```
audserver -start -f <directorypath>\auditlog.conf
```

<directorypath>: Specifies the path to the configuration file (audit log.conf).

To stop audit server logging that starts as a background process in FreeBSD or Linux

Type the following command:

```
audserver -stop
```

To stop audit server logging that starts as a service in Windows

Type the following command:

```
audserver -stopservice
```

Customizing Logging on the NSLOG Server

You can customize logging on the NSLOG server by making additional modifications to the NSLOG server configuration file (log.conf). Use a text editor to modify the log.conf configuration file on the server system.

To customize logging, use the configuration file to define filters and log properties.

- **Log filters.** Filter log information from a NetScaler appliance or a set of NetScaler appliances.
- **Log properties.** Each filter has an associated set of log properties. Log properties define how to store the filtered log information.

This document includes the following details:

- [Creating Filters](#)
- [Specifying Log Properties](#)

Creating Filters

Updated: 2013-11-14

You can use the default filter definition located in the configuration file (audit log.conf), or you can modify the filter or create a new filter. You can create more than one log filter.

Note: For consolidated logging, if a log transaction occurs for which there is no filter definition, the default filter is used (if it is enabled.) The only way you can configure consolidated logging of all the NetScaler appliances is by defining the default filter.

To create a filter

At the command prompt, type the following command in the configuration file (auditlog.conf) :

filter <filterName> [IP <ip>] [NETMASK <mask>] [ON | OFF]

<filterName>: Specify the name of the filter (maximum of 64 alphanumeric characters).

<ip>: Specify the IP addresses.

<mask>: Specify the subnet mask to be used on a subnet.

Specify ON to enable the filter to log transactions, or specify OFF to disable the filter. If no argument is specified, the filter is ON

Examples

```
filter F1 IP 192.168.100.151 ON
```

To apply the filter F2 to IP addresses 192.250.100.1 to 192.250.100.254:

```
filter F2 IP 192.250.100.0 NETMASK 255.255.255.0 ON
```

filterName is a required parameter if you are defining a filter with other optional parameters, such as IP address, or the combination of IP address and Netmask.

Specifying Log Properties

Updated: 2013-11-13

Log properties associated with the filter are applied to all the log entries present in the filter. The log property definition starts with the key word BEGIN and ends with END as illustrated in the following example:

```
BEGIN <filtername>
    logFilenameFormat ...
    logDirectory ...
    logInterval ...
    logFileSizeLimit ....
END
```

Entries in the definition can include the following:

- o **LogFilenameFormat** specifies the file name format of the log file. The name of the file can be of the following types:

Static: A constant string that specifies the absolute path and the file name.

Dynamic: An expression that includes the following format specifiers:

- Date (%{format}t)
- % creates file name with NSIP

Example

```
LogFileNameFormat Ex%{m%d%y}t.log
```

This creates the first file name as Exmddyy.log. New files are named: Exmddyy.log.0, Exmddyy.log.1, and so on. In the following example, the new files are created when the file size reaches 100MB.

Example

```
LogInterval size
LogFileSize 100
LogFileNameFormat Ex%{m%d%y}t
```

Caution: The date format %t specified in the LogFilenameFormat parameter overrides the log interval property for that filter. To prevent a new file being created every day instead of when the specified log file size is reached, do not use %t in the LogFilenameFormat parameter.

- o **logDirectory** specifies the directory name format of the log file. The name of the file can be either of the following:
 - Static: Is a constant string that specifies the absolute path and file name.
 - Dynamic: Is an expression containing the following format specifiers:
 - Date (%{format}t)
 - % creates directory with NSIP

The directory separator depends on the operating system. In Windows, use the directory separator \.

Example:

```
LogDirectory dir1\dir2\dir3
```

In the other operating systems (Linux, FreeBSD, etc.), use the directory separator /.

- o **LogInterval** specifies the interval at which new log files are created. Use one of the following values:
 - Hourly: A file is created every hour. Default value.
 - Daily: A file is created every day at midnight.
 - Weekly: A file is created every Sunday at midnight.
 - Monthly : A file is created on the first day of the month at midnight.
 - None: A file is created only once, when audit server logging starts.
 - Size: A file is created only when the log file size limit is reached.

Example

```
LogInterval Hourly
```

- o **LogFileSizeLimit** specifies the maximum size (in MB) of the log file. A new file is created when the limit is reached.

Note that you can override the loginterval property by assigning size as its value.

The default LogFileSizeLimit is 10 MB.

Example

```
LogFileSizeLimit 35
```

Default Settings for the Log Properties

The following is an example of the default filter with default settings for the log properties:

```
begin default
  logInterval Hourly
  logFileSizeLimit 10
  logFilenameFormat      auditlog%{ %Y%m%d }t.log
end default
```

Following are two examples of defining the default filters:

Example 1

```
Filter f1 IP 192.168.10.1
```

This creates a log file for NSI 192.168.10.1 with the default values of the log in effect.

Example 2

```
Filter f1 IP 192.168.10.1
begin f1
  logFilenameFormat logfiles.log
end f1
```

This creates a log file for NSIP 192.168.10.1. Since the log file name format is specified, the default values of the other log properties are in effect.

Sample Configuration File (audit.conf)

Following is a sample configuration file:

```
#####
# This is the Auditserver configuration file
# Only the default filter is active
# Remove leading # to activate other filters
#####
MYIP <NSAuditserverIP>
MYPORT 3023
#       Filter filter_nsip  IP <Specify the NetScaler IP address to filter on > ON
#       begin filter_nsip
#           logInterval                Hourly
#           logFileSizeLimit           10
#           logDirectory               logdir\%A\
#           logFilenameFormat          nsip%{%d%m%Y}t.log
#       end filter_nsip
Filter default
begin default
    logInterval                Hourly
    logFileSizeLimit           10
    logFilenameFormat          auditlog%{%y%m%d}t.log
end default
```

Web Server Logging

You can use the Web server logging feature to send logs of HTTP and HTTPS requests to a client system for storage and retrieval. This feature has two components:

- The Web log server, which runs on the NetScaler.
- The NetScaler Web Logging (NSWL) client, which runs on the client system.

When you run the NetScaler Web Logging (NSWL) client:

1. It connects to the NetScaler.
2. The NetScaler buffers the HTTP and HTTPS request log entries before sending them to the client.
3. The client can filter the entries before storing them.

To configure Web server logging, you first enable the Web logging feature on the NetScaler and configure the size of the buffer for temporarily storing the log entries. Then, you install NSWL on the client system. You then add the NetScaler IP address (NSIP) to the NSWL configuration file. You are now ready to start the NSWL client to begin logging. You can customize Web server logging by making additional modifications to the NSWL configuration file (log.conf).

Configuring the NetScaler for Web Server Logging

To configure the NetScaler for web server logging you are required to only enable the Web Server Logging feature. Optionally, you can perform the following configurations:

- Modify the size of the buffer (default size is 16 MB) that stores the logged information before it is sent to the NetScaler Web Logging (NSWL) client.
- Specify the custom HTTP headers that you want to export to the NSWL client. You can configure a maximum of two HTTP request and two HTTP response header names.

To configure web server logging by using the command line interface

At the command prompt, perform the following operations:

- Enable the web server logging feature.

```
enable ns feature WL
```

- [Optional] Modify the buffer size for storing the logged information.

```
set ns weblogparam -bufferSizeMB <size>
```

Note: To activate your modification, you must disable and then re-enable the Web server logging feature.

- [Optional] Specify the custom HTTP header names that you want to export.

```
set ns weblogparam [-customReqHdrs <string> ...] [-customRspHdrs <string> ...]
```

Example

```
> set ns weblogparam -customReqHdrs Accept-Encoding X-Forwarded -customRspHdrs (
```

To configure web server logging by using the configuration utility

Navigate to System > Settings and perform the following operations:

- To enable the web server logging feature, click Change Advanced Features and select Web Logging.
- To modify the buffer size, click Change Global System Settings and under Web Logging, enter the buffer size.
- To specify the custom HTTP headers to be exported, click Change Global System Settings and under Web Logging, specify the header values.

Installing the NetScaler Web Logging (NSWL) Client

During installation, the NSWL client executable file (nswl) is installed along with other files. The nswl executable file provides a list of options that you can use. For details, see [Configuring the NSWL Client](#).

Attention: The version of the NSWL client must be the same as that of the NetScaler. For example, if the version of the NetScaler is 10.1 Build 125.9, the NSWL client must also be of the same version.

The following table lists the operating systems on which the NSWL client can be installed.

Table 1. Supported Platforms for the NSWL Client with hardware requirements

Operating system	Version	Hardware requirements	Remarks
Windows	<ul style="list-style-type: none">Windows XP ProfessionalWindows Server 2003Windows 2000/NTWindows Server 2008Windows Server 2008 R2	Processor - Intel x86 ~501 MHz RAM - 512 MB Controller - SCSI	Â
Mac OS	Mac OS 8.6 or later	-	Not supported on NetScaler 10.1 and later releases.
Linux	<ul style="list-style-type: none">RedHat Linux 4 or laterSUSE Linux Enterprise 9.3 or later	Processor - Intel x86 ~501 MHz RAM - 512 MB Controller - SCSI	Â
Solaris	Solaris Sun OS 5.6 or later	Processor - UltraSPARC-IIi 400 MHz RAM - 512 MB Controller - SCSI	Not supported on NetScaler 10.5 and later releases.
FreeBSD	FreeBSD 6.3 or later	Processor - Intel x86 ~501 MHz RAM - 512 MB Controller - SCSI	For NetScaler 10.5, use only FreeBSD 8.4.
AIX	AIX 6.1	-	Not supported on NetScaler 10.5 and later releases.

If the NSWL client system cannot process the log transaction because of a CPU limitation, the Web log buffer overruns and the logging process reinitiates.

Caution: Reinitiation of logging can result in loss of log transactions.

To temporarily solve a NSWL client system bottleneck caused by a CPU limitation, you can tune the Web server logging buffer size on the NetScaler appliance. To solve the problem, you need a client system that can handle the site's throughput.

This document includes the following details:

- [Downloading the NSWL Client](#)

- o Installing the NSWL Client on a Solaris System
- o Installing the NSWL Client on a Linux System
- o Installing the NSWL Client on a FreeBSD System
- o Installing the NSWL Client on a Mac System
- o Installing the NSWL Client on a Windows System
- o Installing the NSWL Client on a AIX System

Downloading the NSWL Client

Updated: 2014-06-25

You can obtain the NSWL client package from either the NetScaler product CD or the Citrix downloads site. Within the package there are separate installation packages for each supported platforms.

To download the NSWL client package from the Citrix site

1. Open the URL: <https://www.citrix.com/downloads.html>.
2. Log in to the site using your credentials.
3. Open the page for the required release number and build.
4. In the page, under Weblog Clients, click Download. The package has the name format as follows: Weblog-<release number>-<build number>.zip.

Installing the NSWL Client on a Solaris System

Updated: 2014-06-25

To install the NSWL client, perform the following operations on the system where you downloaded the package.

1. Extract the nswl_solaris-<release number>-<build number>.tar file from the package.
2. Copy the extracted file to a Solaris system on which you want to install the NSWL client.
3. Extract the files from the tar file with the following command:

```
tar xvf nswl_solaris-9.3-51.5.tar
```

A directory NSweblog is created in the temporary directory, and the files are extracted to the NSweblog directory.

4. Install the package with the following command:

```
pkgadd -d
```

The list of available packages appears. In the following example, one NSweblog package is shown:

```
1 NSweblog NetScaler Weblogging (SunOS,sparc) 7.0
```

5. You are prompted to select the packages. Select the package number of the NSweblog to be installed. After you select the package number and press Enter, the files are extracted and installed in the following directories:
 - o /usr/local/netscaler/etc
 - o /usr/local/netscaler/bin
 - o /usr/local/netscaler/samples
6. To check whether the NSWL package is installed, execute the following command:

```
pkginfo | grep NSweblog
```

Note: To uninstall the NSWL package, execute the following command:

```
pkgrm NSweblog
```

Installing the NSWL Client on a Linux System

Updated: 2014-06-25

To install the NSWL client, perform the following operations on the system where you downloaded the package.

1. Extract the nswl_linux-<release number>-<build number>.rpm file from the package.
2. Copy the extracted file to a system, running Linux OS, on which you want to install the NSWL client.
3. To install the NSWL package, execute the following command:

```
rpm -i nswl_linux-9.3-51.5.rpm
```

This command extracts the files and installs them in the following directories.

- o /usr/local/netscaler/etc
- o /usr/local/netscaler/bin
- o /usr/local/netscaler/samples

Note: To uninstall the NSWL package, execute the following command:

```
rpm -e NSweblog
```

Note: To get more information about the NSweblog RPM file, execute the following command:

```
rpm -qpi *.rpm
```

Note: To view the installed Web server logging files, execute the following command:

```
rpm -qpl *.rpm
```

Installing the NSWL Client on a FreeBSD System

Updated: 2014-06-25

To install the NSWL client, perform the following operations on the system where you downloaded the package.

1. Extract the nswl_bsd-<release number>-<build number>.tgz file from the package.
2. Copy the extracted file to a system, running FreeBSD OS, on which you want to install the NSWL client.
3. To install the NSWL package, execute the following command:

```
pkg_add nswl_bsd-9.3-51.5.tgz
```

This command extracts the files and installs them in the following directories.

- o /usr/local/netscaler/etc
- o /usr/local/netscaler/bin
- o /usr/local/netscaler/samples

Note: To uninstall the NSWL package, execute the following command:

```
pkg_delete NSweblog
```

4. To verify that the package is installed, execute the following command:

```
pkg_info | grep NSweblog
```

Installing the NSWL Client on a Mac System

Updated: 2014-06-25

To install the NSWL client, perform the following operations on the system where you downloaded the package.

1. Extract the nswl_macos-<release number>-<build number>.tgz file from the package.
2. Copy the extracted file to a system, running Mac OS, on which you want to install the NSWL client.
3. To install the NSWL package, execute the following command:

```
pkg_add nswl_macos-9.3-51.5.tgz
```

This command extracts the files and installs them in the following directories:

- o /usr/local/netscaler/etc
- o /usr/local/netscaler/bin
- o /usr/local/netscaler/samples

Note: To uninstall the NSWL package, execute the following command:

```
pkg_delete NSweblog
```

4. To verify that the package is installed, execute the following command:

```
pkg_info | grep NSweblog
```

Installing the NSWL Client on a Windows System

Updated: 2014-09-18

To install the NSWL client, perform the following operations on the system where you downloaded the package.

1. Extract the nswl_win-<release number>-<build number>.zip file from the package.
2. Copy the extracted file to a Windows system on which you want to install the NSWL client.

3. On the Windows system, unzip the file in a directory (referred as <NSWL-HOME>). The following directories are extracted: bin, etc, and samples.
4. At the command prompt, run the following command from the <NSWL-HOME>\bin directory:

```
nswl -install -f <directorypath>\log.conf
```

where,

<directorypath> refers to the path of the configuration file (log.conf). By default, the file is in the <NSWL-HOME>\etc directory. However, you can copy the configuration file to any other directory.

Note: To uninstall the NSWL client, at the command prompt, run the following command from the <NSWL-HOME>\bin directory:

```
> nswl -remove
```

Installing the NSWL Client on a AIX System

Updated: 2014-06-25

To install the NSWL client, perform the following operations on the system where you downloaded the package.

1. Extract the nswl_aix-<release number>-<build number>.rpm file from the package.
2. Copy the extracted file to a system, running AIX OS, on which you want to install the NSWL client.
3. To install the NSWL package, execute the following command:

```
rpm -i nswl_aix-9.3-51.5.rpm
```

This command extracts the files and installs them in the following directories.

- o /usr/local/netscaler/etc
- o /usr/local/netscaler/bin
- o /usr/local/netscaler/samples

Note: To uninstall the NSWL package, execute the following command:

```
rpm -e NSweblog
```

Note: To get more information about the NSweblog RPM file, execute the following command:

```
rpm -qpi *.rpm
```

Note: To view the installed Web server logging files, execute the following command:

```
rpm -qpl *.rpm
```

Configuring the NSWL Client

After installing the NSWL client, you can configure the NSWL client using the nswl executable. These configurations are then stored in the NSWL client configuration file (log.conf).

Note: You can further customize logging on the NSWL client system by making additional modifications to the NSWL configuration file (log.conf). For details, see [Customizing Logging on the NSWL Client System](#).

The following table describes the commands that you can use to configure the NSWL client.

NSWL command	Specifies
nswl -help	The available NSWL help options.
nswl -addns -f <path-to-configuration-file>	The system that gathers the log transaction data. You are prompted to enter the IP address of the NetScaler appliance. Enter a valid user name and password.
nswl -verify -f <path-to-configuration-file>	Check for syntax or semantic errors in the configuration file.
nswl -start -f <path-to-configuration-file>	Start the NSWL client based on the settings in the configuration file. Note: For Solaris and Linux: To start Web server logging as a background process, type the ampersand sign (&) at the end of the command.
nswl -stop (Solaris and Linux only)	Stop the NSWL client if it was started as a background process; otherwise, use CTRL+C to stop Web server logging.
nswl -install -f <path-to-configuration-file> (Windows only)	Install the NSWL client as a service in Windows.
nswl -startservice (Windows only)	Start the NSWL client by using the settings in the configuration file specified in the nswl install option. You can also start NSWL client from Start > Control Panel > Services.
nswl -stopservice (Windows only)	Stops the NSWL client.
nswl -remove	Remove the NSWL client service from the registry.

Run the following commands from the directory in which the NSWL executable is located:

- Windows: \ns\bin
- Solaris and Linux: \usr\local\netscaler\bin

The Web server logging configuration files are located in the following directory path:

- Windows: \ns\etc
- Solaris and Linux: \usr\local\netscaler\etc

The NSWL executable is started as .nswl in Linux and Solaris.

This document includes the following details:

- [Adding the IP Addresses of the NetScaler Appliance](#)
- [Verifying the NSWL Configuration File](#)
- [Running the NSWL Client](#)

Adding the IP Addresses of the NetScaler Appliance

Updated: 2013-07-17

In the NSWL client configuration file (log.conf), add the NetScaler IP address (NSIP) from which the NSWL client will start collecting logs.

To add the NSIP address of the NetScaler appliance

- At the client system command prompt, type:

```
nswl -addns -f <directorypath> \log.conf
```

< directorypath >: Specifies the path to the configuration file (log.conf).

2. At the next prompt, enter the following information:

- o **NSIP:** Specify the IP address of the NetScaler appliance.
- o **Username and Password:** Specify the `nsroot` user credentials of the NetScaler appliance.

Note: If you add multiple NetScaler IP addresses (NSIP), and later you do not want to log all of NetScaler system log details, you can delete the NSIPs manually by removing the NSIP statement at the end of the log.conf file. During a failover setup, you must add both primary and secondary NetScaler IP addresses to the log.conf by using the command. Before adding the IP address, make sure the user name and password exist on the NetScaler appliances.

Verifying the NSWL Configuration File

Updated: 2013-06-20

To make sure that logging works correctly, check the NSWL configuration file (log.conf) on the client system for syntax errors

To verify the configuration in the NSWL configuration file

At the client system command prompt, type:

```
nswl -verify -f <directorypath>\log.conf
```

< directorypath >: Specifies the path to the configuration file (log.conf).

Running the NSWL Client

Updated: 2013-06-20

To start Web server logging

At the client system command prompt, type:

```
nswl -start -f <directorypath>\log.conf
```

<directorypath>: Specifies the path to the configuration file (log.conf).

To stop Web server logging started as a background process on the Solaris or Linux operating systems

At the command prompt, type:

```
nswl -stop
```

To stop Web server logging started as a service on the Windows operating system

At the command prompt, type:

```
nswl -stopservice
```

Customizing Logging on the NSWL Client System

You can customize logging on the NSWL client system by making additional modifications to the NSWL client configuration file (log.conf). Use a text editor to modify the log.conf configuration file on the client system.

To customize logging, use the configuration file to define filters and log properties.

- **Log filters.** Filter log information based on the host IP address, domain name, and host name of the Web servers.
- **Log properties.** Each filter has an associated set of log properties. Log properties define how to store the filtered log information.

This document includes the following details:

- [Creating Filters](#)
- [Specifying Log Properties](#)
- [Understanding the NCSA and W3C Log Formats](#)
- [Creating a Custom Log Format](#)
- [Arguments for Defining a Custom Log Format](#)
- [Time Format Definition](#)
- [Sample Configuration File](#)

Sample Configuration File

Following is a sample configuration file:

```
#####
# This is the NSWL configuration file
# Only the default filter is active
# Remove leading # to activate other filters
#####
#####
# Default filter (default on)
# W3C Format logging, new file is created every hour or on reaching 10MB file size,
# and the file name is Exyyymmdd.log
#####
Filter default
begin default
    logFormat            W3C
    logInterval          Hourly
    logFileSizeLimit     10
    logFilenameFormat    Ex%{ %y%m%d }t.log
end default
#####
# netscaler caches example
# CACHE_F filter covers all the transaction with HOST name www.netscaler.com and the listed
#####
#Filter CACHE_F HOST www.netscaler.com IP 192.168.100.89 192.168.100.95 192.168.100.52 192.1
#####
# netscaler origin server example
# Not interested in Origin server to Cache traffic transaction logging
#####
#Filter ORIGIN_SERVERS IP 192.168.100.64 192.168.100.65 192.168.100.66 192.168.100.67 192.16
100.227 192.168.100.228 OFF
#####
# netscaler image server example
# all the image server logging.
#####
#Filter IMAGE_SERVER HOST www.netscaler.images.com IP 192.168.100.71 192.168.100.72 192.168.
0.171 ON
#####
# NCSA Format logging, new file is created every day midnight or on reaching 20MB file size,
# and the file name is /datadisk5/netscaler/log/NS<hostname>/Nsmmddyy.log.
# Exclude objects that ends with .gif .jpg .jar.
#####
#begin ORIGIN_SERVERS
#    logFormat            NCSA
#    logInterval          Daily
```

```

#       logFileSizeLimit           40
#       logFilenameFormat          /datadisk5/ORIGIN/log/%v/NS{%m%d%y}t.log
#       logExclude                  .gif .jpg .jar
#end ORIGIN_SERVERS

#####
# NCSA Format logging, new file is created every day midnight or on reaching 20MB file size,
# and the file name is /datadisk5/netcaler/log/NS<hostname>/Nsmmddyy.log with log record ti
#####
#begin CACHE_F
#       logFormat                   NCSA
#       logInterval                 Daily
#       logFileSizeLimit            20
#       logFilenameFormat /datadisk5/netcaler/log/%v/NS{%m%d%y}t.log
#       logtime                     GMT
#end CACHE_F

#####
# W3C Format logging, new file on reaching 20MB and the log file path name is
# atadisk6/netcaler/log/server's ip/Exmmyydd.log with log record timestamp as LOCAL.
#####
#begin IMAGE_SERVER
#       logFormat                   W3C
#       logInterval                 Size
#       logFileSizeLimit            20
#       logFilenameFormat /datadisk6/netcaler/log/%AEx{%m%d%y}t
#       logtime                     LOCAL
#end IMAGE_SERVER

#####
# Virtual Host by Name firm, can filter out the logging based on the host name by,
#####

#Filter VHOST_F IP 10.101.2.151 NETMASK 255.255.255.0
#begin VHOST_F
#       logFormat                   W3C
#       logInterval                 Daily
#       logFileSizeLimit            10
logFilenameFormat /ns/prod/vhost/%v/Ex{%m%d%y}t
#end VHOST_F

##### END FILTER CONFIGURATION #####

```

Creating Filters

Updated: 2014-01-06

You can use the default filter definition located in the configuration file (log.conf), or you can modify the filter or create a new filter. You can create more than one log filter.

Note: Consolidated logging, which logs transactions for which no filter is defined, uses the default filter if it is enabled. Consolidated logging of all servers can be done by defining only the default filter.

If the server hosts multiple Web sites and each Web site has its own domain name, and each domain is associated with a virtual server, you can configure Web server logging to create a separate log directory for each Web site. The following table displays the parameters for creating a filter.

Table 1. Parameters for Creating a Filter

Parameter	Specifies
filterName	Name of the filter. The filter name can include alphanumeric characters and cannot be longer than 59 characters. Filter names longer than 59 characters are truncated to 59 characters.
HOST name	Host name of the server for which the transactions are being logged.
IP ip	IP address of the server for which transactions are to be logged (for example, if the server has multiple domains that have one IP address).
IP ip 2...ip n:	Multiple IP addresses (for example, if the server domain has multiple IP addresses).
ip6 ip	IPv6 address of the server for which transactions are to be logged.

IP ip NETMASK mask	IP addresses and netmask combination to be used on a subnet.
ON OFF	Enable or disable the filter to log transactions. If no argument is selected, the filter is enabled (ON).

To create a filter

To create a filter, enter the following command in the log.conf file:

- o `filter <filterName> <HOST name> | [IP<ip>] | [IP<ip 2...ip n>] | <IP ip NETMASK mask> [ON | OFF]`
- o `filter <filterName> <HOST name> | [IP6 ip/<prefix length>] [ON | OFF]`

To create a filter for a virtual server

To create a filter for a virtual server, enter the following command in the log.conf file:

`filter <filterName> <VirtualServer IP address>`

Example

In the following example, you specify an IP address of 192.168.100.0 and netmask of 255.255.255.0. The filter applies to IP addresses 192.168.100.1 through 192.168.100.254.

```
Filter F1 HOST www.netscaler.com ON
Filter F2 HOST www.netscaler.com IP 192.168.100.151 ON
Filter F3 HOST www.netscaler.com IP 192.168.100.151 192.165.100.152 ON
Filter F4 IP 192.168.100.151
Filter F5 IP 192.168.100.151 HOST www.netscaler.com OFF
Filter F6 HOST www.netscaler.com HOST www.xyz.com HOST www.abcxyz.com IP 192.168.100.200 ON
Filter F7 IP 192.250.100.0 NETMASK 255.255.255.0
Filter F8 HOST www.xyz.com IP 192.250.100.0 NETMASK 255.255.255.0 OFF
For creating filters for servers having IPv6 addresses.
Filter F9 2002::8/112 ON
Filter F10 HOST www.abcd.com IP6 2002::8 ON
```

Specifying Log Properties

Log properties are applied to all log entries associated with the filter. The log property definition begins with the keyword BEGIN and ends with END as illustrated in the following example:

```
BEGIN <filtername>
logFormat ...
logFilenameFormat ...
logInterval ...
logFileSize ....
logExclude ....
logTime &#x2013;.
END
```

Entries in the definition can include the following:

- o **LogFormat** specifies the Web server logging feature that supports NCSA, W3C Extended, and custom log file formats.

By default, the logformat property is w3c. To override, enter custom or NCSA in the configuration file, for example:

```
LogFormat NCSA
```

Note: For the NCSA and custom log formats, local time is used to time stamp transactions and for file rotation.

- o **LogInterval** specifies the intervals at which new log files are created. Use one of the following values:
 - Hourly: A file is created every hour.
 - Daily: A file is created every day at midnight. Default value.
 - Weekly: A file is created every Sunday at midnight.

Monthly: A file is created on the first day of the month at midnight.
None: A file is created only once, when Web server logging starts.

Example

```
LogInterval Daily
```

- o **LogFileSizeLimit** specifies the maximum size of the log file in MB. It can be used with any log interval (weekly, monthly, and so on.) A file is created when the maximum file size limit is reached or when the defined log interval time elapses.

To override this behavior, specify the size as the loginterval property so that a file is created only when the log file size limit is reached.

The default LogFileSizeLimit is 10 MB.

Example

```
LogFileSizeLimit 35
```

- o **LogFilenameFormat** specifies the file name format of the log file. The name of the file can be of the following types:

Static: Specifies a constant string that contains the absolute path and file name.

Dynamic: Specifies an expression containing the following format:

- Server IP address (%A)
- Date (%{format}t)
- URL suffix (%x)
- Host name (%v)

Example

```
LogFileNameFormat Ex%{%m%d%y}t.log
```

This command creates the first file name as Exmmdyy.log, then every hour creates a file with file name: Exmmdyy.log.0, Exmmdyy.log.1,..., Exmmdyy.log.n.

Example

```
LogInterval size
LogFileSize 100
LogFileNameFormat Ex%{%m%d%y}t
```

Caution: The date format %t specified in the LogFilenameFormat command overrides the log interval property for that filter. To prevent a new file being created every day instead of when the specified log file size is reached, do not use %t in the LogFilenameFormat.

- o **LogExclude** prevents logging of transactions with the specified file extensions.

Example

```
LogExclude .html
```

This command creates a log file that excludes log transactions for *.html files.

- o **LogTime** specifies log time as either GMT or LOCAL.

The defaults are:

- NCSA log file format: LOCAL
- W3C log file format: GMT.

Understanding the NCSA and W3C Log Formats

The NetScaler supports the following standard log file formats:

- o NCSA Common Log Format
- o W3C Extended Log Format

NCSA Common Log Format

If the log file format is NCSA, the log file displays log information in the following format:

```
Client_IP_address -User_Name [Date:Time -TimeZone] "Method Object HTTP_version"
HTTP_StatusCode BytesSent
```

To use the NCSA Common log format, enter NCSA in the LogFormat argument in the log.conf file.

The following table describes the NCSA Common log format.

Table 2. NCSA Common Log Format

Argument	Specifies
Client_IP_address	The IP address of the client computer.
User Name	The user name.
Date	The date of the transaction.
Time	The time when the transaction was completed.
Time Zone	The time zone (Greenwich Mean Time or local time).
Method	The request method (for example; GET, POST).
Object	The URL.
HTTP_version	The version of HTTP used by the client.
HTTP_StatusCode	The status code in the response.
Bytes Sent	The number of bytes sent from the server.

W3C Extended Log Format

An extended log file contains a sequence of lines containing ASCII characters terminated by either a Line Feed (LF) or the sequence Carriage Return Line Feed (CRLF.) Log file generators must follow the line termination convention for the platform on which they are run.

Log analyzers must accept either LF or CRLF form. Each line may contain either a directive or an entry. If you want to use the W3C Extended log format, enter W3C as the Log-Format argument in the log.conf file.

By default, the standard W3C log format is defined internally as the custom log format, shown as follows:

```
%{%Y-%m-%d%H:%M:%S}t %a %u %S %A %p %m %U %q %s %j %J %T %H %+{user-agent}i %+{cookie} i%+{
```

For a description of the meaning of this each custom format, see "[Appendix A: Arguments for Defining a Custom Log Format](#)." You can also change the order or remove some fields in this W3C log format. For example:

```
logFormat W3C {%Y-%m-%d%H:%M:%S}t %m %U
```

W3C log entries are created with the following format:

```
#Version: 1.0
#Fields: date time cs-method cs-uri
#Date: 12-Jun-2001 12:34
2001-06-12 12:34:23 GET /sports/football.html
2001-06-12 12:34:30 GET /sports/football.html
```

Entries

Entries consist of a sequence of fields relating to a single HTTP transaction. Fields are separated by white space; Citrix recommends the use of tab characters. If a field in a particular entry is not used, a dash (-) marks the omitted field.

Directives

Directives record information about the logging process. Lines beginning with the pound sign (#) contain directives.

The following table describes the directives.

Table 3. Directive Descriptions

Directive	Description
Version: <integer>. <integer>	Displays the version of the extended log file format used. This document defines version 1.0.
Fields: [<specifier>...]	Identifies the fields recorded in the log.
Software: <string>	Identifies the software that generated the log.
Start-Date: <date> <time>	Displays the date and time at which the log was started.
End-Date: <date> <time>	Displays the date and time at which logging finished.
Date: <date> <time>	Displays the date and time when the entry was added.
Remark: <text>	Displays comments. Analysis tools ignore data recorded in this field.

Note: The Version and Fields directives are required. They precede all other entries in the log file.

Example

The following sample log file shows the log entries in W3C Extended log format:

```
#Version: 1.0
#Fields: time cs-method cs-uri
#Date: 12-Jan-1996 00:00:00
00:34:23 GET /sports/football.html
12:21:16 GET /sports/football.html
12:45:52 GET /sports/football.html
12:57:34 GET /sports/football.html
```

Fields

The Fields directive lists a sequence of field identifiers that specify the information recorded in each entry. Field identifiers may have one of the following forms:

- o **identifier:** Relates to the transaction as a whole.
- o **prefix-identifier:** Relates to information transfer between parties defined by the value *prefix*.
- o **prefix (header):** Specifies the value of the HTTP header field header for transfer between parties defined by the value *prefix*. Fields specified in this manner always have the type <string>.

The following table describes defined prefixes.

Table 4. Prefix Descriptions

Prefix	Specifies
c	Client
s	Server
r	Remote
cs	Client to server
sc	Server to client
sr	Server to remote server (prefix used by proxies)
rs	Remote server to server (prefix used by proxies)
x	Application-specific identifier

Examples

The following examples are defined identifiers that use prefixes:

cs-method: The method in the request sent by the client to the server.

sc(Referer): The Referer field in the reply.

c-ip: The IP address of the client.

Identifiers

The following table describes the W3C Extended log format identifiers that do not require a prefix.

Table 5. W3C Extended Log Format Identifiers (No Prefix Required)

Identifier	Description
date	The date on which the transaction was done.
time	The time when the transaction is done.
time-taken	The time taken (in seconds) for the transaction to complete.
bytes	The number of bytes transferred.
cached	Records whether a cache hit has occurred. A zero indicates a cache miss.

The following table describes the W3C Extended log format identifiers that require a prefix.

Table 6. W3C Extended Log Format Identifiers (Requires a Prefix)

Identifier	Description
IP	The IP address and the port number.
dns	The DNS name.
status	The status code.
comment	The comment returned with status code.
method	The method.
url	The URL.
url-stem	The stem portion of the URL.
url-query	The query portion of the URL.

The W3C Extended Log file format allows you to choose log fields. These fields are shown in the following table.

Table 7. W3C Extended Log File Format (Allows Log Fields)

Field	Description
Date	The date on which the transaction is done.
Time	The time when the transaction is done.
Client IP	The IP address of the client.
User Name	The user name.
Service Name	The service name, which is always HTTP.
Server IP	The server IP address.
Server Port	The server port number
Method	The request method (for example; GET, POST).
Url Stem	The URL stem.
Url Query	The query portion of the URL.
Http Status	The status code in the response.
Bytes Sent	The number of bytes sent to the server (request size, including HTTP headers).
Bytes Received	The number of bytes received from the server (response size, including HTTP headers).
Time Taken	The time taken for transaction to complete, in seconds.
Protocol Version	The version number of HTTP being used by the client.
User Agent	The User-Agent field in the HTTP protocol.
Cookie	The Cookie field of the HTTP protocol.
Referer	The Referer field of the HTTP protocol.

Creating a Custom Log Format

Updated: 2013-09-30

You can customize the display format of the log file data manually or by using the NSWL library. By using the custom log format, you can derive most of the log formats that Apache currently supports.

Creating a Custom Log Format by Using the NSWL Library

Use one of the following NSWL libraries depending on whether the NSWL executable has been installed on a Windows or Solaris host computer:

- o **Windows:** The nswl.lib library located in \ns\bin directory on the system manager host computer.
- o **Solaris:** The libnswl.a library located in /usr/local/netcaler/bin.

To create the custom log format by using the NSWL Library

1. Add the following two C functions defined by the system in a C source file:

ns_userDefFieldName() : This function returns the string that must be added as a custom field name in the log record.

ns_userDefFieldVal() : This function implements the custom field value, then returns it as a string that must be added at the end of the log record.

2. Compile the file into an object file.
3. Link the object file with the NSWL library (and optionally, with third party libraries) to form a new NSWL executable.
4. Add a %d string at the end of the logFormat string in the configuration file (log.conf).

Example

```
#####
# A new file is created every midnight or on reaching 20MB file size,
# and the file name is /datadisk5/netcaler/log/NS<hostname>/Nsmmddyy.log and create digital
#signature field for each record.
BEGIN CACHE_F
    logFormat          custom "%a - %{user-agent}i" [%d/%B/%Y %T -%g] "%x" %s %b%{referr
    logInterval        Daily
    logFileSizeLimit    20
    logFilenameFormat   /datadisk5/netcaler/log/%v/NS%{m%dy}t.log
END CACHE_F
```

Creating a Custom Log Format Manually

To customize the format in which log file data should appear, specify a character string as the argument of the LogFormat log property definition. For more information, see ["Appendix A: Arguments for Defining a Custom Log Format"](#). The following is an example where character strings are used to create a log format:

```
LogFormat Custom "%a - %{user-agent}i" "[%d/%m/%Y]t %U %s %b %T"
```

- o The string can contain the `\n` and `\t` type control characters to represent new lines and tabs.
- o Use the `<Esc>` key with literal quotes and backslashes.

The characteristics of the request are logged by placing % directives in the format string, which are replaced in the log file by the values.

If the %v (Host name) or %x (URL suffix) format specifier is present in a log file name format string, the following characters in the file name are replaced by an underscore symbol in the log configuration file name:

```
" * . / : < > ? \ |
```

Characters whose ASCII values lie in the range of 0-31 are replaced by the following:

%<ASCII value of character in hexadecimal>.

For example, the character with ASCII value 22 is replaced by %16.

Caution: If the %v format specifier is present in a log file name format string, a separate file is opened for each virtual host. To ensure continuous logging, the maximum number of files that a process can have open should be sufficiently large. See your operating system documentation for a procedure to change the number of files that can be opened.

Creating Apache Log Formats

You can derive from the custom logs most of the log formats that Apache currently supports. The custom log formats that match Apache log formats are:

NCSA/combined: `LogFormat custom %h %l %u [%t] "%r" %s %B "%{referer}i" "%{user-agent}i"`

NCSA/Common: `LogFormat custom %h %l %u [%t] "%r" %s %B`

Referer Log: `LogFormat custom "%{referer}i" -> %U`

Useragent: `LogFormat custom %{user-agent}i`

Similarly, you can derive the other server log formats from the custom formats.

Arguments for Defining a Custom Log Format

Updated: 2015-04-02

The following table describes the data that you can use as the Log Format argument string:

Table 8. Custom Log Format

Argument	Specifies
%a	Remote IPv4 address.
%A	Local IPv4 address.
%a6	Remote IPv6 address.
%A6	Local IPv6 address.
%B	Bytes sent, excluding the HTTP headers (response size).
%b	Bytes received, excluding the HTTP headers (request size).
%d	User-defined field.
%e1	Value of the first custom HTTP request header.
%e2	Value of the second custom HTTP request header.
%E1	Value of the first custom HTTP response header.
%E2	Value of the second custom HTTP response header.
Note: For instructions on how to export custom HTTP headers, see "Configuring the NetScaler for Web Server Logging."	
%g	Greenwich Mean Time offset (for example, -0800 for Pacific Standard Time).
%h	Remote host.
%H	Request protocol.
%{Foobar}i	Contents of the Foobar: header line(s) in the request sent to the server. The system supports the User-Agent, Referer and cookie headers. The + after the % in this format informs the logging client to use the + as a word separator.

%j	Bytes received, including headers (request size)
%J	Bytes sent, including headers (response size)
%l	Remote log name (from identd, if supplied).
%m	Request method.
%M	Time taken to serve the request (in microseconds)
%{Foobar}o	Contents of Foobar: header line(s) in the reply. USER-AGENT, Referer, and cookie headers (including set cookie headers) are supported.
%p	Canonical port of the server serving the request.
%q	Query string (prefixed with a question mark (?) if a query string exists).
%r	First line of the request.
%s	Requests that were redirected internally, this is the status of the original request.
%t	Time, in common log format (standard English time format).
%{format}t	Time, in the form given by format, must be in the strftime(3) format. For format descriptions, see " Appendix B: Time Format Definition ."
%T	Time taken to serve the request, in seconds.
%u	Remote user (from auth; may be bogus if return status (%s) is 401).
%U	URL path requested.
%v	Canonical name of the server serving the request.
%V	Virtual server IPv4 address in the system, if load balancing, content switching, and/or cache redirection is used.
%V6	Virtual server IPv6 address in the system, if load balancing, content switching, and/or cache redirection is used.

For example, if you define the log format as `%+{user-agent}i`, and if the user agent value is Citrix NetScaler system Web Client, then the information is logged as `NetScaler system+Web+Client`. An alternative is to use double quotation marks. For example, `â€œ%{user-agent}iâ€•` logs it as `â€œCitrix NetScaler system Web Client.â€•` Do not use the <Esc> key on strings from `%. . .r`, `%. . .i` and, `%. . .o`. This complies with the requirements of the Common Log Format. Note that clients can insert control characters into the log. Therefore, you should take care when working with raw log files.

Time Format Definition

Updated: 2015-04-28

The following table lists the characters that you can enter as the format part of the `%{format}t` string described in the Custom Log Format table of "Arguments for Defining a Custom Log Format." Values within brackets ([]) show the range of values that appear. For example, [1,31] in the %d description in the following table shows %d ranges from 1 to 31.

Table 9. Time Format Definition

Argument	Specifies
%%	The same as %.
%a	The abbreviated name of the week day for the locale.
%A	The full name of the week day for the locale.
%b	The abbreviated name of the month for the locale.
%B	The full name of the month for the locale.
%C	The century number (the year divided by 100 and truncated to an integer as a decimal number [1,99]); single digits are preceded by a 0.
%d	The day of month [1,31]; single digits are preceded by 0.
%e	The day of month [1,31]; single digits are preceded by a blank.
%h	The abbreviated name of the month for the locale.
%H	The hour (24-hour clock) [0,23]; single digits are preceded by a 0.
%I	The hour (12-hour clock) [1,12]; single digits are preceded by a 0.
%j	The number of the day in the year [1,366]; single digits are preceded by 0.
%k	The hour (24-hour clock) [0,23]; single digits are preceded by a blank.
%l	The hour (12-hour clock) [1,12]; single digits are preceded by a blank.
%m	The number of the month in the year [1,12]; single digits are preceded by a 0.
%M	The minute [00,59]; leading 0 is permitted but not required.
%n	Inserts a new line.
%p	The equivalent of either a.m. or p.m. for the locale.
%r	The appropriate time representation in 12-hour clock format with %p.
%S	The seconds [00,61]; the range of values is [00,61] rather than [00,59] to allow for the occasional leap second and for the double leap second.
%t	Inserts a tab.
%u	The day of the week as a decimal number [1,7]. 1 represents Sunday, 2 represents Tuesday and so on.
%U	The number of the week in the year as a decimal number [00,53], with Sunday as the first day of week 1.
%w	The day of the week as a decimal number [0,6]. 0 represents Sunday.
%W	Specifies the number of the week in the year as a decimal number [00,53]. Monday is the first day of week 1.
%y	The number of the year within the century [00,99]. For example, 5 would be the fifth year of that century.
%Y	The year, including the century (for example, 1993).

Note: If you specify a conversion that does not correspond to any of the ones described in the preceding table, or to any of the modified conversion specifications listed in the next paragraph, the behavior is undefined and returns 0.

The difference between %U and %W (and also between modified conversions %OU and %OW) is the day considered to be the first day of the week. Week number 1 is the first week in January (starting with a Sunday for %U, or a Monday for %W). Week number 0 contains the days before the first Sunday or Monday in January for %U and %W.

Configuring Call Home

The Call Home feature monitors your NetScaler appliance for critical error conditions. Call Home registers your appliance with the Citrix Technical Support server. If your appliance is successfully registered with the Support server, Call Home automatically uploads system data to that server in the event that one of the conditions occurs. The NetScaler Appliance keeps a full log of all upload events. If you are unable to correct the problem after reviewing the appliance's log, you can contact the Citrix Technical Support team and open a service request. The team can analyze the uploaded system data and recommend possible solutions.

The Call Home feature is supported on all three platforms of NetScaler ADC.

- In NetScaler MPX, Call Home feature is supported on all MPX models.
- In NetScaler SDX, Call Home feature is supported on all VPX instances running on a SDX box 002E

Following is a typical set up for Call Home.

Step 1: Appliance Registration



Step 2: Trigger Based Upload



The process flow for using Call Home can be categorized as follows:

- Registration of the NetScaler appliance to the Citrix Technical Support server.
- Uploading of the appliance's data to the Citrix Technical Support server. The support server has the following URL: <https://callhome.citrix.com/>.
- Opening a Technical Support case (Optional).

Registration of the NetScaler appliance. The appliance has to be registered to the Citrix Technical Support server before Call Home can upload the system data to the server when predefined error conditions occur on the appliance. Enabling the Call Home feature on the NetScaler appliance initiates the registration process. The process flow is as follows:

1. The Call Home process sends the following details to the Citrix Technical Support server:
 - Hardware serial number is shared for NetScaler MPX and SDX models
 - License serial number is shared for NetScaler VPX models.
2. The server checks its database for an active technical support service contract for the appliance.
3. If there is an active technical support service contract, the support server registers the NetScaler appliance for Call Home and sends a successful-registration response to the appliance stating that the feature is successfully enabled. If there is no active technical support service contract, the server sends a registration-failure response to the NetScaler appliance.

The following table lists the error conditions that Call Home currently monitors on a NetScaler Appliance:

Table 1. List of error conditions monitored by Call Home

Error Condition	Indicates	Call Home Monitoring Interval	Corresponding SNMP Alarm Name
Compact flash drive errors	The compact flash drive on the appliance that encountered read or write errors.	24 hours	COMPACT-FLASH-ERRORS
Hard disk drive errors	The hard drives on the appliance that encountered read or write errors.	24 hours	HARD-DISK-DRIVE-ERRORS
Power supply unit failure	One of the power supply units on the NetScaler appliance has failed.	7 seconds	POWER-SUPPLY-FAILURE
SSL card failure	One of the SSL cards on the NetScaler appliance has failed.	7 seconds	SSL-CARD-FAILED

Warm restart	The appliance has warm restarted due to a failure of a system process.	After every restart of the NetScaler appliance.	WARM-RESTART-EVENT
--------------	--	---	--------------------

Note: The Call Home feature do not monitor the power supply unit (PSU) status for VPX models and VPX instances.

Uploading of appliance's data to the Support server. An error condition triggers the following sequence of events:

1. The Call Home process checks the registration status. If the status indicates successful registration, the process advances to the next step.
2. The Call Home process runs a showtech support script that collects all of the system related data in a tar file. The data in the tar file includes configurations, logs, and statistics. Call Home locally saves the tar file at `/var/tmp/support/callhome`.
3. Call Home uploads a copy of the tar file to the Citrix Technical Support server. The Appliance logs the uploading of the tar file in a log file named `callhome.log` located at `/var/log`. You can also configure the CALLHOME-UPLOAD-EVENT SNMP alarm to generate an SNMP alert whenever Call Home uploads happen.
4. If the SNMP alarm related to the error condition is enabled, the SNMP agent on the appliance generates an SNMP trap message and sends it to all of the configured SNMP trap destinations. For more information about configuring SNMP alarms and trap destinations, see "[Configuring the NetScaler to Generate SNMPv1 and SNMPv2 Traps](#)."

Note:

- Call Home creates the Call Home tar file and uploads it to the CITRIX tech support server for only the first occurrence of a particular error condition since the appliance was last restarted. If you want the NetScaler appliance to send you alerts each time a particular error condition occurs, configure the corresponding SNMP alarm for the error condition.
- For the warm restart error condition, the Call Home tar file is uploaded to the server again only if the result of the failure is different from the result of the previous failure.

The Call Home tar file has the following name format:

```
collector_callhome_<NSIP of the appliance>_<P for Primary or standalone, or S for
Secondary>_<date>_<hours, in 24 hr format, according to the local time zone>_<minutes>.
tar.gz. For example, collector_callhome_10.105.13.100_P_2Feb2012_20_30.tar.gz.
```

Opening a Technical Support Service Request . After you review the logs and SNMP trap messages for Call Home upload events, you have the option of contacting the Citrix Technical Support team and opening a service request. For more information about contacting the team and opening a service request, see <http://support.citrix.com/article/CTX132307>.

The Support team can then analyze the system data in the uploaded Call Home tar files and sends recommendations for possible solutions to the administrator's email address.

Before you begin configuring Call Home, do the following:

- Make sure that the NetScaler appliance is connected to the Internet or to a proxy server that has internet connectivity.
- Make sure that you have an active Citrix Technical Support service contract for the appliance.

Configuring Call Home on the NetScaler appliance consists of the following tasks:

1. **Enable the Call Home feature.** When you enable the Call Home feature, the Call Home process registers the appliance with the Citrix Technical Support server. The registration takes some time to complete. During that time, the appliance displays the status as IN PROGRESS. When the registration is complete, the appliance displays the status as SUCCESSFUL.
Note: While upgrading the NetScaler appliance from an older release to release 10.1 or later, the NetScaler appliance prompts you to enable the Call Home feature, if:
 - The Call Home feature is not supported in the older release.
 - The Call Home feature is disabled in the older release.
2. **(Optional) Specify the administrator's email address.** The Call Home process sends the email address to the Support server, where it is stored for future correspondence regarding Call Home uploads.
3. **(Optional) Specify proxy server settings.** NetScaler appliance needs internet connectivity to upload the collector archive to the Citrix Technical Support server. If the appliance does not have internet connectivity, then a proxy server (having internet connectivity) can be configured to upload the data.
4. **(Optional) Enable the CALLHOME-UPLOAD-EVENT SNMP alarm.** The SNMP agent on the NetScaler appliance generates a trap message and sends to all the configured SNMP trap destinations. The message includes the status of uploading of the Call Home tar file by the Call Home process. For more information about configuring SNMP alarms and trap destinations, see "[Configuring the NetScaler to Generate SNMPv1 and SNMPv2 Traps](#)."
5. **(Optional) Enable all of the corresponding SNMP alarms.** Call Home creates and uploads a Call Home tar file for the first occurrence of a monitored error condition since the appliance was last restarted. If you want to be alerted of

these error conditions, you can configure the corresponding SNMP alarm. Table 1 lists all the corresponding SNMP alarms. For more information about configuring SNMP alarms and trap destinations, see "Configuring the NetScaler to Generate SNMPv1 and SNMPv2 Traps."

To enable Call Home by using the command line interface

At the command prompt, type any of the following:

- o enable ns feature ch
- o enable ns feature callhome

To check the status of the appliance's registration to the Support server by using the command line interface

At the command prompt, type:
show callhome

Example

```
> enable ns feature ch
Done
```

```
> show callhome
Callhome feature: ENABLED
Registration with Citrix upload server IN PROGRESS
```

```
E-mail address configured:
Proxy mode:NO    Ipaddress:      Port:0
```

Trigger event	State	First occurrence	Latest occurrence
1) Compact flash errors	Enabled
2) Hard disk drive errors	Enabled
3) Power supply unit failure	Enabled
4) SSL card failure	Enabled
5) Warm restart	Enabled	N/A	..

```
Done
```

```
> show callhome
Callhome feature: ENABLED
Registration with Citrix upload server SUCCESSFUL
```

```
E-mail address configured:
Proxy mode:NO    Ipaddress:      Port:0
```

Trigger event	State	First occurrence	Latest occurrence
1) Compact flash errors	Enabled
2) Hard disk drive errors	Enabled
3) Power supply unit failure	Enabled
4) SSL card failure	Enabled
5) Warm restart	Enabled	N/A	..

```
Done
```

To specify the administrator's email address and proxy server settings by using the command line interface

At the command prompt, type:

- o set callhome -emailAddress <string>
- o set callhome -proxyMode (YES | NO) [-IPAddress <ip_addr|ipv6_addr|*>] [-port <port|*>]
- o show callhome

Example

```
> set callhome -emailAddress exampleadmin@example.com
Done
```

```
> set callhome "proxyMode Yes" "IPAddress 10.102.167.33" "port 80"
```

Done

```
> show callhome
E-mail address configured: exampleadmin@example.com
Proxy mode: YES    Ipaddress: 10.102.167.33    Port: 80
```

Trigger event	State	First occurrence	Latest occurrence
1) Compact flash errors	Enabled
2) Hard disk drive errors	Enabled
3) Power supply unit failure	Enabled
4) SSL card failure	Enabled
5) Warm restart	Enabled	N/A	..

Done

To enable Call Home proxy mode by using the command line interface

At the command prompt, type:

- o set callhome -ipAddress <ipaddress> -port <port> -proxyMode [yes | no]
- o show callhome

Note: Proxy mode is enabled only when the -proxymode parameter is set to YES. If it is set to NO, the proxy functionality does not work, even if the IP address and port are configured. The port number should be for an HTTP service on the proxy server, not for an HTTPS service.

Example

```
> set callhome ipAddress 10.0.0.1 -port 80 -proxyMode yes
Done
```

To enable Call Home by using the configuration utility

Navigate to System > Settings, click Configure Advanced Features and select the Call Home option.

To check the status of the appliance's registration with the Support server by using the configuration utility

Navigation to System > Diagnostics and, in the Technical Support Tools group, select Call Home option to view the status of registration.

To specify the administrator's email address by using the configuration utility

Navigation to System > Diagnostics and, in the Technical Support Tools group, select Call Home option to specify the administrator's email address.

To enable Call Home proxy mode by using the configuration utility

Navigation to System > Diagnostics and, in the Technical Support Tools group, select Call Home option to specify the proxy server's IP address and the port number.

Reporting Tool

Use the Citrix® NetScaler® Reporting tool to view NetScaler performance statistics data as reports. Statistics data are collected by the nscollect utility and are stored in a database. When you want to view certain performance data over a period of time, the Reporting tool pulls out specified data from the database and displays them in charts.

Reports are a collection of charts. The Reporting tool provides built-in reports as well as the option to create custom reports. In a report, you can modify the charts and add new charts. You can also modify the operation of the data collection utility, nscollect, and stop or start its operation.

This document includes the following details:

- Using the Reporting Tool
- Working with Reports
- Working with Charts
- Examples
- Stopping and Starting the Data Collection Utility

Using the Reporting Tool

The Reporting tool is a Web-based interface accessed from the Citrix® NetScaler® appliance. Use the Reporting tool to display the performance statistics data as reports containing graphs. In addition to using the built-in reports, you can create custom reports, which you can modify at any time. Reports can have between one and four charts. You can create up to 256 custom reports.

To invoke the Reporting tool

1. Use the Web browser of your choice to connect to the IP address of the NetScaler (for example, <http://10.102.29.170/>). The Web Logon screen appears.
2. In the User Name text box, type the user name assigned to the NetScaler.
3. In the Password text box, type the password.
4. In the Start in drop-down box, select Reporting.
5. Click Login.

The following screen shots show the report toolbar and the chart toolbar, which are frequently referenced in this documentation.

Figure 1. *Report Toolbar*



Figure 2. *Chart Toolbar*



Working with Reports

Updated: 2013-09-27

You can plot and monitor statistics for the various functional groups configured on the NetScaler over a specified time interval. Reports enable you to troubleshoot or analyze the behavior of your appliance. There are two types of reports: built-in reports and custom reports. Report content for built-in or custom reports can be viewed in a graphical format or a tabular format. The graphical view consists of line, area, and bar charts that can display up to 32 sets of data (counters). The tabular view displays the data in columns and rows. This view is useful for debugging error counters.

The default report that is displayed in the Reporting tool is CPU vs. Memory Usage and HTTP Requests Rate. You can change the default report view by displaying the report you want as your default view, and then clicking Default Report.

Reports can be generated for the last hour, last day, last week, last month, last year, or you can customize the duration. You can do the following with reports:

- Toggle between a tabular view of data and a graphical view of data.
- Change the graphical display type, such as bar chart or line chart.
- Customize charts in a report.
- Export the chart as an Excel comma-separated value (CSV) file.
- View the charts in detail by zooming in, zooming out, or using a drag-and-drop operation (scrolling).
- Set a report as the default report for viewing whenever you log on.
- Add or remove counters.

- Print reports.
- Refresh reports to view the latest performance data.

Using Built-in Reports

The Reporting tool provides built-in reports for frequently viewed data. Built-in reports are available for the following functional groups: System, Network, SSL, Compression, Integrated Cache, NetScaler Gateway, and Citrix NetScaler Application Firewall. By default, the built-in reports are displayed for the last day. However, you can view the reports for the last hour, last week, last month, or last year.

Note: You cannot save changes to built-in reports, but you can save a modified built-in report as a custom report.

To display a built-in report

1. In the left pane of the Reporting tool, under Built-in Reports, expand a group (for example, SSL).
2. Click a report (for example, SSL > All Backend Ciphers).

Creating and Deleting Reports

You can create your own custom reports and save them with user-defined names for reuse. You can plot different counters for different groups based on your requirements. You can create up to 256 custom reports.

You can either create a new report or save a built-in report as a custom report. By default, a newly created custom report contains one chart named System Overview, which displays the CPU Usage counter plotted for the last day. You can customize the interval and set the data source and time zone from the report toolbar. Within a report, you can use the chart toolbars to add, modify, or delete charts, as described in "[Working with Charts](#)."

By default, newly created custom reports contain one chart named System Overview that displays a CPU Usage counter plotted for the last day.

To create a custom report

1. In the Reporting tool, on the report toolbar, click Create, or if you want to create a new custom report based on an existing report, open the existing report, and then click Save As.
2. In Report Name box, type a name for the custom report.
3. Do one of the following:
 - To add the report to an existing folder, in Create in or Save in, click the down arrow to choose an existing folder, and then click OK.
 - To create a new folder to store the report, click the Click to add folder icon, in Folder Name, type the name of the folder, and in Create in, specify where you want the new folder to reside in the hierarchy, and then click OK.

Note: You can create up to 128 folders.

To delete a custom report




1. In the left pane of the Reporting tool, next to Custom Reports, click the Click to manage custom reports icon.
2. Select the check box that corresponds with the report you want to delete, and then click Delete.

Note: When you delete a folder, all the contents of that folder are deleted.

Modifying the Time Interval

By default, built-in reports display data for the last day. However, if you want to change the time interval for a built-in report, you can save the report as a custom report. The new interval applies to all charts in the report. The following table describes the time-interval options.

Table 1. Time Intervals

Time interval	Displays
 Last Hour	Statistics data collected for the last hour.
 Last Day	Statistics data collected for the last day (24 hours).
 Last Week	Statistics data collected for the last week (7 days).

 Last Month	Statistics data collected for the last month (31 days).
 Last Year	Statistics data collected for the last year (365 days).
 Custom	Statistics data collected for a time period that you are prompted to specify.

To modify the time interval

1. In the left pane of the Reporting tool, click a report.
2. On the report toolbar, click Duration, and then click a time interval.

Setting the Data Source and Time Zone

You can retrieve data from different data sources to display them in the reports. You can also define the time zone for the reports and apply the currently displayed report's time selection to all the reports, including the built-in reports.

To set the data source and time zone

1. In the Reporting tool, on the report toolbar, click Settings.
2. In the Settings dialog box, in Data Source, select the data source from which you want to retrieve the counter information.
3. Do one or both of the following:
 - If you want the tool to remember the time period for which a chart is plotted, select the Remember time selection for charts check box.
 - If you want the reports to use the time settings of your NetScaler appliance, select the Use Appliance's time zone check box.

Exporting and Importing Custom Reports

You can share reports with other NetScaler administrators by exporting reports. You can also import reports.

To export or import custom reports

1. In the left pane of the Reporting tool, next to Custom Reports, click the Click to manage custom reports icon.
2. Select the check box that corresponds with the report you want to export or import, and then click Export or Import.
Note: When you export the file, it is exported in a .gz file format.

Working with Charts

Updated: 2013-09-06

Use charts to plot and monitor counters or groups of counters. You can include up to four charts in one report. In each chart, you can plot up to 32 counters. The charts can use different graphical formats (for example, area and bar). You can move the charts up or down within the report, customize the colors and visual display for each counter in a chart, and delete a chart when you do not want to monitor it.

In all report charts, the horizontal axis represents time and the vertical axis represents the value of the counter.

Adding a Chart

When you add a chart to a report, the System Overview chart appears with the CPU Usage counter plotted for the last one day. To plot a different group of statistics or select a different counter, see "[Modifying a Chart](#)."

Note: If you add charts to a built-in report, and you want to retain the report, you must save the report as a custom report.

Use the following procedure to add a chart to a report.

To add a chart to a report

1. In the left pane of the Reporting tool, click a report.
2. Under the chart where you want to add the new chart, click the Add icon.

Modifying a Chart

You can modify a chart by changing the functional group for which the statistics are displayed and by selecting different counters.

To modify a chart

1. In the left pane of the Reporting tool, click a report.
2. Under the chart that you want to modify, click Counters.
3. In the dialog box that appears, in the Title box, type a name for the chart.
4. Next to Plot chart for, do one of the following:
 - To plot counters for global counters, such as Integrated Cache and Compression, click System global statistics.
 - To plot entity counters for entity types, such as Load Balancing and GSLB, click System entities statistics.
5. In Select group, click the desired entity.
6. Under Counters, in Available, click the counter name(s) that you want to plot, and then click the > button.
7. If you selected System entities statistics in step 4, on the Entities tab, under Available, click the entity instance name (s) you want to plot, and then click the > button.
8. Click OK.

Viewing a Chart

You can specify the graphical formats of the plotted counters in a chart. Charts can be viewed as line charts, spline charts, step-line charts, scatter charts, area charts, bar charts, stacked area charts, and stacked bar charts. You can also zoom in, zoom out, or scroll inside the plot area of a chart. You can zoom in or out for all data sources for 1 hour, 1 day, 1 week, 1 month, 1 year, and 3 years.

Other options for customizing the view of a chart include customizing the axes of the charts, changing the background and edge color of the plot area, customizing the color and size of the grids, and customizing the display of each data set (counter) in a chart.

Data set numbers, such as Data Set 1, correspond to the order in which the counters in your graph are displayed at the bottom of the chart. For example, if CPU usage and Memory usage are displayed in first and second order at the bottom of the chart, CPU usage is equal to Data Set 1 and Memory usage is equal to Data Set 2.

Whenever you modify a built-in report, you need to save the report as a custom report to retain your changes.

To change the graph type of a chart

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, under the chart you want to view, on the chart toolbar, click Customize.
3. On the Chart tab, under Category, click Plot type, and then click the graph type you want to display for the chart. If you want to display the graph as 3D, select the Use 3D check box.

To refocus a chart with detailed data

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, on the report toolbar, click Zoom In, and do one or both of the following:
 - To refocus the chart to display data for a specific time window, drag and drop the cursor from the start time to the end time. For example, you can view data for a one-hour period on a certain day.
 - To refocus the chart to display data for a data point, simply click once on chart where you want to zoom in and get more detailed information.
3. Once you have the desired range of time for which you want to view detailed data, on the report toolbar, click Tabular View. Tabular view displays the data in numeric form in rows and columns.

To view numeric data for a graph

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, on the report toolbar, click Tabular View. To return to the graphical view, click Graphical View.
Note: You can also view the numeric data in the graphical view by hovering your cursor over the notches in the gridlines.

To scroll through time in a chart

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, on the report toolbar, click Scroll, and then click inside the chart and drag the cursor in the direction for which you want to see data for a new time period. For example, if you want to view data in the past, click and drag to the left.

To change the background color and text color of a chart

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, under the chart for which you want to customize the axes, click Customize.
3. On the Chart tab, under Category, click one or more of the following:
 - o To change the background color, click Background Color, and then select the options for color, transparency, and effects.
 - o To change the text color, click Text Color, and then select the options for color, transparency, and effects.

To customize the axes of a chart

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, under the chart for which you want to customize the axes, click Customize.
3. On the Chart tab, under Category, click one or more of the following:
 - o To change the scale of the left y-axis, click Left Y-Axis, and then select the scale you want.
 - o To change the scale of the right y-axis, click Right Y-Axis, in Data set to plot, select the data set, and then select the scale you want.

Note: The data set numbers, such as Data Set 1, correspond to the order in which the counters in your graph are displayed at the bottom of the chart. For example, if CPU usage and Memory usage are displayed in first and second order at the bottom of the chart, CPU usage is equal to Data Set 1 and Memory usage is equal to Data Set 2.

 - o To plot each data set in its own hidden y-axis, click Multiple Axes, and then click Enable.

To change the background color, edge color, and gridlines for a plot area of a chart

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, under the chart for which you want to customize the plot area, click Customize.
3. On the Plot Area tab, under Category, click one or more of the following:
 - o To change the background color and edge color of the chart, click Background Color and Edge Color, and then select the options for color, transparency, and effects.
 - o To change the horizontal or vertical grids of the chart, click Horizontal Grids or Vertical Grids, and then select the options for displaying the grids, grid width, grid color, transparency, and effects.

To change the color and graph type of a data set

1. In the left pane of the Reporting tool, select a report.
 2. In the right pane, under the chart for which you want to customize the display of the data set (counters), click Customize.
 3. On the Data Set tab, in Select Data Set, select the data set (counter) for which you want to customize the graphical display.
- Note: The data set numbers, such as Data Set 1, correspond to the order in which the counters in your graph are displayed at the bottom of the chart. For example, if CPU usage and Memory usage are displayed in first and second order at the bottom of the chart, CPU usage is equal to Data Set 1 and Memory usage is equal to Data Set 2.
4. Under Category, do one of more of the following:
 - o To change the background color, click Color, and then select the options for color, transparency, and effects.
 - o To change the graph type, click Plot type, and then select the graph type you want to display for the data set. If you want to display the graph as 3D, select the Use 3D check box.

Exporting Chart Data to Excel

For further data analysis, you can export charts to Excel in a comma-separated value (CSV) format.

To export chart data to Excel

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, under the chart with the data you want to export to Excel, click Export.

Deleting a Chart

If you do not want to use a chart, you can remove it from the report. You can permanently remove charts from custom reports only. If you delete a chart from a built-in report and want to retain the changes, you need to save the report as a custom report.

To delete a chart

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, under the chart that you want to delete, click the Delete icon.

Examples

To display the trend report for CPU usage and memory usage for the last week

1. In the left pane of the Reporting tool, under Built-in Reports, expand System.
2. Click the report CPU vs. Memory Usage and HTTP Requests Rate.
3. In the right pane, on the report toolbar, click Duration, and then click Last Week.

To compare the bytes received rate and the bytes transmitted rate between two interfaces for the last week

1. In the right pane, on the report toolbar, click Create.
2. In the Report Name box, type a name for the custom report (for example, `Custom_Interfaces`), and then click OK.
The report is created with the default System Overview chart, which displays the CPU Usage counter plotted for the last hour.
3. Under System Overview, on the chart toolbar, click Counters.
4. In the counter selection pane, in Title, type a name for the chart (for example, `Interfaces bytes data`).
5. In Plot chart for, click System entities statistics, and then in Select Group, select Interface.
6. On the Entities tab, click the interface name(s) you want to plot (for example, `1/1` and `1/2`), and then click the > button.
7. On the Counters tab, click Bytes received (Rate) and Bytes transmitted (Rate) and then click the > button.
8. Click OK.
9. On the report toolbar, click Duration, and then click Last Week.

Stopping and Starting the Data Collection Utility

Updated: 2014-09-24

The data collection utility, `nscollect`, runs automatically when you start the NetScaler ADC. This utility retrieves the application performance data and stores it in the form of data sources on the ADC. You can create up to 32 data sources. The default data source is `/var/log/db/default`.

The data collection utility creates databases for global counters and entity-specific counters, and uses this data to generate reports. Global-counter databases are created at `/var/log/db/<DataSourceName>`. The entity-specific databases are created based on the entities configured on the NetScaler, and a separate folder is created for each entity type in `/var/log/db/<DataSourceName/EntityNameDB>`.

`Nscollect` retrieves data once every 5 minutes. It retains data in 5-minute granularity for one day, hourly for the last 30 days, and daily for three years.

You might have to stop and restart the data collection utility if data is not updated accurately or the reports display corrupted data.

To stop `nscollect`

At the command prompt, type:
`/netscaler/nscollect stop`

To start `nscollect` on the local system

At the command prompt, type:
`/netscaler/nscollect start`

AppFlow

The Citrix NetScaler appliance is a central point of control for all application traffic in the data center. It collects flow and user-session level information valuable for application performance monitoring, analytics, and business intelligence applications. It also collects web page performance data and database information. AppFlow transmits the information by using the Internet Protocol Flow Information eXport (IPFIX) format, which is an open Internet Engineering Task Force (IETF) standard defined in RFC 5101. IPFIX (the standardized version of Cisco's NetFlow) is widely used to monitor network flow information. AppFlow defines new Information Elements to represent application-level information, web page performance data, and database information.

Using UDP as the transport protocol, AppFlow transmits the collected data, called *flow records*, to one or more IPv4 collectors. The collectors aggregate the flow records and generate real-time or historical reports.

AppFlow provides visibility at the transaction level for HTTP, SSL, TCP, and SSL_TCP flows. You can sample and filter the flow types that you want to monitor.

AppFlow use actions and policies to send records for a selected flow to specific set of collectors. An AppFlow action specifies which set of collectors will receive the AppFlow records. Policies, which are based on Advanced expressions can be configured to select flows for which flow records will be sent to the collectors specified by the associated AppFlow action.

To limit the types of flows, you can enable AppFlow for a virtual server. AppFlow can also provide statistics for the virtual server.

You can also enable AppFlow for a specific service, representing an application server, and monitor the traffic to that application server.

Note: This feature is supported only on NetScaler nCore builds.

This topic includes the following details:

- [How AppFlow Works](#)
- [Configuring the AppFlow Feature](#)
- [Exporting Performance Data of Web Pages to AppFlow Collector](#)

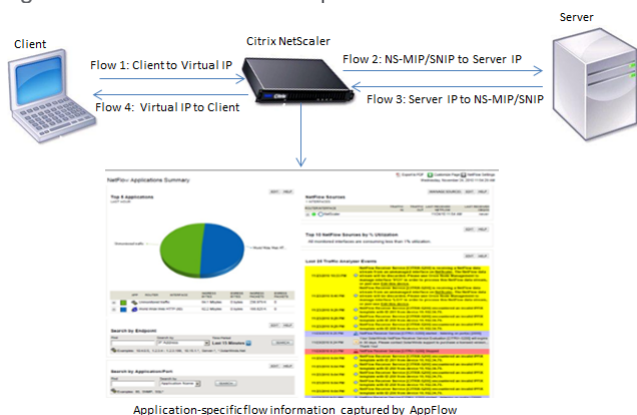
How AppFlow Works

Updated: 2015-05-28

In the most common deployment scenario, inbound traffic flows to a Virtual IP address (VIP) on the NetScaler appliance and is load balanced to a server. Outbound traffic flows from the server to a mapped or subnet IP address on the NetScaler and from the VIP to the client. A flow is a unidirectional collection of IP packets identified by the following five tuples: sourceIP, sourcePort, destIP, destPort, and protocol.

The following figure describes how the AppFlow feature works.

Figure 1. NetScaler Flow Sequence



As shown in the figure, the network flow identifiers for each leg of a transaction depend on the direction of the traffic.

The different flows that form a flow record are:

Flow1: <Client-IP, Client-Port, VIP-IP, VIP-port, Protocol>

Flow2: <NS-MIP/SNIP, NS-port, Server-IP, Server-Port, Protocol>

Flow3: <Server-IP, Server-Port, NS-MIP/SNIP, NS-Port, Protocol>

Flow4: <VIP-IP, VIP-port, Client-IP, Client-Port, Protocol>

To help the collector link all four flows in a transaction, AppFlow adds a custom transactionID element to each flow. For application-level content switching, such as HTTP, it is possible for a single client TCP connection to be load balanced to different backend TCP connections for each request. AppFlow provides a set of records for each transaction.

This topic includes the following details:

- [Flow Records](#)
- [Templates](#)

Flow Records

Updated: 2013-08-20

AppFlow records contain standard NetFlow or IPFIX information, such as time stamps for the beginning and end of a flow, packet count, and byte count. AppFlow records also contain application-level information (such as HTTP URLs, HTTP request methods and response status codes, server response time, and latency), web page performance data (such as page load time, page render time, and time spent on the page), and database information (such as database protocol, database response status and database response size). IPFIX flow records are based on templates that need to be sent before sending flow records.

Templates

AppFlow defines a set of templates, one for each type of flow. Each template contains a set of standard Information Elements (IEs) and Enterprise-specific Information Elements (EIEs). IPFIX templates define the order and sizes of the Information Elements (IE) in the flow record. The templates are sent to the collectors at regular intervals, as described in RFC 5101.

A template can include the following EIEs:

transactionID

An unsigned 32-bit number identifying an application-level transaction. For HTTP, this corresponds to a request and response pair. All flow records that correspond to this request and response pair have the same transaction ID. In the most common case, there are four unidirectional flow records that correspond to this transaction. If the NetScaler generates the response by itself (served from the integrated cache or by a security policy), there may be only two flow records for this transaction.

connectionID

An unsigned 32-bit number identifying a layer-4 connection (TCP or UDP). The NetScaler flows are usually bidirectional, with two separate flow records for each direction of the flow. This information element can be used to link the two flows.

For the NetScaler, connectionID is an identifier for the connection data structure to track the progress of a connection. In an HTTP transaction, for instance, a given connectionID may have multiple transactionID elements corresponding to multiple requests that were made on that connection.

tcpRTT

The round trip time, in milliseconds, as measured on the TCP connection. This can be used as a metric to determine the client or server latency on the network.

httpRequestMethod

An 8-bit number indicating the HTTP method used in the transaction. An options template with the number-to-method mapping is sent along with the template.

httpRequestSize

An unsigned 32-bit number indicating the request payload size.

httpRequestURL

The HTTP URL requested by the client.

httpUserAgent

The source of incoming requests to the Web server.

httpResponseStatus

An unsigned 32-bit number indicating the response status code.

httpResponseSize

An unsigned 32-bit number indicating the response size.

httpResponseTimeToFirstByte

An unsigned 32-bit number indicating the time taken to receive the first byte of the response.

httpResponseTimeToLastByte

An unsigned 32-bit number indicating the time taken to receive the last byte of the response.

flowFlags

An unsigned 64-bit flag used to indicate different flow conditions.

EIEs for web page performance data

clientInteractionStartTime

Time at which the browser receives the first byte of the response to load any objects of the page such as images, scripts, and stylesheets.

clientInteractionEndTime

Time at which the browser received the last byte of response to load all the objects of the page such as images, scripts, and stylesheets.

clientRenderStartTime

Time at which the browser starts to render the page.

clientRenderEndTime

Time at which browser finished rendering the entire page, including the embedded objects.

EIEs for database information

dbProtocolName

An unsigned 8-bit number indicating the database protocol. Valid values are 1 for MS SQL and 2 for MySQL.

dbReqType

An unsigned 8-bit number indicating the database request method used in the transaction. For MS SQL, valid values are 1 is for QUERY, 2 is for TRANSACTION, and 3 is for RPC. For valid values for MySQL, see the MySQL documentation.

dbReqString

Indicates the database request string without the header.

dbRespStatus

An unsigned 64-bit number indicating the status of the database response received from the web server.

dbRespLength

An unsigned 64-bit number indicating the response size.

dbRespStatString

The response status string received from the web server.

Configuring the AppFlow Feature

You configure AppFlow in the same manner as most other policy-based features. First, you enable the AppFlow feature. Then you specify the collectors to which the flow records are sent. After that, you define actions, which are sets of configured collectors. Then you configure one or more policies and associate an action to each policy. The policy tells the NetScaler appliance to select requests the flow records of which are sent to the associated action. Finally, you bind each policy either globally or to specific vservers to put it into effect.

You can further set AppFlow parameters to specify the template refresh interval and to enable the exporting of httpURL, httpCookie, and httpReferer information. On each collector, you must specify the NetScaler IP address as the address of the exporter.

Note: For information about configuring the NetScaler as an exporter on the collector, see the documentation for the specific collector.

The configuration utility provides tools that help users define the policies and actions that determine exactly how the NetScaler appliance export records for a particular flow to a set of collectors(action.) The command line interface provides a corresponding set of CLI-based commands for experienced users who prefer a command line.

This topic includes the following details:

- [Enabling AppFlow](#)
- [Specifying a Collector](#)
- [Configuring an AppFlow Action](#)
- [Configuring an AppFlow Policy](#)
- [Binding an AppFlow Policy](#)
- [Enabling AppFlow for Virtual Servers](#)
- [Enabling AppFlow for a Service](#)
- [Setting the AppFlow Parameters](#)
- [Example: Configuring AppFlow for DataStream](#)

Enabling AppFlow

Updated: 2014-08-07

To be able to use the AppFlow feature, you must first enable it.

Note: AppFlow can be enabled only on nCore NetScaler appliances.

To enable the AppFlow feature by using the command line interface

At the command prompt, type one of the following commands:

```
enable ns feature AppFlow
```

To enable the AppFlow feature by using the configuration utility

Navigate to System > Settings, click Configure Advanced Features and select the AppFlow option.

Specifying a Collector

Updated: 2014-08-07

A collector receives flow records generated by the NetScaler appliance. To be able to send flow records, you must specify at least one collector. You can specify up to four. However, you cannot export the same data to multiple collectors. You can remove unused collectors. By default, the collector listens to IPFIX messages on UDP port 4739. You can change the default port, when configuring the collector. Similarly, by default, NSIP is used as the source IP for appflow traffic. You can change this default source IP to a SNIP or MIP address when configuring a collector.

To specify a collector by using the command line interface

At the command prompt, type the following commands to add a collector and verify the configuration:

- `add appflow collector <name> -IPAddress <ipaddress> -port <port_number> -netprofile <netprofile_name>`
- `show appflow collector <name>`

Example

```
> add appflow collector coll -IPAddress 10.102.29.251 -port 8000 -netprofile n2
```

To specify a collector by using the configuration utility

Navigate to System > AppFlow > Collectors, and create the AppFlow collector.

Configuring an AppFlow Action

Updated: 2014-08-07

An AppFlow action is a set collectors, to which the flow records are sent if the associated AppFlow policy matches.

To configure an AppFlow action by using the command line interface

At the command prompt, type the following commands to configure an Appflow action and verify the configuration:

- o add appflow action <name> --collectors <string> ... [-clientSideMeasurements (Enabled|Disabled)] [-comment <string>]
- o show appflow action

Example

```
> add appflow action apfl-act-collector-1-and-3 -collectors collector-1 collector-3
```

To configure an AppFlow action by using the configuration utility

Navigate to System > AppFlow > Actions, and create the AppFlow action.

Configuring an AppFlow Policy

Updated: 2014-08-07

After you configure an AppFlow action, you must next configure an AppFlow policy. An AppFlow policy is based on a rule, which consists of one or more expressions.

Note: For creating and managing AppFlow policies, the configuration utility provides assistance that is not available at the command line interface.

To configure an AppFlow policy by using the command line interface

At the command prompt, type the following command to add an AppFlow policy and verify the configuration:

- o add appflow policy <name> <rule> <action>
- o show appflow policy <name>

Example

```
> add appflow policy apfl-pol-tcp-dsprt client.TCP.DSTPORT.EQ(22) apfl-act-collector-1-and-3
```

To configure an AppFlow policy by using the configuration utility

Navigate to System > AppFlow > Policies, and create the AppFlow policy.

To add an expression by using the Add Expression dialog box

1. In the Add Expression dialog box, in the first list box choose the first term for your expression.

HTTP

The HTTP protocol. Choose this if you want to examine some aspect of the request that pertains to the HTTP protocol.

SYS

The protected Web site(s). Choose this if you want to examine some aspect of the request that pertains to the recipient of the request.

CLIENT

The computer that sent the request. Choose this if you want to examine some aspect of the sender of the request.

When you make your choice, the rightmost list box lists appropriate terms for the next part of your expression.

2. In the second list box, choose the second term for your expression. The choices depend upon which choice you made in the previous step, and are appropriate to the context. After you make your second choice, the Help window below the Construct Expression window (which was blank) displays help describing the purpose and use of the term you just chose.

3. Continue choosing terms from the list boxes that appear to the right of the previous list box, or typing strings or numbers in the text boxes that appear to prompt you to enter a value, until your expression is finished.

Binding an AppFlow Policy

Updated: 2014-08-07

To put a policy into effect, you must bind it either globally, so that it applies to all traffic that flows through the NetScaler, or to a specific virtual server, so that the policy applies only to the traffic related to that virtual server.

When you bind a policy, you assign it a priority. The priority determines the order in which the policies you define are evaluated. You can set the priority to any positive integer.

In the NetScaler operating system, policy priorities work in reverse order—the higher the number, the lower the priority. For example, if you have three policies with priorities of 10, 100, and 1000, the policy assigned a priority of 10 is performed first, then the policy assigned a priority of 100, and finally the policy assigned an order of 1000.

You can leave yourself plenty of room to add other policies in any order, and still set them to evaluate in the order you want, by setting priorities with intervals of 50 or 100 between each policy when you globally bind it. You can then add additional policies at any time without having to change the priority of an existing policy.

To globally bind an AppFlow policy by using the command line interface

At the command prompt, type the following command to globally bind an AppFlow policy and verify the configuration:

- `bind appflow global <policyName> <priority> [<gotoPriorityExpression> [-type <type>] [-invoke (<labelType> <labelName>)]`
- `show appflow global`

Example

```
bind appflow global af_policy_lb1_10.102.71.190 1 NEXT -type REQ_OVERRIDE -invoke vserver go
```

To bind an AppFlow policy to a specific virtual server by using the command line interface

At the command prompt, type the following command to bind an appflow policy to a specific virtual server and verify the configuration:

```
bind lb vserver <name> -policyname <policy_name> -priority <priority>
```

Example

```
bind lb vserver google -policyname af_policy_google_10.102.19.179 -priority 251
```

To globally bind an AppFlow policy by using the configuration utility

Navigate to System > AppFlow, click AppFlow policy Manager and select the relevant Bind Point (Default Global) and Connection Type, and then bind the AppFlow policy.

To bind an AppFlow policy to a specific virtual server by using the configuration utility

Navigate to Traffic Management > Load Balancing > Virtual Servers, select the virtual server, and click Policies, and bind the AppFlow policy.

Enabling AppFlow for Virtual Servers

Updated: 2014-08-12

If you want to monitor only the traffic through certain virtual servers, enable AppFlow specifically for those virtual servers. You can enable AppFlow for load balancing, content switching, cache redirection, SSL VPN, GSLB, and authentication virtual servers.

To enable AppFlow for a virtual server by using the command line interface

At the command prompt, type:

```
set cs vserver <name> <protocol> <IPAddress> <port> -appflowLog ENABLED
```

Example

```
> set cs vserver Vserver-CS-1 HTTP 10.102.29.161 80 -appflowLog ENABLED
```


To enable AppFlow for a virtual server by using the configuration utility

Navigate to Traffic Management > Content Switching > Virtual Servers, select the virtual server, and enable AppFlow Logging option.

Enabling AppFlow for a Service

Updated: 2014-08-07

You can enable AppFlow for services that are to be bound to the load balancing virtual servers.

To enable AppFlow for a service by using the command line interface

At the command prompt, type:

```
set service <name> -appflowLog ENABLED
```

Example

```
set service ser -appflowLog ENABLED
```

To enable AppFlow for a service by using the configuration utility

Navigate to Traffic Management > Load Balancing > Services, select the service, and enable AppFlow Logging option.

Setting the AppFlow Parameters

Updated: 2014-08-08

You can set AppFlow parameters to customize the exporting of data to the collectors.

To set the AppFlow Parameters by using the command line interface

At the command prompt, type the following commands to set the AppFlow parameters and verify the settings:

- set appflow param [-templateRefresh <secs>] [-appnameRefresh <secs>] [-flowRecordInterval <secs>] [-udpPmtu <positive_integer>] [-httpUrl (**ENABLED** | **DISABLED**)] [-httpCookie (**ENABLED** | **DISABLED**)] [-httpReferer (**ENABLED** | **DISABLED**)] [-httpMethod (**ENABLED** | **DISABLED**)] [-httpHost (**ENABLED** | **DISABLED**)] [-httpUserAgent (**ENABLED** | **DISABLED**)] [-httpXForwardedFor (**ENABLED** | **DISABLED**)] [-clientTrafficOnly (**YES** | **NO**)]
- show appflow Param

Example

```
> set appflow Param -templateRefresh 240 -udpPmtu 128 -httpUrl enabled
```

To set the AppFlow parameters by using the configuration utility

Navigate to System > AppFlow, click Change AppFlow Settings, and specify relevant AppFlow parameters.

Example: Configuring AppFlow for DataStream

Updated: 2013-08-20

The following example illustrates the procedure for configuring AppFlow for DataStream using the command line interface.

```
> enable feature appflow
> add db user sa password freebsd
> add lbvserver lb0 MSSQL 10.102.147.97 1433 -appflowLog ENABLED
> add service sv0 10.103.24.132 MSSQL 1433 -appflowLog ENABLED
> bind lbvserver lb0 sv0
> add appflow collector col0 -IPAddress 10.102.147.90
> add appflow action act0 -collectors col0
> add appflow policy pol0 "mssql.req.query.text.contains(\"select\")" act0
> bind lbvserver lb0 -policyName pol0 -priority 10
```

When the Netscaler appliance receives a database request, the appliance evaluates the request against a configured policy. If a match is found, the details are sent to the AppFlow collector configured in the policy.

Exporting Performance Data of Web Pages to AppFlow Collector

The EdgeSight Monitoring application provides web page monitoring data with which you can monitor the performance of various Web applications served in a Netscaler environment. You can now export this data to AppFlow collectors to get an in-depth analysis of the web page applications. AppFlow, which is based on IPFIX standard, provides more specific information about web application performance than does EdgeSight monitoring alone.

You can configure both load balancing and content switching virtual servers to export EdgeSight Monitoring data to AppFlow collectors. Before configuring a virtual server for AppFlow export, associate an Appflow action with the EdgeSight Monitoring responder policy.

The following web page performance data is exported to AppFlow:

- **Page Load Time.** Elapsed time, in milliseconds, from when the browser starts to receive the first byte of a response until the user starts to interact with the page. At this stage, all the page content might not be loaded.
- **Page Render Time.** Elapsed time, in milliseconds, from when the browser receives the first byte of response until either all page content has been rendered or the page load action has timed out.
- **Time Spent on the Page.** Time spent by users on a page. Represents the period of time from one page request to the next one.

AppFlow transmits the performance data by using the Internet Protocol Flow Information eXport (IPFIX) format, which is an open Internet Engineering Task Force (IETF) standard defined in RFC 5101. The AppFlow templates use the following enterprise-specific Information Elements (IEs) to export the information:

- **Client Load End Time.** Time at which the browser received the last byte of a response to load all the objects of the page such as images, scripts, and stylesheets.
- **Client Load Start Time.** Time at which the browser receives the first byte of the response to load any objects of the page such as images, scripts, and stylesheets.
- **Client Render End Time.** Time at which browser finished rendering the entire page, including the embedded objects.
- **Client Render Start Time.** Time at which the browser started rendering the page.

This topic includes the following details:

- [Prerequisites for Exporting Performance Data of Web Pages to AppFlow Collectors](#)
- [Associating an AppFlow Action with the EdgeSight Monitoring Responder Policy](#)

Prerequisites for Exporting Performance Data of Web Pages to AppFlow Collectors

Updated: 2013-09-13

Before associating the AppFlow action with the AppFlow policy, verify that the following prerequisites have been met:

- The AppFlow feature has been enabled and configured. For instructions, see "[Configuring the AppFlow feature](#)".
- The Responder feature has been enabled. For instructions, see "[Enabling a Responder Feature](#)".
- The EdgeSight Monitoring feature has been enabled. For instructions, see "[Enabling an Application for EdgeSight Monitoring](#)".
- EdgeSight Monitoring has been enabled on the load balancing or content switching virtual servers bound to the services of applications for which you want to collect the performance data. For instructions, see "[Enabling an Application for EdgeSight Monitoring](#)".

Associating an AppFlow Action with the EdgeSight Monitoring Responder Policy

Updated: 2013-10-31

To export the web page performance data to the AppFlow collector, you must associate an AppFlow action with the EdgeSight Monitoring responder policy. An AppFlow action specifies which set of collectors receive the traffic.

To associate an AppFlow action with the EdgeSight Monitoring Responder policy by using the command line interface

At the command prompt, type:

```
set responder policy <name> -appflowAction <action_Name>
```

Example

```
set responder policy pol -appflowAction actn
```

To associate an AppFlow action with the EdgeSight Monitoring Responder policy by using the configuration utility

1. Navigate to AppExpert > Responder > Policies.
2. In the details pane, select an EdgeSight Monitoring responder policy, and then click **Open**.
3. In the **Configure Responder Policy** dialog box, in the **AppFlow Action** drop-down list, select the AppFlow action associated with the collectors to which you want to send the web-page performance data.
4. Click **OK**.

Configuring a Virtual Server to Export EdgeSight Statistics to Appflow Collectors

To export EdgeSight statistics information from a virtual server to the AppFlow collector, you must associate an AppFlow action with the virtual server.

To associate an AppFlow action with a Load Balancing or Content Switching virtual server by using the configuration utility

1. Navigate to Traffic Management > Load Balancing > Virtual Servers or Traffic Management > Content Switching > Virtual Servers.
2. In the details pane, select a virtual server, or multiple virtual servers, and then click Enable EdgeSight Monitoring.
3. In the Enable EdgeSight Monitoring dialog box, select the Export EdgeSight statistics to Appflow check box.
4. From the Appflow Action drop-down list, select the AppFlow action. The AppFlow action defines the list of AppFlow collectors to which it exports EdgeSight Monitoring statistics. If you have selected multiple load balancing virtual servers, the same AppFlow Action will be associated with the responder policies bound to them. You can later change the AppFlow Action configured for each of the selected Load Balancing virtual server individually, if required.
5. Click OK.

AutoScale

Efficient hosting of applications in a cloud requires continuous optimization of application availability. To meet increasing demand, you have to scale network resources upward. When demand subsides, you need to scale down to avoid the unnecessary cost of idle resources. To minimize the cost of running the application by deploying only as many instances as are necessary during any given period of time, you have to constantly monitor traffic. However, monitoring traffic manually is not a feasible option. For the application environment to be able to scale up or down rapidly, you need to automate the processes of monitoring traffic and of scaling resources up and down whenever necessary.

If your organization uses Citrix CloudPlatform to deploy and manage the cloud environment, a Citrix NetScaler appliance can automatically scale users' applications as needed. The CloudPlatform elastic load balancing feature includes a feature called *AutoScale*. A CloudPlatform user can use the AutoScale feature to specify thresholds for various conditions for automatically scaling the application fleet upward and downward. The scale-up and scale-down conditions can vary from simple use cases, such as a server's CPU usage, to complex use cases, such as a combination of a server's CPU usage and responsiveness. CloudPlatform, in turn, configures the NetScaler appliance to load balance traffic to the application virtual machines (VMs), monitor application thresholds and performance, and trigger scale-up and scale-down actions to add or remove VMs to or from the application fleet.

The CloudPlatform user performs all AutoScale configuration tasks by using the CloudPlatform user interface or APIs. The CloudPlatform user:

1. Creates a load balancing rule, with the necessary load balancing algorithm and stickiness.
2. Configures AutoScale parameters by specifying the application instance template, the minimum number of instances to maintain, the maximum number of instances permitted, scale-up and scale-down policies, and other information necessary for the functioning of the feature.
3. Submits the configuration.

For information about configuring a load balancing rule and AutoScale, see *Citrix CloudPlatform 3.0.5 (powered by Apache CloudStack) Administrator's Guide*, at <http://support.citrix.com/article/CTX134823>.


When the CloudPlatform user completes the AutoScale configuration, CloudPlatform uses the NetScaler NITRO API to push all the necessary configuration commands to the NetScaler appliance. As the NetScaler administrator, you do not have to perform any tasks for configuring AutoScale on the NetScaler appliance. However, you might have to be aware of certain prerequisites, and you might have to troubleshoot the configuration if issues arise in the AutoScale configuration. To troubleshoot the configuration, you have to be aware of how CloudPlatform works and what configuration CloudPlatform pushes to the NetScaler appliance. You also need a working knowledge of how to troubleshoot issues on a NetScaler appliance.

How AutoScale Works

When the CloudPlatform user completes the AutoScale configuration, CloudPlatform uses the NetScaler NITRO API to create an AutoScale-related configuration on the NetScaler appliance. For information about the configuration commands that CloudPlatform uses to configure the NetScaler appliance, see "[NetScaler Configuration Details](#)."

The following diagram shows the sequence of operations, beginning with CloudPlatform pushing the AutoScale configuration to the NetScaler appliance. The events are numbered in the order in which they occur, and are described below.

Figure 1. AutoScale Architecture



When the CloudPlatform user submits the AutoScale configuration, the following events occur:

1. CloudPlatform uses the NetScaler NITRO API to push the AutoScale configuration to the NetScaler appliance, creating AutoScale-related entities on the appliance. The entities include a load balancing virtual server, a service group, and monitors.
2. The AutoScale engine on the NetScaler appliance sends API requests to CloudPlatform to initially deploy the minimum number of virtual machines required.
3. CloudPlatform provisions the minimum number of instances (VMs) on the hypervisors (virtualization hosts) that it manages.
4. The NetScaler appliance discovers the IP addresses assigned by CloudPlatform to the newly created VMs and binds them, as services, to the service group representing them. The NetScaler appliance can then load balance traffic to the VMs.
5. NetScaler monitors bound to the service group start monitoring the load by collecting SNMP metrics from the instances.
6. The AutoScale engine on the NetScaler appliance monitors the metrics collected from the VMs and triggers scale-up and scale-down events whenever the metrics breach the configured threshold for the specified period. As part of the scale-up trigger, the NetScaler AutoScale engine sends an API request to CloudPlatform to deploy a new VM. After the virtual machine is deployed, the AutoScale engine binds the service representing the VM (IP address and port) to the service group and, after the configured quiet time, starts forwarding load balanced traffic to the new virtual machine. Likewise, as part of the scale-down trigger, the NetScaler AutoScale engine selects a VM, stops forwarding new requests to that instance, and waits for the configured quiet time (to allow for the processing of current requests to complete) before it sends an API request to CloudPlatform to destroy the chosen instance.

In this way, the NetScaler appliance monitors the application and triggers scale-up and scale-down events on the basis of application load and/or performance.

Supported Environment

AutoScale is supported in the following environment:

- Citrix CloudPlatform 3.0.5.
- Citrix NetScaler MPX/SDX/virtual appliance running Citrix NetScaler release 10.e and later.
- SNMP v1/v2.

Prerequisites

Before you set up AutoScale, do the following:

- Make sure that CloudPlatform is reachable from the NetScaler appliance. You can do so by logging on to the NetScaler appliance and sending ping requests to the CloudPlatform server's IP address.
- Make sure that the network service offering used in CloudPlatform includes the NetScaler appliance as an external load balancing device.
- Use a CloudPlatform and NetScaler release that supports AutoScale. For information about NetScaler releases that support AutoScale, see "[Supported Environment](#)."

If you have to troubleshoot an AutoScale setup, you also have to know the prerequisites for setting up AutoScale in CloudPlatform. See the "Prerequisites" section of "Configuring AutoScale" in the *Citrix CloudPlatform 3.0.5 (powered by Apache CloudStack) Administrator's Guide*, at <http://support.citrix.com/article/CTX134823>.

NetScaler Configuration Details

The following table describes the AutoScale configuration commands that are used by Citrix CloudPlatform to configure a NetScaler appliance.

Table 1. NetScaler Configuration for AutoScale

AutoScale configuration command(s)
<pre>add lb vserver Cloud-VirtualServer-192.0.2.116-22 TCP 192.0.2.116 22 -persistenceType NONE</pre>
<pre>add serviceGroup Clouda35a6b6b76614006b97476e841b80f79 TCP -maxClient 0 -maxReq 0 -cip DISA bind lb vserver Cloud-VirtualServer-192.0.2.116-22 Clouda35a6b6b76614006b97476e841b80f79</pre>
<pre>add server autoscale-internal_server_Clouda35a6b6b76614006b97476e841b80f79 autoscale-interna bind serviceGroup Clouda35a6b6b76614006b97476e841b80f79 autoscale-internal_server_Clouda35a</pre>
<pre>add lb metricTable Cloud-MTb1-192.0.2.116-22 bind lb metricTable Cloud-MTb1-192.0.2.116-22 Linux_User_CPU_-_percentage 1.3.6.1.4.1.2021.1 add lb monitor Cloud-Mon-192.0.2.116-22 LOAD -interval 24 -destPort 161 -snmpCommunity publ bind lb monitor Cloud-Mon-192.0.2.116-22 -metric Linux_User_CPU_-_percentage -metricThresho bind serviceGroup Clouda35a6b6b76614006b97476e841b80f79 -monitorName Cloud-Mon-192.0.2.116-22</pre>
<pre>add autoscale profile Cloud-AutoScale-Profile-192.0.2.116-22 -type CLOUDSTACK -url "http://: xSTRlHsob911600VO1FMxvE1U017uoD6_Z0bkkLaVtK5Y10oBkTzgbTwp3u51CQ</pre>

```
add autoscale action Cloud-AutoScale-ScaleUpAction-192.0.2.116-22 -type SCALE_UP -profileName 208c-47a8-9c16-d582550cf759&displayname=AutoScale-LB-lb&networkids=a3c97129-b729-4c72-994f-
```

```
add autoscale action Cloud-AutoScale-ScaleDownAction-192.0.2.116-22 -type SCALE_DOWN -profileName 192.0.2.116-22
```

```
add autoscale policy Cloud-AutoScale-Policy-Min-192.0.2.116-22 -rule "SYS.VSERVER(\"Cloud-Vi
```

```
add autoscale policy Cloud-AutoScale-Policy-Max-192.0.2.116-22 -rule "SYS.VSERVER(\"Cloud-Vi
```

```
add autoscale policy Cloud-AutoScale-Policy-192.0.2.116-22-35 -rule "SYS.VSERVER(\"Cloud-Vi  
(90))" -action Cloud-AutoScale-ScaleUpAction-192.0.2.116-22
```

```
add autoscale policy Cloud-AutoScale-Policy-192.0.2.116-22-36 -rule "SYS.VSERVER(\"Cloud-Vi  
(30))" -action Cloud-AutoScale-ScaleDownAction-192.0.2.116-22
```

```
add ns timer Cloud-AutoScale-Timer-192.0.2.116-22 -interval 30
```

```
bind ns timer Cloud-AutoScale-Timer-192.0.2.116-22 -policyName Cloud-AutoScale-Policy-Min-192.0.2.116-22
```

```
bind ns timer Cloud-AutoScale-Timer-192.0.2.116-22 -policyName Cloud-AutoScale-Policy-Max-192.0.2.116-22
```

```
bind ns timer Cloud-AutoScale-Timer-192.0.2.116-22 -policyName Cloud-AutoScale-Policy-192.0.2.116-22-35
```

```
bind ns timer Cloud-AutoScale-Timer-192.0.2.116-22 -policyName Cloud-AutoScale-Policy-192.0
```

Troubleshooting

Before you attempt to resolve an AutoScale issue, make sure that the prerequisites have been adhered to, on both the CloudPlatform server and the NetScaler appliance, as described in "Prerequisites." If that does not resolve the issue, your problem could be one of the following.

The AutoScale configuration was successfully configured in CloudPlatform. Yet, the minimum number of VMs has not been created.

- Recommend that the CloudPlatform user deploy one VM manually in the network before configuring AutoScale. Ask the user to remove the AutoScale configuration from the NetScaler appliance or the load balancer from the network, manually deploy one VM (preferably using the template created for the AutoScale configuration), and then create the AutoScale configuration.
- Verify that the CloudPlatform user has configured the VM template in such a way that the VMs that are created from the template can accept traffic without manual intervention. If a provisioned VM cannot accept traffic automatically, the metric remains above the threshold, and the AutoScale configuration continues to provision additional VMs, as designed. To remedy the issue, disable AutoScale from CloudPlatform, fix the template, and then enable AutoScale.
- Verify that the CloudPlatform user has not exceeded the limit for the number of VMs imposed by the user's account.
- Verify that the CloudPlatform server is up and is reachable from the NetScaler appliance.
- Verify that the CloudPlatform log file, management-server.log, has reported the successful creation of the AutoScale configuration in CloudPlatform.
- Verify that the scale-up policy that is responsible for initial scale up (the policy name is prefixed with Cloud-AutoScale-Policy-Min) is receiving hits.

The AutoScale configuration is rapidly spawning a large number of VMs

- Verify that the CloudPlatform user has configured the VM template in such a way that the VMs that are created from the template can accept traffic without manual intervention. If a provisioned VM cannot accept traffic automatically, the metric remains above the threshold, and the AutoScale configuration continues to provision additional VMs, as designed. To remedy the issue, disable AutoScale from CloudPlatform, fix the template, and then enable AutoScale.
- Verify that the quiet time that the CloudPlatform user has configured in the AutoScale configuration is sufficient to ensure even traffic distribution to all the VMs, including the new VM. If the quiet time is too low, and traffic distribution has not stabilized, the metrics might remain above the threshold, and additional VMs might be spawned.

When I ran the top command on my VM, I noticed that the CPU usage on my VM had breached the threshold that was configured for the scale-up action in AutoScale. Yet, the application is not scaling up.

- Verify that the CloudPlatform user has installed an SNMP agent in the VM template, and that the SNMP agent is up and running on every VM.
- Verify that the CloudPlatform user has not exceeded the limit for the number of VMs imposed by the user's account.
- Verify that the CloudPlatform user has correctly configured the SNMP parameters to collect metrics from the VM (for example, the community string and the port).
- Verify that the scale-up or scale-down policy is receiving hits.
- Verify that the CloudPlatform server is up, and that the CloudPlatform server is reachable from the NetScaler appliance.

One or more additional VMs have been created, but they are not accepting traffic (that is, VMs have been created, but the average value of the metrics is still above the threshold)

- Verify that the user has configured the templates in such a way that the VMs created from the templates can start serving traffic without any manual intervention.
- Verify that the service is running on the VMs, on the configured member port.

- Send a ping request to the gateway (virtual router), from the VM that is not accepting traffic.

The AutoScale configuration has been deleted, but the VMs continue to exist

- The VMs might not be deleted immediately after the AutoScale configuration is deleted. Wait for about 5 minutes after you have deleted the AutoScale configuration, and then check again.
- If the destruction of VMs has not commenced after 5 minutes, you might have to delete the VMs manually.

Clustering

A NetScaler cluster is a group of nCore appliances working together as a single system image. Each appliance of the cluster is called a node. The cluster can include as few as 2 or as many as 32 NetScaler nCore hardware or virtual appliances as nodes.

The client traffic is distributed between the nodes to provide high availability, high throughput, and scalability.

To create a cluster, you must add the appliances as cluster nodes, set up communication between the nodes, set up links to the client and server networks, configure the appliances, and configure the distribution of client and server traffic.

NetScaler Feature-level Support in a Cluster

The following table lists the features that are not supported in the NetScaler 10 release and also states the support for the feature in subsequent NetScaler releases.

Note:

- Unless stated otherwise in this table, all NetScaler features are supported in a cluster.
- The entry "Node-level" in the table indicates that the feature is supported only on individual cluster nodes.

NetScaler Feature	10	10.1	10.5	11
SSL (classic policies) Note: SSL - advanced policies are supported from NetScaler 10 onwards.	No	No	No	No
SSL FIPS	No	No	No	No
SSL Certificate Bundle	No	No	No	No
Content switching actions Note: The content switching feature is supported from NetScaler 10 onwards.	No	Yes	Yes	Yes
Policy-based logging for content switching policies	No	Yes	Yes	Yes
Rate limiting	No	Yes	Yes	Yes
Action analytics	No	Yes	Yes	Yes
Branch Repeater load balancing	No	Yes	Yes	Yes
GSLB	No	No	Yes	Yes
RTSP	No	No	No	No
DNSSEC	No	No	No	No
DNS64	No	No	No	No
FTP	No	No	No	Yes
TFTP	No	No	No	No
Connection mirroring	No	No	No	No
Integrated caching	Node-Level	Node-level	Node-level	Node-level
Large shared cache	No	Node-level	Node-level	Node-level
Application firewall	No	No	Node-level	Node-level
HTTP Denial-of-Service Protection (HDOSP)	Node-level	Node-level	Node-level	Node-level
Priority queuing (PQ)	Node-level	Node-level	Node-level	Node-level
Sure connect (SC)	Node-level	Node-level	Node-level	Node-level
AppQoE	NA	Node-level	Yes	Yes
Surge protection	Node-level	Node-level	Node-level	Node-level
MPTCP	No	No	Yes	Yes
MSR	Yes	Yes	Yes	Yes Note: Supported in L2 clusters. Not supported in L3 clusters.
IS-IS (IPv4 and IPv6)	No	Yes	Yes	Yes

IP-IP tunneling	No	Yes	Yes	Yes
Link load balancing	No	No	Yes	Yes
FIS (Failover Interface Set)	No	No	Yes	Yes
Link redundancy (LR)	No	No	Yes	Yes
NAT46	No	No	No	No
NAT64	No	No	No	No
v6 ReadyLogo	No	No	No	No
Traffic domains	No	No	Yes	Yes Note: Supported in L2 clusters. Not supported in L3 clusters.
Route monitor	No	No	No	No
GRE tunneling (CB)	No	No	No	No
Layer 2 mode	No	No	Yes	Yes
Net profiles	No	No	Yes	Yes
HTTPS callout	No	Yes	Yes	Yes
AAA-TM	No	Node-level	Node-level	Node-level
AppFlow	No	Node-level	Node-level	Node-level
Insight	No	No	No	No
HDX Insight	No	No	No	No
VMAC/VRRP	No	No	Yes	Yes
NetScaler Push	No	No	No	No
Stateful Connection Failover	No	No	No	No
Graceful Shutdown	No	No	No	No
DBS AutoScale	No	No	No	No
DSR using TOS	No	No	No	No
Finer Startup-RR Control	Node-level	Node-level	Node-level	Node-level
XML XSM	No	No	No	No
DHCP RA	No	No	No	No
Bridge Group	No	No	Yes Note: Supported from NetScaler 10.5 Build 52.1115.e onwards.	Yes
Network Bridge	No	No	No	No
Web Interface on NetScaler (WlonNS)	No	No	No	No
EdgeSight Monitoring	No	No	No	No
Metrics tables - Local	No	No	No	No
DNS Caching	Node-level	Node-level	Node-level	Node-level
Call Home	Node-level	Node-level	Node-level	Node-level
NetScaler Gateway or SSL VPN	No	No	Node-level	Node-level
CloudBridge Connector	No	No	No	No

Prerequisites for Cluster Nodes

NetScaler appliances that are to be added to a cluster must satisfy the following criteria:

- A NetScaler cluster can only include NetScaler nCore appliances. Clustering of NetScaler Classic appliances is not supported.
- All appliances must have the same software version and build.

- All appliances must be of the same platform type. This means that a cluster can have either all hardware appliances (MPX) or virtual appliances (VPX) or SDX NetScaler instances.

Note: Clustering of SDX NetScaler instances is supported in NetScaler 10.1 and later releases. To create a cluster of SDX NetScaler instances, see "[Setting up a Cluster of NetScaler Instances](#)".

- For a cluster of hardware appliances (MPX), all appliances must be of the same model type.
- All appliances must be on the same network.
- All appliances must have the same licenses. Also, depending on the NetScaler version, there are some additional aspects to address:

For releases prior to NetScaler 10.5 Build 52.x:

- A separate cluster license file is required. This file must be copied to the /nsconfig/license/ directory of the configuration coordinator.
- Because of the separate cluster license file, the cluster feature is available irrespective of the NetScaler license.

- For releases after NetScaler 10.5 Build 52.x:

- No separate cluster license is required.
- Cluster is licensed with the Enterprise and Platinum licenses. Cluster is not available for Standard license.

- Be initially configured and connected to a common client-side and server-side network.


Note: For a cluster of virtual appliances, that has large configurations, it is recommended to use 6 GB RAM for each node of the cluster.

Cluster Overview

A NetScaler cluster is formed by grouping NetScaler appliances together. Based on the network location of the NetScaler appliances that you intend adding to the cluster, you must be aware of the following cluster setups:

Note: Unless specified otherwise, cluster features and configurations are the same for L2 and L3 clusters.

- o **L2 cluster:** In this cluster deployment, all cluster nodes belong to the same network.




- o **L3 cluster:** In this cluster deployment, cluster nodes can belong to different networks.

Note: Not yet supported. Likely to be supported in future NetScaler versions.

Synchronization Across Cluster Nodes

All configurations on a NetScaler cluster are performed on the cluster IP address, which is the management address of the cluster. This cluster IP address is owned by a cluster node that is referred to as the cluster configuration coordinator as shown in the following figure:



The configurations that are available on the configuration coordinator are automatically propagated to the other cluster nodes and therefore all cluster nodes have the same configurations.

Note:

- NetScaler allows only a few configurations to be performed on individual cluster nodes through their NetScaler IP (NSIP) address. These configurations are not propagated across the other cluster nodes. For more information, see "[Operations Supported on Individual Cluster Nodes](#)".
- The following commands when executed on the cluster IP address are not propagated to other cluster nodes:
 - shutdown: Shuts down only the configuration coordinator.
 - reboot: Reboots only the configuration coordinator.
 - rm cluster instance: Removes the cluster instance from the node that you are executing the command or
- If the NetScaler cluster is configured to use a quorum (quorum is mandatory for versions prior to NetScaler 10.5), a command is propagated to the other cluster nodes only when a majority of the nodes are in synch. If a majority of the nodes are not in synch or are in the process of synchronizing, the new commands cannot be accepted and therefore command propagation is temporarily halted.

When a node is added to a cluster, the configurations and the files (SSL certificates, licenses, DNS, and so on) that are available on the cluster configuration coordinator are synchronized to the newly added cluster node. When an existing cluster node, that was intentionally disabled or that had failed, is once again added, the cluster compares the configurations available on the node with the configurations available on the configuration coordinator. If there is a mismatch in configurations, the node is synchronized by using one of the following:

- **Full synchronization.** If the difference between configurations exceeds 255 commands, all the configurations of the configuration coordinator are applied to the node that is rejoining the cluster. The node remains operationally unavailable for the duration of the synchronization.
- **Incremental Synchronization.** If the difference between configurations is less than or equal to 255 commands, only the configurations that are not available are applied to the node that is rejoining the cluster. The operational state of the node remains unaffected.

Note: You can also manually synchronize the configurations and files. For more information, see "[Synchronizing Cluster Configurations](#)" and "[Synchronizing Cluster Files](#)".

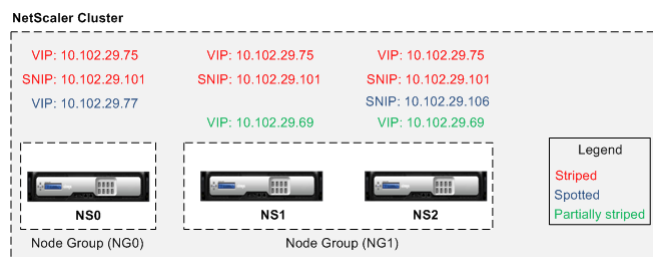
Striped, Partially Striped, and Spotted Configurations

By virtue of command propagation, all nodes in a cluster have the same configurations. However, you may want some configurations to be available only on certain cluster nodes. While you cannot restrict the nodes on which the configurations are available, you can specify the nodes on which the configurations are active.

For example, you can define a SNIP address to be active on only one node, or define a SNIP address to be active on all nodes, or define a VIP address to be active on only one node, or define a VIP address to be active on all nodes, or define a VIP address to be active only on two nodes of a 3-node cluster.

Depending on the number of nodes the configurations are active on, cluster configurations are referred to as striped, partially striped, or spotted configurations.

Figure 1. Three-node cluster with striped, partially striped, and spotted configurations



The following table provides more details on the types of configurations:


Configuration Type	Active on...	Applicable to...	Configurations...
Striped configuration	All the cluster nodes	All entities	No specific configuration required to make an entity striped. By default, all entities defined on a cluster IP address are striped on all the cluster nodes.
Partially striped configuration	A subset of cluster nodes	Refer "Node Groups"	Bind the entities that you want to be partially striped, to a node group. The configuration will be active only on the cluster nodes that belong to the node group.
Spotted configuration	Single cluster node	<ul style="list-style-type: none"> SNIP address SNMP Engine ID Hostname of cluster nodes Entities that can be bound to a node group 	<p>A spotted configuration can be defined using one of two approaches.</p> <ul style="list-style-type: none"> SNIP address. When creating the SNIP address, specify the node on which you want the SNIP address to be active, as the owner node. <p>Example:</p> <pre>add ns ip 10.102.29.106 255.255.255.0 -type SNIP -ownerNode 2 (assuming node NS2 ID is 2)</pre> <p>Note: You cannot change the ownership of a spotted SNIP address at run time. To change the ownership, you must first delete the SNIP address and add it again by specifying the new owner.</p> <ul style="list-style-type: none"> Entities that can be bound to a node group. By binding the entity to a single-member node group.

Note: Citrix recommends that you use spotted SNIP addresses. You can use striped SNIP addresses only if there is a shortage of IP addresses. The use of striped IP addresses can result in ARP flux issues.

Communication in a Cluster Setup

The interfaces of NetScaler appliances that are added to a cluster, are prefixed with a node ID. This helps identify the cluster node to which the interface belongs. Therefore, the interface identifier c/u , where c is the controller number and u is the unit number, now becomes $n/c/u$, where n is the node ID. For example, in the following figure, interface 1/2 of node n1 is represented as 0/1/2, interface 1/1 of node n2 is represented as 1/1/1, and interface 1/4 of node n3 is represented as 2/1/4.

Figure 1. Interface naming convention in a cluster



Server communication

The cluster communicates with the server through the physical connections between the cluster node and the server-side connecting device. The logical grouping of these physical connections is called the server data plane.

Client communication

The cluster communicates with the client through the physical connections between the cluster node and the client-side connecting device. The logical grouping of these physical connections is called the client data plane.

Inter-node communication

The cluster nodes communicate with each other by using the cluster backplane. The backplane is a set of connections in which one interface of each node is connected to a common switch, which is called the cluster backplane switch. Each node of the cluster uses a special MAC address to communicate with other nodes through the backplane.

The following figures show the communication interfaces in L2 clusters and L3 clusters.

Figure 2. Cluster communication interfaces - L2 cluster





Figure 3. Cluster communication interfaces - L3 cluster



Traffic Distribution in a Cluster Setup


In a cluster setup, external networks view the collection of NetScaler appliances as a single entity. So, the cluster must select a single node that must receive the traffic. The cluster does this selection by using Equal Cost Multiple Path (ECMP) or cluster link aggregation traffic distribution mechanism. The selected node is called the flow receiver.

The flow receiver gets the traffic and then, using internal cluster logic determines the node that must process the traffic. This node is called the flow processor. The flow receiver steers the traffic to the flow processor over the backplane.

Note:

- The flow receiver and flow processor must be nodes capable of serving traffic.

Figure 1. Traffic distribution in a cluster



The above figure shows a client request flowing through the cluster. The client sends a request to a virtual IP (VIP) address. A traffic distribution mechanism configured on the client data plane selects one of the cluster nodes as the flow receiver. The flow receiver receives the traffic, determines the node that must process the traffic, and steers the request to that node (unless the flow receiver selects itself as the flow processor).

The flow processor establishes a connection with the server. The server processes the request and sends the response to the subnet IP (SNIP) address that sent the request to the server.

- If the SNIP address is a striped or partially striped IP address, the traffic distribution mechanism configured on the server data plane selects one of the cluster nodes as the flow receiver. The flow receiver receives the traffic, determines the flow processor, and steers the request to the flow processor through the cluster backplane.
- If the SNIP address is a spotted IP address, the node that owns the SNIP address receives the response from the server.


In an asymmetric cluster topology (all cluster nodes are not connected to the external switch), you must use linksets either exclusively or combined with ECMP or cluster link aggregation. For more information, see ["Using Linksets"](#).

Cluster Nodegroups

Note: Nodegroups are supported from NetScaler 10.1 onwards.

As the name indicates, a cluster nodegroup is a group of cluster nodes.

Figure 1. NetScaler cluster with nodegroups



The above figure shows a cluster which has nodegroups NG1 and NG2 that include 3 cluster nodes each. The cluster also has 3 nodes that are not part of any nodegroup.

A nodegroup can be configured for the following:

- To define spotted and partially striped configurations. For more information, see ["Nodegroups for Spotted and Partially-Striped Configurations"](#).
- To configure redundancy of nodegroups. For more information, see ["Configuring Redundancy for Nodegroups"](#).

Note: Supported from NetScaler 10.5 Build 52.1115.e onwards.

Note: The above functions of a nodegroup are mutually exclusive. This means that a nodegroup can provide only one of the above mentioned functionality.

Cluster and Node States

For a cluster to be functional, a majority of the nodes ($n/2 + 1$) must be online and be able to serve traffic by satisfying the following criteria:

- Admin state must be ACTIVE
- Operational state must be ACTIVE
- Health status must be UP

Note: Alternatively, from NetScaler release 10.5, you can configure the cluster to be functional even when the majority criteria is not satisfied. This configuration must be performed when creating a cluster.

The following table describes the states of a cluster.

Type	Description
Admin	<p>An admin state is configured when you add the node to the cluster. It indicates the purpose of the node, which can be in one of the following states:</p> <ul style="list-style-type: none">• Active. Nodes in this state serve traffic if they are operationally active and healthy.• Passive. Nodes in this state do not serve traffic but are in sync with the cluster. This is the default state of a cluster node.• Spare. Nodes in this state do not serve traffic but are in sync with the cluster. Spare nodes act as backup nodes for the cluster. If one of the nodes in the ACTIVE admin state becomes unavailable, a spare node becomes operationally active and starts serving traffic. <p>Note: Whether the spare node remains operationally active depends on the preemption parameter of the add cluster instance command. If preemption is disabled, the spare node continues to serve traffic even if a node in ACTIVE admin state comes back online. If preemption is enabled, when a node in ACTIVE admin state comes back online, it preempts the spare node and starts serving traffic. The spare node goes back to inactive state.</p>
Operational	<p>When a node is part of a cluster, its operational state can change to ACTIVE, INACTIVE, or UNKNOWN. There are a number of reasons for a node being in INACTIVE or UNKNOWN state. Review the ns.log file or error counters to help determine the exact reason.</p> <p>Note: Passive nodes are always operationally INACTIVE. Spare nodes are ACTIVE only when they are serving traffic. Else, they are operationally INACTIVE.</p>
Health	<p>Depending on its health, a node can either be UP or NOT UP. To view the reasons for a node being in NOT UP state, run the show cluster node command for that node.</p>

Routing in a Cluster

Routing in a cluster works in much the same way as routing in a standalone system. A few points to note:

- Routing runs only on spotted SNIP addresses and NSIP addresses.
- All routing configurations must be performed from the cluster IP address and the configurations are propagated to the other cluster nodes.
- Node-specific routing configurations must be performed by using the owner-node argument as follows:

```
!  
interface vlan97  
!  
router ospf  
  owner-node 0  
    ospf router-id 97.131.0.1  
  exit-owner-node  
  owner-node 1  
    ospf router-id 97.131.0.2  
  exit-owner-node  
  owner-node 2  
    ospf router-id 97.131.0.3  
  exit-owner-node  
  redistribute kernel  
  network 97.0.0.0/8 area 0  
!
```

- Retrieve node-specific routing configurations by specifying the node(s) in the owner-node argument as follows:

```
> vtysh  
ns# owner-node 0 1  
ns(node-0 1)# show cluster state  
ns(node-0 1)# exit-owner-node
```

- Clear node-specific routing configurations by specifying the node(s) in the owner-node argument as follows:

```
> vtysh  
ns# owner-node 0 1  
ns(node-0 1)# clear config  
ns(node-0 1)# exit-owner-node
```

- Routing protocol daemons can run and adjacencies can be formed on active and inactive nodes of a cluster.
- Only active nodes advertise host routes to VIP addresses.
- Active and inactive nodes can learn dynamic routes and install them into the routing table.
- Routes learnt on a node are propagated to other nodes in the cluster only if route propagation is configured. This is mostly needed in asymmetric topologies where the unconnected nodes may not be able to form adjacencies.

```
ns(config)# ns route-install propagate
```

Note: Make sure that route propagation is not configured in a symmetric cluster topology as it can result in making the node unavailable to the cluster.

Setting up a NetScaler Cluster

NetScaler appliances that you want to add to the cluster must satisfy the criteria specified in "[Prerequisites for Cluster Nodes](#)". Before actually setting up a cluster, you must be aware of cluster basics. For information, see "[Cluster Overview](#)".

Forming a cluster requires you to set up inter-node communication, create the cluster (by adding the first NetScaler appliance), and then add the other cluster nodes. Each of these steps is explained with relevant details in subsequent topics.

Setting up Inter-Node Communication

The nodes in a cluster communicate with each other through the cluster backplane.

To set up the cluster backplane, do the following for every node

1. Identify the network interface that you want to use for the backplane.
2. Connect an Ethernet or optical cable from the selected network interface to the cluster backplane switch.

For example, to use interface 1/2 as the backplane interface for node 4, connect a cable from the 1/2 interface of node 4 to the backplane switch.

Important points to note when setting up the cluster backplane

- Do not use the appliance's management interface (0/1) as the backplane interface.
- Backplane interfaces must not be used for the client or server data planes.
- Configure a link aggregate (LA) channel to optimize the throughput of the cluster backplane.
- Citrix recommends that you dedicate a separate switch for the backplane, so that large amounts of traffic can be handled seamlessly.
- Backplane interfaces of all nodes of a cluster must be connected to the same switch and bound to the same L2 VLAN. The backplane interfaces, by default, have presence on all L3 VLANs configured on the cluster.
- If you have multiple clusters with the same cluster instance ID, make sure that the backplane interfaces of each cluster are bound to a different VLAN.
- The backplane interface is always monitored, regardless of the HA monitoring settings of that interface.
- The state of MAC spoofing on the different virtualization platforms can affect the steering mechanism on the cluster backplane. Therefore, make sure the appropriate state is configured:
 - XenServer - Disable MAC spoofing
 - HyperV - Enable MAC spoofing
 - ESX - Enable MAC spoofing (also make sure "Forged Transmits" is enabled)
- The Maximum Transmission Unit (MTU) for interfaces of the backplane switch must be greater than or equal to 1578 bytes, if features like MBF, L2 policies, ACLs, routing in CLAG deployments are configured. The MTU on the cluster backplane is automatically updated.

Creating a NetScaler Cluster

To create a cluster, start by taking one of the NetScaler appliances that you want to add to the cluster. On this node, you must create the cluster instance and define the cluster IP address. This node is the first cluster node and is called the cluster configuration coordinator. All configurations that are performed on the cluster IP address are stored on this node and then propagated to the other cluster nodes.

The responsibility of configuration coordination in a cluster is not fixed to a specific node. It can change over time depending on the following factors:

- The priority of the node. The node with the highest priority (lowest priority number) is made the configuration coordinator. Therefore, if a node with a priority number lower than that of the existing configuration coordinator is added, the new node takes over as the configuration coordinator.
Note: Node priority can be configured from NetScaler 10.1 onwards.
- If the current configuration coordinator goes down. The node with the next lowest priority number takes over as the configuration coordinator. If the priority is not set or if there are multiple nodes with the lowest priority number, the configuration coordinator is selected from one of the available nodes.

Note: The configurations of the appliance (including SNIP addresses and VLANs) are cleared by implicitly executing the `clear ns config extended` command. However, the default VLAN and NSVLAN are not cleared from the appliance. Therefore, if you want the NSVLAN on the cluster, make sure it is created before the appliance is added to the cluster.

To create a cluster by using the command line interface

1. Log on to an appliance (for example, appliance with NSIP address 10.102.29.60) that you want to add to the cluster.
2. Add a cluster instance.

```
add cluster instance <cld> -quorumType <NONE | MAJORITY>
```

Note:

- The cluster instance ID must be unique within a LAN.

3. Add the NetScaler appliance to the cluster.

```
add cluster node <nodeId> <IPAddress> -state <state> -backplane <interface_name>
```

Example

Adding a node for an L2 cluster (all cluster nodes are in the same network).

```
> add cluster node 0 10.102.29.60 -state PASSIVE -backplane 0/1/1
```

4. Add the cluster IP address (for example, 10.102.29.61) on this node.

```
add ns ip <IPAddress> <netmask> -type clip
```

Example

```
> add ns ip 10.102.29.61 255.255.255.255 -type clip
```

5. Enable the cluster instance.

```
enable cluster instance <cld>
```

6. Save the configuration.

```
save ns config
```

7. Warm reboot the appliance.

```
reboot -warm
```

Verify the cluster configurations by using the `show cluster instance` command. Verify that the output of the command displays the NSIP address of the appliance as a node of the cluster.

To create a cluster by using the configuration utility

1. Log on to an appliance (for example, an appliance with NSIP address 10.102.29.60) that you intend to add to the cluster.
2. Navigate to System > Cluster.
3. In the details pane, click the Manage Cluster link.

4. In the Cluster Configuration dialog box, set the parameters required to create a cluster. For a description of a parameter, hover the mouse cursor over the corresponding text box.
5. Click Create.
6. In the Configure cluster instance dialog box, make sure that the Enable cluster instance check box is selected.
7. In the Cluster Nodes pane, select the node and click Open.
8. In the Configure Cluster Node dialog box, set the State.
9. Click OK, and then click Save.
10. Warm reboot the appliance.

Adding a Node to the Cluster

You can seamlessly scale the size of a cluster to include a maximum of 32 nodes. When a NetScaler appliance is added to the cluster, the configurations from that appliance are cleared and cluster configurations are synchronized on this node. There can be an intermittent drop in traffic while the synchronization is in progress.

Note:

- The licenses of the appliance are checked against the licenses available on the configuration coordinator. The appliance is added if the licenses match.
- If you use the NetScaler CLI to add a node, the new node does not become a functional part of the cluster until it is explicitly joined to the cluster. Therefore, after adding the node, log on to that node and join the node to the cluster. Alternatively, you can add the node from the command line and use the configuration utility to join the node to the cluster. If you use the configuration utility, you need only log on to the cluster IP address and add the node. The newly added node is automatically joined to the cluster.

Important: Before you add a NetScaler appliance to a cluster:

- Set up the backplane interface for the node.
- If you want the NSVLAN on the cluster, make sure that the NSVLAN is created on the appliance before it is added to the cluster.
- Citrix recommends that you add the node as a passive node. Then, after joining the node to the cluster, complete the node specific configuration from the cluster IP address. Run the force cluster sync command if the cluster has only spotted IP addresses, has L3 VLAN binding, or has static routes.
- When an appliance with a preconfigured link aggregate (LA) channel is added to a cluster, the LA channel continues to exist in the cluster environment. The LA channel is renamed from LA/x to nodeId/LA/x, where LA/x is the LA channel identifier.

To add a node to the cluster by using the command line interface

1. Log on to the cluster IP address and, at the command prompt, do the following:
 - a. Add the appliance (for example, 10.102.29.70) to the cluster.

```
add cluster node <nodeId> <IPAddress> -state <state> -backplane <interface_name>
```

Example

```
> add cluster node 1 10.102.29.70 -state PASSIVE -backplane 1/1/1
```

- b. Save the configuration.

```
save ns config
```

2. Log on to the newly added node (for example, 10.102.29.70) and do the following:
 - a. Join the node to the cluster.

```
join cluster -clip <ip_addr> -password <password>
```

Example

```
> join cluster -clip 10.102.29.61 -password nsroot
```

- b. Save the configuration.

```
save ns config
```

- c. Warm reboot the appliance.

```
reboot -warm
```

To add a node to the cluster by using the configuration utility

1. Log on to the cluster IP address.
2. Navigate to System > Cluster > Nodes.
3. In the details pane, click Add to add the new node (for example, 10.102.29.70).
4. In the Create Cluster Node dialog box, configure the new node. For a description of a parameter, hover the mouse cursor over the corresponding text box.
5. Click Create. When prompted to perform a warm reboot, click Yes.

To join a previously added node to the cluster by using the configuration utility

If you have used the command line to add a node to the cluster, but have not joined the node to the cluster, you can use the following procedure.

Note: When a node joins the cluster, it takes over its share of traffic from the cluster and hence an existing connection can be terminated.

1. Log on to the node that you want to join to the cluster (for example, 10.102.29.70).
2. Navigate to System > Cluster.
3. In the details pane, under Get Started, click the Join Cluster link.
4. In the Join to existing cluster dialog box, set the cluster IP address and the nsroot password of the configuration coordinator. For a description of a parameter, hover the mouse cursor over the corresponding text box.
5. Click OK.

Viewing the Details of a Cluster

You can view the details of the cluster instance and the cluster nodes by logging on to the cluster IP address.

To view details of a cluster instance by using the command line interface

Log on to the cluster IP address and, at the command prompt, type:

```
show cluster instance <cld>
```

Note: When executed from the NSIP address of a cluster node that is not the configuration coordinator, this command displays the status of the cluster on this node.

To view details of a cluster node by using the command line interface

Log on to the cluster IP address and, at the command prompt, type:

```
show cluster node <nodeid>
```

To view details of a cluster instance by using the configuration utility

1. Log on to the cluster IP address.
2. Navigate to System > Cluster.
3. In the details pane, under Get Started, click the Manage Cluster link to view the details of the cluster.

To view details of a cluster node by using the configuration utility

1. Log on to the cluster IP address.
2. Navigate to System > Cluster > Nodes.
3. In the details pane, click the node for which you want to view the details.

Distributing Traffic Across Cluster Nodes

After you have created the NetScaler cluster and performed the required configurations, you must deploy Equal Cost Multiple Path (ECMP) or cluster Link Aggregation (LA) on the client data plane (for client traffic) or server data plane (for server traffic). These mechanisms distribute external traffic across the cluster nodes.

Using Equal Cost Multiple Path (ECMP)

With the Equal Cost Multiple Path (ECMP) mechanism, the router has equal-cost routes to VIP addresses with the next hops as the active nodes of the cluster. The router uses a stateless hash-based mechanism to distribute traffic across the routes.

Note: Routes are limited to the maximum number of ECMP routes supported by the upstream router.

To use ECMP, you must first enable the required routing protocol (OSPF, RIP, BGP, or ISIS) on the cluster IP address. You must bind the interfaces and the spotted IP address (with dynamic routing enabled) to a VLAN. Configure the selected routing protocol and redistribute the kernel routes on the ZebOS by using the vtysh shell.


You must perform similar configurations on the cluster IP address and on the external connecting device.

Note:

- All routing configurations must be done through the cluster IP address. No configurations must be performed on individual cluster nodes.
- Make sure that the licenses on the cluster support dynamic routing, otherwise ECMP does not work.
- ECMP is not supported for wildcard virtual servers since RHI needs a VIP address to advertise to a router and wildcard virtual servers do not have associated VIP addresses.

You must have detailed knowledge of routing protocols to use ECMP. For more information, see ["Configuring Dynamic Routes"](#). For more information on routing in a cluster, see ["Routing in a Cluster"](#).

Figure 1. ECMP topology



As seen in the above figure, the ECMP router can reach the VIP address via SNIP0, SNIP1, or SNIP2.

To configure ECMP on the cluster by using the command line interface

1. Log on to the cluster IP address.
2. Enable the routing protocol.

```
enable ns feature <feature>
```

Example: To enable the OSPF routing protocol.

```
> enable ns feature ospf
```

3. Add a VLAN.

```
add vlan <id>
```

Example

```
> add vlan 97
```

4. Bind the interfaces of the cluster nodes to the VLAN.

```
bind vlan <id> -ifnum <interface_name>
```

Example

```
> bind vlan 97 -ifnum 0/1/2 1/1/2 2/1/2
```

5. Add a spotted SNIP address for each node and enable dynamic routing on it.

```
add ns ip <SNIP> <netmask> -ownerNode <positive_integer> -dynamicRouting ENABLED
```

Example

```
> add ns ip 97.131.0.1 255.0.0.0 -ownerNode 0 -dynamicRouting ENABLED -type SNIP > add
```

6. Bind one of the spotted SNIP addresses to the VLAN. When you bind one spotted SNIP address to a VLAN, all other spotted SNIP addresses defined on the cluster in that subnet are automatically bound to the VLAN.

```
bind vlan <id> -IPAddress <SNIP> <netmask>
```

Example

```
> bind vlan 97 -ipAddress 97.131.0.1 255.0.0.0
```

Note: You can use NSIP addresses of the cluster nodes instead of adding SNIP addresses. If so, you do not have to perform steps 3 - 6.

7. Configure the routing protocol on ZebOS using vtysh shell.

Example: To configure OSPF routing protocol on node IDs 0, 1, and 2.

```
> vtysh      ! interface vlan97      ! router ospf      owner-node 0      ospf router-id 97.131.
```

Note: For VIP addresses to be advertised, RHI setting must be done by using the vserverRHILevel parameter as follows:

```
add ns ip <IPAddress> <netmask> -type VIP -vserverRHILevel <vserverRHILevel>
```

For OSPF specific RHI settings, there are additional settings that can be done as follows:

```
add ns ip <IPAddress> <netmask> -type VIP -ospfLSAType ( TYPE1 | TYPE5 ) -ospfArea <positive_integer>
```

Use the add ns ip6 command to perform the above commands on IPv6 addresses.

8. Configure ECMP on the external switch. The following sample configurations are provided for the Cisco® Nexus 7000 C7010 Release 5.2(1) switch. Similar configurations must be performed on other switches.

```
//For OSPF (IPv4 addresses) Global config: Configure terminal feature ospf      Interf
```

Use Case: ECMP with BGP Routing

To configure ECMP with BGP routing protocol, perform the following steps:

1. Log on to the cluster IP address.
2. Enable BGP routing protocol.

```
> enable ns feature bgp
```

3. Add VLAN and bind the required interfaces.

```
> add vlan 985
> bind vlan 985 -ifnum 0/0/1 1/0/1
```

4. Add the spotted IP address and bind them to the VLAN.

```
> add ns ip 10.100.26.14 255.255.255.0 -ownerNode 1 -dynamicRouting ENABLED
> add ns ip 10.100.26.15 255.255.255.0 -ownerNode 2 -dynamicRouting ENABLED
> bind vlan 985 -ipAddress 10.100.26.10 255.255.255.0
```

5. Configure BGP routing protocol on ZebOS using vtysh shell.

```
> vtysh
conf t
router bgp 65535
neighbor 10.100.26.1 remote-as 65535
```

6. Configure BGP on the external switch. The following sample configurations are provided for the Cisco® Nexus 7000 C7010 Release 5.2(1) switch. Similar configurations must be performed on other switches.

```
router bgp 65535
no synchronization
bgp log-neighbor-changes
neighbor 10.100.26.14 remote-as 65535
neighbor 10.100.26.15 remote-as 65535
  no auto-summary
dont-capability-negotiate
dont-capability-negotiate
no dynamic-capability
```

Using Cluster Link Aggregation

Cluster link aggregation, as the name suggests, is a group of interfaces of cluster nodes. It is an extension of NetScaler link aggregation. The only difference is that, while link aggregation requires the interfaces to be from the same device, in cluster link aggregation, the interfaces are from different nodes of the cluster.


Note: Cluster link aggregation is supported for a cluster of hardware (MPX) appliances or a cluster of virtual (VPX) appliances that are deployed on ESX and KVM hypervisors, with the restriction that the interfaces are not shared with other virtual machines.

For more information about link aggregation, see "[Configuring Link Aggregation](#)".

Cluster link aggregation can be either static or dynamic.

For example, consider a three-node cluster where all three nodes are connected to the upstream switch. A cluster LA channel (CLA/1) is formed by binding interfaces 0/1/2, 1/1/2, and 2/1/2.

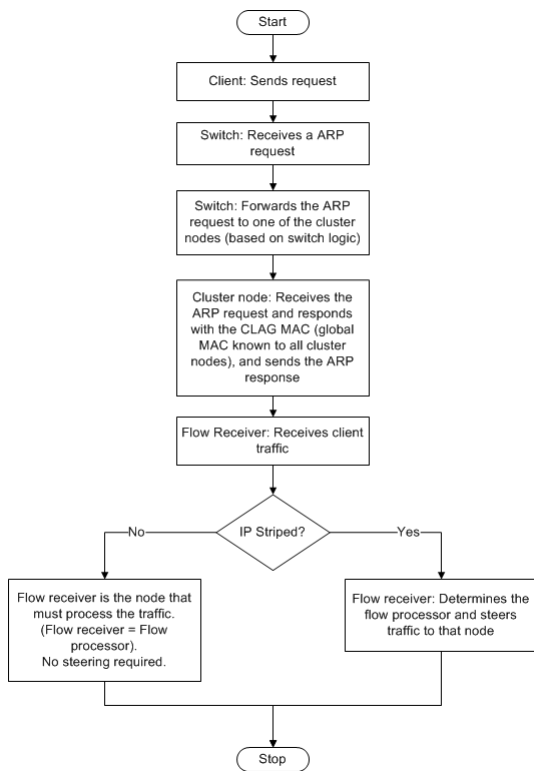
Figure 1. Cluster Link Aggregation topology



A cluster LA channel has the following attributes:

- Each channel has a unique MAC agreed upon by cluster nodes.
- The channel can bind both local and remote nodes' interfaces.
- A maximum of four cluster LA channels are supported in a cluster.
- Backplane interfaces cannot be part of a cluster LA channel.
- When an interface is bound to a cluster LA channel, the channel parameters have precedence over the network interface parameters. A network interface can be bound to one channel only.

Figure 2. Traffic distribution flow using cluster LA



Static Cluster Link Aggregation

You must configure a static cluster LA channel on the cluster IP address and on the external connecting device. If possible, configure the upstream switch to distribute traffic based on IP address or port instead of MAC address.

For more information about configuring a static LA channel, see ["Configuring Link Aggregation Manually"](#).

To configure a static cluster LA channel by using the command line interface

1. Log on to the cluster IP address.

Note: Make sure that you configure the cluster LA channel on the cluster IP address before configuring link aggregation on the external switch. Otherwise, the switch will forward traffic to the cluster even though the cluster LA channel is not configured. This can lead to loss of traffic.

2. Create a cluster LA channel.

```
add channel <id> -speed <speed>
```

Example

```
> add channel CLA/1 -speed 1000
```

Note: You must not specify the speed as AUTO. Rather, you must explicitly specify the speed as 10, 100, 1000, or 10000. Only interfaces that have the speed matching the <speed> attribute in the cluster LA channel are added to the active distribution list.

3. Bind the required interfaces to the cluster LA channel. Make sure that the interfaces are not used for the cluster backplane.

```
bind channel <id> <ifnum>
```

Example

```
> bind channel CLA/1 0/1/2 1/1/2 2/1/2
```

4. Verify the configurations.

```
show channel <id>
```

Example

```
> show channel CLA/1
```

Note: You can bind the cluster LA channel to a VLAN by using the bind vlan command. The interfaces of the channel are automatically bound to the VLAN.

5. Configure static LA on the external switch. The following sample configurations are provided for the Cisco® Nexus 7000 C7010 Release 5.2(1). Similar configurations must be performed on other switches.

Global config:

Configure terminal

Interface level config:

```
interface Ethernet2/47
  switchport
  switchport access vlan 10
  channel-group 7 mode on
  no shutdown
```

```
interface Ethernet2/48
  switchport
  switchport access vlan 10
  channel-group 7 mode on
  no shutdown
```

Dynamic Cluster Link Aggregation

Dynamic cluster LA channel uses Link Aggregation Control Protocol (LACP). For more information about configuring a dynamic LA channel, see ["Configuring Link Aggregation by Using the Link Aggregation Control Protocol"](#).

You must perform similar configurations on the cluster IP address and on the external connecting device. If possible, configure the upstream switch to distribute traffic based on IP address or port instead of MAC address.

Points to remember:

- Enable LACP (by specifying the LACP mode as either ACTIVE or PASSIVE).
Note: Make sure the LACP mode is not set as PASSIVE on both the NetScaler cluster and the external connecting device.
- Specify the same LACP key on each interface that you want to be the part of the channel. For creating a cluster LA channel, the LACP key can have a value from 5 through 8. For example, if you set the LACP key on interfaces 0/1/2, 1/1/2, and 2/1/2 to 5, CLA/1 is created. The interfaces 0/1/2, 1/1/2, and 2/1/2 are automatically bound to CLA/1. Similarly, if you set the LACP key to 6, CLA/2 channel is created.
- Specify the LAG type as Cluster.

To configure a dynamic cluster LA channel by using the command line interface

On the cluster IP address, for each interface that you want to add to the cluster LA channel, type:

```
set interface <id> -lacpMode <lacpMode> -lacpKey <positive_integer> -lagType CLUSTER
```

Example: To configure a cluster LA channel CLA/1 of 3 interfaces.

```
> set interface 0/1/2 -lacpMode active -lacpKey 5 -lagType Cluster
> set interface 1/1/2 -lacpMode active -lacpKey 5 -lagType Cluster
> set interface 2/1/2 -lacpMode active -lacpKey 5 -lagType Cluster
```

Note: Optionally, you can enable [Link Redundancy in a Cluster with LACP](#).

Similarly, configure dynamic LA on the external switch. The following sample configurations are provided for the Cisco® Nexus 7000 C7010 Release 5.2(1). Similar configurations must be performed on other switches.

Global config:

Configure terminal

feature lacp

Interface level config:


```
interface Ethernet2/47
  switchport
  switchport access vlan 10
  channel-group 7 mode active
  no shutdown
```

```
interface Ethernet2/48
  switchport
  switchport access vlan 10
  channel-group 7 mode active
  no shutdown
```


Link Redundancy in a Cluster with LACP

A NetScaler cluster provides link redundancy for LACP to ensure that all nodes have the same partner key.

To understand the need for link redundancy, let us consider the example of the following cluster setup along with the accompanying cases (with attention to case 3):



In this setup, interfaces I1, I2, I3 and I4 are bound to LACP channel with KEY 5. On the partner side, I1 and I2 are connected to Switch 1 to form a single LA channel with KEY 1. Similarly, I3 and I4 are connected to Switch 2 to form a single LA channel with KEY 2.

Now let us consider the following cases to understand the need for link redundancy:

Case 1: Switch 1 is up and Switch 2 is down

In this case, cluster LA on both the nodes would stop receiving LACPDUs from Key2 and would start receiving LACPDUs from Key1. On both the nodes, cluster LA is connected to KEY 1 and I1 and I2 will be UP and channel on both the nodes would be UP.

Case 2: Switch1 goes down and Switch2 becomes UP

In this case, cluster LA on both the nodes would stop receiving LACPDUs from Key1 and would start receiving LACPDUs from Key2. On both the nodes, cluster LA is connected to Key2 and I3 and I4 will be UP and channel on both the nodes would be UP.

Case 3: Both Switch1 and Switch2 are UP

In this case, it is possible that cluster LA on node1 chooses Key1 as its partner and cluster LA on node2 chooses Key2 as its partner. This means that I1 on node1 and I4 on node2 are receiving traffic which is undesirable. This can happen because the LACP state machine is node-level and chooses its partners on first-come first-serve basis. To solve these concerns, link redundancy of dynamic cluster LA is supported. To configure link redundancy on a channel or interface, you must enable it as follows:

```
set channel CLA/1 -linkRedundancy ON
```

Using Linksets

Linksets must be used when some cluster nodes are not physically connected to the external network. In such a cluster topology, the unconnected cluster nodes use the interfaces specified in the linkset to communicate with the external network through the cluster backplane. Linksets are typically used in scenarios when the connecting devices have insufficient ports to connect the cluster nodes.


Note: Linksets are a mandatory configuration in the following scenarios:

- For deployments that require MAC-Based Forwarding (MBF).
- To improve manageability of ACL and L2 policies involving interfaces. You must define a linkset of the interfaces and add ACL and L2 policies based on linksets.

Linksets must be configured only through the cluster IP address.


For example, consider a three node cluster where the upstream switch has only two ports available. Using linksets, you can connect two nodes to the switch and leave the third node unconnected. In the following figure, a linkset (LS/1) is formed by binding the interfaces 0/1/2 and 1/1/2. NS2 is the unconnected node of the cluster.

Figure 1. Linksets topology



The linkset informs NS2 that it can use interfaces 0/1/2 and 1/1/2 to communicate with the network devices. All traffic to and from NS2 is now routed through interfaces 0/1/2 or 1/1/2.

Figure 2. Traffic distribution flow using linksets



To configure a linkset by using the command line interface

1. Log on to the cluster IP address.
2. Create a linkset.

```
add linkset <id>
```

Example

```
> add linkset LS/1
```

3. Bind the required interfaces to the linkset. Make sure the interfaces are not used for the cluster backplane.

```
bind linkset <id> -ifnum <interface_name> ...
```

Example

```
> bind linkset LS/1 -ifnum 0/1/2 1/1/2
```

4. Verify the linkset configurations.

```
show linkset <id>
```

Example

```
> show linkset LS/1
```

Note: You can bind the linkset to a VLAN by using the bind vlan command. The interfaces of the linkset are automatically bound to the VLAN.

To configure a linkset by using the configuration utility

1. Log on to the cluster IP address.
2. Navigate to System > Network > Linksets.
3. In the details pane, click Add.
4. In the Create Linkset dialog box:
 - a. Specify the name of the linkset by setting the Linkset parameter.
 - b. Specify the Interfaces to be added to the linkset and click Add. Repeat this step for each interface you want to add to the linkset.
5. Click Create, and then click Close.

Managing the NetScaler Cluster

After you have created a cluster and configured the required traffic distribution mechanism, the cluster is able to serve traffic. During the lifetime of the cluster, you can perform cluster tasks such as configuring nodegroups, disabling nodes of a cluster, discovering NetScaler appliances, viewing statistics, synchronizing cluster configurations, cluster files, and the time across the nodes, and upgrading or downgrading the software of cluster nodes.

Nodegroups for Spotted and Partially-Striped Configurations

By virtue of the default cluster behavior, all configurations performed on the cluster IP address are available on all nodes of the cluster. However, there might be cases where you need some configurations to be available only on specific cluster nodes.

You can achieve this requirement by defining a nodegroup that includes the specific cluster nodes, and then binding the configuration to that nodegroup. This ensures that the configuration is active only on those cluster nodes. These configurations are called partially-striped or spotted (if active only on a single node). For more information, see [Striped, Partially Striped, and Spotted Configurations](#).

For example, consider a cluster with three nodes. You create a nodegroup NG0 that includes node n1 and another nodegroup NG1 that includes n2 and n3. Bind load balancing virtual servers .77 to NG0 and load balancing virtual server .69 to NG1.

This means that virtual server .77 will be active only on n1 and consequently only n1 will receive traffic that is directed to .77. Similarly, virtual server .69 will be active only on nodes n2 and n3 and consequently only n2 and n3 will receive traffic that is directed to .69.

Figure 1. NetScaler cluster with nodegroups configured for spotted and partial-striped configurations



The entities or configurations that you can bind to a nodegroup are:

- Load balancing, content switching, cache redirection, authentication (AAA) virtual servers
Note: FTP load balancing virtual servers cannot be bound to nodegroups.
- VPN virtual server (Supported from NetScaler 10.5 Build 50.10 onwards)
- Global Server Load Balancing (GSLB) sites and other GSLB entities (Supported from NetScaler 10.5 Build 52.11 onwards)
- Limit identifiers and stream identifiers

Behavior of Nodegroups

Due to the interoperability of nodegroups with different NetScaler features and entities, there are some behavioral aspects to be noted. Nodes in a nodegroup can also be backed up. Read on for more information.

General behavior of a cluster nodegroup

- A nodegroup that has entities bound to it cannot be removed.
- A cluster node that belongs to a nodegroup with entities bound to it, cannot be removed.
- A cluster instance that has nodegroups with entities bound to it, cannot be removed.
- You cannot add an entity that has a dependency on another entity that is not part of the nodegroup. If you need to do so, first remove the dependency. Then, add both the entities to the nodegroup and reassociate the entities.

Examples:

Assume you have a virtual server, VS1, whose backup is virtual server VS2. To add VS1 to a nodegroup, first make sure that VS2 is removed as the backup server of VS1. Then, bind each server individually to the nodegroup, and then configure VS2 as the backup for VS1.

Assume you have a content switching virtual server, CSVS1, whose target load balancing virtual server is LBVS1. To add CSVS1 to a nodegroup, first remove LBVS1 as the target. Then, bind each server individually to the nodegroup, and then configure LBVS1 as the target.

Assume you have a load balancing virtual server, LBVS1, that has a policy which invokes another load balancing virtual server, LBVS2. To add either one of the virtual servers, first remove the association.

Then, bind each server individually to the nodegroup, and then reassociate the virtual servers.

- You cannot bind an entity to a nodegroup that has no nodes and that has the `strict` option enabled. Consequently, you cannot unbind the last node of a nodegroup that has entities bound to it and that has the `strict` option enabled
- The `strict` option cannot be modified for a nodegroup that has no nodes but has entities bound to it.

Backing up Nodes in a Nodegroup

By default, a nodegroup is designed to provide back up nodes for members of a nodegroup. If a nodegroup member goes down, a cluster node that is not a member of the nodegroup dynamically replaces the failed node. This node is called the replacement node.

Note: For a single-member nodegroup, a backup node is automatically preselected when an entity is bound to the nodegroup

When the original member of the nodegroup comes up, the replacement node, by default, is replaced by the original member node.

From NetScaler 10.5 Build 50.10 onwards, however, the NetScaler allows you to change this replacement behavior. When you enable the sticky option, the replacement node is retained even after the original member node comes up. The original node takes over only when the replacement node goes down.

You can also disable the backup functionality. To do this, you must enable the strict option. In this scenario, when a nodegroup member goes down, no other cluster node is picked up as a backup node. The original node continues being part of the nodegroup when it comes up. This option ensures that entities bound to a nodegroup are active only on nodegroup members.

Note: The strict and sticky option can be set only when creating a nodegroup.

Configuring Nodegroups for Spotted and Partially-Striped Configurations

To configure a nodegroup for spotted and partially-striped configurations you must first create a nodegroup and then bind the required nodes to the nodegroup. You must then associate the required entities to that nodegroup. The entities that are bound to the nodegroup will be:

- Spotted - If bound to a nodegroup that has a single node.
- Partially striped - If bound to a nodegroup that has more than one node.

Some points to remember:

- GSLB is supported on a cluster only when GSLB sites are bound to nodegroups that have a single cluster node. For more information, see [Setting Up GSLB in a Cluster](#).
- NetScaler Gateway is supported on a cluster only when the VPN virtual servers are bound to nodegroups that have a single cluster node. The sticky option must be enabled on the nodegroup.
- Application firewall is supported on a cluster only when application firewall profiles are associated with virtual servers that are bound to nodegroups that have a single cluster node. You are not allowed to do the following:
 - Bind application firewall profiles to striped or partially striped virtual servers.
 - Bind the policy to a global bind point or to user-defined policy labels.
 - Unbind, from a nodegroup, a virtual server that has application firewall profiles.

Check [NetScaler Features Supported in a Cluster](#) to see the NetScaler versions from which GSLB, NetScaler Gateway, and application firewall are supported in a cluster.

To configure a nodegroup by using the command line interface

1. Log on to the cluster IP address.
2. Create a nodegroup. Type:

```
add cluster nodegroup <name> -strict (YES | NO)
```

Example

```
> add cluster nodegroup NG0 -strict YES
```

3. Bind the required nodes to the nodegroup. Type the following command for each member of the nodegroup:

```
bind cluster nodegroup <name> -node <nodeId>
```

Example: To bind nodes with IDs 1, 5, and 6.

```
> bind cluster nodegroup NG0 -node 1
> bind cluster nodegroup NG0 -node 5
> bind cluster nodegroup NG0 -node 6
```

4. Bind the entity to the nodegroup. Type the following command once for every entity that you want to bind:

```
bind cluster nodegroup <name> (-vServer <string> | -identifierName <string> | -gsLBsite <string> -service <string>)
```

Note: The `gsLBsite` and `service` parameters are available from NetScaler 10.5 onwards.

Example: To bind virtual servers VS1 and VS2 and rate limit identifier named identifier1.

```
> bind cluster nodegroup NG0 -vServer VS1
> bind cluster nodegroup NG0 -vServer VS2
> bind cluster nodegroup NG0 -identifierName identifier1
```

5. Verify the configurations by viewing the details of the nodegroup. Type:

```
show cluster nodegroup <name>
```

Example

```
> show cluster nodegroup NG0
```

To configure a nodegroup by using the configuration utility

1. Log on to the cluster IP address.
2. Navigate to System > Cluster > Node Groups.

3. In the details pane, click Add.
4. In the Create Node Group dialog box, configure the nodegroup:
 - a. Under Cluster Nodes, click the Add button.
 - The Available list displays the nodes that you can bind to the nodegroup and the Configured list displays the nodes that are bound to the nodegroup.
 - Click the + sign in the Available list to bind the node. Similarly, click the - sign in the Configured list to unbind the node.
 - b. Under Virtual Servers, select the tab corresponding to the type of virtual server that you want to bind to the nodegroup. Click the Add button.
 - The Available list displays the virtual servers that you can bind to the nodegroup and the Configured list displays the virtual servers that are bound to the nodegroup.
 - Click the + sign in the Available list to bind the virtual server. Similarly, click the - sign in the Configured list to unbind the virtual server.

Configuring Redundancy for Nodegroups

Note: Supported from NetScaler 10.5 Build 52.1115.e onwards.

Nodegroups can be configured such that when one nodegroup goes down, another nodegroup can take over and process traffic. For example, when a nodegroup NG1 goes down, NG2 takes over.

Note: This functionality can be used to configure datacenter redundancy where each nodegroup is configured as a datacenter

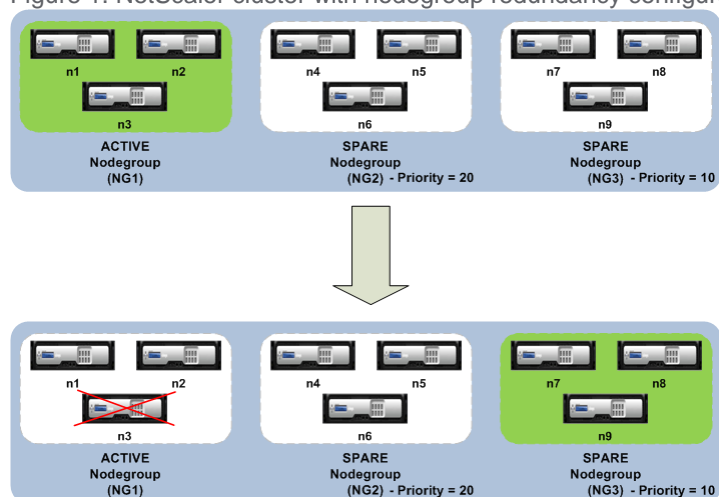
To achieve this use case, cluster nodes must be logically grouped into nodegroups, where some nodegroups must be configured as ACTIVE and others as SPARE. The active nodegroup with the highest priority (that is, the lowest priority number) is made operationally active and therefore serves traffic. When a node from this operationally active nodegroup goes down, the node count of this nodegroup is compared with the node count of the other active nodegroups in order of their priority. If a nodegroup has a higher or equal node count, that nodegroup is made operationally active. Else, the spare nodegroups are checked.

Note:

- Only one state-specific nodegroup can be active at a given point in time.
- A cluster node inherits the state of the nodegroup. So, if a node with "SPARE" state is added to nodegroup with state as "ACTIVE", the node automatically behaves as an active node.
- The preemption parameter that is defined for the cluster instance decides whether the initial active nodegroup will take control when the it comes up again.

The following figure shows a nodegroup setup that has nodegroup redundancy defined. NG1 is initially the active nodegroup. When it loses one of the nodes, the spare nodegroup (NG3) with the highest priority starts serving traffic.

Figure 1. NetScaler cluster with nodegroup redundancy configured



Configuring redundancy for nodegroups

1. Log on to the cluster IP address.
2. Create the active nodegroup and bind the required cluster nodes.

```
add cluster nodegroup NG1 -state ACTIVE
```

```
bind cluster nodegroup NG1 -node n1
```

```
bind cluster nodegroup NG1 -node n2
```

```
bind cluster nodegroup NG1 -node n3
```

3. Create the spare nodegroup and bind the requisite nodes.

```
add cluster nodegroup NG2 -state SPARE -priority 20
```

```
bind cluster nodegroup NG2 -node n4
```

```
bind cluster nodegroup NG2 -node n5
```

```
bind cluster nodegroup NG2 -node n6
```

4. Create another spare nodegroup and bind the requisite nodes.

```
add cluster nodegroup NG3 -state SPARE -priority 10
```

```
bind cluster nodegroup NG3 -node n7
```

```
bind cluster nodegroup NG3 -node n8
```

```
bind cluster nodegroup NG3 -node n9
```

Synchronizing Cluster Configurations

NetScaler configurations that are available on the configuration coordinator are synchronized to the other nodes of the cluster when:

- A node joins the cluster
- A node rejoins the cluster
- A new command is executed through the cluster IP address.

Additionally, you can forcefully synchronize the configurations that are available on the configuration coordinator (full synchronization) to a specific cluster node. Make sure you synchronize one cluster node at a time, otherwise the cluster can get affected.

To synchronize cluster configurations by using the command line interface

At the command prompt of the appliance on which you want to synchronize the configurations, type:

```
force cluster sync
```

To synchronize cluster configurations by using the configuration utility

1. Log on to the appliance on which you want to synchronize the configurations.
2. Navigate to System > Cluster.
3. In the details pane, under Utilities, click Force cluster sync.
4. Click OK.

Synchronizing Time Across Cluster Nodes

The cluster uses Precision Time Protocol (PTP) to synchronize the time across cluster nodes. PTP uses multicast packets to synchronize the time. If there are some issues in time synchronization, you must disable PTP and configure Network Time Protocol (NTP) on the cluster.

To enable/disable PTP by using the command line interface

At the command prompt of the cluster IP address, type:

```
set ptp -state disable
```

To enable/disable PTP by using the configuration utility

1. Log on to the cluster IP address.
2. Navigate to System > Cluster.
3. In the details pane, under Utilities, click Configure PTP Settings.
4. In the Enable/Disable PTP dialog box, select whether you want to enable or disable PTP.
5. Click OK.

Synchronizing Cluster Files

The files available on the configuration coordinator are called cluster files. These files are automatically synchronized on the other cluster nodes when the node is added to the cluster and periodically, during the lifetime of the cluster. Additionally, you can manually synchronize the cluster files.

The directories and files from the configuration coordinator that are synchronized are:

- /nsconfig/ssl/
- /var/netcaler/ssl/
- /var/vpn/bookmark/
- /nsconfig/dns/
- /nsconfig/htmlinjection/
- /netcaler/htmlinjection/ens/
- /nsconfig/monitors/
- /nsconfig/nstemplates/
- /nsconfig/ssh/
- /nsconfig/rc.netcaler
- /nsconfig/resolv.conf
- /nsconfig/inetd.conf
- /nsconfig/syslog.conf
- /nsconfig/snmpd.conf
- /nsconfig/ntp.conf
- /nsconfig/httpd.conf
- /nsconfig/sshd_config
- /nsconfig/hosts
- /nsconfig/enckey
- /var/nslw.bin/etc/krb5.conf
- /var/nslw.bin/etc/krb5.keytab
- /var/lib/likewise/db/
- /var/download/
- /var/wi/tomcat/webapps/
- /var/wi/tomcat/conf/Catalina/localhost/
- /var/wi/java_home/lib/security/cacerts
- /var/wi/java_home/jre/lib/security/cacerts
- /nsconfig/license/
- /nsconfig/rc.conf

To synchronize cluster files by using the command line interface

At the command prompt of the cluster IP address, type:

```
sync cluster files <mode>
```

To synchronize cluster files by using the configuration utility

1. Log on to the cluster IP address.
2. Navigate to System > Cluster.
3. In the details pane, under Utilities, click Synchronize cluster files.
4. In the Synchronize cluster files dialog box, select the files to be synchronized in the Mode drop-down box.
5. Click OK.

Viewing the Statistics of a Cluster

You can view the statistics of a cluster instance and cluster nodes to evaluate the performance or to troubleshoot the operation of the cluster.

To view the statistics of a cluster instance by using the command line interface

At the command prompt of the cluster IP address, type:

```
stat cluster instance <clld>
```

To view the statistics of a cluster node by using the command line interface

At the command prompt of the cluster IP address, type:

```
stat cluster node <nodeid>
```

Note: When executed from the cluster IP address, this command displays the cluster level statistics. However, when executed from the NSIP address of a cluster node, the command displays node level statistics.

To view the statistics of a cluster instance by using the configuration utility

1. Log on to the cluster IP address.
2. Navigate to System > Cluster.
3. In the details pane, in the center of the page, click Statistics.

To view the statistics of a cluster node by using the configuration utility

1. Log on to the cluster IP address.
2. Navigate to System > Cluster > Nodes.
3. In the details pane, select a node and click Statistics to view the statistics of the node. To view the statistics of all the nodes, click Statistics without selecting a specific node.

Discovering NetScaler Appliances

You can discover NetScaler appliances present in the same subnet as the NSIP address of the configuration coordinator. The discovered appliances can then be added to the cluster.

Note: This operation is available only through the configuration utility.

To discover appliances by using the configuration utility

1. Log on to the cluster IP address.
2. Navigate to System > Cluster > Nodes.
3. In the details pane, at the bottom of the page, click Discover NetScalers.
4. In the Discover NetScalers dialog box, set the following parameters:
 - IP address range - Specify the range of IP addresses within which you want to discover appliances. For example, you can search for all NSIP addresses between 10.102.29.4 to 10.102.29.15 by specifying this option as 10.102.29.4 - 15.
 - Backplane interface - Specify the interfaces to be used as the backplane interface. This is an optional parameter. If you do not specify this parameter, you must update it after the node is added to the cluster.
5. Click OK.
6. Select the appliances that you want to add to the cluster.
7. Click OK.

Disabling a Cluster Node

You can temporarily remove a node from a cluster by disabling the cluster instance on that node. A disabled node is not synchronized with the cluster configurations. When the node is enabled again, the cluster configurations are automatically synchronized on it. For more information, see [Cluster Synchronization](#).

A disabled node cannot serve traffic and all existing connections on this node are terminated.

Note: If the configurations of a disabled non-configuration coordinator node are modified (through the NSIP address of the node), the configurations are not automatically synchronized on that node. You must manually synchronize the configurations as described in [Synchronizing Cluster Configurations](#).

To disable a cluster node by using the command line interface

At the command prompt of the node that you want to disable, type:

```
disable cluster instance <cld>
```

Note: To disable the cluster, run the disable cluster instance command on the cluster IP address.

To disable a cluster node by using the configuration utility

1. On the node that you want to disable, navigate to System > Cluster, and click Manage Cluster.
 2. In the Configure cluster instance dialog box, unselect the Enable cluster instance check box.
- Note: To disable the cluster instance on all the nodes, perform the above procedure on the cluster IP address.

Removing a Cluster Node

When a node is removed from the cluster, the cluster configurations are cleared from the node (by internally executing the `clear ns config -extended` command). The SNIP addresses and all VLAN configurations (except the default VLAN and NSVLAN) are also cleared from the appliance.

Note:

- If the deleted node was the cluster configuration coordinator, another node is automatically selected as the cluster configuration coordinator, and the cluster IP address is assigned to that node. All the current cluster IP address sessions will be invalid and you will have to start a new session.
- To delete the whole cluster, you must remove each node individually. When you remove the last node, the cluster IP address(es) are deleted.
- When an active node is removed, the traffic serving capability of the cluster is reduced by one node. Existing connections on this node are terminated.

To remove a cluster node by using the command line interface

For NetScaler 10.1 and later versions

Log on to the cluster IP address and at the command prompt, type:

```
rm cluster node <nodeId>
```

Note: If the cluster IP address is unreachable from the node, execute the `rm cluster instance` command on the NSIP address of that node itself.

For NetScaler 10

1. Log on to the node that you want to remove from the cluster and remove the reference to the cluster instance.

```
rm cluster instance <clId>
```

```
save ns config
```

2. Log on to the cluster IP address and remove the node from which you removed the cluster instance.

```
rm cluster node <nodeId>
```

```
save ns config
```

Make sure you do not run the `rm cluster node` command from the local node as this results in inconsistent configurations between the configuration coordinator and the node.

To remove a cluster node by using the configuration utility

On the cluster IP address, navigate to **System > Cluster > Nodes**, select the node you want to remove and click **Remove**.

Removing a Node from a Cluster Deployed Using Cluster Link Aggregation

To remove a node from a cluster that uses cluster link aggregation as the traffic distribution mechanism, you must make sure that the node is made passive so that it does not receive any traffic and then, on the upstream switch, remove the corresponding interface from the channel.

For detailed information on cluster link aggregation, see [Using Cluster Link Aggregation](#).

To remove a node from a cluster that uses cluster link aggregation as the traffic distribution mechanism

1. Log on to the cluster IP address.
2. Set the state of the cluster node that you want to remove to PASSIVE.

`set cluster node <nodeId> -state PASSIVE`
3. On the upstream switch, remove the corresponding interface from the channel by using switch-specific commands.
Note: You do not have to manually remove the nodes interface on the cluster link aggregation channel. It is automatically removed when the node is deleted in the next step.
4. Remove the node from the cluster.

```
rm cluster node <nodeId>
```

Cluster Setup and Usage Scenarios

This section aims at explaining some scenarios in which the NetScaler cluster can be setup and also how it can be configured for different features and network topologies. These are just some scenarios that we have documented. Provide feedback if you want some other scenarios to be included.

- Creating a Two-Node Cluster
- Migrating an HA Setup to a Cluster Setup
- Migrating an HA Setup to a Cluster Setup without Downtime
- Setting Up GSLB in a Cluster
- Using Cache Redirection in a Cluster
- Using L2 Mode in a Cluster Setup
- Using Cluster LA Channel with Linksets
- Backplane on LA Channel
- Common Interface for Client and Server and Dedicated Interfaces for Backplane
- Common Switch for Client, Server, and Backplane
- Common Switch for Client and Server and Dedicated Switch for Backplane
- Different Switch for Every Node
- Sample Cluster Configurations

Creating a Two-Node Cluster

A two-node cluster is an exception to the rule that a cluster is functional only when a minimum of $(n/2 + 1)$ nodes, where n is the number of cluster nodes, are able to serve traffic. If that formula were applied to a two-node cluster, the cluster would fail if one node went down ($n/2 + 1 = 2$).

A two-node cluster is functional even if only one node is able to serve traffic.

Creating a two node cluster is the same as creating any other cluster. You must add one node as the configuration coordinator and the other node as the other cluster node.

Note: Incremental configuration synchronization is not supported in a two-node cluster. Only full synchronization is supported

Migrating an HA Setup to a Cluster Setup

An existing high availability (HA) setup can be migrated to a cluster setup by first removing the appliances from the HA setup and then creating the NetScaler cluster. This approach will result in a downtime for the application.

Consider an HA setup with appliances NS0 (10.102.97.131) and NS1 (10.102.97.132). NS0 is the primary and NS1 is the secondary appliance of the HA setup.

To convert a HA setup to cluster setup by using the NetScaler command line

1. Log on to each HA node and remove it from the HA setup.

```
rm HA node <id>
```

Example

```
rm HA node 1
```

2. Go to the shell on one of the HA nodes and copy the ns.conf file to another .conf file (for example, ns_backup.conf).
3. On both the nodes, identify the network interfaces to be used for the cluster backplane. Make sure to configure the backplane switch appropriately.
4. Create the cluster on one of the appliances (for example, 10.102.97.131).

```
//On the NSIP address of the first appliance
add cluster instance 1
add cluster node 0 10.102.97.131 -state ACTIVE -backplane 0/1/1
add ns ip 10.102.97.133 255.255.255.255 -type CLIP
enable cluster instance 1
save ns config
reboot -warm
```

5. Add the other appliance to the cluster.

```
//On the cluster IP address
add cluster node 1 10.102.97.132 -state ACTIVE -backplane 1/1/1

//On the NSIP address of the appliance
join cluster -clip 10.102.97.133 -password nsroot
save ns config
reboot -warm
```

6. After the two nodes are up and active, log on to the cluster IP address and modify the backed-up configuration file as follows:
 - a. Remove the features that are not supported on a cluster. For the list of unsupported features, see [NetScaler Features Supported by a Cluster](#). This is an optional step. If you do not perform this step, the execution of unsupported commands will fail.
 - b. Remove the configurations that have interfaces, or update the interface names from the c/u convention to the n/c/u convention.

Example

```
add vlan 10 -ifnum 0/1

should be changed to

add vlan 10 -ifnum 0/0/1 1/0/1
```

- c. The backup configuration file can have SNIP addresses or MIP addresses. These addresses are striped on all the cluster nodes. It is recommended that you add spotted IP addresses for each node.

Example

```
add ns ip 1.1.1.1 255.255.255.0 -ownerNode 0
add ns ip 1.1.1.2 255.255.255.0 -ownerNode 1
```

- d. Update the hostname to specify the owner node.

Example

```
set ns hostname ns0 -ownerNode 0
set ns hostname ns1 -ownerNode 1
```

- e. Change all other relevant networking configuration that depend on spotted IPs. For example, L3 VLAN, RNAT configuration which uses SNIPs as NATIP, INAT rules that refers to SNIPs/MIPs).
- 7. Apply configurations from the backup configuration file to the configuration coordinator through the cluster IP address.

```
batch -fileName <input_filename>
```

Example

```
batch -f ns_backup.conf
```

- 8. Configure appropriate client traffic distribution mechanism (ECMP, cluster LA or linksets).
- 9. Save the configuration.

```
save ns config
```

The appliances of the HA setup are migrated to a cluster setup.

Migrating an HA Setup to a Cluster Setup without Downtime

An existing high availability (HA) setup can be migrated to a cluster setup by first removing the secondary appliance from the HA setup and using that appliance to create a single-node cluster. Then, after the cluster becomes operational and serves traffic, the primary appliance of the HA setup is added to the cluster. This approach will not result in a downtime for the application.

Consider an HA setup with appliances NS0 (10.102.97.131) and NS1 (10.102.97.132). NS0 is the primary and NS1 is the secondary appliance of the HA setup.

To convert a HA setup to cluster setup (without downtime) by using the NetScaler command line

1. Go to the shell on one of the HA nodes and copy the ns.conf file to another .conf file (for example, ns_backup.conf).
Note: Make sure HA pair is stable with respect to configurations.
2. Log on to the secondary appliance NS1 and clear all the configurations. This removes the secondary appliance from the HA setup and makes it a standalone appliance.

```
clear ns config full
```

Note:

- The configurations are cleared to make sure that NS1 does not start owning the VIPs once it becomes a standalone appliance.
 - At this stage, NS0 is still active and continues to serve traffic.
3. Create a cluster on appliance NS1 and configure it as a PASSIVE node.

```
//On the NSIP address of node NS1
add cluster instance 1
add cluster node 0 10.102.97.131 -state PASSIVE -backplane 0/1/1
add ns ip 10.102.97.133 255.255.255.255 -type CLIP
enable cluster instance 1
save ns config
reboot -warm
```

4. Modify the backed-up configuration file.
 - a. Remove the features that are not supported on a cluster. For the list of unsupported features, see [NetScaler Features Supported by a Cluster](#). This is an optional step. If you do not perform this step, the execution of unsupported commands will fail.
 - b. Remove the configurations that have interfaces, or update the interface names from the c/u convention to the n/c/u convention.

Example

```
add vlan 10 -ifnum 0/1
```

should be changed to

```
add vlan 10 -ifnum 0/0/1 1/0/1
```

- c. The backup configuration file can have SNIP addresses or MIP addresses. These addresses are striped on all the cluster nodes. It is recommended that you add spotted IP addresses for each node.

Example

```
add ns ip 1.1.1.1 255.255.255.0 -ownerNode 0
add ns ip 1.1.1.2 255.255.255.0 -ownerNode 1
```

- d. Update the hostname to specify the owner node.

Example

```
set ns hostname ns0 -ownerNode 0
set ns hostname ns1 -ownerNode 1
```

- e. Change all other relevant networking configuration that depend on spotted IPs. For example, L3 VLAN, RNAT configuration which uses SNIPs as NATIP, INAT rules that refers to SNIPs/MIPs).
5. On the cluster, do the following:
 - a. Make the topological changes to the cluster by connecting the cluster backplane, the cluster link aggregation channel, and so on.

- b. Apply configurations from the backup configuration file to the configuration coordinator through the cluster IP address.

```
batch -f ns_backup.conf
```

- c. Configure external traffic distribution mechanisms like ECMP or cluster link aggregation.
6. Switch-over the traffic from the HA setup to the single-node cluster setup.
 - a. Disable all interfaces on the primary appliance NS0.

```
disable interface <interface id>
```

- b. Configure the cluster node as an ACTIVE node.

```
set cluster node 0 -state ACTIVE
```

Note: There can be a small amount (in the order of seconds) of downtime between disabling the interfaces and making the cluster node active.

7. On the primary appliance NS0, do the following:
 - a. Clear all the configurations.

```
clear ns config full
```

- b. Enable all the interfaces.

```
enable interface <interface id>
```

- c. Add the appliance to the cluster.

```
//On the cluster IP address (in this sample, 10.102.97.133)
add cluster node 1 10.102.97.132 -state PASSIVE -backplane 1/1/1
```

```
//On the NSIP address of the appliance
join cluster -clip 10.102.97.133 -password nsroot
save ns config
reboot -warm
```

- d. Perform the required topological and configuration changes.
- e. Configure NS0 as an ACTIVE node.

```
set cluster node 1 -state ACTIVE
```

The appliances of the HA setup are migrated to a cluster setup without any downtime for the application.

Setting Up GSLB in a Cluster

Note: Supported from NetScaler 10.5 Build 52.11 onwards.

To set up GSLB in a cluster you must bind the different GSLB entities to a node group. The node group must have a single member node.

Note:

- The parent-child topology of GSLB is not supported in a cluster.
- If you have configured the static proximity GSLB method, make sure that the static proximity database is present on all the cluster nodes. This happens by default if the database file is available at the default location. However, if the database file is maintained in a directory other than /var/netscaler/locdb/, you must manually synch the file to all the cluster nodes.

To set up GSLB in a cluster by using the command line interface

Log on to the cluster IP address and perform the following operations at the command prompt:

1. Configure the different GSLB entities. For information, see [Configuring Global Server Load Balancing](#).
Note: When creating the GSLB site, make sure that you specify the cluster IP address and public cluster IP address (needed only when the cluster is deployed behind a NAT device). These parameters are required to ensure the availability of the GSLB auto-sync functionality.

```
add gslb site <siteName> <siteType> <siteIPAddress> -publicIP <ip_addr> -clip <ip_addr> <publicCLIP>
```

2. Create a cluster node group.

```
add cluster nodegroup <name> [-sticky ( YES | NO )]
```

Note: Enable the sticky option if you want to set up GSLB based on VPN users.

3. Bind a single cluster node to the node group.

```
bind cluster nodegroup <name> -node <nodeId>
```

4. Bind the local GSLB site to the nodegroup.

```
bind cluster nodegroup <name> -gslbSite <string>
```

Note: Make sure that the IP address of the local GSLB site IP address is striped (available across all cluster nodes).

5. Bind the ADNS (or ADNS-TCP) service or the DNS (or DNS-TCP) load balancing virtual server to the node group.

To bind the ADNS service:

```
bind cluster nodegroup <name> -service <string>
```

To bind the DNS load balancing virtual server:

```
bind cluster nodegroup <name> -vServer <string>
```

6. Bind the GSLB virtual server to the node group.

```
bind cluster nodegroup <name> -vServer <string>
```

7. [Optional] To setup GSLB based on VPN users, bind the VPN virtual vserver to the GSLB node group.

```
bind cluster nodegroup <name> -vServer <string>
```

8. Verify the configurations.

```
show gslb runningConfig
```

To set up GSLB in a cluster by using the graphical user interface

Log on to the cluster IP address and perform the following operations in the Configuration tab:

1. Configure the GSLB entities.

Navigate to Traffic Management > GSLB to perform the required configurations.

2. Create a node group and perform other node group related configurations.

Navigate to System > Cluster > Node Groups to perform the required configurations.

For the detailed configurations to be performed, see the description provided in the CLI procedure mentioned above.

Using Cache Redirection in a Cluster

Cache redirection in a cluster works in the same way as it does on a standalone NetScaler appliance. The only difference is that the configurations are done on the cluster IP address. For more information on cache redirection, see "[Cache Redirection](#)."

Points to remember when using cache redirection in transparent mode on a cluster:

- Before configuring cache redirection, make sure that you have connected all nodes to the external switch and that you have linksets configured. Otherwise, client requests will be dropped.
- When MAC mode is enabled on a load balancing virtual server, make sure MBF mode is enabled on the cluster by using the `enable ns mode MBF` command. Otherwise, the requests are sent to origin server directly instead of being sent to the cache server.

Using L2 Mode in a Cluster Setup

Note: Supported from NetScaler 10.5 and later releases.

To use L2 mode in a cluster setup, you must make sure of the following:

- Spotted IP addresses must be available on all the nodes as required.
- Linksets must be used to communicate with the external network.
- Asymmetric topologies or asymmetric cluster LA groups are not supported.
- Cluster LA group is recommended.
- Traffic is distributed between the cluster nodes only for deployments where services exist.

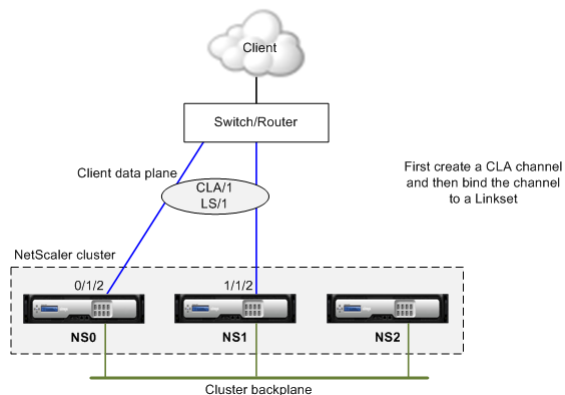
Using Cluster LA Channel with Linksets

In an asymmetric cluster topology, some cluster nodes are not connected to the upstream network. In such a case, you must use linksets. To optimize the performance, you can bind the interfaces that are connected to the switch as a cluster LA channel and then bind the channel to a linkset.

To understand how a combination of cluster LA channel and linksets can be used, consider a three-node cluster for which the upstream switch has only two ports available. You can connect two of the cluster nodes to the switch and leave the other node unconnected.

Note: Similarly, you can also use a combination of ECMP and linksets in an asymmetric topology.

Figure 1. Linksets and cluster LA channel topology



To configure cluster LA channel and linksets by using the NetScaler command line

1. Log on to the cluster IP address.
2. Bind the connected interfaces to a cluster LA channel.


```
add channel CLA/1 -ifnum 0/1/2 1/1/2
```

3. Bind the cluster LA channel to the linkset.

```
add linkset LS/1 -ifnum CLA/1
```

Backplane on LA Channel

In this deployment, LA channels are used for the cluster backplane.



NS0 - nodeId: 0, NSIP: 10.102.29.60

NS1 - nodeId: 1, NSIP: 10.102.29.70

NS2 - nodeId: 2, NSIP: 10.102.29.80

To deploy a cluster with the backplane interfaces as LA channels

1. Create a cluster of nodes NS0, NS1, and NS2.

- a. Log on to the first node that you want to add to the cluster and do the following:

```
create cluster instance 1
add cluster node 0 10.102.29.60 -state ACTIVE
enable cluster instance 1
add ns ip 10.102.29.61 255.255.255.255 -type CLIP
save ns config
reboot -warm
```

- b. Log on to the cluster IP address and do the following:

```
add cluster node 1 10.102.29.70 -state ACTIVE
add cluster node 2 10.102.29.80 -state ACTIVE
```

- c. Log on to the nodes 10.102.29.70 and 10.102.29.80 to join the nodes to the cluster.

```
join cluster -clip 10.102.29.61 -password nsroot
save ns config
reboot -warm
```

As seen in the above commands the interfaces 0/1/1, 1/1/1, and 2/1/1 are configured as the backplane interfaces of the three cluster nodes.

2. Log on to the cluster IP address and do the following:

- a. Create the LA channels for nodes NS0 and NS1.

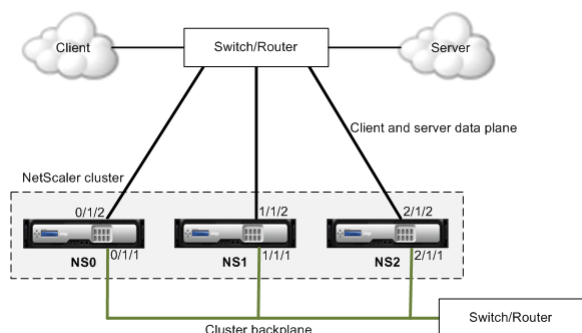
```
add channel 0/LA/1 -ifnum 0/1/1 0/1/2
add channel 1/LA/2 -ifnum 1/1/1 1/1/2
```

- b. Configure the backplane for the cluster nodes.

```
set cluster node 0 -backplane 0/LA/1
set cluster node 1 -backplane 1/LA/2
set cluster node 2 -backplane 2/1/1
```

Common Interfaces for Client and Server and Dedicated Interfaces for Backplane

This is a one-arm deployment of the NetScaler cluster. In this deployment, the client and server networks use the same interfaces to communicate with the cluster. The cluster backplane uses dedicated interfaces for inter-node communication.



NS0 - nodeId: 0, NSIP: 10.102.29.60

NS1 - nodeId: 1, NSIP: 10.102.29.70

NS2 - nodeId: 2, NSIP: 10.102.29.80

To deploy a cluster with a common interface for the client and server and a different interface for the cluster backplane

1. Create a cluster of nodes NS0, NS1, and NS2.

- a. Log on to the first node that you want to add to the cluster and do the following:

```
create cluster instance 1
add cluster node 0 10.102.29.60 -state ACTIVE -backplane 0/1/1
enable cluster instance 1
add ns ip 10.102.29.61 255.255.255.255 -type CLIP
save ns config
reboot -warm
```

- b. Log on to the cluster IP address and do the following:

```
add cluster node 1 10.102.29.70 -state ACTIVE -backplane 1/1/1
add cluster node 2 10.102.29.80 -state ACTIVE -backplane 2/1/1
```

- c. Log on to the nodes 10.102.29.70 and 10.102.29.80 to join the nodes to the cluster.

```
join cluster -clip 10.102.29.61 -password nsroot
save ns config
reboot -warm
```

As seen in the above commands the interfaces 0/1/1, 1/1/1, and 2/1/1 are configured as the backplane interfaces of the three cluster nodes.

2. On the cluster IP address, create VLANs for the backplane interfaces and for the client and server interfaces.

```
//For the backplane interfaces
add vlan 10
bind vlan 10 0/1/1 1/1/1 2/1/1
```

```
//For the interfaces that are connected to the client and server networks.
add vlan 20
bind vlan 20 0/1/2 1/1/2 2/1/2
```

3. On the switch, create VLANs for the interfaces corresponding to the backplane interfaces and the client and server interfaces. The following sample configurations are provided for the Cisco® Nexus 7000 C7010 Release 5.2(1) switch. Similar configurations must be performed on other switches.


```
//For the backplane interfaces. Repeat for each interface...
interface Ethernet2/47
switchport access vlan 100
switchport mode access
```

end

```
//For the interfaces connected to the client and server networks. Repeat for each inter
interface Ethernet2/47
switchport access vlan 200
switchport mode access
end
```


Common Switch for Client, Server, and Backplane

In this deployment, the client, server, and backplane use dedicated interfaces on the same switch to communicate with the NetScaler cluster.



NS0 - nodeId: 0, NSIP: 10.102.29.60

NS1 - nodeId: 1, NSIP: 10.102.29.70

NS2 - nodeId: 2, NSIP: 10.102.29.80

To deploy a cluster with a common switch for the client, server, and backplane

1. Create a cluster of nodes NS0, NS1, and NS2.

a. Log on to the first node that you want to add to the cluster and do the following:

```
create cluster instance 1
add cluster node 0 10.102.29.60 -state ACTIVE -backplane 0/1/1
enable cluster instance 1
add ns ip 10.102.29.61 255.255.255.255 -type CLIP
save ns config
reboot -warm
```

b. Log on to the cluster IP address and do the following:

```
add cluster node 1 10.102.29.70 -state ACTIVE -backplane 1/1/1
add cluster node 2 10.102.29.80 -state ACTIVE -backplane 2/1/1
```

c. Log on to the nodes 10.102.29.70 and 10.102.29.80 to join the nodes to the cluster.

```
join cluster -clip 10.102.29.61 -password nsroot
save ns config
reboot -warm
```

As seen in the above commands the interfaces 0/1/1, 1/1/1, and 2/1/1 are configured as the backplane interfaces of the three cluster nodes.

2. On the cluster IP address, create VLANs for the backplane, client, and server interfaces.

```
//For the backplane interfaces
add vlan 10
bind vlan 10 0/1/1 1/1/1 2/1/1

//For the client-side interfaces
add vlan 20
bind vlan 20 0/1/2 1/1/2 2/1/2

//For the server-side interfaces
```

```
add vlan 30
bind vlan 30 0/1/3 1/1/3 2/1/3
```

3. On the switch, create VLANs for the interfaces corresponding to the backplane interfaces and the client and server interfaces. The following sample configurations are provided for the Cisco® Nexus 7000 C7010 Release 5.2(1) switch. Similar configurations must be performed on other switches.


```
//For the backplane interfaces. Repeat for each interface...
interface Ethernet2/47
switchport access vlan 100
switchport mode access
end
```

```
//For the client interfaces. Repeat for each interface...
interface Ethernet2/48
switchport access vlan 200
switchport mode access
end
```

```
//For the server interfaces. Repeat for each interface...
interface Ethernet2/49
switchport access vlan 300
switchport mode access
end
```

Common Switch for Client and Server and Dedicated Switch for Backplane

In this deployment, the clients and servers use different interfaces on the same switch to communicate with the NetScaler cluster. The cluster backplane uses a dedicated switch for inter-node communication.



NS0 - nodeId: 0, NSIP: 10.102.29.60

NS1 - nodeId: 1, NSIP: 10.102.29.70

NS2 - nodeId: 2, NSIP: 10.102.29.80

To deploy a cluster with the same switch for the clients and servers and a different switch for the cluster backplane

1. Create a cluster of nodes NS0, NS1, and NS2.

- a. Log on to the first node that you want to add to the cluster and do the following:

```
create cluster instance 1
add cluster node 0 10.102.29.60 -state ACTIVE -backplane 0/1/1
enable cluster instance 1
add ns ip 10.102.29.61 255.255.255.255 -type CLIP
save ns config
reboot -warm
```

- b. Log on to the cluster IP address and do the following:

```
add cluster node 1 10.102.29.70 -state ACTIVE -backplane 1/1/1
add cluster node 2 10.102.29.80 -state ACTIVE -backplane 2/1/1
```

- c. Log on to the nodes 10.102.29.70 and 10.102.29.80 to join the nodes to the cluster.

```
join cluster -clip 10.102.29.61 -password nsroot
save ns config
reboot -warm
```

As seen in the above commands the interfaces 0/1/1, 1/1/1, and 2/1/1 are configured as the backplane interfaces of the three cluster nodes.

2. On the cluster IP address, create VLANs for the backplane, client, and server interfaces.

```
//For the backplane interfaces
add vlan 10
bind vlan 10 0/1/1 1/1/1 2/1/1

//For the client-side interfaces
add vlan 20
bind vlan 20 0/1/2 1/1/2 2/1/2
```

```
//For the server-side interfaces
add vlan 30
bind vlan 30 0/1/3 1/1/3 2/1/3
```

3. On the switch, create VLANs for the interfaces corresponding to the backplane interfaces and the client and server interfaces. The following sample configurations are provided for the Cisco® Nexus 7000 C7010 Release 5.2(1) switch. Similar configurations must be performed on other switches.


```
//For the backplane interfaces. Repeat for each interface...
interface Ethernet2/47
switchport access vlan 100
switchport mode access
end
```

```
//For the client interfaces. Repeat for each interface...
interface Ethernet2/48
switchport access vlan 200
switchport mode access
end
```

```
//For the server interfaces. Repeat for each interface...
interface Ethernet2/49
switchport access vlan 300
switchport mode access
end
```

Different Switch for Every Node

In this deployment, each cluster node is connected to a different switch and trunk links are configured between the switches.



The cluster configurations will be the same as the other deployments scenarios. Most of the client-side configurations will be done on the client-side switches.

Sample Cluster Configurations

The following example can be used to configure a four-node cluster with ECMP, cluster LA, or Linksets.

1. Create the cluster.

- a. Log on to first node.
- b. Add the cluster instance.

```
add cluster instance 1
```

- c. Add the first node to the cluster.

```
add cluster node 0 10.102.33.184 -backplane 0/1/1
```

- d. Enable the cluster instance.

```
enable cluster instance 1
```

- e. Add the cluster IP address.

```
add ns ip 10.102.33.185 255.255.255.255 -type CLIP
```

- f. Save the configurations.

```
save ns config
```

- g. Warm reboot the appliance.

```
reboot -warm
```

2. Add the other three nodes to the cluster.

- a. Log on to cluster IP address.
- b. Add the second node to the cluster.

```
add cluster node 1 10.102.33.187 -backplane 1/1/1
```

- c. Add the third node to the cluster.

```
add cluster node 2 10.102.33.188 -backplane 2/1/1
```

- d. Add the fourth node to the cluster.

```
add cluster node 3 10.102.33.189 -backplane 3/1/1
```

3. Join the added nodes to the cluster. This step is not applicable for the first node.

- a. Log on to each newly added node.
- b. Join the node to the cluster.

```
join cluster -clip 10.102.33.185 -password nsroot
```

- c. Save the configuration.

```
save ns config
```

- d. Warm reboot the appliance.

```
reboot -warm
```

4. Configure the NetScaler cluster through the cluster IP address.

```
// Enable load balancing feature
enable ns feature lb
```

```
// Add a load balancing virtual server
add lb vserver first_lbvserver http
....
....
```

5. Configure any one of the following (ECMP, cluster LA, or Linkset) traffic distribution mechanisms for the cluster.

o **ECMP**

- a. Log on to the cluster IP address.

- b. Enable the OSPF routing protocol.

```
enable ns feature ospf
```

- c. Add a VLAN.

```
add vlan 97
```

- d. Bind the interfaces of the cluster nodes to the VLAN.

```
bind vlan 97 -ifnum 0/1/4 1/1/4 2/1/4 3/1/4
```

- e. Add a spotted SNIP on each node and enable dynamic routing on it.

```
add ns ip 1.1.1.10 255.255.255.0 -ownerNode 0 -dynamicRouting ENABLED
add ns ip 1.1.1.11 255.255.255.0 -ownerNode 1 -dynamicRouting ENABLED
add ns ip 1.1.1.12 255.255.255.0 -ownerNode 2 -dynamicRouting ENABLED
add ns ip 1.1.1.13 255.255.255.0 -ownerNode 3 -dynamicRouting ENABLED
```

- f. Bind one of the SNIP addresses to the VLAN.

```
bind vlan 97 -ipAddress 1.1.1.10 255.255.255.0
```

- g. Configure the routing protocol on ZebOS by using vtysh shell.

o Static cluster LA

- a. Log on to the cluster IP address.
b. Add a cluster LA channel.

```
add channel CLA/1 -speed 1000
```

- c. Bind the interfaces to the cluster LA channel.

```
bind channel CLA/1 0/1/5 1/1/5 2/1/5 3/1/5
```

- d. Perform equivalent configuration on the switch.

o Dynamic cluster LA

- a. Log on to the cluster IP address.
b. Add the interfaces to the cluster LA channel.

```
set interface 0/1/5 -lacpmode active -lacpkey 5 -lagtype cluster
set interface 1/1/5 -lacpmode active -lacpkey 5 -lagtype cluster
set interface 2/1/5 -lacpmode active -lacpkey 5 -lagtype cluster
set interface 3/1/5 -lacpmode active -lacpkey 5 -lagtype cluster
```

- c. Perform equivalent configuration on the switch.

- o **Linksets.** Assume that the node with nodeId 3 is not connected to the switch. You must configure a linkset so that the unconnected node can use the other node interfaces to communicate with the switch.

- a. Log on to the cluster IP address.
b. Add a linkset.

```
add linkset LS/1
```

- c. Bind the connected interfaces to the linkset.

```
bind linkset LS/1 -ifnum 0/1/6 1/1/6 2/1/6
```

6. Update the state of the cluster nodes to ACTIVE.

```
set cluster node 0 -state ACTIVE
set cluster node 1 -state ACTIVE
set cluster node 2 -state ACTIVE
set cluster node 3 -state ACTIVE
```

Upgrading or Downgrading the NetScaler Cluster

All the nodes of a NetScaler cluster must be running the same software version. Therefore, to upgrade or downgrade the cluster, you must upgrade or downgrade each NetScaler appliance of the cluster, one node at a time.

A node that is being upgraded or downgraded is not removed from the cluster. The node continues to be a part of the cluster and serves traffic uninterrupted, except for the down-time when the node reboots after it is upgraded or downgraded. However, due to software version mismatch among the cluster nodes, configuration propagation is disabled on the cluster and is enabled only after all the cluster nodes are of the same version. Since configuration propagation is disabled during upgrading or downgrading a cluster, you cannot perform any configurations through the cluster IP address during this time.

Points to note before upgrading or downgrading the cluster

- You cannot add cluster nodes while upgrading or downgrading the cluster software version.
- You can perform node-level configurations through the NSIP address of individual nodes, but you must make sure that you perform the same configurations on all the nodes to maintain them in synch.
- You cannot execute the start nstrace command from the cluster IP address when the cluster is being upgraded. However, you can get the trace of individual nodes by performing this operation on individual cluster nodes through their NetScaler IP (NSIP) address.
- Configurations can be lost during the downgrade of the cluster.
- Owing to changes in cluster licensing that were made in NetScaler 10.5 Build 52.11 (see [license requirements](#)), look into the following:
 - If the cluster is setup in a build prior to NetScaler 10.5 Build 52.11, the cluster will work with the separate cluster license file. No changes are required.
 - If the cluster is setup in NetScaler 10.5 Build 52.11 or later releases and then downgraded to a build prior to NetScaler 10.5 Build 52.11, the downgraded cluster will not work as it now expects a separate cluster license file.
- While upgrading from any NetScaler 10.1 build to a later release, syncookie must be disabled on all TCP profiles (using the "set ns tcpProfile <name> -synCookie DISABLED" command) and after that a striped SNIP must be added on the CLIP subnet. Once upgraded, syncookie can be enabled again.
- While upgrading the NetScaler appliance from a NetScaler 10.1 build to a NetScaler 10.5 build, do not execute the "show audit messages" command as this can cause the NetScaler appliance to crash.
- NetScaler 10.5 54.x and 55.x builds are not suitable for cluster deployment. This is because, for services that need probing, SYN packets are processed locally (on the flow receiver) even though syncookie is disabled.
- When a cluster is being upgraded, it is possible that the upgraded nodes have some additional features activated that are not available on the nodes that are not upgraded. This results in a license mismatch warning while the cluster is being upgraded. This warning will be automatically resolved when all the cluster nodes are upgraded.

To upgrade or downgrade the software of the cluster nodes

1. Make sure the cluster is stable and the configurations are synchronized on all the nodes.
2. For each cluster node perform the following:
 - Note: Citrix recommends that you wait for the previous node to become active before upgrading or downgrading the next node.
 - a. Upgrade or downgrade the cluster node. For detailed information about upgrading and downgrading the software of an appliance, see ["Upgrading or Downgrading the System Software"](#).
 - b. Save the configurations.
 - c. Reboot the appliance.
3. Repeat step 2 for each of the other cluster nodes.

Operations Supported on Individual Cluster Nodes

As a rule, NetScaler appliances that are a part of a cluster cannot be individually configured from their NSIP address. However, there are some operations that are an exception to this rule. These operations, when executed from the NSIP address, are not propagated to other cluster nodes.

The operations are:

- cluster instance (set | rm | enable | disable)
- cluster node (set | rm)
- nstrace (start | show | stop)
- interface (set | enable | disable)
- route (add | rm | set | unset)
- arp (add | rm | send -all)
- force cluster sync
- sync cluster files
- disable ntp sync
- save ns config
- reboot
- shutdown

For example, when you execute the command `disable interface 1/1/1` from the NSIP address of a cluster node, the interface is disabled only on that node. Since the command is not propagated, the interface 1/1/1 remains enabled on all the other cluster nodes.

Clustering FAQs

How many NetScaler appliances can be included in a single NetScaler cluster.

A NetScaler cluster can include as few as 2 or as many as 32 NetScaler nCore hardware or virtual appliances. Each of these nodes must satisfy the criteria specified in "[Prerequisites for Cluster Nodes](#)".

Can a NetScaler appliance be a part of multiple clusters?

No. A NetScaler appliance can belong to one cluster only.

What is a cluster IP address? What is its subnet mask?

The cluster IP address is the management address of a NetScaler cluster. All cluster configurations must be performed by accessing the cluster through this address. The subnet mask of the cluster IP address is fixed at 255.255.255.255.

How can I make a specific cluster node as the cluster configuration coordinator?

To manually set a specific node as the cluster configuration coordinator, you must set the priority of that node to the lowest numeric value (highest priority). To understand, let us consider a cluster with three nodes that have the following priorities:

n1 - 29, n2 - 30, n3 - 31

Here, n1 is the configuration coordinator. If you want to make n2 the configuration coordinator, you must set its priority to a value that is lower than n1, for example, 28. On saving the configuration, n2 becomes the configuration coordinator.

Note: n2 with its original priority value of 30 will also become the configuration coordinator when n1 goes down, as the node with the next lowest priority value is selected in the event that the configuration coordinator goes down.

Why are the network interfaces of a cluster represented in 3-tuple (n/u/c) notation instead of the regular 2-tuple (u/c) notation?

When a NetScaler appliance is part of a cluster, you must be able to identify the node to which the interface belongs. Therefore, the network interface naming convention for cluster nodes is modified from u/c to n/u/c, where n denotes the node Id.

How can I set the hostname for a cluster node?

The hostname of a cluster node must be specified by executing the set ns hostname command through the cluster IP address. For example, to set the hostname of the cluster node with ID 2, the command is:

```
> set ns hostname hostName1 -ownerNode 2
```

Can I automatically detect NetScaler appliances so that I can add them to a cluster?

Yes. The configuration utility allows you to discover appliances that are present in the same subnet as the NSIP address of the configuration coordinator. For more information, see "[Discovering NetScaler Appliances](#)".

Is the traffic serving capability of a cluster affected if a node is removed or disabled or reboot or shutdown or made inactive?

Yes. When any of these operations are performed on an active node of the cluster, the cluster will have one less node to serve traffic. Also, existing connections on this node are terminated.

I have multiple standalone appliances, each of which has different configurations. Can I add them to a single cluster?

Yes. You can add appliances that have different configurations to a single cluster. However, when the appliance is added to the cluster, the existing configurations are cleared. To use the configurations that are available on each of the individual appliances, you must:

1. Create a single *.conf file for all the configurations.
2. Edit the configuration file to remove features that are not supported in a cluster environment.
3. Update the naming convention of interfaces from 2-tuple (u/c) format to 3-tuple (n/u/c) format.
4. Apply the configurations to the configuration coordinator node of the cluster by using the batch command.

Can I migrate the configurations of a standalone NetScaler appliance or an HA setup to the clustered setup?

No. When a node is added to a clustered setup, its configurations are implicitly cleared by using the clear ns config command (with the extended option). In addition, the SNIP addresses and all VLAN configurations (except default VLAN and NSVLAN) are cleared. Therefore, it is recommended that you back up the configurations before adding the appliance to a cluster. Before using the backed-up configuration file for the cluster, you must:

1. Edit the configuration file to remove features that are not supported in a cluster environment.
2. Update the naming convention of interfaces from two-tuple (x/y) format to three-tuple (x/y/z) format.
3. Apply the configurations to the configuration coordinator node of the cluster by using the batch command.

How can I configure a cluster that includes nodes from different networks?

Note: Supported from NetScaler 11 onwards.

A cluster that includes nodes from different networks is called a L3 cluster (sometimes referred to as a cluster in INC mode). In an L3 cluster, all nodes that belong to a single network must be grouped together in a single nodegroup. Therefore, if a cluster includes two nodes each from three different networks, you will have to create 3 nodegroups (one for each network) and associate each of these nodegroups with the nodes that belong to that network. For configuration information, see the steps to setup a cluster.

How can I configure/unconfigure the NSVLAN on a cluster?

- To make the NSVLAN available in a cluster, make sure that each appliance has the same NSVLAN configured before it is added to cluster.

- o To remove the NSVLAN from a cluster node, first remove the node from the cluster and then delete the NSVLAN from the appliance.

I have an cluster setup where some NetScaler nodes are not connected to the external network. Can the cluster still function normally.

Yes. The cluster supports a mechanism called linksets, which allows unconnected nodes to serve traffic by using the interfaces of connected nodes. The unconnected nodes communicate with the connected nodes through the cluster backplane. For more information, see ["Using Linksets"](#).

How can deployments that require MAC-Based Forwarding (MBF) be supported in a clustered setup?

Deployments that use MBF must use linksets. For more information, see ["Using Linksets"](#).

Can I execute commands from the NSIP address of a cluster node?

No. Access to individual cluster nodes through the NetScaler IP (NSIP) addresses is read-only. Therefore, when you log on to the NSIP address of a cluster node you can only view the configurations and the statistics. You cannot configure anything. However, there are some operations you can execute from the NSIP address of a cluster node.

For more information, see ["Operations Supported on Individual Nodes"](#).

Can I disable configuration propagation among cluster nodes?

No, you cannot explicitly disable the propagation of cluster configurations among cluster nodes. However, during a software upgrade or downgrade, a version mismatch can automatically disable configuration propagation.

Can I change the NSIP address or change the NSVLAN of a NetScaler appliance when it is a part of the cluster?

No. To make such changes you must first remove the appliance from the cluster, perform the changes, and then add the appliance to the cluster.

Does the NetScaler cluster support L2 and L3 Virtual Local Area Networks (VLANs)?

Yes. A cluster supports VLANs between cluster nodes. The VLANs must be configured on the cluster IP address.

- o **L2 VLAN.** You can create a layer2 VLAN by binding interfaces that belong to different nodes of the cluster.
- o **L3 VLAN.** You can create a layer3 VLAN by binding IP addresses that belong to different nodes of the cluster. The IP addresses must belong to the same subnet. Make sure that one of the following criteria is satisfied. Otherwise, the L3 VLAN bindings can fail.

All nodes have an IP address on the same subnet as the one bound to the VLAN.

The cluster has a striped IP address and the subnet of that IP address is bound to the VLAN.

When you add a new node to a cluster that has only spotted IPs, the sync happens before spotted IP addresses are assigned to that node. In such cases, L3 VLAN bindings can be lost.

To avoid this loss, either add a striped IP or add the L3 VLAN bindings on the NSIP of the newly added node.

How can I configure SNMP on a NetScaler cluster?

SNMP monitors the cluster, and all the nodes of the cluster, in the same way that it monitors a standalone appliance. The only difference is that SNMP on a cluster must be configured through the cluster IP address.

When generating hardware specific traps, two additional varbinds are included to identify the node of the cluster: node ID and NSIP address of the node.

For detailed information about configuring SNMP, see ["SNMP"](#).

What details must I have available when I contact technical support for cluster-related issues?

The NetScaler appliance provides a `show techsupport -scope cluster` command that extracts configuration data, statistical information, and logs of all the cluster nodes. You must run this command on the cluster IP address.

The output of this command is saved in a file named `collector_cluster_<nsip_CCO>_P_<date-timestamp>.tar.gz` which is available in the `/var/tmp/support/cluster/` directory of the configuration coordinator.

Send this archive to the technical support team to debug the issue.

Can I use striped IP addresses as the default gateway of servers?

In case of cluster deployments, make sure the default gateway of the server points to a striped IP address (if you are using a NetScaler-owned IP address). For example, in case of LB deployments with USIP enabled, the default gateway must be a striped SNIP address.

Can I view routing configurations of a specific cluster node from the cluster IP address?

Yes. You can view and clear the configurations specific to a node by specifying the owner node while entering the vtysh shell.

For example, to view the output of a command on nodes 0 and 1, the command is as follows:

```
> vtysh
ns# owner-node 0 1
ns(node-0 1)# show cluster state
ns(node-0 1)# exit-cluster-node
ns#
```

How can I specify the node for which I want to set the LACP system priority?

Note: Supported from NetScaler 10.1 onwards.

In a cluster, you must set that node as the owner node by using the set lacp command.

For example: To set the LACP system priority for node with ID 2:

```
> set lacp -sysPriority 5 -ownerNode 2
```

How are IP tunnels configured in a cluster setup?

Note: Supported from NetScaler 10.1 onwards.

Configuring IP tunnels in a cluster is the same as on a standalone appliance. The only difference is that in a cluster setup, the local IP address must be a striped SNIP or MIP address. For more information, see "[Configuring IP Tunnels](#)".

How can I add a failover interface set (FIS) on the nodes of a NetScaler cluster?

Note: Supported from NetScaler 10.5 onwards.

On the cluster IP address, specify the ID of the cluster node on which the FIS must be added, using the command as follows:

```
add fis <name> -ownerNode <nodeId>
```

Note:

- The FIS name for each cluster node must be unique.
- A cluster LA channel can be added to a FIS. You must make sure that the cluster LA channel has a local interface as a member interface.

For more information on FIS, see "[Configuring FIS](#)".

How are net profiles configured in a cluster setup?

Note: Supported from NetScaler 10.5 onwards.

You can bind spotted IP addresses to a net profile. This net profile can then be bound to spotted load balancing virtual server or service (that is defined using a nodegroup). The following recommendations must be followed, failing which, the net profile configurations are not honored and the USIP/USNIP settings will be used:

Note:

- If the strict parameter of the nodegroup is set to `Yes`, the net profile must contain a minimum of one IP address from each nodegroup member.
- If the strict parameter of the nodegroup is set to `No`, the net profile must include at least one IP address from each of the cluster nodes.

Troubleshooting the NetScaler Cluster

If a failure occurs in a NetScaler cluster, the first step in troubleshooting is to get information on the cluster instance and the cluster nodes by running the `show cluster instance <clId>` and `show cluster node <nodeId>` commands respectively.

If you are not able to find the issue by using the above two approaches, you can use one of the following:

- o **Isolate the source of the failure.** Try bypassing the cluster to reach the server. If the attempt is successful, the problem is probably with the cluster setup.
- o **Check the commands recently executed.** Run the history command to check the recent configurations performed on the cluster. You can also review the `ns.conf` file to verify the configurations that have been implemented.
- o **Check the ns.log files.** Use the log files, available in the `/var/log/` directory of each node, to identify the commands executed, status of commands, and the state changes.
- o **Check the newnslog files.** Use the newnslog files, available in the `/var/nslog/` directory of each node, to identify the events that have occurred on the cluster nodes. You can view multiple newnslog files as a single file, by copying the files to a single directory, and then running the following command:

```
nsconmsg -K newnslog-node<id> -K newnslog.node<id> -d current
```

If you still cannot resolve the issue, you can try tracing the packets on the cluster or use the `show techsupport -scope cluster` command to send the report to the technical support team.

Tracing the Packets of a NetScaler Cluster

The NetScaler operating system provides a utility called *nstrace* to get a dump of the packets that are received and sent out by an appliance. The utility stores the packets in trace files. You can use these files to debug problems in the flow of packets to the cluster nodes. The trace files must be viewed with the Wireshark application.

Some salient aspects of the *nstrace* utility are:

- Can be configured to trace packets selectively by using classic expressions and default expressions.
- Can capture the trace in multiple formats: *nstrace* format (.cap) and TCP dump format (.pcap).
- Can aggregate the trace files of all cluster nodes on the configuration coordinator.
- Can merge multiple trace files into a single trace file (only for .cap files).

You can use the *nstrace* utility from the NetScaler command line or the NetScaler shell.

To trace packets of a standalone appliance

Run the `start nstrace` command on the appliance. The command creates trace files in the `/var/nstrace/<date-timestamp>` directory. The trace file names are of the form `nstrace<id>.cap`.

You can view the status by executing the `show nstrace` command. You can stop tracing the packets by executing the `stop nstrace` command.

Note: You can also run the *nstrace* utility from the NetScaler shell by executing the `nstrace.sh` file. However, it is recommended that you use the *nstrace* utility through the NetScaler command line interface.

To trace packets of a cluster

You can trace the packets on all the cluster nodes and obtain all the trace files on the configuration coordinator.

Run the `start nstrace` command on the cluster IP address. The command is propagated and executed on all the cluster nodes. The trace files are stored in individual cluster nodes in the `/var/nstrace/<date-timestamp>` directory. The trace file names are of the form `nstrace<id>_node<id>.cap`.

You can use the trace files of each node to debug the nodes operations. But if you want the trace files of all cluster nodes in one location, you must run the `stop nstrace` command on the cluster IP address. The trace files of all the nodes are downloaded on the cluster configuration coordinator in the `/var/nstrace/<date-timestamp>` directory as follows:

```
/var/nstrace/08Mar2012_16_30_25
├── node0
│   ├── nstrace1_node0.cap
│   ├── nstrace2_node0.cap
│   └── nstrace3_node0.cap
├── node1
│   ├── nstrace1_node1.cap
│   └── nstrace2_node1.cap
└── node2
    ├── nstrace1_node2.cap
    └── nstrace2_node2.cap
```

Merge multiple trace files

You can prepare a single file from the trace files (supported only for .cap files) obtained from the cluster nodes. The single trace files gives you a cumulative view of the trace of the cluster packets. The trace entries in the single trace file are sorted based on the time the packets were received on the cluster.

To merge the trace files, at the NetScaler shell, type:

```
nstracemerge.sh -srcdir <DIR> -dstdir <DIR> -filename <name> -filesize <num>
```

where,

- `srcdir` is the directory from which the trace files are merged. All trace files within this directory are merged into a single file.
- `dstdir` is the directory where the merged trace file are created.

- o filename is the name of the trace file that is created.
- o filesize is the size of the trace file.

Examples

Following are some examples of using the nstrace utility to filter packets.

- o To trace the packets on the backplane interfaces of three nodes:

Using classic expressions:

```
start nstrace -filter "INTF == 0/1/1 && INTF == 1/1/1 && INTF == 2/1/1"
```

Using default expressions:

```
start nstrace -filter "CONNECTION.INTF.EQ("0/1/1") && CONNECTION.INTF.EQ("1/1/1"
```

- o To trace the packets from a source IP address 10.102.34.201 or from a system whose source port is greater than 80 and the service name is not "s1":

Using classic expressions

```
start nstrace -filter "SOURCEIP == 10.102.34.201 || (SVCNAME != s1 && SOURCEPORT
```

Using default expressions

```
start nstrace -filter "CONNECTION.SRCIP.EQ(10.102.34.201) || (CONNECTION.SVCNAME
```

Troubleshooting Common Issues

While joining a node to the cluster, I get the following message, "ERROR: Invalid interface name/number." What must I do to resolve this error?

This error occurs if you provided an invalid or incorrect backplane interface while using the add cluster node command to add the node. To resolve this error, verify the interface you provided while adding the node. Make sure that you have not specified the appliance's management interface as the backplane interface, and that the <nodeId> bit of the interface is the same as the node's Id. For example, if the nodeId is 3, the backplane interface must be 3/<c>/<u>.

While joining a node to the cluster, I get the following message, "ERROR: Clustering cannot be enabled, because the local node is not a member of the cluster." What must I do to resolve this error?

This error occurs when you try to join a node without adding the node's NSIP to the cluster. To resolve this error, you must first add the node's NSIP address to the cluster by using the add cluster node command and then execute the join cluster command.

While joining a node to the cluster, I get the following message, "ERROR: Connection refused." What must I do to resolve this error?

This error can occur due to the following reasons:

- **Connectivity problems.** The node cannot connect to the cluster IP address. Try pinging the cluster IP address from the node that you are trying to join.
- **Duplicate cluster IP address.** Check to see if the cluster IP address exists on some non-cluster node. If it does, create a new cluster IP address and try re-joining the cluster.

While joining a node to the cluster, I get the following message, "ERROR: License mismatch between the configuration coordinator and the local node." What must I do to resolve this error?

The appliance that you are joining to the cluster must have the same licenses as the configuration coordinator. This error occurs when the licenses on the node you are joining do not match the licenses on the configuration coordinator. To resolve this error, run the following commands on both the nodes and compare the outputs.

From the command line:

- show ns hardware
- show ns license

From the shell:

- nsconmsg -g feature -d stats
- ls /nsconfig/license
- View the contents of the /var/log/license.log file

What must I do when the configurations of a cluster node are not in synch with the cluster configurations?

In most cases, the configurations are automatically synchronized between all the cluster nodes. However, if you feel that the configurations are not synchronized on a specific node, you must force the synchronization by executing the force cluster sync command from the node that you want to synchronize. For more information, see "[Synchronizing Cluster Configurations](#)".

When configuring a cluster node, I get the following message, "ERROR: Session is read-only; connect to the cluster IP address to modify the configuration."

All configurations on a cluster must be done through the cluster IP address and the configurations are propagated to the other cluster nodes. All sessions established through the NetScaler IP (NSIP) address of individual nodes are read-only.

Why does the node state show "INACTIVE" when the node health shows "UP"?

A healthy node can be in the INACTIVE state for a number of reasons. A scan of ns.log or error counters can help you determine the exact reason.

How can I resolve the health of a node when its health shows "NOT UP"?

Node health "**Not UP**" indicates that there are some issues with the node. To know the root cause, you must run the show cluster node command. This command displays the node properties and the reason for the node failure.

What must I do when the health of a node shows as "NOT UP" and the reason indicates that configuration commands have failed on a node?

This issue arises when some commands are not executed on the cluster nodes. In such cases, you must make sure that the configurations are synchronized using one of the following options:

- o If some of the cluster nodes are in this state, you must perform the force cluster synchronization operation on those nodes. For more information, see ["Synchronizing Cluster Configurations"](#).
- o If all cluster nodes are in this state, you must disable and then enable the cluster instance on all the cluster nodes.

When I run the `set vserver` command, I get the following message, "No such resource." What must I do to resolve this issue?
The `set vserver` command is not supported in clustering. The `unset vserver`, `enable vserver`, `disable vserver`, and `rm vserver` commands are also not supported. However, the `show vserver` command is supported.

I cannot configure the cluster over a Telnet session. What must I do?
Over a telnet session, the cluster IP address can be accessed only in read-only mode. Therefore, you cannot configure a cluster over a telnet session.

I notice a significant time difference across the cluster nodes. What must I do to resolve this issue?
When PTP packets are dropped due to backplane switch or if the physical resources are over-committed in a virtual environment, the time will not get synchronized.

To synchronize the times, you must do the following on the cluster IP address:

1. Disable PTP.

`set ptp -state disable`
2. Configure Network Time Protocol (NTP) for the cluster. For more information, see ["Setting up Clock Synchronization"](#).

What must I do, if there is no connectivity to the cluster IP address and the NSIP address of a cluster node?
If you cannot access to the cluster IP address or the NSIP of a cluster node, you must access the appliance through the serial console. For more information, see ["Using the Command Line Interface"](#).

If the NSIP address is reachable, you can SSH to the cluster IP address from the shell by executing the following command at the shell prompt:

```
# ssh nsroot@<cluster IP address>
```

What must I do to recover a cluster node that has connectivity issues?
To recover a node that has connectivity issues:

1. Disable the cluster instance on that node (since you cannot execute commands from the NSIP of a cluster node).
2. Execute the commands required to recover the node.
3. Enable the cluster instance on that node.

Some nodes of the cluster have two default routes. How can I remove the second default route from the cluster node?
To delete the additional default route, do the following on each node that has the extra route:

1. Disable the cluster instance.

`disable cluster instance <clld>`
2. Remove the route.

`rm route <network> <netmask> <gateway>`
3. Enable the cluster instance.

`enable cluster instance <clld>`

The cluster functionality gets affected when an existing cluster node comes online. What must I do to resolve this issue?
If RPC password of node is changed from the cluster IP address when that node is out of the cluster, then, when the node comes online, there is a mismatch in rpc credentials and this could affect cluster functionality. To solve this issue, use the `set ns rpcNode` command to update the password on the NSIP of the node which has come online.

CloudBridge Connector

The CloudBridge Connector feature of the Citrix NetScaler appliance connects enterprise datacenters to external clouds and hosting environments, making the cloud a secure extension of your enterprise network. Cloud-hosted applications appear as though they are running on one contiguous enterprise network. With Citrix CloudBridge Connector, you can augment your datacenters with the capacity and efficiency available from cloud providers.

The CloudBridge Connector enables you to move your applications to the cloud to reduce costs and increase reliability.

In addition to using CloudBridge Connector between a datacenter and a cloud, you can use it to connect two datacenters for a high-capacity secure and accelerated link.

Understanding CloudBridge Connector

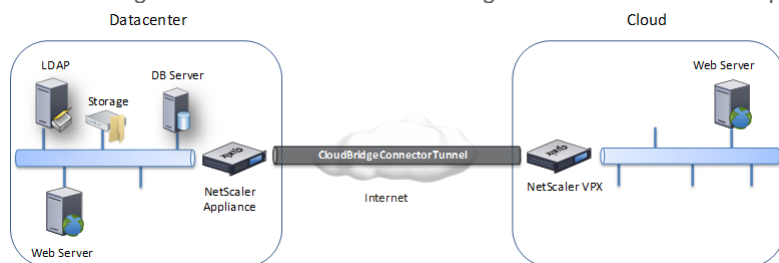
Updated: 2014-04-14

To implement the Citrix CloudBridge Connector solution, you connect a datacenter to another datacenter or an external cloud by setting up a tunnel called the CloudBridge Connector tunnel.

To connect a datacenter to another datacenter, you set up a CloudBridge Connector tunnel between two NetScaler appliances, one in each datacenter.

To connect a datacenter to an external cloud (for example, Amazon AWS cloud), you set up a CloudBridge Connector tunnel between a NetScaler appliance in the datacenter and a virtual appliance (VPX) that resides in the Cloud. The remote end point can be a CloudBridge Connector or a NetScaler VPX with platinum license.

The following illustration shows a CloudBridge Connector tunnel set up between a datacenter and an external cloud.



The appliances between which a CloudBridge Connector tunnel is set up are called the *end points* or *peers* of the CloudBridge Connector tunnel.

A CloudBridge Connector tunnel uses the following protocols:

- Generic Routing Encapsulation (GRE) protocol
- Open-standard IPSec Protocol suite, in transport mode

The GRE protocol provides a mechanism for encapsulating packets, from a wide variety of network protocols, to be forwarded over another protocol. GRE is used to:

- Connect networks running non-IP and non-routable protocols.
- Bridge across a wide area network (WAN).
- Create a transport tunnel for any type of traffic that needs to be sent unchanged across a different network.

The GRE protocol encapsulates packets by adding a GRE header and a GRE IP header to the packets.

The Internet Protocol security (IPSec) protocol suite secures communication between peers in the CloudBridge Connector tunnel.

In a CloudBridge Connector tunnel, IPSec ensures:

- Data integrity
- Data origin authentication

- Data confidentiality (encryption)
- Protection against replay attacks

IPSec uses the transport mode in which the GRE encapsulated packet is encrypted. The encryption is done by the Encapsulating Security Payload (ESP) protocol. The ESP protocol ensures the integrity of the packet by using a HMAC hash function, and ensures confidentiality by using an encryption algorithm. After the packet is encrypted and the HMAC is calculated, an ESP header is generated. The ESP header is inserted after the GRE IP header and, an ESP trailer is inserted at the end of the encrypted payload.

Peers in the CloudBridge Connector tunnel use the Internet Key Exchange version (IKE) protocol (part of the IPSec protocol suite) to negotiate secure communication, as follows:

- The two peers mutually authenticate with each other, using one of the following authentication methods:

Pre-shared key authentication. A text string called a pre-shared key is manually configured on each peer. The pre-shared keys of the peers are matched against each other for authentication. Therefore, for the authentication to be successful, you must configure the same pre-shared key on each of the peers.

Digital certificates authentication. The initiator (sender) peer signs message interchange data by using its private key, and the other receiver peer uses the sender's public key to verify the signature. Typically, the public key is exchanged in messages containing an X.509v3 certificate. This certificate provides a level of assurance that a peer's identity as represented in the certificate is associated with a particular public key.

- The peers then negotiate to reach agreement on:

An encryption algorithm.

Cryptographic keys for encrypting data in one peer and decrypting the data in the other.

This agreement upon the security protocol, encryption algorithm and cryptographic keys is called a Security Association (SA). SAs are one-way (simplex). For example, when two peers, CB1 and CB2, are communicating through a Connector tunnel, CB1 has two Security Associations. One SA is used for processing out-bound packets, and the other SA is used for processing inbound packets.

SAs expire after a specified length of time, which is called the *lifetime*. The two peers use the Internet Key Exchange (IKE) protocol (part of the IPSec protocol suite) to negotiate new cryptographic keys and establish new SAs. The purpose of the limited lifetime is to prevent attackers from cracking a key.

Configuring CloudBridge Connector Tunnel between two Datacenters

You can configure a CloudBridge Connector tunnel between two different datacenters to extend your network without reconfiguring it, and leverage the capabilities of the two datacenters. Having a CloudBridge Connector tunnel configured between the two geographically separated datacenters enables you to implement redundancy and safeguard your setup from failure. The CloudBridge Connector tunnel helps achieve optimal utilization of infrastructure and resources across two datacenters. The applications available across the two datacenters appear as local to the user.

To connect a datacenter to another datacenter, you set up a CloudBridge Connector tunnel between a NetScaler appliance that reside in one datacenter and another NetScaler appliance that reside in the other datacenter.

As an illustration of CloudBridge Connector tunnel between two different datacenters, consider an example in which a CloudBridge Connector tunnel is set up between NetScaler appliance NS_Appliance-1 in datacenter DC1 and NetScaler appliance NS_Appliance-2 in datacenter DC2.

Both NS_Appliance-1 and NS_Appliance-2 function in L2 and L3 mode. They enable communication between private networks in datacenters DC1 and DC2. In L3 mode, NS_Appliance-1 and NS_Appliance-2 enable communication between client CL1 in the datacenter DC1 and server S1 in the datacenter DC2 through the CloudBridge Connector tunnel. Client CL1 and server S1 are on different private networks.

Because client CL1 and server S1 are on different private networks, L3 mode is enabled on NS_Appliance-1 and NS_Appliance-2 and routes are updated as the following:

- CL1 have a route to NS_Appliance-1 for reaching S1
- NS_Appliance-1 have a route to NS_Appliance-2 for reaching S1
- S1 should have a route to NS_Appliance-2 for reaching CL1
- NS_Appliance-2 have a route to NS_Appliance-1 for reaching CL1

The following table lists the settings on NetScaler appliance NS_Appliance-1 in datacenter DC1.

Entity	Name	Details
The NSIP address		198.51.100.12
SNIP address		198.51.100.15
CloudBridge Connector tunnel	Cloud_Connector_DC1-DC2	<ul style="list-style-type: none">◦ Local endpoint IP address of the CloudBridge Connector tunnel = 198.51.100.15◦ Remote endpoint IP address of the CloudBridge Connector tunnel = 203.0.113.133 <p>GRE Tunnel Details</p> <ul style="list-style-type: none">◦ Name = Cloud_Connector_DC1-DC2 <p>IPSec Profile Details</p> <ul style="list-style-type: none">◦ Name = Cloud_Connector_DC1-DC2◦ Encryption algorithm = AES◦ Hash algorithm = HMAC SHA1

The following table lists the settings on NetScaler appliance NS_Appliance-2 in datacenter DC2.

Entity	Name	Details
The NSIP address		203.0.113.131
SNIP address		203.0.113.133
CloudBridge Connector tunnel	Cloud_Connector_DC1-DC2	

	<ul style="list-style-type: none"> Local endpoint IP address of the CloudBridge Connector tunnel = 203.0.113.133 Remote endpoint IP address of the CloudBridge Connector tunnel = 198.51.100.15 <p>GRE Tunnel Details</p> <ul style="list-style-type: none"> Name = Cloud_Connector_DC1-DC2 <p>IPSec Profile Details</p> <ul style="list-style-type: none"> Name = Cloud_Connector_DC1-DC2 Encryption algorithm = AES Hash algorithm = HMAC SHA1
--	--

Following is the traffic flow in the CloudBridge Connector tunnel:

1. Client CL1 sends a request to server S1.
2. The request reaches NetScaler appliance NS-Appliance-1.
3. NS_Appliance-1, checks its routing table and finds that the destination IP address of the request packet belongs to a subnet in datacenter DC2. The appliance decides to forward the packet to be sent across the CC-DC1-DC2 tunnel.
4. NS_Appliance-1 uses the GRE protocol to encapsulate each of the request packets by adding a GRE header and a GRE IP header to the packet. The GRE IP header has the destination IP address set to the IP address of the CloudBridge tunnel (CC-DC1-DC2) end point in DC2 side. This IP Address is a public SNIP address configured on the NetScaler instance running on the NetScaler appliance NS_Appliance-2.
5. For CloudBridge Connector tunnel CC-DC1-DC2, NS_Appliance-1 checks the stored IPSec security association (SA) parameters for processing outbound packets, as agreed between NS_Appliance-1 and NS_Appliance-2 . The IPSec Encapsulating Security Payload (ESP) protocol in NS_Appliance-1 uses these SA parameters for outbound packets, to encrypt the payload of the GRE encapsulated packet.
6. The ESP protocol ensures the packet's integrity and confidentiality by using the HMAC hash function and the encryption algorithm specified for the CloudBridge Connector tunnel CC-DC1-DC2 . The ESP protocol, after encrypting the GRE payload and calculating the HMAC, generates an ESP header and an ESP trailer and inserts them before and at the end of the encrypted GRE payload, respectively.
7. The resulting packet is sent to NS_Appliance-2.
8. NS_Appliance-2 checks the stored IPSec security association (SA) parameters for processing inbound packets, as agreed between NS_Appliance-1 and NS_Appliance-2 for the CloudBridge Connector tunnel CC-DC1-DC2 . The IPSec ESP protocol on NS_Appliance-2 uses these SA parameters for inbound packets, and the ESP header of the request packet, to decrypt the packet.
9. NS_Appliance-2 then decapsulates the packet by removing the GRE header.
10. The resulting packet is the same packet as the one received by NS_Appliance-1 in step 2. This packet has the destination IP address set to the IP address of server S1. NS_Appliance-2 forwards this packet to server S1.
11. S1 processes the request packet and sends out a response packet. The destination IP address in the response packet is the IP address of client CL1, and the source IP address is the IP address of server S1.
12. The response packet reaches NS_Appliance-2.
13. NS_Appliance-2 encapsulates and encrypts the response packet in the same way that NS_Appliance-1 did with the request packet in steps 3-6.
14. NS_Appliance-2 sends the resulting packet to NS_Appliance-1.
15. NS_Appliance-1, upon receiving the packet from NS_Appliance-2, decrypts and decapsulates the packet in the same way that NS_Appliance-2 did with the request packet in steps 9-11.

Prerequisites for Configuring a CloudBridge Connector tunnel between two Datacenters

Before setting up a CloudBridge Connector tunnel, verify that the following tasks have been completed:

1. Deploy and set up a NetScaler appliance in each of the two datacenters.
2. Make sure that the CloudBridge Connector tunnel end-point IP addresses are accessible to each other.

Configuration Steps

To set up a CloudBridge Connector tunnel between a NetScaler appliance that resides in one datacenter and another NetScaler appliance that resides in the other datacenter, use the configuration utility or the command line interface of one of the NetScaler appliance.

When you use the configuration utility, the CloudBridge Connector tunnel configuration created on the first NetScaler appliance, is automatically pushed to the other endpoint (the other NetScaler appliance) of the CloudBridge Connector tunnel. Therefore, you do not have to access the configuration utility of the other NetScaler appliance to create the corresponding CloudBridge Connector tunnel configuration on it.

The CloudBridge Connector tunnel configuration on each of the NetScaler appliance consists of the following entities:

- **IPSec profile.** An IPSec profile entity specifies the IPSec protocol parameters, such as IKE version, encryption algorithm, hash algorithm, and PSK, to be used by the IPSec protocol in the CloudBridge Connector tunnel.
- **GRE tunnel.** An IP tunnel specifies the local IP address (a public SNIP address configured on the local NetScaler appliance), remote IP address (a public SNIP address configured on the remote NetScaler appliance), protocol (GRE) used to set up the CloudBridge Connector tunnel, and an IPSec profile entity.
- **Create a PBR rule and associate the IP tunnel with it** A PBR entity specifies a set of conditions and an IP tunnel entity. The source IP address range and the destination IP range are the conditions for the PBR entity. You must set the source IP address range and the destination IP address range to specify the subnet whose traffic is to traverse the CloudBridge Connector tunnel. For example, consider a request packet that originates from a client on the subnet in the first datacenter and is destined to a server on the subnet in the second datacenter. If this packet matches the source and destination IP address range of the PBR entity on the NetScaler appliance in the first datacenter, it is sent across the CloudBridge Connector tunnel associated with the PBR entity.

To create an IPSEC profile by using the command line interface

At the command prompt, type:

```
add ipsec profile <name> [-ikeVersion ( V1 | V2 )] [-encAlgo ( AES | 3DES ) ...] [-hashAlgo <hashAlgo> ...] [-lifetime <positive_integer>] [-psk | (-publickey <string> -privatekey <string> -peerPublicKey <string>)] [-livenessCheckInterval <positive_integer>] [-replayWindowSize <positive_integer>] [-ikeRetryInterval <positive_integer>] [-retransmissiontime <positive_integer>]
```

Example

```
add ipsec profile Cloud_Connector_DC1-DC2 -encAlgo AES -hashAlgo HMAC_SHA1
```

To create an IP tunnel and bind the IPSEC profile to it by using the command line interface

At the command prompt, type:

```
add ipTunnel <name> <remote> <remoteSubnetMask> <local> [-protocol <protocol>] [-ipsecProfileName <string>]
```

Example

```
add ipTunnel Cloud_Connector_DC1-DC2 203.0.113.133 255.255.255.0 198.51.100.15 "protocol GRE - ipsecProfileName Cloud_Connector_DC1-DC2
```

To create a PBR rule and bind the IPSEC tunnel to it by using the command line interface

At the command prompt, type:

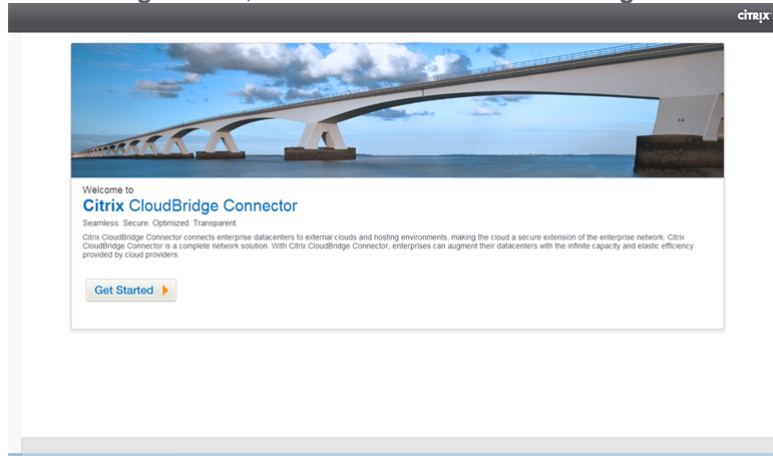
- **add ns pbr** <pbr_name> ALLOW -srcIP = <local_subnet_range> -destIP = <remote_subnet_range> - ipTunnel <tunnel_name>
- **apply ns pbrs**

Example

- **add ns pbr** PBR-DC1-DC2 ALLOW "srcIP 198.51.100.15 "destIP 203.0.113.133 ipTunnel Cloud_Connector_DC1-DC2
- **apply ns pbrs**

To configure a CloudBridge Connector tunnel in a NetScaler appliance by using the configuration utility

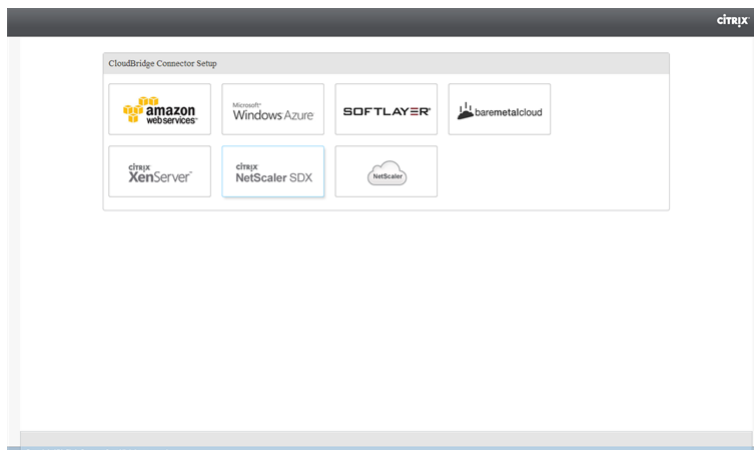
1. Type the NSIP address of a NetScaler appliance in the address line of a web browser.
2. Log on to the configuration utility of the NetScaler appliance by using your account credentials for the appliance.
3. Navigate to **System > CloudBridge Connector**.
4. In the right pane, under **Getting Started**, click **Create/Monitor CloudBridge**.



5. Click **Get Started**.

Note: If you already have any CloudBridge Connector tunnel configured on the NetScaler appliance, this screen does not appear, and you are taken to the **CloudBridge Connector Setup** pane.

6. In the **CloudBridge Connector Setup** pane, click **NetScaler**.



7. In the **NetScaler** pane, provide your account credentials for the remote NetScaler appliance. Click **Continue**.
8. In the **CloudBridge Connector Setting** pane, set the following parameter:
 - o **CloudBridge Connector Name**—Name for the CloudBridge Connector configuration on the local appliance. Must begin with an ASCII alphabetic or underscore (_) character, and must contain only ASCII alphanumeric, underscore, hash (#), period (.), space, colon (:), at (@), equals (=), and hyphen (-) characters. Cannot be changed after the CloudBridge Connector configuration is created.
9. Under **Local Setting**, set the following parameter:
 - o **Subnet IP**—IP address of the local endpoint of the CloudBridge Connector tunnel. Must be a public IP address of type SNIP.
10. Under **Remote Setting**, set the following parameter:
 - o **Subnet IP**—IP address of the peer endpoint of the CloudBridge Connector tunnel. Must be a public IP address of type SNIP.
11. Under **PBR Setting**, set the following parameters:
 - o **Operation**—Either the equals (=) or does not equal (!=) logical operator.
 - o **Source IP Low**—Lower source IP address to match against the source IP address of an outgoing IPv4 packet.
 - o **Source IP High**—Higher source IP address to match against the source IP address of an outgoing IPv4 packet.
 - o **Operation**—Either the equals (=) or does not equal (!=) logical operator.
 - o **Destination IP Low**—Lower destination IP address to match against the destination IP address of an outgoing IPv4 packet.
 - o **Destination IP High**—Higher destination IP address to match against the destination IP address of an outgoing IPv4 packet.
12. (Optional) Under **Security Settings**, set the following IPSec protocol parameters to be used by the IPSec protocol in the CloudBridge Connector tunnel:

- o **Encryption Algorithm** – Encryption algorithm to be used by the IPSec protocol in the CloudBridge tunnel.
- o **Hash Algorithm** – Hash algorithm to be used by the IPSec protocol in the CloudBridge tunnel.
- o **Key** – Select one of the following IPSec authentication methods to be used by the two peers to mutually authenticate.
 - **Auto Generate Key** – Authentication based on a text string, called a pre-shared key (PSK), generated automatically by the local appliance. The PSKs keys of the peers are matched against each other for authentication.
 - **Specific Key** – Authentication based on a manually entered PSK. The PSKs of the peers are matched against each other for authentication.
 - **Pre Shared Security Key** – The text string entered for pre-shared key based authentication.
 - **Upload Certificates** – Authentication based on digital certificates.
 - **Public Key** – A local digital certificate to be used to authenticate the local NetScaler appliance to the peer before establishing IPSec security associations. The same certificate should be present and set for the Peer Public Key parameter in the peer.
 - **Private Key** – Private key of the local digital certificate.
 - **Peer Public Key** – Digital certificate of the peer. Used to authenticate the peer to the local end point before establishing IPSec security associations. The same certificate should be present and set for the Public key parameter in the peer.

13. Click **Done**.

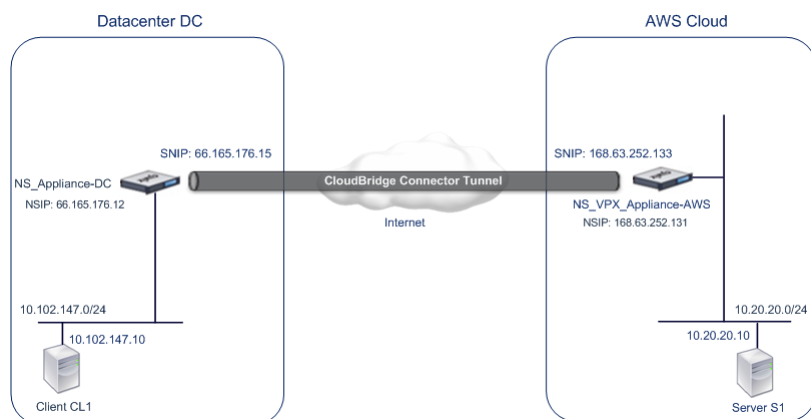
The new CloudBridge Connector tunnel configuration on both the NetScaler appliances appears on the Home tab of the respective configuration utility. The current status of the CloudBridge connector tunnel is indicated in the Configured CloudBridge Connectors pane. A green dot indicates that the tunnel is up. A red dot indicates that the tunnel is down.

Configuring CloudBridge Connector between Datacenter and AWS Cloud

You can configure a CloudBridge Connector tunnel between a datacenter and AWS cloud to leverage the infrastructure and computing capabilities of the data center and the AWS cloud. With AWS, you can extend your network without initial capital investment or the cost of maintaining the extended network infrastructure. You can scale your infrastructure up or down, as required. For example, you can lease more server capabilities when the demand increases.

To connect a datacenter to AWS cloud, you set up a CloudBridge Connector tunnel between a NetScaler appliance that resides in the datacenter and a NetScaler virtual appliance (VPX) that resides in AWS cloud.

As an illustration of a CloudBridge Connector tunnel between a datacenter and Amazon AWS cloud, consider an example in which a CloudBridge Connector tunnel is set up between NetScaler appliance NS_Appliance-DC, in datacenter DC, and NetScaler virtual appliance (VPX) NS_VPX_Appliance-AWS.



Both NS_Appliance-DC and NS_VPX_Appliance-AWS function in L3 mode. They enable communication between private networks in datacenter DC and the AWS cloud. NS_Appliance-DC and NS_VPX_Appliance-AWS enable communication between client CL1 in datacenter DC and server S1 in the AWS cloud through the CloudBridge Connector tunnel. Client CL1 and server S1 are on different private networks.

Note: AWS does not support L2 mode, hence it is necessary to have only L3 mode enabled on both the endpoints.

For proper communication between CL1 and S1, L3 mode is enabled on NS_Appliance-DC and NS_VPX_Appliance-AWS and routes are updated as such:

- CL1 have a route to NS_Appliance-DC for reaching S1.
- NS_Appliance-DC have a route to NS_VPX_Appliance-AWS for reaching S1.
- S1 should have a route to NS_VPX_Appliance-AWS for reaching CL1.
- NS_VPX_Appliance-AWS have a route to NS_Appliance-DC for reaching CL1.

The following table lists the settings on NetScaler appliance NS_Appliance-DC in datacenter DC.

Entity	Name	Details
The NSIP address		66.165.176.12
SNIP address		66.165.176.15
CloudBridge Connector tunnel	CC_Tunnel_DC-AWS	<ul style="list-style-type: none"> ◦ Local endpoint IP address of the CloudBridge Connector tunnel: 66.165.176.15 ◦ Remote endpoint IP address of the CloudBridge Connector tunnel: 168.63.252.133 <p>GRE Tunnel Details</p> <ul style="list-style-type: none"> ◦ Name= CC_Tunnel_DC-AWS <p>IPSec Profile Details</p> <ul style="list-style-type: none"> ◦ Name= CC_Tunnel_DC-AWS ◦ Encryption algorithm= AES ◦ Hash algorithm= HMAC SHA1

The following table lists the settings on NetScaler VPX NS_VPX_Appliance-AWS on AWS cloud.

Entity	Name	Details
NSIP address		10.102.25.30
Public EIP address mapped to the NSIP address		168.63.252.131
SNIP address		10.102.29.30
Public EIP address mapped to the SNIP address		168.63.252.133
CloudBridge Connector tunnel	CC_Tunnel_DC-AWS	<ul style="list-style-type: none"> Local endpoint IP address of the CloudBridge Connector tunnel: 168.63.252.133 Remote endpoint IP address of the CloudBridge Connector tunnel: 66.165.176.15 <p>GRE Tunnel Details</p> <ul style="list-style-type: none"> Name= CC_Tunnel_DC-AWS <p>IPSec Profile Details</p> <ul style="list-style-type: none"> Name= CC_Tunnel_DC-AWS Encryption algorithm= AES Hash algorithm= HMAC SHA1

Prerequisites

Updated: 2015-06-01

Before setting up a CloudBridge Connector tunnel, verify that the following tasks have been completed:

1. Install, configure, and launch an instance of NetScaler Virtual appliance (VPX) on AWS cloud. For instructions on installing NetScaler VPX on AWS, see <http://support.citrix.com/proddocs/topic/netscaler-vpx-10-5/nsvpx-aws-ns-vpxaws-con.html>.
2. Deploy and configure a NetScaler physical appliance, or provisioning and configuring a NetScaler virtual appliance (VPX) on a virtualization platform in the datacenter.
 - For instructions on installing NetScaler virtual appliances on Xenserver, see <http://support.citrix.com/proddocs/topic/netscaler-vpx-10-5/ns-vpx-install-wrapper-con.html>.
 - For instructions on installing NetScaler virtual appliances on VMware ESX or ESXi, see <http://support.citrix.com/proddocs/topic/netscaler-vpx-10-5/ns-vpx-install-on-esx-wrapper-con.html>.
 - For instructions on installing NetScaler virtual appliances on Microsoft Hyper-V, see <http://support.citrix.com/proddocs/topic/netscaler-vpx-10-5/ns-vpx-install-on-msft-hyperv-wrapper-con.html>.
3. Make sure that the CloudBridge Connector tunnel end-point IP addresses are accessible to each other.

NetScaler VPX License

After the initial instance launch, NetScaler VPX for AWS requires a license. If you are bringing your own license (BYOL), see the VPX Licensing Guide at: <http://support.citrix.com/article/CTX122426>.

You have to:

1. Use the licensing portal within MyCitrix to generate a valid license.
2. Upload the license to the instance.

If this is a **paid** marketplace instance, then you do not need to install a license. The correct feature set and performance will activate automatically.

Configuration Steps

To set up a CloudBridge Connector tunnel between a NetScaler appliance that resides in a datacenter and a NetScaler virtual appliance (VPX) that resides on the AWS cloud, use the configuration utility of the NetScaler appliance.

When you use the configuration utility, the CloudBridge Connector tunnel configuration created on the NetScaler appliance, is automatically pushed to the other endpoint or peer (the NetScaler VPX on AWS) of the CloudBridge Connector tunnel. Therefore, you do not have to access the configuration utility (GUI) of the NetScaler VPX on AWS to create the corresponding CloudBridge Connector tunnel configuration on it.

The CloudBridge Connector tunnel configuration on both peers (the NetScaler appliance that resides in the datacenter and the NetScaler virtual appliance (VPX) that resides on the AWS cloud) consists of the following entities:

- **IPSec profile**—An IPSec profile entity specifies the IPSec protocol parameters, such as IKE version, encryption algorithm, hash algorithm, and PSK, to be used by the IPSec protocol in both the peers of the CloudBridge Connector tunnel.
- **GRE tunnel**—An IP tunnel specifies a local IP address (a public SNIP address configured on the local peer), remote IP address (a public SNIP address configured on the remote peer), protocol (GRE) used to set up the CloudBridge Connector tunnel, and an IPSec profile entity.
- **Create a PBR rule and associate the IP tunnel with it**—A PBR entity specifies a set of conditions and an IP tunnel entity. The source IP address range and the destination IP range are the conditions for the PBR entity. You must set the source IP address range and the destination IP address range to specify the subnet whose traffic is to traverse the CloudBridge Connector tunnel. For example, consider a request packet that originates from a client on the subnet in the datacenter and is destined to a server on the subnet in the AWS cloud. If this packet matches the source and destination IP address range of the PBR entity on the NetScaler appliance in the datacenter, it is sent across the CloudBridge Connector tunnel associated with the PBR entity.

To create an IPSEC profile by using the command line interface

At the command prompt, type:

```
add ipsec profile <name> [-ikeVersion ( V1 | V2 )] [-encAlgo ( AES | 3DES ) ...] [-hashAlgo <hashAlgo> ...] [-lifetime <positive_integer>] [-psk | (-publickey <string> -privatekey <string> -peerPublicKey <string>)] [-livenessCheckInterval <positive_integer>] [-replayWindowSize <positive_integer>] [-ikeRetryInterval <positive_integer>] [-retransmissiontime <positive_integer>]
```

Example

```
add ipsec profile CC_Tunnel_DC-AWS -encAlgo AES -hashAlgo HMAC_SHA1
```

To create an IP tunnel and bind the IPSEC profile to it by using the command line interface

At the command prompt, type:

```
add ipTunnel <name> <remote><remoteSubnetMask> <local> [-protocol <protocol>] [-ipsecProfileName <string>]
```

Example

```
add ipTunnel CC_Tunnel_DC-AWS 168.63.252.133 255.255.255.0 66.165.176.15 "protocol GRE -i
```

To create a PBR rule and bind the IPSEC tunnel to it by using the command line interface

At the command prompt, type:

- **add ns pbr** <pbr_name> ALLOW -srcIP = <local_subnet_range> -destIP = <remote_subnet_range> -ipTunnel <tunnel_name>
- **apply ns pbrs**

Example

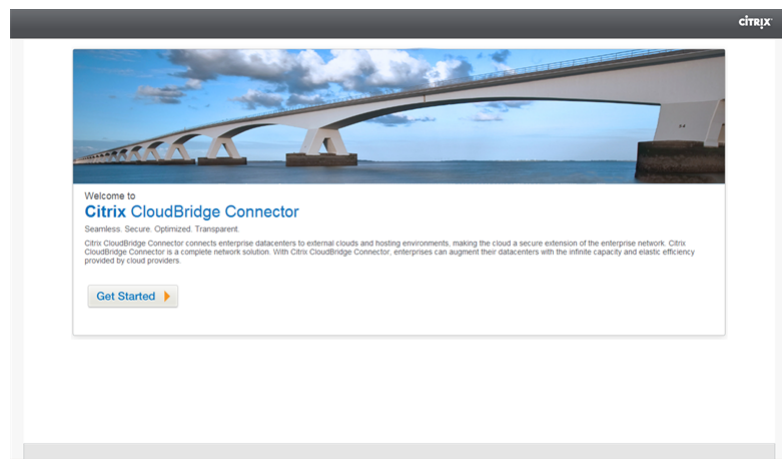
- **add ns pbr** PBR-DC-AWS ALLOW "srcIP 66.165.176.15 "destIP 168.63.252.133 ipTunnel CC_Tunnel_DC-AWS
- **apply ns pbrs**

To configure a CloudBridge Connector tunnel in a NetScaler appliance by using the configuration utility

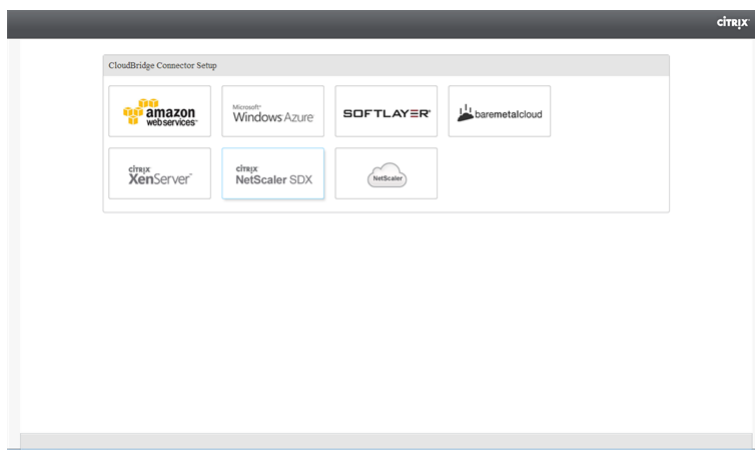
1. Type the NSIP address of a NetScaler appliance in the address line of a web browser.
2. Log on to the configuration utility of the NetScaler appliance by using your account credentials for the appliance.

3. Navigate to System > CloudBridge Connector.
4. In the right pane, under **Getting Started**, click **Create/Monitor CloudBridge**.

The first time you configure a CloudBridge Connector tunnel on the appliance, a Welcome screen appears.



5. On the Welcome screen click **Get Started**.
Note: If you already have a CloudBridge Connector tunnel configured on the NetScaler appliance, the Welcome screen does not appear, so you do not click **Get Started**.
6. In the **CloudBridge Connector Setup** pane, click **amazon web services**.



7. In the **Amazon** pane, provide your AWS account credentials: AWS Access Key ID and AWS Secret Access Key. You can obtain these access keys from the AWS GUI console. Click **Continue**.
8. In the **NetScaler** pane, select the NSIP address of the NetScaler virtual appliance running on AWS. Then, provide your account credentials for the NetScaler virtual appliance. Click **Continue**.
9. In the **CloudBridge Connector Setting** pane, set the following parameter:
 - o **CloudBridge Connector Name**—Name for the CloudBridge Connector configuration on the local appliance. Must begin with an ASCII alphabetic or underscore (_) character, and must contain only ASCII alphanumeric, underscore, hash (#), period (.), space, colon (:), at (@), equals (=), and hyphen (-) characters. Cannot be changed after the CloudBridge Connector configuration is created.
10. Under **Local Setting**, set the following parameter:
 - o **Subnet IP**—IP address of the local endpoint of the CloudBridge Connector tunnel. Must be a public IP address of type SNIP.
11. Under **Remote Setting**, set the following parameter:
 - o **Subnet IP**—IP address of the CloudBridge Connector tunnel end point on the AWS side. Must be an IP address of type SNIP on the NetScaler VPX instance on AWS.
 - o **NAT**—Public IP address (EIP) in AWS that is mapped to the SNIP configured on the NetScaler VPX instance on AWS.
12. Under **PBR Setting**, set the following parameters:
 - o **Operation**—Either the equals (=) or does not equal (!=) logical operator.
 - o **Source IP Low**—Lowest source IP address to match against the source IP address of an outgoing IPv4 packet.
 - o **Source IP High**—Highest source IP address to match against the source IP address of an outgoing IPv4 packet.
 - o **Operation**—Either the equals (=) or does not equal (!=) logical operator.

- o **Destination IP Low** – Lowest destination IP address to match against the destination IP address of an outgoing IPv4 packet.
 - o **Destination IP High** – Highest destination IP address to match against the destination IP address of an outgoing IPv4 packet.
13. (Optional) Under **Security Settings**, set the following IPSec protocol parameters for the CloudBridge Connector tunnel:
- o **Encryption Algorithm** – Encryption algorithm to be used by the IPSec protocol in the CloudBridge tunnel.
 - o **Hash Algorithm** – Hash algorithm to be used by the IPSec protocol in the CloudBridge tunnel.
 - o **Key** – Select one of the following IPSec authentication methods to be used by the two peers to mutually authenticate.
 - **Auto Generate Key** – Authentication based on a text string, called a pre-shared key (PSK), generated automatically by the local appliance. The PSKs keys of the peers are matched against each other for authentication.
 - **Specific Key** – Authentication based on a manually entered PSK. The PSKs of the peers are matched against each other for authentication.
 - **Pre Shared Security Key** – The text string entered for pre-shared key based authentication.
 - **Upload Certificates** – Authentication based on digital certificates.
 - **Public Key** – A local digital certificate to be used to authenticate the local peer to the remote peer before establishing IPSec security associations. The same certificate should be present and set for the Peer Public Key parameter in the peer.
 - **Private Key** – Private key of the local digital certificate.
 - **Peer Public Key** – Digital certificate of the peer. Used to authenticate the peer to the local end point before establishing IPSec security associations. The same certificate should be present and set for the Public key parameter in the peer.
14. Click **Done**.

The new CloudBridge Connector tunnel configuration on the NetScaler appliance in the datacenter appears on the Home tab of the configuration utility.

The corresponding new CloudBridge Connector tunnel configuration on the NetScaler VPX appliance in the AWS cloud appears on the configuration utility.

The current status of the CloudBridge connector tunnel is indicated in the Configured CloudBridge pane. A green dot indicates that the tunnel is up. A red dot indicates that the tunnel is down.

Configuring a CloudBridge Connector Tunnel Between a Datacenter and Azure Cloud

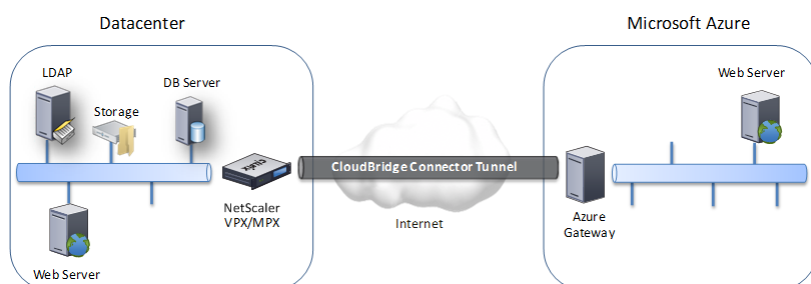
The NetScaler appliance provides connectivity between your enterprise datacenters and the Microsoft cloud hosting provider, Azure, making Azure a seamless extension of the enterprise network. NetScaler encrypts the connection between the enterprise datacenter and Azure cloud so that all data transferred between the two is secure.

This section includes the following:

- How CloudBridge Connector Tunnel Works
- Example of CloudBridge Connector Tunnel Configuration and Data Flow
- Points to Consider for a CloudBridge Connector tunnel Configuration
- Configuring the CloudBridge Connector Tunnel
- Monitoring the CloudBridge Connector Tunnel

How CloudBridge Connector Tunnel Works

To connect a datacenter to AWS cloud, you set up a CloudBridge Connector tunnel between a NetScaler appliance that resides in the datacenter and a gateway that resides in the Azure cloud. The NetScaler appliance in the datacenter and the gateway in Azure cloud are the end points of the CloudBridge Connector tunnel and are called peers of the CloudBridge Connector tunnel.



A CloudBridge Connector tunnel between a datacenter and Azure cloud uses the open-standard Internet Protocol security (IPSec) protocol suite, in tunnel mode, to secure communications between peers in the CloudBridge Connector tunnel. In a CloudBridge Connector tunnel, IPSec ensures:

- Data integrity
- Data origin authentication
- Data confidentiality (encryption)
- Protection against replay attacks

IPSec uses the tunnel mode in which the complete IP packet is encrypted and then encapsulated. The encryption uses the Encapsulating Security Payload (ESP) protocol, which ensures the integrity of the packet by using a HMAC hash function and ensures confidentiality by using an encryption algorithm. The ESP protocol, after encrypting the payload and calculating the HMAC, generates an ESP header and inserts it before the encrypted IP packet. The ESP protocol also generates an ESP trailer and inserts it at the end of the packet.

The IPSec protocol then encapsulates the resulting packet by adding an IP header before the ESP header. In the IP header, the destination IP address is set to the IP address of the CloudBridge Connector peer.

Peers in the CloudBridge Connector tunnel use the Internet Key Exchange version 1 (IKEv1) protocol (part of the IPSec protocol suite) to negotiate secure communication, as follows:

1. The two peers mutually authenticate with each other, using pre-shared key authentication, in which the peers exchange a text string called a pre-shared key (PSK). The pre-shared keys are matched against each other for authentication. Therefore, for the authentication to be successful, you must configure the same pre-shared key on each of the peers.
2. The peers then negotiate to reach agreement on:
 - An encryption algorithm
 - Cryptographic keys for encrypting data on one peer and decrypting it on the other.

This agreement upon the security protocol, encryption algorithm and cryptographic keys is called a Security Association (SA). SAs are one-way (simplex). For example, when a CloudBridge Connector tunnel is set up between a NetScaler appliance in a datacenter and a gateway in an Azure cloud, both the datacenter appliance and the Azure gateway have two SAs. One SA is used for processing out-bound packets, and the other SA is used for processing inbound packets. SAs expire after a specified interval of time, which is called the lifetime.

Example of CloudBridge Connector Tunnel Configuration and Data Flow

As an illustration of CloudBridge Connector Tunnel, consider an example in which a CloudBridge Connector tunnel is set up between NetScaler appliance CB_Appliance-1 in a datacenter and gateway Azure_Gateway-1 in Azure cloud.

CB_Appliance-1 also functions as an L3 router, which enables a private network in the datacenter to reach a private network in the Azure cloud through the CloudBridge Connector tunnel. As a router, CB_Appliance-1 enables communication between client CL1 in the datacenter and server S1 in the Azure cloud through the CloudBridge Connector tunnel. Client CL1 and server S1 are on different private networks.

On CB_Appliance-1, the CloudBridge Connector tunnel configuration includes an IPSec profile entity named CB_Azure_IPSec_Profile, a CloudBridge Connector tunnel entity named CB_Azure_Tunnel, and a policy based routing (PBR) entity named CB_Azure_Pbr.

The IPSec profile entity CB_Azure_IPSec_Profile specifies the IPSec protocol parameters, such as IKE version, encryption algorithm, and hash algorithm, to be used by the IPSec protocol in the CloudBridge Connector tunnel. CB_Azure_IPSec_Profile is bound to IP tunnel entity CB_Azure_Tunnel.

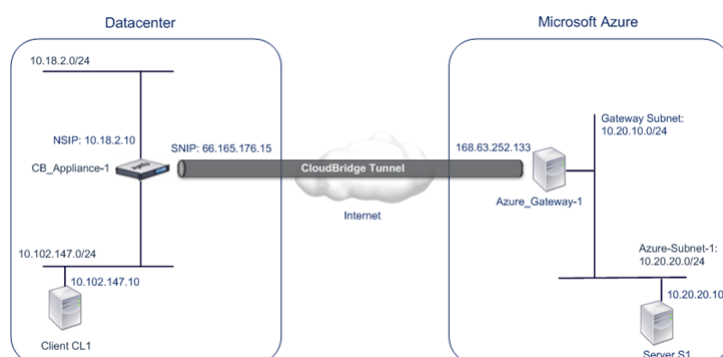
CloudBridge Connector tunnel entity CB_Azure_Tunnel specifies the local IP address (a public IP (SNIP) address configured on the NetScaler appliance), the remote IP address (the IP address of the Azure_Gateway-1), and the protocol (IPSec) used to set up the CloudBridge Connector tunnel. CB_Azure_Tunnel is bound to the PBR entity CB_Azure_Pbr.

The PBR entity CB_Azure_Pbr specifies a set of conditions and a CloudBridge Connector tunnel entity (CB_Azure_Tunnel). The source IP address range and the destination IP address range are the conditions for CB_Azure_Pbr. The source IP address range and the destination IP address range are specified as a subnet in the datacenter and a subnet in the Azure cloud, respectively. Any request packet originating from a client in the subnet in the datacenter and destined to a server in the subnet on the Azure cloud matches the conditions in CB_Azure_Pbr. This packet is then considered for CloudBridge processing and is sent across the CloudBridge Connector tunnel (CB_Azure_Tunnel) bound to the PBR entity.

On Microsoft Azure, the CloudBridge Connector tunnel configuration includes a local network entity named My-Datacenter-Network, a virtual network entity named Azure-Network-for-CloudBridge-Tunnel, and a gateway named Azure_Gateway-1.

The local (local to Azure) network entity My-Datacenter-Network specifies the IP address of the NetScaler appliance on the datacenter side, and the datacenter subnet whose traffic is to traverse the CloudBridge Connector tunnel. The virtual network entity Azure-Network-for-CloudBridge-Tunnel defines a private subnet named Azure-Subnet-1 in Azure. The traffic of the subnet traverses the CloudBridge Connector tunnel. The server S1 is provisioned in this subnet.

The local network entity My-Datacenter-Network is associated with the virtual network entity Azure-Network-for-CloudBridge-Tunnel. This association defines the remote and local network details of the CloudBridge Connector tunnel configuration in Azure. Gateway Azure_Gateway-1 was created for this association to become the CloudBridge end point at the Azure end of the CloudBridge Connector tunnel.



The following table lists the settings used in this example.

Entity	Name	Details
Settings highlight of the CloudBridge Connector tunnel setup		
IP address of the CloudBridge Connector tunnel end point (CB_Appliance-1) in the datacenter side	66.165.176.15	
IP address of the CloudBridge Connector tunnel end point	168.63.252.133	

(Azure_Gateway-1) in the Azure		
Datacenter Subnet, the traffic of which is to traverse the CloudBridge Connector tunnel	10.102.147.0/24	
Azure Subnet, the traffic of which is to traverse the CloudBridge Connector tunnel	10.20.0.0/16	
^		
Settings on NetScaler appliance CB_Appliance-1 in Datacenter		
^	SNIP1(for reference purposes only)	66.165.176.15
IPSec profile	CB_Azure_IPSec_Profile	<ul style="list-style-type: none"> o IKE version = v1 o Encryption algorithm = AES o Hash algorithm = HMAC SHA1
CloudBridge Connector tunnel	CB_Azure_Tunnel	<ul style="list-style-type: none"> o Remote IP = 168.63.252.133 o Local IP= 66.165.176.15 o Tunnel protocol = IPSec o IPSec profile= CB_Azure_IPSec_Profile
Policy based route	CB_Azure_Pbr	<ul style="list-style-type: none"> o Source IP range = Subnet in the datacenter =10.102.147.0-10.102.147.255 o Destination IP range =Subnet in Azure =10.20.0.0-10.20.255.255 o IP Tunnel = CB_Azure_Tunnel
^		
Settings on Microsoft Azure		
Public IP Address of the Azure_Gateway-1	^	168.63.252.133
Local Network	My-Datacenter-Network	<ul style="list-style-type: none"> o VPN Device IP address =SNIP address of the NetScaler appliance = 66.165.176.15 o Address space= Subnet in datacenter =10.102.147.0/24
Virtual Network	Azure-Network-for-CloudBridge-Tunnel	<ul style="list-style-type: none"> o Address Space= 10.20.0.0/16 o Subnet in Azure=Azure-Subnet-1= 10.20.20.0/24 o Local Network=My-Datacenter-Network o Gateway Subnet=10.20.10.0/24

Following is the traffic flow in the CloudBridge Connector tunnel:

1. Client C1 sends a request to server S1.
2. The request reaches NetScaler appliance CB_Appliance-1.
3. The request packet in CB_Appliance-1 matches the condition specified in the PBR entity CB_Azure_Pbr as the source IP address and the destination IP address of the request packet belonging to the source IP range and destination IP range, respectively, set in CB_Azure_Pbr.
4. Because CloudBridge Connector tunnel entity CB_Azure_Tunnel is bound to CB_Azure_Pbr, the appliance prepares the packet to be sent across the CB_Azure_Tunnel.

5. For CloudBridge Connector tunnel CB_Azure_Tunnel, CB_Appliance-1 checks the stored IPSec security association (SA) parameters for processing outbound packets, as agreed between CB_Appliance-1 in the datacenter and Azure_Gateway-1 in the Azure cloud. The IPSec Encapsulating Security Payload (ESP) protocol in the NetScaler appliance uses these SA parameters for outbound packets to encrypt the request packet.
6. The ESP protocol ensures the packet's integrity by using a HMAC hash function and the packet's confidentiality by using the AES encryption algorithm. The ESP protocol, after encrypting the request packet and calculating the HMAC, generates an ESP header and then inserts it before the encrypted IP packet. The ESP protocol also generates an ESP trailer and then inserts it at the end of the encrypted IP packet.
7. The IPSec protocol encapsulates the resulting packet by adding an IP header before the ESP header. The destination address in the IP header is the IP address of Azure-gateway-1, and the source address is the SNIP2 address.
8. The resulting packet is sent to Azure_Gateway-1. There is
9. Azure-gateway-1, upon receiving the packet from CB_Appliance-1, decapsulates the packet by removing the IPSec IF header.
10. Azure-gateway-1 then checks the stored IPSec security association (SA) parameters for processing inbound packets, as agreed between CB_Appliance-1 and Azure_Gateway-1. The IPSec ESP protocol on Azure_Gateway-1 uses these SA parameters for inbound packets, and the ESP header of the decapsulated request packet, to decrypt the packet.
11. The resulting packet is the same packet as the one received by CB_Appliance-1 in step 2. This packet has the destination IP address set to the IP address of server S1. Azure_Gateway-1 forwards this packet to server S1.
12. S1 processes the request packet and sends out a response packet. The destination IP address in the response packet is the IP address of client CL1, and source IP address is the IP address of server S1.
13. The response packet reaches Azure_Gateway-1. Microsoft Azure checks the stored IPSec security association (SA) parameters for processing outbound packets, as agreed between CB_Appliance-1 and Azure_Gateway-1. Microsoft Azure encrypts and encapsulates the response packet in the same way that CB_Appliance-1 encrypted and encapsulated the request packet in steps 5, 6, and 7.
14. Azure_Gateway-1 sends the resulting packet to CB_Appliance-1.
15. CB_Appliance-1, upon receiving the packet from Azure_Gateway-1, decapsulates and decrypts the packet in the same way that Azure_Gateway-1 decapsulated and decrypted the request packet in steps 9 and 10.
16. The resulting packet is the same packet that was received by Azure_Gateway-1 in step 13. This response packet has the destination IP address set to the IP address of server CL1. CB_Appliance-1 forwards the response packet to client CL1.

Points to Consider for a CloudBridge Connector tunnel Configuration

Updated: 2014-04-15

Before configuring a CloudBridge Connector tunnel between a NetScaler appliance in datacenter and Microsoft Azure, consider the following points:

1. The NetScaler appliance must have a public facing IPv4 address (type SNIP) to use as a tunnel end-point address for the CloudBridge Connector tunnel. Also, the NetScaler appliance should not be behind a NAT device.
2. Azure supports the following IPSec settings for a CloudBridge Connector tunnel. Therefore, you must specify the same IPSec settings while configuring the NetScaler for the CloudBridge Connector tunnel.
 - o IKE version = v1
 - o Encryption algorithm = AES
 - o Hash algorithm = HMAC SHA1
3. You must configure the firewall in the datacenter edge to allow the following.
 - o Any UDP packets for port 500
 - o Any UDP packets for port 4500
 - o Any ESP (IP protocol number 50) packets
4. IKE re-keying, which is renegotiation of new cryptographic keys between the CloudBridge Connector tunnel end points to establish new SAs, is not supported. When the Security Associations (SAs) expire, the tunnel goes into the DOWN state. Therefore, you must set a very large value for the lifetimes of SAs.
5. You must configure Microsoft Azure before specifying the tunnel configuration on the NetScaler, because the public IP address of the Azure end (gateway) of the tunnel, and the PSK, are automatically generated when you set up the tunnel configuration in Azure. You need this information for specifying the tunnel configuration on the NetScaler.

Configuring the CloudBridge Connector Tunnel

Updated: 2014-04-15

For setting up a CloudBridge Connector tunnel between your datacenter and Azure, you must install CloudBridge VPX/MPX in your datacenter, configure Microsoft Azure for the CloudBridge Connector tunnel, and then configure the NetScaler appliance in the data center for the CloudBridge Connector tunnel.

Configuring a CloudBridge Connector tunnel between a NetScaler appliance in datacenter and Microsoft Azure consists of the following tasks:

1. **Setting up the NetScaler appliance in the datacenter.** This task involves deploying and configuring a NetScaler physical appliance (MPX), or provisioning and configuring a NetScaler virtual appliance (VPX) on a virtualization platform in the datacenter.
2. **Configuring Microsoft Azure for the CloudBridge Connector tunnel.** This task involves creating local network, virtual network, and gateway entities in Azure. The local network entity specifies the IP address of the CloudBridge Connector tunnel end point (the NetScaler appliance) on the datacenter side, and the datacenter subnet whose traffic is to traverse the CloudBridge Connector tunnel. The virtual network defines a network on Azure. Creating the virtual network includes defining a subnet whose traffic is to traverse the CloudBridge Connector tunnel to be formed. You then associate the local network with the virtual network. Finally, you create a gateway that becomes the end point at the Azure end of the CloudBridge Connector tunnel.
3. **Configuring the NetScaler appliance in the datacenter for the CloudBridge Connector tunnel.** This task involves creating an IPSec profile, an IP tunnel entity, and a PBR entity in the NetScaler appliance in datacenter. The IPSec profile entity specifies the IPSec protocol parameters, such as IKE version, encryption algorithm, hash algorithm, and PSK, to be used in the CloudBridge Connector tunnel. The IP tunnel specifies the IP address of both the CloudBridge Connector tunnel end points (the NetScaler appliance in datacenter and the gateway in Azure) and the protocol to be used in the CloudBridge Connector tunnel. You then associate the IPSec profile entity with the IP tunnel entity. The PBR entity specifies the two subnets, in the datacenter and in the Azure cloud, that are to communicate with each other through the CloudBridge Connector tunnel. You then associate the IP tunnel entity with the PBR entity.

Configuring Microsoft Azure for the CloudBridge Connector tunnel

Updated: 2014-04-15

To create a CloudBridge Connector tunnel configuration on Microsoft Azure, use the Microsoft Windows Azure Management Portal, which is a web based graphical interface for creating and managing resources on Microsoft Azure.

Before you begin the CloudBridge Connector tunnel configuration on Azure cloud, make sure that:

- o You have a user account for Microsoft Azure.
- o You have a conceptual understanding of Microsoft Azure.
- o You are familiar with the Microsoft Windows Azure Management Portal.

Note: The procedures for configuring Microsoft Azure for a CloudBridge Connector tunnel might change over time, depending on the Microsoft Azure release cycle. Citrix recommends the following Microsoft Azure documentation for the latest procedures.

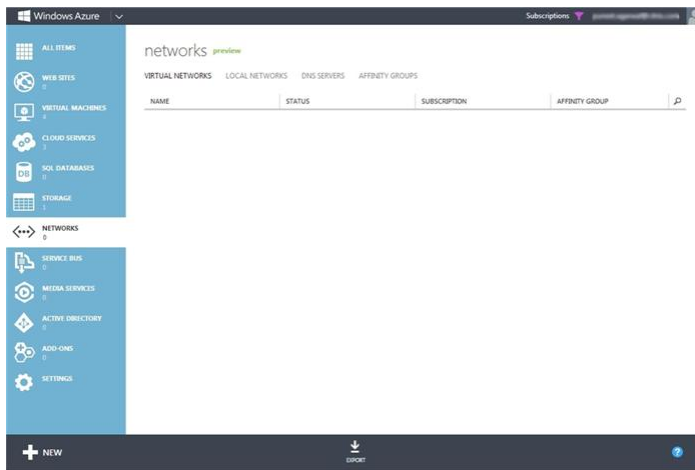
- o <http://www.windowsazure.com/en-us/manage/services/networking/cross-premises-connectivity/>

To configure a CloudBridge Connector tunnel between a datacenter and an Azure cloud, perform the following tasks on Microsoft Azure by using the Microsoft Windows Azure Management Portal:

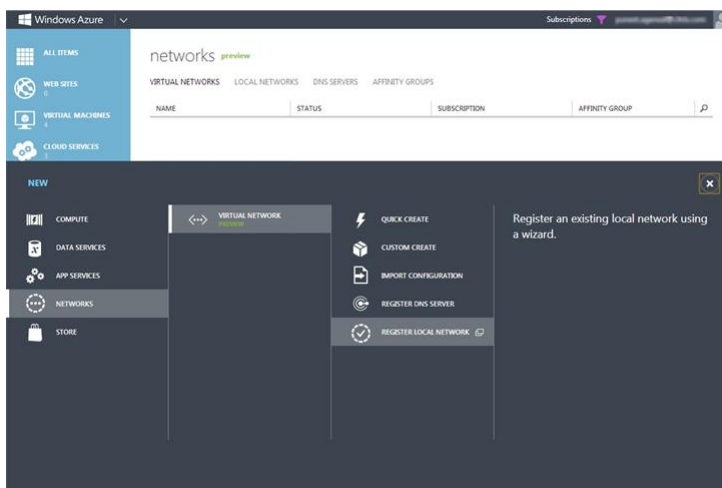
- o **Create a local network entity.** Create a local network entity in Windows Azure for specifying the network details of the datacenter. A local network entity specifies the IP address of the CloudBridge Connector tunnel end point (the NetScaler) on the datacenter side and the datacenter subnet whose traffic is to traverse the CloudBridge Connector tunnel.
- o **Create a Virtual Network.** Create virtual network entity that defines a network on Azure. This task includes defining a private address space, where you provide a range of private addresses and subnets belonging to the range specified in the address space. The traffic of the subnets will traverse the CloudBridge Connector tunnel. You then associate a local network entity with the virtual network entity. This association lets Azure create a configuration for a CloudBridge Connector tunnel between the virtual network and the data center network. A gateway (to be created) in Azure for this virtual network will be the CloudBridge end point at the Azure end of the CloudBridge Connector tunnel. You then define a private subnet for the gateway to be created. This subnet belongs to the range specified in the address space in the virtual network entity.
- o **Create a gateway in Windows Azure.** Create a gateway that becomes the end point at the Azure end of the CloudBridge Connector tunnel. Azure, from its pool of public IP addresses, assigns an IP address to the gateway created.
- o **Gather the public IP address of the gateway and the pre-shared key.** For a CloudBridge Connector tunnel configuration on Azure, the public IP address of the gateway and the pre-shared Key (PSK) are automatically generated by Azure. Make a note of this information. You will need it for configuring the CloudBridge Connector tunnel on the NetScaler in datacenter.

To specify a local network by using the Microsoft Windows Azure Management Portal

1. In the left pane, click NETWORKS.
2. In the lower left-hand corner of the screen, click + NEW.



3. In the NEW navigation pane, click NETWORK, then click VIRTUAL NETWORK, and then click REGISTER LOCAL NETWORK.



4. In the ADD A LOCAL NETWORK wizard, in the specify your local network details screen, set the following parameters:
 - NAME
 - VPN DEVICE IP ADDRESS

ADD A LOCAL NETWORK

Specify your local network details

NAME

My-Datacenter-Network

VPN DEVICE IP ADDRESS

66.165.176.15

→

2

5. In the lower right corner of the screen, click -> (forward arrow mark).
6. On the Specify the address space screen, set the following parameter:
 - ADDRESS SPACE

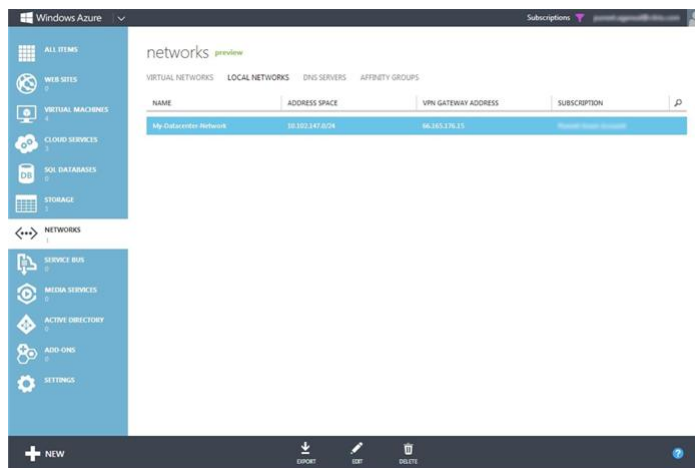
EDIT LOCAL NETWORK

Specify the address space

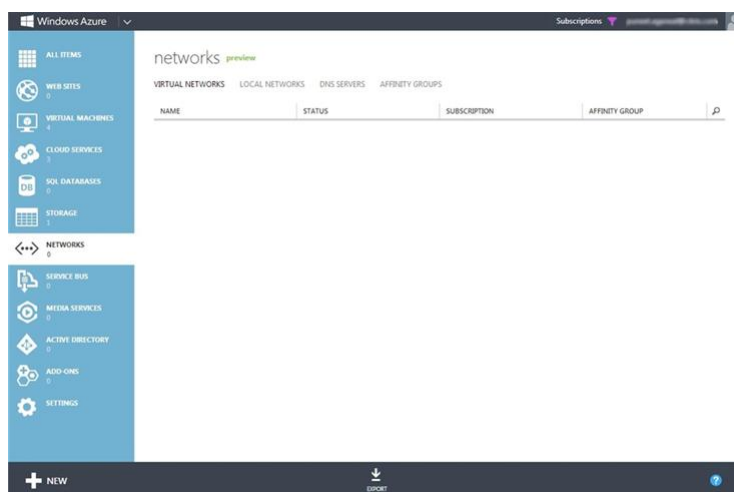
ADDRESS SPACE

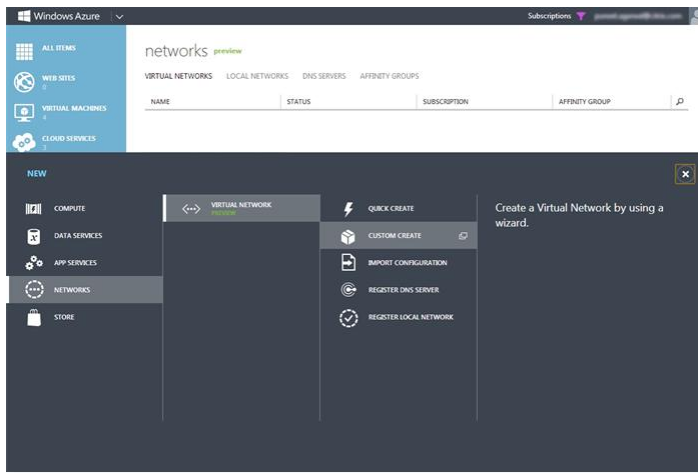
10.102.147.0/24

+



1. In the left pane, click NETWORKS.
2. In the lower left-hand corner of the screen, click + New.





4. In the CREATE A VIRTUAL NETWORK wizard, in the Virtual Network Details screen, set the following parameters:

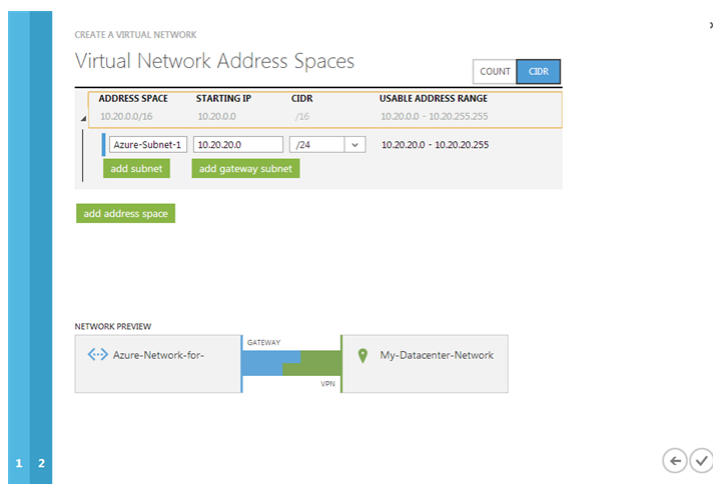
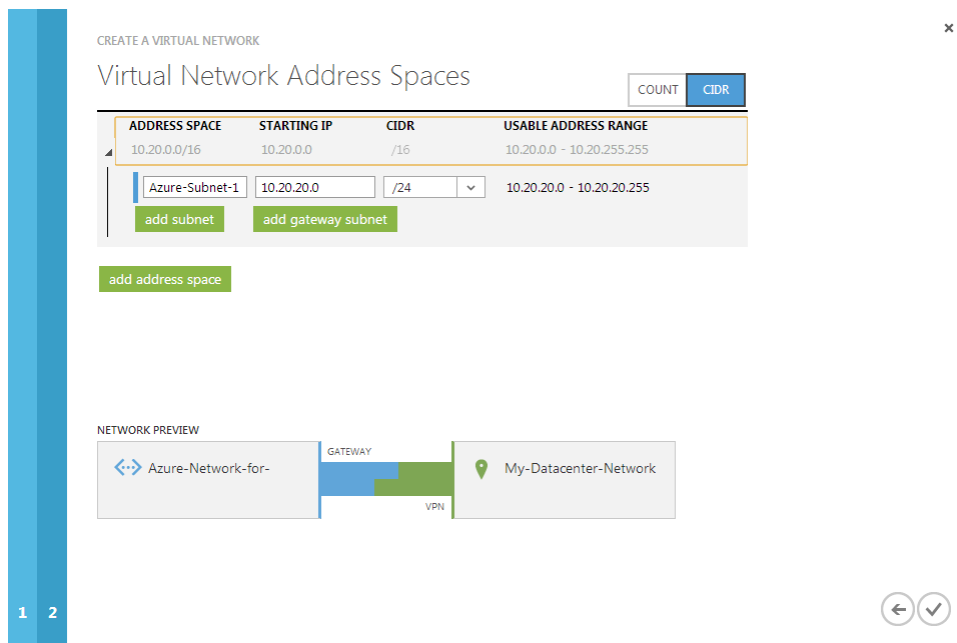
- NAME
- AFFINITY GROUP
- REGION
- AFFINITY GROUP NAME

5. Click -> (forward arrow mark) in the lower right-hand corner of the screen.
6. In the DNS Servers and VPN Connectivity screen, in SITE-TO-SITE CONNECTIVITY, select Configure Site-To-Site VPN and set the following parameter:

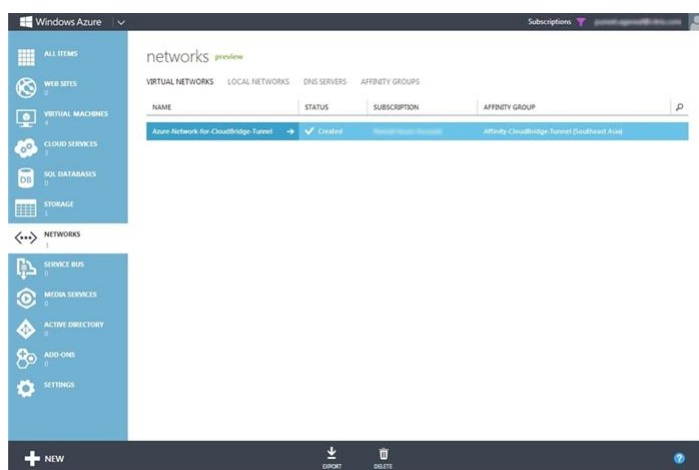
- LOCAL NETWORK

7. In the Address Space and Subnets screen, set the following parameters:

- ADDRESS SPACE
- SUBNETS
- Gateway

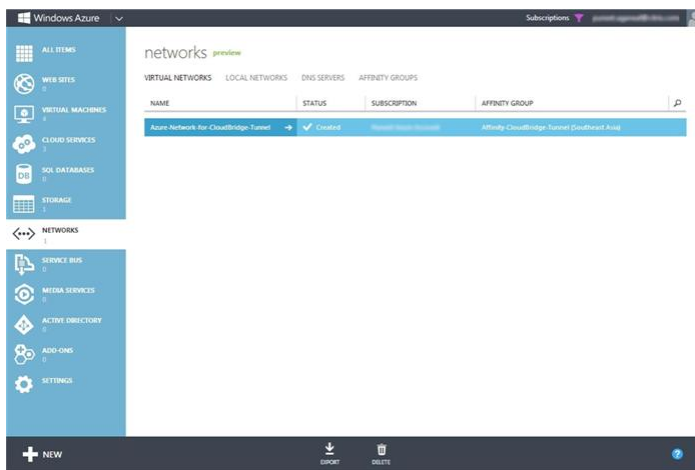


- Click the check mark in the lower right-hand corner of the screen.
- The virtual network is created in Windows Azure and is listed on the VIRTUAL NETWORK tab.

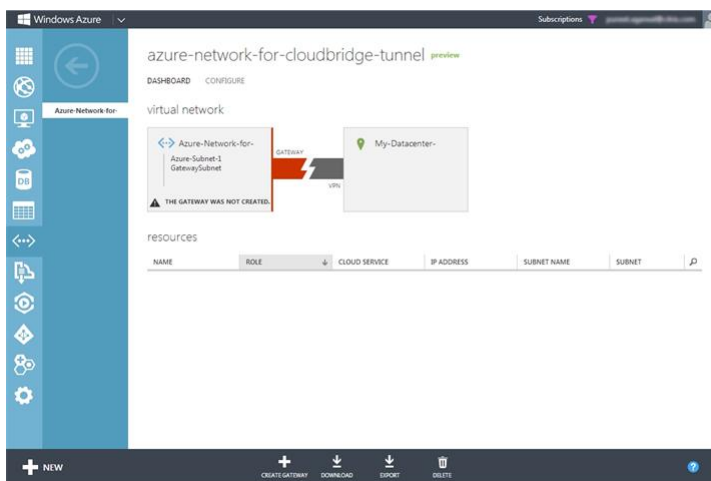


To create a gateway by using the Microsoft Windows Azure Management Portal

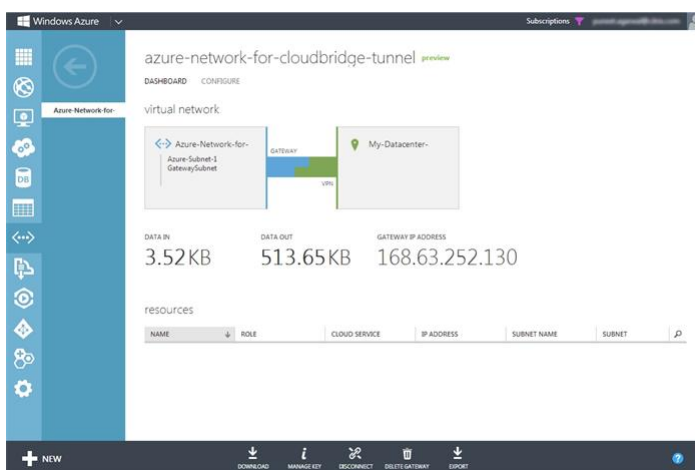
- In the left pane, click NETWORKS.
- On the Virtual Network tab, in the Name column, click the virtual network entity for which you want to create a gateway.



- On the DASHBOARD page of the virtual network, at the bottom of the page, click + Create Gateway.

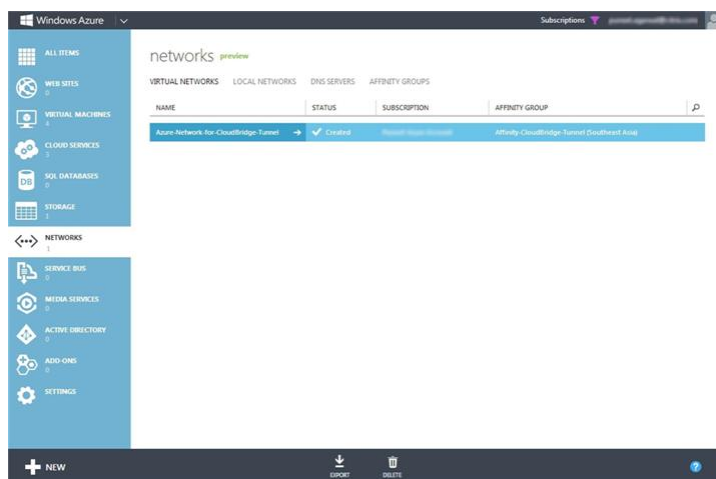


- When prompted to confirm you want the gateway created, click YES. Creating the gateway can take up to 15 minutes.
- When the gateway is created, the DASHBOARD page displays the gateway IP address, which is a public IP address.

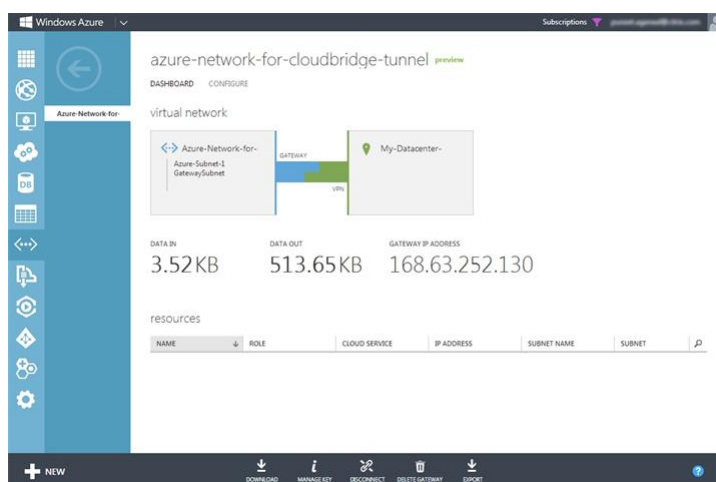


To gather public IP address of the gateway and the pre-shared key information by using the Microsoft Windows Azure Management Portal

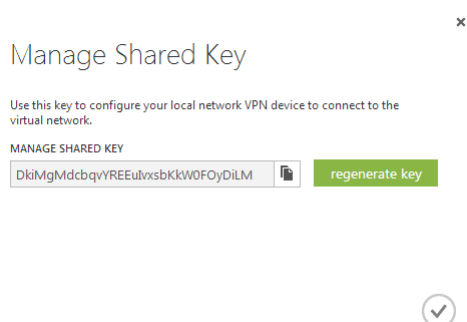
- In the left pane, click NETWORKS.
- On the Virtual Network tab, in the Name column, click the virtual network entity.



3. On the DASHBOARD page of the virtual network, copy the Gateway IP Address.



4. For the Pre Shared Key (PSK), at the bottom of the page, click MANAGE KEY.
5. In the MANAGE SHARED KEY dialog box, copy the SHARED KEY.



Configuring the NetScaler Appliance in the Datacenter for the CloudBridge Connector Tunnel

Updated: 2014-04-15

To configure a CloudBridge Connector tunnel between a datacenter and an Azure cloud, perform the following tasks on the NetScaler in the datacenter. You can use either the NetScaler command line or the configuration utility:

- o **Create an IPSec profile.** An IPSec profile entity specifies the IPSec protocol parameters, such as IKE version, encryption algorithm, hash algorithm, and PSK, to be used by the IPSec protocol in the CloudBridge Connector tunnel.
- o **Create an IP tunnel with IPSec protocol and associate the IPSec profile to it.** An IP tunnel specifies the local IP address (a public SNIP address configured on the NetScaler appliance), remote IP address (the public IP address of the gateway in Azure), protocol (IPSec) used to set up the CloudBridge Connector tunnel, and an IPSec profile entity. The created IP tunnel entity is also called the CloudBridge Connector tunnel entity.
- o **Create a PBR rule and associate the IP tunnel to it.** A PBR entity specifies a set of conditions and an IP tunnel (CloudBridge Connector tunnel) entity. The source IP address range and the destination IP range are the conditions for the PBR entity. You must set the source IP address range to specify the datacenter subne

whose traffic is to traverse the tunnel, and the destination IP address range to specify the Azure subnet whose traffic is to traverse the CloudBridge Connector tunnel. Any request packet originated from a client in the subnet on the datacenter and destined to a server in the subnet on the Azure cloud matches the source and destination IP range of the PBR entity. This packet is then considered for CloudBridge Connector tunnel processing and is sent across the CloudBridge Connector tunnel associated with the PBR entity

The configuration utility combines all these tasks in a single wizard called the CloudBridge Connector wizard.

To create an IPSEC profile by using the NetScaler command line

At the Command prompt, type:

- `add ipsec profile <name> -psk <string> -ikeVersion v1`

To create an IPSEC tunnel and bind the IPSEC profile to it by using the NetScaler command line

At the Command prompt, type:

- `add ipTunnel <name> <remote> <remoteSubnetMask> <local> -protocol IPSEC â€" ipsecProfileName <string>`

To create a PBR rule and bind the IPSEC tunnel to it by using the NetScaler command line

At the Command prompt, type:

- `add pbr <pbrName> ALLOW â€"srcIP <subnet-range> -dstIP <subnet-range> - ipTunnel <tunnelName>`
- `apply pbrs`

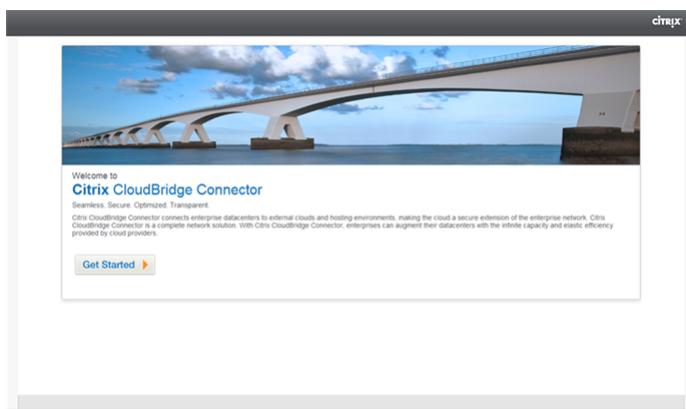
Sample Configuration

The following commands create all settings of NetScaler appliance CB_Appliance-1 used in "Example of CloudBridge Connector Configuration and Data Flow".

```
> add ipsec profile CB_Azure_IPSec_Profile -psk DkiMgMdcBqvYREEuIvxsbKkW0FOyDiLM -ik
Done
> add iptunnel CB_Azure_Tunnel 168.63.252.133 255.255.255.255 66.165.176.15 â€"protoc
Done
> add pbr CB_Azure_Pbr-srcIP 10.102.147.0-10.102.147.255 â€"dstIP 10.20.0.0-10.20.255.2
Done
> add apply pbrs
Done
```

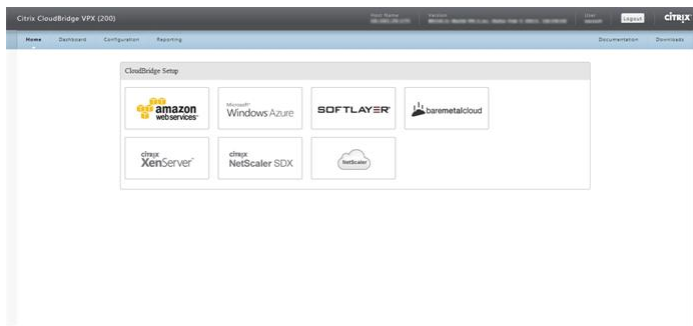
To configure a CloudBridge Connector tunnel in a NetScaler appliance by using the configuration utility

1. Access the configuration utility by using a web browser to connect to the IP address of the NetScaler appliance in the datacenter.
2. Navigate to System > CloudBridge Connector.
3. In the right pane, under Getting Started, click Create/Monitor CloudBridge.
4. Click Get Started.

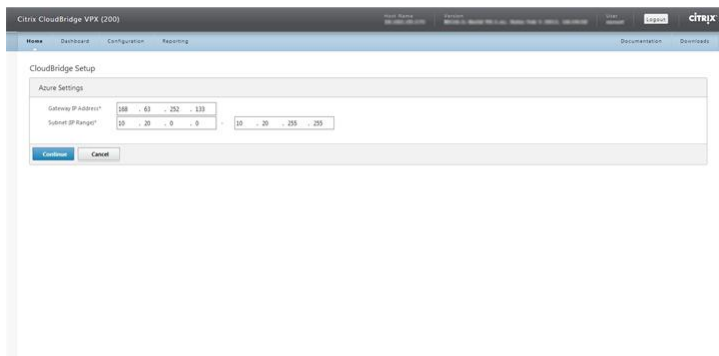


Note: If you already have any CloudBridge Connector tunnel configured on the NetScaler appliance, this screen does not appear, and you are taken to the CloudBridge Connector Setup pane.


5. In the CloudBridge Setup pane, click Microsoft Windows Azure.




6. In the Azure Settings pane, in the Gateway IP Address* field, type the IP address of the Azure gateway. The CloudBridge Connector tunnel is then set up between the NetScaler appliance and the gateway. In the Subnet (IP Range)* text boxes, specify a subnet range (in Azure cloud), the traffic of which is to traverse the CloudBridge Connector tunnel. Click Continue.



7. In the NetScaler Settings pane, from the Local Subnet IP* drop-down list, select a publicly accessible SNIP address configured on the NetScaler appliance. In Subnet (IP Range)* text boxes, specify a local subnet range, the traffic of which is to traverse the CloudBridge Connector tunnel. Click Continue.



8. In the CloudBridge Setting pane, in the CloudBridge Name text box, type a name for the CloudBridge that you want to create.



9. From the Encryption Algorithm and Hash Algorithm drop-down lists, select the AES and HMAC_SHA1 algorithms, respectively. In the Pre Shared Security Key text box, type the security key.
10. Click Done.

Monitoring the CloudBridge Connector Tunnel

Updated: 2014-04-15

You can view statistics for monitoring the performance of a CloudBridge Connector tunnel between the NetScaler appliance in the datacenter and Microsoft Azure. To view CloudBridge Connector tunnel statistics on the NetScaler appliance, use the configuration utility or the NetScaler command line. To view CloudBridge Connector tunnel statistics in Microsoft Azure, use the Microsoft Windows Azure Management Portal.

Displaying CloudBridge Connector tunnel Statistics in the NetScaler appliance

The following table lists the statistical counters available for monitoring CloudBridge Connector tunnels on a NetScaler appliance.

Statistical counter	Specifies
Bytes Received	Total number of bytes received by the NetScaler appliance through all the configured CloudBridge Connector tunnels since the appliance was last started.
Bytes Sent	Total number of bytes sent by the NetScaler appliance through all the configured CloudBridge Connector tunnels since the appliance was last started.
Packets Received	Total number of packets received by the NetScaler appliance through all the configured CloudBridge Connector tunnels since the appliance was last started.
Packets Sent	Total number of packets sent by the NetScaler appliance through all the configured CloudBridge Connector tunnels since the appliance was last started.

All these counters are reset to 0 when the NetScaler appliance is restarted. They do not increment during the following phases:

- Internet Key Exchange (IKE) authentication (pre-shared key) phase on any configured CloudBridge Connector tunnel.
- IKE Security Association (SA) establishment phase on any configured CloudBridge Connector tunnel.

To display CloudBridge Connector tunnel statistics by using the NetScaler command line

At the command prompt, type:

- stat ipsec counters

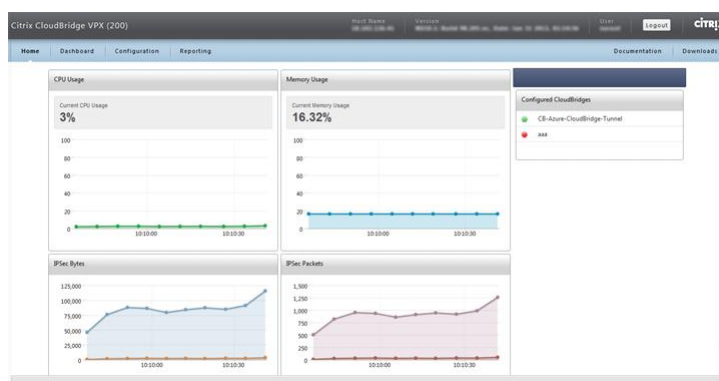
Example

```
> stat ipsec counters
Secure tunnel(s) summary

          Rate (/s)              Total
Bytes Received          0          2811248
Bytes Sent              0          157460630
Packets Received        0          56787
Packets Sent            0          200910
Done
>
```

To display CloudBridge Connector tunnel statistics by using the Configuration utility

- Access the configuration utility by using a web browser to connect to the IP address of the NetScaler appliance
- On the Home tab, the IPSec Bytes and IPSec Packets charts display the statistics of all the CloudBridge Connector tunnels configured on the NetScaler appliance.



Displaying CloudBridge Connector tunnel Statistics in Microsoft Azure

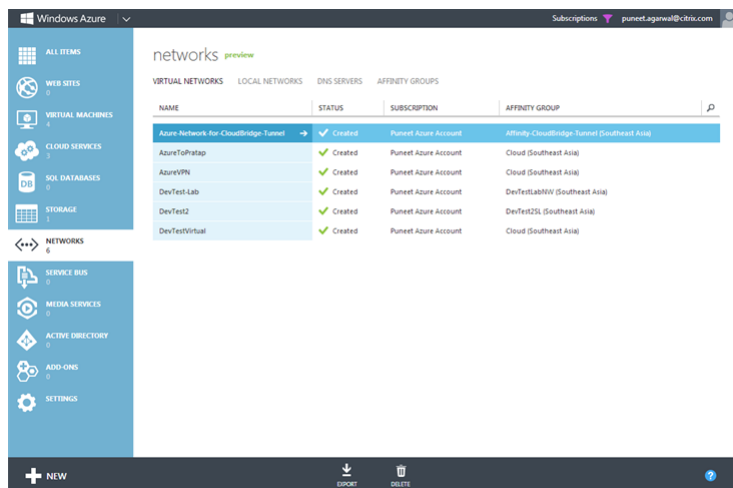
The following table lists the statistical counters available for monitoring CloudBridge Connector tunnels in Microsoft Azure.

Statistical counter	Specifies
---------------------	-----------

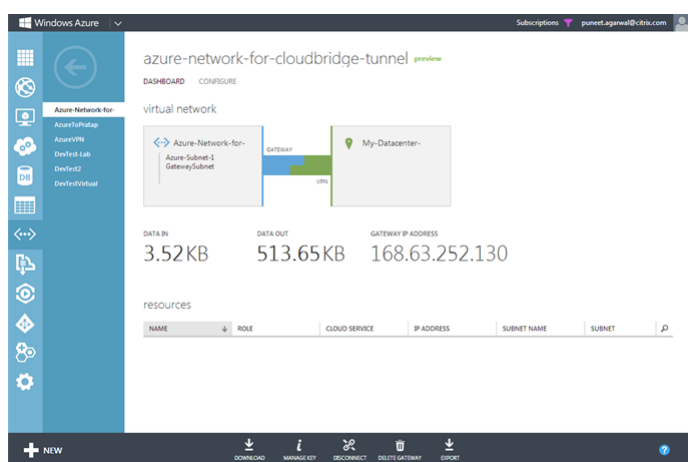
DATA IN	Total number of kilobytes received by the Azure gateway through the CloudBridge Connector tunnel since the gateway was created.
DATA OUT	Total number of kilobytes sent by the Azure gateway through the CloudBridge Connector tunnel since the gateway was created.

To display CloudBridge Connector tunnel statistics by using the Microsoft Windows Azure Management Portal

1. Log on to the Windows Azure Management Portal (<https://manage.windowsazure.com/>) by using your Microsoft Azure account credentials.
2. In the left pane, click NETWORKS.
3. On the Virtual Network tab, in the Name column, select the virtual network entity associated with a CloudBridge Connector tunnel whose statistics you want to display.



4. On the DASHBOARD page of the virtual network, view the DATA IN and DATA OUT counters for the CloudBridge Connector tunnel.



Configuring CloudBridge Connector Tunnel between Datacenter and SoftLayer Enterprise Cloud

The configuration utility includes a wizard that helps you to easily configure a CloudBridge Connector tunnel between a NetScaler appliance in a datacenter and NetScaler VPX instances on the SoftLayer enterprise cloud.

When you use the wizard of the NetScaler appliance in the datacenter, the CloudBridge Connector tunnel configuration created on the NetScaler appliance, is automatically pushed to the other endpoint or peer (the NetScaler VPX on SoftLayer) of the CloudBridge Connector tunnel.

Using the wizard of the NetScaler appliance in the datacenter, you perform the following steps to configure a CloudBridge Connector tunnel.

1. Connect to the Softlayer enterprise cloud by providing the user log on credentials.
2. Select the Citrix XenServer that is running the NetScaler VPX appliance.
3. Select the NetScaler VPX appliance.
4. Provide CloudBridge Connector tunnel parameters to:
 - o Configure a GRE Tunnel.
 - o Configure IPsec on the GRE tunnel.
 - o Create a netbridge, which is a logical representation of the CloudBridge connector, by specifying a name.
 - o Bind the GRE Tunnel to the netbridge.

To configure a CloudBridge Connector tunnel by using the configuration utility

1. Log on to the configuration utility of the NetScaler appliance in the datacenter by using your account credentials for the appliance.
2. Navigate to System > CloudBridge Connector .
3. In the right pane, under Getting Started, click Create/Monitor CloudBridge Connector.
4. Click Get Started.
Note: If you already have any CloudBridge Connector tunnel configured on the NetScaler appliance, this screen does not appear, and you are taken to the CloudBridge Connector Setup pane.
5. In the CloudBridge Connector Setup pane, click Softlayer, and then follow the instructions in the wizard.

High Availability


A high availability (HA) deployment of two Citrix® NetScaler® appliances can provide uninterrupted operation in any transaction. With one appliance configured as the primary node and the other as the secondary node, the primary node accepts connections and manages servers while the secondary node monitors the primary. If, for any reason, the primary node is unable to accept connections, the secondary node takes over.

The secondary node monitors the primary by sending periodic messages (often called heartbeat messages or health checks) to determine whether the primary node is accepting connections. If a health check fails, the secondary node retries the connection for a specified period, after which it determines that the primary node is not functioning normally. The secondary node then takes over for the primary (a process called failover).

After a failover, all clients must reestablish their connections to the managed servers, but the session persistence rules are maintained as they were before the failover.

With Web server logging persistence enabled, no log data is lost due to the failover. For logging persistence to be enabled, the log server configuration must carry entries for both systems in the log.conf file.

The following figure shows a network configuration with an HA pair.
Figure 1. NetScaler Appliances in a High Availability Configuration



To configure HA, you might want to begin by creating a basic setup, with both nodes in the same subnet. You can then customize the intervals at which the nodes communicate health-check information, the process by which nodes maintain synchronization, and the propagation of commands from the primary to the secondary. You can configure fail-safe mode to prevent a situation in which neither node is primary. If your environment includes devices that do not accept NetScaler gratuitous ARP messages, you should configure virtual MAC addresses. When you are ready for a more complex configuration, you can configure HA nodes in different subnets.

To improve the reliability of your HA setup, you can configure route monitors and create redundant links. In some situations, such as when troubleshooting or performing maintenance tasks, you might want to force a node to fail over (assign primary status to the other node), or you might want to force the secondary node to stay secondary or the primary node to stay primary.

Considerations for a High Availability Setup

Note the following requirements for configuring systems in an HA setup:

- In an HA configuration, the primary and secondary NetScaler appliances should be of the same model. Different NetScaler models are not supported in an HA pair (for example, you cannot configure a 10010 model and a 7000 model as an HA pair).
- In an HA setup, both nodes must run the same version of NetScaler, for example, nCore/nCore or classic/classic. If the nodes are running NetScaler classic and you want to migrate to NetScaler nCore of the same NetScaler release, prop and sync are not supported during the migration process. Once migration is complete, prop and sync are auto-enabled. The same applies if you migrate from NetScaler nCore to NetScaler classic.
- Entries in the configuration file (ns.conf) on both the primary and the secondary system must match, with the following exceptions:
 - The primary and the secondary systems must each be configured with their own unique NetScaler IP addresses (NSIPs.)
 - In an HA pair, the node ID and associated IP address of one node must point to the other node. For example, if you have nodes NS1 and NS2, you must configure NS1 with a unique node ID and the IP address of NS2, and you must configure NS2 with a unique node ID and the IP address of NS1.
- If you create a configuration file on either node by using a method that does not go directly through the GUI or the CLI (for example, importing SSL certificates, or changing to startup scripts), you must copy the configuration file to the other node or create an identical file on that node.
- Initially, all NetScaler appliances are configured with the same RPC node password. RPC nodes are internal system entities used for system-to-system communication of configuration and session information. For security, you should change the default RPC node passwords.

One RPC node exists on each NetScaler. This node stores the password, which is checked against the password provided by the contacting system. To communicate with other systems, each NetScaler requires knowledge of those systems, including how to authenticate on those systems. RPC nodes maintain this information, which includes the IP addresses of the other systems, and the passwords they require for authentication.

RPC nodes are implicitly created when adding a node or adding a Global Server Load Balancing (GSLB) site. You cannot create or delete RPC nodes manually.

Note: If the NetScaler appliances in a high availability setup are configured in one-arm mode, you must disable all system interfaces except the one connected to the switch or hub.

- For an IPv6 HA configuration, the following considerations apply:
 - You must install the IPv6PT license on both NetScaler appliances.
 - After installing the IPv6PT license, enable the IPv6 feature by using the configuration utility or the command line interface.
 - Both NetScaler appliances require a global NSIP IPv6 address. In addition, network entities (for example, switches and routers) between the two nodes must support IPv6.


Configuring High Availability

To set up a high availability configuration, you create two nodes, each of which defines the other's NetScaler IP (NSIP) address as a remote node. Begin by logging on to one of the two NetScaler appliances that you want to configure for high availability, and add a node. Specify the other appliance's NetScaler IP (NSIP) address as the address of the new node. Then, log on to the other appliance and add a node that has the NSIP address of the first appliance. An algorithm determines which node becomes primary and which becomes secondary.

Note: The configuration utility provides an option that avoids having to log on to the second appliance.

The following figure shows a simple HA setup, in which both nodes are in same subnet.

Figure 1. Two NetScaler Appliances Connected in a High Availability Configuration



Adding a Remote Node

To add a remote NetScaler appliance as a node in a high availability setup, you specify a unique node ID and the appliance's NSIP. The maximum number of node IDs in an HA setup is 64. When you add an HA node, you must disable the HA monitor for each interface that is not connected or not being used for traffic. For CLI users, this is a separate procedure.

Note: To ensure that each node in the high availability configuration has the same settings, you should synchronize your SSL certificates, startup scripts, and other configuration files with those on the primary node.

To add a node by using the command line interface

At the command prompt, type:

- o add ha node <id> <IPAddress>
- o show ha node

Example

```
> add ha node 3 1000:0000:0000:0000:0005:0600:700a:888b
```

To disable an HA monitor by using the command line interface

At the command prompt, type:

- o set interface <ifNum> [-haMonitor (ON | OFF)]
- o show interface <ifNum>

Example

```
> set interface 1/3 -haMonitor OFF
Done
```


To add a remote node by using the configuration utility

Navigate to System > High Availability and, on the Nodes tab, add a new remote node, or edit an existing node.

Disabling or Enabling a Node

Updated: 2013-08-28

You can disable or enable only a secondary node. When you disable a secondary node, it stops sending heartbeat messages to the primary node, and therefore the primary node can no longer check the status of the secondary. When you enable a node, the node takes part in the high availability configuration.

To disable or enable a node by using the command line interface

At the command prompt, type one of the following commands:

- o set ha node -hastatus DISABLED
- o set ha node -hastatus ENABLED

To disable or enable a node by using the configuration utility

1. Navigate to System > High Availability and, on the Nodes tab, open the node.
2. In the High Availability Status list, select ENABLED (Actively Participate in HA) or DISABLED (Do not participate in HA).

Removing a Node

Updated: 2013-08-28

If you remove a node, the nodes are no longer in high availability configuration.

To remove a node by using the command line interface

At the command prompt, type:

```
rm ha node <id>
```

Example

```
> rm ha node 2
Done
```

To remove a node by using the configuration utility

Navigate to System > High Availability and, on the Nodes tab, delete the node.

Note: You can use the Network Visualizer to view the NetScaler appliances that are configured as a high availability (HA) pair and perform high availability configuration tasks. For more information, see "[Using the Network Visualizer](#)."

Configuring the Communication Intervals

The hello interval is the interval at which the heartbeat messages are sent to the peer node. The dead interval is the time interval after which the peer node is marked DOWN if heartbeat packets are not received. The heartbeat messages are UDP packets sent to port 3003 of the other node in an HA pair.

To set the hello and dead intervals by using the command line interface

At the command prompt, type:

- set HA node [-helloInterval <msecs>] [-deadInterval <secs>]
- show HA node <id>

To set the hello and dead intervals by using the configuration utility

1. Navigate to System > High Availability and, on the Nodes tab, open the node.
2. Set the following parameters:
 - Hello Interval (msecs)
 - Dead Interval (secs)

Configuring Synchronization

Synchronization is a process of duplicating the configuration of the primary node on the secondary node. The purpose of synchronization is to ensure that there is no loss of configuration information between the primary and the secondary nodes, regardless of the number of failovers that occur. Synchronization uses port 3010.

Synchronization is triggered by either of the following circumstances:

- The secondary node in an HA setup comes up after a restart.
- The primary node becomes secondary after a failover.

Automatic synchronization is enabled by default. You can also force synchronization.

Disabling or Enabling Synchronization

Updated: 2013-08-28

Automatic HA synchronization is enabled by default on each node in an HA pair. You can enable or disable it on either node.

To disable or enable automatic synchronization by using the command line interface

At the command prompt, type:

- set HA node -haSync DISABLED
- set HA node -haSync ENABLED

To disable or enable synchronization by using the configuration utility

1. Navigate to System > High Availability and, on the Nodes tab, open the node.
2. Under HA Synchronization, clear or select the Secondary node will fetch the configuration from Primary option.

Forcing the Secondary Node to Synchronize with the Primary Node

Updated: 2013-08-28

In addition to automatic synchronization, the NetScaler supports forced synchronization. You can force the synchronization from either the primary or the secondary node. When you force synchronization from the secondary node, it starts synchronizing its configuration with the primary node.

However, if synchronization is already in progress, forced synchronization fails and the system displays a warning. Forced synchronization also fails in any of the following circumstances:

- You force synchronization on a standalone system.
- The secondary node is disabled.
- HA synchronization is disabled on the secondary node.

To force synchronization by using the command line interface

At the command prompt, type:

force HA sync

To force synchronization by using the configuration utility

1. Navigate to System > High Availability.
2. On the Nodes tab, in the Action list, click Force Synchronization.

Synchronizing Configuration Files in a High Availability Setup

In a high availability setup, you can synchronize various configuration files from the primary node to the secondary node.

To perform the synchronization, you can use the command line interface or the configuration utility at either the primary or the secondary node. Files located on the secondary that are specific to the secondary (not present on the primary) are not deleted during the synchronization.

To synchronize files in a high availability setup by using the command line interface

At the command prompt, type:

```
sync HA files <mode>
```

Example

```
> sync HA files all  
Done
```

To synchronize files in a high availability setup by using the configuration utility

Navigate to System > Diagnostics and, in the Utilities group, click Start HA files synchronization.

Configuring Command Propagation

In an HA setup, any command issued on the primary node propagates automatically to, and is executed on, the secondary before it is executed on the primary. If command propagation fails, or if command execution fails on the secondary, the primary node executes the command and logs an error. Command propagation uses port 3010.

In an HA pair configuration, command propagation is enabled by default on both the primary and secondary nodes. You can enable or disable command propagation on either node in an HA pair. If you disable command propagation on the primary node, commands are not propagated to the secondary node. If you disable command propagation on the secondary node, commands propagated from the primary are not executed on the secondary node.

Note: After reenabling propagation, remember to force synchronization.

If synchronization occurs while you are disabling propagation, any configuration-related changes that you make before the disabling of propagation takes effect are synchronized with the secondary node. This is also true for cases where propagation is disabled while synchronization is in progress.

To disable or enable command propagation by using the command line interface

At the command prompt, type:

- set HA node -haProp DISABLED
- set HA node -haProp ENABLED

To disable or enable command propagation by using the configuration utility

1. Navigate to System > High Availability and, on the Nodes tab, open the node.
2. Clear or select the Primary node will propagate configuration to the Secondary option.

Configuring Fail-Safe Mode

In an HA configuration, fail-safe mode ensures that one node is always primary when both nodes fail the health check. This is to ensure that when a node is only partially available, backup methods are enabled to handle traffic as best as possible. The HA fail-safe mode is configured independently on each node.

The following table shows some of the fail-safe cases. The NOT_UP state means that the node failed the health check yet it is partially available. The UP state means that the node passed the health check.

Table 1. Fail-Safe Mode Cases

Node A (Primary) Health State	Node B (Secondary) Health State	Default HA Behavior	Fail-Safe Enabled HA Behavior	Description
NOT_UP (failed last)	NOT_UP (failed first)	A (Secondary), B (Secondary)	A (Primary), B (Secondary)	If both nodes fail, one after the other, the node that was the last primary remains primary.
NOT_UP (failed first)	NOT_UP(failed last)	A (Secondary), B (Secondary)	A (Secondary), B(Primary)	If both nodes fail, one after the other, the node that was the last primary remains primary.
UP	UP	A (Primary), B (Secondary)	A (Primary), B (Secondary)	If both nodes pass the health check, no change in behavior with fail-safe enabled.
UP	NOT_UP	A(Primary), B (Secondary)	A (Primary), B (Secondary)	If only the secondary node fails, no change in behavior with fail-safe enabled.
NOT_UP	UP	A (Secondary), B(Primary)	A (Secondary), B(Primary)	If only the primary fails, no change in behavior with fail-safe enabled.
NOT_UP	UP (STAYSECONDARY)	A (Secondary), B (Secondary)	A (Primary), B (Secondary)	If the secondary is configured as STAYSECONDARY, the primary remains primary even if it fails.

To enable fail-safe mode by using the command line interface

At the command prompt, type:

```
set HA node [-failSafe ( ON | OFF )]
```

Example

```
set ha node -failsafe ON
```

To enable fail-safe mode by using the configuration utility

1. Navigate to System > High Availability and, on the Nodes tab, open the node.
2. Under Fail-Safe Mode, select the Maintain one Primary node even when both nodes are unhealthy option.

Configuring Virtual MAC Addresses

A Virtual MAC address (VMAC) is a floating entity shared by the primary and the secondary nodes in an HA setup.

In an HA setup, the primary node owns all of the floating IP addresses, such as the MIPs, SNIPs, and VIPs. The primary node responds to Address Resolution Protocol (ARP) requests for these IP addresses with its own MAC address. As a result, the ARP table of an external device (for example, an upstream router) is updated with the floating IP address and the primary node's MAC address.

When a failover occurs, the secondary node takes over as the new primary node. It then uses Gratuitous ARP (GARP) to advertise the floating IP addresses that it acquired from the primary. However, the MAC address that the new primary advertises is the MAC address of its own interface.

Some devices (notably a few routers) do not accept the GARP messages generated by the NetScaler appliance. As a result, some external devices retain the old IP to MAC mapping advertised by the old primary node. This can result in a site going down.

You can overcome this problem by configuring a VMAC on both nodes of an HA pair. Both nodes then possess identical MAC addresses. Therefore, when failover occurs, the MAC address of the secondary node remains unchanged, and the ARP tables on the external devices do not need to be updated.

To create a VMAC, you need to first create a Virtual Router ID (VRID) and bind it to an interface. (In an HA setup, you need to bind the VRID to the interfaces on both nodes.) Once the VRID is bound to an interface, the system generates a VMAC with the VRID as the last octet.

This section includes the following details:

- [Configuring IPv4 VMACs](#)
- [Configuring IPv6 VMACs](#)

Configuring IPv4 VMACs

When you create a IPv4 VMAC address and bind it to a interface, any IPv4 packet sent from the interface uses the VMAC address that is bound to the interface. If there is no IPv4 VMAC bound to an interface, the interface's physical MAC address is used.

The generic VMAC is of the form 00:00:5e:00:01:<VRID>. For example, if you create a VRID with a value of 60 and bind it to an interface, the resulting VMAC is 00:00:5e:00:01:3c, where 3c is the hex representation of the VRID. You can create 255 VRIDs with values from 1 to 255.

Creating or Modifying an IPv4 VMAC

Updated: 2013-08-28

You create an IPv4 virtual MAC by assigning it a virtual router ID. You can then you bind the VMAC to an interface. You cannot bind multiple VRIDs to the same interface. To verify the VMAC configuration, you should display and examine the VMACs and the interfaces bound to the VMACs.

To add a VMAC by using the command line interface

At the command prompt, type:

- `add vrid <id>`
- `bind vrid <id> -ifnum <interface_name>`
- `show vrid`

Example

```
> add vrid 100
Done
> bind vrid 100 -ifnum 1/1 1/2 1/3
Done
```

To unbind interfaces from a VMAC by using the command line interface

At the command prompt, type:

- `unbind vrid <id> -ifnum <interface_name>`

- o show vrid

To configure a VMAC by using the configuration utility

Navigate to System > Network > VMAC and, on the VMAC tab, add a new VMAC, or edit an existing VMAC.

Removing an IPv4 VMAC

Updated: 2013-08-28

To remove an IPv4 virtual MAC, you delete its virtual router ID.

To remove an IPv4 VMAC by using the command line interface

At the command prompt, type:

```
rm vrid <id>
```

Example

```
rm vrid 100s
```

To remove an IPv4 VMAC by using the configuration utility

Navigate to System > Network > VMAC and, on the VMAC tab, delete the IPv4 VMAC.

Configuring IPv6 VMAC6s

The NetScaler supports VMAC6 for IPv6 packets. You can bind any interface to a VMAC6, even if an IPv4 VMAC is bound to the interface. Any IPv6 packet sent from the interface uses the VMAC6 bound to that interface. If there is no VMAC6 bound to an interface, an IPv6 packet uses the physical MAC.

Creating or Modifying a VMAC6

Updated: 2013-08-28

You create an IPv6 virtual MAC by assigning it an IPv6 virtual router ID. You can then you bind the VMAC to an interface. You cannot bind multiple IPv6 VRIDs to an interface. To verify the VMAC6 configuration, you should display and examine the VMAC6s and the interfaces bound to the VMAC6s.

To add a VMAC6 by using the command line interface

At the command prompt, type:

- o add vrid6 <id>
- o bind vrid6 <id> -ifnum <interface_name>
- o show vrid6

Example

```
> add vrid6 100
Done
> bind vrid6 100 -ifnum 1/1 1/2 1/3
Done
```

To unbind interfaces from a VMAC6 by using the command line interface

At the command prompt, type:

- o unbind vrid6 <id> -ifnum <interface_name>
- o show vrid6

To configure a VMAC6 by using the configuration utility

Navigate to System > Network > VMAC and, on the VMAC6 tab, add a new VMAC6, or edit an existing VMAC6.

Removing a VMAC6

Updated: 2013-08-28

To remove an IPv4 virtual MAC, you delete its virtual router ID.

To remove a VMAC6 by using the command line interface

At the command prompt, type:
`rm vrid6 <id>`

Example

```
rm vrid6 100s
```


To remove a VMAC6 by using the configuration utility

Navigate to System > Network > VMAC and, on the VMAC6 tab, delete the virtual router ID.

Configuring High Availability Nodes in Different Subnets

The following figure shows an HA deployment with the two systems located in different subnets:

Figure 1. High Availability over a Routed Network



In the figure, the systems NS1 and NS2 are connected to two separate routers, R3 and R4, on two different subnets. The NetScaler appliances exchange heartbeat packets through the routers. This configuration could be expanded to accommodate deployments involving any number of interfaces.

Note: If you use static routing on your network, you must add static routes between all the systems to ensure that heartbeat packets are sent and received successfully. (If you use dynamic routing on your systems, static routes are unnecessary.)

If the nodes in an HA pair reside on two separate networks, the primary and secondary node must have independent network configurations. This means that nodes on different networks cannot share entities such as MIPs, SNIPs, VLANs, and routes. This type of configuration, where the nodes in an HA pair have different configurable parameters, is known as Independent Network Configuration (INC) or Symmetric Network Configuration (SNC).

The following table summarizes the configurable entities and options for an INC, and shows how they must be set on each node.

Table 1. Behavior of NetScaler Entities and Options in an Independent Network Configuration

NetScaler entities	Options
IPs (NSIP/MIP/SNIPs)	Node-specific. Active only on that node.
VIPs	Floating.
VLANs	Node-specific. Active only on that node.
Routes	Node-specific. Active only on that node. Link load balancing routes are floating.
ACLs	Floating (Common). Active on both nodes.
Dynamic routing	Node-specific. Active only on that node. The secondary node should also run the routing protocols and peer with upstream routers.
L2 mode	Floating (Common). Active on both nodes.
L3 mode	Floating (Common). Active on both nodes.
Reverse NAT (RNAT)	Node-specific. RNAT with VIP, because NATIP is floating.

As in configuring HA nodes in the same subnet, to configure HA nodes in different subnets, you log on to each of the two NetScaler appliances and add a remote node representing the other appliance.

Adding a Remote Node

When two nodes of an HA pair reside on different subnets, each node must have a different network configuration. Therefore, to configure two independent systems to function as an HA pair, you must specify INC mode during the configuration process.

When you add an HA node, you must disable the HA monitor for each interface that is not connected or not being used for traffic. For CLI users, this is a separate procedure.

To add a node by using the command line interface

At the command prompt, type:

- o add ha node <id> <IPAddress> -inc ENABLED
- o show ha node

Example

```
> add ha node 3 10.102.29.170 -inc ENABLED
Done
> add ha node 3 1000:0000:0000:0000:0005:0600:700a:888b
Done
```

To disable an HA monitor by using the command line interface

At the command prompt, type:

- o set interface <ifNum> [-haMonitor (ON | OFF)]
- o show interface <ifNum>

Example

```
> set interface 1/3 -haMonitor OFF
Done
```

To add a remote node by using the configuration utility

1. Navigate to System > High Availability and, on the Nodes tab, add a new remote node.
2. Make sure to select the Turn off HA monitor on interfaces/channels that are down and Turn on INC (Independent Network Configuration) mode on self mode options.

Removing a Node

Updated: 2013-08-28

If you remove a node, the nodes are no longer in high availability configuration.

To remove a node by using the command line interface

At the command prompt, type:

```
rm ha node <id>
```

Example

```
> rm ha node 2
Done
```

To remove a node by using the configuration utility

Navigate to System > High Availability and, on the Nodes tab, delete the node.

Note: You can use the Network Visualizer to view the NetScaler appliances that are configured as a high availability (HA) pair and perform high availability configuration tasks. For more information, see ["Using the Network Visualizer."](#)

Configuring Route Monitors

You can use route monitors to make the HA state dependent on the internal routing table, whether or not the table contains any dynamically learned or static routes. In an HA configuration, a route monitor on each node watches the internal routing table to make sure that a route entry for reaching a particular network is always present. If the route entry is not present, the state of the route monitor changes to DOWN.

When a NetScaler appliance has only static routes for reaching a network, and you want to create a route monitor for the network, you must enable monitored static routes (MSR) for the static routes. MSR removes unreachable static routes from the internal routing table. If MSR is disabled on static routes, an unreachable static route can remain in the internal routing table, defeating the purpose of having the route monitor.

Route Monitors are supported both in non-INC and INC mode.


Route Monitors in HA in non-INC mode	Route Monitors in HA in INC mode
Route monitors are propagated by nodes and exchanged during synchronization.	Route monitors are neither propagated by nodes nor exchanged during synchronization.
Route monitors are active only in the current primary node.	Route monitors are active on both the primary and the secondary node.
The NetScaler appliance always displays the state of a route monitor as UP irrespective of the whether the route entry is present or not in the internal routing table.	The NetScaler appliance displays the state of the route monitor as DOWN if the corresponding route entry is not present in the internal routing table.
<p>A route monitor starts monitoring its route after 180 seconds in the following cases [This is done to allow dynamic routes to get learnt, which may take 180 secs]:</p> <ul style="list-style-type: none"> o reboot o failover o set route6 command for v6 routes o set route msr enable/disable command for v4 routes. o adding a new route monitor 	-

Route monitors are useful in a non-INC mode HA configuration where you want the non-reachability of a gateway from a primary node to be one of the conditions for HA failover.

Consider an example of a non-Inc mode HA setup in a two-arm topology that has NetScaler appliances NS1 and NS2 in the same subnet, with router R1 and switches SW1, SW2, and SW3.

Because R1 is the only router in this setup, you want the HA setup to failover whenever R1 is not reachable from the current primary node. You can configure a route monitor (say, RM1 and RM2, respectively) on each of the nodes to monitor the reachability of R1 from that node.

Figure 1.



With NS1 as the current primary node, the execution flow is as follows:

1. Route monitor RM1 on NS1 monitors NS1's internal routing table for the presence of a route entry for router R1. NS1 and NS2 exchange heartbeat messages through switch SW1 or SW3 at regular intervals.

2. If switch SW1 goes down, the routing protocol on NS1 detects that R1 is not reachable and therefore removes the route entry for R1 from the internal routing table. NS1 and NS2 exchanges heartbeat messages through switch SW3 at regular intervals.
3. Detecting that the route entry for R1 is not present in the internal routing table, RM1 initiates a failover. If route to R1 is down from both NS1 and NS2, failover happens every 180 seconds till one of the appliances is able to reach R1 and restore the connectivity.

Adding a Route Monitor to a High Availability Node

A single procedure creates a route monitor and binds it to an HA node.

To add a route monitor by using the command line interface

At the command prompt, type:

- o bind HA node <id> (-routeMonitor <ip_addr|ipv6_addr> [<netmask>])
- o show HA node

Example

```
> bind HA node 0 -routeMonitor 10.102.71.0 255.255.255.0
Done
> bind HA node 0 -routeMonitor 1000:0000:0000:0000:0005:0600:700a:888b
Done
```

To add a route monitor by using the configuration utility

Navigate to System > High Availability and, on the Route Monitors tab, click Configure.

Removing Route Monitors

Updated: 2013-08-28

To remove a route monitor by using the command line interface

At the command prompt, type:

- o unbind HA node <id> (-routeMonitor <ip_addr|ipv6_addr> [<netmask>])
- o show ha node

Example

```
unbind HA node 3 -routeMonitor 10.102.71.0 255.255.255.0
unbind HA node 3 -routeMonitor 1000:0000:0000:0000:0005:0600:700a:888b
```

To remove a route monitor by using the configuration utility

Navigate to System > High Availability and, on the Route Monitors tab, delete the route monitor.

Limiting Failovers Caused by Route Monitors in non-INC mode

In an HA configuration in non-INC mode, if route monitors fail on both nodes, failover happens every 180 seconds until one of the nodes is able to reach all of the routes monitored by the respective route monitors.

However, for a node, you can limit the number of failovers for a given interval by setting the Maximum Number of Flips and Maximum Flip Time parameters on the nodes. When either limit is reached, no more failovers occur, and the node is assigned as primary even if any route monitor fails on that node. If the node is then able to reach all of the monitored routes, the next monitor failure triggers resetting of the Maximum Number of Flips and Maximum Flip Time parameters on the node and starting the time specified in the Maximum Flip Time parameter.

These parameters are set independently on each node and therefore are neither propagated nor synchronized.

Parameters for limiting the number of failovers

Maximum Number of Flips (maxFlips)

Maximum number of failovers allowed, within the Maximum Flip Time interval, for the node in HA in non INC mode, if the failovers are caused by route-monitor failure.

Maximum Flip Time (maxFlipTime)

Amount of time, in seconds, during which failovers resulting from route-monitor failure are allowed for the node in HA in non INC mode.

To limit the number of failovers by using the command line interface

At the command prompt, type:

- `set HA node [-maxFlips < positive_integer>] [-maxFlipTime <positive_integer>]`
- `show HA node [< id>]`

Example

```
> set ha node -maxFlips 30 -maxFlipTime 60
Done
> sh ha node
1) Node ID: 0
IP: 10.102.169.82 (NS)
Node State: UP
Master State: Primary
Fail-Safe Mode: OFF
INC State: DISABLED
Sync State: ENABLED
Propagation: ENABLED
Enabled Interfaces : 1/1
Disabled Interfaces : None
HA MON ON Interfaces : 1/1
Interfaces on which heartbeats are not seen :None
Interfaces causing Partial Failure:None
SSL Card Status: NOT PRESENT
Hello Interval: 200 msecs
Dead Interval: 3 secs
Node in this Master State for: 0:4:24:1
(days:hrs:min:sec)
2) Node ID: 1
IP: 10.102.169.81
Node State: UP
Master State: Secondary
Fail-Safe Mode: OFF
INC State: DISABLED
Sync State: SUCCESS
Propagation: ENABLED
Enabled Interfaces : 1/1
Disabled Interfaces : None
HA MON ON Interfaces : 1/1
Interfaces on which heartbeats are not seen : None
Interfaces causing Partial Failure: None
SSL Card Status: NOT PRESENT
```

Local node information:
Configured/Completed Flips: 30/0
Configured Flip Time: 60
Critical Interfaces: 1/1

Done

To limit the number of failovers by using the configuration utility

1. Navigate to System > High Availability and, on the Nodes tab, open the local node.
2. Set the following parameters:
 - o Maximum Number of Flips
 - o Maximum Flip Time

Configuring Failover Interface Set

A Failover Interface Set (FIS) is a logical group of interfaces. In an HA configuration, using a FIS is a way to prevent failover by grouping interfaces so that, when one interface fails, other functioning interfaces are still available. A FIS can also be configured for the nodes of a NetScaler cluster.

HA MON interfaces that are not bound to an FIS are known as critical interfaces (CI) because if any of them fails, failover is triggered.

Note

An FIS does not create an active and standby Interfaces or channels. It also does not prevent bridging loops when connecting to links to the same VLAN

Creating or Modifying an FIS

To add an FIS and bind interfaces to it by using the command line interface

At the command prompt, type:

- add fis <name>
- bind fis <name> <ifnum> ...
- show fis <name>

Example

```
> add fis fis1
Done
> bind fis fis1 1/3 1/5
Done
```

An unbound interface becomes a critical interface (CI) if it is enabled and HA MON is on.

To unbind an interface from an FIS by using the command line interface

At the command prompt, type:

- unbind fis <name> <ifnum> ...
- show fis <name>

Example

```
> unbind fis fis1 1/3
Done
```

To configure an FIS by using the configuration utility

Navigate to System > High Availability and, on the Failover Interface Set tab, add a new FIS, or edit an existing FIS.

Removing an FIS

Updated: 2013-08-28

When the FIS is removed, its interfaces are marked as critical interfaces.

To remove an FIS by using the command line interface

At the command prompt, type:

```
rm fis <name>
```

Example


```
> rm fis fis1  
Done
```

To remove an FIS by using the configuration utility

Navigate to System > High Availability and, on the Failover Interface Set tab, delete the FIS.

Understanding the Causes of Failover

The following events can cause failover in an HA configuration:

1. If the secondary node does not receive a heartbeat packet from the primary for a period of time that exceeds the dead interval set on the secondary. (See Note: 1.)
2. The primary node experiences a hardware failure of its SSL card.
3. The primary node does not receive any heartbeat packets on its network interfaces for three seconds.
4. On the primary node, a network interface that is not part of a Failover Interface Set (FIS) or a Link Aggregation (LA) channel and has the HA Monitor (HAMON) enabled, fails. (See Note: 2.)
5. On the primary node, all interfaces in an FIS fail. (See Note: 2.)
6. On the primary node, an LA channel with HAMON enabled fails. (See Note: 2.)
7. On the primary node, all interfaces fail (see Note: 2). In this case, failover occurs regardless of the HAMON configuration.
8. On the primary node, all interfaces are manually disabled. In this case, failover occurs regardless of the HAMON configuration.
9. You force a failover by issuing the force failover command on either node.
10. A route monitor that is bound to the primary node goes DOWN.

Note: 1. For more information about setting the dead interval, see [Configuring the Communication Intervals](#). Possible causes for a node not receiving heartbeat packets from a peer node include:

- A network configuration problem prevents heartbeats from traversing the network between the HA nodes.
- The peer node experiences a hardware or software failure that causes it to freeze (hang), reboot, or otherwise stop processing and forwarding heartbeat packets.

Note: 2. In this case, fail means that the interface was enabled but goes to the DOWN state, as can be seen from the show interface command or from the configuration utility. Possible causes for an enabled interface to be in the DOWN state are LINK DOWN and TXSTALL.

Forcing a Node to Fail Over

You might want to force a failover if, for example, you need to replace or upgrade the primary node. You can force failover from either the primary or the secondary node. A forced failover is not propagated or synchronized. To view the synchronization status after a forced failover, you can view the status of the node.

A forced failover fails in any of the following circumstances:

- You force failover on a standalone system.
- The secondary node is disabled.
- The secondary node is configured to remain secondary.

The NetScaler appliance displays a warning message if it detects a potential issue when you run the force failover command. The message includes the information that triggered the warning, and requests confirmation before proceeding.

You can force a failover on a primary node, secondary node, and when nodes are in listen mode.

◦ Forcing Failover on the Primary Node.

If you force failover on the primary node, the primary becomes the secondary and the secondary becomes the primary. Forced failover is possible only when the primary node can determine that the secondary node is UP.

If the secondary node is DOWN, the force failover command returns the following error message:
"Operation not possible due to invalid peer state. Rectify and retry."

If the secondary system is in the claiming state or inactive, it returns the following error message:
"Operation not possible now. Please wait for system to stabilize before retrying."

◦ Forcing Failover on the Secondary Node.

If you run the force failover command from the secondary node, the secondary node becomes primary and the primary node becomes secondary. A force failover can occur only if the secondary node's health is good and it is not configured to stay secondary.

If the secondary node cannot become the primary node, or if secondary node was configured to stay secondary (using the STAYSECONDARY option), the node displays the following error message:
"Operation not possible as my state is invalid. View the node for more information."

◦ Forcing Failover When Nodes Are in Listen Mode.

When the two nodes of an HA pair are running different versions of the system software, the node running the higher version switches to the listen mode. In this mode, neither command propagation nor synchronization works.

Before upgrading the system software on both nodes, you should test the new version on one of the nodes. To do this, you need to force a failover on the system that has already been upgraded. The upgraded system then takes over as the primary node, but neither command propagation or synchronization occurs. Also, all connections need to be re-established.

To force failover on a node by using the command line interface

At the command prompt, type:

```
force HA failover
```

To force failover on a node by using the configuration utility

Navigate to System > High Availability and, on the Nodes tab, select the node, in the Action list, select Force Failover.

Forcing the Secondary Node to Stay Secondary

In an HA setup, the secondary node can be forced to stay secondary regardless of the state of the primary node.

For example, suppose the primary node needs to be upgraded and the process will take a few seconds. During the upgrade, the primary node may go down for a few seconds, but you do not want the secondary node to take over; you want it to remain the secondary node even if it detects a failure in the primary node.

When you force the secondary node to stay secondary, it will remain secondary even if the primary node goes down. Also, when you force the status of a node in an HA pair to stay secondary, it does not participate in HA state machine transitions. The status of the node is displayed as STAYSECONDARY.

Forcing the node to stay secondary works on both standalone and secondary nodes. On a standalone node, you must use this option before you can add a node to create an HA pair. When you add the new node, the existing node continues to function as the primary node, and the new node becomes the secondary node.

Note: When you force a system to remain secondary, the forcing process is not propagated or synchronized. It affects only the node on which you run the command.

To force the secondary node to stay secondary by using the command line interface

At the command prompt, type:

```
set ha node -hastatus STAYSECONDARY
```

To force the secondary node to stay secondary by using the configuration utility

Navigate to System > High Availability, on the Nodes tab, open the local node, and select STAY SECONDARY.

Forcing the Primary Node to Stay Primary

In an HA setup, you can force the primary node to remain primary even after a failover. You can enable this option either on a primary node in an HA pair or on a standalone system.

On a standalone system, you must run this command before you can add a node to create an HA pair. When you add the new node, it becomes the primary node. The existing node stops processing traffic and becomes the secondary node in the HA pair.

To force the primary node to stay primary by using the command line interface

At the command prompt, type:

```
set ha node -hastatus STAYPRIMARY
```

To force the primary node to stay primary by using the configuration utility

Navigate to System > High Availability, on the Nodes tab, open the local node, and select STAY PRIMARY.

Understanding the High Availability Health Check Computation

The following table summarizes the factors examined in a health check computation:

- State of the CIs
- State of the FISs
- State of the route monitors

The following table summarizes the health check computation.

Table 1. High Availability Health Check Computation

FIS	CI	Route monitor	Condition
N	Y	N	If the system has any CIs, all of those CIs must be UP.
Y	Y	N	If the system has any FISs, all of those FISs must be UP.
Y	Y	Y	If the system has any route monitors configured, all monitored routes must be present in the FIS.

High Availability FAQs

What are the various ports used to exchange the HA-related information between the nodes in an HA configuration?

In an HA configuration, both nodes use the following ports to exchange HA related information:

- UDP Port 3003, to exchange heartbeat packets.
- Port 3010, for synchronization and command propagation.

What are the conditions that trigger synchronization?

Synchronization is triggered by any of the following conditions:

- The incarnation number of the primary node, received by the secondary, does not match that of the secondary node.
Note: Both nodes in an HA configuration maintain a counter called *incarnation number*, which counts the number of configurations in the node's configuration file. Each node sends its incarnation number to each other node in the heartbeat messages. The incarnation number is not incremented for the following commands:
 1. All HA configuration related commands. For example, add ha node, set ha node, and bind ha node.
 2. All Interface related commands. For example, set interface and unset interface.
 3. All channel-related commands. For example, add channel, set channel, and bind channel.
- The secondary node comes up after a restart.
- The primary node becomes secondary after a failover.

What configurations are not synced or propagated in an HA configuration in INC or non-INC mode?

The following commands are neither propagated nor synced to the secondary node:

- All node specific HA configuration commands. For example, add ha node, set ha node, and bind ha node.
- All Interface related configuration commands. For example, set interface and unset interface.
- All channel related configuration commands. For example, add channel, set channel, and bind channel.

What configurations are not synced nor propagated in an HA configuration in INC mode?

The following configurations are not synced or propagated. Each node has its own.

- MIPs
- SNIPs
- VLANs
- Routes (except LLB routes)
- Route monitors
- RNAT rules (except any RNAT rule with VIP as the NAT IP)
- Dynamic routing configurations.

Does a configuration added to the secondary node get synchronized on the primary?

No, a configuration added to the secondary node is not synchronized to the primary.

What could be the reason for both nodes claiming to be the primary in an HA configuration?

The most likely reason is that the primary and secondary nodes are both healthy but the secondary does not receive the heartbeat packets from the primary. The problem could be with the network between the nodes.

Does an HA configuration run into any issues if you deploy the two nodes with different system clock settings?

Different system-clock settings on the two nodes can cause the following issues:

- The time stamps in the log file entries do not match. This situation makes it difficult to analyze the log entries for any issues.
- After a failover, you might have problems with any type of cookie based persistence for load balancing. A significant difference between the times can cause a cookie to expire sooner than expected, resulting in termination of the persistence session.
- Similar considerations apply to any time related decisions on the nodes.

What are the conditions for failure of the *force HA sync* command?

Forced synchronization fails in any of the following circumstances:

- You force synchronization when synchronization is already in progress.
- You force synchronization on a standalone NetScaler appliance.
- The secondary node is disabled.
- HA synchronization is disabled on the current secondary node.
- HA propagation is disabled on the current primary node and you force synchronization from the primary.

What are the conditions for failure of the *sync HA files* command?

Synchronizing configuration files fail in either of the following circumstances:

- On a standalone system.
- With the secondary node disabled.

In an HA configuration, if the secondary node takes over as the primary, does it switch back to secondary status if the original primary comes back online?

No. After the secondary node takes over as the primary, it remains as primary even if the original primary node comes back online again. To interchange the primary and secondary status of the nodes, run the *force failover* command.

What are the conditions for failure of the *force failover* command?

A forced failover fails in any of the following circumstances:

- You force failover on a standalone system.
- The secondary node is disabled.
- The secondary node is configured to remain secondary.
- The primary node is configured to remain primary.
- The state of the peer node is unknown.

Troubleshooting High Availability Issues

The most common high availability issues involve the high availability feature not working at all, or working only intermittently. Following are common high availability issues, and probable causes and resolutions.

o Issue

The inability of the NetScaler appliances to pair the NetScaler appliances in a high availability setup.

Cause

Network connectivity

Resolution

Verify that both the appliances are connected to the switch and the interfaces are enabled.

Cause

Mismatch in the Password for the default Administrator account

Resolution

Verify that the password on both the appliances is the same.

Cause

IP conflict

Resolution

Verify that both the appliances have unique NetScaler IP (NSIP) address. The appliances should not have the same NSIP address.

Cause

Node ID mismatch

Resolution

Verify that the Node ID Configuration on both the appliances is unique. The appliances should not have the same Node ID configuration. Additionally, you must assign value for a Node ID between 1 and 64.

Cause

Mismatch in the password of the RPC node

Resolution

Verify that both the nodes have the same RPC node password.

Cause

An administrator has disabled the remote node

Resolution

Enable the remote node.

Cause

The Firewall application has blocked the heartbeat packets

Resolution

Verify that the UDP port 3003 is allowed.

o Issue

Both the appliances claim to be the primary appliance.

Cause

Missing heartbeat packets between the appliances

Resolution

Verify that the UDP port 3003 is not blocked for communication between the appliances.

o **Issue**

The NetScaler appliance is not able to synchronize the configuration.

Cause

A Firewall application is blocking the required port.

Resolution

Verify that the UDP port 3010 (or UDP port 3008 with secure synchronization) is not blocked for communication between the appliances.

Cause

An administrator has disabled synchronization.

Resolution

Enable synchronization on the appliance that has the issue.

Cause

Different NetScaler releases or builds are installed on appliances.

Resolution

Upgrade the appliances to the same NetScaler release or build.

o **Issue**

Command propagation fails between the appliances.

Cause

A Firewall application is blocking the port.

Resolution

Verify that the UDP port 3011 (or UDP port 3009 with secure propagation) is not blocked for communication between the appliances.

Cause

An administrator has disabled command propagation.

Resolution

Enable command propagation on the appliance that has the issue.

Cause

Different NetScaler releases or builds are installed on appliances.

Resolution

Upgrade the appliances to the same NetScaler release or build.

o **Issue**

The NetScaler appliances in the high availability pair are unable to run the force failover process.

Cause

The Secondary node is disabled.

Resolution

Enable the secondary node.

Cause

The Secondary node is configured to stay secondary.

Resolution

Set the secondary high availability status of the secondary node to Enable from Stay Secondary.

o Issue

The secondary appliance does not receive any traffic after the failover process.

Cause

The upstream router does not understand GARP messages of NetScaler appliance.

Resolution

Configure Virtual MAC (VMAC) address on the secondary appliance.

