



Citrix NetScaler 1000V Administration Guide

Citrix NetScaler 10.1
October 3, 2013

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.

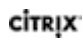
The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

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 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.

© 2013 Cisco Systems, Inc. All rights reserved.

Contents

1 Basic Operations.....	19
Viewing and Saving Configurations.....	20
To view the running configuration by using the command line interface.....	20
To view the running configuration by using the configuration utility.....	20
To find the difference between two configuration files by using the command line interface.....	20
To find the difference between two configuration files by using the configuration utility.....	20
To save configurations by using the command line interface.....	20
To save configurations by using the configuration utility.....	21
To view the saved configurations by using the command line interface.....	21
To view the saved configurations by using the configuration utility.....	21
Clearing the NetScaler Configuration.....	21
To clear the configuration by using the command line interface.....	22
To clear the configuration by using the configuration utility.....	22
Configuring Clock Synchronization.....	22
Setting Up Clock Synchronization.....	23
To add an NTP server by using the command line interface.....	23
To configure an NTP server by using the configuration utility.....	23
Starting or Stopping the NTP Daemon.....	23
To enable or disable NTP synchronization by using the command line interface.....	23
To enable or disable NTP synchronization by using the configuration utility.....	24
Configuring Clock Synchronization Manually.....	24
To enable clock synchronization on your NetScaler by modifying the ntp.conf file.....	24
Viewing the System Date and Time.....	25
To view the system date and time by using the command line interface.....	25
To view the system date and time by using the configuration utility.....	25
Backing up and Restoring the NetScaler Appliance.....	25
Backing up a NetScaler Appliance.....	26

To backup the NetScaler by using the command line interface.....	27
To backup the NetScaler by using the configuration utility.....	28
Restoring the NetScaler Appliance.....	28
To restore the NetScaler by using the command line interface.....	28
To restore the NetScaler by using the configuration utility.....	29
Restarting or Shutting down the Appliance.....	29
To restart the NetScaler by using the command line interface.....	29
To restart the NetScaler by using the configuration utility.....	29
To shut down the NetScaler by using the command line interface.....	30
2 Administration.....	31
Authentication and Authorization.....	33
Configuring Users and Groups.....	33
Configuring User Accounts.....	33
Configuring User Groups.....	34
Configuring Command Policies.....	35
Built-in Command Policies.....	35
Creating Custom Command Policies.....	36
Binding Command Policies to Users and Groups.....	38
Resetting the Default Administrator (nsroot) Password.....	40
To reset the nsroot password.....	40
Example of a User Scenario.....	41
Configuration steps.....	42
Configuring External User Authentication.....	43
Configuring LDAP Authentication.....	44
Configuring RADIUS Authentication.....	47
Configuring TACACS+ Authentication.....	49
Binding the Authentication Policies to the System Global Entity.....	50
SNMP.....	51
Importing MIB Files to the SNMP Manager and Trap Listener.....	51
Configuring the NetScaler to Generate SNMPv1 and SNMPv2 Traps.....	52
Enabling or Disabling an SNMP Alarm.....	52
Configuring Alarms.....	53
Configuring Traps.....	54
Enabling Unconditional SNMP Trap Logging.....	55
Configuring the NetScaler for SNMP v1 and v2 Queries.....	55
Specifying an SNMP Manager.....	56
Specifying an SNMP Community.....	57
Configuring SNMP Alarms for Rate Limiting.....	58
Configuring an SNMP Alarm for Throughput or PPS.....	59

Configuring SNMP Alarm for Dropped Packets.....	60
Configuring the NetScaler for SNMPv3 Queries.....	61
Setting the Engine ID.....	62
Configuring a View.....	62
Configuring a Group.....	63
Configuring a User.....	64
Audit Logging.....	64
Configuring the NetScaler Appliance for Audit Logging.....	65
Configuring Audit Servers.....	66
Configuring Audit Policies.....	67
Binding the Audit Policies Globally.....	68
Configuring Policy-Based Logging.....	68
Installing and Configuring the NSLOG Server.....	69
Installing NSLOG Server on the Linux Operating System.....	70
Installing NSLOG Server on the FreeBSD Operating System.....	71
Installing NSLOG Server Files on the Windows Operating System.....	72
NSLOG Server Command Options.....	73
Adding the NetScaler Appliance IP Addresses on the NSLOG Server.....	75
Verifying the NSLOG Server Configuration File.....	76
Running the NSLOG Server.....	76
To start audit server logging.....	76
To stop audit server logging that starts as a background process in FreeBSD or Linux.....	76
To stop audit server logging that starts as a service in Windows.....	76
Customizing Logging on the NSLOG Server.....	76
Creating Filters.....	77
Specifying Log Properties.....	77
Default Settings for the Log Properties.....	79
Sample Configuration File (audit.conf).....	80
Web Server Logging.....	80
Configuring the NetScaler Appliance for Web Server Logging.....	81
Enabling or Disabling Web Server Logging.....	81
Modifying the Default Buffer Size.....	81
Exporting Custom HTTP Headers.....	82
Installing and Configuring the Client System for Web Server Logging.....	82
Installing NSWL Client on a Solaris Operating System.....	84
Installing NSWL Client on a Linux Operating System.....	85
Installing NSWL Client on a FreeBSD Operating System.....	86
Installing NSWL Client on a Mac OS Operating System.....	87
Installing NSWL Client on a Windows Operating System.....	88

Installing NSWL Client on an AIX Operating System.....	90
NSWL Client Command Options.....	91
Adding the IP Addresses of the NetScaler Appliance.....	92
Verifying the NSWL Configuration File.....	93
Running the NSWL Client.....	93
To start Web server logging.....	93
To stop Web server logging started as a background process on the Solaris or Linux operating systems.....	93
To stop Web server logging started as a service on the Windows operating system.....	93
Customizing Logging on the NSWL Client System.....	93
Creating Filters.....	94
Specifying Log Properties.....	95
Understanding the NCSA and W3C Log Formats.....	97
Creating a Custom Log Format.....	102
Sample Configuration File.....	104
Arguments for Defining a Custom Log Format.....	106
Time Format Definition.....	109
Advanced Configurations.....	112
Configuring TCP Window Scaling.....	112
To configure window scaling by using the command line interface.....	113
To configure window scaling by using the configuration utility.....	113
Configuring Selective Acknowledgment (SACK).....	113
To enable Selective Acknowledgment (SACK) by using the command line interface.....	113
To enable Selective Acknowledgment (SACK) by using the configuration utility.....	114
Viewing the HTTP Band Statistics.....	114
To view HTTP request and response size statistics by using the command line interface.....	114
To view HTTP request and response size statistics by using the configuration utility.....	114
To modify the band range by using the command line interface.....	115
To modify the band range by using the configuration utility.....	115
Configuring HTTP Profiles.....	115
To add an HTTP profile by using the command line interface.....	115
To add an HTTP profile by using the configuration utility.....	116
Configuring WebSocket Connections.....	116
Configuring WebSocket connections by using the command line interface....	116
Configuring WebSocket connections by using the configuration utility.....	117

Configuring TCP Profiles.....	117
To add a TCP profile by using the command line interface.....	118
To add a TCP profile by using the configuration utility.....	119
Configuring a Database Profile.....	119
To create a database profile by using the command line interface.....	119
To create a database profile by using the configuration utility.....	120
To bind a database profile to a load balancing or content switching virtual server by using the command line interface.....	120
To bind a database profile to a load balancing or content switching virtual server by using the configuration utility.....	120
Specifying a TCP Buffer Size.....	120
To set the TCP buffer size in an entity-level TCP profile by using the command line interface.....	121
To set the TCP buffer size in a TCP profile by using the configuration utility...	121
Optimizing the TCP Maximum Segment Size for a Virtual Server Configuration....	122
Specifying the MSS Value in a TCP Profile.....	122
Configuring the NetScaler to Learn the MSS Value from Bound Services.....	123
Reporting Tool.....	124
Using the Reporting Tool.....	124
To invoke the Reporting tool.....	124
Working with Reports.....	125
Working with Charts.....	128
Examples.....	132
Stopping and Starting the Data Collection Utility.....	133
To stop nscollect.....	134
To start nscollect on the local system.....	135
To start nscollect on the remote system.....	135
3 AppFlow.....	137
How AppFlow Works.....	139
Flow Records.....	140
Templates.....	140
EIEs for web page performance data.....	141
EIEs for database information	141
Configuring the AppFlow Feature.....	142
Enabling or Disabling AppFlow.....	142
To enable or disable the AppFlow feature by using the command line interface.....	143
To enable the AppFlow feature by using the configuration utility.....	143
Specifying a Collector.....	143

To specify a collector by using the command line interface.....	143
To specify a collector by using the configuration utility.....	143
Configuring an AppFlow Action.....	144
To configure an AppFlow action by using the command line interface.....	144
To configure an AppFlow action by using the configuration utility.....	144
Configuring an AppFlow Policy.....	145
To configure an AppFlow policy by using the command line interface.....	145
To configure an AppFlow policy by using the configuration utility.....	145
To add an expression by using the Add Expression dialog box.....	146
Binding an AppFlow Policy.....	146
To globally bind an AppFlow policy by using the command line interface.....	147
To bind an AppFlow policy to a specific virtual server by using the command line interface.....	147
To globally bind an AppFlow policy by using the configuration utility.....	147
To bind an AppFlow policy to a specific virtual server by using the configuration utility.....	148
Enabling AppFlow for Virtual Servers.....	148
To enable AppFlow for a virtual server by using the command line interface..	148
To enable AppFlow for a virtual server by using the configuration utility.....	148
Enabling AppFlow for a Service.....	149
To enable AppFlow for a service by using the command line interface.....	149
To enable AppFlow for a service by using the configuration utility.....	149
Setting the AppFlow Parameters.....	149
To set the AppFlow Parameters by using the command line interface.....	150
To set the AppFlow parameters by using the configuration utility.....	150
Example: Configuring AppFlow for DataStream.....	150
Exporting Performance Data of Web Pages to AppFlow Collector.....	151
Prerequisites for Exporting Performance Data of Web Pages to AppFlow Collectors.....	152
Associating an AppFlow Action with the EdgeSight Monitoring Responder Policy..	152
To associate an AppFlow action with the EdgeSight Monitoring Responder policy by using the command line interface.....	152
To associate an AppFlow action with the EdgeSight Monitoring Responder policy by using the configuration utility.....	152
Configuring a Virtual Server to Export EdgeSight Statistics to Appflow Collectors.....	153
4 AutoScale: Automatic Scaling in the Citrix CloudPlatform Environment.....	155
How AutoScale Works.....	157
Supported Environment.....	158

Prerequisites.....	158
NetScaler Configuration Details.....	158
Troubleshooting.....	163
The AutoScale configuration was successfully configured in CloudPlatform. Yet, the minimum number of VMs has not been created.	163
The AutoScale configuration is rapidly spawning a large number of VMs.....	164
When I ran the top command on my VM, I noticed that the CPU usage on my VM had breached the threshold that was configured for the scale-up action in AutoScale. Yet, the application is not scaling up.....	164
One or more additional VMs have been created, but they are not accepting traffic (that is, VMs have been created, but the average value of the metrics is still above the threshold)	164
The AutoScale configuration has been deleted, but the VMs continue to exist.....	165
5 EdgeSight Monitoring for NetScaler.....	167
Configuring EdgeSight Monitoring for NetScaler.....	168
To access the wizard from the NetScaler configuration utility and configure EdgeSight Monitoring.....	168
To configure EdgeSight monitoring from the command line interface and configure EdgeSight Monitoring.....	168
Example.....	168
Enabling an Application for EdgeSight Monitoring.....	170
To enable EdgeSight monitoring on a load balancing or content switching virtual server by using the NetScaler configuration utility.....	170
To enable EdgeSight monitoring on a load balancing or content switching virtual server by using the command line interface.....	171
Example.....	171
Accessing the EdgeSight Monitoring Interface from NetScaler.....	172
6 High Availability.....	173
Considerations for a High Availability Setup.....	175
Configuring High Availability.....	176
Adding a Remote Node.....	177
To add a node by using the command line interface.....	177
To disable an HA monitor by using the command line interface.....	178
To add a remote node by using the configuration utility.....	178
Disabling or Enabling a Node.....	178
To disable or enable a node by using the command line interface.....	179
To disable or enable a node by using the configuration utility.....	179
Removing a Node.....	179

To remove a node by using the command line interface.....	179
To remove a node by using the configuration utility.....	179
Configuring the Communication Intervals.....	180
To set the hello and dead intervals by using the command line interface.....	180
To set the hello and dead intervals by using the configuration utility.....	180
Configuring Synchronization.....	180
Disabling or Enabling Synchronization.....	181
To disable or enable automatic synchronization by using the command line interface.....	181
To disable or enable synchronization by using the configuration utility.....	181
Forcing the Secondary Node to Synchronize with the Primary Node.....	181
To force synchronization by using the command line interface.....	182
To force synchronization by using the configuration utility.....	182
Synchronizing Configuration Files in a High Availability Setup.....	182
To synchronize files in a high availability setup by using the command line interface.....	182
To synchronize files in a high availability setup by using the configuration utility....	182
Configuring Command Propagation.....	183
To disable or enable command propagation by using the command line interface..	183
To disable or enable command propagation by using the configuration utility.....	183
Configuring Fail-Safe Mode.....	184
To enable fail-safe mode by using the command line interface.....	185
To enable fail-safe mode by using the configuration utility.....	185
Configuring Virtual MAC Addresses.....	185
Configuring IPv4 VMACs.....	186
Creating or Modifying an IPv4 VMAC.....	186
Removing an IPv4 VMAC.....	187
Configuring IPv6 VMAC6s.....	188
Creating or Modifying a VMAC6.....	188
Removing a VMAC6.....	189
Configuring High Availability Nodes in Different Subnets.....	189
Adding a Remote Node.....	191
To add a node by using the command line interface.....	191
To disable an HA monitor by using the command line interface.....	192
To add a remote node by using the configuration utility.....	192
Removing a Node.....	192
To remove a node by using the command line interface.....	193
To remove a node by using the configuration utility.....	193
Configuring Route Monitors.....	193
Adding a Route Monitor to a High Availability Node.....	193

To add a route monitor by using the command line interface.....	193
To add a route monitor by using the configuration utility.....	194
Removing Route Monitors.....	194
To remove a route monitor by using the command line interface.....	194
To remove a route monitor by using the configuration utility.....	194
Limiting Failovers Caused by Route Monitors in non-INC mode.....	195
Configuring FIS.....	196
Creating or Modifying an FIS.....	197
To add an FIS and bind interfaces to it by using the command line interface..	197
To unbind an interface from an FIS by using the command line interface.....	197
To configure an FIS by using the configuration utility.....	197
Removing an FIS.....	198
To remove an FIS by using the command line interface.....	198
To remove an FIS by using the configuration utility.....	198
Understanding the Causes of Failover.....	198
Forcing a Node to Fail Over.....	199
Forcing Failover on the Primary Node.....	200
To force failover on the primary node by using the command line interface....	200
To force failover on the primary node by using the configuration utility.....	200
Forcing Failover on the Secondary Node.....	200
To force failover on the secondary node by using the command line interface	200
To force failover on the secondary node by using the configuration utility.....	200
Forcing Failover When Nodes Are in Listen Mode.....	201
To force failover when nodes are in listen mode by using the command line	
interface.....	201
To force failover when nodes are in listen mode by using the configuration	
utility.....	201
Forcing the Secondary Node to Stay Secondary.....	201
To force the secondary node to stay secondary by using the command line	
interface.....	202
To force the secondary node to stay secondary by using the configuration utility....	202
Forcing the Primary Node to Stay Primary.....	202
To force the primary node to stay primary by using the command line interface....	202
To force the primary node to stay primary by using the configuration utility.....	203
Understanding the High Availability Health Check Computation.....	203
High Availability.....	203
Troubleshooting High Availability Issues	206
7 Networking.....	211
IP Addressing.....	212

Configuring NetScaler-Owned IP Addresses	212
Configuring the NetScaler IP Address (NSIP)	212
Configuring and Managing Virtual IP Addresses (VIPs)	213
Configuring ARP response Suppression for Virtual IP addresses (VIPs).....	217
Configuring Subnet IP Addresses (SNIPs)	220
Configuring Mapped IP Addresses (MIPs)	223
Configuring GSLB Site IP Addresses (GSLBIP)	225
Removing a NetScaler-Owned IP Address	225
Configuring Application Access Controls	226
How the NetScaler Proxies Connections	228
How the Destination IP Address Is Selected	228
How the Source IP Address Is Selected	229
Enabling Use Source IP Mode	229
Recommended Usage.....	231
To globally enable or disable USIP mode by using the command line interface.....	231
To enable USIP mode for a service by using the command line interface.....	232
To globally enable or disable USIP mode by using the configuration utility....	232
To enable USIP mode for a service by using the configuration utility.....	232
Configuring Network Address Translation	232
Configuring INAT.....	233
Coexistence of INAT and Virtual Servers	235
Stateless NAT46 Translation.....	236
DNS64.....	240
Stateful NAT64 Translation.....	246
Configuring RNAT.....	251
RNAT in USIP, USNIP, and LLB Modes	256
Configuring RNAT for IPv6 Traffic.....	256
Configuring Prefix-Based IPv6-IPv4 Translation.....	257
Configuring Static ARP	259
To add a static ARP entry by using the command line interface.....	259
To remove a static ARP entry by using the command line interface.....	259
To add a static ARP entry by using the configuration utility.....	259
Setting the Timeout for Dynamic ARP Entries.....	259
To set the time-out for dynamic ARP entries by using the command line interface.....	260
To set the time-out for dynamic ARP entries to its default value by using the command line interface.....	260
To set the time-out for dynamic ARP entries by using the configuration utility	260
Configuring Neighbor Discovery	260

Adding IPv6 Neighbors	261
Removing IPv6 Neighbors	262
Configuring IP Tunnels.....	263
NetScaler as an Encapsulator (Load Balancing with DSR Mode).....	263
NetScaler as a Decapsulator.....	263
Creating IP Tunnels.....	264
Customizing IP Tunnels Globally.....	265
Interfaces.....	266
Configuring MAC-Based Forwarding.....	267
To enable or disable MAC-based forwarding by using the command line interface.....	268
Configuring Network Interfaces.....	269
Setting the Network Interface Parameters.....	269
Enabling and Disabling Network Interfaces.....	270
Resetting Network Interfaces.....	271
Monitoring a Network Interface.....	271
Configuring Forwarding Session Rules.....	272
To create a forwarding session rule by using the command line interface.....	273
To configure a forwarding session rule by using the configuration utility.....	273
Understanding VLANs.....	274
Applying Rules to Classify Frames.....	275
Configuring a VLAN.....	276
Creating or Modifying a VLAN.....	277
Monitoring VLANs.....	278
Configuring VLANs in an HA Setup	278
Configuring VLANs on a Single Subnet	279
Configuring VLANs on Multiple Subnets	279
Configuring Multiple Untagged VLANs across Multiple Subnets	280
Configuring Multiple VLANs with 802.1q Tagging.....	281
Configuring NSVLAN.....	283
To configure NSVLAN by using the command line interface.....	283
To restore the default NSVLAN configuration by using the command line interface.....	283
To configure NSVLAN by using the configuration utility.....	284
Configuring Bridge Groups.....	284
To add a bridge group and bind VLANs by using the command line interface.....	284
To remove a bridge group by using the command line interface.....	285
To configure a bridge group by using the configuration utility	285
Configuring VMACs.....	285
Configuring Link Aggregation.....	286

Configuring Link Aggregation by Using the Link Aggregation Control Protocol.....	286
Binding an SNIP address to an Interface.....	289
To configure the example settings.....	291
Monitoring the Bridge Table and Changing the Aging time.....	293
To change the aging time by using the command line interface.....	294
To change the aging time by using the configuration utility.....	294
To view the statistics of a bridge table by using the command line interface... ..	294
To view the statistics of a bridge table by using the configuration utility.....	294
Understanding NetScaler Appliances in Active-Active Mode Using VRRP.....	294
Health Tracking.....	296
Preemption.....	297
Sharing.....	297
Configuring Active-Active Mode.....	297
Adding a VMAC.....	297
Configuring Send to Master.....	299
An Active-Active Deployment Scenario.....	300
Using the Network Visualizer.....	301
To open the Network Visualizer.....	302
To locate a VLAN or bridge group in the Visualizer.....	302
To modify the network settings of the appliance by using the Visualizer.....	303
To add a channel by using the Visualizer.....	303
To add a VLAN by using the Visualizer.....	303
To add a bridge group by using the Visualizer.....	303
To modify the settings of an interface or channel by using the Visualizer.....	303
To enable or disable an interface or channel by using the Visualizer.....	303
To remove a configured channel, VLAN, or bridge group by using the Visualizer.....	304
To view statistics for a node, channel, interface, or VLAN by using the Visualizer.....	304
To set up an HA deployment by using the Visualizer.....	304
To force the secondary node to take over as the primary by using the Visualizer.....	304
To synchronize the secondary node's configuration with the primary node by using the Visualizer.....	304
To remove the peer node from the HA configuration.....	304
To copy the properties of a node or network entity by using the Visualizer.....	304
Access Control Lists.....	304
ACL Precedence	306
Configuring Simple ACLs	306

Creating Simple ACLs	306
Monitoring Simple ACLs	307
Removing Simple ACLs	308
Configuring Extended ACLs	309
Creating and Modifying an Extended ACL	309
Applying an Extended ACL	310
Disabling and Enabling Extended ACLs	310
Renumbering the priority of Extended ACLs	312
Configuring Extended ACL Logging	313
Monitoring the Extended ACL	314
Removing Extended ACLs	315
Configuring Simple ACL6s	316
Creating Simple ACL6s	316
Monitoring Simple ACL6s	317
Configuring ACL6s	318
Creating and Modifying ACL6s	318
Applying ACL6s	319
Enabling and Disabling ACL6s	320
Renumbering the Priority of ACL6s	321
Monitoring ACL6s	322
Removing ACL6s	323
Terminating Established Connections	323
To terminate all established IPv4 connections that match any of your configured simple ACLs by using the command line interface	324
To terminate all established IPv4 connections that match any of your configured simple ACLs by using the configuration utility	324
To terminate all established IPv6 connections that match any of your configured simple ACL6s by using the command line interface	324
To terminate all established IPv6 connections that match any of your configured simple ACL6s by using the configuration utility	324
IP Routing	325
Configuring Static Routes	325
Weighted Static Routes	325
Null Routes	325
Configuring IPv4 Static Routes	325
Configuring IPv6 Static Routes	327
Configuring Policy-Based Routes	328
Configuring a Policy-Based Routes (PBR) for IPv4 Traffic	329
Configuring a Policy-Based Routes (PBR6) for IPv6 Traffic	337
Internet Protocol version 6 (IPv6)	339

Implementing IPv6 Support.....	341
To enable or disable IPv6 by using the command line interface.....	341
To enable or disable IPv6 by using the configuration utility.....	342
VLAN Support.....	342
Simple Deployment Scenario.....	342
To create IPv4 services by using the command line interface.....	344
To create IPv4 services by using the configuration utility.....	344
To create IPv6 vserver by using the command line interface.....	345
To create IPv6 vserver by using the configuration utility.....	345
To bind a service to an LB vserver by using the command line interface.....	345
To bind a service to an LB vserver by using the configuration utility.....	346
Host Header Modification.....	346
To change the IPv6 address in the host header to an IPv4 address by using the command line interface.....	346
To change the IPv6 address in the host header to an IPv4 address by using the configuration utility.....	346
VIP Insertion.....	347
To configure a mapped IPv6 address by using the command line interface....	347
To configure a mapped IPv6 address by using the configuration utility.....	347
To enable VIP insertion by using the command line interface.....	347
To enable VIP insertion by using the configuration utility.....	348
Traffic Domains.....	348
Benefits of using Traffic Domains.....	348
Default Traffic Domain	349
How Traffic Domains Work.....	349
Supported NetScaler Features in Traffic Domains.....	352
Configuring Traffic Domains.....	353
To create a VLAN and bind interfaces to it by using the command line interface.....	353
To create a traffic domain entity and bind VLANs to it by using the command line interface.....	353
To create a service by using the command line interface.....	354
To create a load balancing virtual server and bind services to it by using the command line interface.....	354
To create a VLAN by using the configuration utility.....	354
To create a traffic domain entity by using the configuration utility.....	354
To create a service by using the configuration utility.....	354
To create a load balancing virtual server and bind services to it by using the configuration utility.....	355

8	Web Interface.....	357
	How Web Interface Works.....	358
	Prerequisites.....	358
	Installing the Web Interface.....	359
	To install the Web interface and JRE tar files by using the command line interface	359
	To install the Web interface and JRE tar files by using the configuration utility.....	360
	Configuring the Web Interface.....	360
	Configuring a Web Interface Site for LAN Users Using HTTP.....	361
	To configure a Web interface site for LAN users using HTTP by using the configuration utility.....	361
	To configure a Web interface site for LAN users using HTTP by using the command line interface.....	363
	Configuring a Web Interface Site for LAN Users Using HTTPS.....	365
	To configure a Web interface site for LAN users using HTTPS by using the configuration utility.....	365
	To configure a Web interface site for LAN users using HTTPS by using the command line.....	368
	Using the WebInterface.conf Dialog Box.....	370
	To search a string in the webinterface.conf file by using the configuration utility....	370
	To save the content of the webinterface.conf to your local system by using the configuration utility.....	371
	Using the config.xml Dialog Box.....	371
	To search a string in the config.xml file by using the configuration utility.....	371
	To save the content of the config.xml to the local system by using the configuration utility.....	372

Chapter 1

Basic Operations

Topics:

- *Viewing and Saving Configurations*
- *Clearing the NetScaler Configuration*
- *Configuring Clock Synchronization*
- *Viewing the System Date and Time*
- *Backing up and Restoring the NetScaler Appliance*
- *Restarting or Shutting down the Appliance*

Any changes you make to the configuration of a NetScaler appliance are temporary until you save the new configuration in the `/nsconfig/ns.conf` directory. An unsaved configuration (the *running configuration*) is replaced by the most recently saved configuration when the appliance restarts.

Viewing and Saving Configurations

Running configurations includes both saved and unsaved configurations. NetScaler configurations must be saved frequently to make sure that the configurations are not lost.

To view the running configuration by using the command line interface

At the command prompt, type:

```
show ns runningConfig
```

To view the running configuration by using the configuration utility

1. Navigate to **System > Diagnostics**.
2. In the details pane, under **View Configuration**, click **Running Configuration**.

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

At the command prompt, type:

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

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

1. Navigate to **System > Diagnostics**.
2. In the details pane, under **View Configuration**, click **Configuration difference**.

To save configurations by using the command line interface

At the command prompt, type:

```
save ns config
```

To save configurations by using the configuration utility

In the configuration utility, click the save icon at the top right corner on the home page of the **Configuration** tab.

To view the saved configurations by using the command line interface

At the command prompt, type:

```
show ns ns.conf
```

To view the saved configurations by using the configuration utility

1. Navigate to **System > Diagnostics**.
2. In the details pane, under **View Configuration**, click **Saved Configuration**.

Clearing the NetScaler Configuration

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

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

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

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

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

Feature and mode settings revert to their default values.

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

To clear the configuration by using the command line interface

At the command prompt, type:

clear ns config -force <level>

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

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

To clear the configuration by using the configuration utility

1. Navigate to **System > Diagnostics**.
2. In the details pane, under **Maintenance**, click **Clear Configuration**.
3. In the **Clear Configuration** dialog box, select the level of configurations to be cleared from the appliance.
4. Click **Run**.

Configuring Clock Synchronization

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

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

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

Setting Up Clock Synchronization

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

To add an NTP server by using the command line interface

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

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

Example

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

To configure an NTP server by using the configuration utility

1. Navigate to **System > NTP Servers**.
2. In the details pane, click **Add**.
3. In the **Create NTP Server** dialog box, configure the NTP server. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create**, and then click **Close**.

Starting or Stopping the NTP Daemon

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

To enable or disable NTP synchronization by using the command line interface

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

- ♦ **enable ntp sync**
- ♦ **disable ntp sync**

To enable or disable NTP synchronization by using the configuration utility

1. Navigate to **System > NTP Servers**.
2. In the details pane, click **Action** and select **NTP Synchronization**.
3. In the **Configure NTP Synchronization** dialog box, enable or disable NTP synchronization for the appliance.
4. Click **OK**.

Configuring Clock Synchronization Manually

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

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

1. Log on to the command line interface.
2. Switch to the shell prompt.
3. Copy the `/etc/ntp.conf` file to `/nsconfig/ntp.conf`, unless the `/nsconfig` directory already contains an `ntp.conf` file.
4. Check the `/nsconfig/ntp.conf` file for the following entries and, if they are present, remove them:
`restrict localhost`
`restrict 127.0.0.2`
5. Add the IP address for the desired NTP server to the `/nsconfig/ntp.conf` file, beneath the file's `server` and `restrict` entries.

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

6. If the `/nsconfig` directory does not contain a file named `rc.netscaler`, create the file.
7. Add the following entry to `/nsconfig/rc.netscaler`:

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

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

This process runs every time the NetScaler is restarted.
8. Reboot the NetScaler to enable clock synchronization.

Note:

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

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

Viewing the System Date and Time

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

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

At the command prompt, type:

```
show ns config
```

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

1. In the navigation pane, expand **System**.
2. In the details pane, select the **System Information** tab.
3. Under **System Information**, view the system date and time.

Backing up and Restoring the NetScaler Appliance

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

Points to remember

- ♦ You cannot use the backup file taken from one appliance to restore a different appliance.
- ♦ You can back up and restore appliances in an HA setup, but make sure that you restore to the same appliance from which the backup file was created. For example, if the backup was taken from the primary appliance of the HA pair, when

restoring make sure that the appliance you are restoring is the same appliance, even if it is no longer the primary appliance.

- ♦ You cannot perform the backup and restore operation on a NetScaler cluster.

Backing up a NetScaler Appliance

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

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

Directory	Sub-Directory or Files
/nsconfig/	<ul style="list-style-type: none"> • ns.conf • ZebOS.conf • rc.netscaler • snmpd.conf • nsbefore.sh • nsafter.sh • monitors
/var/	<ul style="list-style-type: none"> • download/* • log/wicmd.log • wi/tomcat/webapps/* • wi/tomcat/logs/* • wi/tomcat/conf/catalina/localhost/* • nslw.bin/etc/krb.conf • nslw.bin/etc/krb.keytab • netscaler/locdb/* • lib/likewise/db/* • vpn/bookmark/* • netscaler/crl • nstemplates/* • learnt_data/*

Directory	Sub-Directory or Files
/netscaler/	<ul style="list-style-type: none"> • custom.html • vsr.htm

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

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

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

Note:

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

To backup the NetScaler by using the command line interface

At the command prompt, do the following:

1. Save the NetScaler configurations.
save ns config
2. Create the backup file.
create system backup [`<fileName>`] `-level` `<basic | full>` `-comment` `<string>`

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

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

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

3. Verify that the backup file was created.

show system backup

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

To backup the NetScaler by using the configuration utility

1. Navigate to **System > Backup and Restore**.
2. In the details pane, click **Backup**.
3. In the **Backup** screen, specify the details required to backup the appliance.
4. Click **Backup**.

Restoring the NetScaler Appliance

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



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

To restore the NetScaler by using the command line interface

At the command prompt, do the following:

1. Obtain a list of the backup files available on the appliance.
show system backup
2. Restore the appliance by specifying one of the backup files.
restore system backup -fileName <filename>

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

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

3. Reboot the appliance.
reboot

To restore the NetScaler by using the configuration utility

1. Navigate to **System > Backup and Restore**.
2. In the details pane, select the backup file to be restored on the appliance and click **Restore**.
3. Review the details of the backup file and click **Restore** to confirm the operation.

Restarting or Shutting down the Appliance

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

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

You can also restart the appliance by only rebooting the NetScaler software and not rebooting the underlying operating system. This is called a warm reboot.

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

To restart the NetScaler by using the command line interface

At the command prompt, type:

```
reboot [-warm]
```

To restart the NetScaler by using the configuration utility

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

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

3. Click **OK**.

To shut down the NetScaler by using the command line interface

At the command prompt, type:

shutdown

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

Chapter 2

Administration

Topics:

- [Authentication and Authorization](#)
- [SNMP](#)
- [Audit Logging](#)
- [Web Server Logging](#)
- [Advanced Configurations](#)
- [Reporting Tool](#)

The following topics provide a conceptual reference and instructions for managing and monitoring the Citrix NetScaler appliance by using built-in features, such as command policies, Simple Network Management (SNMP), audit logging, web server logging, Network Time Protocol (NTP), and the Reporting tool.

Authentication and Authorization	Configure authentication and authorization to manage access to the NetScaler and different parts of the NetScaler configuration.
SNMP	Learn how SNMP works with NetScaler and how to configure SNMP V1, V2, and V3 on NetScaler.
Audit Logging	Configure the NetScaler audit server log to log and monitor the NetScaler states and status information. Also, learn how to configure audit server logging on a server system and for a deployment scenario.
Web Server Logging	Configure web server log to maintain a history of the page requests that originate from the NetScaler.
Advanced Configurations	Learn how to set advanced configurations, such as NTP, PMTU, and auto detected services, on the NetScaler.

Reporting Tool	Learn how to use the Reporting tool to view performance statistics as reports with graphs that are based on statistics collected by the nscollect utility.
----------------	--

Authentication and Authorization

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

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

Configuring Users and Groups

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

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

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

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

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

Configuring User Accounts

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

To create a user account by using the command line interface

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

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

Example

```
> add system user johnd -promptString user-%u-at-%T
Enter password:
Confirm password:
```

To configure a user account by using the configuration utility

1. Navigate to **System > User Administration > Users**.
2. In the details pane, click **Add**.
3. In the **Create System User** dialog box, configure the user account. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create** and then click **Close**.

Configuring User Groups

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

To create a user group by using the command line interface

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

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

Example

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

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

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

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

Example

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

To configure a user group by using the configuration utility

1. Navigate to **System > User Administration > Groups**.
2. In the details pane, click **Add**.
3. In the **Create System Group** dialog box, configure the user group. For a description of a parameter, hover the mouse cursor over the corresponding field.

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

4. Click **Create** and then click **Close**.

Configuring Command Policies

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

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

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

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

Built-in Command Policies

The following table describes the built-in policies.

Table 2-1. Built-in Command Policies

Policy name	Allows
read-only	Read-only access to all show commands except show ns runningConfig , show ns ns.conf , and the show commands for the NetScaler command group.
operator	Read-only access and access to commands to enable and disable services and servers.
network	Full access, except to the set and unset SSL commands, show ns ns.conf , show ns runningConfig , and show gslb runningConfig commands.
superuser	Full access. Same privileges as the nsroot user.

Creating Custom Command Policies

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

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

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

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

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

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

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

The following table lists examples of regular expressions:

Table 2-2. Examples of Regular Expressions for Command Policies

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

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

Table 2-3. Expressions Used in the Built-in Command Policies

Policy name	Command specification regular expression
read-only	(^man.*) (^show\s+(?!system)(?!configstatus)(?!ns ns\.conf)(?!ns savedconfig)(?!ns runningConfig)(?!gslb runningConfig)(?!audit messages)(?!techsupport).*) (^stat.*)
operator	(^man.*) (^show\s+(?!system)(?!configstatus)(?!ns ns\.conf)(?!ns savedconfig)(?!ns runningConfig)(?!gslb runningConfig)(?!audit messages)(?!techsupport).*) (^stat.*) (^enable disable) (server service).*)

Policy name	Command specification regular expression
network	^(?!clear ns config.*)(?!scp.*)(?!set ssl fips)(?!reset ssl fips)(?!diff ns config)(?!shell)(?!reboot)(?!batch)\S+\s+(?!system)(?!configstatus)(?!ns ns\.conf)(?!ns savedconfig)(?!ns runningConfig)(?!gslb runningConfig)(?!techsupport).*
superuser	.*

To create a command policy by using the command line interface

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

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

Example

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

To configure a command policy by using the configuration utility

1. Navigate to **System > User Administration > Command Policies**.
2. In the details pane, click **Add**.
3. In the **Create Command Policy** dialog box, configure the command policy. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create** and then click **Close**.

Binding Command Policies to Users and Groups

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

Command policies are evaluated in the following order:

- ♦ Command policies bound directly to users and the corresponding groups are evaluated according to priority number. A command policy with a lower priority number is evaluated before one with a higher priority number. Therefore, any privileges the lower-numbered command policy explicitly grants or denies are not overridden by a higher-numbered command policy.

- ♦ When two command policies, one bound to a user account and other to a group, have the same priority number, the command policy bound directly to the user account is evaluated first.

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

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

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

Example

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

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

1. Navigate to **System > User Administration > Users**.
2. In the details pane, select the user to which you want to bind a command policy, and then click **Open**.
3. In the **Configure System User** dialog box, under **Command Policies**, all of the command policies configured on the appliance appear on the list. Select the check box next to the name of the policy you want to bind to this user.
4. In the **Priority** column to the left, modify the default priority as needed to ensure that the policy is evaluated in the proper order.
5. Click **OK**.

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

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

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

Example

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

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

1. Navigate to **System > User Administration > Groups**.

2. In the details pane, select the group to which you want to bind a command policy, and then click **Open**.
3. In the **Configure System Group** dialog box, under **Command Policies**, all the command policies configured on the appliance appear on the list. Select the check box next to the name of the policy you want to bind to this group.
4. In the **Priority** column to the left, modify the default priority as needed to ensure that the policy is evaluated in the proper order.
5. Click **OK**.

Resetting the Default Administrator (nsroot) Password

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

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

To reset the nsroot password

1. Connect a computer to the serial port of the appliance and log on.

Note: You cannot log on by using ssh to perform this procedure; you must connect directly to the appliance.

As the operating system starts, it displays the following message:

```
Hit [Enter] to boot immediately, or any other key for
command prompt.
```

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

2. Press CTRL+C.

The following message appears:

Type '?' for a list of commands, 'help' for more detailed help.

ok

3. Type **boot -s** and press the ENTER key to start the appliance in single user mode.

After the appliance boots, it displays the following message:

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

4. Press the ENTER key to display the # prompt, and type the following commands to mount the file systems:


```
fsck /dev/ad0s1a

mount /dev/ad0s1a/flash
```
5. Using a text editor of your choice, edit the `/nsconfig/ns.conf` file and remove the set system user nsroot entry.
6. Save the file and exit the text editor.
7. Type **reboot** and press the ENTER key to reboot the appliance.

When the appliance completes rebooting, it prompts for the user name and password.
8. Log on with the nsroot user credentials.

Once logged on to the appliance, you will be required to enter a new nsroot user password.
9. Follow the prompts to change the password.
10. Exit the **config ns** menu.

Example of a User Scenario

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

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

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

Table 2-4. Sample Values for Creating Entities

Field	Value	Note
NetScaler host name	ns01.example.net	N/A
User accounts	johnd, mariar, and michaelb	John Doe, IT manager, Maria Ramirez, IT

Field	Value	Note
		administrator and Michael Baldrock, IT administrator.
Groups	Managers and SysOps	All managers and all IT administrators.
Command Policies	read_all, modify_lb, and modify_all	Allow complete read-only access, Allow modify access to load balancing, and Allow complete modify access.

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

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

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

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

Configuration steps

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

The configuration you just created results in the following:

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

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

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

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

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

Configuring External User Authentication

External user authentication is the process of authenticating the users of the Citrix NetScaler appliance by using an external authentication server. The NetScaler supports LDAP, RADIUS, and TACACS+ authentication servers. To configure external user authentication, you must create authentication policies. You can configure one or many authentication policies, depending on your authentication needs. An authentication policy consists of an expression and an action.

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

Note: User accounts must be configured on the NetScaler appliance before users can be externally authenticated. You must first create an onboard system user for all users who will access the appliance, so that you can bind command policies to the user accounts. Regardless of the authentication source, users cannot log on if they are not granted sufficient command authorization through command policies bound to their user accounts or to a group of which they are a member.

Configuring LDAP Authentication

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

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

The port numbers for LDAP connections are:

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

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

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

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

Table 2-5. User Attribute Fields for LDAP Servers

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

LDAP server	User attribute	Case sensitive?
Lotus Domino	CN	Yes
Sun ONE directory (formerly iPlanet)	uid or cn	Yes

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

Table 2-6. Examples of Base Distinguished Name

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

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

Table 2-7. Examples of Bind Distinguished Name

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

To configure LDAP authentication by using the configuration utility

1. Navigate to **System > Authentication**.
2. On the **Policies** tab, click **Add**.
3. In **Name**, type a name for the policy.
4. In **Authentication Type**, select **LDAP**. Next to **Server**, click **New**.
5. In **Name**, type the name of the server.
6. Under **Server**, in **IP Address** and **Port**, type the IP address and port number of the LDAP server.
7. Under **Connection Settings**, provide the following information:

- In **Base DN (location of users)**, type the base DN under which users are located. Base DN is usually derived from the Bind DN by removing the user name and specifying the group where in which are located. Examples of syntax for base DN are:

```
ou=users, dc=ace, dc=com
cn=Users, dc=ace, dc=com
```

- In **Administrator Bind DN**, type the administrator bind DN for queries to the LDAP directory. Examples for syntax of bind DN are:

```
domain/user name
ou=administrator, dc=ace, dc=com
user@domain.name (for Active Directory)
cn=Administrator, cn=Users, dc=ace, dc=com
```

For Active Directory, the group name specified as `cn=groupname` is required. The group name that is defined in the appliance must be identical to the group name that is defined on the LDAP server. For other LDAP directories, the group name either is not required or, if required, is specified as `ou=groupname`.

The appliance binds to the LDAP server, using the administrator credentials, and then searches for the user. After locating the user, the appliance unbinds the administrator credentials and rebinds with the user credentials.

- In **Administrator Password** and **Confirm Administrator Password**, type the administrator password for the LDAP server.
8. To retrieve additional LDAP settings automatically, click **Retrieve Attributes**. The fields under **Other Settings** then populate automatically. If you do not want to do this, skip to Step 12.
 9. Under **Other Settings**, in **Server Logon Name Attribute**, type the attribute under which the appliance should look for user logon names for the LDAP server that you are configuring. The default is `samAccountName`.

10. In **Group Attribute**, leave the default `memberOf` for Active Directory or change it to that of the LDAP server type you are using. This attribute enables the appliance to obtain the groups associated with a user during authorization.
11. In **Security Type**, select the security type.
If you select **PLAINTEXT** or **TLS** for security, use port number 389. If you select **SSL**, use port number 636.
12. To allow users to change their LDAP password, select **Allow Password Change**.
If you select **PLAINTEXT** as the security type, allowing users to change their passwords is not supported.
13. Click **Create**.
14. In the **Create Authentication Policy** dialog box, next to **Named Expressions**, select the expression, click **Add Expression**, click **Create**, and click **Close**.

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

Determining attributes in the LDAP directory

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

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

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

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

Configuring RADIUS Authentication

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

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

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

To configure RADIUS authentication by using the configuration utility

1. Navigate to **System > Authentication**.
2. On the **Policies** tab, click **Add**.
3. In **Name**, type a name for the policy.
4. In **Authentication Type**, select **RADIUS**.
5. Next to **Server**, click **New**.
6. In **Name**, type a name for the server.
7. Under **Server**, in **IP Address**, type the IP address of the RADIUS server.
8. In **Port**, type the port. The default is 1812.
9. Under **Details**, in **Secret Key** and **Confirm Secret Key**, type the RADIUS server secret.
10. In **NAS ID**, type the identifier number, and then click **Create**.
11. In the **Create Authentication Policy** dialog box, next to **Named Expressions**, select the expression, click **Add Expression**, click **Create**, and click **Close**.

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

Choosing RADIUS authentication protocols

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

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

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

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

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

Configuring IP address extraction

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

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

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

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

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

To configure IP address extraction by using the configuration utility

1. Navigate to **System > Authentication**.
2. On the **Policies** tab, select one of the policies and click **Open**.
3. In the **Configure Authentication Policy** dialog box, next to **Server**, click **Modify**.
4. Under **Details**, enter the value for the **Group Vendor Identifier** and **Group Attribute Type** fields.
5. Click **OK** twice.

Configuring TACACS+ Authentication

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

To configure TACACS+ authentication by using the configuration utility

1. Navigate to **System > Authentication**.
2. On the **Policies** tab, click **Add**.
3. In **Name**, type a name for the policy.
4. In **Authentication Type**, select **TACACS**.
5. Next to **Server**, click **New**.
6. In **Name**, type a name for the server.
7. Under **Server**, type the IP address and port number of the TACACS+ server.
8. Under **TACACS server information**, in **TACACS Key** and **Confirm TACACS key**, type the key.
9. In **Authorization**, select **ON** and click **Create**.
10. In the **Create Authentication Policy** dialog box, next to **Named Expressions**, select the expression, click **Add Expression**, click **Create**, and click **Close**.

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

Binding the Authentication Policies to the System Global Entity

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

To bind an authentication policy globally by using the configuration utility

1. Navigate to **System > Authentication**.
2. On the **Policies** tab, click **Global Bindings**.
3. Under **Details**, click **Insert Policy**.
4. Under **Policy Name**, select the policy and click **OK**.

To unbind a global authentication policy by using the configuration utility

1. Navigate to **System > Authentication**.
2. On the **Policies** tab, click **Global Bindings**.
3. In the **Bind/Unbind Authentication Policies** dialog box, in **Policy Name**, select the policy, click **Unbind Policy** and then click **OK**.

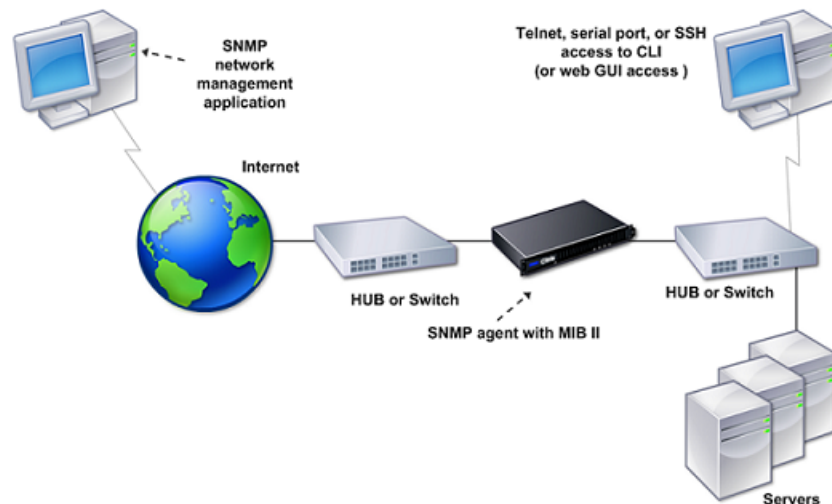
SNMP

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

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

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

Figure 2-1. NetScaler Supporting SNMP



Importing MIB Files to the SNMP Manager and Trap Listener

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

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

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

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

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

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

- ♦ `mib.txt`
- ♦ `traps.txt`

Configuring the NetScaler to Generate SNMPv1 and SNMPv2 Traps

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

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

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

Enabling or Disabling an SNMP Alarm

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

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

To enable or disable an SNMP alarm by using the command line interface

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

- ♦ `enable snmp alarm <alarm name>`
- ♦ `show snmp alarm <alarm name>`

To enable or disable an SNMP alarm by using the configuration utility

1. Navigate to **System > SNMP > Alarms**.
2. In the details pane, select an alarm (for example, **Login-Failure**), and do one of the following:
 - To enable an alarm, click **Enable**.
 - To disable an alarm, click **Disable**.

Configuring Alarms

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

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

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

- ♦ Critical
- ♦ Major
- ♦ Minor
- ♦ Warning
- ♦ Informational

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

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

To configure an SNMP alarm by using the command line interface

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

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

To configure SNMP alarms by using the configuration utility

1. Navigate to **System > SNMP > Alarms**.

2. In the details pane, select an alarm (for example, **Login-Failure**), and then click **Open**.
3. In the **Configure SNMP Alarm** dialog box, configure the SNMP alarm. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **OK**.

Configuring Traps

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

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

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

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

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

To add an SNMP trap by using the command line interface

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

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

Example

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

To configure SNMP Traps by using the configuration utility

1. Navigate to **System > SNMP > Traps**.
2. In the details pane, click **Add**.
3. In the **Create SNMP Trap Destination** dialog box, configure the SNMP trap. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create** and then click **Close**.

Enabling Unconditional SNMP Trap Logging

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

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

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

- ♦ **set snmp option -snmpTrapLogging (ENABLED | DISABLED)**
- ♦ **show snmp option**

To enable or disable unconditional SNMP trap logging by using the configuration utility

1. Navigate to **System > SNMP**.
2. In the details pane, under **Settings**, click **Change SNMP Options**.
3. In the **Configure SNMP Options** dialog box, select the **SNMP Trap Logging** check box.
4. Click **OK**.

Configuring the NetScaler for SNMP v1 and v2 Queries

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

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

- ♦ GET
- ♦ GET NEXT

- ♦ ALL
- ♦ GET BULK

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

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

Specifying an SNMP Manager

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

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

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

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

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

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

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

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

Example

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

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

Important: If you specify the SNMP manager's host name instead of its IP address, you must configure a DNS name server to resolve the host name to the SNMP manager's IP address.

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

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

Example

```
> add nameserver 10.103.128.15  
  
> add snmp manager engwiki.eng.example.net -  
domainResolveRetry 10
```

To add an SNMP manager by using the configuration utility

1. Navigate to **System > SNMP > Managers**.
2. In the details pane, click **Add**.
3. In the **Create SNMP Manager** dialog box, do one of the following:
 - To specify the host name of an SNMP manager, select **Management Host** and set the parameters.

Important: If you specify the SNMP manager's host name instead of its IPv4 address, you must configure a DNS name server to resolve the host name to the SNMP manager's IP address.

Note: The appliance does not support host names for SNMP managers that have IPv6 addresses.

- To specify the IPv4 or IPv6 address of an SNMP manager, select **Management Network** and set the parameters.
4. Click **Create**, and then click **Close**.

Specifying an SNMP Community

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

- ♦ GET
- ♦ GET NEXT
- ♦ ALL
- ♦ GET BULK

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

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

To specify an SNMP community by using the command line interface

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

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

Example

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

To configure an SNMP community string by using the configuration utility

1. Navigate to **System > SNMP > Community**.
2. In the details pane, click **Add**.
3. In the **Create SNMP Community** dialog box, configure the parameters for the SNMP community. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create**, and then click **Close**.

Configuring SNMP Alarms for Rate Limiting

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

Throughput and PPS are monitored every seven seconds. You can configure traps with high-threshold and normal-threshold values, which are expressed as a percentage of the licensed limits. The appliance then generates a trap when throughput or PPS

exceeds the high threshold, and a second trap when the monitored parameter falls to the normal threshold. In addition to sending the traps to the configured destination device, the NetScaler logs the events associated with the traps in the `/var/log/ns.log` file as `EVENT ALERTSTARTED` and `EVENT ALERTENDED`.

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

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

Configuring an SNMP Alarm for Throughput or PPS

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

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

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

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

Example

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

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

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

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

Example

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

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

1. Navigate to **System > SNMP > Alarms**.
2. In the details pane, do one of the following:
 - Select **PF-RL-RATE-THRESHOLD** to configure the SNMP alarm for throughput rate.
 - Select **PF-RL-PPS-THRESHOLD** to configure the SNMP alarm for packets per second.
3. Click **Open**.
4. In the **Configure SNMP Alarm** dialog box, set the parameters for the SNMP alarm and select the **Enable** check box. For a description of a parameter, hover the mouse cursor over the corresponding fields.
5. Click **OK**, and then click **Close**.

Configuring SNMP Alarm for Dropped Packets

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

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

At the command prompt, type:

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

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

At the command prompt, type:

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

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

1. Navigate to **System > SNMP > Alarms**.
2. In the details pane, do one of the following:
 - Select **PF-RL-RATE-PKTS-DROPPED** to configure an SNMP alarm for packets dropped because of excessive throughput.
 - Select **PF-RL-PPS-PKTS-DROPPED** to configure an SNMP alarm for packets dropped because of excessive PPS.
3. Click **Open**.

4. In the **Configure SNMP Alarm** dialog box, set the parameters for the SNMP alarm and select the **Enable** check box. For a description of a parameter, hover the mouse cursor over the corresponding field.
5. Click **OK**, and then click **Close**.

Configuring the NetScaler for SNMPv3 Queries

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

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

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

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

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

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

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

To implement message authentication and access control, you need to:

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

Setting the Engine ID

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

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

To set the engine ID by using the command line interface

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

- ♦ `set snmp engineId <engineId>`
- ♦ `show snmp engineId`

Example

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

To set the engine ID by using configuration utility

1. Navigate to **System > SNMP > Users**.
2. In the details pane, click the **Action** drop-down list and select **Configure Engine ID**.
3. In the **Configure Engine ID** dialog box, in the **Engine ID** text box, type an engine ID (for example, 8000173f0300c095f80c68).
4. Click **OK**.

Configuring a View

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

To add an SNMP view by using the command line interface

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

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

Example

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

To configure an SNMP view by using the configuration utility

1. Navigate to **System > SNMP > Views**.
2. In the details pane, click **Add**.
3. In the **Create SNMP View** dialog box, configure the view. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create**, and then click **Close**.

Configuring a Group

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

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

To add an SNMP group by using the command line interface

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

- ♦ **add snmp group** <name> <securityLevel> -readViewName <string>
- ♦ **show snmp group** <name> <securityLevel>

Example

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

To configure an SNMP group by using the configuration utility

1. Navigate to **System > SNMP > Groups**.

2. In the details pane, click **Add**.
3. In the **Create SNMP Group** dialog box, configure the group. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create**, and then click **Close**.

Configuring a User

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

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

To configure a user by using the command line interface

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

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

Example

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

To configure an SNMP user by using the configuration utility

1. Navigate to **System > SNMP > Users**.
2. In the details pane, click **Add**.
3. In the **Create SNMP User** dialog box, configure the users. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create**, and then click **Close**.

Audit Logging

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

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

Similarly, the native NSLOG protocol has two components— the NSLOG auditing module, which runs on the NetScaler appliance, and the NSLOG server, which can run on the

underlying FreeBSD OS of the NetScaler appliance or on a remote system. NSLOG uses transmission control protocol (TCP) for transfer of data.

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

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

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

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

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

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

Configuring the NetScaler Appliance for Audit Logging

Policies define the SYSLOG or NSLOG protocol, and server actions define what logs are sent where. For server actions, you specify the system information, which runs the SYSLOG or the NSLOG server.

The NetScaler logs the following information related to TCP connections:

- ♦ Source port
- ♦ Destination port
- ♦ Source IP

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

Note: You can enable TCP logging on individual load balancing vservers. You must bind the audit log policy to a specific load balancing vserver that you want to log.

Configuring Audit Servers

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

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

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

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

Example

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

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

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

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

Example

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

To configure an auditing server action by using the configuration utility

1. Navigate to **System > Auditing > Syslog** or **System > Auditing > Nslog**.
2. In the details pane, on the **Servers** tab, click **Add**.
3. In the **Create Auditing Server** dialog box, configure the auditing server. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create**, and then click **Close**.

Configuring Audit Policies

The audit policies define the SYSLOG or NSLOG protocol.

To configure a SYSLOG policy by using the command line interface

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

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

Example

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

To configure an NSLOG policy by using the command line interface

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

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

Example

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

To configure an audit server policy by using the configuration utility

1. Navigate to **System > Auditing > Syslog** or **System > Auditing > Nslog**.
2. In the details pane, on the **Policies** tab, click **Add**.

3. In the **Create Auditing Policy** dialog box, configure the audit policy. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create**, and then click **Close**.

Binding the Audit Policies Globally

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

To configure a SYSLOG policy by using the command line interface

At the command prompt, type:

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

Example

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

To globally bind the audit policy by using the configuration utility

1. Navigate to **System > Auditing > Syslog** or **System > Auditing > Nslog**.
2. In the details pane, on the **Policies** tab, click the **Action** drop-down list and select **Global Bindings**.
3. In the **Bind/Unbind Auditing Global Policies** dialog box, click **Insert Policy**.
4. Select a policy from the drop-down list that appears under **Policy Name**, and click **OK**.

Configuring Policy-Based Logging

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

Pre Requisites

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

- ♦ The related audit policy is bound to system global. For more information about binding audit policies to system global, see ["Binding the Audit Policies Globally."](#)

Configuring an Audit Message Action

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

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

At the command prompt, type:

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

Example

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

To configure an audit message action by using the configuration utility

1. Navigate to **System > Auditing > Message Actions**.
2. In the details pane, click **Add**.
3. In the **Create Message Action** dialog box, configure the message action. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create**, and then click **Close**.

Binding Audit Message Action to a Policy

After you have created an audit message action, you must bind it to a rewrite or responder policy.

Installing and Configuring the NSLOG Server

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

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

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

Table 2-8. Supported Platforms for the NSLOG Server

Operating system	Software requirements
Windows	<ul style="list-style-type: none">♦ Windows XP Professional♦ Windows Server 2003♦ Windows 2000/NT♦ Windows Server 2008♦ Windows Server 2008 R2
Linux	<ul style="list-style-type: none">♦ RedHat Linux 4 or later♦ SUSE Linux Enterprise 9.3 or later
FreeBSD	<ul style="list-style-type: none">♦ FreeBSD 6.3
Mac OS	<ul style="list-style-type: none">♦ Mac OS 8.6 or later

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

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

Installing NSLOG Server on the Linux Operating System

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

To install the NSLOG server package on a Linux operating system

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

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

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

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

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

- `/usr/local/netscaler/etc`
- `/usr/local/netscaler/bin`

- /usr/local/netscaler/samples

To uninstall the NSLOG server package on a Linux operating system

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

```
rpm -e NSauditserver
```

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

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

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

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

*.rpm: Specifies the file name.

Installing NSLOG Server on the FreeBSD Operating System

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

Note: NSLOG server is not supported on the underlying FreeBSD OS of the NetScaler appliance.

To download NSLOG package from www.Citrix.com

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

To install the NSLOG server package on a FreeBSD operating system

1. On the system to which you have downloaded the NSLOG package `AuditServer_<release number>-<build number>.zip` (for example, `AuditServer_9.3-51.5.zip`), extract the FreeBSD NSLOG server package `audserver_bsd-<release number>-<build number>.tgz` (for example, `audserver_bsd-9.3-51.5.tgz`) from the package.

2. Copy the FreeBSD NSLOG server package `audserver_bsd-<release number>-<build number>.tgz` (for example, `audserver_bsd-9.3-51.5.tgz`) to a directory on a system running FreeBSD OS.
3. At a command prompt for the directory into which the FreeBSD NSLOG server package was copied, run the following command to install the package:

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

Example

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

The following directories are extracted:

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

```
pkg_info | grep NSaudserver
```

To uninstall the NSLOG server package on a FreeBSD operating system

At a command prompt, type:

```
pkg_delete NSaudserver
```

Installing NSLOG Server Files on the Windows Operating System

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

To download NSLOG package from www.Citrix.com

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

number>.zip , to your local system (for example, AuditServer_9.3-51.5.zip).

To install NSLOG server on a Windows operating system

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

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

To uninstall the NSLOG server on a Windows operating system

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

```
audserver -remove
```

NSLOG Server Command Options

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

Table 2-9. Audit Server Options

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

Audit server commands	Specifies
<code>audserver -stopservice</code> (Windows Only)	Stop audit server logging.
<code>audserver -remove</code>	Removes the audit server logging service from the registry.

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

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

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

- ♦ On Windows: `\ns\etc`
- ♦ On Linux: `\usr\local\netscaler\etc`

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

Adding the NetScaler Appliance IP Addresses on the NSLOG Server

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

To add the IP addresses of the NetScaler appliance

At a command prompt, type the following command:

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

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

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

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

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

Password: Specifies the password, for example, nsroot.

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

Verifying the NSLOG Server Configuration File

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

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

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

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

Running the NSLOG Server

To start audit server logging

Type the following command at a command prompt:

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

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

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

Type the following command:

```
audserver -stop
```

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

Type the following command:

```
audserver -stopservice
```

Customizing Logging on the NSLOG Server

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

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

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

Creating Filters

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

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

To create a filter

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

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

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

<ip>: Specify the IP addresses.

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

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

Examples

```
filter F1 IP 192.168.100.151 ON
```

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

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

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

Specifying Log Properties

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

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

Entries in the definition can include the following:

- ♦ **LogFilenameFormat** specifies the file name format of the log file. The name of the file can be of the following types:
 - Static: A constant string that specifies the absolute path and the file name.
 - Dynamic: An expression that includes the following format specifiers:
 - ♦ Date ({format}t)
 - ♦ % creates file name with NSIP

Example

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

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

Example

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



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

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

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

Example:

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

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

- ♦ **LogInterval** specifies the interval at which new log files are created. Use one of the following values:

- Hourly: A file is created every hour. Default value.
- Daily: A file is created every day at midnight.
- Weekly: A file is created every Sunday at midnight.
- Monthly : A file is created on the first day of the month at midnight.
- None: A file is created only once, when audit server logging starts.
- Size: A file is created only when the log file size limit is reached.

Example

```
LogInterval Hourly
```

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

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

The default LogFileSizeLimit is 10 MB.

Example

```
LogFileSizeLimit 35
```

Default Settings for the Log Properties

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

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

Following are two examples of defining the default filters:

Example 1

```
Filter f1 IP 192.168.10.1
```

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

Example 2

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

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

Sample Configuration File (audit.conf)

Following is a sample configuration file:

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

Web Server Logging

You can use the Web server logging feature to send logs of HTTP and HTTPS requests to a client system for storage and retrieval. This feature has two components: the Web log server, which runs on the Citrix NetScaler appliance, and the NetScaler Web Logging (NSWL) client, which runs on the client system. When you run the client, it connects to the NetScaler. The NetScaler buffers the HTTP and HTTPS request log entries before sending them to the NSWL client, and the client can filter the entries before storing them. You can log HTTP and HTTPS requests for all of your Web servers on one NSWL client system.

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

Configuring the NetScaler Appliance for Web Server Logging

On the NetScaler appliance you need to enable the Web Server Logging feature, and you can modify the size of the buffer that stores the logged information before sending the logged information to the NetScaler Web Logging (NSWL) client.

Enabling or Disabling Web Server Logging

Web server logging is enabled by default.

To enable or disable Web server logging by using the command line interface

At the command prompt, type the following relevant commands to add or remove Web server logging and verify the configuration:

- ♦ `enable ns feature WL`
- ♦ `disable ns feature WL`
- ♦ `show ns feature`

To enable or disable Web server logging by using the configuration utility

1. Navigate to **System > Settings**.
2. In the details pane, under **Modes and Features**, click **Change advanced features**.
3. In the **Configure Advanced Features** dialog box, select the **Web Logging** check box to enable the Web logging feature, or clear the check box to disable the feature.
4. Click **OK**.

Modifying the Default Buffer Size

You can change the default buffer size of 16 megabytes (MB) for Web server logging to suit your requirements. To activate your modification, you must disable and reenble Web server logging.

To modify the buffer size by using the command line interface

At the command prompt, type the following commands to modify the buffer size and verify the configuration:

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

Example

```
> set weblogparam -bufferSizeMB 32
```

To modify the buffer size by using the configuration utility

1. Navigate to **System > Settings**.
2. In the details pane, under **Settings**, click **Change global system settings**.
3. In the **Configure Global Settings** dialog box, under **Web Logging**, enter a value in the **Buffer Size (in MBytes)** text box (for example, 32).
4. Click **OK**.

Exporting Custom HTTP Headers

The NetScaler can export values of custom HTTP headers to the NSWL client. You can configure up to a maximum of two HTTP request header names and two HTTP response header names.

To export custom HTTP headers by using the command line interface

At the command prompt, type the following commands to export the custom HTTP headers and verify the configuration:

- ♦ **set ns weblogparam** [-customReqHdrs <string> ...] [-customRspHdrs <string> ...]
- ♦ **show ns weblogparam**

Example

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

To export the custom HTTP headers by using the configuration utility

1. Navigate to **System > Settings**.
2. In the details pane, under **Settings**, click **Change global system settings**.
3. In the **Configure Global Settings** dialog box, under **Web Logging**, in the **Custom HTTP Request Header** and **Custom HTTP Response Header** text boxes, enter the HTTP request header name and HTTP response header name.
4. Click **OK**.

Installing and Configuring the Client System for Web Server Logging

During installation, the NSWL client executable file (nswl) is installed along with other files. The nswl executable file includes options for performing several actions on the NSWL client, including running and stopping the NSWL client. In addition, you use the nswl executable to configure the NSWL client with the IP addresses of the NetScaler

appliances from which the NSWL client will start collecting logs. Configuration settings are applied in the NSWL client configuration file (log.conf).

Then, you start the NSWL client by executing the nswl executable. The NSWL client configuration is based on the settings in the configuration file. You can further customize logging on the NSWL client system by making additional modifications to the NSWL configuration file (log.conf).

The following table lists the operating systems on which the NSWL client is supported.

Table 2-10. Supported Platforms for the NSWL Client

Operating system	Version
Windows	<ul style="list-style-type: none"> ♦ Windows XP Professional ♦ Windows Server 2003 ♦ Windows 2000/NT ♦ Windows Server 2008 ♦ Windows Server 2008 R2
Mac OS	Mac OS 8.6 or later
Linux	<ul style="list-style-type: none"> ♦ RedHat Linux 4 or later ♦ SUSE Linux Enterprise 9.3 or later
Solaris	Solaris Sun OS 5.6 or later
FreeBSD	FreeBSD 6.3 or later
AIX	AIX 6.1

The following table describes the minimum hardware specifications for the platform running the NSWL client.

Table 2-11. Minimum Hardware Specification for Platforms Running the NSWL Client

Operating system	Hardware requirements
For Windows / Linux / FreeBSD	<ul style="list-style-type: none"> • Processor- Intel x86 ~501 megahertz (MHz) • RAM - 512 megabytes (MB) • Controller - SCSI

Operating system	Hardware requirements
For Solaris 2.6	<ul style="list-style-type: none"> • Processor - UltraSPARC-IIi 400 MHz • RAM - 512 MB • Controller - SCSI

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



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

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

Installing NSWL Client on a Solaris Operating System

Before installing the NSWL client, you have to copy the NSWL client package from the NetScaler product CD or download it from www.citrix.com. The NSWL client package has the following name format:

Weblog_<release number>-<build number>.zip (for example, Weblog_9.3-51.5.zip). Within the package are separate installation packages for each supported platforms.

To download NSWL client package from www.Citrix.com

1. From any system, open www.citrix.com in the Web browser.
2. In the menu bar, click **Log In**.
3. Enter your login credentials and then click **Log In**.
4. In the menu bar, click **Downloads**.
5. Search to the page of the desired release number and build.
6. On the desired page, under **Weblog Clients**, click **Download** to download a file, having the format Weblog_<release number>-<build number>.zip, to your local system (for example, Weblog_9.3-51.5.zip).

To install the NSWL client package on a Solaris operating system

1. On the system, where you have downloaded the NSWL client package Weblog_<release number>-<build number>.zip (for example, Weblog_9.3-51.5.zip), extract nswl_solaris-<release number>-<build number>.tar (for example, nswl_solaris-9.3-51.5.tar) from the package.
2. Copy the extracted file nswl_solaris-<release number>-<build number>.tar (for example, nswl_solaris-9.3-51.5.tar) to a Solaris system on which you want to install the NSWL client.

3. Extract the files from the `nswl_solaris-<release number>-<build number>.tar` (for example, `nswl_solaris-9.3-51.5.tar` file with the following command:

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

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

4. Install the package with the following command:

```
pkgadd -d
```

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

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

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

```
pkginfo | grep NSweblog
```

To uninstall the NSWL client package on a Solaris operating system

At a command prompt, type:

```
pkgrm NSweblog
```

Installing NSWL Client on a Linux Operating System

Before installing the NSWL client, you have to copy the NSWL client package from the NetScaler product CD or download it from www.citrix.com. The NSWL client package has the following name format:

`Weblog_<release number>-<build number>.zip` (for example, `Weblog_9.3-51.5.zip`). Within the package are separate installation packages for each supported platforms.

To download NSWL client package from www.Citrix.com

1. From any system, open www.citrix.com in the Web browser.
2. In the menu bar, click **Log In**.
3. Enter your login credentials and then click **Log In**.

4. In the menu bar, click **Downloads** .
5. Search to the page of the desired release number and build.
6. On the desired page, under **Weblog Clients**, click **Download** to download a file, having the format `Weblog_<release number>-<build number>.zip`, to your local system (for example, `Weblog_9.3-51.5.zip`).

To install the NSWL client package on a Linux operating system

1. On the system, where you have downloaded the NSWL client package `Weblog_<release number>-<build number>.zip` (for example, `Weblog_9.3-51.5.zip`), **extract** `nswl_linux-<release number>-<build number>.rpm` (for example, `nswl_linux-9.3-51.5.rpm`) from the package.
2. **Copy** the extracted file `nswl_linux-<release number>-<build number>.rpm` (for example, `nswl_linux-9.3-51.5.rpm`) **to a system**, running Linux OS, on which you want to install the NSWL client.
3. To install the NSWL executable, use the following command:

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

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

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

To uninstall the NSWL client package on a Linux operating system

At a command prompt, type:

```
rpm -e NSweblog
```

To get more information about the NSWeblog RPM file

At a command prompt, type:

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

To view the installed Web server logging files

At a command prompt, type:

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

Installing NSWL Client on a FreeBSD Operating System

Before installing the NSWL client, you have to copy the NSWL client package from the NetScaler product CD or download it from www.citrix.com. The NSWL client package has the following name format:

`Weblog_<release number>-<build number>.zip` (for example, `Weblog_9.3-51.5.zip`). Within the package are separate installation packages for each supported platforms.

To download NSWL client package from www.Citrix.com

1. From any system, open www.citrix.com in the Web browser.
2. In the menu bar, click **Log In**.
3. Enter your login credentials and then click **Log In**.
4. In the menu bar, click **Downloads**.
5. Search to the page of the desired release number and build.
6. On the desired page, under **Weblog Clients**, click **Download** to download a file, having the format `Weblog_<release number>-<build number>.zip`, to your local system (for example, `Weblog_9.3-51.5.zip`).

To install the NSWL client package on a FreeBSD operating system

1. On the system, where you have downloaded the NSWL client package `Weblog_<release number>-<build number>.zip` (for example, `Weblog_9.3-51.5.zip`), extract `nswl_bsd-<release number>-<build number>.tgz` (for example, `nswl_bsd-9.3-51.5.tgz`) from the package.
2. Copy the extracted file `nswl_bsd-<release number>-<build number>.tgz` (for example, `nswl_bsd-9.3-51.5.tgz`) to a system, running FreeBSD OS, on which you want to install the NSWL client.
3. Install the package using the following command:

```
pkg_add nswl_bsd-9.3-51.5.tgz
```

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

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

4. To verify that the package is installed, use the following command:

```
pkg_info | grep NSweblog
```

To uninstall the NSWL client package on a FreeBSD operating system

At a command prompt, type:

```
pkg_delete NSweblog
```

Installing NSWL Client on a Mac OS Operating System

Before installing the NSWL client, you have to copy the NSWL client package from the NetScaler product CD or download it from www.citrix.com. The NSWL client package has the following name format:

`Weblog_<release number>-<build number>.zip` (for example, `Weblog_9.3-51.5.zip`).
Within the package are separate installation packages for each supported platforms.

To download NSWL client package from www.Citrix.com

1. From any system, open www.citrix.com in the Web browser.
2. In the menu bar, click **Log In**.
3. Enter your login credentials and then click **Log In**.
4. In the menu bar, click **Downloads**.
5. Search to the page of the desired release number and build.
6. On the desired page, under **Weblog Clients**, click **Download** to download a file, having the format `Weblog_<release number>-<build number>.zip`, to your local system (for example, `Weblog_9.3-51.5.zip`).

To install the NSWL client package on a Mac OS operating system

1. On the system, where you have downloaded the NSWL client package `Weblog_<release number>-<build number>.zip` (for example, `Weblog_9.3-51.5.zip`), extract `nswl_macos-<release number>-<build number>.tgz` (for example, `nswl_macos-9.3-51.5.tgz`) from the package.
2. Copy the extracted file `nswl_macos-<release number>-<build number>.tgz` (for example, `nswl_macos-9.3-51.5.tgz`) to a system, running Mac OS, on which you want to install the NSWL client.
3. To install the package, use the `pkg_add` command:

```
pkg_add nswl_macos-9.3-51.5.tgz
```

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

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

4. To verify that the package is installed, use the following command:

```
pkg_info | grep NSweblog
```

To uninstall the NSWL client package on a Mac OS operating system

At a command prompt, type:

```
pkg_delete NSweblog
```

Installing NSWL Client on a Windows Operating System

Before installing the NSWL client, you have to copy the NSWL client package from the NetScaler product CD or download it from www.citrix.com. The NSWL client package has the following name format:

`Weblog_<release number>-<build number>.zip` (for example, `Weblog_9.3-51.5.zip`). Within the package are separate installation packages for each supported platforms.

To download NSWL client package from www.Citrix.com

1. From any system, open www.citrix.com in the Web browser.
2. In the menu bar, click **Log In**.
3. Enter your login credentials and then click **Log In**.
4. In the menu bar, click **Downloads** .
5. Search to the page of the desired release number and build.
6. On the desired page, under **Weblog Clients**, click **Download** to download a file, having the format `Weblog_<release number>-<build number>.zip`, to your local system (for example, `Weblog_9.3-51.5.zip`).

To install the NSWL client on a Windows system

1. On the system, where you have downloaded the NSWL client package `Weblog_<release number>-<build number>.zip` (for example, `Weblog_9.3-51.5.zip`), extract `nswl_win-<release number>-<build number>.zip` (for example, `nswl_win-9.3-51.5.zip`) from the package.
2. Copy the extracted file `nswl_win-<release number>-<build number>.zip` (for example, `nswl_win-9.3-51.5.zip`) to a Windows system on which you want to install the NSWL client.
3. On the Windows system, unzip the `nswl_<release number>-<build number>.zip` file (for example , `nswl_win-9.3-51.5.zip`). The following directories are extracted:
 - a. <root directory extracted from the Windows NSWL client package zip file>\bin (for example, `C:\nswl_win-9.3-51.5\bin`)
 - b. <root directory extracted from the Windows NSWL client package zip file>\etc (for example, `C:\nswl_win-9.3-51.5\ etc`)
 - c. < root directory extracted from the Windows NSWL client package zip file >\samples (for example, `C:\nswl_win-9.3-51.5\samples`)
4. At a command prompt, run the following command from the <root directory extracted from the Windows NSWL client package zip file>\bin path:


```
nswl -install -f <directorypath> \log.conf
```

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

To uninstall the NSWL client on a Windows system

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

```
nswl -remove
```

Installing NSWL Client on an AIX Operating System

Before installing the NSWL client, you have to copy the NSWL client package from the NetScaler product CD or download it from www.citrix.com. The NSWL client package has the following name format:

Weblog_<release number>-<build number>.zip (for example, Weblog_9.3-51.5.zip). Within the package are separate installation packages for each supported platforms.

To download NSWL client package from www.Citrix.com

1. From any system, open www.citrix.com in the Web browser.
2. In the menu bar, click **Log In**.
3. Enter your login credentials and then click **Log In**.
4. In the menu bar, click **Downloads**.
5. Search to the page of the desired release number and build.
6. On the desired page, under **Weblog Clients**, click **Download** to download a file, having the format Weblog_<release number>-<build number>.zip, to your local system (for example, Weblog_9.3-51.5.zip).

To install the NSWL client package on an AIX operating system

1. On the system, where you have downloaded the NSWL client package Weblog_<release number>-<build number>.zip (for example, Weblog_9.3-51.5.zip), **extract** nswl_aix-<release number>-<build number>.rpm (for example, nswl_aix-9.3-51.5.rpm) from the package.
2. Copy the extracted file nswl_aix-<release number>-<build number>.rpm (for example, nswl_aix-9.3-51.5.rpm) to a system, running AIX OS, on which you want to install the NSWL client.
3. To install the NSWL executable, use the following command:

```
rpm -i nswl_aix-9.3-51.5.rpm
```

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

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

To uninstall the NSWL client package on an AIX operating system

At a command prompt, type:

```
rpm -e NSweblog
```

To get more information about the NSweblog RPM file

At a command prompt, type:

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

To view the installed Web server logging files

At a command prompt, type:

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

NSWL Client Command Options

The following table describes the commands that you can use to configure the NSWL client.

Table 2-12. NSWL Command Options

NSWL command	Specifies
nswl -help	The available NSWL help options.
nswl -addns -f <path to configuration file>	The system that gathers the log transaction data. You are prompted to enter the IP address of the NetScaler appliance. Enter a valid user name and password.
nswl -verify -f <path to configuration file>	Check for syntax or semantic errors in the configuration file (for example, log.conf).
nswl -start -f <path to configuration file>	Start the NSWL client based on the settings in the configuration file (for example, log.conf). For Solaris and Linux: To start Web server logging as a background process, type the ampersand sign (&) at the end of the command.
nswl -stop (Solaris and Linux only)	Stop the NSWL client if it was started as a background process; otherwise, use CTRL+C to stop Web server logging.
nswl -install -f <path to configuration file> (Windows only)	Install the NSWL client as a service in Windows.
nswl -startservice (Windows only)	Start the NSWL client by using the settings in the configuration file (for example, log.conf) specified in the nswl install option. You can also start NSWL

NSWL command	Specifies
	client from Start > Control Panel > Services .
nswl -stopservice (Windows only)	Stops the NSWL client.
nswl -remove	Remove the NSWL client service from the registry.

Run the following commands from the directory in which the NSWL executable is located:

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

The Web server logging configuration files are located in the following directory path:

- ♦ Windows: \ns\etc
- ♦ Solaris and Linux: \usr\local\netscaler\etc

The NSWL executable is started as .nswl in Linux and Solaris.

Adding the IP Addresses of the NetScaler Appliance

In the NSWL client configuration file (log.conf), add the NetScaler IP address (NSIP) from which the NSWL client will start collecting logs.

To add the NSIP address of the NetScaler appliance

1. At the client system command prompt, type:

```
nswl -addns -f < directorypath > \log.conf
```

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

2. At the next prompt, enter the following information:

- **NSIP:** Specify the IP address of the NetScaler appliance.
- **Username and Password:** Specify the `nsroot` user credentials of the NetScaler appliance.

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

Verifying the NSWL Configuration File

To make sure that logging works correctly, check the NSWL configuration file (log.conf) on the client system for syntax errors.

To verify the configuration in the NSWL configuration file

At the client system command prompt, type:

```
nswl -verify -f <directorypath>\log.conf
```

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

Running the NSWL Client

To start Web server logging

At the client system command prompt, type:

```
nswl -start -f <directorypath>\log.conf
```

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

To stop Web server logging started as a background process on the Solaris or Linux operating systems

At the command prompt, type:

```
nswl -stop
```

To stop Web server logging started as a service on the Windows operating system

At the command prompt, type:

```
nswl -stopservice
```

Customizing Logging on the NSWL Client System

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

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

- ♦ **Log filters.** Filter log information based on the host IP address, domain name, and host name of the Web servers.
- ♦ **Log properties.** Each filter has an associated set of log properties. Log properties define how to store the filtered log information.

Creating Filters

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

Note: Consolidated logging, which logs transactions for which no filter is defined, uses the default filter if it is enabled. Consolidated logging of all servers can be done by defining only the default filter.

If the server hosts multiple Web sites and each Web site has its own domain name, and each domain is associated with a virtual server, you can configure Web server logging to create a separate log directory for each Web site. The following table displays the parameters for creating a filter.

Table 2-13. Parameters for Creating a Filter

Parameter	Specifies
filterName	Name of the filter (maximum 64 alphanumeric characters).
HOST name	Host name of the server for which the transactions are being logged.
IP ip	IP address of the server for which transactions are to be logged (for example, if the server has multiple domains that have one IP address).
IP ip 2...ip n:	Multiple IP addresses (for example, if the server domain has multiple IP addresses).
ip6 ip	IPv6 address of the server for which transactions are to be logged.
IP ip NETMASK mask	IP addresses and netmask combination to be used on a subnet.
ON OFF	Enable or disable the filter to log transactions. If no argument is selected, the filter is enabled (ON).

To create a filter

To create a filter, enter the following command in the log.conf file:

- ♦ `filter <filterName> <HOST name> | [IP<ip>] | [IP<ip 2...ip n>] | <IP ip NETMASK mask> [ON | OFF]`

- ♦ `filter <filterName> <HOST name> | [IP6 ip/<prefix length>] [ON | OFF]`

To create a filter for a virtual server

To create a filter for a virtual server, enter the following command in the log.conf file:

```
filter <filterName> <VirtualServer IP address>
```

Example

In the following example, you specify an IP address of 192.168.100.0 and netmask of 255.255.255.0. The filter applies to IP addresses 192.168.100.1 through 192.168.100.254.

```
Filter F1 HOST www.netscaler.com ON
Filter F2 HOST www.netscaler.com IP
192.168.100.151 ON
Filter F3 HOST www.netscaler.com IP
192.168.100.151 192.165.100.152 ON
Filter F4 IP 192.168.100.151
Filter F5 IP 192.168.100.151 HOST
www.netscaler.com OFF
Filter F6 HOST www.netscaler.com HOST www.xyz.com
HOST www.abcxyz.com IP 192.168.100.200 ON
Filter F7 IP 192.250.100.0 NETMASK 255.255.255.0
Filter F8 HOST www.xyz.com IP 192.250.100.0
NETMASK 255.255.255.0 OFF
For creating filters for servers having IPv6
addresses.
Filter F9 2002::8/112 ON
Filter F10 HOST www.abcd.com IP6 2002::8 ON
```

Specifying Log Properties

Log properties are applied to all log entries associated with the filter. The log property definition begins with the keyword **BEGIN** and ends with **END** as illustrated in the following example:

```
BEGIN <filtername>
logFormat ...
logFilenameFormat ...
logInterval ...
logFileSize ....
logExclude ....
logTime ...
END
```

Entries in the definition can include the following:

- ♦ **LogFormat** specifies the Web server logging feature that supports NCSA, W3C Extended, and custom log file formats.

By default, the `logformat` property is `w3c`. To override, enter `custom` or `NCSA` in the configuration file, for example:

```
LogFormat NCSA
```

Note: For the `NCSA` and custom log formats, local time is used to time stamp transactions and for file rotation.

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

Example

```
LogInterval Daily
```

- ♦ **LogFileSizeLimit** specifies the maximum size of the log file in MB. It can be used with any log interval (weekly, monthly, and so on.) A file is created when the maximum file size limit is reached or when the defined log interval time elapses.

To override this behavior, specify the size as the `loginterval` property so that a file is created only when the log file size limit is reached.

The default `LogFileSizeLimit` is 10 MB.

Example

```
LogFileSizeLimit 35
```

- ♦ **LogFilenameFormat** specifies the file name format of the log file. The name of the file can be of the following types:
 - **Static:** Specifies a constant string that contains the absolute path and file name.
 - **Dynamic:** Specifies an expression containing the following format:
 - ♦ Server IP address (%A)
 - ♦ Date (%{format}t)
 - ♦ URL suffix (%x)
 - ♦ Host name (%v)


Example

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

This command creates the first file name as Exmmddyy.log, then every hour creates a file with file name: Exmmddyy.log.0, Exmmddyy.log.1,..., Exmmddyy.log.n.

Example

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

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

- ♦ **LogExclude** prevents logging of transactions with the specified file extensions.

Example

```
LogExclude .html
```

This command creates a log file that excludes log transactions for *.html files.

- ♦ **LogTime** specifies log time as either GMT or LOCAL.

The defaults are:

- NCSA log file format: LOCAL
- W3C log file format: GMT.

Understanding the NCSA and W3C Log Formats

The NetScaler supports the following standard log file formats:

- ♦ NCSA Common Log Format
- ♦ W3C Extended Log Format

NCSA Common Log Format

If the log file format is NCSA, the log file displays log information in the following format:

```
Client_IP_address -User_Name [Date:Time -TimeZone] "Method
Object HTTP_version"
HTTP_StatusCode BytesSent
```

To use the NCSA Common log format, enter NCSA in the LogFormat argument in the log.conf file.

The following table describes the NCSA Common log format.

Table 2-14. NCSA Common Log Format

Argument	Specifies
Client_IP_address	The IP address of the client computer.
User Name	The user name.
Date	The date of the transaction.
Time	The time when the transaction was completed.
Time Zone	The time zone (Greenwich Mean Time or local time).
Method	The request method (for example; GET, POST).
Object	The URL.
HTTP_version	The version of HTTP used by the client.
HTTP_StatusCode	The status code in the response.
Bytes Sent	The number of bytes sent from the server.

W3C Extended Log Format

An extended log file contains a sequence of lines containing ASCII characters terminated by either a Line Feed (LF) or the sequence Carriage Return Line Feed (CRLF.) Log file generators must follow the line termination convention for the platform on which they are run.

Log analyzers must accept either LF or CRLF form. Each line may contain either a directive or an entry. If you want to use the W3C Extended log format, enter W3C as the Log-Format argument in the log.conf file.

By default, the standard W3C log format is defined internally as the custom log format, shown as follows:

```
%{%Y-%m-%d%H:%M:%S}t %a %u %S %A %p %m %U %q %s %j %J %T %H %+
{user-agent}i %+{cookie} i%+{referer}i
```

```
logFormat W3C {%Y-%m-%d%H:%M:%S}t %m %U
```

W3C log entries are created with the following format:

```
#Version: 1.0
#Fields: date time cs-method cs-uri
#Date: 12-Jun-2001 12:34
```

```
2001-06-12 12:34:23 GET /sports/football.html
2001-06-12 12:34:30 GET /sports/football.html
```

Entries

Entries consist of a sequence of fields relating to a single HTTP transaction. Fields are separated by white space; Citrix recommends the use of tab characters. If a field in a particular entry is not used, a dash (-) marks the omitted field.

Directives

Directives record information about the logging process. Lines beginning with the pound sign (#) contain directives.

The following table describes the directives.

Table 2-15. Directive Descriptions

Directive	Description
Version: <integer>.<integer>	Displays the version of the extended log file format used. This document defines version 1.0.
Fields: [<specifier>...]	Identifies the fields recorded in the log.
Software: <string>	Identifies the software that generated the log.
Start-Date: <date> <time>	Displays the date and time at which the log was started.
End-Date: <date> <time>	Displays the date and time at which logging finished.
Date: <date> <time>	Displays the date and time when the entry was added.
Remark: <text>	Displays comments. Analysis tools ignore data recorded in this field.

Note: The Version and Fields directives are required. They precede all other entries in the log file.

Example

The following sample log file shows the log entries in W3C Extended log format:

```
#Version: 1.0
#Fields: time cs-method cs-uri
#Date: 12-Jan-1996 00:00:00
00:34:23 GET /sports/football.html
12:21:16 GET /sports/football.html
```

```
12:45:52 GET /sports/football.html
12:57:34 GET /sports/football.html
```

Fields

The Fields directive lists a sequence of field identifiers that specify the information recorded in each entry. Field identifiers may have one of the following forms:

- ♦ **identifier:** Relates to the transaction as a whole.
- ♦ **prefix-identifier:** Relates to information transfer between parties defined by the value *prefix*.
- ♦ **prefix (header):** Specifies the value of the HTTP header field header for transfer between parties defined by the value *prefix*. Fields specified in this manner always have the type <string>.

The following table describes defined prefixes.

Table 2-16. Prefix Descriptions

Prefix	Specifies
c	Client
s	Server
r	Remote
cs	Client to server
sc	Server to client
sr	Server to remote server (prefix used by proxies)
rs	Remote server to server (prefix used by proxies)
x	Application-specific identifier

Examples

The following examples are defined identifiers that use prefixes:

cs-method: The method in the request sent by the client to the server.

sc(Referer): The Referer field in the reply.

c-ip: The IP address of the client.

Identifiers

The following table describes the W3C Extended log format identifiers that do not require a prefix.

Table 2-17. W3C Extended Log Format Identifiers (No Prefix Required)

Identifier	Description
date	The date on which the transaction was done.
time	The time when the transaction is done.
time-taken	The time taken (in seconds) for the transaction to complete.
bytes	The number of bytes transferred.
cached	Records whether a cache hit has occurred. A zero indicates a cache miss.

The following table describes the W3C Extended log format identifiers that require a prefix.

Table 2-18. W3C Extended Log Format Identifiers (Requires a Prefix)

Identifier	Description
IP	The IP address and the port number.
dns	The DNS name.
status	The status code.
comment	The comment returned with status code.
method	The method.
url	The URL.
url-stem	The stem portion of the URL.
url-query	The query portion of the URL.

The W3C Extended Log file format allows you to choose log fields. These fields are shown in the following table.

Table 2-19. W3C Extended Log File Format (Allows Log Fields)

Field	Description
Date	The date on which the transaction is done.
Time	The time when the transaction is done.

Field	Description
Client IP	The IP address of the client.
User Name	The user name.
Service Name	The service name, which is always HTTP.
Server IP	The server IP address.
Server Port	The server port number
Method	The request method (for example; GET, POST).
Url Stem	The URL stem.
Url Query	The query portion of the URL.
Http Status	The status code in the response.
Bytes Sent	The number of bytes sent to the server (request size, including HTTP headers).
Bytes Received	The number of bytes received from the server (response size, including HTTP headers).
Time Taken	The time taken for transaction to complete, in seconds.
Protocol Version	The version number of HTTP being used by the client.
User Agent	The User-Agent field in the HTTP protocol.
Cookie	The Cookie field of the HTTP protocol.
Referer	The Referer field of the HTTP protocol.

Creating a Custom Log Format

You can customize the display format of the log file data manually or by using the NSWL library. By using the custom log format, you can derive most of the log formats that Apache currently supports.

Creating a Custom Log Format by Using the NSWL Library

Use one of the following NSWL libraries depending on whether the NSWL executable has been installed on a Windows or Solaris host computer:

- ♦ **Windows:** The nswl.lib library located in \ns\bin directory on the system manager host computer.
- ♦ **Solaris:** The libnswl.a library located in /usr/local/netscaler/bin.

To create the custom log format by using the NSWL Library

1. Add the following two C functions defined by the system in a C source file:
 ns_userDefFieldName() : This function returns the string that must be added as a custom field name in the log record.

 ns_userDefFieldVal() : This function implements the custom field value, then returns it as a string that must be added at the end of the log record.
2. Compile the file into an object file.
3. Link the object file with the NSWL library (and optionally, with third party libraries) to form a new NSWL executable.
4. Add a %d string at the end of the logFormat string in the configuration file (log.conf).

Example

```
#####
# A new file is created every midnight or on
# reaching 20MB file size,
# and the file name is /datadisk5/netscaler/log/
NS<hostname>/Nsmdddy.log and create digital
#signature field for each record.
BEGIN CACHE_F
    logFormat      custom "%a - "%{user-agent}i"
[%d/%B/%Y %T -%g] "%x" %s %b%{referrer}i "%{user-
agent}i" "%{cookie}i" %d "
    logInterval      Daily
    logFileSizeLimit      20
    logFilenameFormat      /datadisk5/
netscaler/log/%v/NS%{m%dy}t.log
END CACHE_F
```

Creating a Custom Log Format Manually

To customize the format in which log file data should appear, specify a character string as the argument of the LogFormat log property definition. The following is an example where character strings are used to create a log format:

```
LogFormat Custom ""%a - "%{user-agent}i" "[%d/%m/%Y]t %U %s
%b %T"
```

- ♦ The string can contain the “c” type control characters \n and \t to represent new lines and tabs.

- ♦ Use the <Esc> key with literal quotes and backslashes.

The characteristics of the request are logged by placing % directives in the format string, which are replaced in the log file by the values.

If the %v (Host name) or %x (URL suffix) format specifier is present in a log file name format string, the following characters in the file name are replaced by an underscore symbol in the log configuration file name:

```
" * . / : < > ? \ |
```

Characters whose ASCII values lie in the range of 0-31 are replaced by the following:

%<ASCII value of character in hexadecimal>.

For example, the character with ASCII value 22 is replaced by %16.



Caution: If the %v format specifier is present in a log file name format string, a separate file is opened for each virtual host. To ensure continuous logging, the maximum number of files that a process can have open should be sufficiently large. See your operating system documentation for a procedure to change the number of files that can be opened.

Creating Apache Log Formats

You can derive from the custom logs most of the log formats that Apache currently supports. The custom log formats that match Apache log formats are:

NCSA/combined: `LogFormat custom %h %l %u [%t] "%r" %s %B "%{referer}i" "%{user-agent}i"`

NCSA/Common: `LogFormat custom %h %l %u [%t] "%r" %s %B`

Referer Log: `LogFormat custom "%{referer}i" -> %U`

Useragent: `LogFormat custom %{user-agent}i`

Similarly, you can derive the other server log formats from the custom formats.

Sample Configuration File

Following is a sample configuration file:

```
#####
# This is the NSWL configuration file
# Only the default filter is active
# Remove leading # to activate other filters
#####
#####
# Default filter (default on)
# W3C Format logging, new file is created every hour or on
reaching 10MB file size,
# and the file name is Exyyymmdd.log
```

```

#####
Filter default
begin default
    logFormat                W3C
    logInterval              Hourly
    logFileSizeLimit         10
    logFilenameFormat        Ex{%y%m%d}t.log
end default
#####
# netscaler caches example
# CACHE_F filter covers all the transaction with HOST name
# www.netscaler.com and the listed server ip's
#####
#Filter CACHE_F HOST www.netscaler.com IP 192.168.100.89
192.168.100.95 192.168.100.52 192.168.100.53 ON
#####
# netscaler origin server example
# Not interested in Origin server to Cache traffic transaction
logging
#####
#Filter ORIGIN_SERVERS IP 192.168.100.64 192.168.100.65
192.168.100.66 192.168.100.67 192.168.100.225 192.168.100.226
192.168.
100.227 192.168.100.228 OFF
#####
# netscaler image server example
# all the image server logging.
#####
#Filter IMAGE_SERVER HOST www.netscaler.images.com IP
192.168.100.71 192.168.100.72 192.168.100.169 192.168.100.170
192.168.10
0.171 ON
#####
# NCSA Format logging, new file is created every day midnight
# or on reaching 20MB file size,
# and the file name is /datadisk5/netscaler/log/NS<hostname>/
Nsmddyy.log.
# Exclude objects that ends with .gif .jpg .jar.
#####
#begin ORIGIN_SERVERS
#    logFormat                NCSA
#    logInterval              Daily
#    logFileSizeLimit         40
#    logFilenameFormat        /datadisk5/ORIGIN/log/%v/NS{%m
%d%y}t.log
#    logExclude                .gif .jpg .jar
#end ORIGIN_SERVERS

#####
# NCSA Format logging, new file is created every day midnight
# or on reaching 20MB file size,
# and the file name is /datadisk5/netscaler/log/NS<hostname>/
Nsmddyy.log with log record timestamp as GMT.
#####
#begin CACHE_F
#    logFormat                NCSA
#    logInterval              Daily

```

```

#       logFileSizeLimit      20
#       logFilenameFormat /datadisk5/netcaler/log/%v/NS%{m%d
%y}t.log
#       logtime                GMT
#end CACHE_F

#####
# W3C Format logging, new file on reaching 20MB and the log
file path name is
# atadisk6/netcaler/log/server's ip/Exmmyydd.log with log
record timestamp as LOCAL.
#####
#begin IMAGE_SERVER
#       logFormat              W3C
#       logInterval            Size
#       logFileSizeLimit      20
#       logFilenameFormat /datadisk6/netcaler/log/%AEx%{m%d
%y}t
#       logtime                LOCAL
#end IMAGE_SERVER

#####
# Virtual Host by Name firm, can filter out the logging based
on the host name by,
#####

#Filter VHOST_F IP 10.101.2.151 NETMASK 255.255.255.0
#begin VHOST_F
#       logFormat              W3C
#       logInterval            Daily
#       logFileSizeLimit      10
logFilenameFormat /ns/prod/vhost/%v/Ex%{m%d%y}t
#end VHOST_F

##### END FILTER CONFIGURATION #####

```

Arguments for Defining a Custom Log Format

The following table describes the data that you can use as the Log Format argument string:

Table 2-20. Custom Log Format

Argument	Specifies
%a	Remote IPv4 address.
%A	Local IPv4 address.
%a6	Remote IPv6 address.

%A6	Local IPv6 address.
%B	Bytes sent, excluding the HTTP headers (response size).
%b	Bytes received, excluding the HTTP headers (request size).
%d	User-defined field.
%e1	Value of the first custom HTTP request header.
%e2	Value of the second custom HTTP request header.
%E1	Value of the first custom HTTP response header.
%E2	Value of the second custom HTTP response header.
Note: For instructions on how to export custom HTTP headers, see " Exporting Custom HTTP Headers ."	
%g	Greenwich Mean Time offset (for example, -0800 for Pacific Standard Time).
%h	Remote host.
%H	Request protocol.

<code>%{Foobar}i</code>	Contents of the Foobar: header line(s) in the request sent to the server. The system supports the User-Agent, Referer and cookie headers. The + after the % in this format informs the logging client to use the + as a word separator.
<code>%j</code>	Bytes received, including headers (request size)
<code>%J</code>	Bytes sent, including headers (response size)
<code>%l</code>	Remote log name (from <code>identd</code> , if supplied).
<code>%m</code>	Request method.
<code>%M</code>	Time taken to serve the request (in microseconds)
<code>%{Foobar}o</code>	Contents of Foobar: header line(s) in the reply. USER-AGENT, Referer, and cookie headers (including set cookie headers) are supported.
<code>%p</code>	Canonical port of the server serving the request.
<code>%q</code>	Query string (prefixed with a question mark (?) if a query string exists).
<code>%r</code>	First line of the request.
<code>%s</code>	Requests that were redirected internally, this is the status of the original request.

%t	Time, in common log format (standard English time format).
%{format}t	Time, in the form given by format, must be in the strftime(3) format.
%T	Time taken to serve the request, in seconds.
%u	Remote user (from auth; may be bogus if return status (%s) is 401).
%U	URL path requested.
%v	Canonical name of the server serving the request.
%V	Virtual server IPv4 address in the system, if load balancing, content switching, and/or cache redirection is used.
%V6	Virtual server IPv6 address in the system, if load balancing, content switching, and/or cache redirection is used.

For example, if you define the log format as `%+{user-agent}i`, and if the user agent value is Citrix NetScaler system Web Client, then the information is logged as NetScaler system+Web+Client. An alternative is to use double quotation marks. For example, `"%{user-agent}i"` logs it as "Citrix NetScaler system Web Client." Do not use the <Esc> key on strings from `%. . .r`, `%. . .i` and, `%. . .o`. This complies with the requirements of the Common Log Format. Note that clients can insert control characters into the log. Therefore, you should take care when working with raw log files.

Time Format Definition

The following table lists the characters that you can enter as the format part of the `%{format}t` string described in the Custom Log Format table of ["Arguments for Defining a Custom Log Format."](#) Values within brackets ([]) show the range of values that appear.

For example, [1,31] in the %d description in the following table shows %d ranges from 1 to 31.

Table 2-21. Time Format Definition

Argument	Specifies
%%	The same as %.
%a	The abbreviated name of the week day for the locale.
%A	The full name of the week day for the locale.
%b	The abbreviated name of the month for the locale.
%B	The full name of the month for the locale.
%C	The century number (the year divided by 100 and truncated to an integer as a decimal number [1,99]); single digits are preceded by a 0.
%d	The day of month [1,31]; single digits are preceded by 0.
%e	The day of month [1,31]; single digits are preceded by a blank.
%h	The abbreviated name of the month for the locale.
%H	The hour (24-hour clock) [0,23]; single digits are preceded by a 0.
%I	The hour (12-hour clock) [1,12]; single digits are preceded by a 0.
%j	The number of the day in the year [1,366]; single digits are preceded by 0.
%k	The hour (24-hour clock) [0,23]; single digits are preceded by a blank.
%l	The hour (12-hour clock) [1,12]; single digits are preceded by a blank.
%m	The number of the month in the year [1,12]; single digits are preceded by a 0.

Argument	Specifies
%M	The minute [00,59]; leading 0 is permitted but not required.
%n	Inserts a new line.
%p	The equivalent of either a.m. or p.m. for the locale.
%r	The appropriate time representation in 12-hour clock format with %p.
%S	The seconds [00,61]; the range of values is [00,61] rather than [00,59] to allow for the occasional leap second and for the double leap second.
%t	Inserts a tab.
%u	The day of the week as a decimal number [1,7]. 1 represents Sunday, 2 represents Tuesday and so on.
%U	The number of the week in the year as a decimal number [00,53], with Sunday as the first day of week 1.
%w	The day of the week as a decimal number [0,6]. 0 represents Sunday.
%W	Specifies the number of the week in the year as a decimal number [00,53]. Monday is the first day of week 1.
%y	The number of the year within the century [00,99]. For example, 5 would be the fifth year of that century.
%Y	The year, including the century (for example, 1993).

Note: If you specify a conversion that does not correspond to any of the ones described in the preceding table, or to any of the modified conversion specifications listed in the next paragraph, the behavior is undefined and returns 0.

The difference between %U and %W (and also between modified conversions %OU and %OW) is the day considered to be the first day of the week. Week number 1 is the first week in January (starting with a Sunday for %U, or a Monday for %W). Week number 0 contains the days before the first Sunday or Monday in January for %U and %W.

Advanced Configurations

If you enable path maximum transmission unit (PMTU) discovery, the NetScaler can use it to determine the maximum transmission unit of any Internet channel. For more efficient data transfer, you can configure TCP window scaling and selective acknowledgment. You can view statistics associated with HTTP request and response sizes. For applying a specific HTTP and TCP settings to vservers and services, you can configure HTTP and TCP profiles.

Configuring TCP Window Scaling

The TCP window scaling option, which is defined in RFC 1323, increases the TCP receive window size beyond its maximum value of 65,535 bytes. This option is required for efficient transfer of data over long fat networks (LFNs).

A TCP window determines the amount of outstanding (unacknowledged by the recipient) data a sender can send on a particular connection before receiving any acknowledgment from the receiver. The main purpose of the window is flow control.

The window size field in the TCP header is 16 bits, which limits the ability of the sender to advertise a window size larger than 65535 ($2^{16} - 1$). The TCP window scale extension expands the definition of the TCP window by applying a scale factor to the value in the 16 bit window size field of the TCP header. (Although RFC 1323 describes expanding the definition to up to 30 bits, NetScaler window scaling expands the definition of the TCP window to up to 24 bits.) The scale factor is carried in the new TCP window scale field. This field is sent only in a SYN packet (a segment with the SYN bit on)

To fit a larger window size value into the 16-bit field, the sender right shifts the value by the number of bit positions specified by the scale factor. The receiver left shifts the value by the same number of positions. Therefore, the actual window size is equivalent to:

$$(2^{\text{scale factor}}) * \text{received window size}$$

Before configuring window scaling, make sure that:

- You do not set a high value for the scale factor, because this could have adverse effects on the appliance and the network.
- You have enabled selective acknowledgment (SACK).
- You do not configure window scaling unless you clearly know why you want to change the window size.
- Both hosts in the TCP connection send a window scale option during connection establishment. If only one side of a connection sets this option, window scaling is not used for the connection.
- Each connection for same session is an independent Window Scaling session. For example, when a client's request and the server's response flow through the appliance, it is possible to have window scaling between the client and the appliance without window scaling between the appliance and the server.

By default, window scaling is not enabled.

To configure window scaling by using the command line interface

At the command prompt, type the following commands to configure window scaling and verify the configuration:

- ♦ **set ns tcpParam -WS (ENABLED | DISABLED) -WSVal <positive_integer>**
- ♦ **show ns tcpParam**

Example

```
> set ns tcpParam -WS ENABLED -WSVal 6
```

To configure window scaling by using the configuration utility

1. Navigate to **System > Settings**.
2. In the details pane, under **Settings**, click **Configure TCP Parameters**.
3. In the **Configure TCP Parameters** dialog box, under **TCP Window Scaling**, select the **Windows Scaling** check box to enable window scaling and set the window scaling **Factor**.
4. Click **OK**.

Configuring Selective Acknowledgment (SACK)

NetScaler appliances support Selective Acknowledgment (SACK), as defined in RFC 2018. Using SACK, the data receiver (either a NetScaler appliance or a client) notifies the sender about all the segments that have been received successfully. As a result, the sender (either a NetScaler appliance or a client) needs to retransmit only those segments that were lost during transmission. This improves the performance of data transmission. SACK is important in long fat networks (LFNs). By default, SACK is disabled.

To enable Selective Acknowledgment (SACK) by using the command line interface

At the command prompt, type the following commands to enable Selective Acknowledgment (SACK) and verify the configuration:

- ♦ **set ns tcpParam -SACK (ENABLED | DISABLED)**
- ♦ **show ns tcpParam**

Example

```
> set ns tcpParam -SACK ENABLED
```

To enable Selective Acknowledgment (SACK) by using the configuration utility

1. Navigate to **System > Settings**.
2. In the details pane, under **Settings**, click **Change TCP Parameters**.
3. In the **Configure TCP Parameters** dialog box, under **TCP**, select the **Selective Acknowledgment** check box. For a description of a parameter, hover the mouse cursor over the check box.
4. Click **OK**.

Viewing the HTTP Band Statistics

You can view HTTP band statistics to obtain useful information such as:

- ♦ Average request/response band size.
- ♦ The size range to which most requests/responses belong.
- ♦ Contribution of HTTP pages, in a certain size range, to the overall HTTP traffic.

To view HTTP request and response size statistics by using the command line interface

At the command prompt, type:

```
show protocol httpBand -type (REQUEST|RESPONSE)
```

Example

```
> show protocol httpBand -type REQUEST
```

To view HTTP request and response size statistics by using the configuration utility

1. Navigate to **System > Diagnostics**.
2. In the details pane, under **Troubleshooting Data**, click **HTTP data band statistics**.
3. In the **HTTP Data Band Statistics** dialog box, view the HTTP request and HTTP response size statistics on the **Request** and **Response** tabs, respectively.

You can also modify the band range for HTTP request or response size statistics.

To modify the band range by using the command line interface

At the command prompt, type:

```
set protocol httpBand reqBandSize <value> respBandSize <value>
```

Example

```
> set protocol httpBand reqBandSize 300  
respBandSize 2048
```

To modify the band range by using the configuration utility

1. Navigate to **System > Diagnostics**.
2. In the details pane, under **Troubleshooting Data**, click **HTTP data band statistics**.
3. In the **HTTP Data Band Statistics** dialog box, select the **Request** or the **Response** tab.
4. Click **Configure** and in the dialog box, specify the band size.
5. Click **Close**.

Configuring HTTP Profiles

An HTTP profile is a collection of HTTP parameter settings that can be applied to virtual servers and services. An HTTP profile can be reused on multiple virtual servers or services.

You can use built-in HTTP profiles or configure custom profiles. The following table describes the built-in HTTP profiles.

Table 2-22. Built-in HTTP Profiles

Built-in profile	Description
nshttp_default_strict_validation	Settings for deployments that require strict validation of HTTP requests and responses.
nshttp_default_profile	The default global HTTP settings for the appliance.

To add an HTTP profile by using the command line interface

At the command prompt, type the following commands to add an HTTP profile and verify the configuration:

- ♦ **add ns httpProfile** <name> [-maxReusePool <positive_integer>] [-dropInvalReqs (ENABLED | DISABLED)] [-markHttp09Inval (ENABLED | DISABLED)] [-markConnReqInval (ENABLED | DISABLED)] ...
- ♦ **show ns httpProfile** <name>

Example

```
> add ns httpProfile http_profile1 -maxReusePool
30 -dropInvalReqs ENABLED -markHttp09Inval ENABLED
-markConnReqInval ENABLED
```

To add an HTTP profile by using the configuration utility

1. Navigate to **System > Profiles**.
2. In the details pane, click the **HTTP Profiles** tab, and then click **Add**.
3. In the **Create HTTP Profile** dialog box, configure the parameters for the HTTP profile. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create**.

Configuring WebSocket Connections

WebSocket protocol allows browsers and other clients to create a bi-directional, full duplex TCP connection to the servers. It allows data to be passed back and forth without a request/response model, and so enables live content and creation of real-time games.

The initial handshake is HTTP compliant to satisfy any intermediate devices. During the handshake, the HTTP client connection is upgraded to a WebSocket connection in an HTTP compliant manner. After the upgrade is successful, this connection can be used like a normal TCP connection and does not have to follow any HTTP protocol semantics. When such a connection goes through, the NetScaler appliance tries to interpret any data transfer after the handshake as HTTP, and fails. The WebSocket connections are marked as non-trackable and invalid.

If the HTTP profile that is bound to the virtual server is configured to drop invalid requests, the appliance abruptly closes and resets the connection. However, when the HTTP profile is configured to allow WebSocket connections, the appliance understands the WebSocket handshake. The connection is still marked as non-trackable but it is not marked invalid, and the connection is not dropped.

Configuring WebSocket connections by using the command line interface

At the command prompt, type the following commands to enable websocket connections and verify the configuration:

- ◆ **set ns httpProfile <name> -webSocket (ENABLED | DISABLED)**
- ◆ **show ns httpProfile <name>**

Example: To enable websocket on HTTP profile.

```
> set ns httpProfile http_profile1 -webSocket  
ENABLED
```

Configuring WebSocket connections by using the configuration utility

1. Navigate to **System > Profiles**.
2. In the details pane, click the **HTTP Profiles** tab.
3. Select the HTTP profile for which you want to enable WebSocket connections, and then click **Open**.
4. In the **Configure HTTP Profile** dialog box, select the **Enable WebSocket connections** check box. For a description of the parameter, hover the mouse cursor over the check box.
5. Click **OK**.

Configuring TCP Profiles

A Transmission Control Protocol (TCP) profile is a collection of TCP parameter settings that can be applied to virtual servers and services. A TCP profile can be reused on multiple virtual servers or services. You can use built-in TCP profiles or configure custom profiles. The following table describes the built-in TCP profiles.

Table 2-23. Built-in TCP Profiles

Built-in profile	Description
nstcp_default_tcp_lfp	This profile is useful for long fat pipe networks (WAN) on the client side. Long fat pipe networks have long delay, high bandwidth lines with minimal packet drops.
nstcp_default_tcp_lnp	This profile is useful for long narrow pipe networks (WAN) on the client side. Long narrow pipe networks have considerable packet loss once in a while.

Built-in profile	Description
nstcp_default_tcp_lan	This profile is useful for back-end server connections, where these servers reside on the same LAN as the appliance.
nstcp_default_tcp_lfp_thin_stream	This profile is similar to the nstcp_default_tcp_lfp profile; however, the settings are tuned for small size packet flows.
nstcp_default_tcp_lnp_thin_stream	This profile is similar to the nstcp_default_tcp_lnp profile; however, the settings are tuned for small size packet flows.
nstcp_default_tcp_lan_thin_stream	This profile is similar to the nstcp_default_tcp_lan profile; however, the settings are tuned to small size packet flows.
nstcp_default_tcp_interactive_stream	This profile is similar to the nstcp_default_tcp_lan profile; however, it has a reduced delayed ACK timer and ACK on PUSH packet settings.
nstcp_internal_apps	This profile is useful for internal applications on the appliance (for example, GSLB sitesyncing). This contains tuned window scaling and SACK options for the desired applications. This profile should not be bound to applications other than internal applications.
nstcp_default_profile	This profile represents the default global TCP settings on the appliance.

To add a TCP profile by using the command line interface

At the command prompt, type the following commands to add a TCP profile and verify the configuration:

- ♦ **add ns tcpProfile** <name> [-WS (ENABLED | DISABLED)] [-SACK (ENABLED | DISABLED)] [-WSVal <positive_integer>] [-nagle (ENABLED | DISABLED)] [-ackOnPush (ENABLED | DISABLED)] [-maxBurst <positive_integer>] ...
- ♦ **show ns tcpProfile**

Example

```
> add ns tcpProfile tcp_profile1 -nagle DISABLED -
ackOnPush ENABLED -maxBurst 10 -initialCwnd 6 -
delayedAck 200 -oooQSize 100 -maxPktPerMss 0
```

To add a TCP profile by using the configuration utility

1. Navigate to **System > Profiles**.
2. In the details pane, click on the **TCP Profiles** tab and then click **Add**.
3. In the **Create TCP Profiles** dialog box, configure the parameters for the TCP profile. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create**.

Configuring a Database Profile

A database profile is a named collection of parameters that is configured once but applied to multiple virtual servers that require those particular parameter settings. After creating a database profile, you bind it to load balancing or content switching virtual servers. You can create as many profiles as you need.

To create a database profile by using the command line interface

At the command line, type the following commands to create a database profile and verify the configuration:

- ♦ **add dbProfile** <name> [-interpretQuery (YES | NO)] [-stickiness (YES | NO)] [-kcdAccount <string>]
- ♦ **show dbProfile**

Example

```
> add dbProfile myDBProfile -interpretQuery YES -
stickiness YES -kcdAccount mykcdacct
Done
> show dbProfile myDBProfile
Name: myDBProfile
Interpret Query: YES
Stickyness: YES
KCD Account: mykcdacct
Reference count: 0

Done
>
```

To create a database profile by using the configuration utility

1. Navigate to **System > Profiles**.
2. In the details pane, on the **Database Profiles** tab, do one of the following:
 - To create a database profile, click **Add**.
 - To modify a database profile, click the profile, and then click **Open**.
3. In the **Create Database Profile** or **Configure Database Profile** dialog box, set the following parameters:
 - **Name***
 - **KCD Account**
 - **Interpret Query**
 - **Stickiness**

* Required parameter
4. Click **Create** or **OK**, and then click **Close**.

To bind a database profile to a load balancing or content switching virtual server by using the command line interface

At the command line, type:

```
set (lb | cs) vserver <name> -dbProfileName <string>
```

To bind a database profile to a load balancing or content switching virtual server by using the configuration utility

1. In the navigation pane, expand **Traffic Management**, and then expand **Load Balancing** or **Content Switching**, depending on the type of virtual server to which you want to bind the database profile.
2. Click **Virtual Servers**.
3. In the details pane, select the virtual server, and then click **Open**.
4. In the **Configure Virtual Server (Load Balancing)** or **Configure Virtual Server (Content Switching)** dialog box, on the **Profiles** tab, in the **Database Profile** list, select the database profile.
5. Click **OK**.

Specifying a TCP Buffer Size

You can set the TCP buffer size, both globally and for individual virtual servers and services, through TCP profiles. The value that you set is the minimum value that is advertised by the appliance, and this buffer size is reserved when a client initiates a

connection that is associated with an endpoint-application function, such as compression or SSL. The managed application can request a larger buffer, but if it requests a smaller buffer, the request is not honored, and the specified buffer size is used. If the TCP buffer size is set both at the global level and at the entity level (virtual server or service level), the buffer specified at the entity level takes precedence. If the buffer size that you specify for a service is not the same as the buffer size that you specify for the virtual server to which the service is bound, the appliance uses the buffer size specified for the virtual server for the client-side connection and the buffer size specified for the service for the server-side connection. However, for optimum results, make sure that the values specified for a virtual server and the services bound to it have the same value. The buffer size that you specify is used only when the connection is associated with endpoint-application functions, such as SSL and compression.

You set the TCP buffer size in a custom, entity-level TCP profile by setting the `bufferSize` parameter for the profile. To apply the buffer size setting specified in a custom, entity-level profile, you bind the profile to the virtual server or service. You set the global TCP buffer size by setting the `bufferSize` parameter in the global TCP profile `nstcp_default_profile`. You do not bind `nstcp_default_profile` to an entity. The settings in `nstcp_default_profile` are automatically applied globally.

Note: A high TCP buffer value could limit the number of connections that can be made to the appliance. Additionally, the global TCP parameter `recvBuffSize`, which was set by the use of the `set ns tcpParam` command, has been deprecated. You can now specify the buffer size only through TCP profiles.

To set the TCP buffer size in an entity-level TCP profile by using the command line interface

At the command prompt, type the following commands to set the TCP buffer size and verify the configuration:

- ♦ `set ns tcpProfile <name> -bufferSize <positive_integer>`
- ♦ `show ns tcpProfile <name>`

Example: To set the buffer size to 12000 bytes.

```
> set ns tcpProfile profile1 -bufferSize 12000
```

Note: You can set the TCP buffer size in the global TCP profile by specifying the profile name as `nstcp_default_profile`.

To set the TCP buffer size in a TCP profile by using the configuration utility

1. Navigate to **System > Profiles**.
2. In the details pane, click the **TCP Profiles** tab.

3. Select the profile for which you want to set the TCP buffer size and then click **Open**.

Note: Select `nstcp_default_profile` if you want to set the TCP buffer size in the global TCP profile, .

4. In the **Configure TCP Profile** dialog box, set the TCP buffer size as required. For a description of the parameter, hover the mouse cursor over the text box.
5. Click **OK**.

Optimizing the TCP Maximum Segment Size for a Virtual Server Configuration

You can specify the Maximum Segment Size (MSS) that the NetScaler appliance advertises to a client when the client initiates a connection to a virtual server on the appliance. You can configure the MSS for the virtual servers configured on the appliance in two ways:

- ♦ You can set the MSS for each virtual server to a value of your choice in a TCP profile.
- ♦ You can set the `learnVsvrMSS` global TCP parameter to `ENABLED` to enable MSS learning for all the virtual servers configured on the appliance.

If you know the optimal MSS value for a given virtual server, you can specify the MSS in a TCP profile and bind the profile to the virtual server. When a client initiates a connection with the virtual server, the appliance advertises the specified MSS value to the client. However, if the appliance is also configured to learn the optimum MSS value from bound services (as described in the following section), the learned MSS value takes precedence, and the value specified in the TCP profile is used only until the appliance learns the optimum MSS value. The appliance uses the learned MSS value until the appliance is restarted. If the appliance is restarted, the appliance defaults to the MSS value specified in the virtual server's TCP profile until it learns the MSS value again.

Specifying the MSS Value in a TCP Profile

If you know the optimal MSS value for a given virtual server, you can specify the MSS in a TCP profile and bind the profile to the virtual server. When a client initiates a connection with the virtual server, the NetScaler appliance advertises the specified MSS value to the client.

To specify the MSS value in a TCP profile by using the command line interface

At the command prompt, type the following commands to specify the MSS value in a TCP profile and verify the configuration:

- ♦ `add ns tcpProfile <name> -mss <positive_integer>`

- ♦ **show ns tcpProfile**

```
> add ns tcpProfile tcp_prof1 -mss 1000
```

To specify the MSS value in a TCP profile by using the configuration utility

1. Navigate to **System > Profiles**.
2. In the details pane, do one of the following:
 - To create a TCP profile, click **Add**.
 - To specify the MSS in an existing TCP profile, click the name of the profile, and then click **Open**.
3. In the **Create TCP Profile** or **Configure TCP Profile** dialog box, specify the name and the MSS value. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create** or **OK**.

Configuring the NetScaler to Learn the MSS Value from Bound Services

If you set the global TCP parameter `learnVsvrMSS` to `ENABLED`, the appliance learns the most frequently used MSS value for each configured virtual server. When a client connects to a virtual server, the appliance advertises to the client the MSS value that is optimum for that virtual server. The optimum value is the MSS of the service or subset of bound services that are most frequently selected during load balancing. Consequently, each virtual server configuration uses its own MSS value. This enhancement enables the appliance to optimize the consumption of system resources.

The default value of the `learnVsvrMSS` parameter is `DISABLED`. When enabled, MSS learning is applicable only to virtual servers of type TCP, HTTP, and FTP.

To configure the appliance to learn the MSS for a virtual server by using the command line interface

At the command prompt, type the following commands to configure the appliance to learn the MSS for a virtual server and verify the configuration:

- ♦ **set ns tcpParam -learnVsvrMSS (ENABLED|DISABLED)**
- ♦ **show ns tcpParam**

Example

```
> set ns tcpParam -learnVsvrMSS ENABLED
```

To configure the appliance to learn the MSS for a virtual server by using the configuration utility

1. Navigate to **System > Settings**.
2. In the details pane, under **Settings**, click **Change TCP parameters**.
3. In the **Configure TCP Parameters** dialog box, select the **Learn MSS** check box.

Reporting Tool

Use the Citrix® NetScaler® Reporting tool to view NetScaler performance statistics data as reports. Statistics data are collected by the nscollect utility and are stored in a database. When you want to view certain performance data over a period of time, the Reporting tool pulls out specified data from the database and displays them in charts.

Reports are a collection of charts. The Reporting tool provides built-in reports as well as the option to create custom reports. In a report, you can modify the charts and add new charts. You can also modify the operation of the data collection utility, nscollect, and stop or start its operation.

Using the Reporting Tool

The Reporting tool is a Web-based interface accessed from the Citrix® NetScaler® appliance. Use the Reporting tool to display the performance statistics data as reports containing graphs. In addition to using the built-in reports, you can create custom reports, which you can modify at any time. Reports can have between one and four charts. You can create up to 256 custom reports.

To invoke the Reporting tool

1. Use the Web browser of your choice to connect to the IP address of the NetScaler (for example, <http://10.102.29.170/>).
The Web Logon screen appears.
2. In the **User Name** text box, type the user name assigned to the NetScaler.
3. In the **Password** text box, type the password.
4. In the **Start in** drop-down box, select **Reporting**.
5. Click **Login**.

The following screen shots show the report toolbar and the chart toolbar, which are frequently referenced in this documentation.

Figure 2-2. Report Toolbar

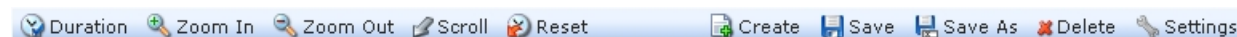


Figure 2-3. Chart Toolbar



Working with Reports

You can plot and monitor statistics for the various functional groups configured on the NetScaler over a specified time interval. Reports enable you to troubleshoot or analyze the behavior of your appliance. There are two types of reports: built-in reports and custom reports. Report content for built-in or custom reports can be viewed in a graphical format or a tabular format. The graphical view consists of line, area, and bar charts that can display up to 32 sets of data (counters). The tabular view displays the data in columns and rows. This view is useful for debugging error counters.

The default report that is displayed in the Reporting tool is CPU vs. Memory Usage and HTTP Requests Rate. You can change the default report view by displaying the report you want as your default view, and then clicking **Default Report**.

Reports can be generated for the last hour, last day, last week, last month, last year, or you can customize the duration.

You can do the following with reports:

- ◆ Toggle between a tabular view of data and a graphical view of data.
- ◆ Change the graphical display type, such as bar chart or line chart.
- ◆ Customize charts in a report.
- ◆ Export the chart as an Excel comma-separated value (CSV) file.
- ◆ View the charts in detail by zooming in, zooming out, or using a drag-and-drop operation (scrolling).
- ◆ Set a report as the default report for viewing whenever you log on.
- ◆ Add or remove counters.
- ◆ Print reports.
- ◆ Refresh reports to view the latest performance data.

Using Built-in Reports

The Reporting tool provides built-in reports for frequently viewed data. Built-in reports are available for the following functional groups: System, Network, SSL, Compression, Integrated Cache, and Citrix NetScaler Application Firewall. By default, the built-in reports are displayed for the last day. However, you can view the reports for the last hour, last week, last month, or last year.

Note: You cannot save changes to built-in reports, but you can save a modified built-in report as a custom report.

To display a built-in report

1. In the left pane of the Reporting tool, under **Built-in Reports**, expand a group (for example, **SSL**).
2. Click a report (for example, **SSL > All Backend Ciphers**).

Creating and Deleting Reports

You can create your own custom reports and save them with user-defined names for reuse. You can plot different counters for different groups based on your requirements. You can create up to 256 custom reports.

You can either create a new report or save a built-in report as a custom report. By default, a newly created custom report contains one chart named **System Overview**, which displays the **CPU Usage** counter plotted for the last day. You can customize the interval and set the data source and time zone from the report toolbar. Within a report, you can use the chart toolbars to add, modify, or delete charts, as described in ["Working with Charts."](#)

By default, newly created custom reports contain one chart named **System Overview** that displays a **CPU Usage** counter plotted for the last day.

To create a custom report

1. In the Reporting tool, on the report toolbar, click **Create**, or if you want to create a new custom report based on an existing report, open the existing report, and then click **Save As**.
2. In **Report Name** box, type a name for the custom report.
3. Do one of the following:
 - To add the report to an existing folder, in **Create in** or **Save in**, click the down arrow to choose an existing folder, and then click **OK**.
 - To create a new folder to store the report, click the **Click to add folder** icon, in **Folder Name**, type the name of the folder, and in **Create in**, specify where you want the new folder to reside in the hierarchy, and then click **OK**.

Note: You can create up to 128 folders.

To delete a custom report







1. In the left pane of the Reporting tool, next to **Custom Reports**, click the **Click to manage custom reports** icon.
2. Select the check box that corresponds with the report you want to delete, and then click **Delete**.

Note: When you delete a folder, all the contents of that folder are deleted.

Modifying the Time Interval

By default, built-in reports display data for the last day. However, if you want to change the time interval for a built-in report, you can save the report as a custom report. The new interval applies to all charts in the report. The following table describes the time-interval options.

Table 2-24. Time Intervals

Time interval	Displays
 Last Hour	Statistics data collected for the last hour.
 Last Day	Statistics data collected for the last day (24 hours).
 Last Week	Statistics data collected for the last week (7 days).
 Last Month	Statistics data collected for the last month (31 days).
 Last Year	Statistics data collected for the last year (365 days).
 Custom	Statistics data collected for a time period that you are prompted to specify.

To modify the time interval

1. In the left pane of the Reporting tool, click a report.
2. On the report toolbar, click **Duration**, and then click a time interval.

Setting the Data Source and Time Zone

You can retrieve data from different data sources to display them in the reports. You can also define the time zone for the reports and apply the currently displayed report's time selection to all the reports, including the built-in reports.

To set the data source and time zone

1. In the **Reporting tool**, on the report toolbar, click **Settings**.
2. In the **Settings** dialog box, in **Data Source**, select the data source from which you want to retrieve the counter information.
3. Do one or both of the following:
 - If you want the tool to remember the time period for which a chart is plotted, select the **Remember time selection for charts** check box.
 - If you want the reports to use the time settings of your NetScaler appliance, select the **Use Appliance's time zone** check box.

Exporting and Importing Custom Reports

You can share reports with other NetScaler administrators by exporting reports. You can also import reports.

To export or import custom reports

1. In the left pane of the Reporting tool, next to **Custom Reports**, click the **Click to manage custom reports** icon.
2. Select the check box that corresponds with the report you want to export or import, and then click **Export** or **Import**.

Note: When you export the file, it is exported in a .gz file format.

Working with Charts

Use charts to plot and monitor counters or groups of counters. You can include up to four charts in one report. In each chart, you can plot up to 32 counters. The charts can use different graphical formats (for example, area and bar). You can move the charts up or down within the report, customize the colors and visual display for each counter in a chart, and delete a chart when you do not want to monitor it.

In all report charts, the horizontal axis represents time and the vertical axis represents the value of the counter.

Adding a Chart

When you add a chart to a report, the **System Overview** chart appears with the **CPU Usage** counter plotted for the last one day. To plot a different group of statistics or select a different counter, see "[Modifying a Chart](#)."

Note: If you add charts to a built-in report, and you want to retain the report, you must save the report as a custom report.

Use the following procedure to add a chart to a report.

To add a chart to a report

1. In the left pane of the **Reporting tool**, click a report.
2. Under the chart where you want to add the new chart, click the **Add** icon.

Modifying a Chart

You can modify a chart by changing the functional group for which the statistics are displayed and by selecting different counters.

To modify a chart

1. In the left pane of the Reporting tool, click a report.
2. Under the chart that you want to modify, click **Counters**.

3. In the dialog box that appears, in the **Title** box, type a name for the chart.
4. Next to **Plot chart for**, do one of the following:
 - To plot counters for global counters, such as Integrated Cache and Compression, click **System global statistics**.
 - To plot entity counters for entity types, such as Load Balancing and GSLB, click **System entities statistics**.
5. In **Select group**, click the desired entity.
6. Under **Counters**, in **Available**, click the counter name(s) that you want to plot, and then click the > button.
7. If you selected **System entities statistics** in step 4, on the **Entities** tab, under **Available**, click the entity instance name(s) you want to plot, and then click the > button.
8. Click **OK**.

Viewing a Chart

You can specify the graphical formats of the plotted counters in a chart. Charts can be viewed as line charts, spline charts, step-line charts, scatter charts, area charts, bar charts, stacked area charts, and stacked bar charts. You can also zoom in, zoom out, or scroll inside the plot area of a chart. You can zoom in or out for all data sources for 1 hour, 1 day, 1 week, 1 month, 1 year, and 3 years.

Other options for customizing the view of a chart include customizing the axes of the charts, changing the background and edge color of the plot area, customizing the color and size of the grids, and customizing the display of each data set (counter) in a chart.

Data set numbers, such as Data Set 1, correspond to the order in which the counters in your graph are displayed at the bottom of the chart. For example, if **CPU usage** and **Memory usage** are displayed in first and second order at the bottom of the chart, **CPU usage** is equal to **Data Set 1** and **Memory usage** is equal to **Data Set 2**.

Whenever you modify a built-in report, you need to save the report as a custom report to retain your changes.

To change the graph type of a chart

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, under the chart you want to view, on the chart toolbar, click **Customize**.
3. On the **Chart** tab, under **Category**, click **Plot type**, and then click the graph type you want to display for the chart. If you want to display the graph is 3D, select the **Use 3D** check box.

To refocus a chart with detailed data

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, on the report toolbar, click **Zoom In**, and do one or both of the following:

- To refocus the chart to display data for a specific time window, drag and drop the cursor from the start time to the end time. For example, you can view data for a one-hour period on a certain day.
 - To refocus the chart to display data for a data point, simply click once on chart where you want to zoom in and get more detailed information.
3. Once you have the desired range of time for which you want to view detailed data, on the report toolbar, click **Tabular View**. Tabular view displays the data in numeric form in rows and columns.

To view numeric data for a graph

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, on the report toolbar, click **Tabular View**. To return to the graphical view, click **Graphical View**.

Note: You can also view the numeric data in the graphical view by hovering your cursor over the notches in the gridlines.

To scroll through time in a chart

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, on the report toolbar, click **Scroll**, and then click inside the chart and drag the cursor in the direction for which you want to see data for a new time period. For example, if you want to view data in the past, click and drag to the left.

To change the background color and text color of a chart

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, under the chart for which you want to customize the axes, click **Customize**.
3. On the **Chart** tab, under **Category**, click one or more of the following:
 - To change the background color, click **Background Color**, and then select the options for color, transparency, and effects.
 - To change the text color, click **Text Color**, and then select the options for color, transparency, and effects.

To customize the axes of a chart

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, under the chart for which you want to customize the axes, click **Customize**.
3. On the **Chart** tab, under **Category**, click one or more of the following:
 - To change the scale of the left y-axis, click **Left Y-Axis**, and then select the scale you want.

- To change the scale of the right y-axis, click **Right Y-Axis**, in **Data set to plot**, select the data set, and then select the scale you want.

Note: The data set numbers, such as Data Set 1, correspond to the order in which the counters in your graph are displayed at the bottom of the chart. For example, if **CPU usage** and **Memory usage** are displayed in first and second order at the bottom of the chart, **CPU usage** is equal to **Data Set 1** and **Memory usage** is equal to **Data Set 2**.

- To plot each data set in its own hidden y-axis, click **Multiple Axes**, and then click **Enable**.

To change the background color, edge color, and gridlines for a plot area of a chart

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, under the chart for which you want to customize the plot area, click **Customize**.
3. On the **Plot Area** tab, under **Category**, click one or more of the following:
 - To change the background color and edge color of the chart, click **Background Color** and **Edge Color**, and then select the options for color, transparency, and effects.
 - To change the horizontal or vertical grids of the chart, click **Horizontal Grids** or **Vertical Grids**, and then select the options for displaying the grids, grid width, grid color, transparency, and effects.

To change the color and graph type of a data set

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, under the chart for which you want to customize the display of the data set (counters), click **Customize**.
3. On the **Data Set** tab, in **Select Data Set**, select the data set (counter) for which you want to customize the graphical display.

Note: The data set numbers, such as Data Set 1, correspond to the order in which the counters in your graph are displayed at the bottom of the chart. For example, if **CPU usage** and **Memory usage** are displayed in first and second order at the bottom of the chart, **CPU usage** is equal to **Data Set 1** and **Memory usage** is equal to **Data Set 2**.

4. Under **Category**, do one of more of the following:
 - To change the background color, click **Color**, and then select the options for color, transparency, and effects.
 - To change the graph type, click **Plot type**, and then select the graph type you want to display for the data set. If you want to display the graph as 3D, select the **Use 3D** check box.

Exporting Chart Data to Excel

For further data analysis, you can export charts to Excel in a comma-separated value (CSV) format.

To export chart data to Excel

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, under the chart with the data you want to export to Excel, click **Export**.

Deleting a Chart

If you do not want to use a chart, you can remove it from the report. You can permanently remove charts from custom reports only. If you delete a chart from a built-in report and want to retain the changes, you need to save the report as a custom report.

To delete a chart

1. In the left pane of the Reporting tool, select a report.
2. In the right pane, under the chart that you want to delete, click the **Delete** icon.

Examples

To display the trend report for CPU usage and memory usage for the last week

1. In the left pane of the Reporting tool, under **Built-in Reports**, expand **System**.
2. Click the report **CPU vs. Memory Usage and HTTP Requests Rate**.
3. In the right pane, on the report toolbar, click **Duration**, and then click **Last Week**.

To compare the bytes received rate and the bytes transmitted rate between two interfaces for the last week

1. In the right pane, on the report toolbar, click **Create**.
2. In the **Report Name** box, type a name for the custom report (for example, `Custom_Interfaces`), and then click **OK**.
The report is created with the default **System Overview** chart, which displays the **CPU Usage** counter plotted for the last hour.
3. Under **System Overview**, on the chart toolbar, click **Counters**.
4. In the counter selection pane, in **Title**, type a name for the chart (for example, `Interfaces bytes data`).
5. In **Plot chart for**, click **System entities statistics**, and then in **Select Group**, select **Interface**.

6. On the **Entities** tab, click the interface name(s) you want to plot (for example, 1/1 and 1/2), and then click the > button.
7. On the **Counters** tab, click **Bytes received (Rate)** and **Bytes transmitted (Rate)** and then click the > button.
8. Click **OK**.
9. On the report toolbar, click **Duration**, and then click **Last Week**.

Stopping and Starting the Data Collection Utility

The performance data is stored in different data sources on the Citrix NetScaler appliance. The default data source is `/var/log/db/default`. You can create up to 32 data sources.

The data collection utility `nscollect` retrieves data from the NetScaler and updates the data source. This utility runs automatically when you start the NetScaler. It creates a database for global counters at `/var/log/db/<DataSourceName>`. The entity-specific databases are created based on the entities configured on the NetScaler. A specific folder is created for each entity type in

```
/var/log/db/<DataSourceName/EntityNameDB>
```

Before creating a database for an entity, `nscollect` allocates a unique number to the entity and creates the database based on that number. It retrieves all the counters available for a group. However, there is a limit on the number of different entities that `nscollect` can retrieve, as described in the following table.

Table 2-25. Limits on Entity Numbers Retrieved by nscollect

Entity name	Limit
Content Switching Virtual Servers	100
Cache Redirection Virtual Servers	50
DOS Policies	100
GSLB Domains	100
GSLB Services	100
GSLB Sites	32
GSLB Virtual Servers	100
Interfaces	8

Entity name	Limit
LB Virtual Servers	100
ACLs	100
ACL6	50
Priority Queuing Policies	100
RNAT IP Addresses	100
SureConnect Policies	100
Services	250
Service Groups	100
System CPU	8
VLAN	25
VPN Virtual Servers	5

The `nscollect` utility retrieves n number of entity counters and creates the entity database. If the first n counters change in the subsequent fetch, the database stores more than n entries for that entity type. However, you need to delete the unused entity counters manually.

Note: The Reporting tool supports only numerical counters.

By default, `nscollect` retrieves data at every 5-minute interval. Data is maintained in 5-minute granularity for one day, hourly for the last 30 days, and daily for three years.

When you start the NetScaler, the `nscollect` utility automatically starts. However, if data is not updated accurately, or there is corrupted data displayed in the reports, you can stop and then restart the utility. You may also want to stop `nscollect` to back up the databases or to create a new data source.

To stop nscollect

At the command prompt, type the following:

```
/netscaler/nscollect stop
```

You can start `nscollect` on either the local system or a remote system.

To start nscollect on the local system

At the command prompt, type the following:

```
/netscaler/nscollect start
```

To start nscollect on the remote system

At the command prompt, type the following:

```
/netscaler/nscollect start -U NS_IP:UserName:Password -ds DataSourceName
```

Example

```
/netscaler/nscollect start -U  
10.102.29.170:nsroot:nsroot -ds default
```

Chapter 3

AppFlow

Topics:

- [How AppFlow Works](#)
- [Configuring the AppFlow Feature](#)
- [Exporting Performance Data of Web Pages to AppFlow Collector](#)

The Citrix NetScaler appliance is a central point of control for all application traffic in the data center. It collects flow and user-session level information valuable for application performance monitoring, analytics, and business intelligence applications. It also collects web page performance data and database information. AppFlow transmits the information by using the Internet Protocol Flow Information eXport (IPFIX) format, which is an open Internet Engineering Task Force (IETF) standard defined in RFC 5101. IPFIX (the standardized version of Cisco's NetFlow) is widely used to monitor network flow information. AppFlow defines new Information Elements to represent application-level information, web page performance data, and database information.

Using UDP as the transport protocol, AppFlow transmits the collected data, called *flow records*, to one or more IPv4 collectors. The collectors aggregate the flow records and generate real-time or historical reports.

AppFlow provides visibility at the transaction level for HTTP, SSL, TCP, and SSL_TCP flows. You can sample and filter the flow types that you want to monitor.

AppFlow use actions and policies to send records for a selected flow to specific set of collectors. An AppFlow action specifies which set of collectors will receive the AppFlow records. Policies, which are based on Advanced expressions can be configured to select flows for which flow records will be sent to the collectors specified by the associated AppFlow action.

To limit the types of flows, you can enable AppFlow for a virtual server. AppFlow can also provide statistics for the virtual server.

You can also enable AppFlow for a specific service, representing an application server, and monitor the traffic to that application server.

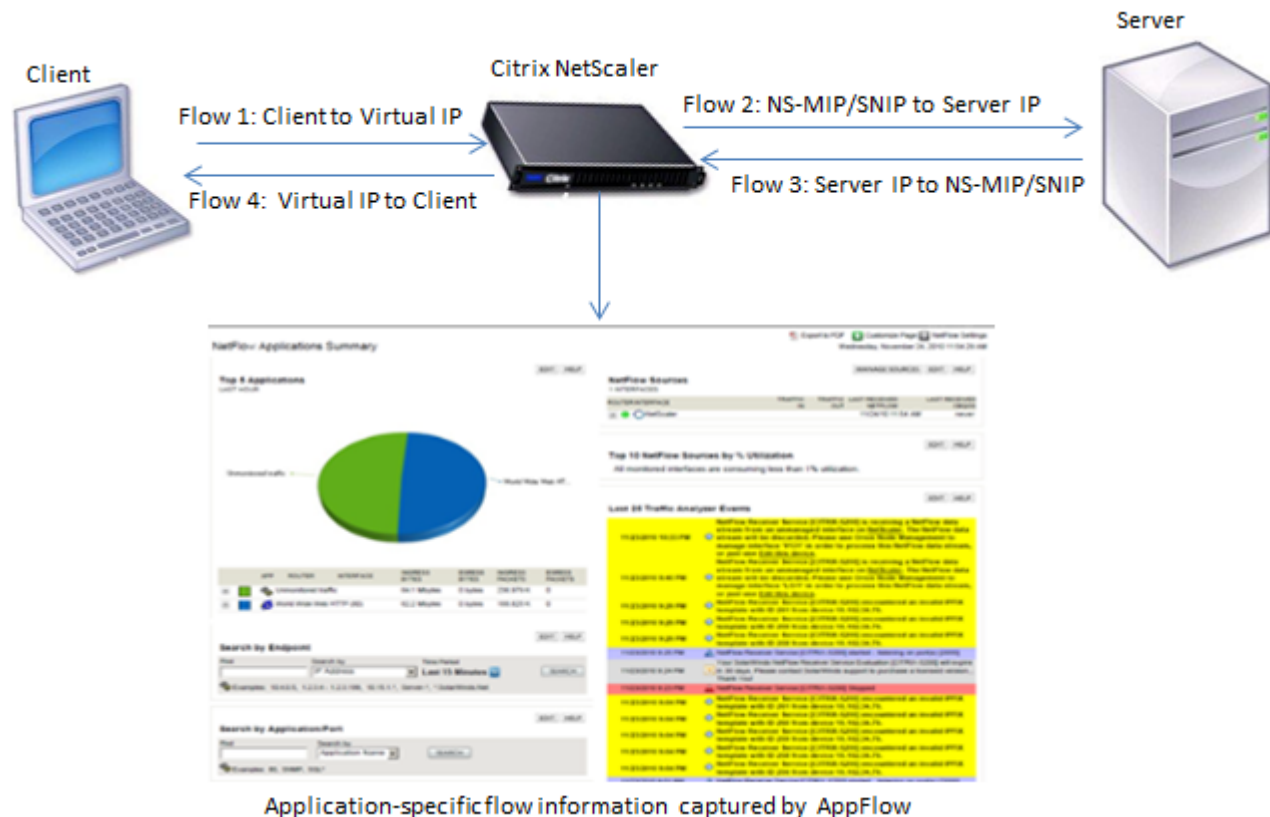
Note: This feature is supported only on NetScaler nCore builds.

How AppFlow Works

In the most common deployment scenario, inbound traffic flows to a Virtual IP address (VIP) on the NetScaler appliance and is load balanced to a server. Outbound traffic flows from the server to a mapped or subnet IP address on the NetScaler and from the VIP to the client. A flow is a unidirectional collection of IP packets identified by the following five tuples: sourceIP, sourcePort, destIP, destPort, and protocol.

The following figure describes how the AppFlow feature works.

Figure 3-1. NetScaler Flow Sequence



As shown in the figure, the network flow identifiers for each leg of a transaction depend on the direction of the traffic.

The different flows that form a flow record are:

Flow1: <Client-IP, Client-Port, VIP-IP, VIP-port, Protocol>

Flow2: <NS-MIP/SNIP, NS-port, Server-IP, Server-Port, Protocol>

Flow3: <Server-IP, Server-Port, NS-MIP/SNIP, NS-Port, Protocol>

Flow4: <VIP-IP, VIP-port, Client-IP, Client-Port, Protocol>

To help the collector link all four flows in a transaction, AppFlow adds a custom `transactionID` element to each flow. For application-level content switching, such as HTTP, it is possible for a single client TCP connection to be load balanced to different backend TCP connections for each request. AppFlow provides a set of records for each transaction.

Flow Records

AppFlow records contain standard NetFlow or IPFIX information, such as time stamps for the beginning and end of a flow, packet count, and byte count. AppFlow records also contain application-level information (such as HTTP URLs, HTTP request methods and response status codes, server response time, and latency), web page performance data (such as page load time, page render time, and time spent on the page), and database information (such as database protocol, database response status and database response size). IPFIX flow records are based on templates that need to be sent before sending flow records.

Templates

AppFlow defines a set of templates, one for each type of flow. Each template contains a set of standard Information Elements (IEs) and Enterprise-specific Information Elements (EIEs). IPFIX templates define the order and sizes of the Information Elements (IE) in the flow record. The templates are sent to the collectors at regular intervals, as described in RFC 5101.

A template can include the following EIEs:

transactionID

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

connectionID

An unsigned 32-bit number identifying a layer-4 connection (TCP or UDP). The NetScaler flows are usually bidirectional, with two separate flow records for each direction of the flow. This information element can be used to link the two flows.

For the NetScaler, `connectionID` is an identifier for the connection data structure to track the progress of a connection. In an HTTP transaction, for instance, a given `connectionID` may have multiple `transactionID` elements corresponding to multiple requests that were made on that connection.

tcpRTT

The round trip time, in milliseconds, as measured on the TCP connection. This can be used as a metric to determine the client or server latency on the network.

httpRequestMethod

An 8-bit number indicating the HTTP method used in the transaction. An options template with the number-to-method mapping is sent along with the template.

httpRequestSize

An unsigned 32-bit number indicating the request payload size.

httpRequestURL

The HTTP URL requested by the client.

httpUserAgent

The source of incoming requests to the Web server.

httpResponseStatus

An unsigned 32-bit number indicating the response status code.

httpResponseSize

An unsigned 32-bit number indicating the response size.

httpResponseTimeToFirstByte

An unsigned 32-bit number indicating the time taken to receive the first byte of the response.

httpResponseTimeToLastByte

An unsigned 32-bit number indicating the time taken to receive the last byte of the response.

flowFlags

An unsigned 64-bit flag used to indicate different flow conditions.

EIEs for web page performance data

clientInteractionStartTime

Time at which the browser receives the first byte of the response to load any objects of the page such as images, scripts, and stylesheets.

clientInteractionEndTime

Time at which the browser received the last byte of response to load all the objects of the page such as images, scripts, and stylesheets.

clientRenderStartTime

Time at which the browser starts to render the page.

clientRenderEndTime

Time at which browser finished rendering the entire page, including the embedded objects.

EIEs for database information

dbProtocolName

An unsigned 8-bit number indicating the database protocol. Valid values are 1 for MS SQL and 2 for MySQL.

dbReqType

An unsigned 8-bit number indicating the database request method used in the transaction. For MS SQL, valid values are 1 is for QUERY, 2 is for TRANSACTION, and 3 is for RPC. For valid values for MySQL, see the MySQL documentation.

dbReqString

Indicates the database request string without the header.

dbRespStatus

An unsigned 64-bit number indicating the status of the database response received from the web server.

dbRespLength

An unsigned 64-bit number indicating the response size.

dbRespStatString

The response status string received from the web server.

Configuring the AppFlow Feature

You configure AppFlow in the same manner as most other policy-based features. First, you enable the AppFlow feature. Then you specify the collectors to which the flow records are sent. After that, you define actions, which are sets of configured collectors. Then you configure one or more policies and associate an action to each policy. The policy tells the NetScaler appliance to select requests the flow records of which are sent to the associated action. Finally, you bind each policy either globally or to specific vservers to put it into effect.

You can further set AppFlow parameters to specify the template refresh interval and to enable the exporting of httpURL, httpCookie, and httpReferer information. On each collector, you must specify the NetScaler IP address as the address of the exporter.

Note: For information about configuring the NetScaler as an exporter on the collector, see the documentation for the specific collector.

The configuration utility provides tools that help users define the policies and actions that determine exactly how the NetScaler appliance export records for a particular flow to a set of collectors(action.) The command line interface provides a corresponding set of CLI-based commands for experienced users who prefer a command line.

Enabling or Disabling AppFlow

To be able to use the AppFlow feature, you must first enable it.

Note: AppFlow can be enabled only on nCore NetScaler appliances.

To enable or disable the AppFlow feature by using the command line interface

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

- ♦ `enable ns feature appflow`
- ♦ `disable ns feature appflow`

To enable the AppFlow feature by using the configuration utility

1. Navigate to **System > Settings**.
2. In the details pane, under **Modes and Features**, click **Configure advanced features**.
3. In the **Configure Advanced Features** dialog box, select the **AppFlow** check box, and then click **OK**.
4. In the **Enable/Disable Feature(s)?** dialog box, click **Yes**.

Specifying a 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. However, you cannot export the same data to multiple collectors. You can remove unused collectors. By default, the collector listens to IPFIX messages on UDP port 4739. You can change the default port, when configuring the collector. Similarly, by default, NSIP is used as the source IP for appflow traffic. You can change this default source IP to a SNIP or MIP address when configuring a collector.

To specify a collector by using the command line interface

At the command prompt, type the following commands to add a collector and verify the configuration:

- ♦ `add appflow collector <name> -IPAddress <ipaddress> -port <port_number> -netprofile <netprofile_name>`
- ♦ `show appflow collector <name>`

Example

```
> add appflow collector coll -IPAddress  
10.102.29.251 -port 8000 -netprofile n2
```

To specify a collector by using the configuration utility

1. Navigate to **System > AppFlow > Collectors**.

2. In the details pane, click **Add**.
3. In the **Create AppFlow Collector** dialog box, specify values for the following parameters, which correspond to parameters described in "Parameters for specifying a collector" as shown:
 - **Name***
 - **IP Address***
 - **Port**
 - **Net Profile**

*A required parameter
4. To remove a collector from the list, select the collector, and then click **Remove**.
5. Click **Create**, and then click **Close**.

Configuring an AppFlow Action

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

To configure an AppFlow action by using the command line interface

At the command prompt, type the following commands to configure an Appflow action and verify the configuration:

- ♦ **add appflow action** <name> --collectors <string> ... [-comment <string>]
- ♦ **show appflow action**

Example

```
> add appflow action apfl-act-collector-1-and-3 -  
collectors collector-1 collector-3
```

To configure an AppFlow action by using the configuration utility

1. Navigate to **System > AppFlow > Actions**.
2. In the details pane, do one of the following:
 - To create a new action, click **Add**.
 - To modify an existing action, select the action, and then click **Open**.
3. In the **Add AppFlow Action** or **Configure AppFlow Action** dialog box, type a name for the new action or the name of an existing action, respectively.

4. Do one of the following to associate collectors with the action:
 - If the collectors that you want are listed, click the corresponding check boxes.
 - If you want to specify all the collectors, click **Activate All**.
 - If you want to specify a new collector, click **Add**.
5. Click **Create** or **OK**, depending on whether you are creating a new action or modifying an existing action.
6. Click **Close**. A message appears in the status bar, stating that the configuration has been successfully implemented.

Configuring an AppFlow Policy

After you configure an AppFlow action, you must next configure an AppFlow policy. An AppFlow policy is based on a rule, which consists of one or more expressions.

Note: For creating and managing AppFlow policies, the configuration utility provides assistance that is not available at the command line interface.

To configure an AppFlow policy by using the command line interface

At the command prompt, type the following command to add an AppFlow policy and verify the configuration:

- ♦ **add appflow policy** <name> <expression> <action>
- ♦ **show appflow policy** <name>

Example

```
> add appflow policy apfl-pol-tcp-dsprt
client.TCP.DSTPORT.EQ(22) apfl-act-collector-1-
and-3
```

To configure an AppFlow policy by using the configuration utility

1. Navigate to **System > AppFlow > Policies**.
2. In the details pane, do one of the following:
 - To create a new policy, click **Add**.
 - To modify an existing policy, select the policy, and then click **Open**.
3. In the **Create AppFlow Policy** or **Configure AppFlow Policy** dialog box, specify the relevant parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.

4. Click **Create** or **OK**, depending on whether you are creating a new policy or modifying an existing policy.
5. Click **Close**. A message appears in the status bar, stating that the configuration has been successfully implemented.

To add an expression by using the Add Expression dialog box

1. In the **Add Expression** dialog box, in the first list box choose the first term for your expression.

HTTP

The HTTP protocol. Choose this if you want to examine some aspect of the request that pertains to the HTTP protocol.

SYS

The protected Web site(s). Choose this if you want to examine some aspect of the request that pertains to the recipient of the request.

CLIENT

The computer that sent the request. Choose this if you want to examine some aspect of the sender of the request.

When you make your choice, the rightmost list box lists appropriate terms for the next part of your expression.

2. In the second list box, choose the second term for your expression.
The choices depend upon which choice you made in the previous step, and are appropriate to the context. After you make your second choice, the **Help** window below the **Construct Expression** window (which was blank) displays help describing the purpose and use of the term you just chose.
3. Continue choosing terms from the list boxes that appear to the right of the previous list box, or typing strings or numbers in the text boxes that appear to prompt you to enter a value, until your expression is finished.

Binding an AppFlow Policy

To put a policy into effect, you must bind it either globally, so that it applies to all traffic that flows through the NetScaler, or to a specific virtual server, so that the policy applies only to the traffic related to that virtual server.

When you bind a policy, you assign it a priority. The priority determines the order in which the policies you define are evaluated. You can set the priority to any positive integer.

In the NetScaler operating system, policy priorities work in reverse order—the higher the number, the lower the priority. For example, if you have three policies with priorities of 10, 100, and 1000, the policy assigned a priority of 10 is performed first, then the policy assigned a priority of 100, and finally the policy assigned an order of 1000.

You can leave yourself plenty of room to add other policies in any order, and still set them to evaluate in the order you want, by setting priorities with intervals of 50 or 100

between each policy when you globally bind it. You can then add additional policies at any time without having to change the priority of an existing policy.

To globally bind an AppFlow policy by using the command line interface

At the command prompt, type the following command to globally bind an AppFlow policy and verify the configuration:

- ♦ `bind appflow global <policyName> <priority> [<gotoPriorityExpression> [-type <type>] [-invoke (<labelType> <labelName>)]`
- ♦ `show appflow global`

Example

```
bind appflow global af_policy_lb1_10.102.71.190 1
NEXT -type REQ_OVERRIDE -invoke vserver google
```

To bind an AppFlow policy to a specific virtual server by using the command line interface

At the command prompt, type the following command to bind an appflow policy to a specific virtual server and verify the configuration:

`bind lb vserver <name> -policyname <policy_name> -priority <priority>`

Example

```
bind lb vserver google -policyname
af_policy_google_10.102.19.179 -priority 251
```

To globally bind an AppFlow policy by using the configuration utility

1. Navigate to **System > AppFlow**.
2. On the **AppFlow** page, click **Policy Manager**.
3. In the **AppFlow Policy Manager** dialog box, in the **Bind Points** menu, select **Default Global**.
4. Click **Insert Policy** to insert a new row and display a drop-down list of all unbound AppFlow policies.
5. Click one of the policies on the list. That policy is inserted into the list of globally bound AppFlow policies.
6. Click **Apply Changes**.

7. Click **Close**. A message appears in the status bar, stating that the configuration has been successfully implemented.

To bind an AppFlow policy to a specific virtual server by using the configuration utility

1. Navigate to **Traffic Management > Load Balancing > Virtual Servers**.
2. On the **Load Balancing Virtual Servers** page, select the virtual server to which you want to bind the AppFlow policy, and then click **Open**.
3. In the **Configure Virtual Server (Load Balancing)** dialog box, select the **Policies** tab to display the policies bound to that particular virtual server.
4. Click **Insert Policy** to insert a new row and display a drop-down list of all unbound AppFlow policies.
5. From the drop-down list that appears under **Policy Name**, select the policy that you want to bind to this virtual server.
6. Click **OK**, and then click **Close**. A message appears in the status bar, stating that the configuration has been successfully implemented.

Enabling AppFlow for Virtual Servers

If you want to monitor only the traffic through certain virtual servers, enable AppFlow specifically for those virtual servers. You can enable AppFlow for load balancing, content switching, cache redirection, SSL VPN, GSLB, and authentication virtual servers.

To enable AppFlow for a virtual server by using the command line interface

At the command prompt, type:

```
set <feature_name> vserver <vServerName> <protocol> <IPAddress> <port> -  
appflowLog ENABLED
```

Example

```
> set cs vserver Vserver-CS-1 HTTP 10.102.29.161  
80 -appflowLog ENABLED
```

To enable AppFlow for a virtual server by using the configuration utility

1. In the navigation pane, expand the feature node for which you want to enable AppFlow, and then click **Virtual Servers**.
For example, to enable AppFlow for a content switching virtual server, navigate to **Traffic Management > Content Switching > Virtual Servers**.

2. In the details pane, do one of the following:
 - To enable AppFlow for a new virtual server, click **Add**.
 - To enable AppFlow for an existing virtual server, select the virtual server, and then click **Open**.
3. In the **Create Virtual Server** (feature_name) dialog box or the **Configure Virtual Server** (feature_name) dialog box, select the **AppFlow Logging** check box.
4. Click **Create** or **OK**, and then click **Close**.

Enabling AppFlow for a Service

You can enable AppFlow for services that are to be bound to the load balancing virtual servers.

To enable AppFlow for a service by using the command line interface

At the command prompt, type:

```
set service <name> -appflowLog ENABLED
```

Example

```
set service ser -appflowLog ENABLED
```

To enable AppFlow for a service by using the configuration utility

1. Navigate to **Traffic Management > Load Balancing > Services**.
2. In the details pane, do one of the following:
 - To enable AppFlow for a new service, click **Add**.
 - To enable AppFlow for an existing service, select the service, and then click **Open**.
3. In the **Create Service** or the **Configure Service** dialog box, select the **AppFlow Logging** check box.
4. Click **OK**, and then click **Close**.

Setting the AppFlow Parameters

You can set AppFlow parameters to customize the exporting of data to the collectors.

To set the AppFlow Parameters by using the command line interface

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

- ♦ **set appflow param** [-templateRefresh <secs>] [-appnameRefresh <secs>] [-flowRecordInterval <secs>] [-udpPmtu <positive_integer>] [-httpUrl (ENABLED | DISABLED)] [-httpCookie (ENABLED | DISABLED)] [-httpReferer (ENABLED | DISABLED)] [-httpMethod (ENABLED | DISABLED)] [-httpHost (ENABLED | DISABLED)] [-httpUserAgent (ENABLED | DISABLED)] [-httpXForwardedFor (ENABLED | DISABLED)][-clientTrafficOnly (YES | NO)]
- ♦ **show appflow Param**

Example

```
> set appflow Param -templateRefresh 240 -udpPmtu
128 -httpUrl enabled
```

To set the AppFlow parameters by using the configuration utility

1. Navigate to **System > AppFlow**.
2. On the **AppFlow** landing page, under **Settings**, click **Change AppFlow Settings**.
3. In the **Configure AppFlow Settings** dialog box, specify relevant parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **OK**, and then click **Close**.

Example: Configuring AppFlow for DataStream

The following example illustrates the procedure for configuring AppFlow for DataStream using the command line interface.

```
> enable feature appflow
> add db user sa password freebsd
> add lbvserver lb0 MSSQL 10.102.147.97 1433 -
appflowLog ENABLED
> add service sv0 10.103.24.132 MSSQL 1433 -
appflowLog ENABLED
> bind lbvserver lb0 sv0
> add appflow collector col0 -IPAddress
10.102.147.90
> add appflow action act0 -collectors col0
> add appflow policy pol0
"mssql.req.query.text.contains(\"select\")" act0
> bind lbvserver lb0 -policyName pol0 -priority 10
```

When the Netscaler appliance receives a database request, the appliance evaluates the request against a configured policy. If a match is found, the details are sent to the AppFlow collector configured in the policy.

Exporting Performance Data of Web Pages to AppFlow Collector

The EdgeSight Monitoring application provides web page monitoring data with which you can monitor the performance of various Web applications served in a Netscaler environment. You can now export this data to AppFlow collectors to get an in-depth analysis of the web page applications. AppFlow, which is based on IPFIX standard, provides more specific information about web application performance than does EdgeSight monitoring alone.

You can configure both load balancing and content switching virtual servers to export EdgeSight Monitoring data to AppFlow collectors. Before configuring a virtual server for AppFlow export, associate an Appflow action with the EdgeSight Monitoring responder policy.

The following web page performance data is exported to AppFlow:

- ♦ **Page Load Time.** Elapsed time, in milliseconds, from when the browser starts to receive the first byte of a response until the user starts to interact with the page. At this stage, all the page content might not be loaded.
- ♦ **Page Render Time.** Elapsed time, in milliseconds, from when the browser receives the first byte of response until either all page content has been rendered or the page load action has timed out.
- ♦ **Time Spent on the Page.** Time spent by users on a page. Represents the period of time from one page request to the next one.

AppFlow transmits the performance data by using the Internet Protocol Flow Information eXport (IPFIX) format, which is an open Internet Engineering Task Force (IETF) standard defined in RFC 5101. The AppFlow templates use the following enterprise-specific Information Elements (IEs) to export the information:

- ♦ **Client Load End Time.** Time at which the browser received the last byte of a response to load all the objects of the page such as images, scripts, and stylesheets.
- ♦ **Client Load Start Time.** Time at which the browser receives the first byte of the response to load any objects of the page such as images, scripts, and stylesheets.
- ♦ **Client Render End Time.** Time at which browser finished rendering the entire page, including the embedded objects.
- ♦ **Client Render Start Time.** Time at which the browser started rendering the page.

Prerequisites for Exporting Performance Data of Web Pages to AppFlow Collectors

Before associating the AppFlow action with the AppFlow policy, verify that the following prerequisites have been met:

- ♦ The AppFlow feature has been enabled and configured. For instructions, see ["Configuring the AppFlow feature"](#).
- ♦ The Responder feature has been enabled.
- ♦ The EdgeSight Monitoring feature has been enabled. For instructions, see ["Enabling an Application for EdgeSight Monitoring."](#)
- ♦ EdgeSight Monitoring has been enabled on the load balancing or content switching virtual servers bound to the services of applications for which you want to collect the performance data. For instructions, see ["Enabling an Application for EdgeSight Monitoring."](#)

Associating an AppFlow Action with the EdgeSight Monitoring Responder Policy

To export the web page performance data to the AppFlow collector, you must associate an AppFlow action with the EdgeSight Monitoring responder policy. An AppFlow action specifies which set of collectors receive the traffic.

To associate an AppFlow action with the EdgeSight Monitoring Responder policy by using the command line interface

At the command prompt, type:

```
set responder policy <policyName> -appflowAction <action_Name>
```

Example

```
set responder policy pol -appflowAction actn
```

To associate an AppFlow action with the EdgeSight Monitoring Responder policy by using the configuration utility

1. Navigate to **AppExpert > Responder > Policies**.
2. In the details pane, select an EdgeSight Monitoring responder policy, and then click **Open**.

3. In the **Configure Responder Policy** dialog box, in the **AppFlow Action** drop-down list, select the AppFlow action associated with the collectors to which you want to send the web-page performance data.
4. Click **OK**.

Configuring a Virtual Server to Export EdgeSight Statistics to Appflow Collectors

To export EdgeSight statistics information from a virtual server to the AppFlow collector, you must associate an AppFlow action with the virtual server.

To associate an AppFlow action with a Load Balancing or Content Switching virtual server by using the configuration utility

1. Navigate to **Traffic Management > Load Balancing > Virtual Servers** or **Traffic Management > Content Switching > Virtual Servers**.
2. In the details pane, select a virtual server, or multiple virtual servers, and then click **Enable EdgeSight Monitoring**.
3. In the **Enable EdgeSight Monitoring** dialog box, select the **Export EdgeSight statistics to Appflow** check box.
4. From the **Appflow Action** drop-down list, select the AppFlow action. The AppFlow action defines the list of AppFlow collectors to which it exports EdgeSight Monitoring statistics. If you have selected multiple load balancing virtual servers, the same AppFlow Action will be associated with the responder policies bound to them.

You can later change the AppFlow Action configured for each of the selected Load Balancing virtual server individually, if required.

5. Click **OK**.

Chapter 4

AutoScale: Automatic Scaling in the Citrix CloudPlatform Environment

Topics:

- [How AutoScale Works](#)
- [Supported Environment](#)
- [Prerequisites](#)
- [NetScaler Configuration Details](#)
- [Troubleshooting](#)

Efficient hosting of applications in a cloud requires continuous optimization of application availability. To meet increasing demand, you have to scale network resources upward. When demand subsides, you need to scale down to avoid the unnecessary cost of idle resources. To minimize the cost of running the application by deploying only as many instances as are necessary during any given period of time, you have to constantly monitor traffic. However, monitoring traffic manually is not a feasible option. For the application environment to be able to scale up or down rapidly, you need to automate the processes of monitoring traffic and of scaling resources up and down whenever necessary.

If your organization uses Citrix CloudPlatform to deploy and manage the cloud environment, a Citrix NetScaler appliance can automatically scale users' applications as needed. The CloudPlatform elastic load balancing feature includes a feature called *AutoScale*. A CloudPlatform user can use the AutoScale feature to specify thresholds for various conditions for automatically scaling the application fleet upward and downward. The scale-up and scale-down conditions can vary from simple use cases, such as a server's CPU usage, to complex use cases, such as a combination of a server's CPU usage and responsiveness. CloudPlatform, in turn, configures the NetScaler appliance to load balance traffic to the application virtual machines (VMs), monitor application thresholds and performance, and trigger scale-up and scale-down actions to add or remove VMs to or from the application fleet.

The CloudPlatform user performs all AutoScale configuration tasks by using the CloudPlatform user interface or APIs. The CloudPlatform user:

1. Creates a load balancing rule, with the necessary load balancing algorithm and stickiness.
2. Configures AutoScale parameters by specifying the application instance template, the minimum number of

instances to maintain, the maximum number of instances permitted, scale-up and scale-down policies, and other information necessary for the functioning of the feature.

3. Submits the configuration.

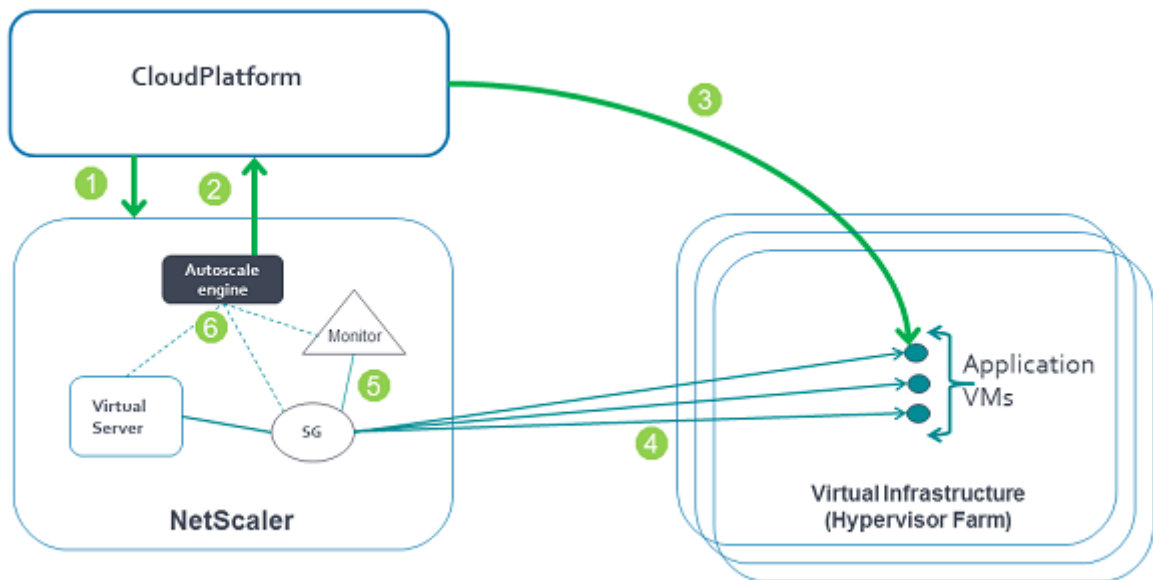
When the CloudPlatform user completes the AutoScale configuration, CloudPlatform uses the NetScaler NITRO API to push all the necessary configuration commands to the NetScaler appliance. As the NetScaler administrator, you do not have to perform any tasks for configuring AutoScale on the NetScaler appliance. However, you might have to be aware of certain prerequisites, and you might have to troubleshoot the configuration if issues arise in the AutoScale configuration. To troubleshoot the configuration, you have to be aware of how CloudPlatform works and what configuration CloudPlatform pushes to the NetScaler appliance. You also need a working knowledge of how to troubleshoot issues on a NetScaler appliance.

How AutoScale Works

When the CloudPlatform user completes the AutoScale configuration, CloudPlatform uses the NetScaler NITRO API to create an AutoScale-related configuration on the NetScaler appliance. For information about the configuration commands that CloudPlatform uses to configure the NetScaler appliance, see "[NetScaler Configuration Details](#)."

The following diagram shows the sequence of operations, beginning with CloudPlatform pushing the AutoScale configuration to the NetScaler appliance. The events are numbered in the order in which they occur, and are described below.

Figure 4-1. AutoScale Architecture



When the CloudPlatform user submits the AutoScale configuration, the following events occur:

1. CloudPlatform uses the NetScaler NITRO API to push the AutoScale configuration to the NetScaler appliance, creating AutoScale-related entities on the appliance. The entities include a load balancing virtual server, a service group, and monitors.
2. The AutoScale engine on the NetScaler appliance sends API requests to CloudPlatform to initially deploy the minimum number of virtual machines required.
3. CloudPlatform provisions the minimum number of instances (VMs) on the hypervisors (virtualization hosts) that it manages.
4. The NetScaler appliance discovers the IP addresses assigned by CloudPlatform to the newly created VMs and binds them, as services, to the service group

representing them. The NetScaler appliance can then load balance traffic to the VMs.

5. NetScaler monitors bound to the service group start monitoring the load by collecting SNMP metrics from the instances.
6. The AutoScale engine on the NetScaler appliance monitors the metrics collected from the VMs and triggers scale-up and scale-down events whenever the metrics breach the configured threshold for the specified period. As part of the scale-up trigger, the NetScaler AutoScale engine sends an API request to CloudPlatform to deploy a new VM. After the virtual machine is deployed, the AutoScale engine binds the service representing the VM (IP address and port) to the service group and, after the configured quiet time, starts forwarding load balanced traffic to the new virtual machine. Likewise, as part of the scale-down trigger, the NetScaler AutoScale engine selects a VM, stops forwarding new requests to that instance, and waits for the configured quiet time (to allow for the processing of current requests to complete) before it sends an API request to CloudPlatform to destroy the chosen instance.

In this way, the NetScaler appliance monitors the application and triggers scale-up and scale-down events on the basis of application load and/or performance.

Supported Environment

AutoScale is supported in the following environment:

- ♦ Citrix CloudPlatform 3.0.5.
- ♦ Citrix NetScaler MPX/SDX/virtual appliance running Citrix NetScaler release 10.e and later.
- ♦ SNMP v1/v2.

Prerequisites

Before you set up AutoScale, do the following:

- ♦ Make sure that CloudPlatform is reachable from the NetScaler appliance. You can do so by logging on to the NetScaler appliance and sending ping requests to the CloudPlatform server's IP address.
- ♦ Make sure that the network service offering used in CloudPlatform includes the NetScaler appliance as an external load balancing device.
- ♦ Use a CloudPlatform and NetScaler release that supports AutoScale. For information about NetScaler releases that support AutoScale, see "[Supported Environment](#)."

NetScaler Configuration Details

The following table describes the AutoScale configuration commands that are used by Citrix CloudPlatform to configure a NetScaler appliance.

Table 4-1. NetScaler Configuration for AutoScale

AutoScale configuration command(s)	Description
<pre>add lb vserver Cloud- VirtualServer-192.0.2.116-22 TCP 192.0.2.116 22 - persistenceType NONE -lbMethod ROUNDROBIN -cltTimeout 9000 - minAutoscaleMembers 2 - maxAutoscaleMembers 5</pre>	<p>Creates a load balancing virtual server to evenly distribute the load on the application instances (ROUND-ROBIN method). The virtual server also specifies the limits for the number of instances to which the application can scale up or down (maxAutoscaleMembers and minAutoscaleMembers, respectively).</p>
<pre>add serviceGroup Clouda35a6b6b76614006b97476e841 b80f79 TCP -maxClient 0 -maxReq 0 -cip DISABLED -usip NO - useproxyport YES -cltTimeout 9000 -svrTimeout 9000 -CKA NO - TCPB NO -CMP NO -autoScale POLICY -memberPort 22 bind lb vserver Cloud- VirtualServer-192.0.2.116-22 Clouda35a6b6b76614006b97476e841 b80f79</pre>	<p>Creates an AutoScale service group for the application instances, with the service group's autoScale parameter set to POLICY. Also specifies the port on which the service group members must receive traffic.</p> <p>The second command binds the service group to the load balancing virtual server.</p>
<pre>add server autoscale- internal_server_Clouda35a6b6b76 614006b97476e841b80f79 autoscale- internal_server_Clouda35a6b6b76 614006b97476e841b80f79 bind serviceGroup Clouda35a6b6b76614006b97476e841 b80f79 autoscale- internal_server_Clouda35a6b6b76 614006b97476e841b80f79 22</pre>	<p>Creates a server entry to represent the application instances.</p> <p>Binds the server entry to the service group.</p>
<pre>add lb metricTable Cloud- MTbl-192.0.2.116-22 bind lb metricTable Cloud- MTbl-192.0.2.116-22 Linux_User_CPU_-_percentage 1.3.6.1.4.1.2021.11.9.0 add lb monitor Cloud- Mon-192.0.2.116-22 LOAD -</pre>	<p>Configures a new SNMP monitor to retrieve the specified metrics.</p>

AutoScale configuration command(s)	Description
<pre>interval 24 -destPort 161 - snmpCommunity public - metricTable Cloud- MTbl-192.0.2.116-22 bind lb monitor Cloud- Mon-192.0.2.116-22 -metric Linux_User_CPU_-_percentage - metricThreshold 2147483647 bind serviceGroup Clouda35a6b6b76614006b97476e841 b80f79 -monitorName Cloud- Mon-192.0.2.116-22 -passive</pre>	
<pre>add autoscale profile Cloud- AutoScale- Profile-192.0.2.116-22 -type CLOUDSTACK -url "http:// 10.102.31.107:8080/client/api" -apiKey t0fEWPTk_ncQYbofjAm1jjlgGTR7UNZ rkZ3sdEpLREBNzBPLSNpNz8qNSbc439 xNtYnEYdWn_MsUC_CUazaIKg - sharedSecret - PrE5h3DP7swHAN12TGB1X- xSTRLHzob9116O0VO1FMxvE1UO17uoD 6_Z0bkkLaVtK5Y10oBkTzgbTwp3u5lC Q</pre>	<p>Creates an AutoScale profile to specify the details required by NetScaler for making API requests to CloudPlatform (URL, API key, and shared secret).</p>
<pre>add autoscale action Cloud- AutoScale- ScaleUpAction-192.0.2.116-22 - type SCALE_UP -profileName Cloud-AutoScale- Profile-192.0.2.116-22 - parameters "command=deployVirtualMachine&z oneid=2ab23590-78cb-4106-8d85-4 412a2f2435f&serviceofferingid=b 9503e47-0d8f-4c89- a88d-04d8b17fe8e9&templateid=1a 4a5084-208c-47a8-9c16- d582550cf759&displayname=AutoSc ale-LB-lb&networkids=a3c97129- b729-4c72-994f-7b918f20ce4d&lbr uleid=f96b7f3b-19ec-4123-891c-6 04f05b032b3" -quietTime 90 -</pre>	<p>Creates a scale-up action, which enables the NetScaler appliance to add virtual machines (instances) to the application fleet.</p>

AutoScale configuration command(s)	Description
vServer Cloud-VirtualServer-192.0.2.116-22	
add autoscale action Cloud-AutoScale-ScaleDownAction-192.0.2.116-22 -type SCALE_DOWN -profileName Cloud-AutoScale-Profile-192.0.2.116-22 -parameters "command=destroyVirtualMachine&lbruleid=f96b7f3b-19ec-4123-891c-604f05b032b3" -vmDestroyGracePeriod 30 -quietTime 90 -vServer Cloud-VirtualServer-192.0.2.116-22	Creates a scale-down action, which enables the NetScaler appliance to remove virtual machines (instances) from the application fleet.
add autoscale policy Cloud-AutoScale-Policy-Min-192.0.2.116-22 -rule "SYS.VSERVER(\"Cloud-VirtualServer-192.0.2.116-22\") .ACTIVESERVICES.LT(SYS.VSERVER(\"Cloud-VirtualServer-192.0.2.116-22\") .MINAUTOSCALEMEMBERS)" -action Cloud-AutoScale-ScaleUpAction-192.0.2.116-22	Creates an AutoScale policy to initially create the specified minimum number of VMs and, later, to ensure that the number of VMs in the fleet does not fall below the required minimum.
add autoscale policy Cloud-AutoScale-Policy-Max-192.0.2.116-22 -rule "SYS.VSERVER(\"Cloud-VirtualServer-192.0.2.116-22\") .ACTIVESERVICES.GT(SYS.VSERVER(\"Cloud-VirtualServer-192.0.2.116-22\") .MAXAUTOSCALEMEMBERS)" -action Cloud-AutoScale-ScaleDownAction-192.0.2.116-22	Creates an AutoScale policy to prevent the number of VMs in the fleet from exceeding the specified maximum.
add autoscale policy Cloud-AutoScale-Policy-192.0.2.116-22-35 -rule "SYS.VSERVER(\"Cloud-VirtualServer-192.0.2.116-22\") .ACTIVESERVICES.LT(SYS.VSERVER(\"Cloud-VirtualServer-192.0.2.116-22\")	Creates an AutoScale policy to evaluate the metrics that are collected and trigger a scale-up action when the metric value breaches the threshold specified for the scale-up policy.

AutoScale configuration command(s)	Description
<code>.MAXAUTOSCALEMEMBERS) && (SYS.VSERVER(\"Cloud-VirtualServer-192.0.2.116-22\") .SNMP_TABLE(0).AVERAGE_VALUE.G T(90))\" -action Cloud-AutoScale- ScaleUpAction-192.0.2.116-22</code>	
<code>add autoscale policy Cloud-AutoScale- Policy-192.0.2.116-22-36 -rule \"SYS.VSERVER(\"Cloud-VirtualServer-192.0.2.116-22\") .ACTIVESERVICES.GT(SYS.VSERVER(\"Cloud-VirtualServer-192.0.2.116-22\") .MINAUTOSCALEMEMBERS) && (SYS.VSERVER(\"Cloud-VirtualServer-192.0.2.116-22\") .SNMP_TABLE(0).AVERAGE_VALUE.L T(30))\" -action Cloud-AutoScale- ScaleDownAction-192.0.2.116-22</code>	Creates an AutoScale policy to evaluate the collected metrics and trigger a scale-down action when the metric value breaches the threshold specified by the scale-down policy.
<code>add ns timer Cloud-AutoScale-Timer-192.0.2.116-22 -interval 30 bind ns timer Cloud-AutoScale-Timer-192.0.2.116-22 - policyName Cloud-AutoScale-Policy-Min-192.0.2.116-22 - priority 1 - gotoPriorityExpression END - sampleSize 1 -threshold 1 bind ns timer Cloud-AutoScale-Timer-192.0.2.116-22 - policyName Cloud-AutoScale-Policy-Max-192.0.2.116-22 - priority 2 - gotoPriorityExpression END - sampleSize 1 -threshold 1 bind ns timer Cloud-AutoScale-Timer-192.0.2.116-22 - policyName Cloud-AutoScale-Policy-192.0.2.116-22-35 - priority 3 -</code>	Creates a timer that enables evaluation of the AutoScale policies at the configured sampling intervals.

AutoScale configuration command(s)	Description
<pre>gotoPriorityExpression END - sampleSize 2 -threshold 2 bind ns timer Cloud-AutoScale- Timer-192.0.2.116-22 - policyName Cloud-AutoScale- Policy-192.0.2.116-22-36 - priority 4 - gotoPriorityExpression END - sampleSize 2 -threshold 2</pre>	

Troubleshooting

Before you attempt to resolve an AutoScale issue, make sure that the prerequisites have been adhered to, on both the CloudPlatform server and the NetScaler appliance, as described in "[Prerequisites](#)." If that does not resolve the issue, your problem could be one of the following.

The AutoScale configuration was successfully configured in CloudPlatform. Yet, the minimum number of VMs has not been created.

- ♦ Recommend that the CloudPlatform user deploy one VM manually in the network before configuring AutoScale. Ask the user to remove the AutoScale configuration from the NetScaler appliance or the load balancer from the network, manually deploy one VM (preferably using the template created for the AutoScale configuration), and then create the AutoScale configuration.
- ♦ Verify that the CloudPlatform user has configured the VM template in such a way that the VMs that are created from the template can accept traffic without manual intervention. If a provisioned VM cannot accept traffic automatically, the metric remains above the threshold, and the AutoScale configuration continues to provision additional VMs, as designed. To remedy the issue, disable AutoScale from CloudPlatform, fix the template, and then enable AutoScale.
- ♦ Verify that the CloudPlatform user has not exceeded the limit for the number of VMs imposed by the user's account.
- ♦ Verify that the CloudPlatform server is up and is reachable from the NetScaler appliance.
- ♦ Verify that the CloudPlatform log file, management-server.log, has reported the successful creation of the AutoScale configuration in CloudPlatform.
- ♦ Verify that the scale-up policy that is responsible for initial scale up (the policy name is prefixed with `Cloud-AutoScale-Policy-Min`) is receiving hits.

The AutoScale configuration is rapidly spawning a large number of VMs

- ♦ Verify that the CloudPlatform user has configured the VM template in such a way that the VMs that are created from the template can accept traffic without manual intervention. If a provisioned VM cannot accept traffic automatically, the metric remains above the threshold, and the AutoScale configuration continues to provision additional VMs, as designed. To remedy the issue, disable AutoScale from CloudPlatform, fix the template, and then enable AutoScale.
- ♦ Verify that the quiet time that the CloudPlatform user has configured in the AutoScale configuration is sufficient to ensure even traffic distribution to all the VMs, including the new VM. If the quiet time is too low, and traffic distribution has not stabilized, the metrics might remain above the threshold, and additional VMs might be spawned.

When I ran the top command on my VM, I noticed that the CPU usage on my VM had breached the threshold that was configured for the scale-up action in AutoScale. Yet, the application is not scaling up.

- ♦ Verify that the CloudPlatform user has installed an SNMP agent in the VM template, and that the SNMP agent is up and running on every VM.
- ♦ Verify that the CloudPlatform user has not exceeded the limit for the number of VMs imposed by the user's account.
- ♦ Verify that the CloudPlatform user has correctly configured the SNMP parameters to collect metrics from the VM (for example, the community string and the port).
- ♦ Verify that the scale-up or scale-down policy is receiving hits.
- ♦ Verify that the CloudPlatform server is up, and that the CloudPlatform server is reachable from the NetScaler appliance.

One or more additional VMs have been created, but they are not accepting traffic (that is, VMs have been created, but the average value of the metrics is still above the threshold)

- ♦ Verify that the user has configured the templates in such a way that the VMs created from the templates can start serving traffic without any manual intervention.

- ♦ Verify that the service is running on the VMs, on the configured member port.
- ♦ Send a ping request to the gateway (virtual router), from the VM that is not accepting traffic.

The AutoScale configuration has been deleted, but the VMs continue to exist

- ♦ The VMs might not be deleted immediately after the AutoScale configuration is deleted. Wait for about 5 minutes after you have deleted the AutoScale configuration, and then check again.
- ♦ If the destruction of VMs has not commenced after 5 minutes, you might have to delete the VMs manually.

Chapter 5

EdgeSight Monitoring for NetScaler

Topics:

- [Configuring EdgeSight Monitoring for NetScaler](#)
- [Enabling an Application for EdgeSight Monitoring](#)
- [Accessing the EdgeSight Monitoring Interface from NetScaler](#)

Citrix EdgeSight for NetScaler is an application to monitor end-user experience with Web applications served in a NetScaler environment. The EdgeSight Monitoring application uses the HTML Injection feature of the NetScaler to provide data with which you can compare the performance of various Web applications across geographical locations.

For EdgeSight monitoring, register the NetScaler appliance with EdgeSight and enable applications in the NetScaler. The EdgeSight application processes the data and displays the information after aggregating it. You can view the data as charts, graphs, or tables.

When EdgeSight monitoring is enabled on a virtual server, NetScaler injects scripts into the responses sent to the clients. Execution or insertion of these scripts does not affect the response to the client. Data injected by the NetScaler is collected by the EdgeSight server through data collector services.

EdgeSight is an agentless application. The EdgeSight server from which you can monitor the user-experience is referred to as *EdgeSight UI server*.

Use the wizard to register the NetScaler appliance with the EdgeSight UI server, select the data collector services, and configure the rate at which data is injected into the response. For more information on the wizard, see "[Configuring EdgeSight Monitoring](#)."

Note: The EdgeSight UI server has a data collector and a Web site. For better scalability, you can install multiple data collectors.

To use EdgeSight monitoring, enable EdgeSight monitoring on the load balancing or content switching virtual servers associated with the selected applications. For instructions, see "[Enabling an Application for EdgeSight Monitoring](#)."

Configuring EdgeSight Monitoring for NetScaler

The configuration utility provides a wizard to assist you with configuration. The wizard for the configuration of EdgeSight monitoring for NetScaler guides you through the configuration procedure in simple steps. The wizard takes care of the following tasks:

- ♦ Enable the EdgeSight Monitoring (HTML Injection) feature.
- ♦ Specify the rate and frequency for injecting data.
- ♦ Bind the EdgeSight server/data collector services to the load balancing virtual server on the NetScaler.
- ♦ Enable the Web applications for EdgeSight monitoring.

Note: You can configure rate or frequency. If you specify the rate to be x, data is injected once after every x responses sent by the service. For example, if you specify the rate to be 500, NetScaler injects the data into the 1st response, 501st response, 1002nd response, and so on. If you specify frequency to be y, data is injected once in every y milliseconds.

To access the wizard from the NetScaler configuration utility and configure EdgeSight Monitoring

1. Navigate to **Sysyem > EdgeSight Monitoring**
2. Click **EdgeSight for NetScaler Wizard**.
3. Follow the instructions presented by the wizard.

To configure EdgeSight monitoring from the command line interface and configure EdgeSight Monitoring

The following example shows the CLI commands executed for configuring EdgeSight monitoring for a NetScaler appliance.

Example

When the EdgeSight for NetScaler overview page link is clicked

```
Jun 9 22:46:21 <local0.info> 10.102.113.114 06/10/2011:02:46:21 GMT
ns PPE-2 : UI CMD_EXECUTED 247 : User nsroot - Remote_ip
10.101.254.143 - Command "show filter action" - Status "Success"
```

Jun 9 22:46:33 <local0.info> 10.102.113.114 06/10/2011:02:46:33 GMT
 ns PPE-2 : UI CMD_EXECUTED 248 : User nsroot - Remote_ip
 10.101.254.143 - Command "show filter postbodyInjection" - Status
 "Success"

Jun 9 22:46:33 <local0.info> 10.102.113.114 06/10/2011:02:46:33 GMT
 ns PPE-2 : UI CMD_EXECUTED 249 : User nsroot - Remote_ip
 10.101.254.143 - Command "show filter htmlinjectionparameter" -
 Status "Success"

Jun 9 22:46:33 <local0.info> 10.102.113.114 06/10/2011:02:46:33 GMT
 ns PPE-2 : UI CMD_EXECUTED 250 : User nsroot - Remote_ip
 10.101.254.143 - Command "show filter htmlinjectionvariable
 EDGESIGHT_SERVER_IP" - Status "Success"

Jun 9 22:46:33 <local0.info> 10.102.113.114 06/10/2011:02:46:33 GMT
 ns PPE-2 : UI CMD_EXECUTED 251 : User nsroot - Remote_ip
 10.101.254.143 - Command "show lb vserver __ESNS_LBVSERVER" -
 Status "Success"

Jun 9 22:46:34 <local0.info> 10.102.113.114 06/10/2011:02:46:34 GMT
 ns PPE-2 : UI CMD_EXECUTED 252 : User nsroot - Remote_ip
 10.101.254.143 - Command "show lb vserver __ESNS_LBVSERVER" -
 Status "Success"

Jun 9 22:46:34 <local0.info> 10.102.113.114 06/10/2011:02:46:34 GMT
 ns PPE-2 : UI CMD_EXECUTED 253 : User nsroot - Remote_ip
 10.101.254.143 - Command "show service" - Status "Success"

When the Register Appliance link is clicked

Jun 9 22:47:52 <local0.info> 10.102.113.114 06/10/2011:02:47:52 GMT
 ns PPE-2 : UI CMD_EXECUTED 254 : User nsroot - Remote_ip
 10.101.254.143 - Command "show ns hostName" - Status "Success"

On completion of the wizard

Jun 9 22:50:01 <local0.info> 10.102.113.114 06/10/2011:02:50:01 GMT
 ns PPE-2 : UI CMD_EXECUTED 255 : User nsroot - Remote_ip
 10.101.254.143 - Command "set filter htmlinjectionvariable
 EDGESIGHT_SERVER_IP -value 10.102.31.147" - Status "Success"

Jun 9 22:50:01 <local0.info> 10.102.113.114 06/10/2011:02:50:01 GMT
 ns PPE-2 : UI CMD_EXECUTED 256 : User nsroot - Remote_ip
 10.101.254.143 - Command "set lb vserver __ESNS_LBVSERVER -
 IPAddress 0.0.0.0 -IPPattern 0.0.0.0 -IPMask * -weight 1 essvc -
 persistenceType NONE -timeout 2 -persistenceBackup NONE -
 backupPersistenceTimeout 2 -lbMethod LEASTCONNECTION -persistMask
 255.255.255.255 -v6persistmasklen 128 -pq OFF -sc OFF -rtspNat OFF -
 m IP -dataOffset 0 -sessionless DISABLED -connfailover DISABLED -
 cltTimeout 180 -soMethod NONE -soPersistence DISABLED -
 soPersistenceTimeOut 2 -redirectPortRewrite DISABLED -
 downStateFlush DISABLED -gt2GB DISABLED -insertVs" - Status "Success"

```
Jun 9 22:50:02 <local0.info> 10.102.113.114 06/10/2011:02:50:02 GMT
ns PPE-2 : UI CMD_EXECUTED 257 : User nsroot - Remote_ip
10.101.254.143 - Command "set filter prebodyInjection "/netscaler/
htmlinjection/ens/prebody.js"" - Status "Success"

Jun 9 22:50:02 <local0.info> 10.102.113.114 06/10/2011:02:50:02 GMT
ns PPE-2 : UI CMD_EXECUTED 258 : User nsroot - Remote_ip
10.101.254.143 - Command "set filter postbodyInjection "/netscaler/
htmlinjection/ens/postbody.js"" - Status "Success"
```

Enabling an Application for EdgeSight Monitoring

EdgeSight monitoring makes it possible to monitor the performance of an application from the end user's perspective.

Note: To enable an application to be monitored, EdgeSight monitoring must be enabled on the virtual servers.

EdgeSight monitoring is possible only if the response from the physical server is not compressed. Therefore, make sure that Compression is enabled globally on the NetScaler. When you use the wizard to configure EdgeSight monitoring, the wizard takes care of enabling compression and other necessary processing.

Note: Before enabling EdgeSight monitoring on a virtual server, make sure that you completed the configuration of EdgeSight monitoring through the wizard.

To enable EdgeSight monitoring on a load balancing or content switching virtual server by using the NetScaler configuration utility

1. Navigate to **Traffic Management > Load Balancing > Virtual Servers**.
2. Select the virtual server, or multiple virtual servers, and then click **Enable EdgeSight Monitoring**.
3. If you want the NetScaler to compress responses after receiving them from the server, select the **Compress response at NetScaler before sending it to the client** check box.
4. Select the **Export EdgeSight statistics to AppFlow** check box.
5. From the **Appflow Action** drop-down list, select the AppFlow action.

The AppFlow action defines the list of AppFlow collectors to which it exports EdgeSight Monitoring statistics. If you have selected multiple load balancing virtual

servers, the same AppFlow action will be configured in the responder policies bound to them.

Note: You can later change the AppFlow action configured for each of the selected Load Balancing virtual server individually, if required.

6. Click **OK**.

To enable EdgeSight monitoring on a load balancing or content switching virtual server by using the command line interface

The following example shows the CLI commands executed for enabling EdgeSight monitoring on a virtual server.

Example

```
Jun 9 22:56:13 <local0.info> 10.102.113.114 06/10/2011:02:56:13 GMT
ns PPE-2 : UI CMD_EXECUTED 266 : User nsroot - Remote_ip
10.101.254.143 - Command "show lb vserver guilb" - Status "Success"

Jun 9 22:56:13 <local0.info> 10.102.113.114 06/10/2011:02:56:13 GMT
ns PPE-2 : UI CMD_EXECUTED 267 : User nsroot - Remote_ip
10.101.254.143 - Command "show lb vserver guilb" - Status "Success"

Jun 9 22:56:14 <local0.info> 10.102.113.114 06/10/2011:02:56:14 GMT
ns PPE-2 : UI CMD_EXECUTED 268 : User nsroot - Remote_ip
10.101.254.143 - Command "show lb vserver guilb" - Status "Success"

Jun 9 22:56:14 <local0.info> 10.102.113.114 06/10/2011:02:56:14 GMT
ns PPE-2 : UI CMD_EXECUTED 269 : User nsroot - Remote_ip
10.101.254.143 - Command "show lb vserver guilb" - Status "Success"

Jun 9 22:56:14 <local0.info> 10.102.113.114 06/10/2011:02:56:14 GMT
ns PPE-2 : UI CMD_EXECUTED 270 : User nsroot - Remote_ip
10.101.254.143 - Command "bind lb vserver guilb -weight 1 -
policyName __ESNS_PREBODY_POLICY" - Status "Success"

Jun 9 22:56:15 <local0.info> 10.102.113.114 06/10/2011:02:56:15 GMT
ns PPE-2 : UI CMD_EXECUTED 271 : User nsroot - Remote_ip
10.101.254.143 - Command "bind lb vserver guilb -weight 1 -
policyName __ESNS_POSTBODY_POLICY" - Status "Success"

Jun 9 22:56:15 <local0.info> 10.102.113.114 06/10/2011:02:56:15 GMT
ns PPE-2 : UI CMD_EXECUTED 272 : User nsroot - Remote_ip
10.101.254.143 - Command "bind lb vserver guilb -weight 1 -
policyName __ESNS_RESPONDER_POLICY -priority 2147483647 -
gotoPriorityExpression END" - Status "Success"
```

```
Jun 9 22:56:15 <local0.info> 10.102.113.114 06/10/2011:02:56:15 GMT
ns PPE-2 : UI CMD_EXECUTED 273 : User nsroot - Remote_ip
10.101.254.143 - Command "bind lb vserver guilb -weight 1 -
policyName __ESNS_REWRITE_POLICY -priority 2147483647 -
gotoPriorityExpression END -type REQUEST" - Status "Success"

Jun 9 22:56:15 <local0.info> 10.102.113.114 06/10/2011:02:56:15 GMT
ns PPE-2 : UI CMD_EXECUTED 274 : User nsroot - Remote_ip
10.101.254.143 - Command "show ns feature" - Status "Success"
```

Accessing the EdgeSight Monitoring Interface from NetScaler

You can view the data presented by the EdgeSight monitoring application from the NetScaler appliance. The configuration utility of the NetScaler displays a login screen to access the EdgeSight UI server and upon successful login, displays the data.

To access EdgeSight for NetScaler by using the NetScaler configuration utility

1. Navigate to **System > EdgeSight Monitoring**.
2. Click **Access EdgeSight for NetScaler**.
3. Enter the credentials for accessing the EdgeSight UI server.

Chapter 6

High Availability

Topics:

- [*Considerations for a High Availability Setup*](#)
- [*Configuring High Availability*](#)
- [*Configuring the Communication Intervals*](#)
- [*Configuring Synchronization*](#)
- [*Synchronizing Configuration Files in a High Availability Setup*](#)
- [*Configuring Command Propagation*](#)
- [*Configuring Fail-Safe Mode*](#)
- [*Configuring Virtual MAC Addresses*](#)
- [*Configuring High Availability Nodes in Different Subnets*](#)
- [*Configuring Route Monitors*](#)
- [*Limiting Failovers Caused by Route Monitors in non-INC mode*](#)
- [*Configuring FIS*](#)
- [*Understanding the Causes of Failover*](#)
- [*Forcing a Node to Fail Over*](#)
- [*Forcing the Secondary Node to Stay Secondary*](#)
- [*Forcing the Primary Node to Stay Primary*](#)

A high availability (HA) deployment of two Citrix® NetScaler® appliances can provide uninterrupted operation in any transaction. With one appliance configured as the primary node and the other as the secondary node, the primary node accepts connections and manages servers while the secondary node monitors the primary. If, for any reason, the primary node is unable to accept connections, the secondary node takes over.

The secondary node monitors the primary by sending periodic messages (often called heartbeat messages or health checks) to determine whether the primary node is accepting connections. If a health check fails, the secondary node retries the connection for a specified period, after which it determines that the primary node is not functioning normally. The secondary node then takes over for the primary (a process called failover).

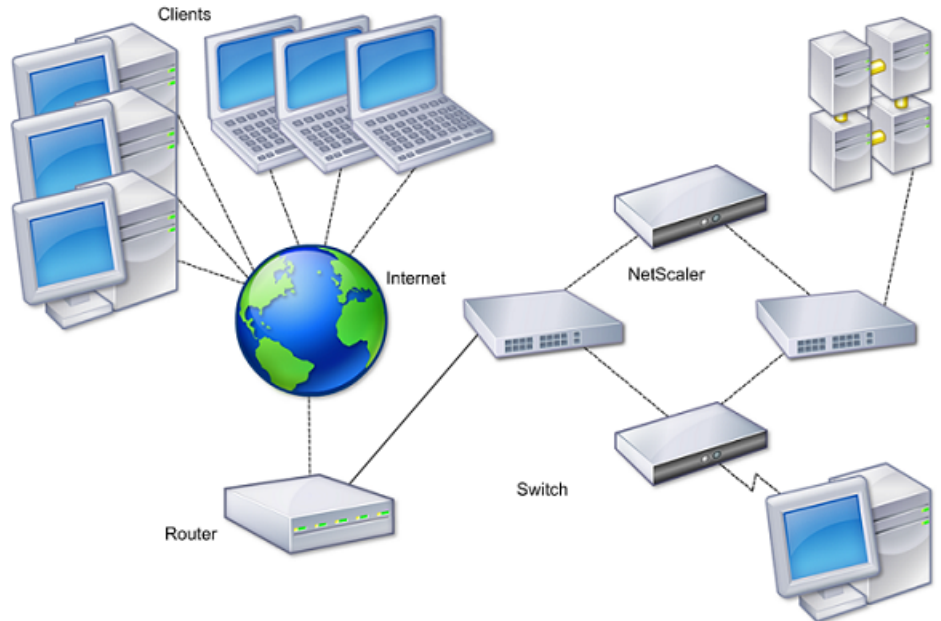
After a failover, all clients must reestablish their connections to the managed servers, but the session persistence rules are maintained as they were before the failover.

With Web server logging persistence enabled, no log data is lost due to the failover. For logging persistence to be enabled, the log server configuration must carry entries for both systems in the log.conf file.

The following figure shows a network configuration with an HA pair.

- *Understanding the High Availability Health Check Computation*
- *High Availability*
- *Troubleshooting High Availability Issues*

Figure 6-1. NetScaler Appliances in a High Availability Configuration



To configure HA, you might want to begin by creating a basic setup, with both nodes in the same subnet. You can then customize the intervals at which the nodes communicate health-check information, the process by which nodes maintain synchronization, and the propagation of commands from the primary to the secondary. You can configure fail-safe mode to prevent a situation in which neither node is primary. If your environment includes devices that do not accept NetScaler gratuitous ARP messages, you should configure virtual MAC addresses. When you are ready for a more complex configuration, you can configure HA nodes in different subnets.

To improve the reliability of your HA setup, you can configure route monitors and create redundant links. In some situations, such as when troubleshooting or performing maintenance tasks, you might want to force a node to fail over (assign primary status to the other node), or you might want to force the secondary node to stay secondary or the primary node to stay primary.

Considerations for a High Availability Setup

Note the following requirements for configuring systems in an HA setup:

- ♦ In an HA configuration, the primary and secondary NetScaler appliances should be of the same model. Different NetScaler models are not supported in an HA pair (for example, you cannot configure a 10010 model and a 7000 model as an HA pair).
- ♦ In an HA setup, both nodes must run the same version of NetScaler, for example, nCore/nCore or classic/classic. If the nodes are running NetScaler classic and you want to migrate to NetScaler nCore of the same NetScaler release, prop and sync are not supported during the migration process. Once migration is complete, prop and sync are auto-enabled. The same applies if you migrate from NetScaler nCore to NetScaler classic.
- ♦ Entries in the configuration file (ns.conf) on both the primary and the secondary system must match, with the following exceptions:
 - The primary and the secondary systems must each be configured with their own unique NetScaler IP addresses (NSIPs.)
 - In an HA pair, the node ID and associated IP address of one node must point to the other node. For example, if you have nodes NS1 and NS2, you must configure NS1 with a unique node ID and the IP address of NS2, and you must configure NS2 with a unique node ID and the IP address of NS1.
- ♦ If you create a configuration file on either node by using a method that does not go directly through the GUI or the CLI (for example, importing SSL certificates, or changing to startup scripts), you must copy the configuration file to the other node or create an identical file on that node.
- ♦ Initially, all NetScaler appliances are configured with the same RPC node password. RPC nodes are internal system entities used for system-to-system communication of configuration and session information. For security, you should change the default RPC node passwords.

One RPC node exists on each NetScaler. This node stores the password, which is checked against the password provided by the contacting system. To communicate with other systems, each NetScaler requires knowledge of those systems, including how to authenticate on those systems. RPC nodes maintain this information, which includes the IP addresses of the other systems, and the passwords they require for authentication.

RPC nodes are implicitly created when adding a node or adding a Global Server Load Balancing (GSLB) site. You cannot create or delete RPC nodes manually.

Note: If the NetScaler appliances in a high availability setup are configured in one-arm mode, you must disable all system interfaces except the one connected to the switch or hub.

- ♦ For an IPv6 HA configuration, the following considerations apply:

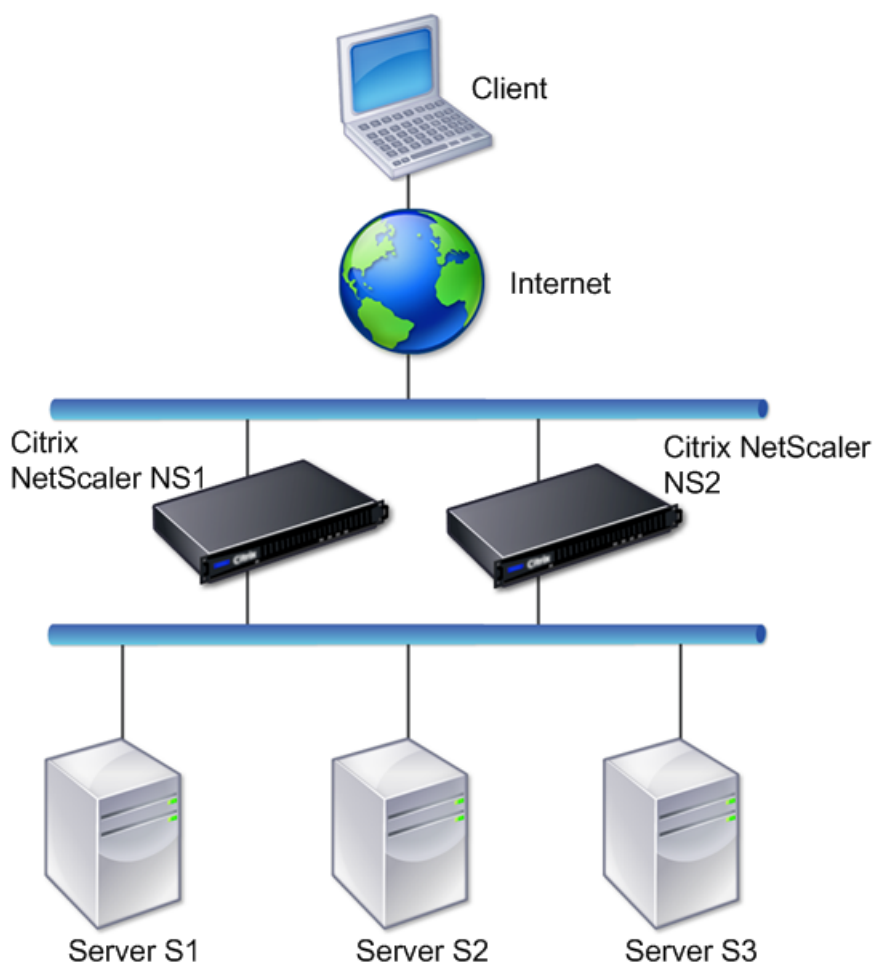
- You must install the IPv6PT license on both NetScaler appliances.
- After installing the IPv6PT license, enable the IPv6 feature by using the configuration utility or the command line interface.
- Both NetScaler appliances require a global NSIP IPv6 address. In addition, network entities (for example, switches and routers) between the two nodes must support IPv6.

Configuring High Availability

To set up a high availability configuration, you create two nodes, each of which defines the other's NetScaler IP (NSIP) address as a remote node. Begin by logging on to one of the two NetScaler appliances that you want to configure for high availability, and add a node. Specify the other appliance's NetScaler IP (NSIP) address as the address of the new node. Then, log on to the other appliance and add a node that has the NSIP address of the first appliance. An algorithm determines which node becomes primary and which becomes secondary.

Note: The configuration utility provides an option that avoids having to log on to the second appliance.

The following figure shows a simple HA setup, in which both nodes are in same subnet.

Figure 6-2. Two NetScaler Appliances Connected in a High Availability Configuration

Adding a Remote Node

To add a remote NetScaler appliance as a node in a high availability setup, you specify a unique node ID and the appliance's NSIP. The maximum number of node IDs in an HA setup is 64. When you add an HA node, you must disable the HA monitor for each interface that is not connected or not being used for traffic. For CLI users, this is a separate procedure.

Note: To ensure that each node in the high availability configuration has the same settings, you should synchronize your SSL certificates, startup scripts, and other configuration files with those on the primary node.

To add a node by using the command line interface

At the command prompt, type:

- ♦ `add ha node <id> <IPAddress>`

- ♦ **show ha node**

Example

```
> add ha node 3 1000:0000:0000:0000:0005:0600:700a:888b
```

To disable an HA monitor by using the command line interface

At the command prompt, type:

- ♦ **set interface** <ifNum> [-haMonitor (ON | OFF)]
- ♦ **show interface** <ifNum>

Example

```
> set interface 1/3 -haMonitor OFF
Done
```

To add a remote node by using the configuration utility

1. Navigate to **System > High Availability**.
2. In the details pane, select the **Nodes** tab, and then click **Add**.
3. In the **High Availability Setup** dialog box, in the **Remote Node IP Address** text box, type the NSIP address of the NetScaler that is to be added as the remote node. If the NSIP is an IPv6 address, select the **IPv6** check box before entering the address.
4. If you want to add the local node to the remote node automatically, select the **Configure remote system to participate in High Availability setup** check box. If you do not select this option, you will have to log in to the appliance represented by the remote node and add the node that you are currently configuring.
5. Make sure that the **Turn off HA monitor on interfaces/channels that are down** check box is selected.
6. Click **OK**. The **Nodes** page displays both of the nodes in your HA configuration (the local node and the remote node).

Disabling or Enabling a Node

You can disable or enable only a secondary node. When you disable a secondary node, it stops sending heartbeat messages to the primary node, and therefore the primary node can no longer check the status of the secondary. When you enable a node, the node takes part in the high availability configuration.

To disable or enable a node by using the command line interface

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

- ♦ `set ha node -hastatus DISABLED`
- ♦ `set ha node -hastatus ENABLED`

To disable or enable a node by using the configuration utility

1. Navigate to **System > High Availability**.
2. In the details pane, on the **Nodes** tab, select the local node, and then click **Open**.
3. In the **Configure Node** dialog box, under **High Availability Status**, do one of the following:
 - To enable the node, select the **DISABLED (Do not participate in HA)** check box.
 - To enable the node, select the **ENABLED (Do not participate in HA)** check box.
4. Click **OK**. A message appears in the status bar, stating that the node has been configured successfully.

Removing a Node

If you remove a node, the nodes are no longer in high availability configuration.

To remove a node by using the command line interface

At the command prompt, type:

`rm ha node <id>`

Example

```
> rm ha node 2
Done
```

To remove a node by using the configuration utility

1. In the navigation pane, expand **System**, and then click **High Availability**.
2. On the **High Availability** page, select the **Nodes** tab.
3. On the **Nodes** page, select the node that you want to remove, and click **Remove**.
4. On the **Remove** dialog box, click **Yes**.

Note: You can use the Network Visualizer to view the NetScaler appliances that are configured as a high availability (HA) pair and perform high availability configuration tasks. For more information, see ["Using the Network Visualizer."](#)

Configuring the Communication Intervals

The hello interval is the interval at which the heartbeat messages are sent to the peer node. The dead interval is the time interval after which the peer node is marked DOWN if heartbeat packets are not received. The heartbeat messages are UDP packets sent to port 3003 of the other node in an HA pair.

To set the hello and dead intervals by using the command line interface

At the command prompt, type:

- ♦ `set HA node [-helloInterval <msecs>] [-deadInterval <secs>]`
- ♦ `show HA node <id>`

To set the hello and dead intervals by using the configuration utility

1. Navigate to **System > High Availability**.
2. In the details pane, on the **Nodes** tab, select the local node, and then click **Open**.
3. In the **Configure Node** dialog box, under **Intervals**, set the following parameters.
 - Hello Interval (msecs)
 - Dead Interval (secs)
4. Click **OK**.

Configuring Synchronization

Synchronization is a process of duplicating the configuration of the primary node on the secondary node. The purpose of synchronization is to ensure that there is no loss of configuration information between the primary and the secondary nodes, regardless of the number of failovers that occur. Synchronization uses port 3010.

Synchronization is triggered by either of the following circumstances:

- ♦ The secondary node in an HA setup comes up after a restart.
- ♦ The primary node becomes secondary after a failover.

Automatic synchronization is enabled by default. You can also force synchronization.

Disabling or Enabling Synchronization

Automatic HA synchronization is enabled by default on each node in an HA pair. You can enable or disable it on either node.

To disable or enable automatic synchronization by using the command line interface

At the command prompt, type:

- ♦ `set HA node -haSync DISABLED`
- ♦ `set HA node -haSync ENABLED`

To disable or enable synchronization by using the configuration utility

1. Navigate to **System > High Availability**.
2. In the details pane, on the **Nodes** tab, select the local node, and then click **Open**.
3. In the **Configure Node** dialog box, under **HA Synchronization**, do one of the following:
 - To disable HA synchronization, clear the **Secondary node will fetch the configuration from Primary** check box.
 - To enable HA synchronization, select the **Secondary node will fetch the configuration from Primary** check box.
4. Click **OK**. A message appears in the status bar, stating that the node has been configured successfully.

Forcing the Secondary Node to Synchronize with the Primary Node

In addition to automatic synchronization, the NetScaler supports forced synchronization. You can force the synchronization from either the primary or the secondary node. When you force synchronization from the secondary node, it starts synchronizing its configuration with the primary node.

However, if synchronization is already in progress, forced synchronization fails and the system displays a warning. Forced synchronization also fails in any of the following circumstances:

- ♦ You force synchronization on a standalone system.
- ♦ The secondary node is disabled.
- ♦ HA synchronization is disabled on the secondary node.

To force synchronization by using the command line interface

At the command prompt, type:

force HA sync

To force synchronization by using the configuration utility

1. In the navigation pane, expand **System**, and then click **High Availability**.
2. In the details pane, on the **Nodes** tab, click **Force Synchronization**.

Synchronizing Configuration Files in a High Availability Setup

In a high availability setup, you can synchronize various configuration files from the primary node to the secondary node.

To perform the synchronization, you can use the command line interface or the configuration utility at either the primary or the secondary node. Files located on the secondary that are specific to the secondary (not present on the primary) are not deleted during the synchronization.

To synchronize files in a high availability setup by using the command line interface

At the command prompt, type:

sync HA files <mode>

Example

```
> sync HA files all
Done
```

To synchronize files in a high availability setup by using the configuration utility

1. Navigate to **System > Diagnostics**.
2. In the details pane, under **Utilities**, click **Start file synchronization**.
3. In the **Start file synchronization** dialog box, in the **Mode** drop-down list, select the appropriate type of synchronization (for example, **Everything except licenses and rc.conf**), and then click **OK**.

Configuring Command Propagation

In an HA setup, any command issued on the primary node propagates automatically to, and is executed on, the secondary before it is executed on the primary. If command propagation fails, or if command execution fails on the secondary, the primary node executes the command and logs an error. Command propagation uses port 3010.

In an HA pair configuration, command propagation is enabled by default on both the primary and secondary nodes. You can enable or disable command propagation on either node in an HA pair. If you disable command propagation on the primary node, commands are not propagated to the secondary node. If you disable command propagation on the secondary node, commands propagated from the primary are not executed on the secondary node.

Note: After reenabling propagation, remember to force synchronization.

If synchronization occurs while you are disabling propagation, any configuration-related changes that you make before the disabling of propagation takes effect are synchronized with the secondary node. This is also true for cases where propagation is disabled while synchronization is in progress.

To disable or enable command propagation by using the command line interface

At the command prompt, type:

- ♦ `set HA node -haProp DISABLED`
- ♦ `set HA node -haProp ENABLED`

To disable or enable command propagation by using the configuration utility

1. Navigate to **System > High Availability**.
2. In the details pane, on the **Nodes** tab, select the local node, and then click **Open**.
3. In the **Configure Node** dialog box, under **HA Propagation**, do one of the following:
 - To disable HA Propagation, clear the **Primary node will propagate configuration to the Secondary** check box.
 - To enable HA Propagation, select the **Primary node will propagate configuration to the Secondary** check box.
4. Click **OK**. A message appears in the status bar, stating that the node has been configured successfully.

Configuring Fail-Safe Mode

In an HA configuration, fail-safe mode ensures that one node is always primary when both nodes fail the health check. This is to ensure that when a node is only partially available, backup methods are enabled to handle traffic as best as possible. The HA fail-safe mode is configured independently on each node.

The following table shows some of the fail-safe cases. The NOT_UP state means that the node failed the health check yet it is partially available. The UP state means that the node passed the health check.

Table 6-1. Fail-Safe Mode Cases

Node A (Primary) Health State	Node B (Secondary) Health State	Default HA Behavior	Fail-Safe Enabled HA Behavior	Description
NOT_UP(failed last)	NOT_UP (failed first)	A (Secondary), B (Secondary)	A (Primary), B (Secondary)	If both nodes fail, one after the other, the node that was the last primary remains primary.
NOT_UP (failed first)	NOT_UP(failed last)	A (Secondary), B (Secondary)	A(Secondary), B(Primary)	If both nodes fail, one after the other, the node that was the last primary remains primary.
UP	UP	A (Primary), B (Secondary)	A (Primary), B (Secondary)	If both nodes pass the health check, no change in behavior with fail-safe enabled.
UP	NOT_UP	A(Primary), B(Secondary)	A (Primary), B (Secondary)	If only the secondary node fails, no change in behavior with fail-safe enabled.
NOT_UP	UP	A(Secondary), B(Primary)	A(Secondary), B(Primary)	If only the primary fails, no change in behavior with fail-safe enabled.

Node A (Primary) Health State	Node B (Secondary) Health State	Default HA Behavior	Fail-Safe Enabled HA Behavior	Description
NOT_UP	UP (STAYSECONDARY)	A (Secondary), B (Secondary)	A (Primary), B (Secondary)	If the secondary is configured as STAYSECONDARY, the primary remains primary even if it fails.

To enable fail-safe mode by using the command line interface

At the command prompt, type:

```
set HA node [-failSafe ( ON | OFF )]
```

Example

```
set ha node -failsafe ON
```

To enable fail-safe mode by using the configuration utility

1. In the navigation pane, expand **System**, and then click **High Availability**.
2. In the details pane, on the **Nodes** tab, select the local node, and then click **Open**.
3. In the **Configure Node** dialog box, under **Fail-Safe Mode**, select the **Maintain one Primary node even when both nodes are unhealthy** check box.
4. Click **OK**. A message appears in the status bar, stating that the node has been configured successfully.

Configuring Virtual MAC Addresses

A Virtual MAC address (VMAC) is a floating entity shared by the primary and the secondary nodes in an HA setup.

In an HA setup, the primary node owns all of the floating IP addresses, such as the MIPs, SNIPs, and VIPs. The primary node responds to Address Resolution Protocol (ARP) requests for these IP addresses with its own MAC address. As a result, the ARP table of an external device (for example, an upstream router) is updated with the floating IP address and the primary node's MAC address.

When a failover occurs, the secondary node takes over as the new primary node. It then uses Gratuitous ARP (GARP) to advertise the floating IP addresses that it acquired from the primary. However, the MAC address that the new primary advertises is the MAC address of its own interface.

Some devices (notably a few routers) do not accept the GARP messages generated by the NetScaler appliance. As a result, some external devices retain the old IP to MAC mapping advertised by the old primary node. This can result in a site going down.

You can overcome this problem by configuring a VMAC on both nodes of an HA pair. Both nodes then possess identical MAC addresses. Therefore, when failover occurs, the MAC address of the secondary node remains unchanged, and the ARP tables on the external devices do not need to be updated.

To create a VMAC, you need to first create a Virtual Router ID (VRID) and bind it to an interface. (In an HA setup, you need to bind the VRID to the interfaces on both nodes.) Once the VRID is bound to an interface, the system generates a VMAC with the VRID as the last octet.

Configuring IPv4 VMACs

When you create a IPv4 VMAC address and bind it to a interface, any IPv4 packet sent from the interface uses the VMAC address that is bound to the interface. If there is no IPv4 VMAC bound to an interface, the interface's physical MAC address is used.

The generic VMAC is of the form 00:00:5e:00:01:<VRID>. For example, if you create a VRID with a value of 60 and bind it to an interface, the resulting VMAC is 00:00:5e:00:01:3c, where 3c is the hex representation of the VRID. You can create 255 VRIDs with values from 1 to 255.

Creating or Modifying an IPv4 VMAC

You create an IPv4 virtual MAC by assigning it a virtual router ID. You can then you bind the VMAC to an interface. You cannot bind multiple VRIDs to the same interface. To verify the VMAC configuration, you should display and examine the VMACs and the interfaces bound to the VMACs.

To add a VMAC by using the command line interface

At the command prompt, type:

- ◆ **add vrID** <id>
- ◆ **bind vrid** <id> -ifnum <interface_name>
- ◆ **show vrid**

Example

```
> add vrID 100
Done
```

```
> bind vrid 100 -ifnum 1/1 1/2 1/3
Done
```

To unbind interfaces from a VMAC by using the command line interface

At the command prompt, type:

- ♦ **unbind vrid** <id> -ifnum <interface_name>
- ♦ **show vrid**

To configure a VMAC by using the configuration utility

1. Navigate to **System > Network > VMAC**.
2. In the details pane, on the **VMAC** tab, do one of the following:
 - To create a new VMAC, click **Add**.
 - To modify an existing VMAC, click **Open**.
3. In the **Create VMAC** or **Configure VMAC** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Under **Associate Interfaces**, do one of the following:
 - To bind interfaces to the VMAC, select the desired interfaces from the **Available Interfaces** table, and click **Add**.
 - To unbind interfaces from the VMAC, select the desired interfaces from the **Configured Interfaces** table, and click **Remove**.
5. Click **OK**.

Removing an IPv4 VMAC

To remove an IPv4 virtual MAC, you delete its virtual router ID.

To remove an IPv4 VMAC by using the command line interface

At the command prompt, type:

```
rm vrid <id>
```

Example

```
rm vrid 100s
```

To remove an IPv4 VMAC by using the configuration utility

1. Navigate to **System > Network > VMAC**.

2. In the details pane, on the **VMAC** tab, select the virtual router ID that you want to remove, and then click **Remove**. A message appears in the status bar, stating that the VMAC has been successfully removed.

Configuring IPv6 VMAC6s

The NetScaler supports VMAC6 for IPv6 packets. You can bind any interface to a VMAC6, even if an IPv4 VMAC is bound to the interface. Any IPv6 packet sent from the interface uses the VMAC6 bound to that interface. If there is no VMAC6 bound to an interface, an IPv6 packet uses the physical MAC.

Creating or Modifying a VMAC6

You create an IPv6 virtual MAC by assigning it an IPv6 virtual router ID. You can then you bind the VMAC to an interface. You cannot bind multiple IPv6 VRIDs to an interface. To verify the VMAC6 configuration, you should display and examine the VMAC6s and the interfaces bound to the VMAC6s.

To add a VMAC6 by using the command line interface

At the command prompt, type:

- ♦ **add vrID6** <id>
- ♦ **bind vrID6** <id> -ifnum <interface_name>
- ♦ **show vrID6**

Example

```
> add vrID6 100
Done
> bind vrID6 100 -ifnum 1/1 1/2 1/3
Done
```

To unbind interfaces from a VMAC6 by using the command line interface

At the command prompt, type:

- ♦ **unbind vrID6** <id> -ifnum <interface_name>
- ♦ **show vrID6**

To configure a VMAC6 by using the configuration utility

1. Navigate to **System > Network > VMAC**.
2. In the details pane, on the **VMAC6** tab, do one of the following:
 - To create a new VMAC6, click **Add**.

- To modify an existing VMAC6, click **Open**.
- 3. In the **Create VMAC6** or **Configure VMAC6** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
- 4. Under **Associate Interfaces**, do one of the following:
 - To bind interfaces to the VMAC6, select the desired interfaces from the **Available Interfaces** table, and click **Add**.
 - To unbind interfaces from the VMAC6, select the desired interfaces from the **Configured Interfaces** table, and click **Remove**.
- 5. Click **OK**.

Removing a VMAC6

To remove an IPv4 virtual MAC, you delete its virtual router ID.

To remove a VMAC6 by using the command line interface

At the command prompt, type:

```
rm vrid6 <id>
```

Example

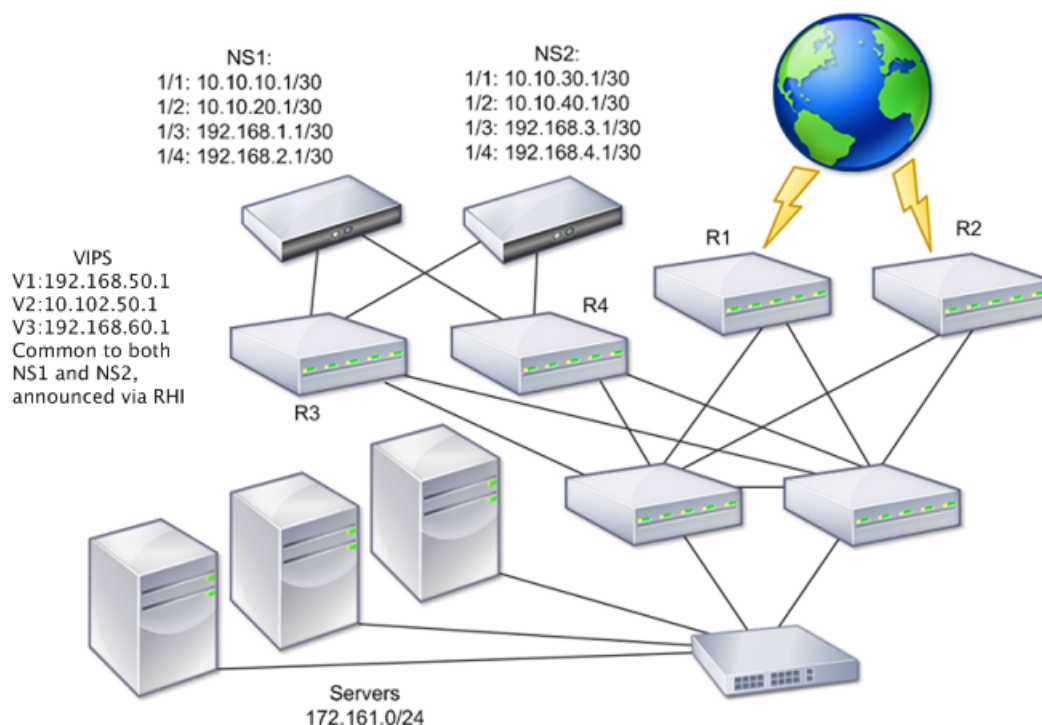
```
rm vrid6 100s
```

To remove a VMAC6 by using the configuration utility

1. Navigate to **System > Network > VMAC**.
2. In the details pane, on the **VMAC6** tab, select the virtual router ID that you want to remove, and then click **Remove**. A message appears in the status bar, stating that the VMAC6 has been successfully removed.

Configuring High Availability Nodes in Different Subnets

The following figure shows an HA deployment with the two systems located in different subnets:

Figure 6-3. High Availability over a Routed Network

In the figure, the systems NS1 and NS2 are connected to two separate routers, R3 and R4, on two different subnets. The NetScaler appliances exchange heartbeat packets through the routers. This configuration could be expanded to accommodate deployments involving any number of interfaces.

Note: If you use static routing on your network, you must add static routes between all the systems to ensure that heartbeat packets are sent and received successfully. (If you use dynamic routing on your systems, static routes are unnecessary.)

If the nodes in an HA pair reside on two separate networks, the primary and secondary node must have independent network configurations. This means that nodes on different networks cannot share entities such as MIPs, SNIPs, VLANs, and routes. This type of configuration, where the nodes in an HA pair have different configurable parameters, is known as Independent Network Configuration (INC) or Symmetric Network Configuration (SNC).

The following table summarizes the configurable entities and options for an INC, and shows how they must be set on each node.

Table 6-2. Behavior of NetScaler Entities and Options in an Independent Network Configuration

NetScaler entities	Options
IPs (NSIP/MIP/SNIPs)	Node-specific. Active only on that node.

NetScaler entities	Options
VIPs	Floating.
VLANs	Node-specific. Active only on that node.
Routes	Node-specific. Active only on that node. Link load balancing routes are floating.
ACLs	Floating (Common). Active on both nodes.
Dynamic routing	Node-specific. Active only on that node. The secondary node should also run the routing protocols and peer with upstream routers.
L2 mode	Floating (Common). Active on both nodes.
L3 mode	Floating (Common). Active on both nodes.
Reverse NAT (RNAT)	Node-specific. RNAT with VIP, because NATIP is floating.

As in configuring HA nodes in the same subnet, to configure HA nodes in different subnets, you log on to each of the two NetScaler appliances and add a remote node representing the other appliance.

Adding a Remote Node

When two nodes of an HA pair reside on different subnets, each node must have a different network configuration. Therefore, to configure two independent systems to function as an HA pair, you must specify INC mode during the configuration process.

When you add an HA node, you must disable the HA monitor for each interface that is not connected or not being used for traffic. For CLI users, this is a separate procedure.

To add a node by using the command line interface

At the command prompt, type:

- ♦ **add ha node <id> <IPAddress> -inc ENABLED**
- ♦ **show ha node**

Example

```
> add ha node 3 10.102.29.170 -inc ENABLED
Done
> add ha node 3 1000:0000:0000:0000:0005:0600:700a:
888b
Done
```

To disable an HA monitor by using the command line interface

At the command prompt, type:

- ♦ **set interface** <ifNum> [-haMonitor (ON | OFF)]
- ♦ **show interface** <ifNum>

Example

```
> set interface 1/3 -haMonitor OFF
Done
```

To add a remote node by using the configuration utility

1. Navigate to **System > High Availability**.
2. In the details pane, select the **Nodes** tab, and then click **Add**.
3. In the **High Availability Setup** dialog box, in the **Remote Node IP Address** text box, type the NSIP address of the NetScaler that is to be added as the remote node. If the NSIP is an IPv6 address, select the **IPv6** check box before entering the address.
4. If you want to add the local node to the remote node automatically, select the **Configure remote system to participate in High Availability setup** check box. If you do not select this option, you will have to log in to the appliance represented by the remote node and add the node that you are currently configuring.
5. Make sure that the **Turn off HA monitor on interfaces/channels that are down** check box is selected.
6. Select the **Turn on INC (Independent Network Configuration) mode on self mode** check box.
7. Click **OK**. The **Nodes** page displays both of the nodes in your HA configuration (the local node and the remote node).

Removing a Node

If you remove a node, the nodes are no longer in high availability configuration.

To remove a node by using the command line interface

At the command prompt, type:

rm ha node <id>

Example

```
> rm ha node 2
Done
```

To remove a node by using the configuration utility

1. In the navigation pane, expand **System**, and then click **High Availability**.
2. On the **High Availability** page, select the **Nodes** tab.
3. On the **Nodes** page, select the node that you want to remove, and click **Remove**.
4. On the **Remove** dialog box, click **Yes**.

Note: You can use the Network Visualizer to view the NetScaler appliances that are configured as a high availability (HA) pair and perform high availability configuration tasks. For more information, see "[Using the Network Visualizer](#)."

Configuring Route Monitors

You can use route monitors to make the HA state dependent on the internal routing table, whether or not the table contains any dynamically learned or static routes. In an HA configuration, a route monitor on each node watches the internal routing table to make sure that a route entry for reaching a particular network is always present. If the route entry is not present, the state of the route monitor changes to DOWN.

Adding a Route Monitor to a High Availability Node

A single procedure creates a route monitor and binds it to an HA node.

To add a route monitor by using the command line interface

At the command prompt, type:

- ♦ **bind HA node** <id> (-routeMonitor <ip_addr|ipv6_addr> [<netmask>])
- ♦ **show HA node**

Example

```
> bind HA node 3 -routeMonitor 10.102.71.0
255.255.255.0
```

```
Done
> bind HA node 3 -routeMonitor
1000:0000:0000:0000:0005:0600:700a:888b
Done
```

To add a route monitor by using the configuration utility

1. Navigate to **System > High Availability**.
2. In the details pane, on the **Route Monitors** tab, click **Configure**.
3. In **Bind / Unbind Route Monitor(s)** dialog box, in the **Network** text box, do one of the following:
 - For a IPv4 network, type an IPv4 network address (for example, 10.102.29.30) and in the **Netmask** text box, type a subnet mask (for example, 255.255.255.0).
 - For a IPv6 network, select the **IPv6** check box and type a IPv6 network address (for example, 1000:0000:0000:0000:0005:0600:700a:888b).
4. Click **Add**. The Route Monitor is added and appears in the **Configured Route Monitors** table.
5. Click **OK**.

Removing Route Monitors

To remove a route monitor by using the command line interface

At the command prompt, type:

- ♦ **unbind HA node <id> (-routeMonitor <ip_addr|ipv6_addr> [<netmask>])**
- ♦ **show ha node**

Example

```
unbind HA node 3 -routeMonitor 10.102.71.0
255.255.255.0
unbind HA node 3 -routeMonitor
1000:0000:0000:0000:0005:0600:700a:888b
```

To remove a route monitor by using the configuration utility

1. In the navigation pane, expand **System**, and then click **High Availability**.
2. In the details pane, on the **Route Monitors** tab, click **Configure**.
3. In the **Bind / Unbind Route Monitor(s)** dialog box, under **Configured Route Monitors**, select a route monitor to remove and click **Remove**.

4. Click **OK**.

Limiting Failovers Caused by Route Monitors in non-INC mode

In an HA configuration in non-INC mode, if route monitors fail on both nodes, failover happens every 180 seconds until one of the nodes is able to reach all of the routes monitored by the respective route monitors.

However, for a node, you can limit the number of failovers for a given interval by setting the Maximum Number of Flips and Maximum Flip Time parameters on the nodes. When either limit is reached, no more failovers occur, and the node is assigned as primary even if any route monitor fails on that node. If the node is then able to reach all of the monitored routes, the next monitor failure triggers resetting of the Maximum Number of Flips and Maximum Flip Time parameters on the node and starting the time specified in the Maximum Flip Time parameter.

These parameters are set independently on each node and therefore are neither propagated nor synchronized.

Parameters for limiting the number of failovers

Maximum Number of Flips (maxFlips)

Maximum number of failovers allowed, within the Maximum Flip Time interval, for the node in HA in non INC mode, if the failovers are caused by route-monitor failure.

Maximum Flip Time (maxFlipTime)

Amount of time, in seconds, during which failovers resulting from route-monitor failure are allowed for the node in HA in non INC mode.

To limit the number of failovers by using the command line interface

At the command prompt, type:

- ◆ `set HA node [-maxFlips < positive_integer>] [-maxFlipTime <positive_integer>]`
- ◆ `show HA node [< id>]`

Example

```
> set ha node -maxFlips 30 -maxFlipTime 60
Done
> sh ha node
1) Node ID: 0
IP: 10.102.169.82 (NS)
Node State: UP
Master State: Primary
Fail-Safe Mode: OFF
INC State: DISABLED
```

```

Sync State: ENABLED
Propagation: ENABLED
Enabled Interfaces : 1/1
Disabled Interfaces : None
HA MON ON Interfaces : 1/1
Interfaces on which heartbeats are not seen :None
Interfaces causing Partial Failure:None
SSL Card Status: NOT PRESENT
Hello Interval: 200 msec
Dead Interval: 3 secs
Node in this Master State for: 0:4:24:1
(days:hrs:min:sec)
2) Node ID: 1
IP: 10.102.169.81
Node State: UP
Master State: Secondary
Fail-Safe Mode: OFF
INC State: DISABLED
Sync State: SUCCESS
Propagation: ENABLED
Enabled Interfaces : 1/1
Disabled Interfaces : None
HA MON ON Interfaces : 1/1
Interfaces on which heartbeats are not seen : None
Interfaces causing Partial Failure: None
SSL Card Status: NOT PRESENT

Local node information:
Configured/Completed Flips: 30/0
Configured Flip Time: 60
Critical Interfaces: 1/1

Done

```

To limit the number of failovers by using the configuration utility

1. In the navigation pane, expand **System** and click **High Availability**.
2. In the details pane, on the **Nodes** tab, select the local node, and then click **Open**.
3. In the **Configure Node** dialog box, under **Intervals**, set the following parameters:
 - Maximum Number of Flips
 - Maximum Flip Time
4. Click **OK**.

Configuring FIS

Link redundancy is a way to prevent failover by grouping interfaces so that, when one interface fails, other functioning interfaces are still available. The link redundancy feature allows you to group the two interfaces into a failover interface set (FIS), which

prevents the failure of a single link from causing failover to the secondary system unless all of the interfaces on the primary system are nonfunctional.

Each interface in an FIS maintains independent bridge entries. HA MON interfaces that are not bound to an FIS are known as critical interfaces (CI) because if any of them fails, failover is triggered.

Creating or Modifying an FIS

To add an FIS and bind interfaces to it by using the command line interface

At the command prompt, type:

- ♦ **add fis** <name>
- ♦ **bind fis** <name> <ifnum> ...
- ♦ **show fis** <name>

Example

```
> add fis fis1
Done
> bind fis fis1 1/3 1/5
Done
```

An unbound interface becomes a critical interface (CI) if it is enabled and HA MON is on.

To unbind an interface from an FIS by using the command line interface

At the command prompt, type:

- ♦ **unbind fis** <name> <ifnum> ...
- ♦ **show fis** <name>

Example

```
> unbind fis fis1 1/3
Done
```

To configure an FIS by using the configuration utility

1. In the navigation pane, expand **System**, and then click **High Availability**.
2. In the details pane, on the **Failover Interface Set** tab, do one of the following:

- To create a new FIS, click **Add**.
 - To modify an existing FIS, click **Open**.
3. In the **Create FIS** or **Configure FIS** dialog box, in the **Name** text box, type the name of the FIS.
 4. Select an available interface and click **Add** to bind it to the FIS. Repeat to bind additional interfaces.
 5. Click **OK**.

Removing an FIS

When the FIS is removed, its interfaces are marked as critical interfaces.

To remove an FIS by using the command line interface

At the command prompt, type:

```
rm fis <name>
```

Example

```
> rm fis fis1
Done
```

To remove an FIS by using the configuration utility

1. In the navigation pane, expand **System**, and then click **High Availability**.
2. In the details pane, on the **Failover Interface Set** tab, select the FIS that you want to remove and click **Remove**.
3. In the **Remove** dialog box, click **Yes**.

Understanding the Causes of Failover

The following events can cause failover in an HA configuration:

1. If the secondary node does not receive a heartbeat packet from the primary for a period of time that exceeds the dead interval set on the secondary. (See Note: 1.)
2. The primary node experiences a hardware failure of its SSL card.
3. The primary node does not receive any heartbeat packets on its network interfaces for three seconds.
4. On the primary node, a network interface that is not part of a Failover Interface Set (FIS) or a Link Aggregation (LA) channel and has the HA Monitor (HAMON) enabled, fails. (See Note: 2.)

5. On the primary node, all interfaces in an FIS fail. (See Note: 2.)
6. On the primary node, an LA channel with HAMON enabled fails. (See Note: 2.)
7. On the primary node, all interfaces fail (see Note: 2). In this case, failover occurs regardless of the HAMON configuration.
8. On the primary node, all interfaces are manually disabled. In this case, failover occurs regardless of the HAMON configuration.
9. You force a failover by issuing the force failover command on either node.
10. A route monitor that is bound to the primary node goes DOWN.

Note: 1. For more information about setting the dead interval, see [Configuring the Communication Intervals](#). Possible causes for a node not receiving heartbeat packets from a peer node include:

- A network configuration problem prevents heartbeats from traversing the network between the HA nodes.
- The peer node experiences a hardware or software failure that causes it to freeze (hang), reboot, or otherwise stop processing and forwarding heartbeat packets.

Note: 2. In this case, fail means that the interface was enabled but goes to the DOWN state, as can be seen from the **show interface** command or from the configuration utility. Possible causes for an enabled interface to be in the DOWN state are LINK DOWN and TXSTALL.

Forcing a Node to Fail Over

You might want to force a failover if, for example, you need to replace or upgrade the primary node. You can force failover from either the primary or the secondary node. A forced failover is not propagated or synchronized. To view the synchronization status after a forced failover, you can view the status of the node.

A forced failover fails in any of the following circumstances:

- ♦ You force failover on a standalone system.
- ♦ The secondary node is disabled.
- ♦ The secondary node is configured to remain secondary.

The NetScaler appliance displays a warning message if it detects a potential issue when you run the force failover command. The message includes the information that triggered the warning, and requests confirmation before proceeding.

Forcing Failover on the Primary Node

If you force failover on the primary node, the primary becomes the secondary and the secondary becomes the primary. Forced failover is possible only when the primary node can determine that the secondary node is UP.

If the secondary node is DOWN, the force failover command returns the following error message: "Operation not possible due to invalid peer state. Rectify and retry."

If the secondary system is in the claiming state or inactive, it returns the following error message: "Operation not possible now. Please wait for system to stabilize before retrying."

To force failover on the primary node by using the command line interface

At the command prompt, type:

```
force HA failover
```

To force failover on the primary node by using the configuration utility

1. In the navigation pane, expand **System**, and then click **High Availability**.
2. In the details pane, on the **Nodes** tab, click **Force Failover**.
3. In the **Warning** dialog box, click **Yes**.

Forcing Failover on the Secondary Node

If you run the force failover command from the secondary node, the secondary node becomes primary and the primary node becomes secondary. A force failover can occur only if the secondary node's health is good and it is not configured to stay secondary.

If the secondary node cannot become the primary node, or if secondary node was configured to stay secondary (using the STAYSECONDARY option), the node displays the following error message: "Operation not possible as my state is invalid. View the node for more information."

To force failover on the secondary node by using the command line interface

At the command prompt, type:

```
force HA failover
```

To force failover on the secondary node by using the configuration utility

1. In the navigation pane, expand **System**, and then click **High Availability**.
2. In the details pane, on the **Nodes** tab, click **Force Failover**.

3. In the **Warning** dialog box, click **Yes**.

Forcing Failover When Nodes Are in Listen Mode

When the two nodes of an HA pair are running different versions of the system software, the node running the higher version switches to the listen mode. In this mode, neither command propagation nor synchronization works.

Before upgrading the system software on both nodes, you should test the new version on one of the nodes. To do this, you need to force a failover on the system that has already been upgraded. The upgraded system then takes over as the primary node, but neither command propagation or synchronization occurs. Also, all connections need to be re-established.

To force failover when nodes are in listen mode by using the command line interface

At the command prompt, type:

```
force HA failover
```

To force failover when nodes are in listen mode by using the configuration utility

1. In the navigation pane, expand **System**, and then click **High Availability**.
2. In the details pane, on the **Nodes** tab, click **Force Failover**.
3. In the **Warning** dialog box, click **Yes**.

Forcing the Secondary Node to Stay Secondary

In an HA setup, the secondary node can be forced to stay secondary regardless of the state of the primary node.

For example, suppose the primary node needs to be upgraded and the process will take a few seconds. During the upgrade, the primary node may go down for a few seconds, but you do not want the secondary node to take over; you want it to remain the secondary node even if it detects a failure in the primary node.

When you force the secondary node to stay secondary, it will remain secondary even if the primary node goes down. Also, when you force the status of a node in an HA pair to stay secondary, it does not participate in HA state machine transitions. The status of the node is displayed as **STAYSECONDARY**.

Forcing the node to stay secondary works on both standalone and secondary nodes. On a standalone node, you must use this option before you can add a node to create an HA pair. When you add the new node, the existing node continues to function as the primary node, and the new node becomes the secondary node.

Note: When you force a system to remain secondary, the forcing process is not propagated or synchronized. It affects only the node on which you run the command.

To force the secondary node to stay secondary by using the command line interface

At the command prompt, type:

```
set ha node -hastatus STAYSECONDARY
```

To force the secondary node to stay secondary by using the configuration utility

1. Navigate to **System > High Availability**.
2. In the details pane, on the **Nodes** tab, select the local node, and then click **Open**.
3. In the **Configure Node** dialog box, under **High Availability Status**, select **STAY SECONDARY**.
4. Click **OK**.

Forcing the Primary Node to Stay Primary

In an HA setup, you can force the primary node to remain primary even after a failover. You can enable this option either on a primary node in an HA pair or on a standalone system.

On a standalone system, you must run this command before you can add a node to create an HA pair. When you add the new node, it becomes the primary node. The existing node stops processing traffic and becomes the secondary node in the HA pair.

To force the primary node to stay primary by using the command line interface

At the command prompt, type:

```
set ha node -hastatus STAYPRIMARY
```

To force the primary node to stay primary by using the configuration utility

1. Navigate to **System > High Availability**.
2. In the details pane, on the **Nodes** tab, select the local node, and then click **Open**.
3. In the **Configure Node** dialog box, under **High Availability Status**, select **STAY PRIMARY**.
4. Click **OK**.

Understanding the High Availability Health Check Computation

The following table summarizes the factors examined in a health check computation:

- ♦ State of the CIs
- ♦ State of the FISs
- ♦ State of the route monitors

The following table summarizes the health check computation.

Table 6-3. High Availability Health Check Computation

FIS	CI	Route monitor	Condition
N	Y	N	If the system has any CIs, all of those CIs must be UP.
Y	Y	N	If the system has any FISs, all of those FISs must be UP.
Y	Y	Y	If the system has any route monitors configured, all monitored routes must be present in the FIS.

High Availability

What are the various ports used to exchange the HA-related information between the nodes in an HA configuration?

In an HA configuration, both nodes use the following ports to exchange HA related information:

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

What are the conditions that trigger synchronization?

Synchronization is triggered by any of the following conditions:

- ♦ The incarnation number of the primary node, received by the secondary, does not match that of the secondary node.

Note: Both nodes in an HA configuration maintain a counter called *incarnation number*, which counts the number of configurations in the node's configuration file. Each node sends its incarnation number to each other node in the heartbeat messages. The incarnation number is not incremented for the following commands:

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

- ♦ The secondary node comes up after a restart.
- ♦ The primary node becomes secondary after a failover.

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

The following commands are neither propagated nor synced to the secondary node:

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

What configurations are not synced nor propagated in an HA configuration in INC mode?

The following configurations are not synced or propagated. Each node has its own.

- ♦ MIPs
- ♦ SNIPs
- ♦ VLANs
- ♦ Routes (except LLB routes)

- ♦ Route monitors
- ♦ RNAT rules (except any RNAT rule with VIP as the NAT IP)
- ♦ Dynamic routing configurations.

Does a configuration added to the secondary node get synchronized on the primary?

No, a configuration added to the secondary node is not synchronized to the primary.

What could be the reason for both nodes claiming to be the primary in an HA configuration?

The most likely reason is that the primary and secondary nodes are both healthy but the secondary does not receive the heartbeat packets from the primary. The problem could be with the network between the nodes.

Does an HA configuration run into any issues if you deploy the two nodes with different system clock settings?

Different system-clock settings on the two nodes can cause the following issues:

- ♦ The time stamps in the log file entries do not match. This situation makes it difficult to analyze the log entries for any issues.
- ♦ After a failover, you might have problems with any type of cookie based persistence for load balancing. A significant difference between the times can cause a cookie to expire sooner than expected, resulting in termination of the persistence session.
- ♦ Similar considerations apply to any time related decisions on the nodes.

What are the conditions for failure of the *force HA sync* command?

Forced synchronization fails in any of the following circumstances:

- ♦ You force synchronization when synchronization is already in progress.
- ♦ You force synchronization on a standalone NetScaler appliance.
- ♦ The secondary node is disabled.
- ♦ HA synchronization is disabled on the current secondary node.
- ♦ HA propagation is disabled on the current primary node and you force synchronization from the primary.

What are the conditions for failure of the *sync HA files* command?

Synchronizing configuration files fail in either of the following circumstances:

- ♦ On a standalone system.
- ♦ With the secondary node disabled.

In an HA configuration, if the secondary node takes over as the primary, does it switch back to secondary status if the original primary comes back online?

No. After the secondary node takes over as the primary, it remains as primary even if the original primary node comes back online again. To interchange the primary and secondary status of the nodes, run the *force failover* command.

What are the conditions for failure of the *force failover* command?

A forced failover fails in any of the following circumstances:

- ♦ You force failover on a standalone system.
- ♦ The secondary node is disabled.
- ♦ The secondary node is configured to remain secondary.
- ♦ The primary node is configured to remain primary.
- ♦ The state of the peer node is unknown.

Troubleshooting High Availability Issues

The most common high availability issues involve the high availability feature not working at all, or working only intermittently. Following are common high availability issues, and probable causes and resolutions.

- ♦ **Issue**

The inability of the NetScaler appliances to pair the NetScaler appliances in a high availability setup.

- **Cause**

Network connectivity

Resolution

Verify that both the appliances are connected to the switch and the interfaces are enabled.

- **Cause**

Mismatch in the Password for the default Administrator account

Resolution

Verify that the password on both the appliances is the same.

- **Cause**

IP conflict

Resolution

Verify that both the appliances have unique NetScaler IP (NSIP) address. The appliances should not have the same NSIP address.

- **Cause**
Node ID mismatch
Resolution
Verify that the Node ID Configuration on both the appliances is unique. The appliances should not have the same Node ID configuration. Additionally, you must assign value for a Node ID between 1 and 64.
- **Cause**
Mismatch in the password of the RPC node
Resolution
Verify that both the nodes have the same RPC node password.
- **Cause**
An administrator has disabled the remote node
Resolution
Enable the remote node.
- **Cause**
The Firewall application has blocked the heartbeat packets
Resolution
Verify that the UDP port 3003 is allowed.
- ♦ **Issue**
Both the appliances claim to be the primary appliance.
 - **Cause**
Missing heartbeat packets between the appliances
Resolution
Verify that the UDP port 3003 is not blocked for communication between the appliances.
- ♦ **Issue**
The NetScaler appliance is not able to synchronize the configuration.
 - **Cause**
A Firewall application is blocking the required port.
Resolution
Verify that the UDP port 3010 (or UDP port 3008 with secure synchronization) is not blocked for communication between the appliances.
 - **Cause**
An administrator has disabled synchronization.
Resolution

Enable synchronization on the appliance that has the issue.

- **Cause**

Different NetScaler releases or builds are installed on appliances.

Resolution

Upgrade the appliances to the same NetScaler release or build.

- ♦ **Issue**

Command propagation fails between the appliances.

- **Cause**

A Firewall application is blocking the port.

Resolution

Verify that the UDP port 3011 (or UDP port 3009 with secure propagation) is not blocked for communication between the appliances.

- **Cause**

An administrator has disabled command propagation.

Resolution

Enable command propagation on the appliance that has the issue.

- **Cause**

Different NetScaler releases or builds are installed on appliances.

Resolution

Upgrade the appliances to the same NetScaler release or build.

- ♦ **Issue**

The NetScaler appliances in the high availability pair are unable to run the force failover process.

- **Cause**

The Secondary node is disabled.

Resolution

Enable the secondary node.

- **Cause**

The Secondary node is configured to stay secondary.

Resolution

Set the secondary high availability status of the secondary node to Enable from Stay Secondary.

- ♦ **Issue**

The secondary appliance does not receive any traffic after the failover process.

- **Cause**

The upstream router does not understand GARP messages of NetScaler appliance.

Resolution

Configure Virtual MAC (VMAC) address on the secondary appliance.

Chapter 7

Networking

Topics:

- [IP Addressing](#)
- [Interfaces](#)
- [Access Control Lists](#)
- [IP Routing](#)
- [Internet Protocol version 6 \(IPv6\)](#)
- [Traffic Domains](#)

The following topics provide a conceptual reference and instructions for configuring the various networking components on the NetScaler appliance.

IP Addressing	Learn the various types of NetScaler-owned IP addresses and how to create, customize, and remove them.
Interfaces	Configure some of the basic network configurations that must be done to get started.
Access Control Lists (ACLs)	Configure the different types of Access Control Lists and how to create, customize, and remove them.
IP Routing	Learn and configure the routing functionality of the NetScaler appliance, both static and dynamic.
Internet Protocol version 6 (IPv6)	Learn how the NetScaler appliance supports IPv6.
Traffic Domains	Learn and configure traffic domains to segment network traffic for different applications.

IP Addressing

Before you can configure the NetScaler appliance, you must assign the NetScaler IP Address (NSIP), also known as the Management IP address. You can also create other NetScaler-owned IP addresses for abstracting servers and establishing connections with the servers. In this type of configuration, the appliance serves as a proxy for the abstracted servers. You can also proxy connections by using network address translations (INAT and RNAT). When proxying connections, the appliance can behave either as a bridging (Layer 2) device or as a packet forwarding (Layer 3) device. To make packet forwarding more efficient, you can configure static ARP entries. For IPv6, you can configure neighbor discovery (ND).

Configuring NetScaler-Owned IP Addresses

The NetScaler-owned IP Addresses—NetScaler IP Address (NSIP), Virtual IP Addresses (VIPs), Subnet IP Addresses (SNIPs), Mapped IP Addresses (MIPs), and Global Server Load Balancing Site IP Addresses (GSLBIPs)—exist only on the NetScaler appliance. The NSIP uniquely identifies the NetScaler on your network, and it provides access to the appliance. A VIP is a public IP address to which a client sends requests. The NetScaler terminates the client connection at the VIP and initiates a connection with a server. This new connection uses a SNIP or a MIP as the source IP address for packets forwarded to the server. If you have multiple data centers that are geographically distributed, each data center can be identified by a unique GSLBIP.

You can configure some NetScaler-owned IP addresses to provide access for management applications.

Configuring the NetScaler IP Address (NSIP)

The NetScaler IP (NSIP) address is the IP address at which you access the NetScaler for management purposes. The NetScaler can have only one NSIP, which is also called the Management IP address. You must add this IP address when you configure the NetScaler for the first time. If you modify this address, you must reboot the NetScaler. You cannot remove an NSIP address. For security reasons, NSIP should be a non-routable IP address on your organization's LAN.

Note: Configuring the NetScaler IP address is mandatory.

To create the NetScaler IP address by using the command line interface

At the command prompt, type:

- ♦ `set ns config [-IPAddress <ip_addr> -netmask <netmask>]`
- ♦ `show ns config`

Example

```
> set ns config -ipaddress 10.102.29.170 -netmask  
255.255.255.0  
Done
```

To configure the NetScaler IP address by using the configuration utility

1. In the navigation pane, click **System**.
2. On the **System Information** tab, click **Setup Wizard**.
3. In the **Setup Wizard** dialog box, click **Next**.
4. Under **System Configuration**, set the following parameters:
 - **IP Address**
 - **Netmask**
5. Follow the instructions in the **Setup Wizard** to complete the configuration.

Configuring and Managing Virtual IP Addresses (VIPs)

Configuration of a virtual server IP address (VIP) is not mandatory during initial configuration of the NetScaler. When you configure load balancing, you assign VIPs to virtual servers.

In some situations, you need to customize VIP attributes or enable or disable a VIP. A VIP is usually associated with a virtual server, and some of the attributes of the VIP are customized to meet the requirements of the virtual server. You can host the same virtual server on multiple NetScaler appliances residing on the same broadcast domain, by using ARP and ICMP attributes. After you add a VIP (or any IP address), the NetScaler sends, and then responds to, ARP requests. VIPs are the only NetScaler-owned IP addresses that can be disabled. When a VIP is disabled, the virtual server using it goes down and does not respond to ARP, ICMP, or L4 service requests.

As an alternative to creating VIPs one at a time, you can specify a consecutive range of VIPs.

To create a VIP address by using the command line interface

At the command prompt, type:

- ♦ **add ns ip** <IPAddress> <netmask> -type <type>
- ♦ **show ns ip** <IPAddress>

Example

```
> add ns ip 10.102.29.59 255.255.255.0 -type VIP
Done
```

To create a range of VIP addresses by using the command line interface

At the command prompt, type:

- ♦ `add ns ip <IPAddress> <netmask> -type <type>`
- ♦ `show ns ip <IPAddress>`

Example

```
> add ns ip 10.102.29.[60-64] 255.255.255.0 -type
VIP
ip "10.102.29.60" added
ip "10.102.29.61" added
ip "10.102.29.62" added
ip "10.102.29.63" added
ip "10.102.29.64" added
Done
```

To configure a VIP address by using the configuration utility

1. Navigate to **System > Network > IPs**.
2. In the details pane, do one of the following:
 - To create a new IP, click **Add**.
 - To modify an existing IP, select the IP, and then click **Open**.
3. In the **Create IP** or **Configure IP** dialog box, set the following parameters:
 - **IP Address***
 - **Netmask***
 - **IP Type:** Select **VIP**.
 - **ARP Response**
 - **ICMP Response**
 - **ARP**
 - **Virtual Server**

- **Dynamic Routing**
- **Host Route**
- **Gateway IP***
- **Metric**
- **V Server RHI Level**
- **OSPF LSA Type**
- **Area**

*A required parameter

4. Click **Create** or **OK**, and then click **Close**. The IP address that you configured appears in the details pane.

To create a range of VIP addresses by using the configuration utility

1. Navigate to **System > Network > IPs**.
2. In the details pane, click **Add Range**.
3. In the **Create IP - Range** dialog box, set the following parameters:

- **IP Address***
- **Netmask***
- **Type**—type. Select **VIP**.
- **IP Type**
- **ARP**
- **ICMP Response**
- **Virtual Server**
- **Dynamic Routing**
- **Host Route**
- **Gateway IP***
- **Metric**
- **V Server RHI Level**
- **OSPF LSA Type**
- **Area**

*A required parameter

4. Click **Create**, and then click **Close**. The range of IP addresses that you created appears in the details pane.

To enable or disable an IPv4 VIP address by using the command line interface

At the command prompt, type one of the following sets of commands to enable or disable a VIP and verify the configuration:

- ♦ **enable ns ip <IPAddress>**
- ♦ **show ns ip <IPAddress>**
- ♦ **disable ns ip <IPAddress>**
- ♦ **show ns ip <IPAddress>**

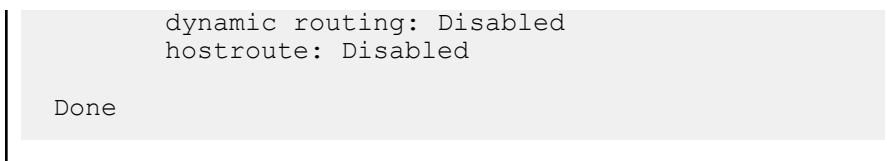
Example

```
> enable ns ip 10.102.29.79
Done
> show ns ip 10.102.29.79

      IP: 10.102.29.79
      Netmask: 255.255.255.255
      Type: VIP
      state: Enabled
      arp: Enabled
      icmp: Enabled
      vserver: Enabled
      management access: Disabled
         telnet: Disabled
         ftp: Disabled
         ssh: Disabled
         gui: Disabled
         snmp: Disabled
      Restrict access: Disabled
      dynamic routing: Disabled
      hostroute: Disabled

Done
> disable ns ip 10.102.29.79
Done
> show ns ip 10.102.29.79

      IP: 10.102.29.79
      Netmask: 255.255.255.255
      Type: VIP
      state: Disabled
      arp: Enabled
      icmp: Enabled
      vserver: Enabled
      management access: Disabled
         telnet: Disabled
         ftp: Disabled
         ssh: Disabled
         gui: Disabled
         snmp: Disabled
      Restrict access: Disabled
```



To enable or disable a VIP address by using the configuration utility

1. Navigate to **System > Network > IPs**.
2. In the details pane, on the **IPv4s** tab, select the VIP address and do one of the following:
 - To enable the selected IP address, click **Enable**.
 - To disable the selected IP address, click **Disable**.
3. In the details pane, verify that the VIP address is enabled or disabled, as appropriate.

Configuring ARP response Suppression for Virtual IP addresses (VIPs)

You can configure the NetScaler appliance to respond or not respond to ARP requests for a Virtual IP (VIP) address on the basis of the state of the virtual servers associated with that VIP.

For example, if virtual servers V1, of type HTTP, and V2, of type HTTPs, share VIP address 10.102.29.45 on a NetScaler appliance, you can configure the appliance to not respond to any ARP request for VIP 10.102.29.45 if both V1 and V2 are in the DOWN state.

The following three options are available for configuring ARP-response suppression for a virtual IP address.

- ♦ **NONE.** The NetScaler appliance responds to any ARP request for the VIP address, irrespective of the state of the virtual servers associated with the address.
- ♦ **ONE VSERVER.** The NetScaler appliance responds to any ARP request for the VIP address if at least one of the associated virtual servers is in UP state.
- ♦ **ALL VSERVER.** The NetScaler appliance responds to any ARP request for the VIP address if all of the associated virtual servers are in UP state.

Following table shows the sample behavior of NetScaler appliance for a VIP configured with two virtual servers:

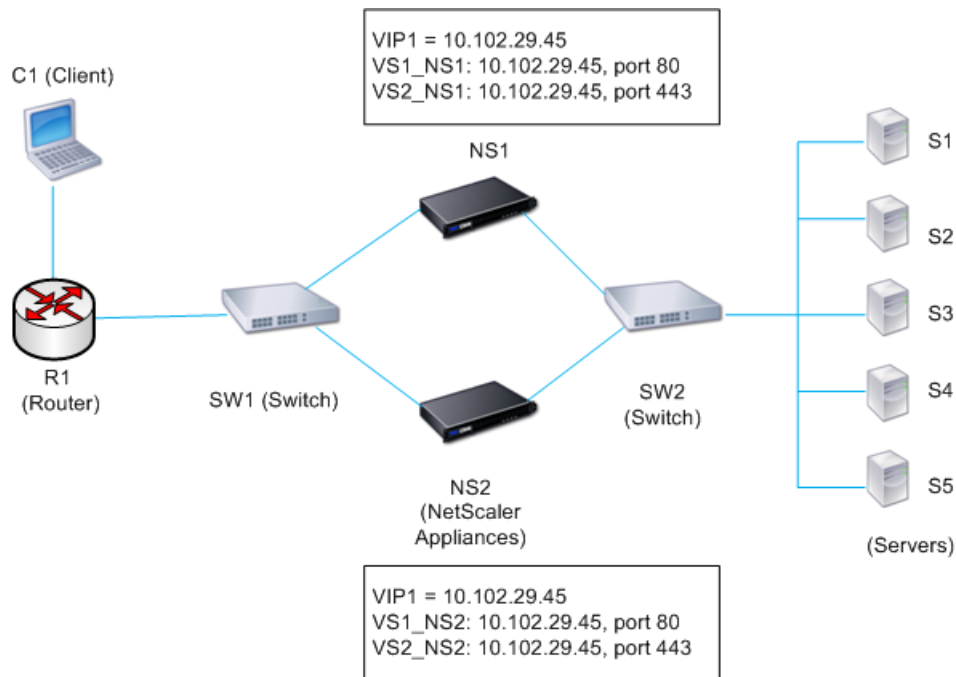
Associated virtual servers for a VIP	STATE 1	STATE 2	STATE 3	STATE 4
NONE				
V1	UP	UP	DOWN	DOWN

Associated virtual servers for a VIP	STATE 1	STATE 2	STATE 3	STATE 4
V2	UP	DOWN	UP	DOWN
Respond to an ARP request for this VIP?	Yes	Yes	Yes	Yes
ONE VSERVER				
V1	UP	UP	DOWN	DOWN
V2	UP	DOWN	UP	DOWN
Respond to an ARP request for this VIP?	Yes	Yes	Yes	No
ALL VSERVER				
V1	UP	UP	DOWN	DOWN
V2	UP	DOWN	UP	DOWN
Respond to an ARP request for this VIP?	Yes	No	No	No

Consider an example where you want to test the performance of two virtual servers, V1 and V2, which have the same VIP address but are of different types and are each configured on NetScaler appliances NS1 and NS2. Let's call the shared VIP address *VIP1*.

V1 load balances servers S1, S2, and S3. V2 load balances servers S4 and S5.

On both NS1 and NS2, for VIP1, the ARP suppression parameter is set to ALL_VSERVER. If you want to test the performance of V1 and V2 on NS1, you must manually disable V1 and V2 on NS2, so that NS2 does not respond to any ARP request for VIP1.

Figure 7-1.

The execution flow is as follows:

1. Client C1 sends a request to V1. The request reaches R1.
2. R1 does not have an APR entry for the IP address (VIP1) of V1, so R1 broadcasts an ARP request for VIP1.
3. NS1 replies with source MAC address MAC1 and source IP address VIP1. NS2 does not reply to the ARP request.
4. SW1 learns the port for VIP1 from the ARP reply and updates its bridge table, and R1 updates the ARP entry with MAC1 and VIP1.
5. R1 forwards the packet to address VIP1 on NS1.
6. NS1's load balancing algorithm selects server S2, and NS1 opens a connection between one of its SNIP or MIP addresses and S2. When S2 sends a response to the client, the response returns by the same path.
7. Now you want to test the performance of V1 and V2 on NS2, so you enable V1 and V2 on NS2 and disable them on NS1. NS2 now broadcasts an ARP message for VIP1. In the message, MAC2 is the source MAC address and VIP1 is the source IP address.
8. SW1 learns the port number for reaching MAC2 from the ARP broadcast and updates its bridge table to send subsequent client requests for VIP1 to NS2. R1 updates its ARP table.
9. Now suppose the ARP entry for VIP1 times out in the ARP table of R1, and client C1 sends a request for V1. Because R1 does not have an APR entry for VIP1, it broadcasts an ARP request for VIP1.

10. NS2 replies with a source MAC address and VIP1 as the source IP address. NS1 does not reply to the ARP request.

To configure ARP response suppression by using the command line interface

At the command prompt, type:

- ♦ `set ns ip -arpResponse <arpResponse>]`
- ♦ `show ns ip <IPAddress>`

Example

```
> set ns ip 10.102.29.96 -arpResponse ALL_VSERVERS
Done
```

To configure ARP response suppression by using the configuration utility

1. Navigate to **System > Network > IPs > IPv4**.
2. In the details pane, select the IP, and then click **Open**.
3. In the **Configure IP** dialog box, set the **ARP Response** parameter.
4. Click **OK**, and then click **Close**.

Configuring Subnet IP Addresses (SNIPs)

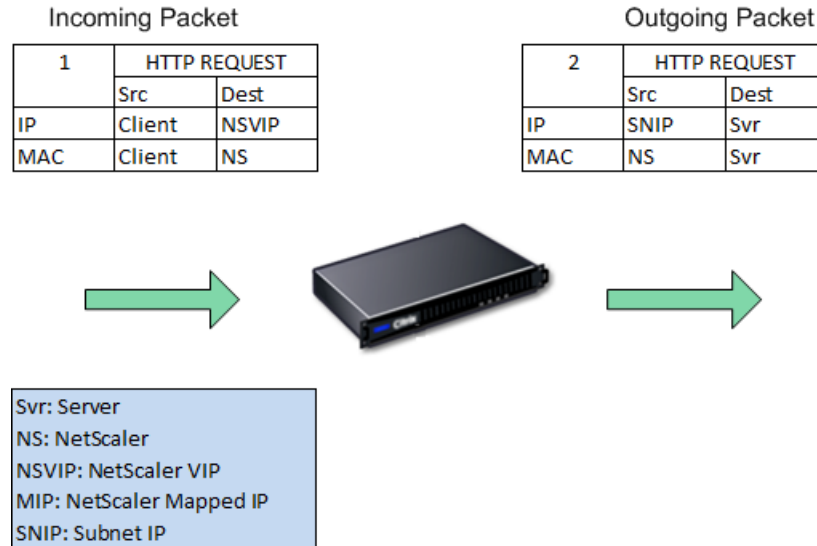
A subnet IP (SNIP) address is used in connection management and server monitoring. It is not mandatory to specify a SNIP when you initially configure the NetScaler appliance. In a multiple-subnet scenario, the NetScaler IP (NSIP) address, the mapped IP (MIP) address, and the IP address of a server can exist on different subnets. To eliminate the need to configure additional routes on devices such as servers, you can configure subnet IP addresses (SNIPs) on the NetScaler. With Use SNIP (USNIP) mode enabled, a SNIP is the source IP address of a packet sent from the NetScaler to the server, and the SNIP is the IP address that the server uses to access the NetScaler. This mode is enabled by default.

The SNIP enables the NetScaler appliance to connect to the subnet, which is different than that of the MIP and NSIP addresses, similar to local network of the appliance. This functionality is very useful in the topology where backend servers are connected directly to the NetScaler appliance through an L2 switch and are in different subnets than that of MIP and NSIP addressed servers.

When you add a SNIP, a route corresponding to the SNIP is added to the routing table. The NetScaler determines the next hop for a service from the routing table, and if the IP address of the hop is within the range of a SNIP, the NetScaler uses the SNIP to source traffic to the service. When multiple SNIPs cover the IP addresses of the next hops, the SNIPs are used in round robin manner.

The following figure illustrates USNIP mode.

Figure 7-2. USNIP Mode



As an alternative to creating SNIPs one at a time, you can specify a consecutive range of SNIPs.

To configure a SNIP address by using the command line interface

At the command prompt, type:

- `add ns ip <IPAddress> <netmask> -type <type>`
- `show ns ip <IPAddress>`

Example

```
> add ns ip 10.102.29.203 255.255.255.0 -type SNIP
Done
```

To create a range of SNIP addresses by using the command line interface

At the command prompt, type:

- `add ns ip <IPAddress> <netmask> -type <type>`
- `show ns ip <IPAddress>`

Example

```
> add ns ip 10.102.29.[205-209] 255.255.255.0 -
```

```
type SNIP
ip "10.102.29.205" added
ip "10.102.29.206" added
ip "10.102.29.207" added
ip "10.102.29.208" added
ip "10.102.29.209" added
Done
```

To configure a SNIP address by using the configuration utility

1. Navigate to **System > Network > IPs > IPv4**.
2. Navigate to **Network > IPs > IPv4**.
3. In the details pane, do one of the following:
 - To create a new IP address, click **Add**.
 - To modify an existing IP address, select the address, and then click **Open**.
4. In the **Create IP** or **Configure IP** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
5. Click **Create** or **OK**, and then click **Close**. The IP address that you configured appears in the details pane.

To create a range of SNIP addresses by using the configuration utility

1. Navigate to **System > Network > IPs > IPv4**.
2. In the details pane, click **Add Range**.
3. In the **Create IP - Range** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create**, and then click **Close**. The range of IP addresses that you created appears in the details pane.

To enable or disable USNIP mode by using the command line interface

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

- ♦ `enable ns mode usnip`
- ♦ `disable ns mode usnip`

To enable or disable USNIP mode by using the configuration utility

1. Navigate to **System > Settings**.
2. In the details pane, in the **Modes and Features** group, click **Change modes**.
3. In the **Configure Modes** dialog box, do one of the following:
 - To enable USNIP, select the **Use Subnet IP** check box.

- To disable USNIP, clear the **Use Subnet IP** check box.
4. Click **OK**.
 5. In the **Enable/Disable Feature(s)?** dialog box, click **Yes**.

Configuring Mapped IP Addresses (MIPs)

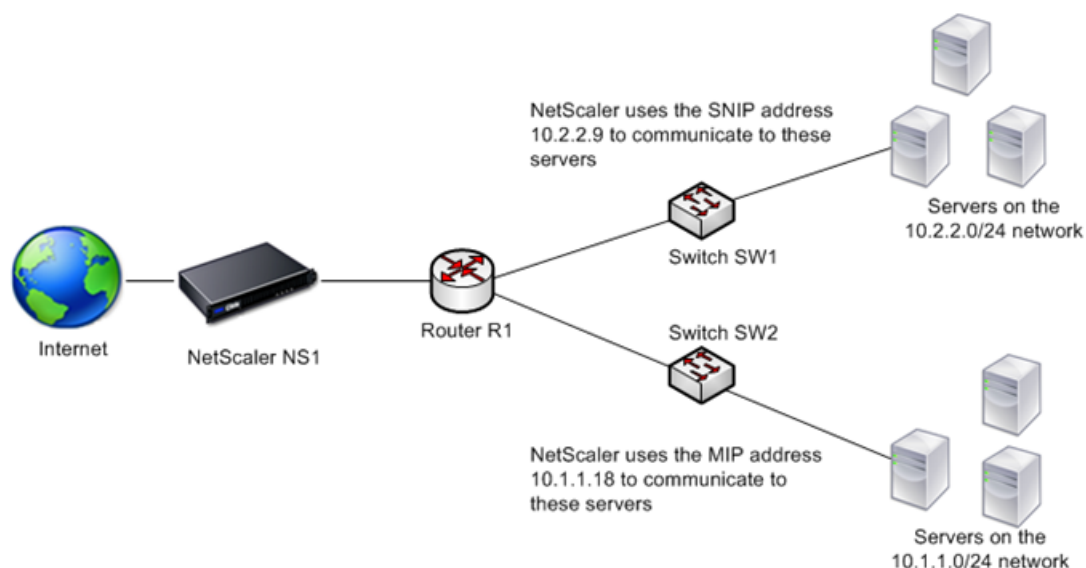
Mapped IP addresses (MIP) are used for server-side connections. A MIP can be considered a default Subnet IP (SNIP) address, because MIPs are used when a SNIP is not available or Use SNIP (USNIP) mode is disabled.

If the mapped IP address is the first in the subnet, the NetScaler appliance adds a route entry, with this IP address as the gateway to reach the subnet. You can create or delete a MIP during run time without rebooting the appliance.

As an alternative to creating MIPs one at a time, you can specify a consecutive range of MIPs.

The following diagram shows the use of the MIP and SNIP addresses in a NetScaler appliance that connects to the backend servers across the subnets.

Figure 7-3. MIP and SNIP addresses



In the setup, if the NetScaler appliance and the backend servers are in the 10.1.1.0/24 subnet, then the appliance uses the MIP address to communicate to the servers. However, if the setup has backend servers on additional subnets, such as 10.2.2.0/24, and there is no router between the NetScaler appliance and the subnet, then you can configure a SNIP address that has a range of 10.2.2.x/24, such as 10.2.2.9 in this case, to communicate to the additional subnet.

You can enable the NetScaler appliance to use MIP to communicate to the additional subnet. However, if the setup has a Firewall application between the appliance and the server, then the Firewall might prevent the traffic other than 10.2.2.0/24. In such cases, you need a SNIP address to communicate to the servers.

To create a MIP address by using the command line interface

At the command prompt, type:

- ♦ `add ns ip <IPAddress> <netmask> -type <type>`
- ♦ `show ns ip <IPAddress>`

Example

```
> add ns ip 10.102.29.171 255.255.255.0 -type MIP
Done
```

To create a range of MIP addresses by using the command line interface

At the command prompt, type:

- ♦ `add ns ip <IPAddress> <netmask> -type <type>`
- ♦ `show ns ip <IPAddress>`

Example

```
> add ns ip 10.102.29.[173-175] 255.255.255.0 -
type MIP
ip "10.102.29.173" added
ip "10.102.29.174" added
ip "10.102.29.175" added
Done
```

To configure a MIP address by using the configuration utility

1. Navigate to **System > Network > IPs > IPv4**.
2. In the details pane, do one of the following:
 - To create a new IP address, click **Add**.
 - To modify an existing IP address, select the address, and then click **Open**.
3. In the **Create IP** or **Configure IP** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create** or **OK**, and then click **Close**. The IP address that you configured appears in the details pane.

To create a range of MIP addresses by using the configuration utility

1. Navigate to **System > Network > IPs > IPv4**.
2. In the details pane, click **Add Range**.
3. In the **Create IP - Range** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create**, and then click **Close**. The range of IP addresses that you created appears in the details pane.

Configuring GSLB Site IP Addresses (GSLBIP)

A GSLB site IP (GSLBIP) address is an IP address associated with a GSLB site. It is not mandatory to specify a GSLBIP address when you initially configure the NetScaler appliance. A GSLBIP address is used only when you create a GSLB site.

Removing a NetScaler-Owned IP Address

You can remove any IP address except the NSIP. The following table provides information about the processes you must follow to remove the various types of IP addresses. Before removing a VIP, remove the associated virtual server.

Table 7-1. Implications of Removing a NetScaler-Owned IP Address

IP address type	Implications
Subnet IP address (SNIP)	If IP address being removed is the last IP address in the subnet, the associated route is deleted from the route table. If the IP address being removed is the gateway in the corresponding route entry, the gateway for that subnet route is changed to another NetScaler-owned IP address.
Mapped IP address (MIP)	<p>If a SNIP exists, you can remove the MIPs. The NetScaler uses NSIP and SNIPs to communicate with the servers when the MIP is removed. Therefore, you must also enable use SNIP (USNIP) mode.</p> <p>For information about enabling and disabling USNIP mode, see "Configuring Subnet IP Addresses (SNIPs)."</p>
Virtual Server IP address (VIP)	Before removing a VIP, you must first remove the vserver associated with it.

IP address type	Implications
GSLB-Site-IP address	Before removing a GSLB site IP address, you must remove the site associated with it.

To remove an IP address by using the command line interface

At the command prompt, type:

```
rm ns ip <IPaddress>
```

Example

```
rm ns ip 10.102.29.54
```

To remove an IP address by using the configuration utility

1. Navigate to **System > Network > IPs**.
2. On the **IPs** page, on the **IPv4s** tab, select the IP address that you want to remove, and then click **Remove**.
3. In the **Remove** dialog box, click **Yes**. A message appears in the status bar, stating that the IP address has been removed successfully.

Configuring Application Access Controls

Application access controls, also known as management access controls, form a unified mechanism for managing user authentication and implementing rules that determine user access to applications and data. You can configure MIPs and SNIPs to provide access for management applications. Management access for the NSIP is enabled by default and cannot be disabled. You can, however, control it by using ACLs.

For information about using ACLs, see "[Access Control Lists \(ACLs\)](#)."

The NetScaler appliance does not support management access to VIPs.

The following table provides a summary of the interaction between management access and specific service settings for Telnet.

Management Access	Telnet (State Configured on the NetScaler)	Telnet (Effective State at the IP Level)
Enable	Enable	Enable
Enable	Disable	Disable
Disable	Enable	Disable

Management Access	Telnet (State Configured on the NetScaler)	Telnet (Effective State at the IP Level)
Disable	Disable	Disable

The following table provides an overview of the IP addresses used as source IP addresses in outbound traffic.

Application/ IP	NSIP	MIP	SNIP	VIP
ARP	Yes	Yes	Yes	No
Server side traffic	No	Yes	Yes	No
RNAT	No	Yes	Yes	Yes
ICMP PING	Yes	Yes	Yes	No
Dynamic routing	Yes	No	Yes	Yes

The following table provides an overview of the applications available on these IP addresses.

Application/ IP	NSIP	MIP	SNIP	VIP
SNMP	Yes	Yes	Yes	No
System access	Yes	Yes	Yes	No

You can access and manage the NetScaler by using applications such as Telnet, SSH, GUI, and FTP.

Note: Telnet and FTP are disabled on the NetScaler for security reasons. To enable them, contact the customer support. After the applications are enabled, you can apply the controls at the IP level.

To configure the NetScaler to respond to these applications, you need to enable the specific management applications. If you disable management access for an IP address, existing connections that use the IP address are not terminated, but no new connections can be initiated.

Also, the non-management applications running on the underlying FreeBSD operating system are open to protocol attacks, and these applications do not take advantage of the NetScaler appliance's attack prevention capabilities.

You can block access to these non-management applications on a MIP, SNIP, or NSIP. When access is blocked, a user connecting to a NetScaler by using the MIP, SNIP, or NSIP is not be able to access the non-management applications running on the underlying operating system.

To configure management access for an IP address by using the command line interface

At the command prompt, type:

```
set ns ip <IPAddress> -mgmtAccess <value> -telnet <value> -ftp <value> -gui <value> -  
ssh <value> -snmp <value> -restrictAccess (ENABLED | DISABLED)
```

Example

```
> set ns ip 10.102.29.54 -mgmtAccess enabled -  
restrictAccess ENABLED  
Done
```

To enable management access for an IP address by using the configuration utility

1. Navigate to **System > Network > IPs > IPv4**.
2. In the details pane, select the IP address that you want to modify (for example, **10.102.29.54**), and then click **Open**.
3. In the **Configure IP** dialog box, under **Application Access Control**, select the **Enable Management Access control to support the below listed applications** check box.
4. Select the application or applications that you want to enable.
5. To block access to non-management applications on an IP address, select the **Allow access only to management applications** check box.
6. Click **OK**.

How the NetScaler Proxies Connections

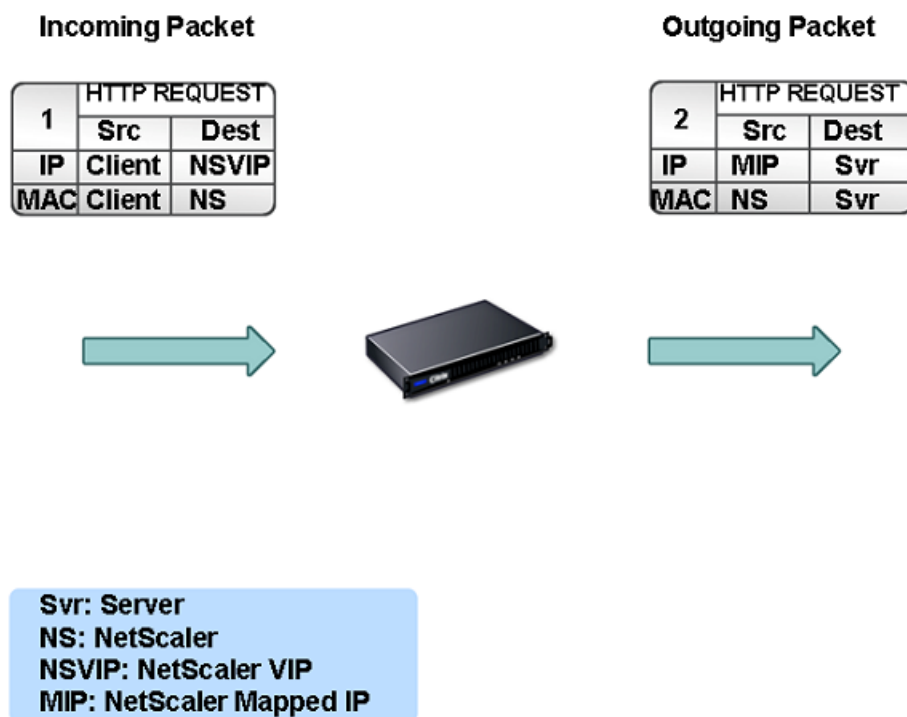
When a client initiates a connection, the NetScaler appliance terminates the client connection, initiates a connection to an appropriate server, and sends the packet to the server. The appliance does not perform this action for service type UDP or ANY.

You can configure the NetScaler to process the packet before initiating the connection with a server. The default behavior is to change the source and destination IP addresses of a packet before sending the packet to the server. You can configure the NetScaler to retain the source IP address of the packets by enabling Use Source IP mode.

How the Destination IP Address Is Selected

Traffic sent to the NetScaler appliance can be sent to a virtual server or to a service. The appliance handles traffic to virtual servers and services differently. The NetScaler terminates traffic received at a virtual server IP (VIP) address and changes the destination IP address to the IP address of the server before forwarding the traffic to the server, as shown in the following diagram.

Figure 7-4. Proxying Connections to VIPs



Packets destined for a service are sent directly to the appropriate server, and the NetScaler does not modify the destination IP addresses. In this case, the NetScaler functions as a proxy.

How the Source IP Address Is Selected

When the NetScaler appliance communicates with the physical servers or peer devices, by default, it does not use the IP address of the client. NetScaler maintains a pool of mapped IP addresses (MIPs) and subnet IP addresses (SNIPs), and selects an IP address from this pool to use as the source IP address of a connection to the physical server. Depending on the subnet in which the physical server is placed, NetScaler decides whether a MIP should be used or SNIP.

Note: If the Use Source IP (USIP) option is enabled, NetScaler uses the IP address of the client.

Enabling Use Source IP Mode

When the NetScaler appliance communicates with the physical servers or peer devices, by default, it uses one of its own IP addresses as the source IP. The appliance maintains a pool of mapped IP addresses (MIPs) and subnet IP addresses (SNIPs), and selects an IP address from this pool to use as the source IP address for a connection to the physical

server. The decision of whether to select a MIP or a SNIP depends on the subnet in which the physical server resides.

If necessary, you can configure the NetScaler appliance to use the client's IP address as source IP. Some applications need the actual IP address of the client. The following use cases are a few examples:

- ♦ Client's IP address in the web access log is used for billing purposes or usage analysis.
- ♦ Client's IP address is used to determine the country of origin of the client or the originating ISP of the client. For example, many search engines such as Google provide content relevant to the location to which the user belongs.
- ♦ The application must know the client's IP address to verify that the request is from a trustworthy source.
- ♦ Sometimes, even though an application server does not need the client's IP address, a firewall placed between the application server and the NetScaler may need the client's IP address for filtering the traffic.

Enable **Use Source IP** mode (USIP) mode if you want NetScaler to use the client's IP address for communication with the servers. By default, USIP mode is disabled. USIP mode can be enabled globally on the NetScaler or on a specific service. If you enable it globally, USIP is enabled by default for all subsequently created services. If you enable USIP for a specific service, the client's IP address is used only for the traffic directed to that service.

As an alternative to USIP mode, you have the option of inserting the client's IP address (CIP) in the request header of the server-side connection for an application server that needs the client's IP address.

In earlier NetScaler releases, USIP mode had the following source-port options for server-side connections:

- ♦ Use the client's port. With this option, connections cannot be reused. For every request from the client, a new connection is made with the physical server.
- ♦ Use proxy port. With this option, connection reuse is possible for all requests from the same client. Before NetScaler release 8.1 this option imposed a limit of 64000 concurrent connections for all server-side connections.

In the later NetScaler releases, if USIP is enabled, the default is to use a proxy port for server-side connections and not reuse connections. Not reusing connections may not affect the speed of establishing connections.

By default, the **Use Proxy Port** option is enabled if the **USIP** mode is enabled.

Note: If you enable the **USIP** mode, it is recommended to enable the **Use Proxy Port** option.

The following figure shows how the NetScaler uses IP addresses in USIP mode.

Figure 7-5. IP Addressing in USIP Mode

Recommended Usage

Enable **USIP** in the following situations:

- ♦ Load balancing of Intrusion Detection System (IDS) servers
- ♦ Stateless connection failover
- ♦ Sessionless load balancing
- ♦ If you use the Direct Server Return (DSR) mode

Note: When USIP is required in the one-arm mode installation of the NetScaler appliance, make sure that the server's gateway is one of the IP addresses owned by the NetScaler.

- ♦ If you enable USIP, set the idle timeout for server connections to a value lower than the default value, so that idle connections are cleared quickly on the server side.
- ♦ For transparent cache redirection, if you enable USIP, enable **L2CONN** also.
- ♦ Because HTTP connections are not reused when USIP is enabled, a large number of server-side connections may accumulate. Idle server connections can block connections for other clients. Therefore, set limits on maximum number of connections to a service. Citrix also recommends setting the HTTP server time-out value, for a service on which USIP is enabled, to a value lower than the default, so that idle connections are cleared quickly on the server side.

To globally enable or disable USIP mode by using the command line interface

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

- ♦ `enable ns mode usip`
- ♦ `disable ns mode usip`

To enable USIP mode for a service by using the command line interface

At the command prompt, type:

set service <ServiceName> -usip (YES | NO)

Example

```
set service Service-HTTP-1 -usip YES
```

To globally enable or disable USIP mode by using the configuration utility

1. In the navigation pane, expand **System** and click **Settings**.
2. On the **Settings** page, under **Modes and Features**, click **Configure modes**.
3. In the **Configure Modes** dialog box, do one of the following:
 - To enable Use Source IP mode, select the **Use Source IP** check box.
 - To disable Use Source IP mode, clear the **Use Source IP** check box.
4. Click **OK**.
5. In the **Enable/Disable Feature(s)?** dialog box, click **Yes**.

To enable USIP mode for a service by using the configuration utility

1. Navigate to **Traffic Management > Load Balancing > Services**.
2. In the details pane, select the service for which you want to enable the USIP mode, and then click **Open**.
3. In the **Configure Service** dialog box, click the **Advanced** tab.
4. Under **Settings**, select the **Use Source IP** check box.
5. Click **OK**.

Configuring Network Address Translation

Network address translation (NAT) involves modification of the source and/or destination IP addresses and/or the TCP/UDP port numbers of IP packets that pass through the NetScaler appliance. Enabling NAT on the appliance enhances the security of your private network, and protects it from a public network such as the Internet, by modifying your networks source IP addresses when data passes through the NetScaler. Also, with the help of NAT entries, your entire private network can be represented by a few shared public IP addresses. The NetScaler supports the following types of network address translation:

- ♦ Inbound NAT (INAT), in which the NetScaler replaces the destination IP address in the packets generated by the client with the private IP address of the server.
- ♦ Reverse NAT (RNAT), in which the NetScaler replaces the source IP address in the packets generated by the servers with the public NAT IP addresses.

Configuring INAT

When a client sends a packet to a NetScaler appliance that is configured for Inbound Network Address Translation (INAT), the appliance translates the packet's public destination IP address to a private destination IP address and forwards the packet to the server at that address.

The following configurations are supported:

- ♦ **IPv4-IPv4 Mapping:** A public IPv4 address on the NetScaler appliance listens to connection requests on behalf of a private IPv4 server. The NetScaler appliance translates the packet's public destination IP address to the destination IP address of the server and forwards the packet to the server at that address.
- ♦ **IPv4-IPv6 Mapping:** A public IPv4 address on the NetScaler appliance listens to connection requests on behalf of a private IPv6 server. The NetScaler appliance creates an IPv6 request packet with the IP address of the IPv6 server as the destination IP address.
- ♦ **IPv6-IPv4 Mapping:** A public IPv6 address on the NetScaler appliance listens to connection requests on behalf of a private IPv4 server. The NetScaler appliance creates an IPv4 request packet with the IP address of the IPv4 server as the destination IP address.
- ♦ **IPv6-IPv6 Mapping:** A public IPv6 address on the NetScaler appliance listens to connection requests on behalf of a private IPv6 server. The NetScaler appliance translates the packet's public destination IP address to the destination IP address of the server and forwards the packet to the server at that address.

When the appliance forwards a packet to a server, the source IP address assigned to the packet is determined as follows:

- ♦ If use subnet IP (USNIP) mode is enabled and use source IP (USIP) mode is disabled, the NetScaler uses a subnet IP address (SNIP) as the source IP address.
- ♦ If USNIP mode is disabled and USIP mode is disabled, the NetScaler uses a mapped IP address (MIP) as the source IP address.
- ♦ If USIP mode is enabled, and USNIP mode is disabled the NetScaler uses the client IP (CIP) address as the source IP address.
- ♦ If both USIP and USNIP modes are enabled, USIP mode takes precedence.
- ♦ You can also configure the NetScaler to use a unique IP address as the source IP address, by setting the proxyIP parameter.
- ♦ If none of the above modes are enabled and a unique IP address has not been specified, the NetScaler attempts to use a MIP as the source IP address.
- ♦ If both USIP and USNIP modes are enabled and a unique IP address has been specified, the order of precedence is as follows: USIP-unique IP-USNIP-MIP-Error.

To protect the NetScaler from DoS attacks, you can enable TCP proxy. However, if other protection mechanisms are used in your network, you may want to disable them.

You can create, modify, or remove an INAT entry.

To create an INAT entry by using the command line interface

At the command prompt, type the following commands to create an INAT entry and verify its configuration:

- ♦ **add inat** <name> <publicIP> <privateIP> [-tcpproxy (**ENABLED** | **DISABLED**)] [-ftp (**ENABLED** | **DISABLED**)] [-usip (**ON** | **OFF**)] [-usnip (**ON** | **OFF**)] [-proxypIP <ip_addr|ipv6_addr>]
- ♦ **show inat** [<name>]

Example

```
> add inat ip4-ip4 172.16.1.2 192.168.1.1 -proxypip
10.102.29.171
Done
```

To modify an INAT entry by using the command line interface

To modify an INAT entry, type the **set inat** command, the name of the entry, and the parameters to be changed, with their new values.

To remove an INAT configuration by using the command line interface

At the command prompt, type:

rm inat <name>

Example

```
> rm inat ip4-ip4
Done
```

To configure an INAT entry by using the configuration utility

1. Navigate to **System > Network > Routes**.
2. On the **Routes** page, on the **INAT** tab, do one of the following:
 - To create a new INAT entry, click **Add**.
 - To modify an existing INAT entry, select the entry, and then click **Open**.
3. In the **Create INAT** or **Configure INAT** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.

4. Click **Create** or **OK**, and then click **Close**.

To remove an INAT configuration by using the configuration utility

1. Navigate to **System > Network > Routes**.
2. On the **INAT** tab, select the name of the INAT configuration that you want to remove.
3. Click **Remove**, and then click **Close**.

Coexistence of INAT and Virtual Servers

If both INAT and RNAT are configured, the INAT rule takes precedence over the RNAT rule. If RNAT is configured with a network address translation IP (NAT IP) address, the NAT IP address is selected as the source IP address for that RNAT client.

The default public destination IP in an INAT configuration is the virtual IP (VIP) address of the NetScaler device. virtual servers also use VIPs. When both INAT and a virtual server use the same IP address, the Vserver configuration overrides the INAT configuration.

Following are a few sample configuration setup scenarios and their effects.

Case	Result
You have configured a virtual server and a service to send all data packets received on a specific NetScaler port to the server directly. You have also configured INAT and enabled TCP. Configuring INAT in this manner sends all data packets received through a TCP engine before sending them to the server.	All packets received on the NetScaler, except those received on the specified port, pass through the TCP engine.
You have configured a virtual server and a service to send all data packets of service type TCP, that are received on a specific port on the NetScaler, to the server after passing through the TCP engine. You have also configured INAT and disabled TCP. Configuring INAT in this manner sends the data packets received directly to the server.	Only packets received on the specified port pass through the TCP engine.
You have configured a virtual server and a service to send all data packets received to either of two servers. You are attempting to configure INAT to send all data packets received to a different server.	The INAT configuration is not allowed.

Case	Result
You have configured INAT to send all received data packets directly to a server. You are attempting to configure a virtual server and a service to send all data packets received to two different servers.	The vserver configuration is not allowed.

Stateless NAT46 Translation

The stateless NAT46 feature enables communication between IPv4 and IPv6 networks through IPv4 to IPv6 packet translation, and vice versa, without maintaining any session information on the NetScaler appliance.

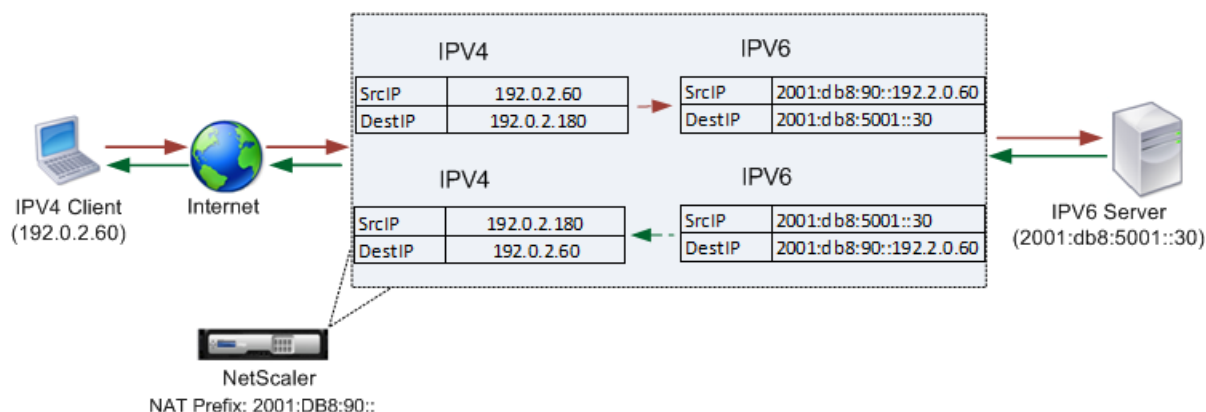
For a stateless NAT46 configuration, the appliance translates an IPv4 packet to IPv6 or an IPv6 packet to IPv4 as defined in RFCs 6145 and 2765.

Note: This feature is supported only on NetScaler 10.e and later.

A stateless NAT46 configuration on the NetScaler appliance has the following components:

- ♦ **IPv4-IPv6 INAT entry**—An INAT entry defining a 1:1 relationship between an IPv4 address and an IPv6 address. In other words, an IPv4 address on the appliance listens to connection requests on behalf of an IPv6 server. An IPv4 request packet for this IPv4 address is translated into an IPv6 packet, and then the IPv6 packet is sent to the IPv6 server.
The appliance translates an IPv6 response packet into an IPv4 response packet with its source IP address field set as the IPv4 address specified in the INAT entry. The translated packet is then sent to the client.
- ♦ **NAT46 IPv6 prefix**—A global IPv6 prefix of length 96 bits ($128-32=96$) configured on the appliance. During IPv4 packet to IPv6 packet translation, the appliance sets the source IP address of the translated IPv6 packet to a concatenation of the NAT46 IPv6 prefix [96 bits] and the IPv4 source address [32 bits] that was received in the request packet.
During IPv6 packet to IPv4 packet translation, the appliance sets the destination IP address of the translated IPv4 packet to the last 32 bits of the destination IP address of the IPv6 packet.

Consider an example in which an enterprise hosts site `www.example.com` on server S1, which has an IPv6 address. To enable communication between IPv4 clients and IPv6 server S1, NetScaler appliance NS1 is deployed with a stateless NAT46 configuration that includes an IPv4-IPv6 INAT entry for server S1, and a NAT46 Prefix. The INAT entry includes an IPv4 address at which the appliance listens to connection requests from IPv4 clients on behalf of the IPv6 server S1.



The following table lists the settings used in this example:

Entities	Name	Value
IP address of the client	Client_IPv4 (for reference purposes only)	192.0.2.60
IPv6 address of the server	Sevr_IPv6 (for reference purposes only)	2001:DB8:5001::30
IPv4 address defined in the INAT entry for IPv6 server S1	Map-Sevr-IPv4 (for reference purposes only)	192.0.2.180
IPv6 prefix for NAT 46 translation	NAT46_Prefix (for reference purposes only)	2001:DB8:90::

Following is the traffic flow in this example:

1. IPv4 Client CL1 sends a request packet to the Map-Sevr-IPv4 (192.0.2.180) address on the NetScaler appliance.
2. The appliance receives the request packet and searches the NAT46 INAT entries for the IPv6 address mapped to the Map-sevr-IPv4 (192.0.2.180) address. It finds the Sevr-IPv6 (2001:DB8:5001::30) address.
3. The appliance creates a translated IPv6 request packet with:
 - Destination IP address field = Sevr-IPv6 = 2001:DB8:5001::30
 - Source IP address field = Concatenation of NAT Prefix (First 96 bits) and Client_IPv4 (last 32 bits) = 2001:DB8:90::192.0.2.60
4. The appliance sends the translated IPv6 request to Sevr-IPv6.
5. The IPv6 server S1 responds by sending an IPv6 packet to the NetScaler appliance with:
 - Destination IP address field = Concatenation of NAT Prefix (First 96 bits) and Client_IPv4 (last 32 bits) = 2001:DB8:90::192.0.2.60

- Source IP address field = Sevr-IPv6 = 2001:DB8:5001::30
- 6. The appliance receives the IPv6 response packet and verifies that its destination IP address matches the NAT46 prefix configured on the appliance. Because the destination address matches the NAT46 prefix, the appliance searches the NAT46 INAT entries for the IPv4 address associated with the Sevr-IPv6 address (2001:DB8:5001::30). It finds the Map-Sevr-IPv4 address (192.0.2.180).
- 7. The appliance creates an IPv4 response packet with:
 - Destination IP address field = The NAT46 prefix stripped from the destination address of the IPv6 response = Client_IPv4 (192.0.2.60)
 - Source IP address field = Map-Sevr-IPv4 address (192.0.2.180)
- 8. The appliance sends the translated IPv4 response to client CL1.

Configuring Stateless NAT46

Creating the required entities for stateless NAT46 configuration on the NetScaler appliance involves the following procedures:

1. Create an IPv4-IPv6 mapping INAT entry with stateless mode enabled.
2. Add a NAT46 IPv6 prefix.

To configure an INAT mapping entry by using the command line interface

At the command prompt, type:

- ♦ **add inat** <name> <publicIPv4> <privateIPv6> -mode STATELESS
- ♦ **show inat** <name>

To add an NAT46 prefix by using the command line interface

At the command prompt, type:

- ♦ **set inatparam** -nat46v6Prefix <ipv6_addr|*>
- ♦ **show inatparam**

Example

```
> add inat exmpl-com-stls-nat46 192.0.2.180
2001:DB8:5001::30 -mode stateless
Done

> set inatparam -nat46v6Prefix 2001:DB8:90::/96
Done
```

To configure an INAT mapping entry by using the configuration utility

1. Navigate to **System > Network > Routes**.
2. In the details pane, on the **INAT** tab, do one of the following:
 - To create a new INAT entry, click **Add**.

- To modify an existing INAT entry, select the entry, and then click **Open**.
3. In the **Create INAT** or **Configure INAT** dialog box, set the following parameters:
 - Name*
 - Public IP Address*
 - Private IP Address* (Select the IPv6 check box and enter the address in IPv6 format.)
 - Mode (Select **Stateless** from the drop down list.)
- * A required parameter
4. Click **Create** or **OK**, and then click **Close**.

To add a NAT46 prefix by using the configuration utility

1. Navigate to **System > Network**.
2. In the details pane, under **Settings**, click **Configure INAT Parameters**.
3. In the **Configure INAT Parameters** dialog box, set the **NAT46 Prefix** parameter.
4. Click **OK**.

Setting Global Parameters for Stateless NAT46

The appliance provides some optional global parameters for stateless NAT46 configurations.

To set global parameters for stateless NAT46 by using the command line interface

At the command prompt, type:

- ♦ **set inatparam** [-nat46IgnoreTOS (YES | NO)] [-nat46ZeroChecksum (ENABLED | DISABLED)] [-nat46v6Mtu <positive_integer>] [-nat46FragHeader (ENABLED | DISABLED)]
- ♦ **show inatparam**

Example

```
> set inatparam -nat46IgnoreTOS YES -nat46ZeroChecksum
DISABLED -nat46v6Mtu 1400 -nat46FragHeader DISABLED
Done
```

To set global parameters for stateless NAT46 by using the configuration utility

1. Navigate to **System > Network**.
2. In the details pane, under **Settings**, click **Configure INAT Parameters**.
3. In the **Configure INAT Parameters** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.

4. Click **OK**.

Limitations of Stateless NAT46

The following limitations apply to stateless NAT46:

- ◆ Translation of IPv4 options is not supported.
- ◆ Translation of IPv6 routing headers is not supported.
- ◆ Translation of hop-by-hop extension headers of IPv6 packets is not supported.
- ◆ Translation of ESP and EH headers of IPv4 packets is not supported.
- ◆ Translation of multicast packets is not supported.
- ◆ Translation of destination option headers and source routing headers is not supported.
- ◆ Translation of fragmented IPv4 UDP packets that do not contain UDP checksum is not supported.

DNS64

The NetScalerDNS64 feature responds with a synthesized DNS AAAA record to an IPv6 client sending an AAAA request for an IPv4-only domain. The DNS64 feature is used with the NAT64 feature to enable seamless communication between IPv6-only clients and IPv4-only servers. DNS64 enables discovery of the IPv4 domain by the IPv6 only clients, and NAT64 enables communication between the clients and servers.

For synthesizing an AAAA record, the NetScaler appliance fetches a DNS A record from a DNS server. The DNS64 prefix is a 96-bit IPv6 prefix configured on the NetScaler appliance. The NetScaler appliance synthesizes the AAAA record by concatenation of the DNS64 Prefix (96 bits) and the IPv4 address (32 bits).

For enabling communication between IPv6 clients and IPv4 servers, a NetScaler appliance with DNS64 and NAT64 configuration can be deployed either on the IPv6 client side or on the IPv4 server side. In both cases, the DNS64 configuration on the NetScaler appliance is similar and includes a load balancing virtual server acting as a proxy server for DNS servers. If the NetScaler appliance is deployed on the client side, the load balancing virtual server must be specified, on the IPv6 client, as the nameserver for a domain.

Consider an example where a NetScaler appliance with DNS64 and NAT64 configuration is configured on the IPv4 side. In this example, an enterprise hosts site `www.example.com` on server S1, which has an IPv4 address. To enable communication between IPv6 clients and IPv4 server S1, NetScaler appliance NS1 is deployed with a DNS64 and stateful NAT64 configuration.

The DNS64 configuration includes DNS load balancing virtual server LBVS-DNS64-1, on which the DNS64 option is enabled. A DNS64 policy named DNS64-Policy-1, and an associated DNS64 action named DNS64-Action-1, are also configured on NS1, and DNS64-Policy-1 is bound to LBVS-DNS64-1. LBVS-DNS64-1 acts as a DNS proxy server for DNS servers DNS-1 and DNS-2.

When traffic arriving at LBVS-DNS64-1 matches the conditions specified in DNS64-Policy-1, the traffic is processed according to the settings in DNS64-Action-1. DNS64-Action-1 specifies the DNS64 prefix used, with the A record received from a DNS server, to synthesize an AAAA record.

The global DNS parameter cacherecords is enabled on the NetScaler appliance, so the appliance caches DNS records. This setting is necessary for the DNS64 to work properly.

The following table lists the settings used in the above example:

Entity	Name	Value
IPv6 client	CL1 (for reference purposes only)	<ul style="list-style-type: none"> IP address = 2001:DB8:5001::30
DNS64 Prefix		<ul style="list-style-type: none"> 2001:DB8:300::
Service on NS representing DNS server DNS-1	SVC-DNS-1	<ul style="list-style-type: none"> IP address = 203.0.113.50 Port = 53
Service on NS representing DNS server DNS-2	SVC-DNS-2	<ul style="list-style-type: none"> IP address = 203.0.113.60 Port = 53
DNS64 action	DNS64-Action-1	<ul style="list-style-type: none"> DNS64 Prefix=2001:DB8:300::
DNS64 policy	DNS64-Policy-1	<ul style="list-style-type: none"> DNS64 action = DNS64-Action-1 Rule= CLIENT.IP.SRC.IN_SUBNET(2001:DB8:5001::/64)
DNS load balancing virtual server	LBVS-DNS64-1	<ul style="list-style-type: none"> IP address=2001:DB8:9999::99 Bound DNS services= SVC-DNS-1, SVC-DNS-2 DNS64=Enabled Bound DNS64 policy= DNS64-Policy-1

Following is the traffic flow in this example:

1. IPv6 client CL1 sends a DNS AAAA request for the IPv6 address of the site `www.example.com`.
2. The request is received by the DNS load balancing virtual server LBVS-DNS64-1 on NetScaler appliance NS1.
3. NS1 checks its DNS cache records for the requested AAAA record and finds that AAAA record for the site `www.example.com` does not exist in the DNS cache.
4. LBVS-DNS64-1's load balancing algorithm selects DNS server DNS-1 and forwards the AAAA request to it.
5. Because the site `www.example.com` is hosted on an IPv4 server, the DNS server DNS-1 does not have any AAAA record for the site `www.example.com`.
6. DNS-1 sends either an empty DNS AAAA response or an error message to LBVS-DNS64-1.
7. Because DNS64 option is enabled on LBVS-DNS64-1 and the AAAA request from CL1 matches the condition specified in DNS64-Policy-1, NS1 sends a DNS A request to DNS-1 for the IPv4 address of `www.example.com`.
8. DNS-1 responds by sending the DNS A record for `www.example.com` to LBVS-DNS64-1. The A record includes the IPv4 address for `www.example.com`.
9. NS1 synthesizes an AAAA record for the site `www.example.com` with:
 - IPv6 address for site `www.example.com` = Concatenation of DNS64 Prefix (96 bits) specified in the associated DNS64action, and IPv4 address of DNS A record (32 bits) = `2001:DB8:300::192.0.2.60`
10. NS1 sends the synthesized AAAA record to IPv6 client CL1. NS1 also caches the A record into its memory. NS1 uses the cached A record to synthesize AAAA records for subsequent AAAA requests.

Points to Consider for a DNS64 Configuration

Before configuring DNS64 on a NetScaler appliance, consider the following points:

- ♦ The DNS64 feature of the NetScaler appliance is compliant with RFC 6174.
- ♦ The DNS64 feature of the NetScaler appliance does not support DNSSEC. The NetScaler appliance does not synthesize an AAAA record from a DNSSEC response received from a DNS server. A response is classified as a DNSSEC response, only if it contains RRSIG records.
- ♦ The NetScaler appliance supports DNS64 prefix of length of only 96 bits.
- ♦ Though the DNS64 feature is used with the NAT64 feature, the DNS64 and NAT64 configurations are independent on the NetScaler appliance. For a particular flow, you must specify the same IPv6 prefix value for the DNS64 prefix and the NAT64 prefix parameters, so that the synthesized IPv6 addresses received by the client are routed to the particular NAT64 configuration. For more information on configuring NAT64 on a NetScaler appliance, See "[Stateful NAT64](#)".
- ♦ The following are the different cases of DN64 processing by the NetScaler appliance:
 - If the AAAA response from the DNS server includes AAAA records, then each record in the response is checked for the set of exclusion rule configured on the

NetScaler appliance for the particular DNS64 configuration. The NetScaler removes the IPv6 addresses, whose prefix matches the exclusion rule, from the response. If the resulting response includes at least one IPv6 record, the NetScaler appliance forwards this response to the client, else, the appliance synthesizes a AAAA response from the A record of the domain and sends it to the IPv6 client.

- If the AAAA response from the DNS server is an empty answer response, the appliance requests for A resource records with the same domain name or searches in its own records if the appliance is an authentic domain name server for the domain. If the request results in an empty answer or error, the same is forwarded to the client.
- If the response from the DNS server includes RCODE=1 (format error), the NetScaler appliance forwards the same to the client. If there is no response before the timeout, the NetScaler appliance sends a response with RCODE=2 (server failure) to the client.
- If the response from the DNS server includes a CNAME, the chain is followed until the terminating A or AAAA record is reached. If the CNAME does not have any AAAA resource records, the NetScaler appliance fetches the DNS A record to be used for synthesizing AAAA record. The CNAME chain is added to the answer section along with the synthesized AAAA record and then sent to the client.
- ♦ The DNS64 feature of the NetScaler appliance also supports responding to PTR request. When a PTR request for a domain of an IPv6 address is received on the appliance and the IPv6 address matches any of the configured DNS64 prefix, the appliance creates a CNAME record mapping the IP6-ARPA domain into the corresponding IN-ADDR.ARPA domain and the newly formed IN-ADDR.ARPA domain is used for resolution. The appliance searches the local PTR records and if the records are not present, the appliance sends a PTR request for IN-ADDR.ARPA domain to the DNS server. The NetScaler appliance uses the response from the DNS server to synthesize response for the initial PTR request.

Configuration Steps

Creating the required entities for stateful NAT64 configuration on the NetScaler appliance involves the following procedures:

- ♦ **Add DNS services.** DNS services are logical representation of DNS servers for which the NetScaler appliance acts as a DNS proxy server.
- ♦ **Add DNS64 action and DNS64 policy and then bind the DNS64 action to the DNS64 policy.** A DNS64 policy specifies conditions to be matched against traffic for DNS64 processing according to the settings in the associated DNS64 action. The DNS64 action specifies the mandatory DNS64 prefix and the optional exclude rule and mapped rule settings.
- ♦ **Create a DNS load balancing virtual server and bind the DNS services and the DNS64 policy to it.** The DNS load balancing virtual server acts as a DNS proxy server for DNS servers represented by the bound DNS services. Traffic arriving at the virtual server is matched against the bound DNS64 policy for DNS64 processing.

Note: The command line interface has separate commands for these two tasks, but the configuration utility combines them in a single dialog box.

- ♦ **Enable caching of DNS records.** Enable the global parameter for the NetScaler appliance to cache DNS records, which are obtained through DNS proxy operations.

To create a service of type DNS by using the command line interface

At the command prompt, type:

- ♦ **add service** <name> <IP> <serviceType> <port> ...

To create a DNS64 action by using the command line interface

At the command prompt, type:

- ♦ **add dns action64** <actionName> -Prefix <ipv6_addr|*> [-mappedRule <expression>] [-excludeRule <expression>]

To create a DNS64 policy by using the command line interface

At the command prompt, type:

- ♦ **add dns policy64** <name> -rule <expression> -action <string>

To create a DNS load balancing virtual server by using the command line interface

At the command prompt, type:

- ♦ **add lb vserver** <name> DNS <IPAddress> <port> -dns64 (ENABLED | DISABLED) [-bypassAAAA (YES | NO)] ...

To bind the DNS services and the DNS64 policy to the DNS load balancing virtual server by using the command line interface

At the command prompt, type:

- ♦ **bind lb vserver** <name> <serviceName> ...
- ♦ **bind lb vserver** <name> -policyName <string> -priority <positive_integer> ...

Example

```
> add service SVC-DNS-1 203.0.113.50 DNS 53
Done

> add service SVC-DNS-2 203.0.113.60 DNS 53
Done

> add dns Action64 DNS64-Action-1 -Prefix 2001:DB8:300::/96
Done

> add dns Policy64 DNS64-Policy-1 -rule
"CLIENT.IPv6.SRC.IN_SUBNET(2001:DB8:5001::/64)"
-action DNS64-Action-1
Done

> add lb vserver LBVS-DNS64-1 DNS 2001:DB8:9999::99 53 -dns64
ENABLED
Done
```

```
> bind lb vserver LBVS-DNS64-1 SVC-DNS-1
Done

> bind lb vserver LBVS-DNS64-1 SVC-DNS-2
Done

> bind lb vserver LBVS-DNS64-1 -policyname DNS64-Policy-1 -
priority 2
Done
```

To create a service of type DNS by using the configuration utility

1. Navigate to **Traffic Management > Load Balancing > Services**.
2. In the details pane, click **Add**.
3. In the **Create Service** dialog box, set the following parameters:
 - Service Name*
 - Server*
 - Protocol*
 - Port*
4. Click **Create**.
5. Repeat steps 3-4 to create another service.
6. Click **Close**.

To create a DNS64 action by using the command line interface

1. Navigate to **Traffic Management > DNS > Actions**.
2. On the **DNS Actions V6** page, click **Add**.
3. In the **Create DNS Action64** dialog box, set the following parameters:
 - Name*
 - Prefix*
 - Mapped Rule
 - Exclude Rule
4. Click **create**, and then **Close**.

To create a DNS64 policy by using the command line interface

1. Navigate to **Traffic Management > DNS > Policies**.
2. On the **DNS Policies V6** page, click **Add**.

3. In the **Create DNS Policy64** dialog box, set the following parameters:

- Name*
- Action*
- Rule*

4. Click **create**, and then **Close**.

To create a DNS load balancing virtual server and bind the DNS services and the DNS64 policy to it by using the configuration utility

1. Navigate to **Traffic Management > Load Balancing > Virtual Servers**.

2. In the details pane, click **Add**.

3. In the **Create Virtual Servers (Load Balancing)** dialog box, set the following parameters:

- Name*
- IP Address*
- Protocol*
- Port*

4. Select the **Enable DNS64** check box.

5. Under the **Services** tab, in the **Active** column, select the check box for the service that you want to bind to the virtual server.

6. Under the **Policies** tab, click **Insert Policy** to insert a new row and display a drop-down list of all unbound policies.

7. From the drop-down list that appears under **Policy Name**, select the DNS64 policy that you want to bind to this virtual server.

8. Click **create**, and then **Close**.

Stateful NAT64 Translation

The stateful NAT64 feature enables communication between IPv4 clients and IPv6 servers through IPv6 to IPv4 packet translation, and vice versa, while maintaining session information on the NetScaler appliance.

A stateful NAT64 configuration on the NetScaler appliance has the following components:

- ♦ **NAT64 rule**— An entry consisting of an ACL6 rule and a netprofile, which consists of a pool of NetScaler owned SNIP Addresses.
- ♦ **NAT64 IPv6 Prefix**— A global IPv6 prefix of length 96 bits (128-32=96) configured on the appliance.

Note: Currently the NetScaler appliance supports only one prefix to be used commonly with all NAT 64 rules.

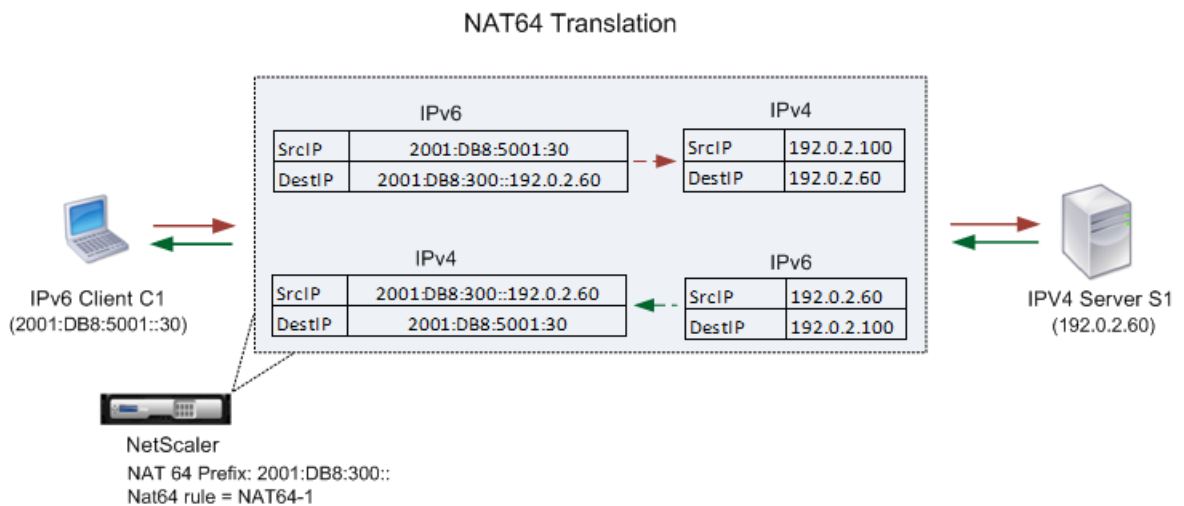
The NetScaler appliance considers an incoming IPv6 packet for NAT64 translation when all of the following conditions are met:

- ♦ The incoming IPv6 packet matches the ACL6 rule bound to a NAT64 rule.
- ♦ The destination IP address of the IPv6 packet matches the NAT64 IPv6 prefix.

When an IPv6 request packet received by the NetScaler appliance matches an ACL6 defined in a NAT64 rule and the destination IP of the packet matches the NAT64 IPv6 prefix, the NetScaler appliance considers the IPv6 packet for translation.

The appliance translates this IPv6 packet to an IPv4 packet with a source IP address matching one of the IP address bound to the netprofile defined in the NAT64 rule, and a destination IP address consisting of the last 32 bits of the destination IPv6 address of the IPv6 request packet. The NetScaler appliance creates a NAT64 session for this particular flow and forwards the packet to the IPv4 server. Subsequent responses from the IPv4 server and requests from the IPv6 client are translated accordingly by the appliance, on the basis of information in the particular NAT64 session.

Consider an example in which an enterprise hosts site `www.example.com` on server S1, which has an IPv4 address. To enable communication between IPv6 clients and IPv4 server S1, NetScaler appliance NS1 is deployed with a stateful NAT64 configuration that includes a NAT64 rule and a NAT64 prefix. A mapped IPv6 address of server S1 is formed by concatenating the NAT64 IPv6 prefix [96 bits] and the IPv4 source address [32 bits]. This mapped IPv6 address is then manually configured in the DNS servers. The IPv6 clients get the mapped IPv6 address from the DNS servers to communicate with IPv4 server S1.



The following table lists the settings used in this example:

Entities	Name	Value
IPv6 address of client CL1	Client_IPv6 (for reference purposes only)	2001:DB8:5001::30

IPv4 address of server S1	Sevr_IPv4 (for reference purposes only)	192.0.2.60
IPv6 prefix for NAT64 translation	NAT64_Prefix (for reference purposes only)	2001:DB8:300::
Mapped IPv6 address (NAT64_Prefix + Sevr_IPv4) of server S1 for IPv6 clients to reach server S1	Map-Sevr-IPv6 (for reference purposes only)	2001:DB8:300::192.0.2.60
ACL6 rule	ACL6-1	<ul style="list-style-type: none"> ♦ Action = ALLOW ♦ Source IP address = 2001:DB8:5001::30
IPset	IPset-1	IP addresses bound (of type SNIPs) = 192.0.2.100 and 192.0.2.102
Netprofile	Netprofile-1	Source IP address = IPset-1
NAT64 rule	NAT64-1	ACL6 rule = ACL6-1 Netprofile = Netprofile-1

Following is the traffic flow in this example:

1. IPv6 client CL1 sends a request packet to Map-Sevr-IPv6 (2001:DB8:300::192.0.2.60) address.
2. The NetScaler appliance receives the request packet. If the request packet matches the ACL6 defined in the NAT64 rule, and the destination IP address of the packet matches the NAT64 IPv6 prefix, the NetScaler considers the IPv6 packet for translation.
3. The appliance creates a translated IPv4 request packet with:
 - Destination IP address field containing the NAT64 prefix stripped from the destination address of the IPv6 request (Sevr_IPv4 = 192.0.2.60)
 - Source IP address field containing one of the IPv4 address bound to Netprofile-1 (in this case, 192.0.2.100)
4. The NetScaler appliance creates a NAT64 session for this flow and sends the translated IPv4 request to server S1.
5. IPv6 server S1 responds by sending an IPv4 packet to the NetScaler appliance with:
 - Destination IP address field containing 192.0.2.100
 - Source IP address field containing the address of Sevr_IPv4 (192.0.2.60)
6. The appliance receives the IPv4 response packet, searches all the session entries, and finds that the IPv6 response packet matches the NAT64 session entry created in step 4. The appliance considers the IPv4 packet for translation.

7. The appliance creates a translated IPv6 response packet with:
 - Destination IP address field=Client_IPv6=2001:DB8:5001::30
 - Source IP address field = Concatenation of NAT64 Prefix (First 96 bits) and Sevr_IPv4 (last 32 bits) =2001:DB8:300::192.0.2.60
8. The appliance sends the translated IPv6 response to client CL1.

Limitations of Stateful NAT64

The following limitations apply to stateful NAT64 translation:

- ♦ Translation of IPv4 options is not supported.
- ♦ Translation of IPv6 routing headers is not supported.
- ♦ Translation of hop-by-hop extension headers of IPv6 packets is not supported.
- ♦ Translation of ESP and EH headers of IPv6 packets is not supported.
- ♦ Translation of multicast packets is not supported.
- ♦ Packets of Stream Control Transmission Protocol (SCTP), Datagram Congestion Control Protocol (DCCP), and IPSec, are not translated.

Configuring Stateful NAT64

Creating the required entities for stateful NAT64 configuration on the NetScaler appliance involves the following procedures:

1. Add an ACL6 rule with action ALLOW.
2. Add an ipset, which binds multiple IP addresses.
3. Add a netprofile and bind the ipset to it. If you want to bind only one IP address, you need not create an ipset entity. In that case, bind the IP address directly to the netprofile.
4. Add a NAT64 rule, which includes binding the ACL6 rule and the netprofile to the NAT 64 rule.
5. Add a NAT64 IPv6 prefix.

To add an ACL6 rule by using the command line interface

At the command prompt, type:

- ♦ **add ns acl6** <acl6name> <acl6action> ...

To add an IPset and bind multiple IPs to it by using the command line interface

At the command prompt, type:

- ♦ **add ipset** <name>
- ♦ **bind ipset** <name> <IPaddress ...>

To add a netprofile by using the command line interface

At the command prompt, type:

- ♦ **add netprofile** <name> -srcIP <IPAddress or IPset>

To add a NAT64 rule by using the command line interface

At the command prompt, type:

- ♦ **add nat64** <name> <acl6name> -netProfile <string>

To add a NAT64 prefix by using the command line interface

At the command prompt, type:

- ♦ **set ipv6** -natprefix <ipv6_addr|*>

Example

```
> add acl6 ACL6-1 ALLOW -srcIPv6 2001:DB8:5001::30
Done

> apply acls6
Done

> add ip 192.0.2.100 255.255.255.0 -type SNIP
Done

> add ip 192.0.2.102 255.255.255.0 -type SNIP
Done

> add ipset IPset-1
Done

> bind ipset IPset-1 192.0.2.100 192.0.2.102
IPAddress "192.0.2.100" bound
IPAddress "192.0.2.102" bound
Done

> add netprofile Netprofile-1 -srcIP IPset-1
Done

> add nat64 NAT64-1 ACL6-1 -netprofile Netprofile-1
Done

> set ipv6 -natprefix 2001:DB8:300::/96
Done
```

To add a NAT64 rule by using the configuration utility

1. Navigate to **System > Network > Routes**.
2. On the **Routes** page, click the **NAT64** tab.
3. In the **Create NAT64** dialog box, set the following parameters:
 - Name
 - ACL6 Name

- Net Profile
4. Click **OK**.

To add a NAT64 prefix by using the configuration utility

1. Navigate to **System > Network**.
2. In the details pane, in the **Settings** group, click **Change IPv6 Settings**.
3. In the **Configure Configuration for IPv6** dialog box, set the following parameter:
 - IPv6 NAT prefix
4. Click **OK**.

Configuring RNAT

In Reverse Network Address Translation (RNAT), the NetScaler appliance replaces the source IP addresses in the packets generated by the servers with public NAT IP addresses. By default, the appliance uses a Mapped IP address (MIP) as the NAT IP address. You can also configure the appliance to use a unique NAT IP address for each subnet. You can also configure RNAT by using Access Control Lists (ACLs). Use Source IP (USIP), Use Subnet IP (USNIP), and Link Load Balancing (LLB) modes affect the operation of RNAT. You can display statistics to monitor RNAT.

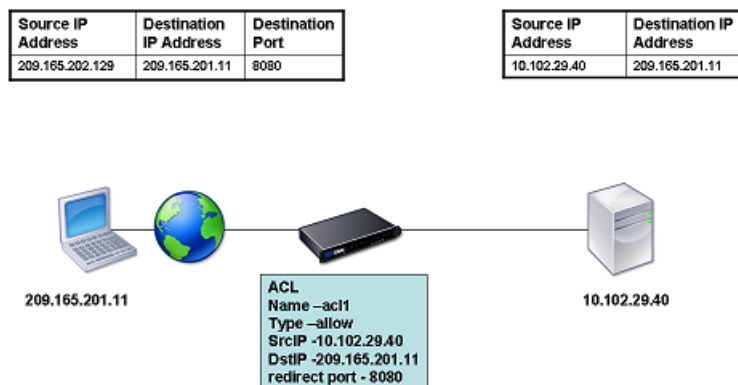
Note: The ephemeral port range for RNAT on the NetScaler appliance is 1024-65535.

You can use either a network address or an extended ACL as the condition for an RNAT entry:

- ♦ **Using a Network address.** When you use a network address, RNAT processing is performed on all of the packets coming from the specified network.
- ♦ **Using Extended ACLs.** When you use ACLs, RNAT processing is performed on all packets that match the ACLs. To configure the NetScaler appliance to use a unique IP address for traffic that matches an ACL, you must perform the following three tasks:
 - a. Configure the ACL.
 - b. Configure RNAT to change the source IP address and Destination Port.
 - c. Apply the ACL.

The following diagram illustrates RNAT configured with an ACL.

Figure 7-6. RNAT with an ACL

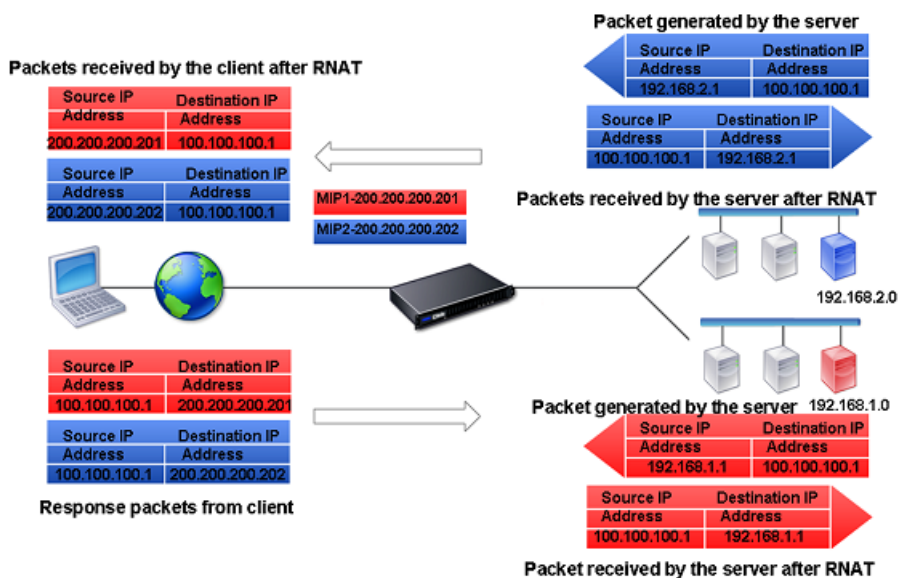


You have the following basic choices for the type of NAT IP address:

- ♦ **Using a MIP or SNIP as the NAT IP Address.** When using a MIP as the NAT IP address, the NetScaler appliance replaces the source IP addresses of server-generated packets with the a MIP. Therefore, the MIP address must be a public IP address. If Use Subnet IP (USNIP) mode is enabled, the NetScaler can use a subnet IP address (SNIP) as the NAT IP address.
- ♦ **Using a Unique IP Address as the NAT IP Address.** When using a unique IP address as the NAT IP address, the NetScaler appliance replaces the source IP addresses of server-generated packets with the unique IP address specified. The unique IP address must be a public NetScaler-owned IP address. If multiple NAT IP addresses are configured for a subnet, NAT IP selection uses the round robin algorithm.

This configuration is illustrated in the following diagram.

Figure 7-7. Using a Unique IP Address as the NAT IP Address



Creating an RNAT Entry

The following instructions provide separate command-line procedures for creating RNAT entries that use different conditions and different types of NAT IP addresses. In the configuration utility, all of the variations can be configured in the same dialog box, so there is only one procedure for configuration utility users.

To create an RNAT entry by using the command line interface

At the command prompt, type one the following commands to create, respectively, an RNAT entry that uses a network address as the condition and a MIP or SNIP as the NAT IP address, an RNAT entry that uses a network address as the condition and a unique IP address as the NAT IP address, an RNAT entry that uses an ACL as the condition and a MIP or SNIP as the NAT IP address, or an RNAT entry that uses an ACL as a condition and a unique IP address as the NAT IP address:

- ♦ **set rnat** <IPAddress> <netmask>
- ♦ **set rnat** IPAddress <netMask> -natip <NATIPAddress>
- ♦ **set rnat** <aclname> [-redirectPort <port>]
- ♦ **set rnat** <aclname> [-redirectPort <port>] -natIP <NATIPAddress>

Use the following command to verify the configuration:

- ♦ **show rnat**

Examples

A network address as the condition and a MIP or SNIP as the NAT IP address:

```
> set rnat 192.168.1.0 255.255.255.0
Done
```

A network address as the condition and a unique IP address as the NAT IP address:

```
> set rnat 192.168.1.0 255.255.255.0 -natip
10.102.29.50
Done
```

If instead of a single NAT IP address you specify a range, RNAT entries are created with all the NetScaler-owned IP addresses, except the NSIP, that fall within the range specified:

```
> set rnat 192.168.1.0 255.255.255.0 -natIP
10.102.29.[50-110]
Done
```

An ACL as the condition and a MIP or SNIP as the NAT IP address:

```
> set rnat acl1
Done

An ACL as a condition and a unique IP address as
the NAT IP address:

> set rnat acl1 -natIP 209.165.202.129
Done

If instead of a single NAT IP address you specify
a range, RNAT entries are created with all the
NetScaler-owned IP addresses, except the NSIP,
that fall within the range specified:

> set rnat acl1 -natIP 10.102.29.[50-70]
Done
```

To create an RNAT entry by using the configuration utility

1. Navigate to **System > Network > Routes**.
2. On the **Routes** page, click the **RNAT** tab.
3. In the details pane, click **Configure RNAT**.
4. In the **Configure RNAT** dialog box, do one of the following:
 - If you want to use the network address as a condition for creating an RNAT entry, click **Network** and set the following parameters:
 - ♦ **Network**
 - ♦ **Netmask**
 - If you want to use an extended ACL as a condition for creating an RNAT entry, click **ACL** and set the following parameters:
 - ♦ **ACL Name**
 - ♦ **Redirect Port**
5. To set a MIP or SNIP as a NAT IP, jump to Step 7.
6. To set a unique IP address as a NAT IP, in the **Available NAT IP (s)** list, select the IP address that you want to set as the NAT IP, and then click **Add**. The NAT IP you selected appears in the **Configured NAT IP(s)** list.
7. Click **Create**, and then **Close**.

Monitoring RNAT

You can display RNAT statistics to troubleshoot issues related to IP address translation.

To view RNAT statistics by using the command line interface

At the command prompt, type:

stat rnat**Example**

```

> stat rnat

RNAT summary

s)              Total              Rate (/
Bytes Received      0
Bytes Sent          0
Packets Received    0
Packets Sent        0
Syn Sent            0
Current RNAT sessions
--                  0
Done
>

```

The following tables describes the statistics associated with RNAT and RNAT IP.

Table 7-2. RNAT Statistics

Statistic	Description
Bytes received	Bytes received during RNAT sessions
Bytes sent	Bytes sent during RNAT sessions
Packets received	Packets received during RNAT sessions
Packets sent	Packets sent during RNAT sessions
Syn sent	Requests for connections sent during RNAT sessions
Current sessions	Currently active RNAT sessions

To monitor RNAT by using the configuration utility

1. Navigate to **System > Network > Routes**.
2. In the details pane, on the **RNAT** tab, click **Statistics**. The **Statistics** dialog box appears, displaying the RNAT statistics.

RNAT in USIP, USNIP, and LLB Modes

When RNAT and Use Source IP (USIP) are both configured, RNAT takes precedence. When RNAT and USNIP are configured, selection of the source IP address is based on the state of USNIP, as follows:

- ♦ If USNIP is off, the NetScaler appliance uses the mapped IP addresses.
- ♦ If USNIP is on, the NetScaler uses a SNIP as the NAT IP address.

This behavior does not apply when a unique NAT IP address is used.

In a topology where the NetScaler appliance performs both Link Load Balancing (LLB) and RNAT for traffic originating from the server, the appliance selects the source IP address based on the router. The LLB configuration determines selection of the router.

Configuring RNAT for IPv6 Traffic

Reverse Network Address Translation (RNAT) rules for IPv6 packets are called *RNAT6s*. When an IPv6 packet generated by a server matches the conditions specified in the RNAT6 rule, the appliance replaces the source IPv6 address of the IPv6 packet with a configured NAT IPv6 address before forwarding it to the destination. The NAT IPv6 address is one of the NetScaler owned SNIP6 or VIP6 addresses.

When configuring an RNAT6 rule, you can specify either an IPv6 prefix or an ACL6 as the condition:

- ♦ **Using a IPv6 network address.** When you use an IPv6 prefix, the appliance performs RNAT processing on those IPv6 packets whose IPv6 address matches the prefix.
- ♦ **Using ACL6s.** When you use an ACL6, the appliance performs RNAT processing on those IPv6 packets that match the conditions specified in the ACL6.

You have one of the following options to set the NAT IP address:

- ♦ Specify a set of NetScaler owned SNIP6 and VIP6 addresses for an RNAT6 rule. The NetScaler appliance uses any one of the IPv6 addresses from this set as a NAT IP address for each session. The selection is based on the round robin algorithm and is done for each session.
- ♦ Do not specify any NetScaler owned SNIP6 or VIP6 address for an RNAT6 rule. The NetScaler appliance uses any one of the NetScaler owned SNIP6 or VIP6 addresses as a NAT IP address. The selection is based on the next hop network to which an IPv6 packet that matches the RNAT rule is destined.

To create an RNAT6 rule by using the command line interface

At the command prompt, to create the rule and verify the configuration, type:

- ♦ **add rnat6** <name> (<network> | (<acl6name> [-redirectPort <port>]))
- ♦ **bind rnat6** <name> <natIP6> ...
- ♦ **show rnat6**

To modify or remove an RNAT6 rule by using the command line interface

- ♦ To modify an RNAT6 rule whose condition is an ACL6, type the **set rnat6** <name> command, followed by a new value for the redirectPort parameter.
- ♦ To remove an RNAT6 rule, type the **clear rnat6** <name> command.
- ♦ **show rnat6**

To configure an RNAT6 rule by using the configuration utility

1. Navigate to **System > Network > Routes**.
2. On the **Routes** page, click the **RNAT6** tab.
3. In the details pane, do one of the following:
 - To create an RNAT6 rule, click **Add**.
 - To modify an RNAT6 rule whose condition is an ACL6, select the RNAT6 rule, and then click **Open**.
4. In the **Create RNAT6** dialog box, do one of the following:
 - If you want to use an IPv6 network address as the condition for an RNAT6 rule, click **Network** and set the **Network** parameter.
 - If you want to use an ACL6 as the condition for an RNAT6 rule, click **ACL**, and then set the **ACL6 Name** parameter and, optionally, the **Redirect Port** parameter.
5. In the **Configure RNAT6** dialog box, set the **Redirect Port** parameter.
6. To set a NetScaler owned IPv6 address as a NAT IP address, in the **Available** list, select the IPv6 address that you want to set as the NAT IP address, and then click **+**. The NAT IP address you selected appears in the **Configured** list.
7. Click **Create** or **OK**, and then **Close**.

Configuring Prefix-Based IPv6-IPv4 Translation

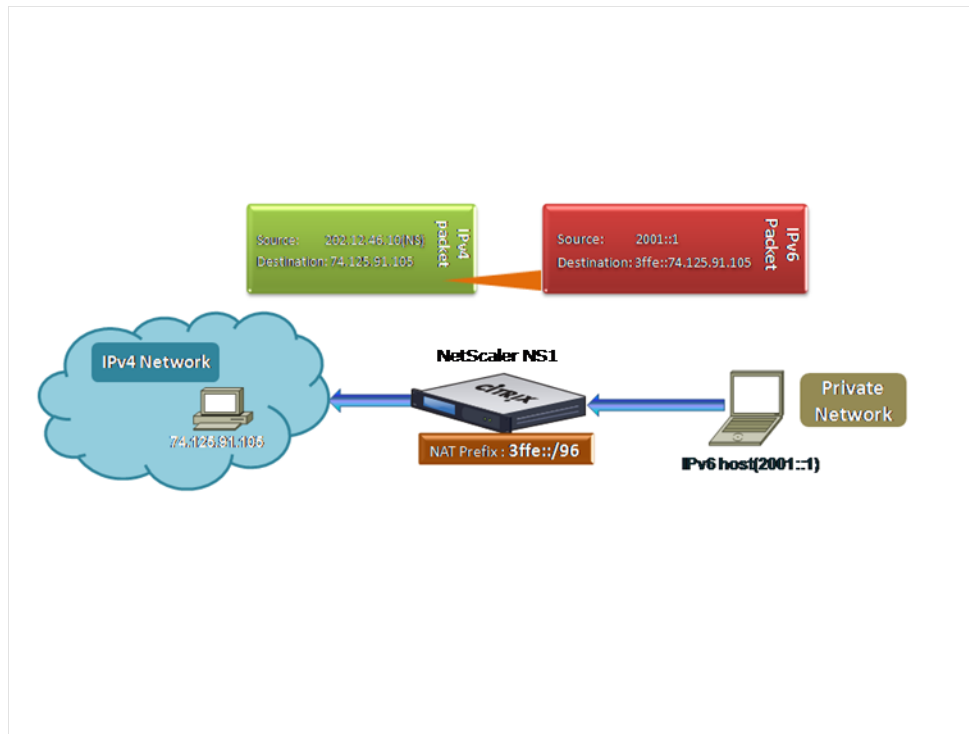
Prefix-based translation is a process of translating packets sent from private IPv6 servers into IPv4 packets, using an IPv6 prefix configured in the NetScaler appliance. This prefix has a length of 96 bits (128-32=96). The IPv6 servers embed the destination IP address of the IPv4 servers or hosts in the last 32 bits of the destination IP address field of the IPv6 packets. The first 96 bits of the destination IP address field are set as the IPv6 NAT prefix.

The NetScaler appliance compares the first 96 bits of the destination IP address of all the incoming IPv6 packets to the configured prefix. If there is a match, the NetScaler appliance generates an IPv4 packet and sets the destination IP address as the last 32 bits of the destination IP address of the matched IPv6 packet. IPv6 packets addressed to this prefix have to be routed to the NetScaler so that the IPv6-IPv4 translation is done by the NetScaler.

In the following diagram, 3ffe::/96 is configured as the IPv6 NAT prefix on NetScaler NS1. The IPv6 host sends an IPv6 packet with destination IP address 3ffe::

74.125.91.105. NS1 compares the first 96 bits of the destination IP address of all the incoming IPv6 packets to the configured prefix, and they match. NS1 then generates an IPv4 packet and sets the destination IP address as 74.125.91.105.

Figure 7-8. IPv6-IPv4 Prefix-Based Translation



To configure prefix-based IPv6-IPv4 translation by using the command line interface

At the command prompt, type the following commands to set a NAT prefix and verify its configuration:

- ♦ **set ipv6 [-natprefix <ipv6_addr|*>]**
- ♦ **show ipv6**

Example

```
> set ipv6 -natprefix 3ffe::/96
Done
```

To configure prefix-based IPv6-IPv4 translation by using the configuration utility

1. Navigate to **System > Network**.

2. In the details pane, in the **Settings** group, click **Change IPv6 Settings**.
3. In the **Configure IPv6 settings** dialog box, set the **IPv6 NAT prefix** parameter.
4. Click **OK**.

Configuring Static ARP

You can add static ARP entries to and remove static ARP entries from the ARP table. After adding an entry, you should verify the configuration. If the IP address, port, or MAC address changes after you create a static ARP entry, you must remove or manually adjust the static entry. Therefore, creating static ARP entries is not recommended unless necessary.

To add a static ARP entry by using the command line interface

At the command prompt, type:

- ♦ **add arp** -IPAddress <ip_addr> -mac<mac_addr> -ifnum <interface_name>
- ♦ **show arp** <IPAddress>

Example

```
> add arp -ip 10.102.29.6 -mac 00:24:e8:73:ca:ec -  
ifnum 1/1  
Done
```

To remove a static ARP entry by using the command line interface

At the command prompt, type the **rm arp** command and the IP address.

To add a static ARP entry by using the configuration utility

1. Navigate to **System > Network > ARP Table**.
2. On the **ARP Table** page, in the details pane, click **Add**.
3. In the **Create ARP entry** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create** or **OK**, and then click **Close**.

Setting the Timeout for Dynamic ARP Entries

You can globally set an aging time (time-out value) for dynamically learned ARP entries. The new value applies only to ARP entries that are dynamically learned after the new value is set. Previously existing ARP entries expire after the previously configured aging time.

You can specify an ARP time-out value of from 1 through 1200 seconds.

To set the time-out for dynamic ARP entries by using the command line interface

At the command prompt, type the following commands to set the time-out for dynamic ARP entries and verify its configuration:

- ♦ `set arpparam -timeout <positive_integer>]`
- ♦ `show arpparam`

Example

```
> set arpparam -timeout 500
Done
```

To set the time-out for dynamic ARP entries to its default value by using the command line interface

At the command prompt, type the following commands to set the time-out for dynamic ARP entries to its default value and verify its configuration:

- ♦ `unset arpparam`
- ♦ `show arpparam`

Example

```
> unset arpparam
Done
```

To set the time-out for dynamic ARP entries by using the configuration utility

1. Navigate to **System > Network**.
2. In the details pane, in the **Settings** group, click **Configure ARP Global Parameters**.
3. In the **Configure ARP Global Parameters** dialog box, type a value for **ARP Table Entry Timeout**.
4. Click **OK**.

Configuring Neighbor Discovery

Neighbor discovery (ND) is one of the most important protocols of IPv6. It is a message-based protocol that combines the functionality of the Address Resolution Protocol (ARP), Internet Control Message Protocol (ICMP), and Router Discovery. ND allows nodes

to advertise their link layer addresses and obtain the MAC addresses or link layer addresses of the neighboring nodes. This process is performed by the Neighbor Discovery protocol (ND6).

Neighbor discovery can perform the following functions:

Router Discovery

Enables a host to discover the local routers on an attached link and automatically configure a default router.

Prefix Discovery

Enables the host to discover the network prefixes for local destinations.

Note: Currently, the NetScaler does not support Prefix Discovery.

Parameter Discovery

Enables a host to discover additional operating parameters, such as MTU and the default hop limit for outbound traffic.

Address Autoconfiguration

Enables hosts to automatically configure IP addresses for interfaces both with and without stateful address configuration services such as DHCPv6. The NetScaler does not support Address Autoconfiguration for Global IPv6 addresses.

Address Resolution

Equivalent to ARP in IPv4, enables a node to resolve a neighboring node's IPv6 address to its link-layer address.

Neighbor Unreachability Detection

Enables a node to determine the reachability state of a neighbor.

Duplicate Address Detection

Enables a node to determine whether an NSIP address is already in use by a neighboring node.

Redirect

Equivalent to the IPv4 ICMP Redirect message, enables a router to redirect the host to a better first-hop IPv6 address to reach a destination.

Note: The NetScaler does not support IPv6 Redirect.

To enable neighbor discovery, you create entries for the neighbors.

Adding IPv6 Neighbors

Adding IPv6 neighbors enables neighbor discovery.

To add an IPv6 neighbor by using the command line interface

At the command prompt, type:

- ♦ **add nd6** <neighbor> <mac> <ifnum> [-vlan <integer>]

♦ **show nd6****Example**

```

> add nd6 2001::1 00:04:23:be:3c:06 1/1 -vlan 1
Done
> show nd6
Neighbor          MAC-Address (Vlan,
Interface)        State      TIME
-----
-----
1) ::1            00:d0:68:0b:
58:da( 1, LO/1)   REACHABLE  PERMANENT
2) fe80::2d0:68ff:fe0b:58da 00:d0:68:0b:
58:da( 1, LO/1)   REACHABLE  PERMANENT
3) 2001::1        00:04:23:be:3c:
06( 1, 1/1)       REACHABLE  STATIC
Done

```

To add an IPv6 neighbor by using the configuration utility

1. Navigate to **System > Network > IPv6 Neighbors**.
2. In the details pane, click **Add**.
3. In the **CreateIPv6 Neighbor** dialog box, in the **Neighbor** and **MAC Address** text boxes, respectively, type IPv6 address and MAC Address of the neighbor (for example, 3ffe:100:100::1, 00:d0:68:0b:58:da).
4. If the neighbor is part of a VLAN, in the **VLAN** field, type the VLAN ID (for example, 1).
5. In the **Interface** list box, select the interface of the neighbor (for example, LO/1).
6. Click **Create**, and click **Close**.

Removing IPv6 Neighbors**To remove a neighbor discovery entry by using the command line interface**

At the command prompt, type:

```
rm nd6 <Neighbor> -vlan <VLANID>
```

Example

```
rm nd6 3ffe:100:100::1 -vlan 1
```

To remove all neighbor discovery entries by using the command line interface

At the command prompt, type:

```
clear nd6
```

To remove a neighbor discovery entry by using the configuration utility

1. Navigate to **System > Network > IPv6 Neighbor**.
2. In the details pane, select the neighbor entry that you want to remove (for example, 3ffe:100:100::1).
3. Click **Remove**.

To remove all neighbor discovery entries by using the configuration utility

1. Navigate to **System > Network > IPv6 Neighbor**.
2. In the **IPv6 Neighbors** page, click **Clear**.

Configuring IP Tunnels

An IP Tunnel is a communication channel, that can be created by using encapsulation technologies, between two networks that do not have a routing path. Every IP packet that is shared between the two networks is encapsulated within another packet and then sent via the tunnel.

The NetScaler appliance implements IP Tunneling in the following ways:

- ♦ NetScaler as an Encapsulator (Load Balancing with DSR mode)
- ♦ NetScaler as a Decapsulator

NetScaler as an Encapsulator (Load Balancing with DSR Mode)

Consider an organization that has multiple data centers across different countries, where the NetScaler maybe at one location and the back-end servers are located in a different country. In essence, the NetScaler and the back-end servers are on different networks and are connected via a router.

When you configure Direct Server Return (DSR) on this NetScaler, the packet sent from the source subnet is encapsulated by the NetScaler and sent via a router and tunnel to the appropriate back-end server. The back-end server decapsulates the packet and responds directly to the client, without allowing the packet to pass via the NetScaler.

NetScaler as a Decapsulator

Consider an organization having multiple data centers each having NetScalers and back-end servers. When a packet is sent from data center A to data center B it is usually

sent via an intermediary, say a router or another NetScaler. The NetScaler processes the packet and then forwards the packet to the back-end server. However, if an encapsulated packet is sent, the NetScaler must be able to decapsulate the packet before sending it to the back-end servers. To enable the NetScaler to function as a decapsulator, a tunnel is added between the router and the NetScaler. When the encapsulated packet, with additional header information, reaches the NetScaler, the data packet is decapsulated i.e. the additional header information is removed, and the packet is then forwarded to the appropriate back-end servers.

The NetScaler can also be used as a decapsulator for the Load Balancing feature, specifically in scenarios when the number of connections on a vserver exceeds a threshold value and all the new connections are then diverted to a back-up vserver.

Creating IP Tunnels

To create an IP tunnel by using the command line interface

At the command prompt type:

- ♦ **add iptunnel** <name> <remotelp> <remoteSubnetMask> <localIp> -type -protocol (ipoverip | GRE) -ipsecprofile <name>
- ♦ **show iptunnel**

Note: While configuring an IP tunnel in a cluster setup, the local IP address must be a striped SNIP or MIP address. Clustering of NetScaler 1000V appliances is not supported.

To remove an IP tunnel by using the command line interface

To remove an IP tunnel, type the **rm iptunnel** command and the name of the tunnel.

To create an IP Tunnel by using the configuration utility

1. Navigate to **System > Network > IP Tunnels**.
2. In the details pane, click **Add**.
3. In the **Add IP Tunnel** dialog box, specify values for the following parameters:
 - **Name***—name
 - **Remote IP***—remotelp
 - **Remote Mask***—remoteSubnetMask
 - **Local IP Type***—localIp (in the **local IP Type** drop down list, select one of the IP type (Mapped IP, Subnet IP, and Virtual). All the configured IPs of the selected IP type will be populated in the **Local IP** drop down list. Select the desired IP from the list.)
 - **Protocol**—protocol and ipsecProfileName from the corresponding field when you select protocol as GRE.

*A required parameter.

4. Click **Create**, and then click **Close**.

To create an IPv6 tunnel by using the command line interface

At the command prompt type:

- ♦ **add ip6tunnel** <name> <remotelp> <local>
- ♦ **show ip6tunnel**

To remove an IPv6 tunnel by using the command line interface

To remove an IPv6 tunnel, type the **rm ip6tunnel** command and the name of the tunnel.

To create an IPv6 Tunnel by using the configuration utility

1. Navigate to **System > Network > IP Tunnels**.
2. On the **IPv6 Tunnels** tab, click **Add**.
3. In the **Create IPv6 Tunnel** dialog box, set the following parameters:
 - **Name***
 - **Remote IP***
 - **Local IP Type*** (In the **local IP Type** drop down list, select one of the IP type (SNIP6 or VIP6). All the configured IPv6 addresses of the selected IPv6 type are be populated in the **Local IP** drop down list. Select the desired IP from the list.)

*A required parameter.

4. Click **Create**, and then click **Close**.

Customizing IP Tunnels Globally

By globally specifying the source IP address, you can assign a common source IP address across all tunnels. Also, because fragmentation is CPU-intensive, you can globally specify that the NetScaler appliance drop any packet that requires fragmentation. Alternatively, if you would like to fragment all packets as long as a CPU threshold value is not reached, you can globally specify the CPU threshold value.

To globally customize IP tunnels by using the command line interface

At the command prompt, type the following commands to globally customize IP tunnels and verify the configuration:

- ♦ **set iptunnelparam** -srcIP <sourceIPAddress> -srcIPRoundRobin (YES | NO)-dropFrag [YES | NO] -dropFragCpuThreshold <Positive integer>
- ♦ **show iptunnelparam**

Example

```
> set iptunnelparam -srcIP 12.12.12.22 -dropFrag
Yes -dropFragCpuThreshold 50
```

```

Done
> set iptunnelparam -srcIPRoundRobin YES -dropFrag
Yes -dropFragCpuThreshold 50
Done

```

Note: To create a new MIP or SNIP address to use as the global source IP address, use the **add ns ip** command before you type the **set iptunnelparam** command.

To globally customize IP tunnels by using the configuration utility

1. Navigate to **System > Network**.
2. In the details pane, in the **Settings** group, click **IPv4 Tunnel Global Settings**.
3. In the **Configure IP Tunnel Global Parameters** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **OK** and then click **Close**.

To globally customize IPv6 tunnels by using the command line interface

At the command prompt, type the following commands to globally customize IPv6 tunnels and verify the configuration:

- ♦ **set ip6tunnelparam** -srcIP <IPv6Address> -srcIPRoundRobin (YES | NO)-dropFrag [YES | NO] -dropFragCpuThreshold <Positive integer>
- ♦ **show ip6tunnelparam**

Note: To create a new VIP6 or SNIP6 address to use as the global source IP address, use the **add ns ip6** command before you type the **set ip6tunnelparam** command.

To globally customize IPv6 tunnels by using the configuration utility

1. Navigate to **System > Network**.
2. In the details pane, in the **Settings** group, click **IPv6 Tunnel Global Settings**.
3. In the **Configure IPv6 Tunnel Global Parameters** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **OK** and then click **Close**.

Interfaces

Before you begin configuring interfaces, decide whether your configuration can use MAC-based forwarding mode, and either enable or disable this system setting accordingly. The number of interfaces in your configuration is different for the different models of the Citrix NetScaler appliance. In addition to configuring individual

interfaces, you can logically group interfaces, using VLANs to restrict data flow within a set of interfaces, and you can aggregate links into channels. In a high availability setup, you can configure a virtual MAC (VMAC) address if necessary. If you use L2 mode, you might want to modify the aging of the bridge table.

When your configuration is complete, decide whether you should enable the system setting for path MTU discovery. NetScaler appliances can be deployed in active-active mode using VRRP. An active-active deployment, in addition to preventing downtime, makes efficient use of all the NetScaler appliances in the deployment. You can use the Network Visualizer tool to view the network configuration of a NetScaler deployment and configure interfaces, channels, VLANs, and bridge groups.

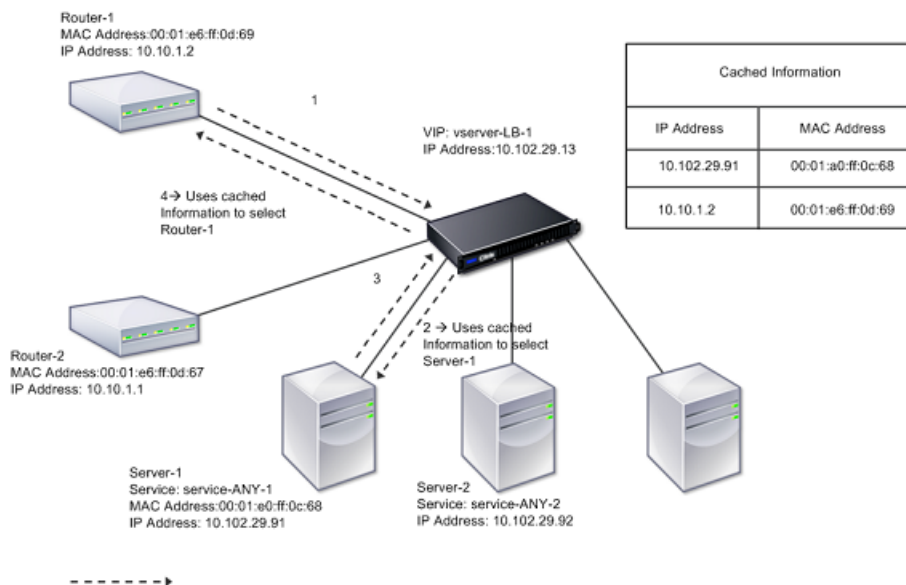
Configuring MAC-Based Forwarding

With MAC-based forwarding (MBF) enabled, when a request reaches the NetScaler appliance, the appliance remembers the source MAC address of the frame and uses it as the destination MAC address for the resulting replies. MAC-based forwarding can be used to avoid multiple-route/ARP lookups and to avoid asymmetrical packet flows. MAC-based forwarding may be required when the NetScaler is connected to multiple stateful devices, such as VPNs or firewalls, because it ensures that the return traffic is sent to the same device that the initial traffic came from.

MAC-based forwarding is useful when you use VPN devices, because it guarantees that all traffic flowing through a VPN passes back through the same VPN device.

The following topology diagram illustrates the process of MAC-based forwarding.

Figure 7-9. MAC-Based Forwarding Mode



When MAC-based forwarding (MBF) is enabled, the NetScaler caches the MAC address of:

- ♦ The source (a transmitting device such as router, firewall, or VPN device) of the inbound connection.
- ♦ The server that responds to the requests.

When a server replies through the NetScaler appliance, the appliance sets the destination MAC address of the response packet to the cached address, ensuring that the traffic flows in a symmetric manner, and then forwards the response to the client. The process bypasses the route table lookup and ARP lookup functions. However, when the NetScaler initiates a connection, it uses the route and ARP tables for the lookup function. In a direct server return configuration, you must enable MAC-based forwarding.

Some deployment topologies may require the incoming and outgoing paths to flow through different routers. MAC-based forwarding would break this topology design.

MBF should be disabled in the following situations:

- ♦ **When you configure link load balancing.** In this case, asymmetric traffic flows are desirable because of link costs.
- ♦ **When a server uses network interface card (NIC) teaming without using LACP (802.1ad Link Aggregation).** To enable MAC-based forwarding in this situation, you must use a layer 3 device between the NetScaler and server.

Note: MBF can be enabled when the server uses NIC teaming with LACP, because the virtual interface uses one MAC address.

When MBF is disabled, the NetScaler uses L2 or L3 connectivity to forward the responses from servers to the clients. Depending on the route table, the routers used for outgoing connection and incoming connection can be different. In the case of reverse traffic (response from the server):

- ♦ If the source and destination are on different IP subnets, the NetScaler uses the route lookup to locate the destination.
- ♦ If the source is on the same subnet as the destination, the NetScaler looks up the ARP table to locate the network interface and forwards the traffic to it. If the ARP table does not exist, the NetScaler requests the ARP entries.

To enable or disable MAC-based forwarding by using the command line interface

At the command prompt, type:

- ♦ `enable ns mode mbf`
- ♦ `disable ns mode mbf`

To enable or disable MAC-based forwarding by using the configuration utility

1. Navigate to **System > Settings**.
2. In the details pane, in the **Modes and Features** group, click **Configure modes**.

3. In the **Configure Modes** dialog box, do one of the following:
 - To enable MAC-based forwarding, select the **MAC-based forwarding** check box.
 - To disable MAC-based forwarding, clear the **MAC-based forwarding** check box.
4. Click **OK**.
5. In the **Enable/Disable Feature(s)?** dialog box, click **Yes**.

Configuring Network Interfaces

Network interfaces in the NetScaler appliance are numbered in <slot>/<port> notation. After configuring your interfaces, you should display the interfaces and their settings to verify the configuration. You can also display this information to troubleshoot a problem in the configuration.

To manage the network interfaces, you might have to enable some interfaces and disable others. You can reset an interface to renegotiate its settings. You can clear the accumulated statistics for an interface. To verify the configuration, you can display the interface settings. You can display the statistics for an interface to evaluate its health.

Setting the Network Interface Parameters

The network interface configuration is neither synchronized nor propagated. For an HA pair, you must perform the configuration on each unit independently.

Network interface parameters include Link Aggregate Control Protocol (LACP) settings. For more information about Link Aggregate Control Protocol (LACP), see "[Configuring Link Aggregation Using the Link Aggregate Channel Protocol](#)."

Note: Configuring speed, duplex, and auto negotiation parameters of an interface is not supported on NetScaler 1000V

To set the network interface parameters by using the command line interface

At the command prompt, type:

- ♦ **set interface** <id> [-flowControl <flowControl>] [-haMonitor (ON | OFF)] [(ON | OFF)] [-tagall (ON | OFF)] [-lacpMode <lacpMode>] [-lacpKey<positive_integer>] [-lacpPriority <positive_integer>] [-lacpTimeout (LONG | SHORT)] [-ifAlias <string>] [-throughput <positive_integer>][-bandwidthHigh <positive_integer> [-bandwidthNormal <positive_integer>]]
- ♦ **show interface** [<id>]

Example

```
> set interface 1/8 -duplex full
Done
```

To set the network interface parameters by using the configuration utility

1. Navigate to **System > Network > Interfaces**.
2. On the **Interfaces** pane, select the network interface that you want to modify (for example, **1/8**), and then click **Open**.
3. In the **Configure Interface** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **OK**.

Enabling and Disabling Network Interfaces

By default, the network interfaces are enabled. You must disable any network interface that is not connected to the network, so that it cannot send or receive packets. Disabling a network interface that is connected to the network in a high availability setup can cause failover.

For more information about high availability, see "[High Availability](#)."

To enable or disable a network interface by using the command line interface

At the command prompt, type one of the following pairs of commands to enable or disable an interface and verify the setting:

- ♦ **enable interface** <interface_num>
- ♦ **show interface** <interface_num>
- ♦ **disable interface** <interface_num>
- ♦ **show interface** <interface_num>

Example

```
> enable interface 1/8
Done
> show interface 1/8
Interface 1/8 (Gig Ethernet 10/100/1000
Mbits) #2
  flags=0x4004000 <ENABLED, DOWN, BOUND to
LA/1, down, autoneg, 802.1q>
  MTU=1514, MAC=00:d0:68:15:fd:3d, downtime
906h58m40s
  Requested: media UTP, speed AUTO, duplex
FULL, fctl OFF, throughput 0
  RX: Pkts(0) Bytes(0) Errs(0) Drops(0)
Stalls(0)
  TX: Pkts(0) Bytes(0) Errs(0) Drops(0)
Stalls(0)
  NIC: InDisc(0) OutDisc(0) Fctls(0)
Stalls(0) Hangs(0) Muted(0)
```

```

Done      Bandwidth thresholds are not set.

```

To enable or disable a network interface by using the configuration utility

1. Navigate to **System > Network > Interfaces**.
2. On the **Interfaces** pane, select the network interface that you want to enable or disable, and do one of the following:
 - To enable a network interface, click **Enable**.
 - To disable a network interface, click **Disable**.

A message appears in the status bar, stating that the network interface has been enabled or disabled successfully.

Resetting Network Interfaces

Network interface settings control properties such as duplex and speed. To renegotiate the settings of a network interface, you must reset it.

Note: Configuring speed, duplex, and auto negotiation parameters of an interface is not supported on NetScaler 1000V.

To reset a network interface by using the command line interface

At the command prompt, type the following commands to reset an interface and verify the setting:

- ♦ **reset interface** <interface_num>
- ♦ **show interface** <interface_num>

Example

```

> reset interface 1/8
Done

```

To reset a network interface by using the configuration utility

1. Navigate to **System > Network > Interfaces**.
2. On the **Interfaces** pane, select the network interface that you want to reset (for example, **1/8**).
3. Click **Reset Interface**.

Monitoring a Network Interface

You can display network interface statistics to monitor parameters such as packets sent and packets received, throughput, Link Aggregate Control Protocol (LACP) data units,

and errors, and use the information to check the health of the network interface. You can clear the statistics of a network interface to monitor its statistics from the time the statistics are cleared.

To display the statistics of the network interfaces by using the command line interface

At the command prompt, type:

```
stat interface <interface_num>
```

To display the statistics of an Interface by using the configuration utility

1. Navigate to **System > Network > Interfaces**.
2. On the **Interfaces** page, select the network interface whose statistics you want to display (for example, **1/8**).
3. Click **Statistics**.

To clear a network interface's statistics by using the command line interface

At the command prompt, type:

```
clear interface <interface_num>
```

Example

```
> clear interface 1/8
Done
```

To clear a network interface's statistics by using the configuration utility

1. Navigate to **System > Network > Interfaces**.
2. On the **Interfaces** pane, select the network interface whose statistics you want to clear (for example, **1/8**).
3. Click **Clear Statistics**.

Configuring Forwarding Session Rules

By default, the NetScaler appliance does not create session entries for traffic that it only forwards (L3 mode). For a case in which a client request that the appliance forwards to a server results in a response that has to return by the same path, you can create a forwarding-session rule. A forwarding-session rule creates forwarding-session entries for traffic that originates from or is destined for a particular network and is forwarded by the NetScaler.

To create a forwarding session rule by using the command line interface

At the command prompt, type the following commands to create a forwarding-session rule and verify the configuration:

- ♦ **add forwardingSession** <name> [<network> <netmask>] | [-aclname <string>] - connfailover (ENABLED | DISABLED)
- ♦ **show forwardingSession**

Example

A network address as the condition:

```
> add forwardingSession fs-nw-1 10.102.105.51  
255.255.255.255  
Done
```

An ACL as the condition:

```
> add forwardingSession fs-acl-1 acl1  
Done
```

To configure a forwarding session rule by using the configuration utility

1. Navigate to **System > Network > Forwarding Sessions**.
2. In the details pane, click **Add**.
3. In the **Create Forwarding Session** dialog box, set the **Name** parameter:
4. Do one of the following:
 - If you want to use the network address as a condition for creating a forwarding session rule, click **Subnet** and set the following parameters:
 - ♦ **Subnet IP**
 - ♦ **Netmask**
 - If you want to use an extended ACL as a condition for creating a forwarding session rule, click **ACL** and set the **ACL Name** parameter:
5. If the appliance is configured as a high availability node, and you want to synchronize the forwarding session's connection information with the secondary node, select **Connection Failover**.
6. Click **Create**, and then **Close**.

Understanding VLANs

A NetScaler appliance supports Layer 2 port and IEEE 802.1q tagged VLANs. VLAN configurations are useful when you need to restrict traffic to certain groups of stations. You can configure a network interface as a part of multiple VLANs by using IEEE 802.1q tagging.

You can configure VLANs and bind them to IP subnets. The NetScaler then performs IP forwarding between these VLANs (if it is configured as the default router for the hosts on these subnets).

The NetScaler supports the following types of VLANs:

Port-Based VLANs. The membership of a port-based VLAN is defined by a set of network interfaces that share a common, exclusive Layer 2 broadcast domain. You can configure multiple port-based VLANs. By default, all network interfaces on the NetScaler are members of VLAN 1.

If you apply 802.1q tagging to the port, the network interface belongs to a port-based VLAN. Layer 2 traffic is bridged within a port-based VLAN, and Layer 2 broadcasts are sent to all members of the VLAN if Layer 2 mode is enabled. When you add an untagged network interface as a member of a new VLAN, it is removed from its current VLAN.

Default VLAN. By default, the network interfaces on the NetScaler are included in a single, port-based VLAN as untagged network interfaces. This VLAN is the default VLAN. It has a VLAN ID (VID) of 1. This VLAN exists permanently. It cannot be deleted, and its VID cannot be changed.

When you add a network interface to a different VLAN as an untagged member, the network interface is automatically removed from the default VLAN. If you unbind a network interface from its current port-based VLAN, it is added to the default VLAN again.

Tagged VLANs. 802.1q tagging (defined in the IEEE 802.1q standard) allows a networking device (such as the NetScaler) to add information to a frame at Layer 2 to identify the VLAN membership of the frame. Tagging allows network environments to have VLANs that span multiple devices. A device that receives the packet reads the tag and recognizes the VLAN to which the frame belongs. Some network devices do not support receiving both tagged and untagged packets on the same network interface—in particular, Force10 switches. In such cases, you need to contact customer support for assistance.

The network interface can be a tagged or untagged member of a VLAN. Each network interface is an untagged member of one VLAN only (its native VLAN). This network interface transmits the frames for the native VLAN as untagged frames. A network interface can be a part of more than one VLAN if the other VLANs are tagged.

When you configure tagging, be sure to match the configuration of the VLAN on both ends of the link. The port to which the NetScaler connects must be on the same VLAN as the NetScaler network interface.

Note: This VLAN configuration is neither synchronized nor propagated, therefore you must perform the configuration on each unit in an HA pair independently.

Applying Rules to Classify Frames

VLANs have two types of rules for classifying frames:

Ingress rules. Ingress rules classify each frame as belonging only to a single VLAN. When a frame is received on a network interface, the following rules are applied to classify the frame:

- ♦ If the frame is untagged, or has a tag value equal to 0, the VID of the frame is set to the port VID (PVID) of the receiving interface, which is classified as belonging to the native VLAN. (PVIDs are defined in the IEEE 802.1q standard.)
- ♦ If frame has a tag value equal to FFF, the frame is dropped.
- ♦ If the VID of the frame specifies a VLAN of which the receiving network interface is not a member, the frame is dropped. For example, if a packet is sent from a subnet associated with VLAN ID 12 to a subnet associated with VLAN ID 10, the packet is dropped. If an untagged packet with VID 9 is sent from the subnet associated with VLAN ID 10 to a network interface PVID 9, the packet is dropped.

Egress Rules. The following egress rules are applied:

- ♦ If the VID of the frame specifies a VLAN of which the transmission network interface is not a member, the frame is discarded.
- ♦ During the learning process (defined by the IEEE 802.1q standard), the Src MAC and VID are used to update the bridge lookup table of the NetScaler.
- ♦ A frame is discarded if its VID specifies a VLAN that does not have any members. (You define members by binding network interfaces to a VLAN.)

VLANs and Packet Forwarding on the NetScaler

The forwarding process on the NetScaler appliance is similar to that on any standard switch. However, the NetScaler performs forwarding only when Layer 2 mode is on. The key features of the forwarding process are:

- ♦ Topology restrictions are enforced. Enforcement involves selecting each network interface in the VLAN as a transmission port (depending on the state of the network interface), bridging restrictions (do not forward on the receiving network interface), and MTU restrictions.
- ♦ Frames are filtered on the basis of information in the bridge table lookup in the forwarding database (FDB) table of the NetScaler. The bridge table lookup is based on the destination MAC and the VID. Packets addressed to the MAC address of the NetScaler are processed at the upper layers.
- ♦ All broadcast and multicast frames are forwarded to each network interface that is a member of the VLAN, but forwarding occurs only if L2 mode is enabled. If L2 mode is disabled, the broadcast and multicast packets are dropped. This is also true for MAC addresses that are not currently in the bridging table.

- ♦ A VLAN entry has a list of member network interfaces that are part of its untagged member set. When forwarding frames to these network interfaces, a tag is not inserted in the frame.
- ♦ If the network interface is a tagged member of this VLAN, the tag is inserted in the frame when the frame is forwarded.

When a user sends any broadcast or multicast packets without the VLAN being identified, that is, during duplicate address detection (DAD) for NSIP or ND6 for the next hop of the route, the packet is sent out on all the network interfaces, with appropriate tagging based on either the Ingress and Egress rules. ND6 usually identifies a VLAN, and a data packet is sent on this VLAN only. Port-based VLANs are common to IPv4 and IPv6. For IPv6, the NetScaler supports prefix-based VLANs.

Configuring a VLAN

You can implement VLANs in the following environments:

- ♦ Single subnet
- ♦ Multiple subnets
- ♦ Single LAN
- ♦ VLANs (no tagging)
- ♦ VLANs (802.1q tagging)

If you configure VLANs that have only untagged network interfaces as their members, the total number of possible VLANs is limited to the number of network interfaces available in the NetScaler. If more IP subnets are required with a VLAN configuration, 802.1q tagging must be used.

When you bind a network interface to a VLAN, the network interface is removed from the default VLAN. If the network interfaces need to be a part of more than one VLAN, you can bind the network interfaces to the VLANs as tagged members.

You can configure the NetScaler to forward traffic between VLANs at Layer 3. In this case, a VLAN is associated with a single IP subnet. The hosts in a VLAN that belong to a single subnet use the same subnet mask and one or more default gateways connected to that subnet. Configuring Layer 3 for a VLAN is optional. Layer 3 is used for IP forwarding (inter-VLAN routing). Each VLAN has a unique IP address and subnet mask that define an IP subnet for the VLAN. In an HA configuration, this IP address is shared with the other NetScaler appliances. The NetScaler forwards packets between configured IP subnets (VLANs).

When you configure the NetScaler, you must not create overlapping IP subnets. Doing so impedes Layer 3 functionality.

Each VLAN is a unique Layer 2 broadcast domain. Two VLANs, each bound to separate IP subnets, cannot be combined into a single broadcast domain. Forwarding traffic between two VLANs requires a Layer 3 forwarding (routing) device, such as the NetScaler appliance.

Creating or Modifying a VLAN

To configure a VLAN, you create a VLAN entity, and then bind network interfaces and IP addresses to the VLAN. If you remove a VLAN, its member interfaces are added to the default VLAN.

To create a VLAN by using the command line interface

At the command prompt, type:

add vlan <id> [-aliasName <string>] [-ipv6DynamicRouting (ENABLED|DISABLED)]

Example

```
> add vlan 2 -aliasName "Network A"
Done
```

To bind an interface to a VLAN by using the command line interface

At the command prompt, type:

bind vlan <id> -ifnum <slot/port>

Example

```
> bind vlan 2 -ifnum 1/8
Done
```

To bind an IP address to a VLAN by using the command line interface

At the command prompt, type:

bind vlan <id> -IPAddress <IPAddress> <netMask>

Example

```
> bind vlan 2 -IPAddress 10.102.29.54 255.255.255.0
Done
```

To remove a VLAN by using the command line interface

At the command prompt, type:

rm vlan <id>

To configure a VLAN by using the configuration utility

1. Navigate to **System > Network > VLANs**.
2. In the details pane, do one of the following:
 - To create a new VLAN, click **Add**.
 - To modify an existing VLAN, click **Open**.
3. In the **Add VLAN** or **Configure/Modify VLAN** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. To bind an IP address to a VLAN, under **IPs**, select the **Active** check box corresponding to the IP address that you want to bind to the VLAN (for example, **10.102.29.54**). The **Type** column displays the IP address type (such as mapped IP, virtual IP, or subnet IP) for each IP address in the **IP Addresses** column.
5. To bind a network interface to a VLAN, under **Interfaces**, select the **Active** check box corresponding to the interface that you want to bind to the VLAN (for example, **1/8**).
6. Click **Create** or **OK**, and then click **Close**.

Monitoring VLANs

You can display VLAN statistics such as packets received, bytes received, packets sent, and bytes sent, and use the information to identify anomalies and or debug a VLAN.

To view the statistics of a VLAN by using the command line interface

At the command prompt, type:

```
stat vlan <vlanID>
```

Example

```
stat vlan 2
```

To view the statistics of a VLAN by using the configuration utility

1. Navigate to **System > Network > VLANs**.
2. On the **VLANs** page, select the VLAN whose statistics you want to view (for example, **2**).
3. Click **Statistics**.

Configuring VLANs in an HA Setup

VLAN configuration for a high-availability setup requires that the NetScaler appliances have the same hardware configuration, and the VLANs configured on them must be mirror images.

The correct VLAN configuration is implemented automatically when the configuration is synchronized between the NetScaler appliances. The result is identical actions on all the appliances. For example, adding network interface 0/1 to VLAN2 adds this network interface to VLAN 2 on all the appliances participating in the high-availability setup.

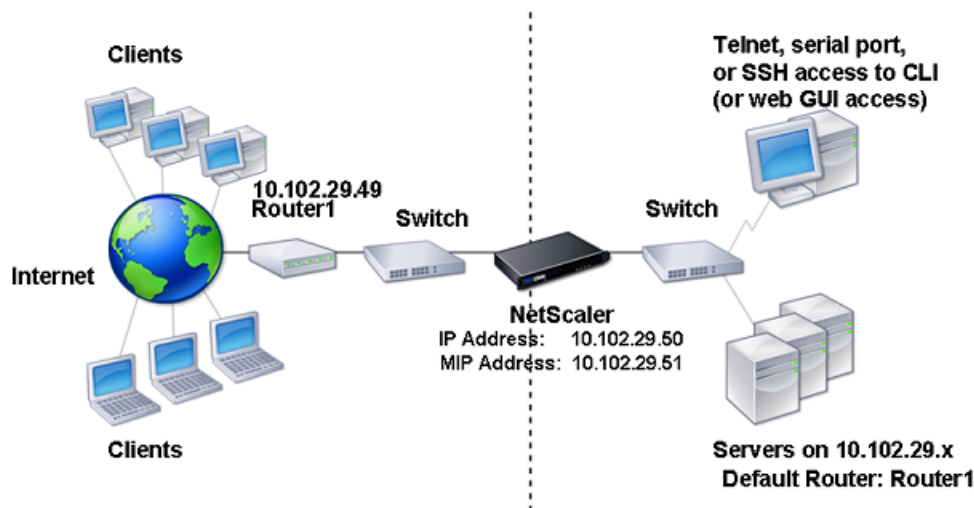
Note: If you use network-interface-specific commands in an HA setup, the configurations you create are not propagated to the other NetScaler appliance. You must perform these commands on each appliance in an HA pair to ensure that the configuration of the two appliances in the HA pair remains synchronized.

Configuring VLANs on a Single Subnet

Before configuring a VLAN on a single subnet, make sure that Layer 2 Mode is enabled.

The following figure shows a single subnet environment

Figure 7-10. VLAN on a Single Subnet



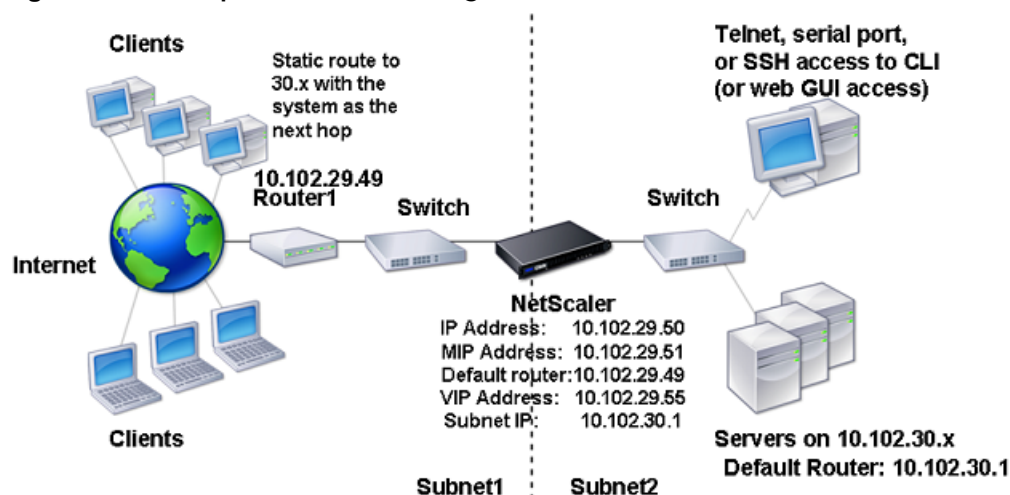
In the above figure:

1. The default router for the NetScaler and the servers is Router 1.
2. Layer 2 mode must be enabled on the NetScaler for the NetScaler to have direct access to the servers.
3. For this subnet, a virtual server can be configured for load balancing on the NetScaler.

To configure a VLAN on a single subnet, follow the procedures described in "[Creating or Modifying a VLAN](#)." VLAN configuration parameters are not required, because the network interfaces are members of this VLAN.

Configuring VLANs on Multiple Subnets

To configure a single VLAN across multiple subnets, you must add a VIP for the VLAN and configure the routing appropriately. The following figure shows a single VLAN configured across multiple subnets.

Figure 7-11. Multiple Subnets in a Single VLAN

To configure a single VLAN across multiple subnets, perform the following tasks:

1. Disable Layer 2 mode.
2. Add a VIP.

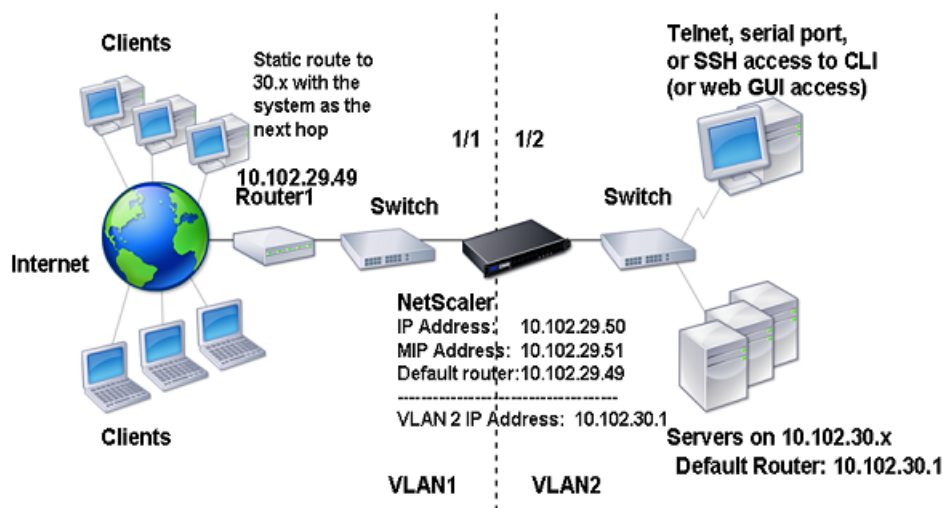
For the procedure to add a VIP, see "[Configuring and Managing Virtual IP Addresses \(VIPs\)](#)."

3. Configure RNAT ID.

For the procedure to configure the RNAT ID, see "[Configuring RNAT](#)."

Configuring Multiple Untagged VLANs across Multiple Subnets

In environments with multiple untagged VLANs across multiple subnets, a VLAN is configured for each IP subnet. A network interface is bound to one VLAN only. The following figure shows this configuration.

Figure 7-12. Multiple Subnets with VLANs - No Tagging

To implement the configuration shown in the above figure, perform the following tasks:

1. Add VLAN 2.

For the procedure to create a VLAN, see "[Creating or Modifying a VLAN](#)."

2. Bind the 1/2 network interface of the NetScaler to VLAN 2 as an untagged network interface.

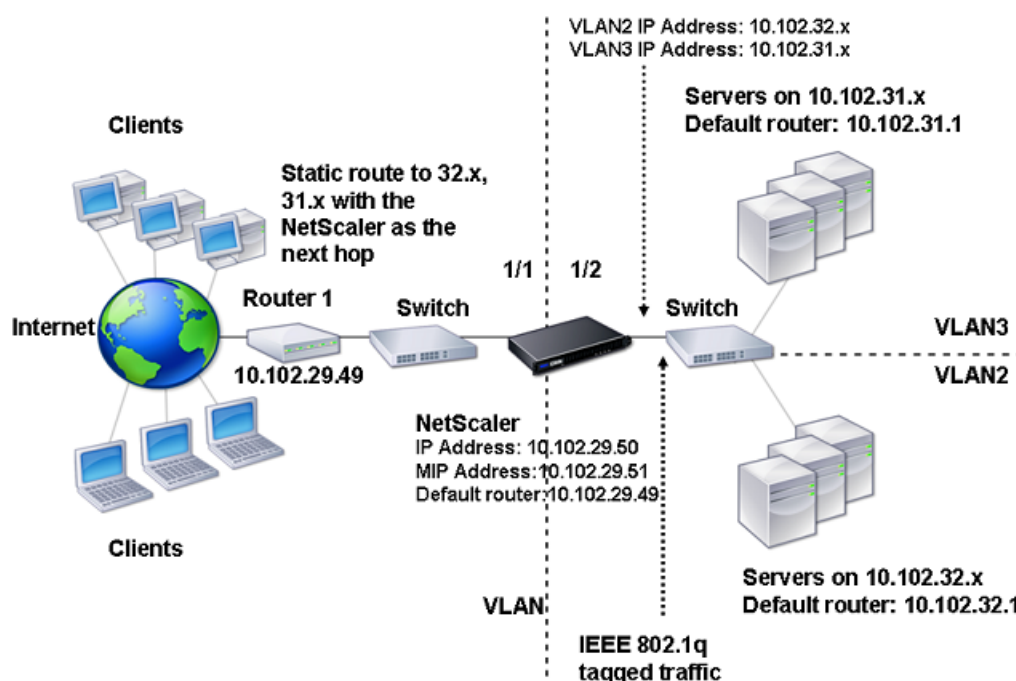
For the procedure to bind a network interface to a VLAN, see "[Creating or Modifying a VLAN](#)."

3. Bind the IP address and subnet mask to VLAN 2.

For the procedure to bind a network interface to a VLAN, see "[Creating or Modifying a VLAN](#)."

Configuring Multiple VLANs with 802.1q Tagging

For multiple VLANs with 802.1q tagging, each VLAN is configured with a different IP subnet. Each network interface is in one VLAN. One of the VLANs is set up as tagged. The following figure shows this configuration.

Figure 7-13. Multiple VLANs with IEEE 802.1q Tagging

To implement the configuration shown in the above figure, perform the following tasks:

1. Add VLAN 2.

For the procedure to create a VLAN, see "[Creating or Modifying a VLAN.](#)"

2. Bind the 1/2 network interface of the NetScaler to VLAN 2 as an untagged network interface.

For the procedure to bind a network interface to a VLAN, see "[Creating or Modifying a VLAN.](#)"

3. Bind the IP address and netmask to VLAN 2.

For the procedure to bind an IP address to a VLAN, see "[Creating or Modifying a VLAN.](#)"

4. Add VLAN 3.

For the procedure to create a VLAN, see "[Creating or Modifying a VLAN.](#)"

5. Bind the 1/2 network interface of the NetScaler to VLAN 3 as a tagged network interface.

For the procedure to bind a network interface to a VLAN, see "[Creating or Modifying a VLAN.](#)"

For the procedure to bind a tagged network interface, see "[Creating or Modifying a VLAN.](#)"

6. Bind the IP address and netmask to VLAN 3.

For the procedure to bind an IP address to a VLAN, see "[Creating or Modifying a VLAN](#)."

Configuring NSVLAN

NSVLAN is a VLAN to which the NetScaler management IP (NSIP) address's subnet is bound. The NSIP subnet is available only on interfaces that are associated with NSVLAN. By default, NSVLAN is VLAN1, but you can designate a different VLAN as NSVLAN. If you do so, you must reboot the NetScaler appliance for the change to take effect. After the reboot, NSIP subnet traffic is restricted to the new NSVLAN.

The traffic from the NetScaler IP subnet can be tagged (802.1q) with the VLAN ID specified for NSVLAN. You must configure the attached switch interface to tag and allow this same VLAN ID on the connected interface.

If you remove your NSVLAN configuration, the NSIP subnet is automatically bound to VLAN1, restoring the default NSVLAN.

To configure NSVLAN by using the command line interface

At the command prompt, type:

- ♦ **set ns config -nsvlan** <positive_integer> -ifnum <interface_name> ... [-tagged (YES|NO)]
- ♦ **show ns config**

Note: The configuration takes effect after the NetScaler appliance is rebooted.

Example

```
> set ns config -nsvlan 300 -ifnum 1/1 1/2 1/3 -  
tagged NO  
Done  
  
> save config  
Done
```

To restore the default NSVLAN configuration by using the command line interface

At the command prompt, type:

- ♦ **unset ns config -nsvlan**
- ♦ **show ns config**

Example

```
> unset ns config -nsvlan  
Done
```

To configure NSVLAN by using the configuration utility

1. Navigate to **System > Settings**.
2. In the details pane, under **Settings**, click **Change NSVLAN Settings**.
3. In the **Configure NSVLAN Settings** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Under **Interfaces**, select interfaces from the **Available Interfaces** list and click **Add** to move them to the **Configured Interfaces** list.
5. Click **OK**. In the **Warning** dialog box, click **OK**. The configuration takes effect after the NetScaler appliance is restarted.

Configuring Bridge Groups

Typically, when you want to merge two or more VLANs into a single domain, you change the VLAN configuration on all the devices in the separate domains. This can be a tedious task. To more easily merge multiple VLANs into a single broadcast domain, you can use bridge groups.

The bridge groups feature works the same way as a VLAN. Multiple VLANs can be bound to a single bridge group, and all VLANs bound to same bridge group form a single broadcast domain. You can bind only Layer 2 VLANs to a bridge group. For Layer 3 functionality, you must assign an IP address to a bridge group.

In Layer 2 mode, a broadcast packet received on an interface belonging to a particular VLAN is bridged to other VLANs that belong to the same bridge group. In the case of a unicast packet, the NetScaler appliance searches its bridge table for the learned MAC addresses of all the VLANs belonging to same bridge group.

In Layer 3 forwarding mode, an IP subnet is bound to a bridge group. The NetScaler accepts incoming packets belonging to the bound subnet and forwards the packets only on VLANs that are bound to the bridge group.

IPv6 routing can be enabled on a configured bridge group.

To add a bridge group and bind VLANs by using the command line interface

To add a bridge group and bind VLANs and verify the configuration, type the following commands:

- ♦ **add bridgegroup <id> [-ipv6DynamicRouting (ENABLED | DISABLED)]**
- ♦ **show bridgegroup <id>**

- ♦ **bind bridgegroup** <id> -vlan <positive_integer>
- ♦ **show bridgegroup** <id>

Example

```
> add bridgegroup 12
Done
```

To remove a bridge group by using the command line interface

At the command prompt, type:

```
rm bridgegroup <id>
```

Example

```
rm bridgegroup 12
```

To configure a bridge group by using the configuration utility

1. Navigate to **System > Network > Bridge Groups**.
2. In the details pane, do one of the following:
 - To create a new bridge group, click **Add**.
 - To modify an existing bridge group, click **Open**.
3. In the **Create Bridge Group** or **Configure Bridge Group** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. To bind a VLAN to a bridge group, under **VLANs**, select the **Active** check box corresponding to the interface that you want to bind to the bridge group (for example, **1/8**).
5. To bind an IP address to a bridge group, under **IPs**, select the **Active** check box corresponding to the IP address that you want to bind to the bridge group (for example, **10.102.29.54**).
6. Click **Create** or **OK**, and then click **Close**.

Configuring VMACs

The primary and secondary nodes in a high availability (HA) setup share the Virtual MAC address (VMAC) floating entity. The primary node owns the floating IP addresses (such

as MIP, SNIP, and VIP) and responds to ARP requests for these IP addresses with its own MAC address. Therefore, the ARP table of an external device, such as an upstream router, is updated with the floating IP address and the MAC address of the primary node.

When a failover occurs, the secondary node takes over as the new primary node. The former secondary node uses Gratuitous ARP (GARP) to advertise the floating IP addresses that it had learned from the old primary node. The MAC address that the new primary node advertises is the MAC address of its own network interface. Some devices (a few routers) do not accept these GARP messages. Therefore, these external devices retain the IP address-to-MAC address mapping that the old primary node had advertised. This can result in a GSLB site going down.

Therefore, you must configure a VMAC on both nodes of an HA pair. This means that both nodes have identical MAC addresses. When a failover occurs, the MAC address of the secondary node remains unchanged, and the ARP tables on the external devices do not need to be updated.

For the procedures to configure a VMAC, see "[High Availability](#)."

Configuring Link Aggregation

Link aggregation combines data coming from multiple ports into a single high-speed link. Configuring link aggregation increases the capacity and availability of the communication channel between the NetScaler appliance and other connected devices. An aggregated link is also referred to as a "channel." You can configure the channels manually, or you can use Link Aggregation Control Protocol (LACP). You cannot apply LACP to a manually configured channel, nor can you manually configure a channel created by LACP.

Note: Configuring link aggregation manually is not supported on NetScaler 1000V.

When a network interface is bound to a channel, the channel parameters have precedence over the network interface parameters. (That is, the network interface parameters are ignored.) A network interface can be bound only to one channel.

When a network interface is bound to a channel, it drops its VLAN configuration. When network interfaces are bound to a channel, either manually or by LACP, they are removed from the VLANs that they originally belonged to and added to the default VLAN. However, you can bind the channel back to the old VLAN, or to a new one. For example, if you bind the network interfaces 1/2 and 1/3 to a VLAN with ID 2, and then bind them to a channel LA/1, the network interfaces are moved to the default VLAN, but you can bind them back to VLAN 2.

Configuring Link Aggregation by Using the Link Aggregation Control Protocol

The Link Aggregation Control Protocol (LACP) enables network devices to exchange link aggregation information by exchanging LACP Data Units (LACPDUs). Therefore, you cannot enable LACP on network interfaces that are members of a channel that you created manually.

When using LACP to configure link aggregation, you use different commands and parameters for modifying link aggregation channels than you do for creating link aggregation channels. To remove a channel, you must disable LACP on all interfaces that are part of the channel.

Note: In an High Availability configuration, LACP configurations are neither propagated nor synchronized.

Creating Link Aggregation Channels

For creating a link aggregation channel by using LACP, you need to enable LACP and specify the same LACP key on each interface that you want to be the part of the channel. For example, if you enable LACP and set the LACP Key to 3 on interfaces 1/1 and 1/2, a link aggregation channel LA/3 is created and interfaces 1/1 and 1/2 are automatically bound to it.

Note: When enabling LACP on a network interface, you must specify the LACP Key.

By default, LACP is disabled on all network interfaces.

To create an LACP channel by using the command line interface

At the command prompt, type:

- ♦ **set interface** <id> [-lacpMode <lacpMode>] [-lacpKey<positive_integer>] [-lacpPriority <positive_integer>] [-lacpTimeout (LONG | SHORT)]
- ♦ **show interface** [<id>]

To create an LACP channel by using the configuration utility

1. Navigate to **System > Network > Interfaces**.
2. On the **Interfaces** pane, select the network interface that you want to modify (for example, 1/8), and then click **Open**.
3. In the **Configure interface** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **OK**.

Modifying Link aggregation Channels

After you have created an LACP channel by specifying interfaces, you can modify properties of the channel.

To modify an LACP channel using the command line interface

At the command prompt, type:

- ♦ **set channel** <id> [-ifnum <interfaceName> ...] [-state (ENABLED | DISABLED)] [-speed <speed>] [-flowControl <flowControl>] [-haMonitor (ON | OFF)] [-ifAlias <string>] [-throughput <positive_integer>] [-tagall (ON | OFF)] [-bandwidthHigh <positive_integer>] [-bandwidthNormal <positive_integer>]]

- ♦ **show channel**

Example

```
> set channel LA/3 -state ENABLED -speed 10000
Done
```

To modify an LACP channel by using the configuration utility

1. Navigate to **System > Network > Channels**.
2. In the details pane, select an LACP channel and then click **Open**.
3. In the **Configure Channel** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **OK**.

Removing a Link Aggregation Channel

To remove a link aggregation channel that was created by using LACP, you need to disable LACP on all the interfaces that are part of the channel.

To remove an LACP channel by using the command line interface

At the command prompt, type:

- ♦ **set interface <id> -lacpMode Disable**
- ♦ **show interface [<id>]**

To remove an LACP channel by using the configuration utility

1. Navigate to **System > Network > Interfaces**.
2. On the **Interfaces** pane, select the network interface that you want to modify (for example, **1/8**), and then click **Open**.
3. In the **Configure Interface** dialog box, clear the **Enable LACP** check box.
4. Click **OK**.

Configuring the LACP System Priority

The LACP system priority determines which peer device of an LACP LA channel can have control over the LA channel. This number is globally applied to all LACP channels on the appliance. The lower the value, the higher the priority.

To configure the LACP system priority by using the command line interface

At the command prompt, type the following commands to set the priority for a standalone appliance and verify the configuration:

- ♦ **set lacp -sysPriority <positive_integer>**
- ♦ **show lacp**

Example:

```
set lacp -sysPriority 50
```

To configure the LACP system priority by using the configuration utility

1. Navigate to **System > Network > Interfaces**.
2. In the details pane, click the **Action** drop-down list and select **Set LACP....**
3. In the **Configure LACP** dialog box, specify the system priority and the owner node (applicable only for a cluster setup).

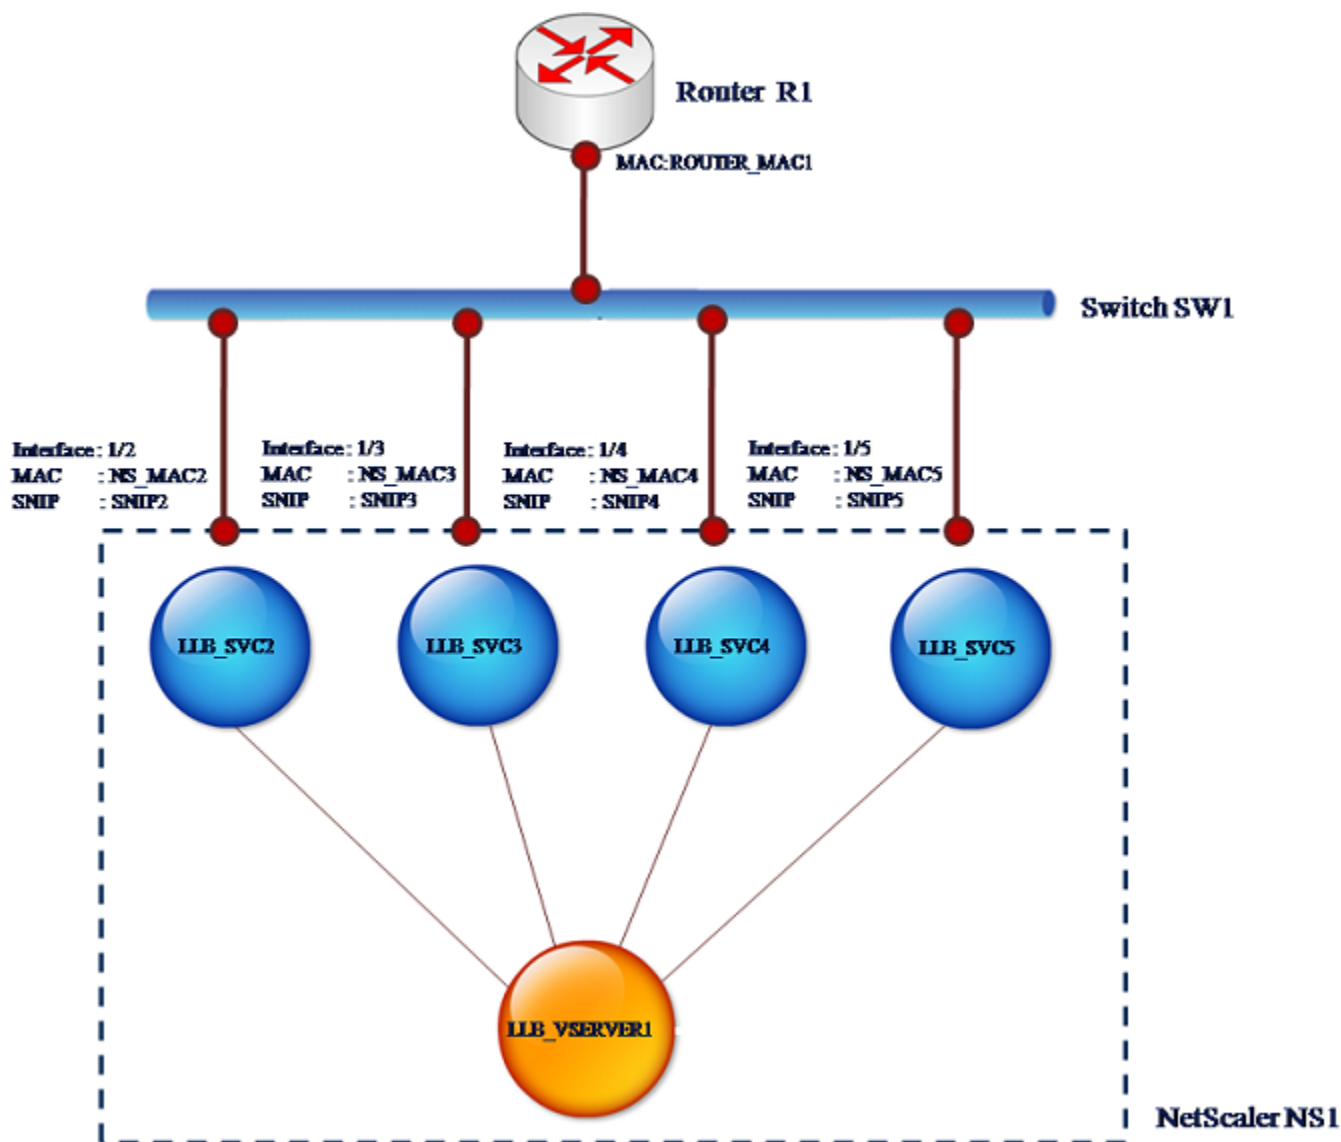
Note: Clustering of NetScaler 1000V appliances is not supported.

4. Click **OK**.
5. To verify the configuration, click the **Action** drop-down list and select **View LACP Details....**

Binding an SNIP address to an Interface

You can now bind a NetScaler owned SNIP address to an interface without using Layer 3 VLANs. Any packets related to the SNIP address will go only through the bound interface.

This feature can be useful in a scenario where the upstream switch does not support Link Aggregation channels and you want the NetScaler appliance to load balance traffic, originated from a server, across the four links to the upstream switch as shown in the following illustration.



The following tables describe the example settings for the scenario:

Entity	Name	Value
SNIP addresses on NS1	SNIP2 (for reference purpose only)	10.10.10.2
	SNIP3 (for reference purpose only)	10.10.10.3
	SNIP4 (for reference purpose only)	10.10.10.4

Entity	Name	Value
	SNIP5 (for reference purpose only)	10.10.10.5
LLB virtual server on NS1	LLB_VSERVER1	-
Transparent monitor on NS1	TRANS_MON	-
LLB services on NS1	LLB_SVC2	10.10.10.240
	LLB_SVC3	10.10.10.120
	LLB_SVC4	10.10.10.60
	LLB_SVC5	10.10.10.30
MAC address of interface 1/2 on NS1	NS_MAC_2 (for reference purpose only)	00:e0:ed:0f:bc:e0
MAC address of interface 1/3 on NS1	NS_MAC_3 (for reference purpose only)	00:e0:ed:0f:bc:df
MAC address of interface 1/4 on NS1	NS_MAC_4 (for reference purpose only)	00:e0:ed:0f:bc:de
MAC address of interface 1/5 on NS1	NS_MAC_5 (for reference purpose only)	00:e0:ed:1c:89:53
IP address of Router R1	Router_IP (for reference purpose only)	10.10.10.1
MAC address of interface of R1	ROUTER_MAC1 (for reference purpose only)	00:21:a1:2d:db:cc

To configure the example settings

1. Add four different SNIPs in different subnet ranges. This is for ARP to be resolved on four different links. For more information on creating a SNIP address, see ["Configuring Subnet IP Addresses \(SNIPs\)."](#)

Command Line Interface example

```
> add ns ip 10.10.10.2 255.255.255.0 -type SNIP
Done
> add ns ip 10.10.10.3 255.255.255.128 -type SNIP
Done
> add ns ip 10.10.10.4 255.255.255.192 -type SNIP
Done
> add ns ip 10.10.10.5 255.255.255.224 -type SNIP
Done
```

2. Add four different dummy services in the added SNIP subnets. This is to ensure that the traffic is sent out with source IP as one of the four configured SNIPs.

Command Line Interface example

```
> add service LLB_SVC2 10.10.10.240 any *
Done
> add service LLB_SVC3 10.10.10.120 any *
Done
> add service LLB_SVC4 10.10.10.60 any *
Done
> add service LLB_SVC5 10.10.10.30 any *
Done
```

3. Add a transparent ping monitor for monitoring the gateway. Bind the monitor to each of the configured dummy services. This is to make the state of the services as UP.

Command Line Interface example

```
> add monitor TRANS_MON ping -destIP 10.10.10.1 -
transparent YES
Done
> bind monitor TRANS_MON LLB_SVC2
Done
> bind monitor TRANS_MON LLB_SVC3
Done
> bind monitor TRANS_MON LLB_SVC4
Done
> bind monitor TRANS_MON LLB_SVC5
Done
```

4. Add a link load balancing (LLB) virtual server and bind the dummy services to it.

Command Line Interface example

```
> add lb vserver LLB_VSERVER1 any
Done
> set lb vserver LLB_VSERVER1 -lbmethod ROUNDROBIN
Done
> bind lb vserver LLB_VSERVER1 LLB_SVC2
Done
> bind lb vserver LLB_VSERVER1 LLB_SVC2
Done
> bind lb vserver LLB_VSERVER1 LLB_SVC2
Done
> bind lb vserver LLB_VSERVER1 LLB_SVC2
Done
```

5. Add the LLB virtual server as the default LLB route.

Command Line Interface example

```
> add lb route 0.0.0.0 0.0.0.0 LLB_VSERVER1
Done
```

6. Add an ARP entry for each of the dummy services with the MAC address of the gateway. This way the gateway is reachable through these dummy services. For more information on adding an ARP entry, see ["Configuring Static ARP."](#)

Command Line Interface example

```
> add arp -ipaddress 10.10.10.240 -mac 00:21:a1:2d:db:cc -
ifnum 1/2
Done
> add arp -ipaddress 10.10.10.120 -mac 00:21:a1:2d:db:cc -
ifnum 1/3
Done
> add arp -ipaddress 10.10.10.60 -mac 00:21:a1:2d:db:cc -
ifnum 1/4
Done
> add arp -ipaddress 10.10.10.30 -mac 00:21:a1:2d:db:cc -
ifnum 1/5
Done
```

7. Bind a specific interface to an SNIP by adding an ARP entry for each of these SNIPs. This is to ensure that the response traffic will reach the same interface through which the request went out. For more information on adding an ARP entry, see ["Configuring Static ARP."](#)

Command Line Interface example

```
> add arp -ipAddress 10.10.10.2 -mac 00:e0:ed:0f:bc:e0 -
ifnum 1/2
Done
> add arp -ipAddress 10.10.10.3 -mac 00:e0:ed:0f:bc:df -
ifnum 1/3
Done
> add arp -ipAddress 10.10.10.4 -mac 00:e0:ed:0f:bc:de -
ifnum 1/4
Done
> add arp -ipAddress 10.10.10.5 -mac 00:e0:ed:1c:89:53 -
ifnum 1/5
Done
```

Monitoring the Bridge Table and Changing the Aging time

NetScaler appliance bridges frames on the basis of bridge table lookup of the destination MAC address and the VLAN ID. However, the appliance performs forwarding only when Layer 2 mode is enabled.

The bridge table is dynamically generated, but you can display it, modify the aging time for the bridge table, and view bridging statistics.

To change the aging time by using the command line interface

At the command prompt, type:

- ♦ `set bridgetable -bridgeAge <positive_integer>`
- ♦ `show bridgetable`

Example

```
> set bridgetable -bridgeage 70
Done
```

To change the aging time by using the configuration utility

1. Navigate to **System > Network > Bridge Table**.
2. In the details pane, click **Change Ageing Time**.
3. In the **Change Ageing Time** dialog box, in the **Ageing Time (seconds)** text box, type the aging time (for example, **70**).
4. Click **OK**. All the MAC entries in the bridge table are updated with the aging time.

To view the statistics of a bridge table by using the command line interface

At the command prompt, type:

```
stat bridge
```

To view the statistics of a bridge table by using the configuration utility

1. On the **Bridge Table** page, select the MAC address for which you want to view the statistics (for example, **00:12:01:0a:5f:46**).
2. Click **Statistics**.

Understanding NetScaler Appliances in Active-Active Mode Using VRRP

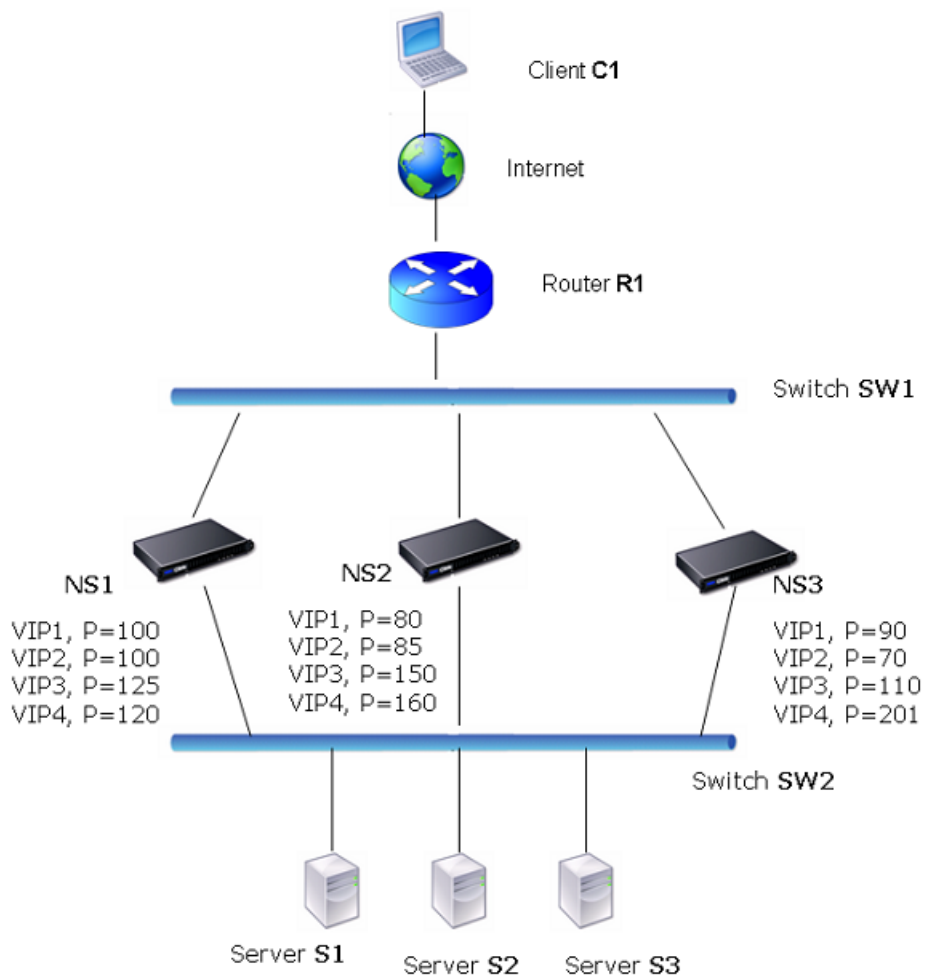
An active-active deployment, in addition to preventing downtime, makes efficient use of all the NetScaler appliances in the deployment. In active-active deployment mode, the same VIPs are configured on all NetScaler appliances in the configuration, but with different priorities, so that a given VIP can be active on only one appliance at a time.

Note: This feature is supported only on NetScaler nCore builds.

The active VIP is called the master VIP, and the corresponding VIPs on the other NetScaler appliances are called the backup VIPs. If a master VIP fails, the backup VIP with the highest priority takes over and becomes the master VIP. All the NetScaler appliances in an active-active deployment use the Virtual Router Redundancy Protocol (VRRP) protocol to advertise their VIPs and the corresponding priorities at regular intervals.

NetScaler appliances in active-active mode can be configured so that no NetScaler is idle. In this configuration, different sets of VIPs are active on each NetScaler. For example, in the following diagram, VIP1, VIP2, VIP3, and VIP4 are configured on appliances NS1, NS2, and NS3. Because of their priorities, VIP1 and VIP2 are active on NS1, VIP3 is active on NS2 and VIP4 is active on NS3. If, for example, NS1 fails, VIP1 on NS3 and VIP2 on NS2 become active.

Figure 7-14. An Active-Active Configuration



The NetScaler appliances in the above diagram process traffic as follows:

1. Client C1 sends a request to VIP1. The request reaches R1.
2. R1 does not have an ARP entry for VIP1, so it broadcasts an ARP request for VIP1.
3. VIP1 is active in NS1, so NS1 replies with a source MAC address as the VMAC (for example VMAC1) associated with VIP1, and VIP1 as the source IP address.
4. SW1 learns the port for VIP1 from the ARP reply and updates its bridge table.
5. R1 updates the ARP entry with VMAC1 and VIP1.
6. R1 forwards the packet to the VIP1 on NS1.
7. NS1's load balancing algorithm selects server S2, and NS1 opens a connection between one of its SNIP or MIP addresses and S2.
8. S2 replies to the SNIP or MIP on the NetScaler.
9. NS1 sends S2's reply to the client. In the reply, NS1 inserts MAC address of the physical interface as the source MAC address and VIP1 as the source IP address.
10. Should NS1 fail, the NetScaler appliances use the VRRP protocol to select the VIP1 with the highest priority. In this case, VIP1 on NS3 becomes active, and the following two steps update the active-active configuration.
11. NS3 broadcasts a GARP message for VIP1. In the message, VMAC1 is the source MAC address and VIP1 is the source IP address.
12. SW1 learns the new port for VMAC1 from the GARP broadcast and updates its bridge table to send subsequent client requests for VIP1 to NS3. R1 updates its ARP table.

The priority of a VIP can be modified by health tracking. If you enable health tracking, you should make sure that preemption is also enabled, so that a VIP whose priority is lowered can be preempted by another VIP.

In some situations, traffic might reach a backup VIP. To avoid dropping such traffic, you can enable sharing, on a per-node basis, as you create an active-active configuration. Or you can enable the global send to master option. On a node on which sharing is enabled, it takes precedence over send to master.

Health Tracking

Base priority (BP-range 1-255) ordinarily determines which VIP is the master VIP, but effective priority (EP) can also affect the determination.

For example, if a VIP on NS1 has a priority of 101 and same VIP on NS2 has a priority of 99, the VIP on NS1 is active. However, if two vservers are using the VIP on NS1 and one of them goes DOWN, health tracking can reduce the EP of VIP on NS1. VRRP then makes the VIP on NS2 the active VIP.

Following are the health tracking options for modifying EP:

- ♦ **NONE.** No tracking. EP = BP
- ♦ **ALL.** If all virtual servers are UP, then EP = BP. Otherwise, EP = 0.

- ♦ **ONE.** If at least one virtual server is UP, then EP = BP. Otherwise, EP = 0.
- ♦ **PROGRESSIVE.** If ALL virtual servers are UP, then EP = BP. If ALL virtual servers are DOWN then EP = 0. Otherwise EP = BP (1 - K/N), where N is the total number of virtual servers associated with the VIP and k is the number of virtual servers that are down.

Note: If you specify a value other than NONE, preemption should be enabled, so that the backup VIP with the highest priority becomes active if the priority of the master VIP is downgraded.

Preemption

Preemption of an active VIP by another VIP that attains a higher priority is enabled by default, and normally should be enabled. In some cases, however, you may want to disable it. Preemption is a per-node setting for each VIP.

Preemption can occur in the following situations:

- ♦ An active VIP goes down and a VIP with a lower priority takes its place. If the VIP with the higher priority comes back online, it preempts the currently active VIP.
- ♦ Health tracking causes the priority of a backup VIP to become higher than that of the active VIP. The backup VIP then preempts the active VIP.

Sharing

In the event that traffic reaches a backup VIP, the traffic is dropped unless the sharing option is enabled on the backup VIP. This behavior is a per node setting for each VIP and is disabled by default.

In the figure "[An Active-Active Configuration](#)," VIP1 on NS1 is active and VIP1 VIPs on NS2 and NS3 are backups. Under certain circumstances, traffic may reach VIP1 on NS2. If Sharing is enabled on NS2, this traffic is processed instead of dropped.

Configuring Active-Active Mode

On each NetScaler appliance that you want to deploy in active-active mode, you must add a VMAC and bind the VMAC to a VIP. The VMAC for a given VIP must be same on each appliance. For example, if VIP 10.102.29.5, is created on the appliances, a virtual router ID must be created on each NetScaler and bound to VIP 10.102.29.5 on each NetScaler. When you bind a VMAC to a VIP, the NetScaler sends VRRP advertisements to each VLAN that is bound to that VIP. The VMAC can be shared by different VIPs configured on the same NetScaler.

Adding a VMAC

To add a VMAC for an active-active configuration, you create a virtual router ID. To bind a VMAC to a VIP, you associate the VMAC's virtual router ID with the VIP.

To add a VMAC by using the command line interface

At the command prompt, type:

add vrID <value> -priority <value> -preemption (ENABLED|DISABLED) -sharing (ENABLED | DISABLED) -tracking (NONE|ONE|ALL|PROGRESSIVE)

Example

```
add vrID 125 -priority 100 -sharing ENABLED -  
tracking ONE
```

To add a VMAC by using the configuration utility

1. Navigate to **System > Network > VMAC**.
2. On the **VMAC** page, click **Add**.
3. In the **Add VMAC** dialog box, in **Virtual Router ID** text box, type a number (for example, **125**) to assign as the VMAC ID.
4. In the **Priority** text box, enter a priority number (for example, **100**) that will associated with VIPs bound this VMAC.
5. In the **Tracking** drop down box, select a health tracking option (for example, **ONE**).
6. Select or clear the **Preemption** check box to disable or enable preemption on VIPs that are bound to this VMAC.
7. Select or clear the **Sharing** check box to enable or disable sharing on VIPs that are bound to this VMAC.
8. Click **Create**.

To bind a VMAC by using the command line interface

At the command prompt, type:

set ns ip <VIP address> -vrid <value>

Example

```
set ns ip 10.102.29.5 -vrid 125
```

To bind a VMAC to a VIP by using the NetScaler configuration utility

1. Navigate to **System > Network > IPs**.
2. In the details pane, on the **IPv4s** tab, select the VIP address (for example, **10.102.29.5**) that you want to bind to a VMAC, and then click **Open**.
3. In the **Configure IP** dialog box, in the **Virtual Router Id** drop down box, select a virtual router ID (for example, **125**).

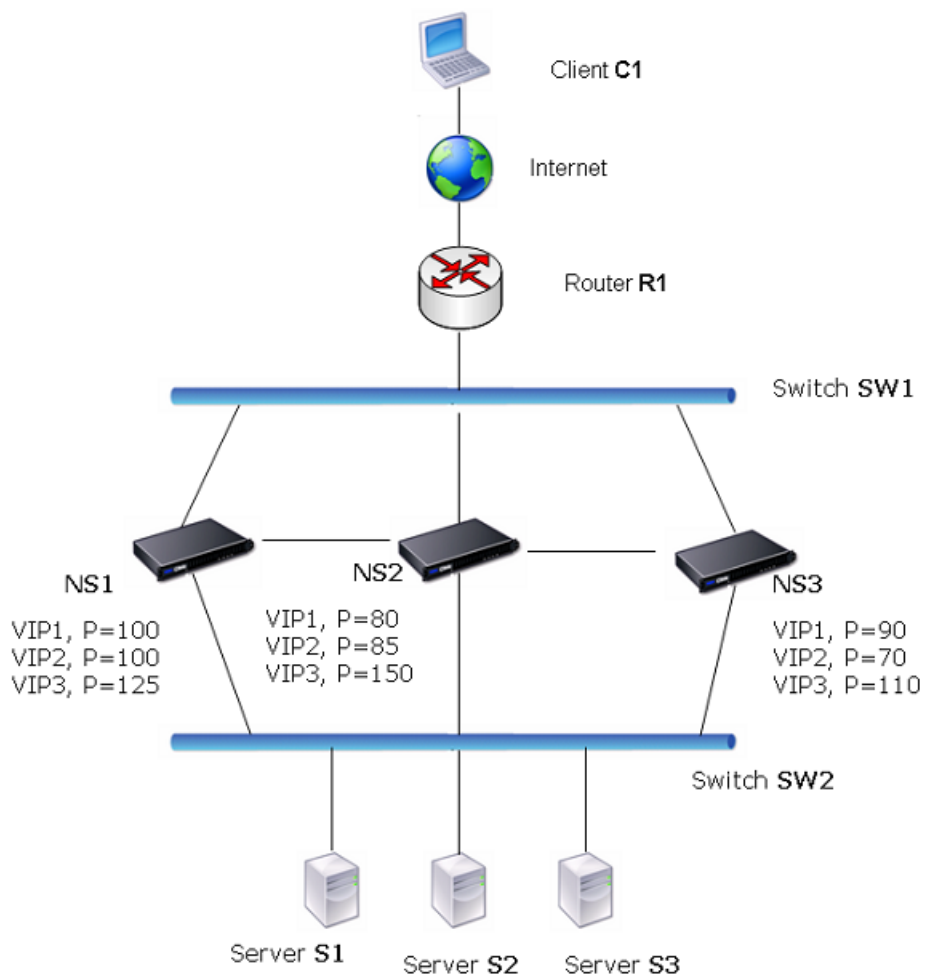
4. Click **OK**.

Configuring Send to Master

Usually, the traffic destined to a VIP reaches the NetScaler appliance on which the VIP is active, because an ARP request with the VIP and a VMAC on that appliance has reached the upstream router. But in some cases, such as static routes configured on the upstream router for the VIP subnet, or a topology that blocks this route, the traffic can reach a NetScaler appliance on which the VIP is in backup state. If you want this appliance to forward the data packets to the appliance on which the VIP is active, you need to enable the send to master option. This behavior is a per node setting and is disabled by default.

For example, in the following diagram, VIP1 is configured on NS1, NS2, and NS3 and is active on NS1. Under certain circumstances, traffic for VIP1 (active on NS1) may reach VIP1 on NS3. When the send to master option is enabled on NS3, NS3 forwards the traffic to NS1 through NS2 by using route entries for NS1.

Figure 7-15. An Active-Active Configuration with Send to Master Option Enabled



To enable send to master by using the command line interface

At the command prompt, type:

set vrIDParam -sendToMaster (ENABLED|DISABLED)

Example

```
> set vrIDParam -sendToMaster ENABLED
Done
```

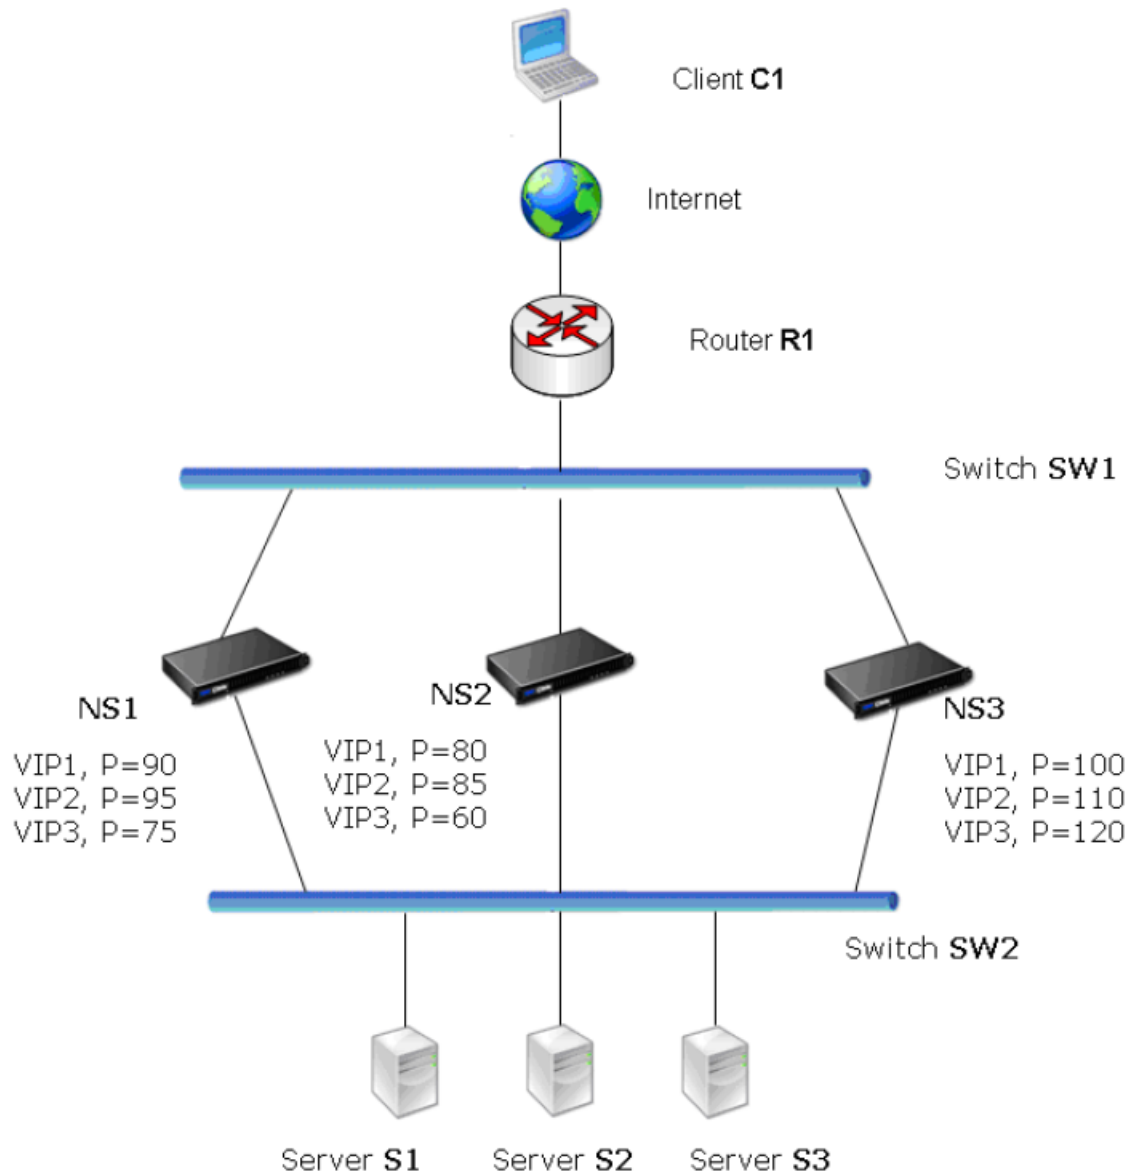
To enable send to master by using the configuration utility

1. Navigate to **System > Network**.
2. In the details pane, under **Settings**, click **Virtual Router Parameters**.
3. In the **Virtual Router Parameters** dialog box, select **Send to Master** option.
4. Click **OK**.

An Active-Active Deployment Scenario

Following is an example of a possible active-active deployment scenario.

In the following diagram, VIP1, VIP 2 and VIP3 are configured on all three appliances, NS1, NS2, and NS3. Base Priorities for each VIPs are as shown in the diagram. Health tracking is disabled for each VIP. The priorities of VIPs are set so that VIP1, VIP2, and VIP3 are active on NS3. If NS3 fails, VIP1, VIP2, and VIP3 become active on NS1.

Figure 7-16. An Active-Active Deployment Scenario

Using the Network Visualizer

The Network Visualizer is a tool that you can use to view the network configuration of a NetScaler node, including the network configuration of the nodes in a high availability (HA) deployment. You can also modify the configuration of VLANs, interfaces, channels, and bridge groups, and perform HA configuration tasks.

In an HA deployment, you can both view and configure network entities on the node to which you are logged on, but you can view the details of only the network entities that

are configured on the peer node. However, you can perform certain tasks, such as viewing details and statistics of the peer node and forcing a failover.

When you are logged on to a standalone appliance, you can use the Network Visualizer to do the following:

- ♦ View a consolidated graphical summary of key network components, such as VLANs, interfaces, channels, and bridge groups. You can also view the individual details of various network components.
- ♦ Modify appliance settings.
- ♦ Add, modify, and enable and disable interfaces and channels that are configured on the NetScaler appliance.
- ♦ Add and modify VLANs and bridge groups.
- ♦ Configure an HA deployment (add a node).
- ♦ View node details, node statistics, and statistics for VLANs and interfaces.
- ♦ Copy the properties of a network entity to a document or spreadsheet.

When you are logged on to an appliance in an HA deployment, you can perform the above tasks only on the appliance to which you are logged on. Following are additional tasks that you can perform in the Network Visualizer when you are logged on to one of the appliances in an HA pair:

- ♦ View the configuration details and high availability details of both nodes in an HA pair.
- ♦ Perform HA configuration tasks, such as synchronization and force failover.
- ♦ Remove the peer node from the HA configuration.
- ♦ View statistics for the peer node.
- ♦ Copy the properties of the peer node to a document or spreadsheet.

To open the Network Visualizer

1. Navigate to **System > Network**.
2. In **Monitor Connections**, click **Network Visualizer**.

To locate a VLAN or bridge group in the Visualizer

Open the **Network Visualizer**, and then do the following:

- To locate a VLAN or bridge group, in the **Search** text field, begin typing the ID of the VLAN or the bridge group that you want to locate.

Alternatively, begin typing the IP address of a bound subnet or the ID of a bound interface. The VLANs or bridge groups whose names match the typed characters are highlighted.

To highlight multiple entities simultaneously, separate the IDs and IP addresses with white spaces. Entities whose IDs or IP addresses match any of the typed IDs and IP addresses are highlighted.

- To clear the Search field, click the x adjacent to the field.

To modify the network settings of the appliance by using the Visualizer

1. Open the **Network Visualizer** and click the icon representing the appliance to which you are logged on.
2. In **Related Tasks**, click **Open**.

To add a channel by using the Visualizer

1. Open the **Network Visualizer** and click a network interface.
2. In **Related Tasks**, click **Add Channel**.

To add a VLAN by using the Visualizer

Open the **Network Visualizer**, click the appliance to which you are logged on, and then do one of the following:

- Click an existing **VLAN**, and then, in **Related Tasks**, click **Add**.
- Click an existing bridge group, and then, in **Related Tasks**, click **Add VLAN**.

To add a bridge group by using the Visualizer

Open the **Network Visualizer**, click the appliance to which you are logged on, and then do one of the following:

- Click an existing bridge group, and then, in **Related Tasks**, click **Add**.
- Click an existing **VLAN**, and then, in **Related Tasks**, click **Add Bridge Group**.

To modify the settings of an interface or channel by using the Visualizer

1. Open the **Network Visualizer** and click the interface whose settings you want to modify.
2. In **Related Tasks**, click **Open**.

To enable or disable an interface or channel by using the Visualizer

1. Open the **Network Visualizer** and click the interface or channel that you want to enable or disable.
2. In **Related Tasks**, do one of the following.
 - To enable the interface or channel, click **Enable**.
 - To disable the interface or channel, click **Disable**.

To remove a configured channel, VLAN, or bridge group by using the Visualizer

1. Open the **Network Visualizer** and click the channel, VLAN, or bridge group that you want to remove from the configuration.
2. In **Related Tasks**, click **Remove**.

To view statistics for a node, channel, interface, or VLAN by using the Visualizer

1. Open the **Network Visualizer** and click the node, interface, or VLAN whose statistics you want to view.
2. In **Related Tasks**, click **Statistics**.

To set up an HA deployment by using the Visualizer

1. Open the **Network Visualizer** and click the appliance.
2. In **Related Tasks**, click **HA Setup**.

To force the secondary node to take over as the primary by using the Visualizer

1. Open the **Network Visualizer** and click one of the nodes.
2. In **Related Tasks**, click **Force Failover**.

To synchronize the secondary node's configuration with the primary node by using the Visualizer

1. Open the **Network Visualizer** and click one of the nodes.
2. In **Related Tasks**, click **Force Synchronization**.

To remove the peer node from the HA configuration

1. Open the **Network Visualizer** and click the peer node.
2. In **Related Tasks**, click **Remove**.

To copy the properties of a node or network entity by using the Visualizer

1. Open the **Network Visualizer** and click the appliance or network entity whose properties you want to copy to a document or spreadsheet.
2. In **Related Tasks**, click **Copy Properties**.

Access Control Lists

Access Control Lists (ACLs) filter IP traffic and secure your network from unauthorized access. An ACL consists of a set of conditions that the NetScaler® appliance uses to

allow or deny access. Consider a small organization that consists of 3 departments, Finance, HR, and Documentation, where no department wants another to access its data. The administrator of the organization can configure ACLs on the NetScaler to allow or deny access. When the NetScaler receives a data packet, it compares the information in the data packet with the conditions specified in the ACL and allows or denies access. The NetScaler supports simple ACLs, extended ACLs, and ACL6s. If both simple and extended ACLs are configured, incoming packets are compared to the simple ACLs first.

Simple ACLs filter packets on the basis of their source IP address and, optionally, their destination port and/or their protocol. Any packet that has the characteristics specified in the ACL is dropped. You can create up to 200,000 simple ACLs.

Extended ACLs filter data packets on the basis of various parameters, such as source IP address, source port, action, and protocol. An extended ACL defines the conditions that a packet must satisfy for the NetScaler to process the packet, bridge the packet, or drop the packet. These actions are known as "processing modes." You can create up to 10,000 extended ACLs.

The processing modes are:

- ♦ ALLOW - The NetScaler processes the packet.
- ♦ BRIDGE - The NetScaler bridges the packet to the destination without processing it.
- ♦ DENY - The NetScaler drops the packet.

The NetScaler processes an IP packet directly when both of the following conditions exist:

- ♦ ACLs are configured on the NetScaler.
- ♦ The IP packet does not match any of the ACLs.

Simple ACL6s filter IPv6 packets on the basis of their source IPv6 address and, optionally, their destination port and/or their protocol. Any packet that has the characteristics specified in the simple ACL6 is dropped. You can create up to 200,000 simple ACL6s.

ACL6s are ACLs created specifically for IPv6 addresses. ACL6s filter packets on the basis of packet parameters, such as source IP address, source port, action, and so on. An ACL6 defines the condition that a packet must satisfy for the NetScaler to process the packet, bridge the packet, or drop the packet. These actions are known as "processing modes." You can create up to 8,000 ACL6s.

The processing modes are:

- ♦ ALLOW - The NetScaler processes the packet.
- ♦ BRIDGE - The NetScaler bridges the packet to the destination without processing it.
- ♦ DENY - The NetScaler drops the packet.

The NetScaler processes an IP packet directly when both of the following conditions exist:

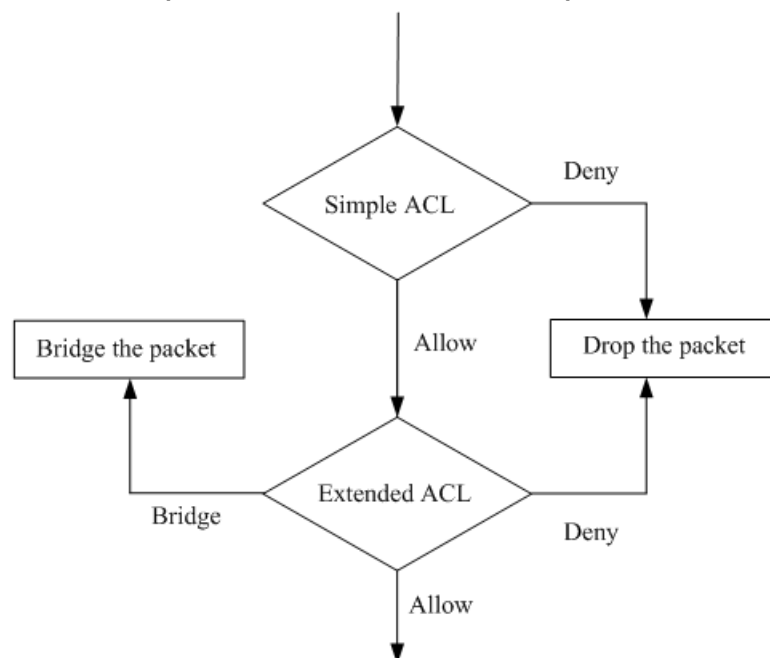
- ♦ ACL6s are configured on the NetScaler.

- ♦ The IP packet does not match any of the ACL6s.

ACL Precedence

An IPv4 packet that matches the conditions specified in a simple ACL is dropped. If the packet does not match any simple ACL, the NetScaler compares the packet's characteristics to those specified in any configured extended ACLs. If the packet matches an extended ACL, the NetScaler applies the action specified in the Extended ACL, as shown in the following diagram.

Figure 7-17. Simple and Extended ACLs Flow Sequence



IPv6 packets are compared only to ACL6s.

Configuring Simple ACLs

A simple ACL, which uses few parameters, cannot be modified once created. When creating a simple ACL, you can specify a time to live (TTL), in seconds, after which the ACL expires. ACLs with TTLs are not saved when you save the configuration. You can also remove a simple ACL manually. You can display simple ACLs to verify their configuration, and you can display statistics to monitor their performance.

Creating Simple ACLs

Use either of the following procedures to create a simple ACL.

To create a simple ACL by using the command line interface

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

- ♦ **add ns simpleacl** <aclname> DENY -srcIP <ip_addr> [-destPort<port> -protocol (TCP | UDP)] [-TTL <positive_integer>]
- ♦ **show ns simpleacl** [<aclname>]

Example

```
> add simpleacl rule1 DENY -srcIP 10.102.29.5 -TTL
600
Done
```

To create a simple ACL by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the **ACLs** pane, on the **Simple ACLs** tab, click **Add**.
3. In the **Add Simple ACL** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create**, and then click **Close**.

Monitoring Simple ACLs

You can display the simple ACL statistics, which include the number of hits, the number of misses, and the number of simple ACLs configured.

To view simple ACL statistics by using the command line interface

At the command prompt, type:

stat ns simpleacl

Example

```
>stat ns simpleacl

Total                                     Rate (/s)
Deny SimpleACL hits                      0
0
SimpleACL hits                           0
0
SimpleACL misses                          0
11
SimpleACLs count                         --
1
Done
```

The following table describes statistics you can display for simple ACLs.

Table 7-3. Simple ACL Statistics

Statistic	Indicates
Deny SimpleACL hits	Packets dropped because they match deny simple ACL
SimpleACL hits	Packets matching a simple ACL
SimpleACL misses	Packets not matching any simple ACL
SimpleACL count	Number of simple ACLs configured

To display simple ACL statistics by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the details pane, select the ACL whose statistics you want to display (for example, rule1).
3. Click **Statistics**.
4. View the ACL statistics in the new window that opens.

Removing Simple ACLs

If you need modify a simple ACL, you must remove it and create a new one.

To remove a single simple ACL by using the command line interface

At the command prompt, type:

- ♦ `rm ns simpleacl <aclname>`
- ♦ `show ns simpleacl`

To remove all simple ACLs by using the command line interface

At the command prompt, type:

- ♦ `clear ns simpleacl`
- ♦ `show ns simpleacl`

To remove a single simple ACL by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the details pane, on the **Simple ACLs** tab, select the simple ACL that you want to remove (for example, rule1).
3. Click **Remove**.
4. In the **Remove** dialog box, click **Yes**.
5. In the details pane, on the **Simple ACLs** tab, verify that the entry for rule1 has been removed

To remove all simple ACLs by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the details pane, on the **Simple ACLs** tab, click **Clear**.
3. In the **Clear Simple ACL (s)** dialog box, click **Yes**.
4. In the details pane, verify that there are no entries in the **Simple ACLs** tab.

Configuring Extended ACLs

To configure extended ACLs, many users first create extended ACLs and then modify them.

For any of the following actions to take effect, they must be applied, by clicking the **Commit** button:

- ♦ Activate
- ♦ Remove
- ♦ Disable
- ♦ Change the Priority

Other actions include:

- ♦ Configure logging
- ♦ Verify the configuration
- ♦ Monitor ACL statistics

Note: If you configure both simple and extended ACLs, simple ACLs take precedence over extended ACLs.

Parameters of Extended ACLs can be configured during creation. Additionally, the following actions can be performed on Extended ACLs: Modify, Remove, Apply, Disable, Enable and Renummer the priority of Extended ACLs.

You can collect statistics of packets using Extended ACLs by enabling logging.

Creating and Modifying an Extended ACL

To create an extended ACL by using the command line interface

At the command prompt, type:

- ♦ **add ns acl** <aclname> <aclaction> [-srcIP [<operator>] <srcIPVal>] [-srcPort [<operator>] <srcPortVal>] [-destIP [<operator>] <destIPVal>] [-destPort [<operator>] <destPortVal>] [-TTL <positive_integer>] [-srcMac <mac_addr>] [(-protocol <protocol> [-established]) | -protocolNumber <positive_integer>] [-vlan <positive_integer>] [-interface <interface_name>] [-icmpType <positive_integer>] [-

```
icmpCode <positive_integer>]] [-priority <positive_integer>] [-state ( ENABLED |
DISABLED )] [-logstate ( ENABLED | DISABLED )] [-ratelimit <positive_integer>]]
```

- ♦ **show ns acl** [<aclname>]

Example

```
> add ns acl restrict DENY -srcport 45-1024 -
destIP 192.168.1.1 -protocol TCP
Done
```

To configure an extended ACL by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the details pane, on the **Extended ACLs** tab, click **Add**.
3. In the **Create ACL** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create**, and then click **Close**.

Applying an Extended ACL

After you create or modify an extended ACL, you must activate it by using one of the following procedures. These procedures reapply all the ACLs.

For example, if you have created the ACLs rule1 through rule10, and then you create an ACL called rule11, and apply it, all of the ACLs (rule1 through rule11) are applied afresh.

If a session has a DENY ACL related to it, that session is terminated.

To apply an ACL by using the command line interface

At the command prompt, type:

- ♦ **apply ns acls**
- ♦ **show ns acl**

To apply an ACL by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. Click **Commit**.
3. In the **Apply ACL(s)** dialog box, click **Yes**.
4. Verify the information on the **Extended ACLs** tab.

Disabling and Enabling Extended ACLs

By default, ACLs are enabled. This means when ACLs are applied, the NetScaler appliance compares incoming packets against the ACLs.

Disable an ACL if it will not be used for a certain period. After the ACLs are applied, the NetScaler does not compare incoming packets against disabled ACLs.

To disable or enable an extended ACL by using the command line interface

At the command prompt, type one of the following pairs of commands to disable or enable an ACL and verify the result:

- ♦ **disable ns acl <aclname>**
- ♦ **show ns acl [<aclname>]**
- ♦ **enable ns acl <aclname>**
- ♦ **show ns acl [<aclname>]**

Example

```
> disable ns acl restrict
Done

> show ns acl restrict
      Name: restrict
Action: DENY      Hits: 0
      srcIP
      destIP = 192.168.1.1
      srcMac:
Protocol: TCP
      srcPort = 45-1024
destPort
      Vlan:
Interface:
      Active Status: DISABLED
Applied Status: NOTAPPLIED
      Priority: 10
NAT: NO
      TTL:
      Log Status: DISABLED
Done

> enable ns acl restrict
Done

> show ns acl restrict
      Name: restrict
Action: DENY      Hits: 0
      srcIP
      destIP = 192.168.1.1
      srcMac:
Protocol: TCP
      srcPort = 45-1024
destPort
      Vlan:
Interface:
      Active Status: ENABLED
```

```
Applied Status: APPLIED
Priority: 10
NAT: NO
TTL:
Log Status: DISABLED
Done
```

To disable or enable an extended ACL by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the details pane, on the **Extended ACLs** tab, select the ACL (for example, rule1) and click **Open**.
3. In the **Configure ACL** dialog box, select the **Enable ACL** check box to enable, or clear the check box to disable, the ACL.
4. Click **OK**.
5. If you want to apply the new setting, which reapplies all ACLs, click **Commit**, and then, in the **Apply ACL(s)** dialog box, click **Yes**.
6. In the details pane, on the **Extended ACLs** tab, view the list to verify the changed status under the column **Active Status**.

Renumbering the priority of Extended ACLs

The renumber procedure resets the priorities of the ACLs to multiples of 10. The priority (an integer value) defines the order in which the NetScaler appliance evaluates ACLs. All priorities are multiples of 10, unless you configure a specific priority to an integer value. When you create an ACL without specifying a priority, the NetScaler automatically assigns a priority that is a multiple of 10.

If a packet matches the condition defined by the ACL, the NetScaler performs an action. If the packet does not match the condition defined by the ACL, the NetScaler compares the packet against the ACL with the next-highest priority.

Consider the following example. Two ACLs, rule1 and rule2, are automatically assigned priorities 20 and 30 when they are created. You need to add a third ACL, rule3, to be evaluated immediately after rule1. Rule3 must have a priority between 20 and 30. In this case, you can specify the priority as 25. Later, you can easily renumber the ACLs with priorities that are multiples of 10, without affecting the order in which the ACLs are applied.

To renumber the ACLs by using the command line interface

At the command prompt, type:

```
renumber ns acls
```

To renumber the ACLs by using the configuration utility

1. Navigate to **System > Network > ACLs**.

2. In the details pane, on the **Extended ACLs** tab, click **Renumber Priority (s)**.
3. In the **Renumber Priority (s) ACL(s)** dialog box, click **Yes**.
4. In the details pane, on the **Extended ACLs** tab, verify the changed priority.

Configuring Extended ACL Logging

You can configure the NetScaler appliance to log details for packets that match an extended ACL. In addition to the ACL name, the logged details include packet-specific information such as the source and destination IP addresses. The information is stored either in the syslog file or in the nslog file, depending on the type of global logging (syslog or nslog) enabled.

Logging can be enabled at both the global level and the ACL level. The global setting takes precedence.

To optimize logging, when multiple packets from the same flow match an ACL, only the first packet's details are logged, and the counter is incremented for every packet that belongs to the same flow. A flow is defined as a set of packets that have the same values for the following parameters:

- ♦ Source IP address
- ♦ Destination IP address
- ♦ Source port
- ♦ Destination port
- ♦ Protocol

If the packet is not from the same flow, or if the time duration is beyond the meantime, a new flow is created. Mean time is the time during which packets of the same flow do not generate additional messages (although the counter is incremented).

Note: The total number of different flows that can be logged at any given time is limited to 10,000.

To configure ACL Logging by using the command line interface

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

- ♦ **set ns acl** <aclName> [-logState (ENABLED | DISABLED)] [-rateLimit <positive_integer>]
- ♦ **show ns acl** [<aclName>]

Example

```
> set ns acl restrict -logstate ENABLED -ratelimit 120
Warning: ACL modified, apply ACLs to activate change
```

```
> apply ns acls
Done
```

To configure ACL Logging by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the details pane, click the **Extended ACLs** tab, and then select the ACL for which you want to configure logging (for example, rule1).
3. Click **Open**.
4. In the **Configure ACL** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
 - **Log State**
 - **Log Rate Limit**
5. Click **OK**.
6. In the **ACL modified, apply ACLs to activate change** dialog box, click **OK**.

Monitoring the Extended ACL

You can display statistics for monitoring the performance of an extended ACL.

To display the statistics of an extended ACL by using the command line interface

At the command prompt, type:

```
stat ns acl
```

Example

```
>stat ns acl rule1

ACL: rule1
s)              Total              Rate (/
Hits for this ACL
0                  0
Done
```

The following table lists the statistics associated with extended ACLs and their descriptions.

Table 7-4. Extended ACL Statistics

Statistic	Specifies
Allow ACL hits	Packets matching ACLs with processing mode set to ALLOW. NetScaler processes these packets.
NAT ACL hits	Packets matching a NAT ACL, resulting in a NAT session.
Deny ACL hits	Packets dropped because they match ACLs with processing mode set to DENY.
Bridge ACL hits	Packets matching a bridge ACL, which in transparent mode bypasses service processing.
ACL hits	Packets matching an ACL.
ACL misses	Packets not matching any ACL.

To display the statistics of an extended ACL by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the details pane, on the **Extended ACLs** tab, select the ACL whose statistics you want to view (for example, rule1).
3. Click **Statistics**.
4. View the statistics in the new window that opens.

Removing Extended ACLs

You can remove a single extended ACL or all extended ACLs.

To remove a single extended ACL by using the command line interface

At the command prompt, type:

- ♦ `rm ns acl <aclName>`
- ♦ `show ns acl`

To remove all extended ACLs by using the command line interface

At the command prompt, type:

- ♦ `clear ns acls`
- ♦ `show ns acl`

To remove a single extended ACL by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the details pane, on the **Extended ACLs** tab, select the ACL that you want to remove (for example, rule1).
3. Click **Remove**.
4. In the **Remove** dialog box, click **Yes**.

To remove all extended ACLs by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the details pane, on the **Extended ACLs** tab, click **Clear**.
3. In the **Clear ACL (s)** dialog box, click **Yes**.
4. In the details pane, on the **Extended ACLs** tab, verify that no ACLs are listed.

Configuring Simple ACL6s

A simple ACL6, which uses few parameters, cannot be modified once created. Instead, you must remove the simple ACL6 and create a new one. When creating a simple ACL6, you must specify its name, and a source IP address value against which to match packets. Optionally, you can specify a destination port and a time to live (TTL) value. A TTL is the number of seconds after which the simple ACL6 expires. ACL6s with TTLs are not saved when you save the configuration. Simple ACL6s can traverse the extension headers (if present) of all the incoming IPv6 packets to identify the layer 4 protocol and take a specified action.

Creating Simple ACL6s

To create a simple ACL6, you must specify its name and source IP address. You can also specify a destination port and time to live (TTL).

To create a simple ACL6 by using the command line interface

At the command prompt, type the following commands to create a simple ACL6 and verify the configuration:

- ♦ **add ns simpleacl6** <aclname> DENY -srcIPv6 <ipv6_addr|null> [-destPort<port> -protocol (TCP | UDP)] [-TTL <positive_integer>]
- ♦ **show ns simpleacl6** [<aclname>]

Example

```
> add ns simpleacl6 rule1 DENY -srcIPv6 3ffe:
192:168:215::82 -destPort 80 -Protocol TCP -TTL
9000
Done
```

To create a simple ACL6 by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the **ACLs** pane, on the **Simple ACL6s** tab, click **Add**.
3. In the **Add Simple ACL6** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create**, and then click **Close**.

To remove a single simple ACL6 by using the command line interface

At the command prompt, type:

- ♦ `rm ns simpleacl6 <aclname>`
- ♦ `show ns simpleacl6`

To remove all simple ACL6s by using the command line interface

At the command prompt, type:

- ♦ `clear ns simpleacl6`
- ♦ `show ns simpleacl6`

To remove one or all simple ACL6s by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the details pane, on the **Simple ACL6s** tab, do one of the following:
 - Select the simple ACL6 that you want to remove, and then click **Remove**.
 - To remove all simple ACL6s, click **Clear**.
3. In the **Proceed or Clear Simple ACL6(s)** dialog box, click **Yes**.

Monitoring Simple ACL6s

You can display the following simple ACL6 statistics:

Table 7-5. Simple ACL6 Statistics

Statistic	Indicates
Deny simpleACL6 hits	Packets dropped because they match a simple deny ACL6
Simple ACL6 hits	Packets matching a simple ACL6
Simple ACL6 misses	Packets not matching any simple ACL6
Simple ACL6 count	Number of simple ACL6s configured

To display simple ACL6 statistics by using the command line interface

At the command prompt, type:

```
stat ns simpleacl6
```

To display simple ACL6 statistics by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the details pane, on the **Simple ACL6s** tab, select the simple ACL6 whose statistics you want to display.
3. Click **Statistics**.

Configuring ACL6s

ACL6s can be configured during creation. Additionally, the following actions can be performed on ACL6s: Modify, Apply, Disable, Enable, Renumber and Remove the priority of ACL6s. Log files of ACL6s can be configured to collect statistics of packets. If a packet matches the condition defined by the ACL6, the NetScaler performs an action. If the packet does not match the condition defined by the ACL6, the NetScaler compares the packet against the ACL6 with the next-highest priority. ACL6s can traverse the extension headers (if present) of all the incoming IPv6 packets to identify the layer 4 protocol and take a specified action.

Creating and Modifying ACL6s

To create an ACL6 by using the command line interface

At the command prompt, type:

- ♦ **add ns acl6** <acl6name> <acl6action> [-srcIPv6 [<operator>] <srcIPv6Val>] [-srcPort [<operator>] <srcPortVal>] [-destIPv6 [<operator>] <destIPv6Val>] [-destPort [<operator>] <destPortVal>] [-TTL <positive_integer>] [-srcMac <mac_addr>] [(-protocol <protocol> [-established]) | -protocolNumber <positive_integer>] [-vlan <positive_integer>] [-interface <interface_name>] [-icmpType <positive_integer> [-icmpCode <positive_integer>]] [-priority <positive_integer>] [-state (ENABLED | DISABLED)]
- ♦ **show ns acl6** [<acl6name>]

Example

```
Example
> add ns acl6 rule6 DENY -srcport 45-1024 -
destIPv6 2001::45 -protocol TCP
Done
```

To modify or remove an ACL6 by using the command line interface

- ♦ To modify an ACL6, type the **set ns ACL6** command, the name of the ACL6, and the parameters to be changed, with their new values.
- ♦ To remove an ACL6, type the **rm ns ACL6** command and the name of the <entity>.

To configure an ACL6 by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the details pane, on the ACL6s tab, do one of the following:
 - To create a new ACL6, click **Add**.
 - To modify an existing ACL6, select the ACL6, and then click **Open**.
3. In the **Create ACL6** or **Configure ACL6** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create** or **OK**, and click **Close**.

Applying ACL6s

After you create an ACL6, you must activate it. The following procedures reapply all the ACL6s.

For example, if you have created the ACL6s rule1 through rule10, and then you create an ACL6 called rule11 and apply it, all of the ACL6s (rule1 through rule11) are applied afresh.

If a session has a DENY ACL related to it, the session is destroyed.

You must apply one of the following procedures after every action you perform on an ACL6 (for example, after disabling an ACL6). However, you can add or modify more than one ACL6 and apply all of them at the same time.

Note: ACL6s created on the NetScaler do not work until they are applied.

To apply ACL6s by using the command line interface

At the command prompt, type:

```
apply ns acls6
```

To apply ACL6s by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. Click **Commit**.
3. In the **Apply ACL(s)** dialog box that appears, click **Yes**.
4. Verify that the settings displayed on the **ACL6s** tab are correct.

Enabling and Disabling ACL6s

By default, ACL6s are enabled. Therefore, after the ACL6s are applied, the NetScaler appliance compares incoming packets against the configured ACL6s.

If an ACL6 is not required to be part of the lookup table but needs to be retained in the configuration, it must be disabled before the ACL6s are applied. After the ACL6s are applied, the NetScaler does not compare incoming packets against disabled ACL6s.

To disable or enable an ACL6 by using the command line interface

At the command prompt, type:

- ♦ **enable ns acl6** <acl6name>
- ♦ **show ns acl6** [<acl6name>]
- ♦ **disable ns acl6** <acl6name>
- ♦ **show ns acl6** [<acl6name>]

Note: ACL6s created on the NetScaler do not work until they are applied.

Example

```
> enable ns acl6 rule6
Done

> show ns acl6 rule6
      Name: rule6
Action: DENY
      srcIPv6
destIPv6 = 2001::45
      srcMac:
Protocol: TCP
      srcPort = 45-1024
destPort
      Vlan:
Interface:
      Active Status: ENABLED
Applied Status: NOTAPPLIED
      Priority: 10
Hits: 0
      TTL:
Done

> disable ns acl6 rule6
Done

> show ns acl6 rule6
      Name: rule6
Action: DENY
      srcIPv6
destIPv6 = 2001::45
```

```

srcMac:
Protocol: TCP
srcPort = 45-1024
destPort
Vlan:
Interface:
Active Status: DISABLED
Applied Status: NOTAPPLIED
Priority: 10
Hits: 0
TTL:
Done

```

To disable or enable an ACL6 by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the details pane, on the **ACL6s** tab, select the ACL (for example, rule1) and do one of the following:
 - To disable the extended ACL6, click **Disable**.
 - To enable the extended ACL6, click **Enable**.
3. If you want to apply the new setting, which reapplies all ACLs, click **Commit**, and then, in the **Apply ACL6(s)** dialog box, click **Yes**.
4. In the details pane, on the **Extended ACL6s** tab, view the list to verify the changed status under the column **Active Status**.

Renumbering the Priority of ACL6s

The renumber procedure resets the priorities of the ACL6s to multiples of 10. The priority (an integer value) defines the order in which the NetScaler appliance evaluates ACL6s. All priorities are multiples of 10, unless you configure a specific priority to an integer value. When you create an ACL6 without specifying a priority, the NetScaler automatically assigns a priority that is a multiple of 10.

If a packet matches the condition defined by the ACL6, the NetScaler performs an action. If the packet does not match the condition defined by the ACL6, the NetScaler compares the packet against the ACL6 with the next-highest priority.

Consider the following example. Two ACL6s, rule1 and rule2, are automatically assigned priorities 20 and 30 when they are created. You need to add a third ACL, rule3, to be evaluated immediately after rule1. Rule3 must have a priority between 20 and 30. In this case, you can specify the priority as 25. Later, you can easily renumber the ACL6s with priorities that are multiples of 10, without affecting the order in which the ACL6s are applied.

To renumber the priorities of the ACL6s by using the command line interface

At the command prompt, type:

```
renumber ns acls6
```

Example

```
> renumber ns acl6
Done
```

To renumber the priority of ACL6s by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the details pane, on the ACL6s tab, click **Renumber Priority (s) ACL(s)**.
3. In the **Renumber Priority (s) ACL(s)** dialog box, click **Yes**.
4. Verify the action in the details pane.

Monitoring ACL6s

You can display statistics for monitoring the performance of an ACL6.

To display the statistics for an ACL6s by using the command line interface

At the command prompt, type:

```
stat ns acl6 <acl6name>
```

The following table lists the statistics associated with ACL6s and their descriptions.

Table 7-6. ACL6 Statistics

Statistic	Specifies
Allow ACL6 hits	Packets matching IPv6 ACLs with processing mode set to ALLOW. The NetScaler processes these packets.
NAT ACL6 hits	Packets matching a NAT ACL6, resulting in a NAT session.
Deny ACL6 hits	Packets dropped because they match IPv6 ACLs with processing mode set to DENY.
Bridge ACL6 hits	Packets matching a bridge IPv6 ACL, which in transparent mode bypasses service processing.
ACL6 hits	Packets matching an IPv6 ACL.
ACL6 misses	Packets not matching any IPv6 ACL.

To display the statistics for an ACL6 by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the details pane, on the **ACL6s** tab, select the ACL whose statistics you want to view (for example, rule1).
3. Click **Statistics**.
4. View the statistics in the new window that opens.

Removing ACL6s

You can remove a single ACL6 or all ACL6s.

To remove an extended ACL6 by using the command line interface

At the command prompt, type:

- ♦ `rm ns acl6 <acl6name>`
- ♦ `show ns acl6`

To remove all extended ACL6s by using the command line interface

At the command prompt, type:

`clear ns acls6`

To remove an extended ACL6 by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the details pane, on the **ACL6s** tab, select the ACL that you want to remove (for example, rule1).
3. Click **Remove**.
4. In the **Remove** dialog box, click **Yes**.
5. In the details pane, on the **ACL6s** tab, verify that the ACL6 is not listed

To remove all extended ACLs by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the details pane, on the **ACL6s** tab, click **Clear**.
3. In the **Clear ACL (s)** dialog box, click **Yes**.
4. In the details pane, on the **Extended ACLs** tab, verify that no ACLs are listed.

Terminating Established Connections

For a simple ACL, the NetScaler appliance blocks any new connections that match the conditions specified in the ACL. The appliance does not block any packets related to existing connections that were established before the ACL was created.

However, you can immediately terminate the established connections by running a flush operation from the command line interface or the configuration utility.

Flush can be useful in the following cases:

- ♦ You receive a list of blacklisted IP addresses and want to completely block those IP addresses from accessing the NetScaler appliance. In this case, you create simple ACLs to block any new connections from these IP addresses, and then run flush to terminate any existing connections.
- ♦ You want to terminate a large number of connections from a particular network without taking the time to terminate them one by one.

When you run flush, the appliance searches through all of its established connections and terminates those that match conditions specified in any of the simple ACLs configured on the appliance.

Note: If you plan to create more than one simple ACL and flush existing connections that match any of them, you can minimize the effect on performance by first creating all of the simple ACLs and then running flush only once.

To terminate all established IPv4 connections that match any of your configured simple ACLs by using the command line interface

At the command prompt, type:

```
flush simpleacl -estSessions
```

To terminate all established IPv4 connections that match any of your configured simple ACLs by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the ACLs pane, on the Simple ACLs tab, click **Flush**.

To terminate all established IPv6 connections that match any of your configured simple ACL6s by using the command line interface

At the command prompt, type:

```
flush simpleacl6 -estSessions
```

To terminate all established IPv6 connections that match any of your configured simple ACL6s by using the configuration utility

1. Navigate to **System > Network > ACLs**.
2. In the ACLs pane, on the Simple ACL6s tab, click **Flush**.

IP Routing

NetScaler appliances support both dynamic and static routing. Because simple routing is not the primary role of a NetScaler, the main objective of running dynamic routing protocols is to enable route health injection (RHI), so that an upstream router can choose the best among multiple routes to a topographically distributed virtual server.

Most NetScaler implementations use some static routes to reduce routing overhead. You can create backup static routes and monitor routes to enable automatic switchover in the event that a static route goes down. You can also assign weights to facilitate load balancing among static routes, create null routes to prevent routing loops, and configure IPv6 static routes. You can configure policy based routes (PBRs), for which routing decisions are based on criteria that you specify.

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

Configuring Static Routes

Static routes are manually created to improve the performance of your network. You can monitor static routes to avoid service disruptions. Also, you can assign weights to ECMP routes, and you can create null routes to prevent routing loops.

Weighted Static Routes

When the NetScaler appliance makes routing decisions involving routes with equal distance and cost, that is, Equal Cost Multi-Path (ECMP) routes, it balances the load between them by using a hashing mechanism based on the source and destination IP addresses. For an ECMP route, however, you can configure a weight value. The NetScaler then uses both the weight and the hashed value for balancing the load.

Null Routes

If the route chosen in a routing decision is inactive, the NetScaler appliance chooses a backup route. If all the backup routes become inaccessible, the appliance might reroute the packet to the sender, which could result in a routing loop leading to network congestion. To prevent this situation, you can create a null route, which adds a null interface as a gateway. The null route is never the preferred route, because it has a higher administrative distance than the other static routes. But it is selected if the other static routes become inaccessible. In that case, the appliance drops the packet and prevents a routing loop.

Configuring IPv4 Static Routes

You can add a simple static route or a null route by setting a few parameters, or you can set additional parameters to configure a monitored or monitored and weighted static route. You can change the parameters of a static route. For example, you might want to assign a weight to an unweighted route, or you might want to disable monitoring on a monitored route.

Note: Monitored static route is not supported on NetScaler 1000V.

To create a static route by using the command line interface

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

- ♦ **add route** <network> <netmask> <gateway>[-cost <positive_integer>] [-advertise (DISABLED | ENABLED)]
- ♦ **show route** [<network> <netmask> [<gateway>]] [<routeType>] [-detail]

Example

```
> add route 10.102.29.0 255.255.255.0 10.102.29.2 -  
cost 2 -advertise ENABLED  
Done
```

To create a null route by using the command line interface

At the command prompt type:

- ♦ **add route** <network> <netmask> null
- ♦ **show route** <network> <netmask>

Example

```
> add route 10.102.29.0 255.255.255.0 null  
Done
```

To remove a static route by using the command line interface

At the command prompt, type:

rm route <network> <netmask> <gateway>

Example

```
> rm route 10.102.29.0 255.255.255.0 10.102.29.3  
Done
```

To configure a static route by using the configuration utility

1. Navigate to **System > Network > Routes**.
2. In the details pane, on the **Basic** tab, do one of the following:
 - To create a new static route, click **Add**.
 - To modify an existing static route, click **Open**.
3. In the **Create Route** or **Configure Route** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create** or **OK**, and then click **Close**.

To remove a route by using the configuration utility

1. Navigate to **System > Network > Routes**.
2. On the **Routes** pane, click the **Basic** tab, select the route you want to remove (for example, **192.168. 20.2**), and then click **Remove**.
3. In the **Remove** dialog box, click **Yes**.

Configuring IPv6 Static Routes

You can configure a maximum of six default IPv6 static routes. IPv6 routes are selected on the basis of whether the MAC address of the destination device is reachable. This can be determined by using the IPv6 Neighbor Discovery feature. Routes are load balanced and only source/destination-based hash mechanisms are used. Therefore, route selection mechanisms such as round robin are not supported. The next hop address in the default route need not belong to the NSIP subnet.

Note: Monitored static route is not supported on NetScaler 1000V.

To create an IPv6 route by using the command line interface

At the command prompt, type the following commands to create an IPv6 route and verify the configuration:

- ♦ **add route6** <network> <gateway> [-vlan <positive_integer>]
- ♦ **show route6** [<network> [<gateway>]]

Example

```
> add route6 ::/0 FE80::67 -vlan 5
Done
```

To remove an IPv6 route by using the command line interface

At the command prompt, type:

rm route6 <network> <gateway>

Example

```
> rm route6 ::/0 FE80::67
Done
```

To configure an IPv6 route by using the configuration utility

1. Navigate to **System > Network > Routes**.
2. In the details pane, on the **IPv6** tab, do one of the following:
 - To create a new route, click **Add**.
 - To modify an existing route, click **Open**.
3. In the **Create IPv6 Route** or **Configure IPv6 Route** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create** or **OK**, and then click **Close**.

To remove an IPv6 route by using the configuration utility

1. Navigate to **System > Network > Routes**.
2. On the **Routes** pane, click the **IPv6** tab.
3. Select the network, from which you want to remove the route (for example, **::/0**), and then click **Remove**.
4. In the **Remove** dialog box, click **Yes**.

Configuring Policy-Based Routes

Policy-based routing bases routing decisions on criteria that you specify. A policy-based route (PBR) specifies criteria for selecting packets and, typically, a next hop to which to send the selected packets. For example, you can configure the NetScaler appliance to route outgoing packets from a specific IP address or range to a particular next hop router. Each packet is matched against each configured PBR, in the order determined by the specified priorities, until a match is found. If no match is found, or if the matching PBR specifies a DENY action, the NetScaler applies the routing table for normal destination-based routing.

A PBR bases routing decisions for the data packets on parameters such as source IP address, source port, destination IP address, destination port, protocol, and source MAC address. A PBR defines the conditions that a packet must satisfy for the NetScaler to route the packet. These actions are known as "processing modes." The processing modes are:

- ♦ ALLOW - The NetScaler sends the packet to the designated next-hop router.
- ♦ DENY - The NetScaler applies the routing table for normal destination-based routing.

The NetScaler process PBRs before processing the RNAT rules.

You can create PBRs for outgoing IPv4 and IPv6 traffic.

Many users begin by creating PBRs and then modifying them. To activate a new PBR, you must apply it. To deactivate a PBR, you can either remove or disable it. You can change the priority number of a PBR to give it a higher or lower precedence.

Configuring a Policy-Based Routes (PBR) for IPv4 Traffic

Configuring PBRs involves the following tasks:

- ♦ Create a PBR.
- ♦ Apply PBRs.
- ♦ (Optional) Disable or enable a PBR.
- ♦ (Optional) Renumber the priority of the PBR.

Creating or Modifying a PBR

You cannot create two PBRs with the same parameters. If you attempt to create a duplicate, an error message appears.

You can configure the priority of a PBR. The priority (an integer value) defines the order in which the NetScaler appliance evaluates PBRs. When you create a PBR without specifying a priority, the NetScaler automatically assigns a priority that is a multiple of 10.

If a packet matches the condition defined by the PBR, the NetScaler performs an action. If the packet does not match the condition defined by the PBR, the NetScaler compares the packet against the PBR with the next highest priority.

Instead of sending the selected packets to a next hop router, you can configure the PBR to send them to a link load balancing virtual server to which you have bound multiple next hops. This configuration can provide a backup if a next hop link fails.

Consider the following example. Two PBRs, p1 and p2, are configured on the NetScaler and automatically assigned priorities 20 and 30. You need to add a third PBR, p3, to be evaluated immediately after the first PBR, p1. The new PBR, p3, must have a priority between 20 and 30. In this case, you can specify the priority as 25.

To create a PBR by using the command line interface

At the command prompt, type:

- ♦ **add ns pbr** <name> <action> [-srcIP [<operator>] <srcIPVal>] [-srcPort [<operator>] <srcPortVal>] [-destIP [<operator>] <destIPVal>] [-destPort [<operator>] <destPortVal>] [-nextHop <nextHopVal>] [-srcMac <mac_addr>] [-protocol <protocol> | -protocolNumber <positive_integer>] [-vlan <positive_integer>] [-interface <interface_name>] [-priority <positive_integer>] [-state (ENABLED | DISABLED)]

- ◆ **sh ns pbr**

Example

```
> add ns pbr pbr1 allow -srcip 10.102.37.252 -
destip 10.10.10.2 -nexthop 10.102.29.77
Done
```

To modify the priority of a PBR by using the command line interface

At the command prompt, type the following commands to modify the priority and verify the configuration:

- ◆ **set ns pbr <name> [-action (ALLOW | DENY)] [-srcIP [<operator>] <srcIPVal>] [-srcPort [<operator>] <srcPortVal>] [-destIP [<operator>] <destIPVal>] [-destPort [<operator>] <destPortVal>] [-nextHop <nextHopVal>] [-srcMac <mac_addr>] [-protocol <protocol> | -protocolNumber <positive_integer>] [-vlan <positive_integer>] [-interface <interface_name>] [-priority <positive_integer>] [-state (ENABLED | DISABLED)]**
- ◆ **show ns pbr [<name>]**

Example

```
> set ns pbr pbr1 -priority 23
Done
```

To remove one or all PBRs by using the command line interface

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

- ◆ **rm ns pbr <name>**
- ◆ **clear ns PBRs**

Example

```
> rm ns pbr pbr1
Done
> clear ns PBRs
Done
```

To create a PBR by using the configuration utility

1. Navigate to **System > Network > PBRs**.

2. In the details pane, do one of the following:
 - To create a new PBR, click **Add**.
 - To modify an existing PBR, click **Open**.
3. In the **Create PBR** or **Configure PBR** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create** or **OK**, and then click **Close**.

To remove one or all PBRs by using the configuration utility

1. Navigate to **System > Network > PBRs**.
2. To remove a single PBR, in the details pane, select the PBR that you want to remove (for example, **p1**), and then click **Remove**.
3. To remove all PBRs, click **Clear**.

Applying a PBR

You must apply a PBR to activate it. The following procedure reapplies all PBRs that you have not disabled. The PBRs constitute a memory tree (lookup table). For example, if you create 10 PBRs (p1 - p10), and then you create another PBR (p11) and apply it, all of the PBRs (p1 - p11) are freshly applied and a new lookup table is created. If a session has a DENY PBR related to it, the session is destroyed.

You must apply this procedure after every modification you make to any PBR. For example, you must follow this procedure after disabling a PBR.

Note: PBRs created on the NetScaler appliance do not work until they are applied.

To apply a PBR by using the command line interface

At the command prompt, type:

```
apply ns PBRs
```

To apply a PBR by using the configuration utility

1. Navigate to **System > Network > PBRs**.
2. In the details pane, select the PBR that you want to apply (for example, **p1**).
3. Click **Commit**.
4. In the **Apply PBR(s)** dialog box, click **Yes**.

Enabling or Disabling PBRs

By default, the PBRs are enabled. This means that when PBRs are applied, the NetScaler appliance automatically compares incoming packets against the configured PBRs. If a PBR is not required in the lookup table, but it needs to be retained in the configuration, it must be disabled before the PBRs are applied. After the PBRs are applied, the NetScaler does not compare incoming packets against disabled PBRs.

To enable or disable a PBR by using the command line interface

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

- ♦ **enable ns pbr <name>**
- ♦ **disable ns pbr <name>**

Examples

```
> enable ns PBR pbr1
Done
> show ns PBR pbr1
1)      Name: pbr1
        Action: ALLOW
Hits: 0
        srcIP = 10.102.37.252
        destIP = 10.10.10.2
        srcMac:
Protocol:
        Vlan:
Interface:
        Active Status: ENABLED
Applied Status: APPLIED
        Priority: 10
        NextHop: 10.102.29.77

Done

> disable ns PBR pbr1
Warning: PBR modified, use 'apply pbrs' to commit
this operation
> apply pbrs
Done
> show ns PBR pbr1
1)      Name: pbr1
        Action: ALLOW
Hits: 0
        srcIP = 10.102.37.252
        destIP = 10.10.10.2
        srcMac:
Protocol:
        Vlan:
Interface:
        Active Status: DISABLED
Applied Status: NOTAPPLIED
        Priority: 10
        NextHop: 10.102.29.77

Done
```

To enable or disable a PBR by using the configuration utility

1. Navigate to **System > Network > PBRs**.
2. In the details pane, select the PBR (for example, p1) and do one of the following:
 - To enable the PBR, click **Enable**.
 - To disable the PBR, click **Disable**.

A message appears in the status bar, stating that the PBR has been successfully enabled or disabled.

Renumbering PBRs

You can automatically renumber the PBRs to set their priorities to multiples of 10.

To renumber PBRs by using the command line interface

At the command prompt, type:

```
renumber ns pbrs
```

To renumber PBRs by using the configuration utility

1. Navigate to **System > Network > PBRs**.
2. In the details pane, click **Renumber Priority (s)**.
3. In the **Renumber Priority(s) PBR(s)** dialog box, click **Yes**.

Use Case - PBR with Multiple Hops

Consider a scenario in which two PBRs, PBR1 and PBR2, are configured on NetScaler appliance NS1. PBR1 routes all the outgoing packets, with source IP address as 10.102.29.30, to next hop router R1. PBR2 routes all the outgoing packets, with source IP address as 10.102.29.90, to next hop router R2. R3 is another next hop router connected to NS1.

If router R1 fails, all the outgoing packets that matched against PBR1 are dropped. To avoid this situation, you can specify a link load balancing (LLB) virtual server in the next hop field while creating or modifying a PBR. Multiple next hops are bound to the LLB virtual server as services (for example R1, R2, and R3). Now, if R1 fails, all the packets that matched against PBR1 are routed to R2 or R3 as determined by the LB method configured on the LLB virtual server.

The NetScaler appliance throws an error if you attempt to create a PBR with an LLB virtual server as the next hop in the following cases:

- ♦ Adding another PBR with the same LLB virtual server.
- ♦ Specifying a nonexistent LLB virtual server.
- ♦ Specifying an LLB virtual server for which the bound services are not next hops.
- ♦ Specifying an LLB virtual server for which the LB method is not set to one of the following:

- LEASTPACKETS
- LEASTBANDWIDTH
- DESTIPHASH
- SOURCEIPHASH
- WEIGHTDRR
- SRCIPDESTIP_HASH
- LTRM
- CUSTOM LOAD
- ♦ Specifying an LLB virtual server for which the LB persistence type is not set to one of the following:
 - DESTIP
 - SOURCEIP
 - SRCDSTIP

The following table lists the names and values of the entities configured on the NetScaler appliance:

Table 7-7. Sample Values for Creating Entities

Entity Type	Name	IP Address
Link load balancing virtual server	LLB1	NA
Services (next hops)	Router1	1.1.1.254
	Router2	2.2.2.254
	Router3	3.3.3.254
PBRs	PBR1	NA
	PBR2	NA

To implement the configuration described above, you need to:

1. Create services Router1, Router2, and Router3 that represent next hop routers R1, R2, and R3.
2. Create link load balancing virtual server LLB1 and bind services Router1, Router2, and Router3 to it.
3. Create PBRs PBR1 and PBR2, with next hop fields set as LLB1 and 2.2.2.254 (IP address of the router R2), respectively.

To create a service by using the command line interface

At the command prompt, type:

- ♦ **add service** <name> <IP> <serviceType> <port>
- ♦ **show service** <name>

Example

```
> add service Router1 1.1.1.254 ANY *
Done
> add service Router2 2.2.2.254 ANY *
Done
> add service Router3 3.3.3.254 ANY *
Done
```

To create services by using the configuration utility

1. Navigate to **Traffic Management > Load Balancing > Services**.
2. In the details pane, click **Add**.
3. In the **Create Service** dialog box, specify values for the following parameters:
 - **Service Name***—name
 - **Server**—IP
 - **Protocol***—serviceType (Select **ANY** from the drop-down list.)
 - **Port***—port

* A required parameter
4. Click **Create**.
5. Repeat Steps 2-4 to create another service.
6. Click **Close**.
7. In the **Services** pane, select the services that you just configured and verify that the settings displayed at the bottom of the screen are correct.

To create a link load balancing virtual server and bind a service by using the command line interface

At the command prompt, type:

- ♦ **add lb vserver** <name> <serviceType>
- ♦ **bind lb vserver** < name> <serviceName>
- ♦ **show lb vserver** < name>

Example

```
> add lb vserver LLB1 ANY
Done
```

```
> bind lb vserver LLB1 Router1 Router2 Router3
Done
```

To create a link load balancing virtual server and bind a service by using the configuration utility

1. Navigate to **Traffic Management > Load Balancing > Virtual Servers**.
2. In the **Load Balancing Virtual Servers** pane, click **Add**.
3. In the **Create Virtual Servers (Load Balancing)** dialog box, specify values for the following parameters:

- **Name***—name
- **Protocol***—serviceType (Select **ANY**.)

* A required parameter

Note: Make sure **Directly Addressable** is unchecked.

4. Under the **Services** tab, in the **Active** column, select the check box for the service that you want to bind to the virtual server.
5. Click **Create**, and then click **Close**.
6. In the **Load Balancing Virtual Servers** tab, select the virtual server that you just created, and verify that the settings displayed in the **Details** pane are correct.

To create a PBR by using the command line interface

At the command prompt, type:

- ♦ **add ns pbr** <name> <action> [-srcIP [<operator>] <srcIPVal>] [-nextHop <nextHopVal>]
- ♦ **show ns pbr**

Example

```
> add pbr PBR1 ALLOW -srcIP 10.102.29.30 -nextHop
LLB1
Done
> add pbr PBR2 ALLOW -srcIP 10.102.29.90 -nextHop
2.2.2.254
Done
```

To create a PBR by using the configuration utility

1. Navigate to **System > Network > PBRs**.
2. In the details pane, do one of the following:

- To create a new PBR, click **Add**.
 - To modify an existing PBR, click **Open**.
3. In the **Create PBR** or **Configure PBR** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
 4. Click **Create** or **OK**, and then click **Close**.

Configuring a Policy-Based Routes (PBR6) for IPv6 Traffic

Configuring PBR6s involves the following tasks:

- ♦ Create a PBR6.
- ♦ Apply PBR6s.
- ♦ (Optional) Disable or enable a PBR6.
- ♦ (Optional) Renumber the priority of the PBR6.

Creating or Modifying a PBR6

You cannot create two PBR6s with the same parameters. If you attempt to create a duplicate, an error message appears.

You can configure the priority of a PBR6. The priority (an integer value) defines the order in which the NetScaler appliance evaluates PBR6s. When you create a PBR6 without specifying a priority, the NetScaler automatically assigns a priority that is a multiple of 10.

If a packet matches the condition defined by the PBR6, the NetScaler performs an action. If the packet does not match the condition defined by the PBR6, the NetScaler compares the packet against the PBR6 with the next highest priority.

To create a PBR6 by using the command line interface

At the command prompt, type:

- ♦ **add ns pbr6** <name> <action> [-srcIPv6 [<operator>] <srcIPv6Val>] [-srcPort [<operator>] <srcPortVal>] [-destIPv6 [<operator>] <destIPv6Val>] [-destPort [<operator>] <destPortVal>] [-srcMac <mac_addr>] [-protocol <protocol> | -protocolNumber <positive_integer>] [-vlan <positive_integer>] [-interface <interface_name>] [-priority <positive_integer>] [-state (ENABLED | DISABLED)][[-nextHop <nextHopVal>] [-nextHopVlan <positive_integer>]
- ♦ **show ns pbr**

To modify or remove a PBR6 by using the command line interface

To modify a PBR6, type the **set pbr6** <name> command and the parameters to be changed, with their new values.

To remove one or all PBR6s by using the command line interface

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

- ♦ `rm ns pbr6 <name>`
- ♦ `clear ns pbr6`

To create or modify a PBR6 by using the configuration utility

1. Navigate to **System > Network > PBRs**.
2. On the **PBR6s** tab, do one of the following:
 - To create a new PBR6, click **Add**.
 - To modify an existing PBR6, click **Open**.
3. In the **Create PBR6** or **Configure6 PBR** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. Click **Create** or **OK**, and then click **Close**. A message appears in the status bar, stating that the PBR6 has been configured successfully.

To remove one or all PBR6s by using the configuration utility

1. Navigate to **System > Network > PBRs**.
2. To remove a single PBR6, on the **PBR6s** tab, select the PBR6 that you want to remove, and then click **Remove**.
3. To remove all PBR6s, click **Clear**.

Applying PBR6s

You must apply a PBR6 to activate it. The following procedure reapplies all PBR6s that you have not disabled. The PBR6s constitute a memory tree (lookup table). For example, if you create 10 PBR6s (p6_1 - p6_10), and then you create another PBR6 (p6_11) and apply it, all of the PBR6s (p6_1 - p6_11) are freshly applied and a new lookup table is created. If a session has a DENY PBR6 related to it, the session is destroyed.

You must apply this procedure after every modification you make to any PBR6. For example, you must follow this procedure after disabling a PBR6.

Note: PBR6s created on the NetScaler appliance do not work until they are applied.

To apply PBR6s by using the command line interface

At the command prompt, type:

```
apply ns PBR6
```

To apply PBR6s by using the configuration utility

1. Navigate to **System > Network > PBRs**.
2. On the **PBR6s** tab, click **Commit**.
3. In the **Confirm** dialog box, click **Yes**.

Enabling or Disabling a PBR6

By default, the PBR6s are enabled. This means that when PBR6s are applied, the NetScaler appliance automatically compares outgoing IPv6 packets against the configured PBR6s. If a PBR6 is not required in the lookup table, but it needs to be retained in the configuration, it must be disabled before the PBR6s are applied. After the PBR6s are applied, the NetScaler does not compare incoming packets against disabled PBR6s.

To enable or disable a PBR6 by using the command line interface

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

- ♦ **enable ns pbr <name>**
- ♦ **disable ns pbr <name>**

To enable or disable a PBR6 by using the configuration utility

1. Navigate to **System > Network > PBR6s**.
2. On the **PBR6s** tab, select the PBR6 (for example, **p1_6**) and do one of the following:
 - To enable the PBR6, click **Enable**.
 - To disable the PBR6, click **Disable**.

A message appears in the status bar, stating that the PBR6 has been successfully enabled or disabled.

Renumbering PBR6s

You can automatically renumber the PBR6s to set their priorities to multiples of 10.

To renumber PBR6s by using the command line interface

At the command prompt, type:

```
renumber ns pbr6
```

To renumber PBR6s by using the configuration utility

1. Navigate to **System > Network > PBRs**.
2. On the **PBR6s** tab, **Renumber Priority (s)**.
3. In the **Renumber Priority(s) PBR(s)** dialog box, click **Yes**.

Internet Protocol version 6 (IPv6)

A NetScaler appliance supports both server-side and client-side IPv6 and can therefore function as an IPv6 node. It can accept connections from IPv6 nodes (both hosts and routers) and from IPv4 nodes, and can perform Protocol Translation (RFC 2765) before

sending traffic to the services. You have to license the IPv6 feature before you can implement it.

The following table lists some of the IPv6 features that the NetScaler appliance supports.

Table 7-8. Some Supported IPv6 Features

IPv6 features
IPv6 addresses for SNIPs (NSIP6, VIP6, and SNIP6)
Neighbor Discovery (Address Resolution, Duplicated Address Detection, Neighbor Unreachability Detection, Router Discovery)
Management Applications (ping6, telnet6, ssh6)
Static Routing and Dynamic routing (OSPF)
Port Based VLANs
Access Control Lists for IPv6 addresses (ACL6)
IPv6 Protocols (TCP6, UDP6, ICMP6)
Server Side Support (IPv6 addresses for vservers, services)
USIP (Use source IP) and DSR (Direct Server Return) for IPv6
SNMP and CVPN for IPv6
HA with native IPv6 node address
IPv6 addresses for MIPs
Path-MTU discovery for IPv6

The following table lists NetScaler components that support IPv6 addresses and provides references to the PDF documentation of the components.

Table 7-9. NetScaler Components That Support IPv6 Addresses and the Corresponding Documentation

NetScaler component	Section that documents IPv6 support	Document title
Network	Adding, Customizing, Removing, Removing all, and Viewing routes.	Citrix NetScaler Networking Guide
SSL Offload	Creating IPv6 vservers for SSL Offload	Citrix NetScaler Traffic Management Guide

NetScaler component	Section that documents IPv6 support	Document title
SSL Offload	Specifying IPv6 SSL Offload Monitors	Citrix NetScaler Traffic Management Guide
SSL Offload	Creating IPv6 SSL Offload Servers	Citrix NetScaler Traffic Management Guide
Load Balancing	Creating IPv6 vservers for Load Balancing	Citrix NetScaler Traffic Management Guide
Load Balancing	Specifying IPv6 Load Balancing Monitors	Citrix NetScaler Traffic Management Guide
Load Balancing	Creating IPv6 Load Balancing Servers	Citrix NetScaler Traffic Management Guide
DNS	Creating AAAA Records	Citrix NetScaler Traffic Management Guide

You can configure IPv6 support for the above features after implementing the IPv6 feature on your NetScaler appliance. You can configure both tagged and prefix-based VLANs for IPv6. You can also map IPv4 addresses to IPv6 addresses.

Implementing IPv6 Support

IPv6 support is a licensed feature, which you have to enable before you can use or configure it. If IPv6 is disabled, the NetScaler does not process IPv6 packets. It displays the following warning when you run an unsupported command:

"Warning: Feature(s) not enabled [IPv6PT]"

The following message appears if you attempt to run IPv6 commands without the appropriate license:

"ERROR: Feature(s) not licensed"

After licensing the feature, use either of the following procedures to enable or disable IPv6.

To enable or disable IPv6 by using the command line interface

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

- ♦ `enable ns feature ipv6pt`
- ♦ `disable ns feature ipv6pt`

To enable or disable IPv6 by using the configuration utility

1. Navigate to **System > Settings**.
2. On the **Settings** page, under **Modes and Features**, click **change advanced features**.
3. In the **Configure Advanced Features** dialog box, do one of the following:
 - To enable IPv6, select the **IPv6 Protocol Translation** check box.
 - To disable IPv6, clear the **IPv6 Protocol Translation** check box.
4. Click **OK**.
5. In the **Enable/Disable Feature(s)?** dialog box, click **Yes**.

VLAN Support

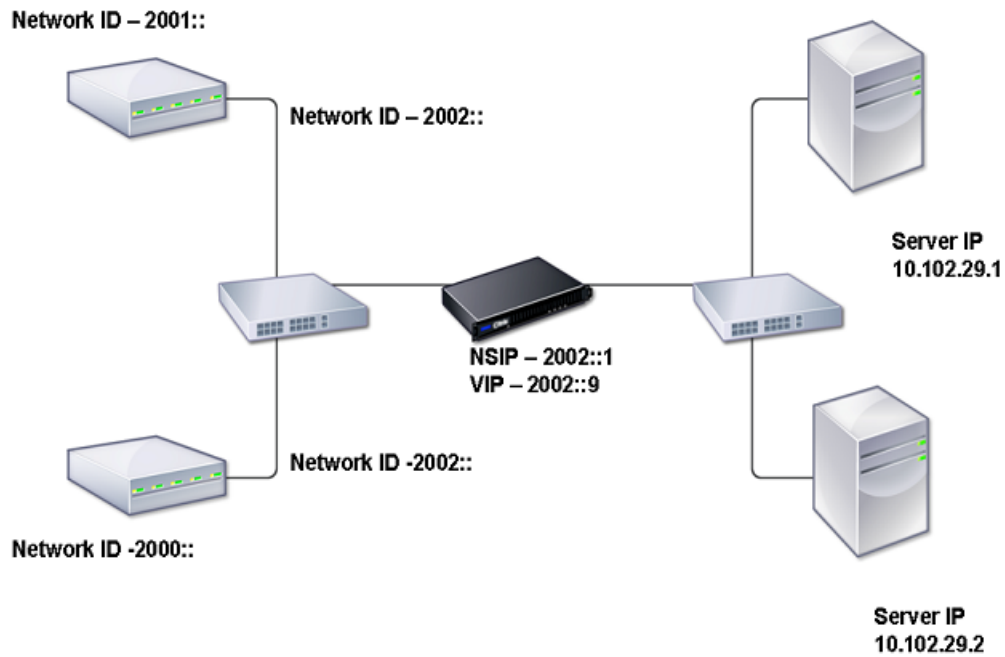
If you need to send broadcast or multicast packets without identifying the VLAN (for example, during DAD for NSIP, or ND6 for the next hop of the route), you can configure the NetScaler appliance to send the packet on all the interfaces with appropriate tagging. The VLAN is identified by ND6, and a data packet is sent only on the VLAN.

For more information about ND6 and VLANs, see "[Configuring Neighbor Discovery](#)."

Port-based VLANs are common for IPv4 and IPv6. Prefix-based VLANs are supported for IPv6.

Simple Deployment Scenario

Following is an example of a simple load balancing set-up consisting of an IPv6 vserver and IPv4 services, as illustrated in the following topology diagram.

Figure 7-18. IPv6 Sample Topology

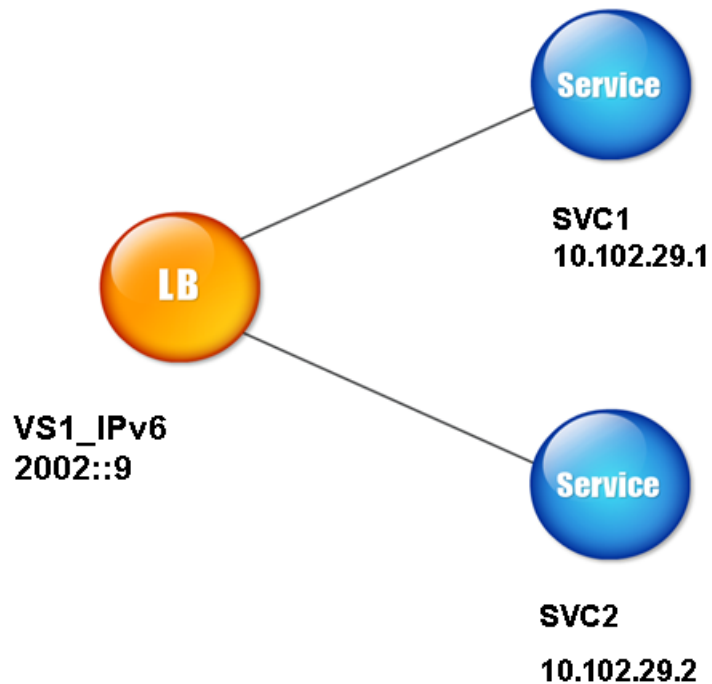
The following table summarizes the names and values of the entities that must be configured on the NetScaler.

Table 7-10. Sample Values for Creating Entities

Entity type	Name	Value
LB Vserver	VS1_IPv6	2002::9
Services	SVC1	10.102.29.1
	SVC2	10.102.29.2

The following figure shows the entities and values of the parameters to be configured on the NetScaler.

Figure 7-19. IPv6 Entity Diagram



To configure this deployment scenario, you need to do the following:

1. Create an IPv6 service.
2. Create an IPv6 LB vserver.
3. Bind the services to the vserver.

To create IPv4 services by using the command line interface

At the command prompt, type:

add service <Name> <IPAddress> <Protocol> <Port>

Example

```
add service SVC1 10.102.29.1 HTTP 80
add service SVC2 10.102.29.2 HTTP 80
```

To create IPv4 services by using the configuration utility

1. Navigate to **Traffic Management > Load Balancing > Services**.
2. On the **Services** page, click **Add**.

3. In the **Create Service** dialog box, in the **Service Name**, **Server**, and **Port** text boxes, type the name, IP address, and port of the service (for example, SVC1, 10.102.29.1, and 80).
4. In the **Protocol** drop-down list box, select the type of the service (for example, HTTP).
5. Click **Create** and click **Close**.
6. Repeat Steps 1-5 to create a service SVC2 with IP address 10.102.29.2 and port 80.

To create IPv6 vserver by using the command line interface

At the command prompt, type:

```
add lb vserver <Name> <IPAddress> <Protocol> <Port>
```

Example

```
add lb vserver VS1_IPv6 2002::9 HTTP 80
```

To create IPv6 vserver by using the configuration utility

1. Navigate to **Traffic Management > Load Balancing > Virtual Servers**.
2. In the **Load Balancing Virtual Servers** page, click **Add**.
3. In the **Create Virtual Servers (Load Balancing)** dialog box, select the **IPv6** check box.
4. In the **Name**, **Port**, and **IP Addresses** text boxes, type the name, port, and IP address of the vserver (for example, VS1_IPv6, 80, and 2002::9).
5. Click **Create** and click **Close**.

To bind a service to an LB vserver by using the command line interface

At the command prompt, type:

```
bind lb vserver <name> <service>
```

Example

```
bind lb vserver VS1_IPv6 SVC1
```

The vservers receive IPv6 packets and the NetScaler performs Protocol Translation (RFC 2765) before sending traffic to the IPv4-based services.

To bind a service to an LB vserver by using the configuration utility

1. Navigate to **Traffic Management > Load Balancing > Virtual Servers**.
2. In the **Load Balancing Virtual Servers** page, select the vserver for which you want to bind the service (for example, VS1_IPv6).
3. Click **Open**.
4. In the **Configure Virtual Server (Load Balancing)** dialog box, on the **Services** tab, select the **Active** check box corresponding to the service that you want to bind to the vserver (for example, SVC1).
5. Click **OK**.
6. Repeat Steps 1-4 to bind the service (for example, SVC2 to the vserver).

Host Header Modification

When an HTTP request has an IPv6 address in the host header, and the server does not understand the IPv6 address, you must map the IPv6 address to an IPv4 address. The IPv4 address is then used in the host header of the HTTP request sent to the vserver.

To change the IPv6 address in the host header to an IPv4 address by using the command line interface

At the command prompt, type:

```
set ns ip6 <IPv6Address> -map <IPAddress>
```

Example

```
set ns ip6 2002::9 -map 200.200.200.200
```

To change the IPv6 address in the host header to an IPv4 address by using the configuration utility

1. Navigate to **System > Network > IPs**.
2. In the **IPs** page, click the **IPv6s** tab and select the IP address for which you want to configure a mapped IP address, for example, 2002:0:0:0:0:0:9.
3. Click **Open**.
4. In the **Configure IP6** dialog box, in the **Mapped IP** text box, type the mapped IP address that you want to configure, for example, 200.200.200.200.
5. Click **OK**.

VIP Insertion

If an IPv6 address is sent to an IPv4-based server, the server may not understand the IP address in the HTTP header, and may generate an error. To avoid this, you can map an IPv4 address to the IPv6 VIP and enable VIP insertion.

To configure a mapped IPv6 address by using the command line interface

At the command prompt, type:

```
set ns ip6 <IPv6Address> -map <IPAddress>
```

Example

```
> set ns ip6 2002::9 -map 200.200.200.200
Done
```

To configure a mapped IPv6 address by using the configuration utility

1. Navigate to **System > Network > IPs**.
2. In the **IPs** page, click the **IPv6s** tab and select the IP address for which you want to configure a mapped IP address (for example, 2002:0:0:0:0:0:9).
3. Click **Open**.
4. In the **Configure IP6** dialog box, in the **Mapped IP** text box, type the mapped IP address that you want to configure (for example, 200.200.200.200).
5. Click **OK**.

Use either of the following procedures to enable insertion of an IPv4 VIP address and port number in the HTTP requests sent to the servers.

To enable VIP insertion by using the command line interface

At the command prompt, type:

```
set lb vserver <name> -insertVserverIPPort <Value>
```

Example

```
> set lb vserver VS1_IPv6 -insertVserverIPPort ON
Done
```

To enable VIP insertion by using the configuration utility

1. Navigate to **Traffic Management > Load Balancing > Virtual Servers**.
2. In the **Load Balancing Virtual Servers** page, in the **Load Balancing Virtual Servers** page, select the vserver that you want to enable port insertion (for example, VS1_IPv6).
3. Click **Open**.
4. In the **Configure Virtual Server (Load Balancing)** dialog box, click the **Advanced** tab.
5. In the **Vserver IP Port Insertion** drop-down list box, select **VIPADDR**.
6. In the **Vserver IP Port Insertion text box**, type the vip header.

Traffic Domains

Traffic domains are a way to segment network traffic for different applications. You can use traffic domains to create multiple isolated environments within a NetScaler appliance. An application belonging to a specific traffic domain communicates with entities and processes traffic within that domain. The traffic belonging to one traffic domain cannot cross the boundary of another traffic domain.

Benefits of using Traffic Domains

The main benefits of using traffic domains on a NetScaler appliance are the following:

- ♦ **Use of duplicate IP addresses in a Network.** Traffic domains allow you to use duplicate IP address on the network. You can assign the same IP address or network address to multiple devices on a network, or multiple entities on a NetScaler appliance, as long as each of the duplicate address belongs to a different traffic domain.
- ♦ **Use of Duplicate entities on the NetScaler appliance.** Traffic domains also allow you to use duplicate NetScaler feature entities on the appliance. You can create entities with the same settings as long as each entity is assigned to a separate traffic domain.

Note: Duplicate entities with same name is not supported.

- ♦ **Multitenancy.** Using traffic domains, you can provide hosting services for multiple customers by isolating each customer's type of application traffic within a defined address space on the network.

A traffic domain is uniquely identified by an identifier, which is an integer value. Each traffic domain needs a VLAN or a set of VLANs. The isolation functionality of the traffic domain depends on the VLANs bound to the traffic domain. More than one VLAN can be bound to a traffic domain, but the same VLAN cannot be a part of multiple traffic

domains. Therefore, the maximum number of traffic domains that can be created depends on the number of VLANs configured on the appliance.

Default Traffic Domain

A NetScaler appliance has a preconfigured traffic domain, called the *default traffic domain*, which has an ID of 0. All factory settings and configurations are part of the default traffic domain. You can create other traffic domains and then segment traffic between the default traffic domain and each of the other traffic domains. You cannot remove the default traffic domain from the NetScaler appliance. Any feature entity that you create without setting the traffic domain ID is automatically associated with the default traffic domain.

Note: Some features and configurations are supported only in the default traffic domain. They do not work in nondefault traffic domains. For a list of the features supported in all traffic domains, see ["Supported NetScaler Features in Traffic Domains."](#)

How Traffic Domains Work

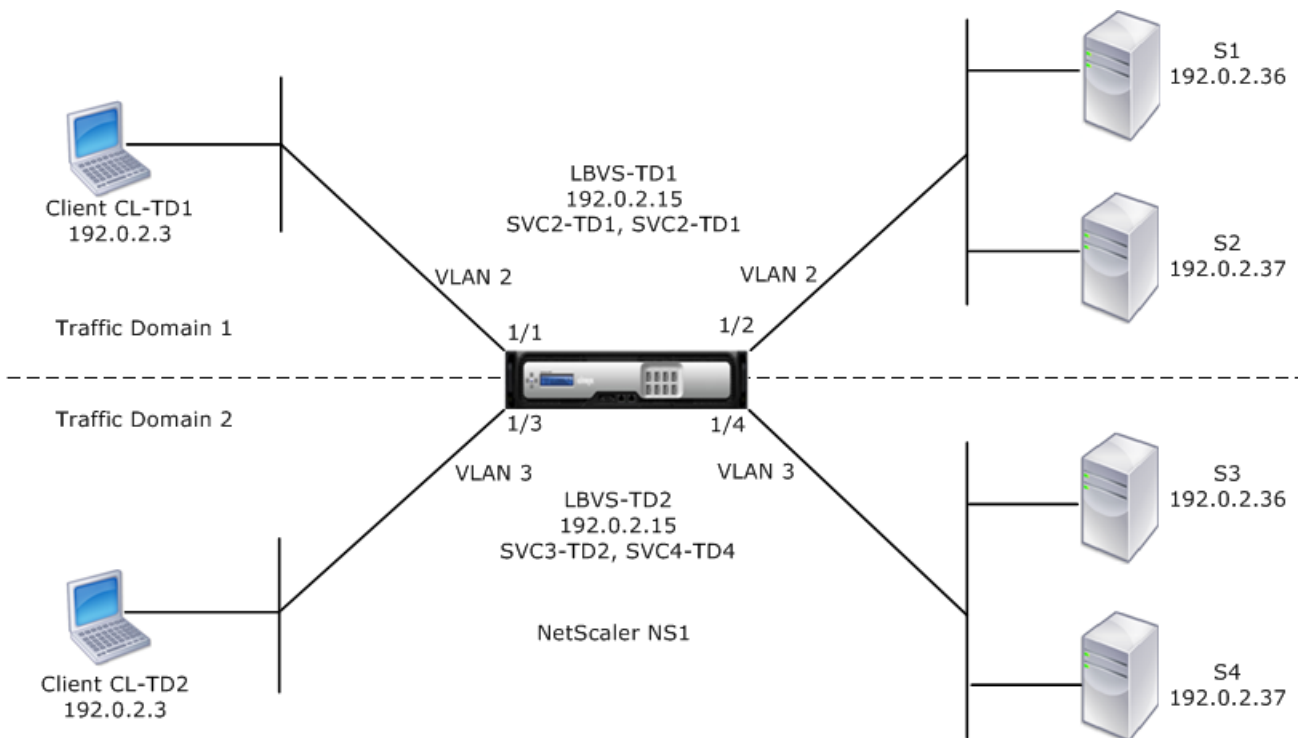
As an illustration of traffic domains, consider an example in which two traffic domains, with IDs 1 and 2, are configured on NetScaler appliance NS1.

In traffic domain 1, load balancing virtual server LBVS-TD1 is configured to load balance traffic across servers S1 and S2. On the NetScaler appliance, servers S1 and S2 are represented by services SVC1-TD1 and SVC2-TD1, respectively. Servers S1 and S2 are connected to NS1 through L2 switch SW2-TD1. Client CL-TD1 is on a private network connected to NS1 through L2 switch SW1-TD1. SW1-TD1 and SW2-TD1 are connected to VLAN 2 of NS1. VLAN 2 is bound to traffic domain 1, which means that client CL-TD1 and servers S1 and S2 are part of traffic domain 1.

Similarly in traffic domain 2, load balancing virtual server LBVS-TD2 is configured to load balance traffic across S3 and S4. On the NetScaler appliance, servers S3 and S4 are represented by services SVC3-TD2 and SVC4-TD2, respectively. Servers S3 and S4 are connected to NS1 through L2 switch SW2-TD2. Client CL-TD2 is on a private network connected to NS1 through L2 switch SW1-TD2. SW1-TD2 and SW2-TD2 are connected to VLAN 3 of NS1. VLAN 3 is bound to traffic domain 2, which means that client CL-TD2 and servers S3 and S4 are part of traffic domain 2.

On the NetScaler appliance, entities LBVS-TD1 and LBVS-TD2 share the same settings, including the IP address. The same is true for SVC1-TD1 and SVC3-TD2, and for SVC2-TD1 and SVC4-TD2. This is possible because these entities are in different traffic domains.

Similarly, servers S1 and S3, S2 and S4 share the same IP address, and clients CL-TD1 and CL-TD2 each have the same IP address.

Figure 7-20. How traffic domains work

The following table lists the settings used in the example.

Entity	Name	Details
Settings in traffic domain 1		
VLANs bound to traffic domain 1	VLAN 2	VLAN Id: 2 Interfaces bound: 1/1, 1/2
Client connected to TD1	CL-TD1 (for reference purposes only)	IP address: 192.0.2.3
Load balancing virtual server in TD1	LBVS-TD1	IP address: 192.0.2.15
Service bound to virtual server LBVS-TD1	SVC1-TD1	IP address: 192.0.2.36
	SVC2-TD1	IP address: 192.0.2.37
SNIP	SNIP-TD1 (for reference purposes only)	IP address: 192.0.2.27
Settings in traffic domain 2		

Entity	Name	Details
VLAN bound to traffic domain 2	VLAN 3	VLAN Id: 3 Interfaces bound: 1/3, 1/4
Client connected to TD2	CL-TD2 (for reference purposes only)	IP address: 192.0.2.3
Load balancing virtual server in TD2	LBVS-TD2	IP address: 192.0.2.15
Service bound to virtual server LBVS-TD2	SVC3-TD2	IP address: 192.0.2.36
	SVC4-TD2	IP address: 192.0.2.37
SNIP in TD2	SNIP-TD2 (for reference purposes only)	IP address: 192.0.2.29

Following is the traffic flow in traffic domain 1:

1. Client CL-TD1 broadcasts an ARP request for the IP address of 192.0.2.15.
2. The ARP request reaches NS1 on interface 1/1, which is bound to VLAN 2. Because VLAN 2 is bound to traffic domain 1, NS1 updates the ARP table of traffic domain 1 for the IP address of client CL-TD1.
3. Because the ARP request is received on traffic domain 1, NS1 looks for an entity configured on traffic domain 1 that has an IP address of 192.0.2.15. NS1 finds that a load balancing virtual server LBVS-TD1 is configured on traffic domain 1 and has the IP address 192.0.2.15.
4. NS1 sends an ARP response with the MAC address of interface 1/1.
5. The ARP reply reaches CL-TD1. CL-TD1 updates its ARP table for the IP address of LBVS-TD1 with the MAC address of interface 1/1 of NS1.
6. Client CL-TD1 sends a request to 192.0.2.15. The request is received by LBVS-TD1 on port 1/1 of NS1.
7. LBVS-TD1's load balancing algorithm selects server S2, and NS1 opens a connection between a SNIP in traffic domain 1 (192.0.2.27) and S2.
8. S2 replies to SNIP 192.0.2.27 on NS1.
9. NS1 sends S2's reply to client CL-TD1.

Following is the traffic flow in traffic domain 2:

1. Client CL-TD2 broadcasts an ARP request for the IP address of 192.0.2.15.
2. The ARP request reaches NS1 on interface 1/3, which is bound to VLAN 3. Because VLAN 3 is bound to traffic domain 2, NS1 updates traffic-domain 2's ARP-table entry for the IP address of client CL-TD2, even though an ARP entry for the same IP address (CL-TD1) is already present in the ARP table of traffic domain 1.

3. Because the ARP request is received in traffic domain 2, NS1 searches traffic domain 2 for an entity that has an IP address of 192.0.2.15. NS1 finds that load balancing virtual server LBVS-TD2 is configured in traffic domain 2 and has the IP address 192.0.2.15. NS1 ignores LBVS-TD1 in traffic domain 1, even though it has the same IP address as LBVS-TD2.
4. NS1 sends an ARP response with the MAC address of interface 1/3.
5. The ARP reply reaches CL-TD2. CL-TD2 updates its ARP table entry for the IP address of LBVS-TD2 with the MAC address of interface 1/3 of NS1.
6. Client CL-TD2 sends a request to 192.0.2.15. The request is received by LBVS-TD2 on interface 1/3 of NS1.
7. LBVS-TD2's load balancing algorithm selects server S3, and NS1 opens a connection between a SNIP in traffic domain 2 (192.0.2.29) and S3.
8. S2 replies to SNIP 192.0.2.29 on NS1.
9. NS1 sends S2's reply to client CL-TD2.

Supported NetScaler Features in Traffic Domains

The NetScaler features in the following list are supported in all traffic domains.

Supported features in traffic domains	
<ul style="list-style-type: none"> ♦ ARP table ♦ ND6 table ♦ Bridge table ♦ All types of IPv4 and IPv6 addresses ♦ IPv4 and IPv6 routes ♦ ACL and ACL6 ♦ PBR & PBR6 ♦ INAT ♦ RNAT ♦ Net profiles ♦ SNMP MIBs ♦ Fragmentation ♦ Monitors ♦ Content Switching ♦ Cache Redirection 	<ul style="list-style-type: none"> ♦ Persistency ♦ Service ♦ Servicegroup ♦ Policies (*) ♦ PING ♦ TRACEROUTE ♦ PMTU ♦ High Availability (connection mirroring is not supported) ♦ Cookie Persistency ♦ MSS ♦ Logging ♦ Priority Queuing ♦ Surge Protection ♦ HTTP DOSP (*) ♦ Load balancing of SSL servers

Supported features in traffic domains	
	<ul style="list-style-type: none"> ♦ Load balancing of authentication servers

Any NetScaler feature not listed above is supported only in the default traffic domain. Traffic domains are not supported in a cluster configuration.

Configuring Traffic Domains

Configuring a traffic domain on the NetScaler appliance consists of the following tasks:

- ♦ **Add VLANs.** Create VLANs and bind specified interfaces to them.
- ♦ **Create a traffic domain entity and bind VLANs to it.** This involves the following two tasks:
 - Create a traffic domain entity uniquely identified by an ID, which is an integer value.
 - Bind the specified VLANs to the traffic domain entity. All the interfaces that are bound to the specified VLANs are associated with the traffic domain. More than one VLAN can be bound to a traffic domain, but a VLAN cannot be a part of multiple traffic domains.
- ♦ **Create feature entities on the traffic domain.** Create the required feature entities in the traffic domain. The CLI commands and configuration dialog boxes of all the supported features in a nondefault traffic domain include a parameter called a *traffic domain identifier* (td). When configuring a feature entity, if you want the entity to be associated with a particular traffic domain, you must specify the td. Any feature entity that you create without setting the td is automatically associated with the default traffic domain.

To give you an idea of how feature entities are associated with a traffic domain, this topic covers the procedures for configuring all the entities mentioned in the figure titled "[How traffic domains work](#)."

The command line interface has two commands for these two tasks, but the configuration utility combines them in a single dialog box.

To create a VLAN and bind interfaces to it by using the command line interface

At the command prompt, type:

- ♦ `add vlan <id>`
- ♦ `bind vlan <id> -ifnum <slot/port>`
- ♦ `show vlan <id>`

To create a traffic domain entity and bind VLANs to it by using the command line interface

At the command prompt, type:

- ♦ **add trafficdomain** <td>
- ♦ **bind trafficdomain** <td> -vlan <id>
- ♦ **show trafficdomain** <td>

To create a service by using the command line interface

At the command prompt, type:

- ♦ **add service** <serviceName> <IP> <serviceType> <port> -td <id>
- ♦ **show service** <name>

To create a load balancing virtual server and bind services to it by using the command line interface

At the command prompt, type:

- ♦ **add lb vserver** <name> <serviceType> <ip> <port> -td <id>
- ♦ **bind lb vserver** <vserverName> <serviceName>
- ♦ **show lb vserver** <name>

To create a VLAN by using the configuration utility

1. Navigate to **System > Network > VLANs**.
2. In the details pane, click **Add**.
3. In the **Add VLAN** dialog box, set the parameters. For a description of a parameter, hover the mouse cursor over the corresponding field.
4. To bind a network interface to a VLAN, under **Interfaces**, select the **Active** check box corresponding to the interface that you want to bind to the VLAN.
5. Click **Create**, and then click **Close**.

To create a traffic domain entity by using the configuration utility

1. Navigate to **System > Network > Traffic Domains**.
2. In the details pane, click **Add**.
3. In the **Create Traffic Domain** dialog box, set the Traffic Domain parameter.
4. On the **VLAN** tab, select an available VLAN and click **Add** to bind it to the traffic domain. Repeat to bind additional VLANs.
5. Click **Create**, and then click **Close**.

To create a service by using the configuration utility

1. Navigate to **Traffic Management > Load Balancing > Services**.
2. In the details pane, click **Add**.

3. In the Create Service dialog box, set the following parameters.

- Service Name*
- Server*
- Protocol*
- Port*
- Traffic Domain ID

* A required parameter

4. Click **Create**.

5. Repeat steps 3-4 to create another service.

6. Click **Close**

To create a load balancing virtual server and bind services to it by using the configuration utility

1. Navigate to **Traffic Management > Load Balancing > Virtual Servers**.

2. In the **Load Balancing Virtual Servers** pane, click **Add**.

3. In the **Create Virtual Servers (Load Balancing)** dialog box, set the following parameters.

- Name*
- IP Address*
- Protocol*
- Port*
- Traffic Domain ID

* A required parameter

4. Under the **Services** tab, in the **Active** column, select the check box for the service that you want to bind to the virtual server.

5. Click **Create**.

6. Repeat steps 3-5 to create another virtual server.

7. Click **Close**.

Chapter 8

Web Interface

Topics:

- [How Web Interface Works](#)
- [Prerequisites](#)
- [Installing the Web Interface](#)
- [Configuring the Web Interface](#)
- [Using the WebInterface.conf Dialog Box](#)
- [Using the config.xml Dialog Box](#)

The Web Interface on Citrix NetScaler appliances is based on Java Server Pages (JSP) technology and provides access to Citrix XenApp and Citrix XenDesktop applications. Users access resources through a standard Web browser or by using the Citrix XenApp plug-in.

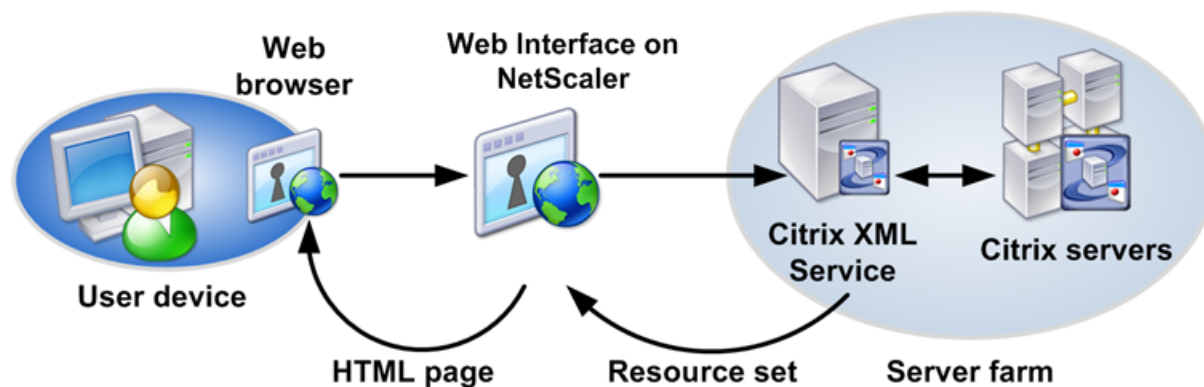
The Web Interface runs as a service on port 8080 on the NetScaler appliance. To create Web Interface sites, Java is executed on Apache Tomcat Web server version 6.0.35 on the NetScaler appliance. The Web Interface sites provide user access to the XenApp and XenDesktop resources, which include applications, content, and desktops.

The Web Interface installation includes installing the Web Interface tar file and JRE tar file on the NetScaler appliance. To configure the Web Interface, you create a Web Interface site and bind one or more XenApp or XenDesktop farms to it.

How Web Interface Works

The following figure illustrates a basic Web interface session.

Figure 8-1. A Basic Web Interface Session



Following is a typical set of interactions among a user device, a NetScaler running the Web interface, and a server farm.

1. A user authenticates to the Web interface through a Web browser or by using the XenApp plug-in.
2. The Web interface reads the user's credentials and forwards the information to the Citrix XML Service running on servers in the server farm.
3. The Citrix XML Service on the designated server retrieves from the servers a list of resources that the user can access. These resources constitute the user's resource set and are retrieved from the Independent Management Architecture (IMA) system.
4. The Citrix XML Service then returns the user's resource set to the Web interface running on the NetScaler.
5. The user clicks an icon that represents a resource on the HTML page.
6. The Web interface queries the Citrix XML Service for the least busy server.
7. The Citrix XML Service returns the address of this server to the Web interface.
8. The Web interface sends the connection information to the Web browser.
9. The Web browser initiates a session with the server.

Prerequisites

The following prerequisites are required before you begin installing and configuring the Web interface.

- ♦ XenApp or XenDesktop farms are set up and running in your environment.

- ♦ Conceptual knowledge of the Web interface.

Installing the Web Interface

To install the Web interface, you need to install the following files:

- ♦ **Web interface tar file.** The setup file for installing the Web interface on the NetScaler appliance. This tar file also includes Apache Tomcat Web server version 6.0.35. The file name has the following format: `nswi-<version number>.tgz` (for example, `nswi-1.5.tgz`).
- ♦ **JRE tar file.** The JRE tarball. You can use the Diablo Latte JRE version 1.6.0-7 for 64-bit FreeBSD 6.x/amd64 platform available on FreeBSD Foundation Web site at <http://www.freebsdoundation.org/java/java16>. Alternatively, you can use OpenJDK6 package for FreeBSD 6.x/amd63. You can download `openjdk6-b17_2.tbz` from <http://ftp.riken.jp/pub/FreeBSD/ports/amd64/packages-6-stable/java/>.

Note: On a high availability setup, make sure that both the Web interface tar file and the JRE tar file are installed on both the primary and the secondary appliances.

Copy the tar files to a local workstation or to the `/var` directory of the appliance.

These files install all the Web interface components and JRE on the hard drive and configure automatic startup of the Tomcat Web server with Web interface at appliance startup time. Both tar files are internally expanded in the `/var/wi` directory on the hard drive.

Note: After installing web interface on the appliance and before creating a web interface site, you must place the client plugin in the appliance by using the appropriate **Upload Plugins** utility provided on the web interface details pane.

To install the Web interface and JRE tar files by using the command line interface

At the command prompt, type:

`install wi package -wi <URL> -jre <URL>`

Examples

```
> install wi package -wi sftp://
username:password@10.102.29.12/var/nswi-1.5.tgz -
jre ftp://username:password@10.102.29.14/tmp/
diablojre- freebsd6.amd64.1.6.0.07.02.tbz

> install wi package -wi ftp://
username:password@10.102.29.15/var/nswi-1.5.tgz -
```

```
jre file:///var/diablo-  
jrefreebsd6.amd64.1.6.0.07.02.tbz
```

To install the Web interface and JRE tar files by using the configuration utility

1. Navigate to **System > Web Interface**.
2. In the details pane, under **Getting Started**, click **Install Web Interface**.
3. In the **Install Web Interface** dialog box, in the **Web Interface tar file path** text box, type the complete path to the Web interface tar file. You can also use the **browse** button to locate the file on your local system or the NetScaler hard drive.
4. In the **JRE tar file path** text box, type the complete path to the JRE tar file. You can also use the **browse** button to locate the file on your local system or the NetScaler hard drive.
5. Click **Install**.

Configuring the Web Interface

To configure the web interface, you create a web interface site and bind one or more XenApp or XenDesktop farms to it. You then configure the web interface to work behind an HTTP or an HTTPS virtual server.

- ♦ **Using an HTTP or an HTTPS virtual server.** You create an HTTP or an HTTPS virtual server on the NetScaler appliance and bind the web interface service, running on port 8080 of the NetScaler appliance, to the virtual server. Clients on the LAN use the virtual server IP address to access the web interface. When using this access method, the URL format for the web interface site is as follows:

```
<HTTP or HTTPS>://<HTTP or HTTPS vserver IP address>:<vserver  
port number>/<web interface site path>
```

The following access methods are available for clients accessing the web interface site when it is configured using an HTTP or an HTTPS virtual server:

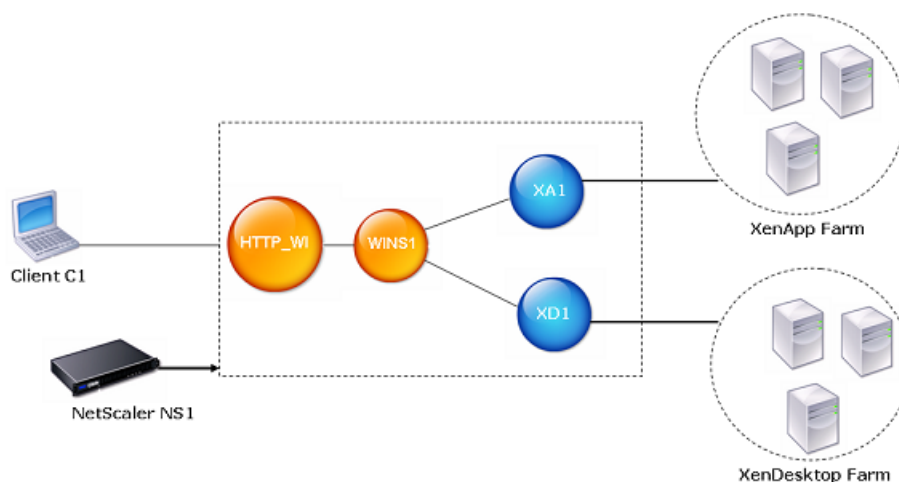
- **Direct.** Actual address of a XenApp or XenDesktop server is sent to the clients.
- **Alternate.** Alternate address of a XenApp or XenDesktop server is sent to the clients.
- **Translated.** Translated address, from the defined internal addresses to external addresses and ports mapping table, is sent to the clients from a specified network. When you use this option, you have to define internal address to external address and port mappings.

Configuring a Web Interface Site for LAN Users Using HTTP

In this scenario, user and the Web interface setup are on the same enterprise LAN. The enterprise has both a XenApp and a XenDesktop farm. Users access the Web interface by using an HTTP vserver. The Web interface exposes its own login page for authentication. The vserver IP address is used to access the Web interface.

The following figure illustrates the Web interface running on the NetScaler appliance NS1. A Web interface site WINS1 is created and a XenApp farm XA1 and a XenDesktop farm XD1 are bound to it. An HTTP vserver HTTP_WI is also created. Client C1 uses the IP address of the HTTP_WI vserver to access the WINS1 site.

Figure 8-2. A Web Interface Site Configured for LAN Users Using HTTP



To configure a Web interface site for LAN users using HTTP by using the configuration utility

1. Navigate to **System > Web Interface**.
2. In the details pane, click **Web Interface Wizard**.
3. On the wizard **Introduction** page, click **Next**.
4. On the wizard **Configure Web Interface Site** page, configure the following parameters:
 - **Site Path*** (You cannot change the name of an existing Web interface site.)

- **Site Type**
- **Published Resource Type**
- **Kiosk Mode**
- **Authentication Methods**
- **Login Title**
- **Web Session Timeout**
- **Enable access through receiver client**

* A required parameter.

5. In **Default Access Methods**, select the **Direct** or **Alternate** or **Translated** option and configure the following parameters:

- **Virtual Server**
- **Protocol** (select **HTTP**)
- **IP Address**
- **Port**

Note:

When you create the HTTP vserver by using the configuration utility, the configuration utility automatically creates a service, which logically represents the Web interface service running on the NetScaler appliance, and binds the service to the HTTP virtual server.

6. Click **Next**.
7. On the wizard's **Configure Access Methods** page, do one of the following:
 - To set an access method for a client IP address or network, click **Add**.
 - To change the access method for a client IP address or network, select the association, and then click **Open**.
8. In the **Configure Access Method** dialog box, configure the following parameters:
 - **Client IP Address*** (You cannot change this parameter after setting it.)
 - **Netmask** (You cannot change this parameter after setting it.)
 - **Access Method**

* A required parameter.

Note: Before you configure access method based on the client IP address, you must enable USIP mode on the web interface service to make the client's IP address available with the web interface.

9. Click **Next**.

10. On the wizard's **Configure Address Translations** page, click **Add** for adding a mapping between an Internal IP address and an external IP address.

Note: The **Configure Address Translations** page appears on the wizard when you set the Translated access method for a Client's IP address or network.

11. In the **Configure Address Translations** dialog box, configure the following parameters:
 - **Internal IP Address**
 - **Internal Port**
 - **External IP Address**
 - **External Port**
 - **Access Type**

* A required parameter.
12. Click **Next**.
13. On the wizard's **Configure XenApp/XenDesktop Farm** page, do one of the following:
 - To add a XenApp or XenDesktop farm, click **Add**.
 - To modify an existing XenApp or XenDesktop farm, select the farm, and then click **Open**.
14. In the **Create XenApp/XenDesktop Farm** or **Configure XenApp/XenDesktop Farm** dialog box, configure the following parameters:
 - **Name*** (You cannot change the name of an existing XenApp or XenDesktop farm.)
 - **XML Service Addresses***
 - **XML Service Port**
 - **Transport**
 - **Load Balance**

* A required parameter.
15. Click **Next**, and then click **Finish**.
16. Verify that the Web interface site you configured is correct by selecting the site and viewing the **Details** section at the bottom of the pane. To view the Web interface site, in the navigation pane, expand **System**, expand **Web Interface**, and then click **Sites**.

To configure a Web interface site for LAN users using HTTP by using the command line interface

1. Add a Web interface site. Set **Direct** or **Alternate** or **Translated** for the defaultAccessMethod parameter. At the command prompt, type:

```
add wi site <sitePath> -siteType ( XenAppWeb | XenAppServices ) -
publishedResourceType ( Online | Offline | DualMode ) -kioskMode ( ON | OFF ) -
wiAuthenticationMethods ( Explicit | Anonymous ) -webSessionTimeout
<positive_integer> -defaultAccessMethod <defaultAccessMethod> -loginTitle
<string>
```

Example

```
> add wi site WINS1 -siteType XenAppWeb -
publishedResourceType Online -kioskMode ON -
defaultAccessMethod Direct
```

2. (Optional) Set an access method for a Client's IP address or network. At the command prompt, type:
bind wi site <sitePath> -accessMethod <accessMethod> -clientIpAddress <ip_addr> -clientNetMask <netmask>
3. If you have set the **Translated** access method for a Client's IP address or network then provide Internal IP and external IP address mappings. At the command prompt, type:
bind wi site <sitePath> -translationInternalIp <ip_addr> -translationInternalPort <port|*> -translationExternalIp <ip_addr> -translationExternalPort <port|*> [-accessType <accessType>]
4. Bind XenApp or XenDesktop farms to the Web interface site. At the command prompt, type:
bind wi site <sitePath> <farmName> <xmlServerAddresses> -xmlPort <value> -transport (HTTP | HTTPS) -loadBalance (ON | OFF)

Example

```
> bind wi site WINS1 XA1 10.102.46.6 -xmlPort 80 -transport
HTTP -LoadBalance OFF
> bind wi site WINS1 XD1 10.102.46.50 -xmlPort 80 -
transport HTTP -LoadBalance OFF
```

5. Create a service that is a logical representation of the Web interface service running on the NetScaler appliance. At the command prompt, type:
add service <name> <IP address> <serviceType> <port>

Example

```
> add service WI_Loopback_Service 127.0.0.1 HTTP 8080
```

6. Add an HTTP vserver. At the command prompt, type:
add lb vserver <virtualServerName> <protocol> <IPAddress> <port>

Example

```
> add lb vserver HTTP_WI HTTP 10.102.29.5 80
```

7. Bind the Web interface service to the HTTP vserver. At the command prompt, type:

```
bind lb vserver <virtualServerName> <serviceName>
```

Example

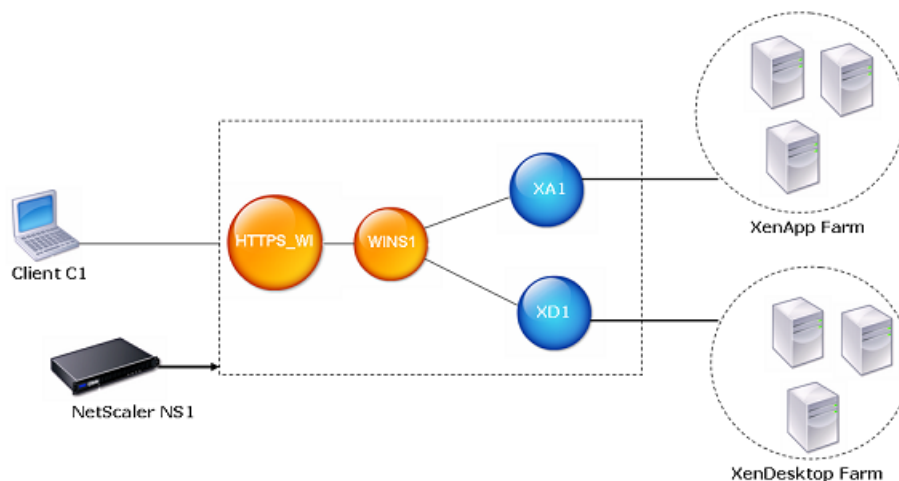
```
> bind lb vserver HTTP_WI WI_Loopback_Service
```

Configuring a Web Interface Site for LAN Users Using HTTPS

In this scenario, user accounts and the Web interface setup are on the same enterprise LAN. Users access the Web interface by using an SSL-based (HTTPS) vserver. The Web interface exposes its own login page for authentication. SSL offloading is done by this vserver on the NetScaler. The vserver IP address is used to access the Web interface instead of the NetScaler IP address (NSIP).

The following figure illustrates the Web interface running on the NetScaler appliance NS1. A Web interface site WINS1 is created and a XenApp farm XA1 and a XenDesktop farm XD1 are bound to it. An HTTPS vserver HTTPS_WI is also created. Client C1 uses the IP address of the HTTPS_WI vserver to access the WINS1 site.

Figure 8-3. A Web Interface Site Configured for LAN Users Using HTTPS



To configure a Web interface site for LAN users using HTTPS by using the configuration utility

1. Navigate to **System > Web Interface**.

2. In the details pane, click **Web Interface Wizard**.
3. On the wizard **Introduction** page, click **Next**.
4. On the wizard **Configure Web Interface Site** page, configure the following parameters:
 - **Site Path*** (You cannot change the name of an existing Web interface site.)
 - **Site Type**
 - **Published Resource Type**
 - **Kiosk Mode**
 - **Authentication Methods**
 - **Login Title**
 - **Web Session Timeout**
 - **Enable access through receiver client**

* A required parameter.

5. In **Default Access Methods**, select the **Direct** or **Alternate** or **Translated** option and configure the following parameters:
 - **Virtual Server**
 - **Protocol** (select HTTPs)
 - **IP Address**
 - **Port**

Note:

When you create the HTTPS vserver by using the configuration utility, the configuration utility automatically creates a service, which logically represents the Web interface service running on the NetScaler appliance, and binds the service to the HTTPS virtual server.

6. Click **Next**.
7. On the wizard's **Specify a server Certificate** page, you create or specify an existing SSL certificate-key pair. The SSL certificate-key pair is automatically bound to the HTTPS vserver.
8. Click **Next**.
9. On the wizard's **Configure Access Methods** page, do one of the following:
 - To set an access method for a client IP address or network, click **Add**.
 - To change the access method for a client IP address or network, select the association, and then click **Open**.
10. In the **Configure Access Method** dialog box, configure the following parameters:

- **Client IP Address*** (You cannot change this parameter after setting it.)
- **Netmask** (You cannot change this parameter after setting it.)
- **Access Method**

* A required parameter.

Note: Before you configure access method based on the client IP address, you must enable USIP mode on the web interface service to make the client's IP address available with the web interface.

11. Click **Next**.
12. On the wizard's **Configure Address Translations** page, click **Add** for adding a mapping between an Internal IP address and an external IP address.

Note: The **Configure Address Translations** page appears on the wizard when you set the Translated access method for a Client's IP address or network.

13. In the **Configure Address Translations** dialog box, configure the following parameters:
 - **Internal IP Address**
 - **Internal Port**
 - **External IP Address**
 - **External Port**
 - **Access Type**

* A required parameter.
14. On the wizard's **Configure XenApp/XenDesktop Farm** page, do one of the following:
 - To add a XenApp or XenDesktop farm, click **Add**.
 - To modify an existing XenApp or XenDesktop farm, select the farm, and then click **Open**.
15. In the **Create XenApp/XenDesktop Farm** or **Configure XenApp/XenDesktop Farm** dialog box, configure the following parameters:
 - **Name*** (You cannot change the name of an existing XenApp or XenDesktop farm.)
 - **XML Service Addresses***
 - **XML Service Port**
 - **Transport**
 - **Load Balance**

* A required parameter.

16. Click **Next**, and then click **Finish**.
17. Verify that the Web interface site you configured is correct by selecting the site and viewing the **Details** section at the bottom of the pane. To view the Web interface site, in the navigation pane, expand **System**, expand **Web Interface**, and then click **Sites**.

To configure a Web interface site for LAN users using HTTPS by using the command line

1. Add a Web interface site. Set **Direct** or **Alternate** or **Translated** for the defaultAccessMethod parameter. At the command prompt, type:

```
add wi site <sitePath> -siteType ( XenAppWeb | XenAppServices ) -
publishedResourceType ( Online | Offline | DualMode ) -kioskMode ( ON | OFF ) -
wiAuthenticationMethods ( Explicit | Anonymous ) -webSessionTimeout
<positive_integer> -defaultAccessMethod <defaultAccessMethod> -loginTitle
<string>
```

Example

```
> add wi site WINS1 -siteType XenAppWeb -
publishedResourceType Online -kioskMode ON -
defaultAccessMethod Direct
```

2. (Optional) Set an access method for a Client's IP address or network. At the command prompt, type:


```
bind wi site <sitePath> -accessMethod <accessMethod> -clientIpAddress <ip_addr>
-clientNetMask <netmask>
```
3. If you have set the **Translated** access method for a Client's IP address or network then provide Internal IP and external IP address mappings. At the command prompt, type:


```
bind wi site <sitePath> -translationInternalIp <ip_addr> -translationInternalPort
<port|*> -translationExternalIp <ip_addr> -translationExternalPort <port|*> [-
accessType <accessType>]
```
4. Bind XenApp or XenDesktop farms to the Web interface site. At the command prompt, type:


```
bind wi site <sitePath> <farmName> <xmlServerAddresses> -xmlPort <value> -
transport ( HTTP | HTTPS ) -loadBalance ( ON | OFF )
```

Example

```
> bind wi site WINS1 XA1 10.102.46.6 -xmlPort 80 -transport
HTTP -LoadBalance OFF
> bind wi site WINS1 XD1 10.102.46.50 -xmlPort 80 -
transport HTTP -LoadBalance OFF
```

5. Create a service that is a logical representation of the Web interface service running on the NetScaler appliance. At the command prompt, type:


```
add service <name> <IPAddress> <serviceType> <port>
```

Example

```
> add service WI_Loopback_Service 127.0.0.1 HTTP 8080
```

6. Add an HTTPS vserver. At the command prompt, type:

```
add lb vserver <virtualServerName> <protocol> <IPAddress> <port>
```

Example

```
> add lb vserver HTTPS_WI SSL 10.102.29.3 443
```

7. Bind the Web interface service to the HTTPS vserver. At the command prompt, type:

```
bind lb vserver <virtualServerName> <serviceName>
```

Example

```
> bind lb vserver HTTPS_WI WI_Loopback_Service
```

8. Create an SSL certificate key pair. At the command prompt, type:

```
add ssl certkey <certificate-KeyPairName> -cert <certificateFileName> -key  
<privateKeyFileName>
```

Example

```
> add ssl certkey SSL-Certkey-1 -cert /nsconfig/ssl/  
test1.cer -key /nsconfig/ssl/test1
```

9. Bind the SSL certificate key pair to the HTTPS vserver. At the command prompt, type:

```
bind ssl vserver <vserverName> -certkeyName <certificate- KeyPairName>
```

Example

```
> bind ssl vserver HTTPS_WI -certkeyName SSL-Certkey-1
```

10. Add a rewrite action. At the command prompt, type:

```
add rewrite action <name> <type> <target> [<stringBuilderExpr>] [(-pattern  
<expression>]
```

Example

```
> add rewrite action Replace_HTTP_to_HTTPS INSERT AFTER  
"HTTP.RES.HEADER(\"Location\").Value(0).Prefix(4)\" \"s\""
```

11. Create a rewrite policy and bind the rewrite action to it. At the command prompt, type:

```
add rewrite policy <name> <expression> <rewriteAction>
```

Example

```
> add rewrite policy rewrite_location "HTTP.RES.STATUS ==
302 && HTTP.RES.HEADER(\"Location
\").Value(0).startswith(\"http:\")" Replace_HTTP_to_HTTPS
```

12. Bind the rewrite policy to the HTTPS vserver. At the command prompt, type:
bind lb vserver <VserverName> -policyname <rewritePolicyName> -priority <value>
 -type response

Example

```
> bind lb vserver HTTPS_WI -policyname rewrite_location -
priority 10 -type response
```

Using the WebInterface.conf Dialog Box

The WebInterface.conf dialog box in the configuration utility displays the content of the webinterface.conf file for a Web Interface site.

You can do the following from this dialog box:

- ♦ Edit the WebInterface.conf file and save the changes.
- ♦ Search the file's content for instances of a text string.
- ♦ Easily save the WebInterface.conf file to your local computer.

To search a string in the webinterface.conf file by using the configuration utility

1. Navigate to **System > Web Interface > Sites**.
2. In the details pane, select the web interface site, and then click **WebInterface.conf**.
3. In the **WebInterface.conf** dialog box, use the following controls:
 - **Find**. Displays the following search options that you can use to find one or more instances of a text string in a configuration:
 - ♦ **Look for**. Provides a space for you to type the text string that you want to locate in the configuration. As you type the text, the first instance is displayed. If the word you are looking for is not in the file, the **Look for** text box will change color.
 - ♦ **Next**. Finds and highlights the next occurrence of the text string you typed in **Look for**.
 - ♦ **Previous**. Finds and highlights the previous occurrence of the text string you typed in **Look for**.

- ♦ **Mark All.** Highlights all instances of the text string at one time you typed in **Look for**. Scroll to review each highlighted instance.

4. Click **Close**.

To save the content of the `webinterface.conf` to your local system by using the configuration utility

1. Navigate to **System > Web Interface > Sites**.
2. In the details pane, select the web interface site, and then click **WebInterface.conf**.
3. In the **WebInterface.conf** dialog box, click **Save output text to a file**.
4. Click **Close**.

Using the `config.xml` Dialog Box

The `Config.xml` dialog box in the configuration utility displays the content of the `config.xml` file for a Web Interface site of site type XenApp/XenDesktop Services Site.

You can do the following from this dialog box:

- ♦ Edit the `config.xml` file and save the changes.
- ♦ Search the file's content for instances of a text string.
- ♦ Easily save the `config.xml` file to your local computer.

To search a string in the `config.xml` file by using the configuration utility

1. Navigate to **System > Web Interface > Sites**.
2. In the details pane, select the web interface site of site type XenApp/XenDesktop Services Site, and then click **Config.xml**.
3. In the **Config.xml** dialog box, use the following controls:
 - **Find.** Displays the following search options that you can use to find one or more instances of a text string in a configuration:
 - ♦ **Look for.** Provides a space for you to type the text string that you want to locate in the configuration. As you type the text, the first instance is displayed. If the word you are looking for is not in the file, the **Look for** text box will change color.
 - ♦ **Next.** Finds and highlights the next occurrence of the text string you typed in **Look for**.

- ♦ **Previous.** Finds and highlights the previous occurrence of the text string you typed in **Look for**.
- ♦ **Mark All.** Highlights all instances of the text string at one time you typed in **Look for**. Scroll to review each highlighted instance.

4. Click **Close**.

To save the content of the config.xml to the local system by using the configuration utility

1. Navigate to **System > Web Interface > Sites**.
2. In the details pane, select the web interface site of site type XenApp/XenDesktop Services Site, and then click **Config.xml**.
3. In the **Config.xml** dialog box, click **Save output text to a file**.
4. Click **Close**.