



Citrix NetScaler 1000V Frequently Asked Questions

Citrix NetScaler 10.5
December 11, 2014

Cisco Systems, Inc.
www.cisco.com

Cisco has more than 200 offices worldwide.
Addresses, phone numbers, and fax numbers
are listed on the Cisco website at
www.cisco.com/go/offices.

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If the equipment causes interference to radio or television reception, which can be determined by turning the equipment off and on, users are encouraged to try to correct the interference by using one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/TV technician for help.

Modifications to this product not authorized by Cisco could void the FCC approval and negate your authority to operate the product.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

CITRIX Citrix and other Citrix product names referenced herein are trademarks of Citrix Systems, Inc. and/or one of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. All other product names, company names, marks, logos, and symbols are trademarks of their respective owners.

© 2014 Cisco Systems, Inc. All rights reserved.

Contents

1	AppFlow	5
2	AutoScale	9
3	Clustering	11
4	Configuration Utility - Frequently Asked Questions	17
5	Content Switching	25
6	Debugging	31
7	High Availability	33
8	Integrated Caching	37
	Content Groups.....	37
	Cache policy.....	38
	Memory Requirements.....	38
	Verification commands.....	40
	Flushing Objects.....	41
	Flash Cache.....	42
	Default Behaviour.....	42
	Interoperability with other features.....	43
	Miscellaneous.....	43
9	Load Balancing	47
10	Migration	51

11 SSL.....	53
Basic Questions.....	53
Certificates and Keys.....	56
Ciphers.....	58
Certificates.....	62
OpenSSL.....	69
System Limits.....	70

Chapter 1

AppFlow

Which build of NetScaler supports AppFlow?

AppFlow is supported on NetScaler appliances running version 9.3 and above with nCore build.

What is the format used by AppFlow to transmit data?

AppFlow transmits information in 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.

What do AppFlow records contain?

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). IPFIX flow records are based on templates that must be sent before sending flow records.

After an upgrade to NetScaler Version 9.3 Build 48.6 CI, why does an attempt to open a virtual server from the GUI result in the error message "The AppFlow feature is only available on Citrix Netscaler Ncore"

AppFlow is supported only on nCore appliances. When you open the virtual server configuration tab, clear the **AppFlow** checkbox.

What does the transaction ID in an AppFlow records contain?

A transaction ID is an unsigned 32-bit number identifying an application-level transaction. For HTTP, a transaction corresponds to a request and response pair. All flow records that correspond to this request and response pair have the same transaction ID. A typical transaction has four uniflow records. If the NetScaler generates the response by itself (served from the integrated cache or by a security policy), there might be only two flow records for the transaction.

What is an AppFlow action ?

An Appflow action is a set of collectors to which the flow records are sent if the associated AppFlow policy matches.

What commands can I run on the NetScaler appliance to verify that the AppFlow action is a hit?

The show appflow action. For example:

```
> show appflow action
1)      Name: aFL-act-collector-1
        Collectors: collector-1
        Hits: 0
        Action Reference Count: 2
2)      Name: apfl-act-collector-2-and-3
        Collectors: collector-2,
collector-3
        Hits: 0
        Action Reference Count: 1
3)      Name: apfl-act-collector-1-and-3
        Collectors: collector-1,
collector-3
        Hits: 0
        Action Reference Count: 1
```

What is an AppFlow collector?

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. You can remove unused collectors.

What NetScaler version is required for using AppFlow?

Use NetScaler version 9.3.49.5 or higher, and remember that AppFlow is available in only the nCore builds.

What transport protocol does AppFlow use?

AppFlow uses UDP as the transport protocol.

What ports needs to be opened if I have a firewall in the network?

Port 4739. It is the default UDP port the AppFlow collector uses for listening on IPFIX messages. If the user changes the default port, that port should be opened on the firewall.

How can I change the default port AppFlow uses?

When you add an AppFlow collector by using the **add appflowCollector** command, you can specify the port to be used.

```
> add appflowCollector coll1 -IPAddress  
10.102.29.251 -port 8000  
Done
```

What does setting clientTrafficOnly do?

NetScaler generates AppFlow records only for client-side traffic.

How many collectors can be configured at a time?

You can configure up to four AppFlow collectors at a time on the NetScaler appliance. Please note that the maximum number of collectors that can be configured on a NetScaler appliance is four.

Chapter 2

AutoScale

Can the CloudPlatform AutoScale feature be used without a NetScaler appliance?

No. The NetScaler appliance is currently required for the AutoScale feature to work. If the CloudPlatform administrator configures AutoScale in a network that does not include a NetScaler appliance, CloudPlatform throws an error.

What happens if the AutoScale feature is used with a NetScaler release that does not support AutoScale?

If the AutoScale feature is used with a NetScaler release that does not support AutoScale, the CloudPlatform user interface throws an error. CloudPlatform also writes a message to the log file, indicating that the configured NetScaler does not support AutoScale.

In a load balancing rule, can manually provisioned virtual machine instances coexist with instances provisioned by the AutoScale feature?

No. The CloudPlatform virtual machine group in a load balancing rule can contain only manually provisioned instances or only instances provisioned by the AutoScale feature. They cannot coexist.

Is there a limit on the number of virtual machine instances to which we can scale up by using AutoScale?

Yes. The CloudPlatform administrator specifies the maximum number of members to which the configuration can scale up. When the limit is reached, virtual machines are not provisioned even if the scale-up condition is satisfied. The upper limit prevents uncontrolled spawning of VMs due to misconfiguration of the AutoScale feature or unexpected load conditions.

Are AutoScale events observable?

The events generated for deploying or destroying virtual machines are observable. These events are logged in the NetScaler logs (ns.log) and in the CloudPlatform logs (management-server.log). However, you cannot observe the metric values collected by NetScaler monitors.

What metrics can be used in AutoScale policies?

In an AutoScale policy, you can use any metric that is exposed through SNMP, or any NetScaler statistics associated with the load balancing virtual server used in the AutoScale configuration. For example, you can use metrics associated with CPU, memory, or disk usage, and NetScaler metrics such as throughput or response time.

What should a CloudPlatform administrator do before performing maintenance tasks on a CloudPlatform network in which AutoScale is configured?

The CloudPlatform administrator should disable the AutoScale configuration from the CloudPlatform user interface. Disabling the AutoScale configuration temporarily disables any scale-up or scale-down events. However, disabling AutoScale for an application, in CloudPlatform, does not affect the ability of the NetScaler appliance to serve traffic to existing virtual machines.

With AutoScale configured, are any configured VM limits enforced on the user account?

The NetScaler appliance works in the context of an AutoScale user account. Therefore, any limits that the CloudPlatform administrator has imposed on the number of VMs that can be created by the account are automatically enforced when the NetScaler appliance attempts to create more VMs than are permitted.

Is AutoScale supported in a high availability (HA) NetScaler pair?

No. Currently, HA mode is not supported for AutoScale.

Chapter 3

Clustering

How many NetScaler appliances can I have in a cluster?

A NetScaler cluster can include as few as 2 or as many as 32 NetScaler nCore hardware or virtual appliances.

Can a cluster have NetScaler appliances from different networks?

No. The current cluster implementation requires that all cluster nodes be in the same network.

Can a NetScaler appliance be a part of multiple clusters?

No. An appliance can belong to only one cluster.

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
```

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.

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.

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 an appliance is part of a cluster, you must be able to identify the node to which the network 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.

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/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.
- ◆ To remove the NSVLAN from a cluster node, first remove the node from the cluster and then delete the NSVLAN from the appliance.

Can a cluster node that is not connected to the client or server network still serve traffic?

Yes. The cluster supports a traffic distribution 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.

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.

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.

- ◆ **L2 VLAN.** You can create a layer2 VLAN by binding interfaces that belong to different nodes of the cluster.
- ◆ **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.

What details must I have available when I contact technical support for cluster-related issues?

The NetScaler 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?

Applicable for NetScaler 10.1 and later releases.

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 can I configure IP tunnels in a cluster?

Applicable for NetScaler 10.1 and later releases.

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.

How can I add a failover interface set (FIS) on the nodes of a NetScaler cluster?

Applicable for NetScaler 10.5 and later releases.

On the cluster IP address, specify the ID of the cluster node on which the FIS must be added.

```
add fis <name> -ownerNode <nodeld>
```

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.

Are Net Profiles supported on a cluster?

Applicable for NetScaler 10.5 and later releases.

Net profiles are now supported on a NetScaler cluster. You can bind spotted IP addresses to a net profile which can then be bound to spotted lbvserver or service (defined using a node group) with the following recommendations:

Note:

- ◆ If the "strict" parameter of the node group is "Yes", the net profile must contain a minimum of one IP address from each node of the node group member.

- ◆ If the "strict" parameter of the node group is "No", the net profile must include at least one IP address from each of the cluster nodes.
- ◆ If the above recommendations are not followed, the net profile configurations will not be honored and the USIP/USNIP settings will be used.

Chapter 4

Configuration Utility - Frequently Asked Questions

Q: I am using a MAC Safari browser to upgrade a NetScaler ADC. On the upgrade wizard, when I click the **Browse** button to choose the build file from the appliance, the dialog box does not show any files or folders. Also, when I navigate back to the root folder, the dialog box displays the top level folder, but I cannot browse it. What should I do?

A: On the Safari browser, click the **Settings** icon and navigate to **Preferences > Security > Manage Website Settings > Java**, and then change value of the **When visiting other websites** setting to **Run in unsafe mode**.

Q: After I upgraded the JRE on my appliance to version 7.51, the graphical user interface (GUI) applet does not load when I access the NetScaler configuration utility. The applet download stops at 1%, and Java generates a security error. What should I do?

A: You can use the configuration utility with JRE 7.51 if, in the Java Control Panel, you lower the security level or add the NetScaler appliance's URL to the Exception Site List.

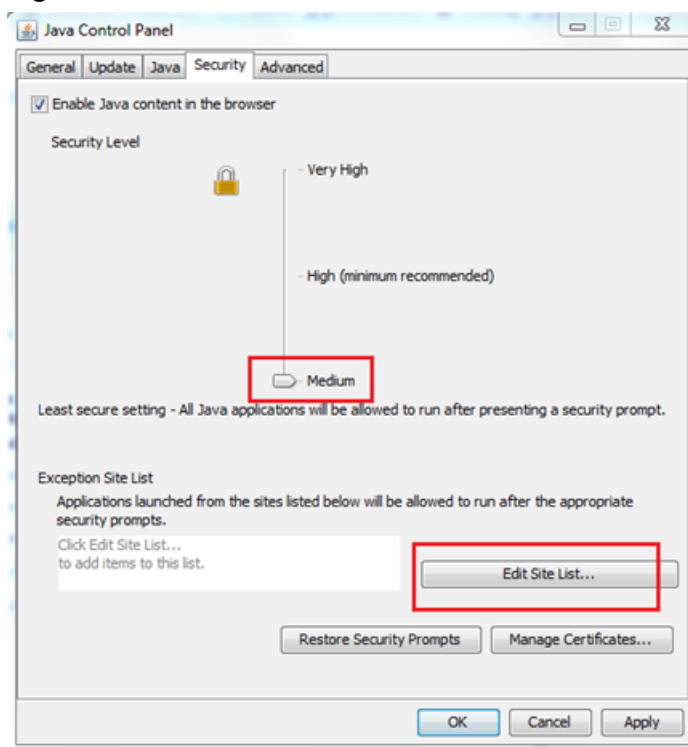
If using a Windows computer, navigate to **Control Panel** and click **Java**.

If using a MAC computer, navigate to **System Preferences** and click **Java**.

Then, in the **Java Control Panel** dialog box, do either of the following:

- ◆ Click the **Security** tab, and then set the **Security Level** to medium.
- ◆ Click the **Security** tab, click the **Edit Site List** button, and then add the URL of the NetScaler appliance in the **Exception Site List** box.

Figure 4-1. Java Control Panel



Q: After I upgraded the JRE on my appliance to version 7.45, the graphical user interface (GUI) applet does not load when I access the NetScaler configuration utility. The applet download stops at 1%, and Java throws a security error. What should I do?

A: The NetScaler GUI is not compatible with JRE version 7.45. Citrix recommends using a JRE version earlier than 7.45 to access the configuration utility. If you have already upgraded to version 7.45 and do not want to downgrade, you can configure Java to not keep temporary files. However, you will have to download the JAR files every time you access the configuration utility.

To downgrade JRE 7.45 to an earlier version:

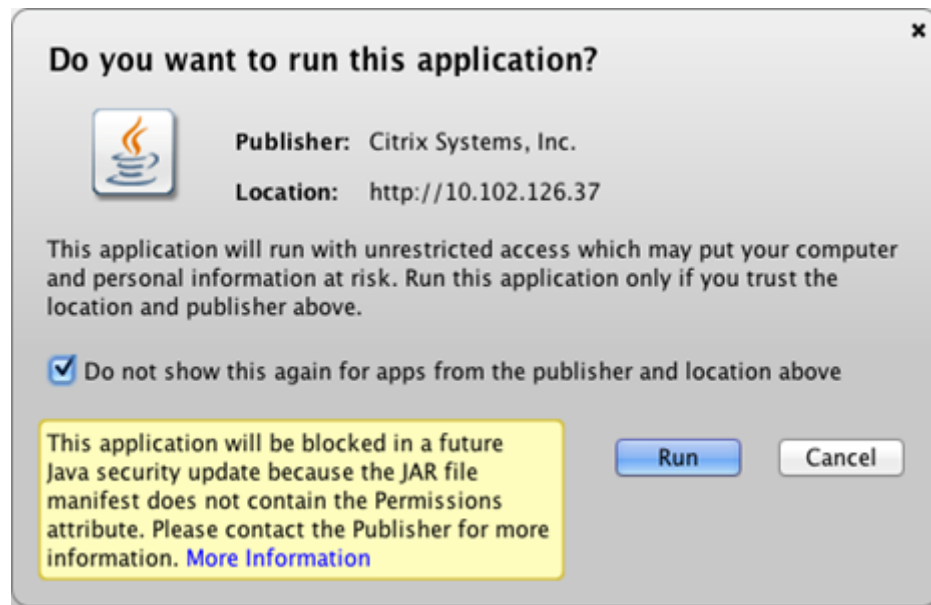
1. Uninstall JRE 7.45 by following the instructions at http://www.java.com/en/download/help/uninstall_java.xml.
2. Download an earlier version of JRE from the following location: <http://www.oracle.com/technetwork/java/javase/archive-139210.html>.

To continue using JRE 7.45:

If you are using a Windows system, navigate to **Control Panel** and click **Java**. The **Java Control Panel** opens.

If you are using a MAC system, navigate to **System Preferences** and click **Java**. The **Java Control Panel** opens.

1. Click the **General** tab.
2. Under **Temporary Internet Files**, click **Settings**.
3. Clear the **Keep temporary files on my computer** option.
4. Close the browser and re-launch the GUI.
5. Open the configuration utility and click **Run** and **Allow**, respectively, when the following warnings appear:





Note:

- ◆ The Jar files will not be cached and you will have to download the files every time you access the configuration utility.
- ◆ This problem is fixed in the following releases:
 - Release 9.3, build 65.x and later
 - Release 10.0, build 78.x and later
 - Release 10.1, build 122.x and later
 - Release 10.1.e, build 120.13xx.e

Q: What should I do before accessing the NetScaler configuration utility?

A: Before accessing a new version of the NetScaler software:

- ◆ Clear your browser cache.
- ◆ Make Sure that JavaScript, Java, and plug-ins are enabled in your browser. For help with enabling Java for your browser, see http://java.com/en/download/help/enable_browser.xml.
- ◆ Clear the “Temporary internet files” in the Java console.
- ◆ On the Java tab of the Java console, in Java Runtime Environment Settings, make sure that the latest version of JRE is present and is enabled.

Q: I am using HTTP to access the configuration utility. Which port should I open?

A: Open TCP port 3010 when using HTTP to access the configuration utility.

Q: I am using HTTPS to access the configuration utility. Which port should I open?

A: Open TCP port 3008 when using HTTPS to access the configuration utility.

Q: After entering the IP address of the NetScaler appliance in the address bar, I get the following error: "Java Applet could not be loaded". What should I do?

A: Verify that Java is installed properly. You can download Java from www.java.com. If you are using a MAC Safari browser, Java is disabled if it is not used for 35 days.

To enable Java plug-in in Safari, follow these steps:

1. In the Safari browser, choose **Safari > Preferences** or press Command-comma (⌘-,) on your keyboard.
2. Click **Security**, and then select **Enable Java**.



3. Close the Safari Preferences window.

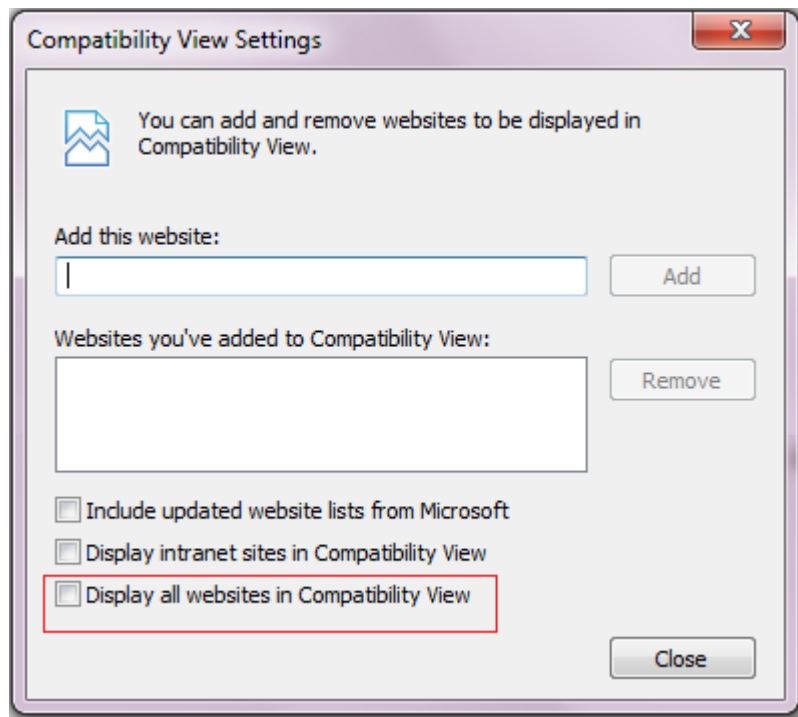
Q: With which browsers is the configuration utility compatible for different operating systems?

A: The following table lists the compatible browsers:

Operating System	Browser	Versions
Windows 7	Internet Explorer	8 and 9
	Mozilla Firefox	3.6.25 and above
	Chrome	15 and above
Windows 64 bit	Internet Explorer	8 and 9
	Chrome	15 and above
MAC OS	Mozilla Firefox	12 and above
	Safari	5.1.3

Q: When I access the NetScaler configuration utility by using Internet Explorer version 8 or 9, the browser displays only a grey bar at the top of the screen. What should I do?

A: The browser might be set in compatibility mode. To disable compatibility mode, go to **Tools > Compatibility View Settings** and clear the **Display all websites in Compatibility View** check box.



Q: Even after I disable compatibility mode in Internet Explorer version 8 or 9, the configuration utility does not appear. What should I do?

A: Make sure that the browser mode and document mode in the browser are set to the same version. To view the configuration, press F12. Set the values to either Internet Explorer 8 or Internet Explorer 9.

Q: When I access the NetScaler configuration utility by using Internet Explorer version 9, the utility displays the following error message: "You are not logged in. Please login." What should I do?

A: Make sure that the cookies are not blocked in your Internet Explorer settings. Go to **Tools > Internet Options**. Click the **Privacy** tab, and then under **Settings**, make sure that the slider is set to **Medium** or any lower value.

Q: I am using a MAC OS with JRE 1.7. After logging on to the configuration utility, I am not able to enter value in any of the text fields. What should I do?

A: Install Java 7, update 21 or higher.

Q: I am using a MAC OS. When I click outside a dialog window, the screen goes out of focus. Now, my browser looks disabled and hung. What should I do?

A: Click on the Java icon in the system dock, and in the JRE Security Warning window, click **Don't Block**. For details, see <http://www.oracle.com/technetwork/java/javase/7u21-relnotes-1932873.html>



Q: Is there any compatibility issue in using JRE version 7_11 with the latest version of browsers?

A: Yes. Internet Explorer 9 and Firefox 18.0.1 block Java on computers running JRE version 7_11. You have to manually activate Java in the browser or upgrade to a later version of JRE (JRE 7_13).

Chapter 5

Content Switching

I have installed a non-NetScaler load balancing appliance on the network. However, I would like to use the content switching feature of the NetScaler appliance to direct the client requests to the load balancing appliance. Is it possible to use the Content switching feature of the NetScaler appliance with a non-NetScaler load balancing appliance?

Yes. You can use the Content switching feature of the NetScaler appliance with the load balancing feature of the NetScaler appliance or a non-NetScaler load balancing appliance. However, when using the non-NetScaler load balancing appliance, make sure that you create a load balancing virtual server on the NetScaler appliance and bind it to the non-NetScaler load balancing appliance as a service.

How is a Content switching virtual server different from a load balancing virtual server?

A Content switching virtual server is capable only of sending the client requests to other virtual servers. It does not communicate with the servers.

A Load balancing virtual server balances the client load among servers and communicates with the servers. It monitors server availability and can be used to apply different load balancing algorithms to distribute the traffic load.

Content switching is a method used to direct client requests for specific types of content to targeted servers by way of load balancing virtual servers. You can direct the client requests to the servers best suited to handle them. This result in reduced overheads to process the client requests on the servers.

I want to implement the Content switching feature of the NetScaler appliance to direct the client requests. What types of client request can I direct by using the Content switching feature?

You can direct only HTTP, HTTPS, FTP, TCP, Secure TCP, and RTSP client requests by using the Content switching

feature. To direct HTTPS client requests, you must configure the SSL offload feature on the appliance.

I want to create Content switching rules on the NetScaler appliance. What are the various elements of the client request on which I can create a content switching rule?

You can create the content switching rules based on the following elements and their values in the client request:

- ◆ URL
- ◆ URL tokens
- ◆ HTTP version
- ◆ HTTP Headers
- ◆ Source IP address of the client
- ◆ Client version
- ◆ Destination TCP port

I understand that the content switching feature of the NetScaler appliance helps enhance the performance of the network. Is this correct?

Yes. You can direct the client requests you the servers best suited to handle them. The result is reduced overhead for processing the client requests on the servers.

Which feature of the NetScaler appliance should I configure on the NetScaler appliance to enhance the site manageability and response time to the client requests?

You can configure the content switching feature of the NetScaler appliance to enhance the site manageability and response time to the client request. This feature enables you to create content groups within the same domain name and IP address. This approach is flexible, unlike the common approach of explicitly partitioning the content into different domain names and IP addresses, which are visible to the user.

Multiple partitions dividing a Web site into various domain names and IP addresses force the browser to create a separate connection for each domain it finds when rendering and fetching the content of a web page. These additional WAN connections degrade the response time for the web page.

I have hosted a web site on a web server farm. What advantages does the NetScaler content switching feature offer for this type of setup?

The content switching feature provides the following advantages on a NetScaler appliance in a site that is based in a web server farm:

- ◆ Manage the site content by creating a content group within the same domain and IP address.
- ◆ Enhance the response time to client requests by using the content group within the same domain and IP address.
- ◆ Avoid the need for full content replication across domains.
- ◆ Enable application-specific content partitioning. For example, you can direct client requests to a server that handles only dynamic content or only static content, as appropriate for the request.
- ◆ Support multi-homing of multiple domains on the same server and use the same IP address.
- ◆ Reuse connections to the servers.

I want to implement the content switching feature on the NetScaler appliance. I want to direct the client requests to the various servers after evaluating the various parameters of each request. What approach should I follow to implement this setup when configuring the content switching feature?

You can use policy expressions to create policies for the content switching feature. An expression is a condition evaluated by comparing the qualifiers of the client request to an operand by using an operator. You can use the following parameters of the client request to create an expression:

- ◆ **Method**- HTTP request method.
- ◆ **URL**- URL in the HTTP header.
- ◆ **URL TOKENS**- Special tokens in the URL.
- ◆ **VERSION**- HTTP request version.
- ◆ **URL QUERY**- Contains the URL Query LEN, URL LEN, and HTTP header.
- ◆ **SOURCEIP**- IP address of the client.

Following is a complete list of the operators that you can use to create an expression:

- ◆ == (equals)
- ◆ != (not equals)
- ◆ EXISTS
- ◆ NOT EXISTS
- ◆ CONTAINS
- ◆ NOT CONTAINS
- ◆ GT (greater than)
- ◆ LT (less than)

You can also create various rules, which are logical aggregations of a set of expressions. You can combine multiple expressions to create rules. To combine expressions, you can use the && (AND) and || (OR) operators. You can also use parenthesis to create nested and complex rules.

I want to configure a rule based policy along with a URL based policy for the same content switching virtual server. Is it possible to create both types of policies for the same content switching virtual server?

Yes. You can create both type of policies for the same content switching virtual server. However, be sure to assign priorities to set an appropriate precedence for the policies.

I want to create content switching policies that evaluate the domain name, along with a prefix and suffix of a URL, and direct the client requests accordingly. Which type of content switching policy should I create?

You can create a Domain and Exact URL policy. When this type of policy is evaluated, the NetScaler appliance selects a content group if the complete domain name and the URL in the client request match the ones configured. The client request must match the configured domain name and exactly match the prefix and suffix of the URL if they are configured.

I want to create content switching policies that evaluate the domain name, along with a partial prefix and suffix of URL, and direct the client requests accordingly. Which type of content switching policy should I create?

You can create a Domain and Wildcard URL policy for the content switching virtual server. When this type of policy is evaluated, the NetScaler appliance selects a content group if the request matches the complete domain name and partially matches the URL prefix.

What is a Wildcard URL policy?

You can use wildcards to evaluate partial URLs in client requests to the URL you have configured on the NetScaler

appliance. You can use wildcards in the following types of URL-based policies:

- ◆ **Prefix only.** For example, the `/sports/*` expression matches all URLs available under the `/sports` URL. Similarly, the `/sports*` expression matches all URLs whose prefix is `/sports`.
- ◆ **Suffix only.** For example the `/*.jsp` expression matches all URLs with a file extension of `jsp`.
- ◆ **Prefix and Suffix.** For example, the `/sports/*.jsp` expression matches all URLs under the `/sports/` URL that also have the `jsp` file extension. Similarly, the `/sports*.jsp` expression matches all URLs with a prefix of `/sports*` and a file extension of `jsp`.

What is a Domain and Rule policy?

When you create a Domain and Rule policy, the client request must match the complete domain and the rule configured on the NetScaler appliance.

What is the default precedence set for evaluating policies?

By default, the rule based policies are evaluated first.

If some of the content is the same for all client requests, what type of precedence should I use for evaluating policies?

If some of the content is same for all the users and different content should be served on the basis of client attributes, you can use URL-based precedence for policy evaluation.

What policy expression syntaxes are supported in content switching?

Content switching supports two types of policy expressions:

- ◆ **Classic Syntax-** Classic syntax in content switching starts with the keyword `REQ` and is more advanced than the default syntax. Classic policies cannot be bound to an action. Therefore, the target load balancing virtual server can be added only after binding the content switching virtual server.
- ◆ **Default Syntax:** Default syntax generally starts with key word `HTTP` and is easier to configure. A target load balancing virtual server action can be bound to a Default Syntax policy, and the policy can be used on multiple content switching virtual servers.

Can I bind a single content switching policy to multiple virtual servers?

Yes. You can bind a single content switching policy to multiple virtual servers by using policies with defined

actions. Content switching policies that use an action can be bound to multiple content switching virtual servers because the target load balancing virtual server is no longer specified in the content switching policy. The ability to bind a single policy to multiple content switching virtual servers helps to further reduce the size of the content switching configuration.

Can I create an action based policy using classic expressions?

No. As of now NetScaler does not support policies using classic syntax expressions with actions. The target load balancing virtual server should be added when binding the policy instead of defining it in an action.

Chapter 6

Debugging

How can I determine the interface (CLI, GUI, or API) through which an operation was performed?

The NetScaler keeps track of the interfaces through which operations are performed. You can view this information in syslogs (in the NetScaler GUI, navigate to **Configuration > System > Auditing > Audit Messages > Syslog messages**) or in the ns.log (located at the /var/log/ directory) file.

For example, operations that are performed through the API are flagged as "API CMD_EXECUTED."

Chapter 7

High Availability

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 HArealted information:

- ◆ UDP Port 3003, to exchange heartbeat packets
- ◆ Port 3010, for synchronization and command propagation

What configurations are not synced or propagated in an HA configuration in either INC or non-INC mode?

Configurations implemented with 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 or propagated in an HA configuration in INC mode?

The following configurations are neither synced nor 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.

Note: Dynamic routing is not supported in NetScaler 1000V.

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:

- All HA configuration related commands. For example, add ha node, set ha node, and bind ha node.
- All Interface related commands. For example, set interface and unset interface.
- 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.

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.
- ◆ 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 if the secondary node is 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:

- ◆ 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.

Chapter 8

Integrated Caching

Content Groups

How is a DEFAULT content group different from other content groups?

The behavior of the DEFAULT content group is exactly the same as any other group. The only attribute that makes the DEFAULT content group special is that if an object is being cached and no content group has been created, the object is cached in the DEFAULT group.

What is the 'cache-Control' option of the content group level?

You can send any cache-control header the browser. There is a content group level option, `-cacheControl`, which enables you to specify the cache-control header that you want to be inserted in the response to the browser.

What is the 'Minhit' option in content group level?

Minhit is an integer value specifying the minimum number of hits to a cache policy before the object is cached. This value is configurable at the content group level. Following is the syntax to configure this value from the command line interface.

```
add/set cache contentGroup <Content_Group_Name> [-minHits <Integer>]
```

What is the use of the expireAtLastByte option?

The `expireAtLastByte` option enables the integrated cache to expire the object as soon as it has been downloaded. Only requests that are outstanding requests at that time are served from cache. any new requests are sent to the server. This setting is useful when the object is frequently modified, as in the case of stock quotes. This expiry mechanism works along with the Flash Cache feature. To configure `expireAtLastByte` option, run the following command from the command line interface:

```
add cache contentGroup <Group_Name> -  
expireAtLastByte YES
```

Cache policy

What is a caching policy?

Policies determine which transactions are cacheable and which are not. Additionally, policies add or override the standard HTTP caching behavior. Policies determine an action, such as CACHE or NOCACHE, depending on the specific characteristics of the request or response. If a response matches policy rules, the object in the response is added to the content group configured in the policy. If you have not configured a content group, the object is added to the DEFAULT content group.

What is a policy hit?

A hit occurs when a request or response matches a cache policy.

What is a miss?

A miss occurs when a request or response does not match any cache policy. A miss can also occur if the request or response matches a cache policy but some override of RFC behavior prevents the object from being stored in the cache.

I have configured Integrated Caching feature of the NetScaler appliance. When adding the following policy, an error message appears. Is there any error in the command?

```
add cache policy image_caching -rule expl |  
ns_ext_not_jpeg -action cache  
  
> ERROR: No such command
```

In the preceding command, the expression should be within the quotation marks. Without quotation marks, the operator is considered to be the pipe operator.

Memory Requirements

What are the commands that I can run on the NetScaler appliance to check the memory allocated to cache?

To display the memory allocated for cache in the NetScaler appliance, run any of the following commands from the command line interface:

- ◆ show cache parameter

In the output, check the value of the **Memory usage limit** parameter. This is the maximum memory allocated for cache.

- ◆ show cache <Content_Group_Name>

In the output, check the values of the Memory usage and Memory usage limit parameters indicating the memory used and allocated for the individual content group.

My NetScaler appliance has 2 GB of memory. Is there any recommended memory limit for cache?

For any model of the NetScaler appliance, you can allocate half of the memory to the cache. However, Citrix recommends allocating a little less than half of the memory, because of internal memory dependency. You can run the following command to allocate 1 GB of memory to cache:

```
set cache parameter -memLimit 1024
```

Is it possible to allocate memory for individual content groups?

Yes. Even though you allocate memory for the integrated cache globally by running the **set cache parameter - memlimit<Integer>**, you can allocate memory to individual content groups by running the **set cache <Content_Group_Name> -memLimit <Integer>** command. The maximum memory you can allocate to content groups (combined) cannot exceed the memory you have allocated to the integrated cache.

What is the dependency of memory between integrated cache and TCP buffer?

If the NetScaler appliance has 2 GB memory, then the appliance reserves approximately 800 to 900 MB of memory and remaining is allocated to FreeBSD operating system. Therefore, you can allocate up to 512 MB of memory to integrated cache and the rest is allocated to TCP buffer.

Does it affect the caching process if I do not allocate global memory to the integrated cache?

If you do not allocate memory to integrated cache, all requests are sent to the server. To make sure that you have allocated memory to the integrated cache, run the **show cache parameter** command. Actually no objects will be cached if global memory is 0, so this needs to be set first.

Verification commands

What are the options for displaying cache statistics?

You can use either of the following options to display the statistics for cache:

- ◆ `stat cache`
To display the summary of the cache statistics.
- ◆ `stat cache -detail`
To display the full details of the cache statistics.

What are the options for displaying the cached content?

To display the cached content, you can run the `show cache object` command.

What is the command that I can run to display the characteristics of an object stored in cache?

If the object stored in the cache is, for example, `GET // 10.102.12.16:80/index.html`, you can display the details about the object by running the following command from the command line interface of the appliance:

```
show cache object -url '/index.html' -host  
10.102.3.96 -port 80
```

Is it mandatory to specify the group name as a parameter to display the parameterized objects in cache?

Yes. It is mandatory to specify the group name as a parameter to display the parameterized objects in cache. For example, consider that you have added the following policies with the same rule:

```
add cache policy p2 -rule  
ns_url_path_cgibin -action CACHE -  
storeInGroup g1  
add cache policy p1 -rule  
ns_url_path_cgibin -action CACHE -  
storeInGroup g2
```

In this case, for the multiple requests, if policy p1 is evaluated, its hit counter is incremented and the policy stores the object in the g1 group, which has hit parameters. Therefore, you have to run the following command to display the objects from the cache:

```
show cache object -url "/cgi-bin/setCookie.pl" -host  
10.102.18.152 groupName g1
```

Similarly, for another set of multiple requests, if policy p2 is evaluated, its hit counter is incremented and the policy

stores the object in the g2 group, which does not have hit parameters. Therefore, you have to run the following command to display the objects from the cache:

```
show cache object -url "/cgi-bin/setCookie2.pl" -host
10.102.18.152
```

I notice that there are some blank entries in the output of the nscachemgr command. What are those entries?

Consider the following sample output of the nscachemgr command. The blank entries in this output are highlighted in bold face for your reference:

```
root@ns# /netscaler/nscachemgr -a
//10.102.3.89:80/image8.gif
//10.102.3.97:80/staticdynamic.html
//10.102.3.97:80/
//10.102.3.89:80/image1.gif
//10.102.3.89:80/file5.html
//10.102.3.96:80/
//10.102.3.97:80/bg_logo_segue.gif
//10.102.3.89:80/file500.html
//10.102.3.92:80/
//10.102.3.96:80/cgi-bin/rfc/
ccProxyReval.pl
Total URLs in IC = 10
```

The blank entries in the output are due to the default caching properties for GET / HTTP/1.1.

Flushing Objects

How can I flush a selective object from the cache?

You can identify an object uniquely by its complete URL. To flush such object, you can perform any of the following tasks:

- ◆ Flush cache
- ◆ Flush content group
- ◆ Flush the specific object

To flush the specific object, you have to specify the query parameters. You specify the invalParam parameter to flush the object. This parameter applies only to a query.

Does any change in the cache configuration trigger flushing of cache?

Yes. When you make any changes to the cache configuration, all the SET cache commands inherently flush the appropriate content groups.

I have updated the objects on the server. Do I need to flush the cached objects?

Yes. When you update objects on the server, you must flush the cached objects, or at least the relevant objects and content groups. The integrated cache is not affected by an update to the server. It continues to serve the cached objects until they expire.

Flash Cache

What is Flash Cache feature of the NetScaler appliance?

The phenomenon of Flash crowds occurs when a large number of clients access the same content. The result is a sudden surge in traffic toward the server. The Flash Cache feature enables the NetScaler appliance to improve performance in such situation by sending only one request to the server. All other requests are queued on the appliance and the single response is served to all of the requests. You can use either of the following commands to enable the Fast Cache feature:

- ◆ `add cache contentGroup <Group_Name> -flashCache YES`
- ◆ `set cache contentGroup <Group_Name> -flashCache YES`

What is the limit for Flash Cache clients?

The number of Flash Cache clients depends on the availability of resources on the NetScaler appliance.

Default Behaviour

Does the NetScaler appliance proactively receive objects upon expiry?

The NetScaler appliance never proactively receives objects on expiry. This is true even for the negative objects. The first access after expiry triggers a request to the server.

Does the integrated cache add clients to the queue for serving even before it starts receiving the response?

Yes. The integrated cache adds clients to the queue for serving even before it starts receiving the response.

What is the default value for the Verify cached object using parameter of the cache configuration?

HOSTNAME_AND_IP is the default value.

Does the NetScaler appliance create log entries in the log files?

Yes. The NetScaler appliance creates log entries in the log files.

Are compressed objects stored in the cache?

Yes. Compressed objects are stored in the cache.

Interoperability with other features

What happens to objects that are currently stored in cache and are being accessed through SSL VPN?

Objects stored in the cache and accessed regularly are served as cache hits when accessed through SSL VPN.

What happens to objects stored in the cache when accessed through SSL VPN and later accessed through a regular connection?

The objects stored through the SSL VPN access are served as a hit when accessed through the regular connection.

When using weblogging, how do I differentiate entries that indicate response served from cache from those served by the server?

For responses served from the integrated cache, the server log field contains the value IC. For responses served from a server, the server log field contains the value sent by the server. Following is a sample log entry for an integrated caching transaction:

```
"10.102.1.52 - "Mozilla/4.0 (compatible; MSIE 6.0; Windows NT 4.0; .NET CLR 1.0.3705)" "GET /" 200 0 "IC" 10.102.1.45"
```

Along with a client request, the response logged is the one sent to the client and not necessarily the one sent by the server.

Miscellaneous

What do you mean by configuring relexpiry and absexpiry?

By configuring relexpiry and absexpiry, it means that you are overriding the header irrespective of what appears in the header. You can configure different expiry setting and the content group level. With relexpiry, expiration of the header is based on the time at which the object was received by the NetScaler. With absexpiry, expiration is

based on the time configured on the NetScaler. Relexpiry is configured in terms of seconds. Absexpiry is a time of day.

What do you mean by configuring weakpos and heuristic?

The weakpos and heuristic are like fall back values. If there is an expiry header, it is considered only if the last-modified header is present. The NetScaler appliance sets expiry on the basis of the last-modified header and the heuristic parameter. The heuristic expiry calculation determines the time to expiry by checking the last-modified header. Some percentage of the duration since the object was last modified is used as time to expiry. The heuristic of an object that remains unmodified for longer periods of time and is likely to have longer expiry periods. The -heurExpiryParam specifies what percentage value to use in this calculation. Otherwise, the appliance uses the weakpos value.

What should I consider before configuring dynamic caching?

If there is some parameter that is in name-value form and does not have the full URL query, or the appliance receives the parameter in a cookie header or POST body, consider configuring dynamic caching. To configure dynamic caching, you have to configure hitParams parameter.

How is hexadecimal encoding supported in the parameter names?

On the NetScaler appliance, the %HEXHEX encoding is supported in the parameter names. In the names that you specify for hitParams or invalParams, you can specify a name that contains %HEXHEX encoding in the names. For example, name, nam%65, and n%61m%65 are equivalent.

What is the process for selecting a hitParam parameter?

Consider the following excerpt of an HTTP header for a POST request:

```
How do we select a hitparam?  
POST /data2html.asp?  
param1=value1&param2=&param3&param4=value4  
HTTP/1.1  
Accept: image/gif, image/x-xbitmap, image/  
jpeg, image/pjpeg,  
application/vnd.ms-powerpoint, application/  
vnd.ms-excel,  
application/msword, application/x-  
shockwave-flash, */*  
Referer: http://10.102.3.97/forms.html  
Accept-Language: en-us  
Content-Type: application/x-www-form-  
urlencoded  
Accept-Encoding: gzip, deflate  
User-Agent: Mozilla/4.0 (compatible; MSIE  
6.0; Windows NT 5.1)
```

```
Host: 10.102.3.97
Content-Length: 153
Connection: Keep-Alive
Cache-Control: no-cache
Cookie:
ASPSESSIONIDQGQGRNY=NNLLKDADEENOAFLLCCDGFDMO
S1=This+text+is+only+text%2C+not+more+and
+not+less%2C+%0D%0Ajust+text+to+be+itself
%2C+namely+%22Text%22+to+be+posted+as+text+
%28what+else...%29&B1=Submit
```

In the above request, you can use S1 and B1, highlighted in bold face for your reference, as hitParams depending on your requirements. Additionally, if you use -matchCookies YES in the ASPSESSIONIDQGQGRNY content group, then you can also use these parameters as hitParams.

What happens to the queued clients if the response is not cacheable?

If the response is not cacheable, all of the clients in the queue receive the same response that the first client receives.

Can I enable the Poll every time (PET) and Flash Cache features on the same content group?

No. You cannot enable PET and Flash Cache on the same content group. The integrated cache does not perform AutoPET function on Flash Cache content groups. The PET feature ensures that the integrated cache does not serve a stored object without consulting the server. You can configure PET explicitly for a content group.

When are the log entries created for the queued clients?

The log entries are created for the queued clients soon after the appliance receives the response header. The log entries are created only if the response header does not make the object non-cacheable.

What is the meaning of the DNS, HOSTNAME, and HOSTNAME_AND_IP values of the Verify cached object using parameter of the cache configuration?

The meanings are as follows:

- ◆ **set cache parameter -verifyUsing HOSTNAME**
This ignores the destination IP address.
- ◆ **set cache parameter -verifyUsing HOSTNAME_AND_IP**
This matches the destination IP address.
- ◆ **set cache parameter -verifyUsing DNS**
This uses the DNS server.

I have set weakNegRelExpiry to 600, which is 10 minutes. I noticed that 404 responses are not getting cached. What is the reason ?

This completely depends on your configuration. By default, 404 responses are cached for 10 minutes. If you want all 404 responses to be fetched from the server, specify -weakNegRelExpiry 0. You can fine tune the -weakNegRelExpiry to a desired value, such as higher or lower to get the 404 responses cached appropriately. If you have configured -absExpiry for positive responses, then it might not yield desired results.

When the user accesses the site by using the Mozilla Firefox browser, the updated content is served. However, when the user accesses the site by using the Microsoft Internet Explorer browser, stale content is served. What could be the reason?

The Microsoft Internet Explorer browser might be taking the content from its local cache instead of the NetScaler integrated cache. The reason could be that the Microsoft Internet Explorer browser is not respecting the expiry related header in the response.

To resolve this issue, you can disable the local cache of the Internet Explorer and clear the offline content. After clearing the offline content, the browser should display the updated content

What if Hits are zero?

Check to see if the server time and NS time are in sync. And the weakPosrelexpiry limit set should bear the time difference between NS and server as shown below

```
root@ns180# date
```

```
Tue May 15 18:53:52 IST 2012
```

Why are policies getting hits but nothing is being cached?

Verify that memory is allocated to the integrated cache and that the allocation is greater than zero.

Is it possible to zero the cache counters?

There is no command line or GUI option for setting the cache counters to zero, and flushing the cache does not do so either. Rebooting the box automatically sets these counters to zero.

Chapter 9

Load Balancing

What are the various load balancing policies I can create on the NetScaler appliance?

You can create the following types of load balancing policies on the NetScaler appliance:

- ◆ Least Connections
- ◆ Round Robin
- ◆ Least response time
- ◆ Least bandwidth
- ◆ Least packets
- ◆ URL hashing
- ◆ Domain name hashing
- ◆ Source IP address hashing
- ◆ Destination IP address hashing
- ◆ Source IP - Destination IP hashing
- ◆ Token
- ◆ LRTM

Can I achieve the Web farm security by implementing load balancing using the NetScaler appliance?

Yes. You can achieve Web farm security by implementing load balancing using the NetScaler appliance. NetScaler appliance enables you to implement the following options of the load balancing feature:

- ◆ IP Address hiding: Enables you to install the actual servers to be on private IP address space for security reasons and for IP address conservation. This process is transparent to the end-user because the NetScaler appliance accepts requests on behalf of the server. While in the address hiding mode, the appliance completely isolates the two networks. Therefore, a client can access a service running on the private

subnet, such as FTP or a Telnet server, through a different VIP on the appliance for that service.

- ◆ Port Mapping: Enables the actual TCP services to be hosted on non-standard ports for security reasons. This process is transparent to the end-user as the NetScaler appliance accepts requests on behalf of the server on the standard advertised IP address and port number.

What are various devices that I can use to load balance with a NetScaler appliance?

You can load balance following devices with a NetScaler appliance:

- ◆ Server farms
- ◆ Caches or Reverse Proxies
- ◆ Firewall devices
- ◆ Intrusion detection systems
- ◆ SSL offload devices
- ◆ Compression devices
- ◆ Content Inspection servers

Why should I implement the load balancing feature for the website?

You can implement the Load balancing feature for the website to take the following advantages:

- ◆ Reduce the response time: When you implement the load balancing feature for the website, one of the major benefits is the boost you can look forward to in load time. With two or more servers sharing the load of the web traffic, each of the servers runs less traffic load than a single server alone. This means there are more resources available to fulfill the client requests. This results in a faster website.
- ◆ Redundancy: Implementing the load balancing feature introduces a bit of redundancy. For example, if the website is balanced across three servers and one of them does not respond at all, the other two can keep running and the website visitors do not even notice any downtime. Any load balancing solution immediately stops sending traffic to the backend server that is not available.

Why do I need to disable the Mac Based Forwarding (MBF) option for Link Load Balancing (LLB)?

- ◆ If you enable the MBF option, the NetScaler appliance considers that the incoming traffic from the client and

the outgoing traffic to the same client flow through the same upstream router. However, the LLB feature requires that the best path be chosen for the return traffic.

- ◆ Enabling the MBF option breaks this topology design by sending the outgoing traffic through the router that forwarded the incoming client traffic.

Chapter 10

Migration

How is the rollback procedure performed on a NetScaler appliance?

The rollback procedure is similar to the basic upgrade procedure. Select the target build that you want to roll back to and perform the downgrade.

Before rolling back to a different release, Citrix recommends that you create a copy of your current configuration files. To downgrade from release 10.1, see [Downgrading from Release 10.1](#) or to downgrade to an earlier build in 10.1, see [Downgrading to an Earlier Build within Release 10.1](#).

Can I upgrade a NetScaler appliance directly from release 10.0 to 10.5, or should I upgrade to 10.1 first and then to 10.5?

You can directly upgrade a NetScaler appliance regardless of the version or the build number.

Is the procedure for upgrading a NetScaler from classic to nCore different?

Yes. Only version 9.3 or earlier supports classic builds. For details, see [Upgrading from a Classic to an nCore Release](#).

Can the primary appliance and secondary appliance have separate builds?

Recommended practice is to use the same version and build number on both the primary and the secondary appliance.

Can both the appliances in an High Availability (HA) pair be upgraded at the same time?

No. In an HA pair, first upgrade the secondary node and then upgrade the primary node. For details, refer [Upgrading a High Availability Pair](#) or [Upgrading a NetScaler High Availability Pair to a Later Build](#).

Does Citrix support firmware upgrades in the amazon AWS cloud?

Yes.

Can I use the FTP server to upgrade the NetScaler appliance?

No. You must first download the firmware from the Citrix download site, save it on your local computer and then upgrade the appliance.

Is the procedure for upgrading the NetScaler appliance with GSLB configurations different from an upgrade of an appliance that is not involved in GSLB?

No. The upgrade procedure is similar to the basic upgrade procedure. The only difference is that you can upgrade the standalone or HA appliances on different sites in a phased manner.

Chapter 11

SSL

Basic Questions

HTTPS access to the NetScaler configuration utility fails on a VPX instance. How do I gain access?

A certificate-key pair is required for HTTPS access to the NetScaler configuration utility. On a NetScaler ADC, a certificate-key pair is automatically bound to the internal services. On an MPX or SDX appliance, the default key size is 1024 bytes, and on a VPX instance, the default key size is 512 bytes. However, most browsers today do not accept a key that is less than 1024 bytes. As a result, HTTPS access to the VPX configuration utility is blocked.

Citrix recommends that you install a certificate-key pair of at least 1024 bytes and bind it to the internal service for HTTPS access to the configuration utility or update the ns-server-certificate to 1024 bytes. You can use HTTP access to the configuration utility or the NetScaler command line to install the certificate.

What are the various steps involved in setting up a secure channel for an SSL transaction?

Setting up a secure channel for an SSL transaction involves the following steps:

1. The client sends an HTTPS request for a secure channel to the server.
2. After selecting the protocol and cipher, the server sends its certificate to the client.
3. The client checks the authenticity of the server certificate.
4. If any of the checks fail, the client displays the corresponding feedback.
5. If the checks pass or the client decides to continue even if a check fails, the client creates a temporary, disposable key called the *pre-master secret* and

encrypts it by using the public key of the server certificate.

6. The server, upon receiving the pre-master secret, decrypts it by using the server's private key and generates the session keys. The client also generates the session keys from the pre-master secret. Thus both client and server now have a common session key, which is used for encryption and decryption of application data.

I understand that SSL is a CPU-intensive process. What is the CPU cost associated with the SSL process?

The following two stages are associated with the SSL process:

- ◆ The initial handshake and secure channel setup by using the public and private key technology.
- ◆ Bulk data encryption by using the asymmetric key technology.

Both of the preceding stages can affect server performance, and they require intensive CPU processing for of the following reasons:

1. The initial handshake involves public-private key cryptography, which is very CPU intensive because of large key sizes (1024bit, 2048bit, 4096bit).
2. Encryption/decryption of data is also computationally expensive, depending on the amount of data that needs to be encrypted or decrypted.

What are the various entities of an SSL configuration?

An SSL configuration has the following entities:

- ◆ Server certificate
- ◆ Certificate Authority (CA) certificate
- ◆ Cipher suite that specifies the protocols for the following tasks:
 - Initial key exchange
 - Server and client authentication
 - Bulk encryption algorithm
 - Message authentication
- ◆ Client authentication
- ◆ CRL

- ◆ SSL Certificate Key Generation Tool that enables you to create the following files:
 - Certificate request
 - Self signed certificate
 - RSA and DSA keys
 - DH parameters

I want to use the SSL offloading feature of the Citrix NetScaler appliance. What are the various options for receiving an SSL certificate?

You must receive an SSL certificate before you can configure the SSL setup on the Citrix NetScaler appliance. You can use any of the following methods to receive an SSL certificate:

- ◆ Request a certificate from an authorized CA.
- ◆ Use the existing server certificate.
- ◆ Create a certificate-key pair on the Citrix NetScaler appliance.

Note: This is a test certificate signed by the test Root-CA generated by the NetScaler. Test certificates signed by this Root-CA are not accepted by browsers. The browser throws a warning message stating that the server's certificate cannot be authenticated.

- ◆ For anything other than test purposes, you must provide a valid CA certificate and CA key to sign the server certificate.

What are the minimum requirements for an SSL setup?

The minimum requirements for configuring an SSL setup are as follows:

- ◆ Obtain the certificates and keys.
- ◆ Create a load balancing SSL virtual server.
- ◆ Bind HTTP or SSL services to the SSL virtual server.
- ◆ Bind certificate-key pair to the SSL virtual server.

What are the limits for the various components of SSL?

SSL components have the following limits:

- ◆ Bit size of SSL certificates: 4096.
- ◆ Number of SSL certificates: Depends on the available memory on the appliance.

- ◆ Maximum linked intermediate CA SSL certificates: 9 per chain.
- ◆ CRL revocations: Depends on the available memory on the appliance.

What are the various steps involved in the end-to-end data encryption on a Citrix NetScaler appliance?

The steps involved in the server-side encryption process on a Citrix NetScaler appliance are as follows:

1. The client connects to the SSL VIP configured on the Citrix NetScaler appliance at the secure site.
2. After receiving the secure request, the appliance decrypts the request, applies layer 4-7 content switching techniques and load balancing policies, and selects the best available backend Web server for the request.
3. The Citrix NetScaler appliance creates an SSL session with the selected server.
4. After establishing the SSL session, the appliance encrypts the client request and sends it to the Web server by using the secure SSL session.
5. When the appliance receives the encrypted response from the server, it decrypts and re-encrypts the data, and sends the data to the client by using the client side SSL session.

The multiplexing technique of the Citrix NetScaler appliance enables the appliance to reuse SSL sessions that have been established with the Web servers. Therefore, the appliance avoids the CPU intensive key exchange, known as *full handshake*. This process reduces the overall number of SSL sessions on the server and maintains end-to-end security.

Certificates and Keys

Can I place the certificate and key files at any location? Is there any recommended location to store these files?

You can store the certificate and key files on the Citrix NetScaler appliance or a local computer. However, Citrix recommends that you store the certificate and key files in the `/nsconfig/ssl` directory of the Citrix NetScaler appliance. The `/etc` directory exists in the flash memory of the Citrix NetScaler appliance. This provides portability and

facilitates backup and restoration of the certificate files on the appliance.

Note: Make sure that the certificate and the key files are stored in the same directory.

What is the maximum size of the DH parameter supported on the Citrix NetScaler appliance?

The Citrix NetScaler appliance supports a DH parameter of maximum 2048 bits.

What is the maximum certificate-chain length, that is, the maximum number of certificates in a chain, supported on a Citrix NetScaler appliance?

A Citrix NetScaler appliance can send a maximum of 10 certificates in a chain when sending a server certificate message. A chain of the maximum length includes the server certificate and nine intermediate CA certificates.

What are the various certificate and key formats supported on the Citrix NetScaler appliance?

The Citrix NetScaler appliance supports the following certificate and key formats:

- ◆ Privacy Enhanced Mail (PEM)
- ◆ Distinguished Encoding Rule (DER)

Is there a limit for the number of certificates and keys that I can install on the Citrix NetScaler appliance?

No. The number of certificates and keys that can be installed is limited only by the available memory on the Citrix NetScaler appliance.

I have saved the certificate and key files on the local computer. I want to transfer these files to the Citrix NetScaler appliance by using the FTP protocol. Is there any preferred mode for transferring these files to the Citrix NetScaler appliance?

Yes. If using the FTP protocol, you should use binary mode to transfer the certificate and key files to the Citrix NetScaler appliance.

Note: By default, FTP is disabled. Citrix recommends using the SCP protocol for transferring certificate and key files. The configuration utility implicitly uses SCP to connect to the appliance.

What is the default directory path for the certificate and key?

The default directory path for the certificate and key is /nsconfig/ssl.

When adding a certificate and key pair, what happens if I do not specify an absolute path to the certificate and key files?

When adding a certificate and key pair, if you do not specify an absolute path to the certificate and key files, the Citrix NetScaler appliance searches the default directory, /nsconfig/ssl, for the specified files and attempts to load them to the kernel. For example, if the cert1024.pem and rsa1024.pem files are available in the /nsconfig/ssl directory of the appliance, both of the following commands are successful:

```
add ssl certKey cert1 -cert cert1204.pem -key
rsa1024.pem
```

```
add ssl certKey cert1 -cert /nsconfig/ssl/
cert1204.pem -key /nsconfig/ssl/rsa1024.pem
```

I have configured a high availability setup. I want to implement the SSL feature on the setup. How should I handle the certificate and key files in a high availability setup?

In a high availability setup, you must store the certificate and key files on both the primary and the secondary Citrix NetScaler appliance. The directory path for the certificate and key files must be the same on both appliances before you add an SSL certificate-key pair on the primary appliance.

Ciphers

What is a NULL-Cipher?

Ciphers with no encryption are known as NULL-Ciphers. For example, NULL-MD5 is a NULL-Cipher.

Are the NULL-Ciphers enabled by default for an SSL VIP or an SSL service?

No. NULL-Ciphers are not enabled by default for an SSL VIP or an SSL service.

What is the procedure to remove NULL-Ciphers?

To remove the NULL-Ciphers from an SSL VIP, run the following command:

```
bind ssl cipher <SSL_VIP> REM NULL
```

To remove the NULL-Ciphers from an SSL Service, run the following command:

```
bind ssl cipher <SSL_Service> REM NULL -service
```

What are the various cipher aliases supported on the Citrix NetScaler appliance?

The Citrix NetScaler appliance supports the following cipher aliases:

1. Alias Name: ALL
Description: All NetScaler-supported ciphers, excluding NULL ciphers
2. Alias Name: DEFAULT
Description: Default cipher list with encryption strength >= 128bit
3. Alias Name: kRSA
Description: Ciphers with RSA key exchange algorithm
4. Alias Name: kEDH
Description: Ciphers with Ephemeral-DH key exchange algorithm
5. Alias Name: DH
Description: Ciphers with DH key exchange algorithm
6. Alias Name: EDH
Description: Ciphers with DH key exchange algorithm and authentication algorithm
7. Alias Name: aRSA
Description: Ciphers with RSA authentication algorithm
8. Alias Name: aDSS
Description: Ciphers with DSS authentication algorithm
9. Alias Name: aNULL
Description: Ciphers with NULL authentication algorithm
10. Alias Name: DSS
Description: Ciphers with DSS authentication algorithm
11. Alias Name: DES
Description: Ciphers with DES encryption algorithm
12. Alias Name: 3DES
Description: Ciphers with 3DES encryption algorithm
13. Alias Name: RC4

- Description: Ciphers with RC4 encryption algorithm
- 14. Alias Name: RC2
Description: Ciphers with RC2 encryption algorithm
- 15. Alias Name: eNULL
Description: Ciphers with NULL encryption algorithm
- 16. Alias Name: MD5
Description: Ciphers with MD5 message authentication code (MAC) algorithm
- 17. Alias Name: SHA1
Description: Ciphers with SHA-1 MAC algorithm
- 18. Alias Name: SHA
Description: Ciphers with SHA MAC algorithm
- 19. Alias Name: NULL
Description: Ciphers with NULL encryption algorithm
- 20. Alias Name: RSA
Description: Ciphers with RSA key exchange algorithm and authentication algorithm
- 21. Alias Name: ADH
Description: Ciphers with DH key exchange algorithm, and NULL authentication algorithm
- 22. Alias Name: SSLv2
Description: SSLv2 protocol ciphers
- 23. Alias Name: SSLv3
Description: SSLv3 protocol ciphers
- 24. Alias Name: TLSv1
Description: SSLv3/TLSv1 protocol ciphers
- 25. Alias Name: TLSv1_ONLY
Description: TLSv1 protocol ciphers
- 26. Alias Name: EXP
Description: Export ciphers
- 27. Alias Name: EXPORT
Description: Export ciphers
- 28. Alias Name: EXPORT40
Description: Export ciphers with 40-bit encryption
- 29. Alias Name: EXPORT56

Description: Export ciphers with 56-bit encryption

30. Alias Name: LOW

Description: Low strength ciphers (56-bit encryption)

31. Alias Name: MEDIUM

Description: Medium strength ciphers (128-bit encryption)

32. Alias Name: HIGH

Description: High strength ciphers (168-bit encryption)

33. Alias Name: AES

Description: AES ciphers

34. Alias Name: FIPS

Description: FIPS-approved ciphers

What is the command to display all the predefined ciphers of the Citrix NetScaler appliance?

To display all the predefined ciphers of the Citrix NetScaler appliance, at the NetScaler command line, type:

```
show ssl cipher
```

What is the command to display the details of an individual cipher of the Citrix NetScaler appliance?

To display the details of an individual cipher of the Citrix NetScaler appliance, at the NetScaler command line, type:

```
show ssl cipher <Cipher_Name/Cipher_Alias_Name/  
Cipher_Group_Name>
```

Example:

```
> show cipher SSL3-RC4-SHA  
1) Cipher Name: SSL3-RC4-SHA  
Description: SSLv3 Kx=RSA Au=RSA  
Enc=RC4(128)  
Mac=SHA1  
Done
```

What is the significance of adding the predefined ciphers of the Citrix NetScaler appliance?

Adding the predefined ciphers of the Citrix NetScaler appliance causes the NULL-Ciphers to get added to an SSL VIP or an SSL service.

Certificates

Why do I need to bind the server certificate?

Binding the server certificates is the basic requirement for enabling the SSL configuration to process SSL transactions.

To bind the server certificate to an SSL VIP, at the NetScaler command line, type:

```
bind ssl vservice <vServerName> -certkeyName  
<cert_name>
```

To bind the server certificate to an SSL service, at the NetScaler command line, type:

```
bind ssl service <serviceName> -certkeyName <cert_name>
```

How many certificates can I bind to an SSL VIP or an SSL service?

On a NetScaler virtual appliance, you can bind a maximum of two certificates to an SSL VIP or an SSL service, one each of type RSA and type DSA.

Does SNI support Subject Alternative Name (SAN) certificates?

No. On a NetScaler appliance, SNI is not supported with a SAN extension certificate.

What happens if I unbind or overwrite a server certificate?

When you unbind or overwrite a server certificate, all the connections and SSL sessions created by using the existing certificate are terminated. When you overwrite an existing certificate, the following message appears:

```
ERROR:
```

```
Warning: Current certificate replaces the  
previous binding.
```

I want to create a server certificate on a Citrix NetScaler appliance to test and evaluate the product. What is the procedure to create a server certificate?

Perform the following procedure to create a test certificate.

Note: A certificate created with this procedure cannot be used to authenticate all the users and browsers. After using the certificate for testing, you should obtain a server certificate signed by an authorized Root CA.

To create a self-signed server certificate:

1. To create a Root CA certificate, at the NetScaler command line, type:
create ssl rsakey /nsconfig/ssl/test-ca.key 1024
create ssl certreq /nsconfig/ssl/test-ca.csr -keyfile /nsconfig/ssl/test-ca.key

Enter the required information when prompted, and then type the following command:
create ssl cert /nsconfig/ssl/test-ca.cer /nsconfig/ssl/test-ca.csr ROOT_CERT -keyfile /nsconfig/ssl/test-ca.key
2. Perform the following procedure to create a server certificate and sign it with the root CA certificate that you just created
 - a. To create the request and the key, at the NetScaler command line, type:
create ssl rsakey /nsconfig/ssl/test-server.key 1024
create ssl certreq /nsconfig/ssl/test-server.csr -keyfile /nsconfig/ssl/test-server.key
 - b. Enter the required information when prompted.
 - c. To create a serial-number file, at the NetScaler command line, type:

```
shell
# echo '01' >
/nsconfig/ssl/serial.txt
# exit
```
3. To create a server certificate signed by the root CA certificate created in step 1, at the NetScaler command line, type:
create ssl cert /nsconfig/ssl/test-server.cer /nsconfig/ssl/test-server.csr SRVR_CERT -CAcert /nsconfig/ssl/test-ca.cer -CAkey /nsconfig/ssl/test-ca.key -CAserial /nsconfig/ssl/serial.txt
4. To create a Citrix NetScaler certkey, which is the in-memory object that holds the server certificate information for SSL handshakes and bulk encryption, at the NetScaler command line, type:
add ssl certkey test-certkey -cert /nsconfig/ssl/test-server.cer -key /nsconfig/ssl/test-server.key
5. To bind the certkey object to the SSL virtual server, at the NetScaler command line, type:

```
bind ssl vserver <vServerName> -certkeyName  
<cert_name>
```

I have received a Citrix NetScaler appliance on which Citrix NetScaler software release 9.0 is installed. I have noticed an additional license file on the appliance. Is there any change in the licensing policy starting with Citrix NetScaler software release 9.0?

Yes. Starting with Citrix NetScaler software release 9.0, the appliance might not have a single license file. The number of license files depends on the Citrix NetScaler software release edition. For example, if you have installed the Enterprise edition, you might need additional license files for the full functionality of the various features. However, if you have installed the Platinum edition, the appliance has only one license file.

How do I export the certificate from Internet Information Service (IIS)?

There are many ways to do this, but by using the following method the appropriate certificate and private key for the Web site are exported. This procedure **must** be performed on the actual IIS server.

1. Open the Internet Information Services (IIS) Manager administration tool.
2. Expand the **Web Sites** node and locate the SSL-enabled Web site that you want to serve through the Citrix NetScaler.
3. Right-click this Web site and click **Properties**.
4. Click the **Directory Security** tab and, in the Secure Communications section of the window, select the **View Certificate** box.
5. Click the **Details** tab, and then click **Copy to File**.
6. On the **Welcome to the Certificate Export Wizard** page, click **Next**.
7. Select **Yes, export the private key** and click **Next**.

Note: The private key **MUST** be exported for SSL Offload to work on the Citrix NetScaler

8. Make sure that the **Personal Information Exchange - PKCS #12** radio button is selected, and select *only* the **Include all certificates in the certification path if possible** check box. Click **Next**.
9. Enter a password and click **Next**.

10. Enter a file name and location, and then click **Next**. Give the file an extension of **.PFX**.
11. Click **Finish**.

How do I convert the PKCS#12 certificate and install it on the Citrix NetScaler?

1. Move the exported **.PFX** certificate file to a location from where it may be copied to the Citrix NetScaler (that is, to a machine that permits SSH access to the management interface of a Citrix NetScaler appliance). Copy the certificate to the appliance by using a secure copy utility such as SCP.

2. Access the BSD shell and convert the certificate (for example, **cert.PFX**) to **.PEM** format:

```
root@ns# openssl pkcs12 -in cert.PFX -out cert.PEM
```

3. To make sure that the converted certificate is in correct **x509** format, verify that the following command produces no error:

```
root@ns# openssl x509 -in cert.PEM -text
```

4. Verify that the certificate file contains a private key. Begin by issuing the following command:

```
root@ns# cat cert.PEM
```

Verify that the output file includes an **RSA PRIVATE KEY** section.

```
-----BEGIN RSA PRIVATE KEY-----
Mkm^s9KMs9023pz/s...
-----END RSA PRIVATE KEY-----
```

The following is another example of an **RSA PRIVATE KEY** section:

```
Bag Attributes
1.3.6.1.4.1.311.17.2: <No Values>
localKeyID: 01 00 00 00
Microsoft CSP Name: Microsoft RSA
SChannel Cryptographic
Provider
friendlyName:
4b9cef4cc8c9b849ff5c662fd3e0ef7e_76267e3
e-6183-4d45-886e-6e067297b38f

Key Attributes
X509v3 Key Usage: 10
-----BEGIN RSA PRIVATE KEY-----
Proc-Type: 4, ENCRYPTED
DEK-Info: DES-EDE3-CBC, 43E7ACA5F4423968
```

```
pZJ2SfsSVqMbRRf6ug37Clua5gY0Wld4frPIxFXy
JquUhr31dilW5ta3hbIaQ+Rg

... (more random characters)
v8dMugeRplkaH2Uwt/
mWBk4t71Yv7GeHmcmjafK8H8iW80ooPO3D/
ENV8X4U/tlh

5eU6ky3WYZ1BTy6thxxLlwAullynVXZEflNLxqlo
X+ZYl6djgjE3qg==
-----END RSA PRIVATE KEY-----
```

The following is a **SERVER CERTIFICATE** section:

```
Bag Attributes
localKeyID: 01 00 00 00
friendlyName: AG Certificate
subject=/C=AU/ST=NSW/L=Wanniassa/O=Dave
Mother
Asiapacific/OU=Support/
CN=davemother.food.lan
issuer=/DC=lan/DC=food/CN=hotdog
-----BEGIN CERTIFICATE-----
MIIFiTCCBHGgAwIBAgIKGryDgAAAAAAHzANBgkq
hkiG9w0BAQUFADA8MRMwEQYK

... (more random characters)
5pLDWYVHhLkA1pSxvFjNJHRSIydWHc5ltGyKqIUc
BezVaXyel94pNSUYx07NpPV/

MY2ovQyQZM8gGe3+lGFum0VHbv/y/gB9HhFesog=
-----END CERTIFICATE-----
```

The following is an **INTERMEDIATE CA CERTIFICATE** section:

```
Bag Attributes: <Empty Attributes>
subject=/DC=lan/DC=food/CN=hotdog
issuer=/DC=lan/DC=food/CN=hotdog
-----BEGIN CERTIFICATE-----
MIIESDCCAzCgAwIBAgIQah20fCRYTY9LRXYMIRaK
GjANBgkqhkiG9w0BAQUFADA8

... (more random characters)
Nt0nksawDnbKo86rQcNnY5xUs7c7pj2zxj/
IOsgNHUp5W6dDI9pQoqFFaDk=

-----END CERTIFICATE-----
```

Further Intermediate CA certificates may follow, depending on the certification path of the exported certificate.

5. Open the .PEM file in a text editor

6. Locate the first line of the .PEM file and the first instance of the following line, and copy those two lines and all the lines between them:

```
-----END CERTIFICATE-----
```

Note: Make sure that last copied line is the first -----END CERTIFICATE----- line in the .PEM file.

7. Paste the copied lines into a new file. Call the new file something intuitive, such as `cert-key.pem`. This is the certificate-key pair for the server hosting the HTTPS service. This file should contain both the section labeled `RSA PRIVATE KEY` and the section labeled `SERVER CERTIFICATE` in the example above.

Note: The certificate-key pair file contains the private key and must therefore be kept secure.

8. Locate any subsequent sections beginning with -----BEGIN CERTIFICATE----- and ending with ---END CERTIFICATE-----, and copy each such section to a separate new file.

These sections correspond to certificates of trusted CAs that have been included in the certification path. These sections should be copied and pasted into new individual files for these certificates. For example, the `INTERMEDIATE CA CERTIFICATE` section of the example above should be copied and pasted into a new file).

For multiple intermediate CA certificates in the original file, create new files for each intermediate CA certificate in the order in which they appear in the file. Keep track (using appropriate filenames) of the order in which the certificates appear, as they need to be linked together in the correct order in a later step.

9. Copy the certificate-key file (`cert-key.pem`) and any additional CA certificate files into the `/nsconfig/ssl` directory on the Citrix NetScaler.
10. Exit the BSD shell and access the Citrix NetScaler prompt.
11. Follow the steps in "Install the certificate-key files on the appliance" to install the key/certificate once uploaded on the device.

How do I convert the PKCS#7 certificate and install it on the NetScaler appliance?

You can use OpenSSL to convert a PKCS #7 Certificate to a format recognizable by the NetScaler appliance. The

procedure is identical to the procedure for PKCS #12 certificates, except that you invoke OpenSSL with different parameters. The steps for converting PKCS #7 certificates are as follows:

1. Copy the certificate to the appliance by using a secure copy utility, such as SCP.
2. Convert the certificate (for example, cert.P7B) to PEM format:

```
> openssl pkcs7 -inform DER -in cert.p7b -  
print_certs -text -out cert.pem
```

3. Follow steps 3 through 7 as described in the answer to Q32 for PKCS #12 certificates.

Note: Before loading the converted PKCS #7 certificate to the appliance, be sure to verify that it contains a private key, exactly as described in step 3 for the PKCS #12 procedure. PKCS #7 certificates, particularly those exported from IIS, do not typically contain a private key.

When I bind a cipher to a virtual server or service by using the bind cipher command, I see the error message "Command deprecated."

The command for binding a cipher to a virtual server or service has changed.

Use the **bind ssl vserver** <vservername> -ciphername <ciphername> command to bind an SSL cipher to an SSL virtual server.

Use the **bind ssl service** <serviceName> -ciphername <ciphername> command to bind an SSL cipher to an SSL service.

Note: New ciphers and cipher groups are added to the existing list and not replaced.

Why can't I create a new cipher group and bind ciphers to it by using the add cipher command?

The add cipher command functionality has changed in release 10. The command only creates a cipher group. To add ciphers to the group, use the bind cipher command.

OpenSSL

How do I use OpenSSL to convert certificates between PEM and DER?

To use OpenSSL, you must have a working installation of the OpenSSL software and be able to execute Openssl from the command line.

x509 certificates and RSA keys can be stored in a number of different formats.

Two common formats are DER (a binary format used primarily by Java and Macintosh platforms) and PEM (a base64 representation of DER with header and footer information, which is used primarily by UNIX and Linux platforms). There is also an obsolete NET (Netscape server) format that was used by earlier versions of IIS (up to and including 4.0) and various other less common formats that are not covered in this article.

A key and the corresponding certificate, as well as the root and any intermediate certificates, can also be stored in a single PKCS#12 (.P12, .PFX) file.

Procedure

Use the Openssl command to convert between formats as follows:

1. To convert a certificate from PEM to DER:

```
x509 -in input.crt -inform PEM -out  
output.crt -outform DER
```

2. To convert a certificate from DER to PEM:

```
x509 -in input.crt -inform DER -out  
output.crt -outform PEM
```

3. To convert a key from PEM to DER:

```
rsa -in input.key -inform PEM -out  
output.key -outform DER
```

4. To convert a key from DER to PEM:

```
rsa -in input.key -inform DER -out  
output.key -outform PEM
```

Note: If the key you are importing is encrypted with a supported symmetric cipher, you are prompted to enter the pass-phrase.

Note: To convert a key to or from the obsolete NET (Netscape server) format, substitute NET for PEM or DER as appropriate. The stored key is encrypted in a weak unsalted RC4 symmetric cipher, so a pass-phrase will be requested. A blank pass-phrase is acceptable.

System Limits

What are the important numbers to remember?

1. Create Certificate Request:
 - Request File Name: Maximum 63 characters
 - Key File Name: Maximum 63 characters
 - PEM Passphrase (For Encrypted Key): Maximum 31 characters
 - Common Name: Maximum 63 characters
 - City: Maximum 127 characters
 - Organization Name: Maximum 63 characters
 - State/Province Name: Maximum 63 characters
 - Email Address: Maximum 39 Characters
 - Organization Unit: Maximum 63 characters
 - Challenge Password: Maximum 20 characters
 - Company Name: Maximum 127 characters
2. Create Certificate:
 - Certificate File Name: Maximum 63 characters
 - Certificate Request File Name: Maximum 63 characters
 - Key File Name: Maximum 63 characters
 - PEM Passphrase: Maximum 31 characters
 - Validity Period: Maximum 3650 days
 - CA Certificate File Name: Maximum 63 characters
 - CA Key File Name: Maximum 63 characters
 - PEM Passphrase: Maximum 31 characters
 - CA Serial Number File: Maximum 63 characters

3. Create and Install a Server Test Certificate:
 - Certificate File Name: Maximum 31 characters
 - Fully Qualified Domain Name: Maximum 63 characters
4. Create Diffie-Hellman (DH) key:
 - DH Filename (with path): Maximum 63 characters
 - DH Parameter Size: Maximum 2048 bits
5. Import PKCS12 key:
 - Output File Name: Maximum 63 characters
 - PKCS12 File Name: Maximum 63 characters
 - Import Password: Maximum 31 characters
 - PEM Passphrase: Maximum 31 characters
 - Verify PEM Passphrase: Maximum 31 characters
6. Export PKCS12
 - PKCS12 File Name: Maximum 63 characters
 - Certificate File Name: Maximum 63 characters
 - Key File Name: Maximum 63 characters
 - Export Password: Maximum 31 characters
 - PEM Passphrase: Maximum 31 characters
7. CRL Management:
 - CA Certificate File Name: Maximum 63 characters
 - CA Key File Name: Maximum 63 characters
 - CA Key File Password: Maximum 31 characters
 - Index File Name: Maximum 63 characters
 - Certificate File Name: Maximum 63 characters
8. Create RSA Key:
 - Key Filename: Maximum 63 characters
 - Key Size: Maximum 4096 bits
 - PEM Passphrase: Maximum 31 characters
 - Verify Passphrase: Maximum 31 characters
9. Create DSA Key:
 - Key Filename: Maximum 63 characters

- Key Size: Maximum 4096 bits
 - PEM Passphrase: Maximum 31 characters
 - Verify Passphrase: Maximum 31 characters
10. Change advanced SSL settings:
- Maximum CRL memory size: Maximum 1024 Mbytes
 - Encryption trigger timeout (10 mS ticks): Maximum 200
 - Encryption trigger packet count: Maximum 50
 - OCSP cache size: Maximum 512 Mbytes
11. Install Certificate:
- Certificate-Key pair Name: Maximum 31 characters
 - Certificate File Name: Maximum 63 characters
 - Private Key File Name: Maximum 63 characters
 - Password: Maximum 31 characters
 - Notification Period: Maximum 100
12. Create Cipher Group:
- Cipher Group Name: Maximum 39 characters
13. Create CRL:
- CRL Name: Maximum 31 characters
 - CRL File: Maximum 63 characters
 - URL: Maximum 127 characters
 - Base DN: Maximum 127 characters
 - Bind DN: Maximum 127 characters
 - Password: Maximum 31 characters
 - Day(s): Maximum 31
14. Create SSL Policy:
- Name: Maximum 127 characters
15. Create SSL Action:
- Name: Maximum 127 characters
16. Create OCSP Responder:
- Name: Maximum 32 characters
 - URL: Maximum 128 characters

- Batching Depth: Maximum 8
 - Batching Delay: Maximum 10000
 - Produced At Time Skew: Maximum 86400
 - Request Time-out: Maximum 120000
17. Create Virtual Server:
- Name: Maximum 127 characters
 - Redirect URL: Maximum 127 characters
 - Client Time-out: Maximum 31536000 secs
18. Create Service:
- Name: Maximum 127 characters
 - Idle Time-out (secs):
Client: Maximum 31536000
Server: Maximum 31536000
19. Create Service Group:
- Service Group Name: Maximum 127 characters
 - Server ID: Maximum 4294967295
 - Idle Time-out (secs):
Client: Maximum value 31536000
Server: Maximum 31536000
20. Create Monitor:
- Name: Maximum 31 characters
21. Create Server:
- Server Name: Maximum 127 characters
 - Domain Name: Maximum 255 characters
 - Resolve Retry: Maximum 20939 secs