

Getting Started with Citrix NetScaler

Oct 13, 2015

Getting Started with Citrix NetScaler

Intended for system and network administrators who install and configure complex networking equipment, this section of the library describes initial set-up and basic configuration of the NetScaler, including the following topics.

- Understanding the NetScaler
- Processing Order of Features
- Where Does a NetScaler Appliance Fit in the Network?
- How a NetScaler Communicates with Clients and Servers
- Introduction to the Citrix NetScaler Product Line
- Installing the NetScaler Hardware
- Accessing a Citrix NetScaler
- Configuring a NetScaler for the First Time
- Configuring a High Availability Pair for the First Time
- Configuring a FIPS Appliance for the First Time
- Understanding Common Network Topologies
- Configuring System Management Settings
- Load Balancing Traffic on a NetScaler Appliance
- Accelerating Load Balanced Traffic by Using Compression
- Securing Load Balanced Traffic by Using SSL
- Features at a Glance

Understanding the NetScaler

The Citrix NetScaler product is an application switch that performs application-specific traffic analysis to intelligently distribute, optimize, and secure Layer 4-Layer 7 (L4â€”L7) network traffic for web applications. For example, a NetScaler bases load balancing decisions on individual HTTP requests instead of on long-lived TCP connections, so that the failure or slowdown of a server is managed much more quickly and with less disruption to clients. The NetScaler feature set can be broadly categorized as consisting of switching features, security and protection features, and server-farm optimization features.

Switching Features

When deployed in front of application servers, a NetScaler ensures optimal distribution of traffic by the way in which it directs client requests. Administrators can segment application traffic according to information in the body of an HTTP or TCP request, and on the basis of L4â€”L7 header information such as URL, application data type, or cookie. Numerous load balancing algorithms and extensive server health checks improve application availability by ensuring that client requests are directed to the appropriate servers.

Security and Protection Features

NetScaler security and protection features protect web applications from Application Layer attacks. A NetScaler allows legitimate client requests and can block malicious requests. It provides built-in defenses against denial-of-service (DoS) attacks and supports features that protect against legitimate surges in application traffic that would otherwise overwhelm the servers. An available built-in firewall protects web applications from Application Layer attacks, including buffer overflow exploits, SQL injection attempts, cross-site scripting attacks, and more. In addition, the firewall provides identity theft protection by securing confidential corporate information and sensitive customer data.

Optimization Features

Optimization features offload resource-intensive operations, such as Secure Sockets Layer (SSL) processing, data compression, client keep-alive, TCP buffering, and the caching of static and dynamic content from servers. This improves the performance of the servers in the server farm and therefore speeds up applications. A NetScaler supports several transparent TCP optimizations, which mitigate problems caused by high latency and congested network links, accelerating the delivery of applications while requiring no configuration changes to clients or servers.

Understanding Policies and Expressions

A policy defines specific details of traffic filtering and management on a NetScaler. It consists of two parts: the expression and the action. The expression defines the types of requests that the policy matches. The action tells the NetScaler what to do when a request matches the expression. As an example, the expression might be to match a specific URL pattern to a type of security attack, with the action being to drop or reset the connection. Each policy has a priority, and the priorities determine the order in which the policies are evaluated.

When a NetScaler receives traffic, the appropriate policy list determines how to process the traffic. Each policy on the list contains one or more expressions, which together define the criteria that a connection must meet to match the policy.

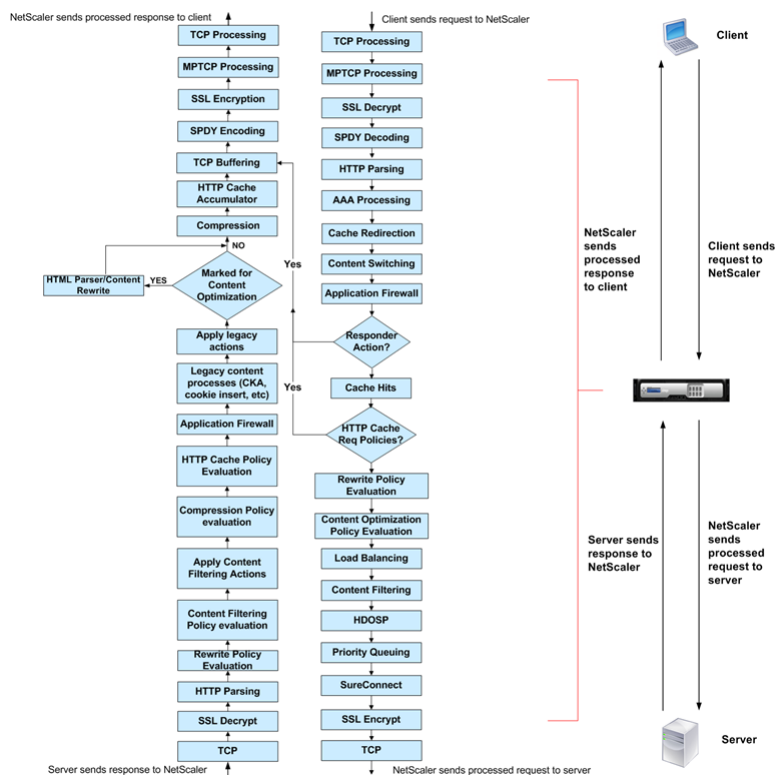
For all policy types except Rewrite policies, a NetScaler implements only the first policy that a request matches, not any additional policies that it might also match. For Rewrite policies, the NetScaler evaluates the policies in order and, in the case of multiple matches, performs the associated actions in that order. Policy priority is important for getting the results you want.

Processing Order of Features

Depending on requirements, you can choose to configure multiple features. For example, you might choose to configure both compression and SSL offload. As a result, an outgoing packet might be compressed and then encrypted before being sent to the client.

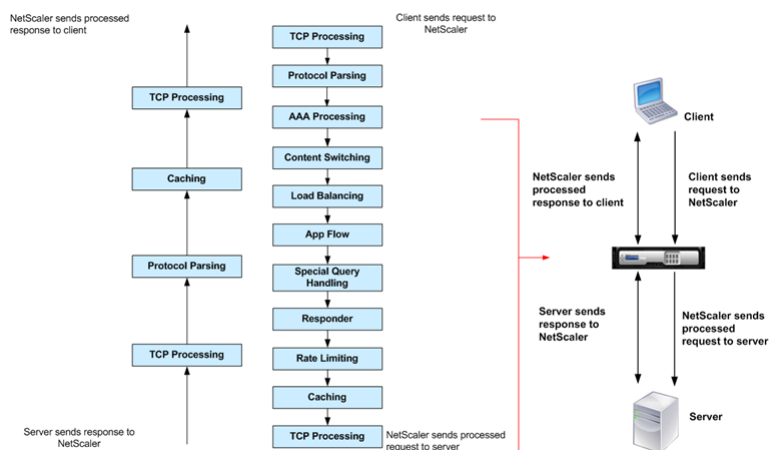
The following figure shows the L7 packet flow in the NetScaler.

Figure 1. L7 Packet Flow Diagram



The following figure shows the DataStream packet flow in the NetScaler. DataStream is supported for MySQL and MS SQL databases. For information about the DataStream feature, see "[DataStream](#)."

Figure 2. DataStream Packet Flow Diagram



Where Does a NetScaler Appliance Fit in the Network?


A NetScaler appliance resides between the clients and the servers, so that client requests and server responses pass through it. In a typical installation, virtual servers configured on the appliance provide connection points that clients use to access the applications behind the appliance. In this case, the appliance owns public IP addresses that are associated with its virtual servers, while the real servers are isolated in a private network. It is also possible to operate the appliance in a transparent mode as an L2 bridge or L3 router, or even to combine aspects of these and other modes.

Physical Deployment Modes

Updated: 2013-09-04

A NetScaler appliance logically residing between clients and servers can be deployed in either of two physical modes: inline and one-arm. In inline mode, multiple network interfaces are connected to different Ethernet segments, and the appliance is placed between the clients and the servers. The appliance has a separate network interface to each client network and a separate network interface to each server network. The appliance and the servers can exist on different subnets in this configuration. It is possible for the servers to be in a public network and the clients to directly access the servers through the appliance, with the appliance transparently applying the L4-L7 features. Usually, virtual servers (described later) are configured to provide an abstraction of the real servers. The following figure shows a typical inline deployment.

Figure 1. Inline Deployment



In one-arm mode, only one network interface of the appliance is connected to an Ethernet segment. The appliance in this case does not isolate the client and server sides of the network, but provides access to applications through configured virtual servers. One-arm mode can simplify network changes needed for NetScaler installation in some environments.

For examples of inline (two-arm) and one-arm deployment, see "[Understanding Common Network Topologies](#)."

Citrix NetScaler as an L2 Device

Updated: 2013-09-04

A NetScaler functioning as an L2 device is said to operate in L2 mode. In L2 mode, the NetScaler forwards packets between network interfaces when all of the following conditions are met:

- The packets are destined to another device's media access control (MAC) address.
- The destination MAC address is on a different network interface.
- The network interface is a member of the same virtual LAN (VLAN).

By default, all network interfaces are members of a pre-defined VLAN, VLAN 1. Address Resolution Protocol (ARP) requests and responses are forwarded to all network interfaces that are members of the same VLAN. To avoid bridging loops, L2 mode must be disabled if another L2 device is working in parallel with the NetScaler.

For information about how the L2 and L3 modes interact, see ["Configuring Modes of Packet Forwarding."](#)

For information about configuring L2 mode, see ["Enabling and Disabling Layer 2 Mode."](#)

Citrix NetScaler as a Packet Forwarding Device

Updated: 2014-03-14

A NetScaler appliance can function as a packet forwarding device, and this mode of operation is called L3 mode. With L3 mode enabled, the appliance forwards any received unicast packets that are destined for an IP address that does not belong to the appliance, if there is a route to the destination. The appliance can also route packets between VLANs.

In both modes of operation, L2 and L3, the appliance generally drops packets that are in:

- Multicast frames
- Unknown protocol frames destined for an appliance's MAC address (non-IP and non-ARP)
- Spanning Tree protocol (unless BridgeBPDUs is ON)

For information about how the L2 and L3 modes interact, see ["Configuring Modes of Packet Forwarding."](#)

For information about configuring the L3 mode, see ["Enabling and Disabling Layer 3 Mode."](#)

How a NetScaler Communicates with Clients and Servers

A NetScaler appliance is usually deployed in front of a server farm and functions as a transparent TCP proxy between clients and servers, without requiring any client-side configuration. This basic mode of operation is called Request Switching technology and is the core of NetScaler functionality. Request Switching enables an appliance to multiplex and offload the TCP connections, maintain persistent connections, and manage traffic at the request (application layer) level. This is possible because the appliance can separate the HTTP request from the TCP connection on which the request is delivered.

Depending on the configuration, an appliance might process the traffic before forwarding the request to a server. For example, if the client attempts to access a secure application on the server, the appliance might perform the necessary SSL processing before sending traffic to the server.

To facilitate efficient and secure access to server resources, an appliance uses a set of IP addresses collectively known as NetScaler-owned IP addresses. To manage your network traffic, you assign NetScaler-owned IP addresses to virtual entities that become the building blocks of your configuration. For example, to configure load balancing, you create virtual servers to receive client requests and distribute them to services, which are entities representing the applications on your servers.

Understanding NetScaler-Owned IP Addresses

Updated: 2014-03-12

To function as a proxy, a NetScaler appliance uses a variety of IP addresses. The key NetScaler-owned IP addresses are:

NetScaler IP (NSIP) address

The NSIP address is the IP address for management and general system access to the appliance itself, and for communication between appliances in a high availability configuration.

Virtual server IP (VIP) address

A VIP address is the IP address associated with a virtual server. It is the public IP address to which clients connect.

An appliance managing a wide range of traffic may have many VIPs configured.

Subnet IP (SNIP) address

A SNIP address is used in connection management and server monitoring. You can specify multiple SNIP addresses for each subnet. SNIP addresses can be bound to a VLAN.

IP Set

An IP set is a set of IP addresses, which are configured on the appliance as SNIP. An IP set is identified with a meaningful name that helps in identifying the usage of the IP addresses contained in it.

Net Profile

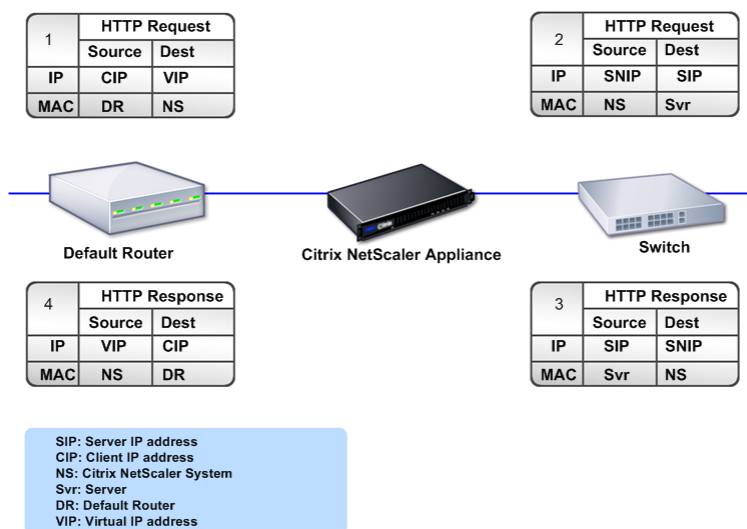
A net profile (or network profile) contains an IP address or an IP set. A net profile can be bound to load balancing or content switching virtual servers, services, service groups, or monitors. During communication with physical servers or peers, the appliance uses the addresses specified in the profile as source IP addresses.

How Traffic Flows Are Managed

Updated: 2014-03-12

Because a NetScaler appliance functions as a TCP proxy, it translates IP addresses before sending packets to a server. When you configure a virtual server, clients connect to a VIP address on the NetScaler instead of directly connecting to a server. As determined by the settings on the virtual server, the appliance selects an appropriate server and sends the client's request to that server. By default, the appliance uses a SNIP address to establish connections with the server, as shown in the following figure.

Figure 1. Virtual Server Based Connections



In the absence of a virtual server, when an appliance receives a request, it transparently forwards the request to the server. This is called the transparent mode of operation. When operating in transparent mode, an appliance translates the source IP addresses of incoming client requests to the SNIP address but does not change the destination IP address. For this mode to work, L2 or L3 mode has to be configured appropriately.

For cases in which the servers need the actual client IP address, the appliance can be configured to modify the HTTP header by inserting the client IP address as an additional field, or configured to use the client IP address instead of a SNIP address for connections to the servers.


Traffic Management Building Blocks

Updated: 2013-06-24

The configuration of a NetScaler appliance is typically built up with a series of virtual entities that serve as building blocks for traffic management. The building block approach helps separate traffic flows. Virtual entities are abstractions, typically representing IP addresses, ports, and protocol handlers for processing traffic. Clients access applications and resources through these virtual entities. The most commonly used entities are virtual servers and services. Virtual servers represent groups of servers in a server farm or remote network, and services represent specific applications on each server.

Most features and traffic settings are enabled through virtual entities. For example, you can configure an appliance to compress all server responses to a client that is connected to the server farm through a particular virtual server. To configure the appliance for a particular environment, you need to identify the appropriate features and then choose the right mix of virtual entities to deliver them. Most features are delivered through a cascade of virtual entities that are bound to each other. In this case, the virtual entities are like blocks being assembled into the final structure of a delivered application. You can add, remove, modify, bind, enable, and disable the virtual entities to configure the features. The following figure shows the concepts covered in this section.

Figure 2. How Traffic Management Building Blocks Work




A Simple Load Balancing Configuration

Updated: 2013-08-30

In the example shown in the following figure, the NetScaler appliance is configured to function as a load balancer. For this configuration, you need to configure virtual entities specific to load balancing and bind them in a specific order. As a load balancer, an appliance distributes client requests across several servers and thus optimizes the utilization of resources.

The basic building blocks of a typical load balancing configuration are services and load balancing virtual servers. The services represent the applications on the servers. The virtual servers abstract the servers by providing a single IP address to which the clients connect. To ensure that client requests are sent to a server, you need to bind each service to a virtual server. That is, you must create services for every server and bind the services to a virtual server. Clients use the VIP address to connect to a NetScaler appliance. When the appliance receives client requests sent to the VIP address, it sends them to a server determined by the load balancing algorithm. Load balancing uses a virtual entity called a monitor to track whether a specific configured service (server plus application) is available to receive requests.

Figure 3. Load Balancing Virtual Server, Services, and Monitors



In addition to configuring the load balancing algorithm, you can configure several parameters that affect the behavior and performance of the load balancing configuration. For example, you can configure the virtual server to maintain persistence based on source IP address. The appliance then directs all requests from any specific IP address to the same server.

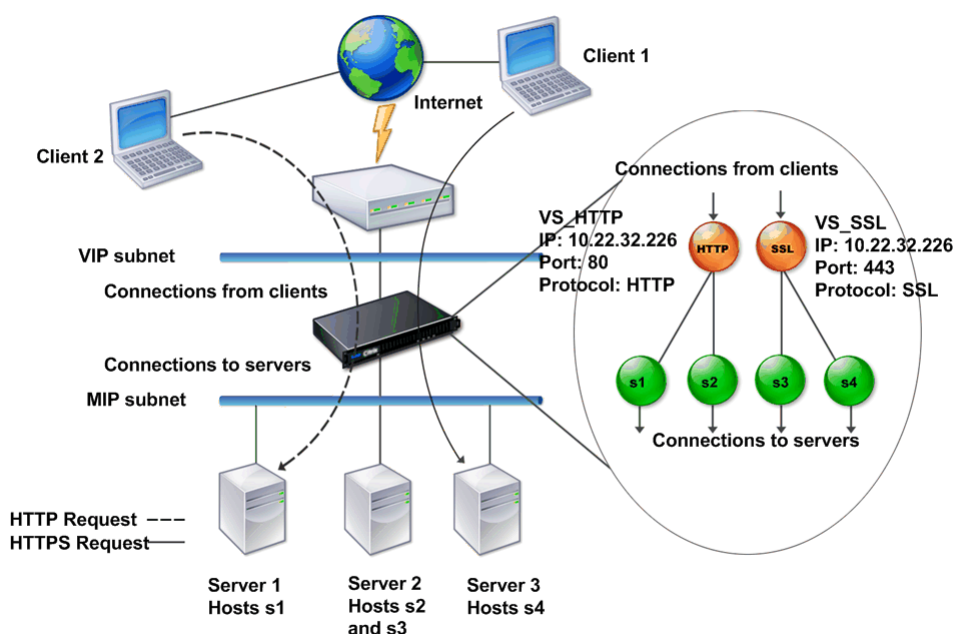
Understanding Virtual Servers

A virtual server is a named NetScaler entity that external clients can use to access applications hosted on the servers. It is represented by an alphanumeric name, virtual IP (VIP) address, port, and protocol. The name of the virtual server is of only local significance and is designed to make the virtual server easier to identify. When a client attempts to access applications on a server, it sends a request to the VIP instead of the IP address of the physical server. When the appliance receives a request at the VIP address, it terminates the connection at the virtual server and uses its own connection with the server on behalf of the client. The port and protocol settings of the virtual server determine the applications that the virtual server represents. For example, a web server can be represented by a virtual server and a service whose port and protocol are set to 80 and HTTP, respectively. Multiple virtual servers can use the same VIP address but different protocols and ports.

Virtual servers are points for delivering features. Most features, like compression, caching, and SSL offload, are normally enabled on a virtual server. When the appliance receives a request at a VIP address, it chooses the appropriate virtual server by the port on which the request was received and its protocol. The appliance then processes the request as appropriate for the features configured on the virtual server.

In most cases, virtual servers work in tandem with services. You can bind multiple services to a virtual server. These services represent the applications running on physical servers in a server farm. After the appliance processes requests received at a VIP address, it forwards them to the servers as determined by the load balancing algorithm configured on the virtual server. The following figure illustrates these concepts.

Figure 4. Multiple Virtual Servers with a Single VIP Address



The preceding figure shows a configuration consisting of two virtual servers with a common VIP address but different ports and protocols. Each of the virtual servers has two services bound to it. The services s1 and s2 are bound to VS_HTTP and represent the HTTP applications on Server 1 and Server 2. The services s3 and s4 are bound to VS_SSL and represent the SSL applications on Server 2 and Server 3 (Server 2 provides both HTTP and SSL applications). When the appliance receives an HTTP request at the VIP address, it processes the request as specified by the settings of VS_HTTP and sends it to either Server 1 or Server 2. Similarly, when the appliance receives an HTTPS request at the VIP address, it processes it as specified by the settings of VS_SSL and it sends it to either Server 2 or Server 3.

Virtual servers are not always represented by specific IP addresses, port numbers, or protocols. They can be represented by wildcards, in which case they are known as wildcard virtual servers. For example, when you configure a virtual server with a wildcard instead of a VIP, but with a specific port number, the appliance intercepts and processes all traffic conforming to that protocol and destined for the predefined port. For virtual servers with wildcards instead of VIPs and port numbers, the appliance intercepts and processes all traffic conforming to the protocol.

Virtual servers can be grouped into the following categories:

Load balancing virtual server

Receives and redirects requests to an appropriate server. Choice of the appropriate server is based on which of the various load balancing methods the user configures.

Cache redirection virtual server

Redirects client requests for dynamic content to origin servers, and requests for static content to cache servers.

Cache redirection virtual servers often work in conjunction with load balancing virtual servers.

Content switching virtual server

Directs traffic to a server on the basis of the content that the client has requested. For example, you can create a content switching virtual server that directs all client requests for images to a server that serves images only. Content switching virtual servers often work in conjunction with load balancing virtual servers.

Virtual private network (VPN) virtual server

Decrypts tunneled traffic and sends it to intranet applications.

SSL virtual server

Receives and decrypts SSL traffic, and then redirects to an appropriate server. Choosing the appropriate server is similar to choosing a load balancing virtual server.

Understanding Services

Updated: 2014-03-12

Services represent applications on a server. While services are normally combined with virtual servers, in the absence of a virtual server, a service can still manage application-specific traffic. For example, you can create an HTTP service on a NetScaler appliance to represent a web server application. When the client attempts to access a web site hosted on the web server, the appliance intercepts the HTTP requests and creates a transparent connection with the web server.

In service-only mode, an appliance functions as a proxy. It terminates client connections, uses a SNIP address to establish a connection to the server, and translates the destination IP addresses of incoming client requests to a SNIP address. Although the clients send requests directly to the IP address of the server, the server sees them as coming from the SNIP address. The appliance translates the IP addresses, port numbers, and sequence numbers.

A service is also a point for applying features. Consider the example of SSL acceleration. To use this feature, you must create an SSL service and bind an SSL certificate to the service. When the appliance receives an HTTPS request, it decrypts the traffic and sends it, in clear text, to the server. Only a limited set of features can be configured in the service-only case.

Services use entities called monitors to track the health of applications. Every service has a default monitor, which is based on the service type, bound to it. As specified by the settings configured on the monitor, the appliance sends probes to the application at regular intervals to determine its state. If the probes fail, the appliance marks the service as down. In such cases, the appliance responds to client requests with an appropriate error message or re-routes the request as determined by the configured load balancing policies.

Introduction to the Citrix NetScaler Product Line

The Citrix NetScaler product line optimizes delivery of applications over the Internet and private networks, combining application-level security, optimization, and traffic management into a single, integrated appliance. You install a NetScaler appliance in your server room and route all connections to your managed servers through it. The NetScaler features that you enable and the policies you set are then applied to incoming and outgoing traffic.

A NetScaler can be integrated into any network as a complement to existing load balancers, servers, caches, and firewalls. It requires no additional client or server side software, and can be configured using the NetScaler web-based GUI and CLI configuration utilities.

NetScaler appliances are available in a variety of hardware platforms that have a range of specifications, including multicore processors.

The NetScaler operating system is the base operating system for all NetScaler hardware platforms. The NetScaler operating system is available in three editions: Standard, Enterprise, and Platinum.

Citrix NetScaler Hardware Platforms

Updated: 2013-11-08

NetScaler hardware is available in a variety of platforms that have a range of hardware specifications, including multicore processors. All hardware platforms support some combination of Fast Ethernet, Gigabit Ethernet, and 10 Gigabit Ethernet interfaces.

The following platforms are available for NetScaler .

- Citrix NetScaler MPX 5500
- Citrix NetScaler MPX 5550/5650
- Citrix NetScaler MPX 7500/9500
- Citrix NetScaler MPX 8200/8400/8600
- Citrix NetScaler MPX 9700/10500/12500/15500
- Citrix NetScaler MPX 11500/13500/14500/16500/18500/20500
- Citrix NetScaler MPX 15000
- Citrix NetScaler MPX 17000
- Citrix NetScaler MPX 17500/19500/21500
- Citrix NetScaler MPX 17550/19550/20550/21550

For more information about the hardware platform specifications, see ["Introduction to the Hardware Platforms."](#)

The following tables list different editions of the NetScaler and the hardware platforms on which they are available.

Table 1. Product Editions and MPX Hardware Platforms

Hardware	MPX 5500	MPX 5550/5650	MPX 7500/9500	MPX 8200/8400/8600	MPX 15000	MPX 17000
Platinum Edition	Yes	Yes	Yes	Yes	Yes	Yes
Enterprise Edition	Yes	Yes	Yes	Yes	Yes	Yes
Standard Edition	Yes	Yes	Yes	Yes	Yes	Yes

Table 2. Product Editions and MPX Hardware Platforms (contd.)

Hardware	MPX 9700/10500/12500/15500	MPX 11500/13500/14500/16500/18500/20500	MPX 17500/19500/21500	MPX 17550/19550/20550/21550
Platinum Edition	Yes	Yes	Yes	Yes
Enterprise Edition	Yes	Yes	Yes	Yes
Standard Edition	Yes	Yes	Yes	Yes

Citrix NetScaler Editions

The NetScaler operating system is available in Standard, Enterprise, and Platinum editions. The Enterprise and Standard editions have limited features available. Feature licenses are required for all editions.

For instructions on how to obtain and install licenses, see "<http://support.citrix.com/article/ctx121062>."

The Citrix NetScaler editions are described as follows:

- *Citrix NetScaler, Standard Edition.* Provides small and medium enterprises with comprehensive Layer 4-Layer 7 (L4-L7) traffic management, enabling increased web application availability.
- *Citrix NetScaler, Enterprise Edition.* Provides web application acceleration and advanced L4-L7 traffic management, enabling enterprises to increase web application performance and availability and reduce datacenter costs.
- *Citrix NetScaler, Platinum Edition.* Provides a web application delivery solution that reduces data center costs and accelerates application performance, with end-to-end visibility of application performance, and provides advanced application security.

The following table summarizes the features supported by each edition in the Citrix NetScaler product line:

Table 3. Citrix NetScaler Application Delivery Product Line Features

Key Features	Platinum Edition	Enterprise Edition	Standard Edition
Application availability	Â	Â	Â
Layer 4 load balancing	Yes	Yes	Yes
Layer 7 content switching	Yes	Yes	Yes
AppExpert rate controls	Yes	Yes	Yes
IPv6 support	Yes	Yes	Yes
Global server load balancing (GSLB)	Yes	Yes	Optional
Dynamic routing protocols	Yes	Yes	No
Surge protection	Yes	Yes	No
Priority queuing	Yes	Yes	No
Application acceleration	Â	Â	Â
Client and server TCP optimizations	Yes	Yes	Yes
Citrix AppCompress for HTTP	Yes	Yes	Optional
Citrix AppCache	Yes	Optional	No
Citrix Branch Repeater client	Yes	No	No
Application security	Â	Â	Â
Layer 4 DoS defenses	Yes	Yes	Yes
Layer 7 content filtering	Yes	Yes	Yes
HTTP/URL Rewrite	Yes	Yes	Yes
NetScaler Gateway, EE SSL VPN	Yes	Yes	Yes
Layer 7 DoS Defenses	Yes	Yes	No
AAA security	Yes	Yes	No
Application firewall with XML security	Yes	Optional	No
Simple manageability	Â	Â	Â
AppExpert visual policy builder	Yes	Yes	Yes
AppExpert service callouts	Yes	Yes	Yes
AppExpert templates	Yes	Yes	Yes
Role-based administration	Yes	Yes	Yes
Configuration wizards	Yes	Yes	Yes

Citrix Command Center	Yes	Yes	No
Citrix EdgeSight for NetScaler	Yes	Optional	No
Web 2.0 optimization	Â	Â	Â
Rich Internet application support	Yes	Yes	Yes
Advanced server offload	Yes	Yes	No
Lower total cost of ownership (TCO)	Â	Â	Â
TCP buffering	Yes	Yes	Yes
TCP multiplexing	Yes	Yes	Yes
SSL offload and acceleration	Yes	Yes	Yes
Cache redirection	Yes	Yes	No
Citrix EasyCall	Yes	No	No

Note: While we have taken care to ensure absolute accuracy when compiling this information, it might change. For the latest information, see Citrix Support at "<http://www.citrix.com>."

Supported Releases on NetScaler Hardware

Updated: 2014-06-30

The following table lists the earliest NetScaler builds for releases that are supported on the NetScaler MPX platforms.

Hardware	Software Release	Software Build #
MPX 5500	10.5	All
	10.1	All
	10.0	All
	9.3	All
MPX 5550/5650	10.5	All
	10.1	All
	10.0	71.6.nc and later
	9.3	59.5.nc and later
MPX 7500/9500	10.5	All
	10.1	All
	10.0	All
	9.3	All
MPX 8005/8015	10.5	All
	10.1	122.17.nc and later
	9.3	65.8.nc and later
MPX 8200/8400/8600	10.5	All
	10.1	All
	10.0	70.7.nc and later
	9.3	58.5.nc and later
MPX 9700/10500/12500	10.5	All
	10.1	All
	10.0	All
	9.3	All
MPX 9700/10500/12500 10G	10.5	All
	10.1	All
	10.0	All

	9.3	All
MPX 15500	10.5	All
	10.1	All
	10.0	All
	9.3	All
MPX 15500 10G	10.5	All
	10.1	All
	10.0	All
	9.3	All
MPX 11500/13500/14500/16500/18500/20500	10.5	All
	10.1	All
	10.0	All
	9.3	52.3.nc and later
MPX 11515/11520/11530/11540/11542	10.5	All
	10.1	123.11.nc and later
	9.3	65.8.nc and later
MPX 15000	10.5	All
	10.1	All
	10.0	All
	9.3	All
MPX 17000	10.5	All
	10.1	All
	10.0	All
	9.3	All
MPX 17500/19500/21500	10.5	All
	10.1	All
	10.0	All
	9.3	All
MPX 17550/19550/20550/21550	10.5	All
	10.1	All
	10.0	All
	9.3	53.5.nc and later
MPX 22040/22060/22080/22100/22120	10.5	51.10.nc and later
	10.1	123.11.nc and later
	9.3	65.8.nc and later
MPX 24100/24150	10.5	51.10.nc and later
	10.1	129.11.nc and later

Supported Browsers

Updated: 2014-06-24

To access the configuration utility and Dashboard, your workstation must have a supported web browser and version 1.6 or above of the Java applet plug-in installed.

Operating System	Browser	Versions
Windows 7	Internet Explorer	8, 9, and 10

	Mozilla Firefox	3.6.25 and above
	Google Chrome	15 and above
Windows 64 bit	Internet Explorer	8 and 9
	Google Chrome	15 and above
MAC	Mozilla Firefox	12 and above
	Safari	5.1.3
	Google Chrome	15 and above

Installing the NetScaler Hardware

Before installing a NetScaler appliance, review the pre-installation checklist. A NetScaler is typically mounted in a rack, and all models ship with rack-rail hardware. All models except the 7000 support small form factor pluggable SFP, XFP, or SFP+ transceivers. After mounting the appliance and installing the transceivers, connect the NetScaler to your network. Use a console cable to connect the NetScaler to a personal computer so that you can perform an initial configuration. After connecting everything else, connect the NetScaler to a power source.

This document includes the following:

- [Unpacking the Appliance](#)
- [Rack Mounting the Appliance](#)
- [Installing and Removing 1G SFP Transceivers](#)
- [Installing and Removing XFP and 10G SFP+ Transceivers](#)
- [Connecting the Cables](#)

Unpacking the Appliance

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

Use the following list to verify that you received everything that should have been included in the box.

- The appliance you ordered
 - One RJ-45 to DB-9 adapter
 - One 6 ft RJ-45/DB-9 cable
 - The following list specifies the number of power cables included for each appliance model:
 - One power cable for the MPX 5500, MPX 5550/5650, MPX 7500/9500, and MPX 8005/8015/8200/8400/8600/8800 appliances
 - Two power cables for the MPX 15000, MPX 17000, MPX 9700/10500/12500/15500, MPX 11500/13500/14500/16500/18500/20500, MPX 11515/11520/11530/11540/11542, MPX 14000, MPX 17500/19500/21500, and MPX 25100T/25160T appliances
 - Four power cables for the MPX 22040/22060/22080/22100/22120 and MPX 24100/24150 appliances
- Note: Make sure that a power outlet is available for each cable.
- Note: For Brazilian customers, Citrix does not ship a power cable. Use a cable that conforms to the **ABNT NBR 14136:2002** standard.
- One standard 4-post rail kit
- Note: If the kit that you received does not fit your rack, contact your Citrix sales representative to order the appropriate kit.

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

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

Rack Mounting the Appliance

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

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

The following table lists the different hardware platforms and the rack units required for each platform.

Table 1. *Height Requirements For Each Platform*

Platform	Number of rack units
MPX 5500	One rack unit
MPX 5550/5650	One rack unit
MPX 7500/9500	One rack unit
MPX 8005/8015/8200/8400/8600/8800	One rack unit
MPX 9700/10500/12500/15500	Two rack units
MPX 14000	One rack unit
MPX 15000, MPX 17000	Two rack units
MPX 11500/13500/14500/16500/18500/20500	Two rack units
MPX 11515/11520/11530/11540/11542	Two rack units
MPX 17500/19500/21500	Two rack units
MPX 17550/19550/20550/21550	Two rack units
MPX 22040/22060/22080/22100/22120	Two rack units
MPX 24100/24150	Two rack units
MPX 25100T/25160T	Two rack units

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

Note: The same rail kit is used for both square-hole and round-hole racks. See "[Installing the Rail Assembly to the Rack](#)" for specific instructions for threaded, round-hole racks.

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

Perform the following tasks to mount the appliance:

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

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

To remove the inner rails from the rail assembly

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

To attach the inner rails to the appliance

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

Figure 1. Attaching inner rails



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

To install the rack rails on the rack

1. If you have a round-hole, threaded rack, skip to step 3.
2. Install square nut retainers into the front post and back post of the rack as shown in the following figures. Before inserting a screw, be sure to align the square nut with the correct hole for your 1U or 2U appliance. The three holes are not evenly spaced.

Figure 2. Installing Retainers into the Front Rack Posts






Figure 3. Installing Retainers into the Rear Rack Posts



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


Figure 4. Installing the Rail Assembly to the Rack



To install the appliance in the rack

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

Figure 5. Rack Mounting the Appliance



Installing and Removing 1G SFP Transceivers

Note: This section applies to the MPX 8005/8015/8200/8400/8600/8800, MPX 9700/10500/12500/15500, MPX 11500/13500/14500/16500/18500/20500, MPX 11515/11520/11530/11540/11542, MPX 22040/22060/22080/22100/22120, and MPX 24100/24150 appliances.

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

Note: The 1G SFP transceiver is hot-swappable from release 9.3 build 47.5 and later on the NetScaler appliances that use the e1k interface. The following platforms support 1G SFP transceivers:

- o MPX 7500/9500
- o MPX 8005/8015/8200/8400/8600/8800
- o MPX 9700/10500/12500/15500
- o MPX 11500/13500/14500/16500/18500/20500
- o MPX 11515/11520/11530/11540/11542
- o MPX 22040/22060/22080/22100/22120
- o MPX 24100/24150

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

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


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

To install a 1G SFP transceiver

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

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

Figure 1. Installing a 1G SFP transceiver



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

To remove a 1G SFP transceiver

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

Installing and Removing XFP and 10G SFP+ Transceivers

Note: This section applies to the MPX 8005/8015/8200/8400/8600/8800, MPX 9700/10500/12500/15500, MPX 15000, MPX 17000, MPX 11500/13500/14500/16500/18500/20500, MPX 11515/11520/11530/11540/11542, MPX 14000, MPX 17500/19500/21500, MPX 17550/19550/20550/21550, MPX 22040/22060/22080/22100/22120, MPX 24100/24150, and MPX 25100T/25160T appliances.

A 10-Gigabit Small Form-Factor Pluggable (XFP or SFP+) is a compact optical transceiver that can operate at speeds of up to 10 gigabits per second. The MPX 15000 and MPX 17000 appliances use XFP transceivers and the MPX 8005/8015/8200/8400/8600/8800, MPX 9700/10500/12500/15500, MPX 11500/13500/14500/16500/18500/20500, MPX 17500/19500/21500, MPX 17550/19550/20550/21550, MPX 22040/22060/22080/22100/22120, and MPX 24100/24150 appliances use 10G SFP+ transceivers. Autonegotiation is enabled by default on the XFP/10G SFP+ ports into which you insert your XFP/10G SFP+ transceiver. As soon as a link between the port and the network is established, the mode is matched on both ends of the cable and for 10G SFP+ transceivers, the speed is also autonegotiated.

Note: An XFP transceiver is **not hot-swappable** on the NetScaler appliances. You must restart a NetScaler appliance after you insert an XFP transceiver.

However, the 10G SFP+ transceiver is hot-swappable from release 9.3 build 57.5 and later on the NetScaler appliances that use the ixgbe (ix) interface. The following platforms support 10G SFP+ transceivers:

- o MPX 8005/8015/8200/8400/8600/8800
- o MPX 9700/10500/12500/15500 10G and 10G FIPS
- o MPX 11500/13500/14500/16500/18500/20500
- o MPX 11515/11520/11530/11540/11542
- o MPX 14000
- o MPX 17500/19500/21500
- o MPX 17550/19550/20550/21550
- o MPX 22040/22060/22080/22100/22120
- o MPX 24100/24150
- o MPX 25100T/25160T

The following platforms support XFP transceivers:

- o MPX 15000
- o MPX 17000


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

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

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

To install an XFP/10G SFP+ transceiver

1. Remove the XFP/10G SFP+ transceiver carefully from its box.
Danger: Do not look directly into fiber optic transceivers and cables. They emit laser beams that can damage your eyes.
2. Align the XFP/10G SFP+ transceiver to the front of the XFP/10G SFP+ transceiver port on the front panel of the appliance.
3. Hold the XFP/10G SFP+ transceiver between your thumb and index finger and insert it into the XFP/10G SFP+ transceiver port, pressing it in until you hear the transceiver snap into place.
4. Move the locking hinge to the DOWN position as shown in the following figure.
Figure 1. Locking an XFP transceiver



5. Verify that the LED is green and blinks twice, which indicates that the transceiver is functioning correctly.
6. Do not remove the dust caps attached to the transceiver and cable until you are ready to insert the cable.

To remove an XFP/10G SFP+ transceiver

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

Connecting the Cables

When the appliance is securely mounted on the rack, you are ready to connect the cables. Ethernet cables and the optional console cable are connected first. Connect the power cable last.

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

Connecting the Ethernet Cables

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

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

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

Figure 1. Inserting an Ethernet cable



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

To connect the Ethernet cable to a 1G SFP fiber, 10G SFP+, or XFP transceiver

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


Connecting the Console Cable

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

To connect the console cable to a computer or terminal

1. Insert the DB-9 connector at the end of the cable into the console port that is located on the front panel of the appliance, as shown in the following figure.

Figure 2. Inserting a console cable



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

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


Connecting the Power Cable

An MPX 5500, MPX 5550/5650, MPX 7500/9500, MPX 8005/8015/8200/8400/8600/8800 appliance has one power cable. All the other appliances come with two power cables, but they can also operate if only one power cable is connected, except the MPX 22040/22060/22080/22100/22120 and MPX 24100/24150 platforms which come with four power cables and require two power cables for proper operation. A separate ground cable is not required, because the three-prong plug provides grounding.

To connect the appliance to the power source

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

Figure 3. Inserting a power cable



2. Connect the other end of the power cable to a standard 110V/220V power outlet.
 3. If a second power supply is provided, repeat steps 1 and 2 to connect the second power supply.
- Note: The MPX 9700/10500/12500/15500, MPX 11500/13500/14500/16500/18500/20500, MPX 11515/11520/11530/11540/11542, MPX 17500/19500/21500, and MPX 17550/19550/20550/21550 appliance emit a high-pitched alert if one power supply fails or if you connect only one power cable to the appliance. To silence the alarm, you can press the small red button located on the back panel of the appliance.

Accessing a Citrix NetScaler

A NetScaler appliance has both a command line interface (CLI) and a graphical user interface (GUI). The GUI includes a configuration utility for configuring the appliance and a statistical utility, called Dashboard. For initial access, all appliances ship with the default NetScaler IP address (NSIP) of 192.168.100.1 and default subnet mask of 255.255.0.0. You can assign a new NSIP and an associated subnet mask during initial configuration.

If you encounter an IP address conflict when deploying multiple NetScaler units, check for the following possible causes:

- Did you select an NSIP that is an IP address already assigned to another device on your network?
- Did you assign the same NSIP to multiple NetScaler appliances?
- The NSIP is reachable on all physical ports. The ports on a NetScaler are host ports, not switch ports.

The following table summarizes the available access methods.

Table 1. Methods for Accessing a NetScaler appliance

Access Method	Port	Default IP Address Required? (Y/N)
CLI	Console	N
CLI and GUI	Ethernet	Y

Using the Command Line Interface

Updated: 2013-09-04

You can access the CLI either locally, by connecting a workstation to the console port, or remotely, by connecting through secure shell (SSH) from any workstation on the same network.

Logging on to the Command Line Interface through the Console Port

The appliance has a console port for connecting to a computer workstation. To log on to the appliance, you need a serial crossover cable and a workstation with a terminal emulation program.

To log on to the CLI through the console port

1. Connect the console port to a serial port on the workstation, as described in .
2. On the workstation, start HyperTerminal or any other terminal emulation program. If the logon prompt does not appear, you may need to press ENTER one or more times to display it.
3. Log on by using the administrator credentials. The command prompt (>) appears on the workstation monitor.

Logging on to the Command Line Interface by using SSH

The SSH protocol is the preferred remote access method for accessing an appliance remotely from any workstation on the same network. You can use either SSH version 1 (SSH1) or SSH version 2 (SSH2.)

If you do not have a working SSH client, you can download and install any of the following SSH client programs:

- PuTTY

Open Source software supported on multiple platforms. Available at:

["http://www.chiark.greenend.org.uk/~sgtatham/putty/"](http://www.chiark.greenend.org.uk/~sgtatham/putty/)

- Vandyke Software SecureCRT

Commercial software supported on the Windows platform. Available at:

["http://www.vandyke.com/products/securecrt/"](http://www.vandyke.com/products/securecrt/)

These programs have been tested by the Citrix NetScaler team, which has verified that they work correctly with a NetScaler appliance. Other programs may also work correctly, but have not been tested.

To verify that the SSH client is installed properly, use it to connect to any device on your network that accepts SSH connections.

To log on to a NetScaler by using an SSH client

1. On your workstation, start the SSH client.
2. For initial configuration, use the default NetScaler IP address (NSIP), which is 192.168.100.1. For subsequent access, use the NSIP that was assigned during initial configuration. Select either SSH1 or SSH2 as the protocol.
3. Log on by using the administrator credentials. For example:

```
login as: nsroot
Using keyboard-interactive authentication.
Password:
Last login: Tue Jun 16 10:37:28 2009 from 10.102.29.9

Done
>
```

Using the Graphical User Interface

Updated: 2014-06-30

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

Additionally, if a license is not present on an MPX appliance when it starts, and you add a license later and restart the appliance, you might lose the certificate binding.

Citrix recommends that you install a certificate-key pair of at least 1024 bytes on a NetScaler ADC for HTTPS access to the configuration utility, and that you install an appropriate license before starting the ADC.

The graphical user interface includes a configuration utility and a statistical utility, called Dashboard, either of which you access through a workstation connected to an Ethernet port on the appliance. If your computer does not have a supported Java plug-in installed, the utility prompts you to download and install the plug-in the first time you log on. If automatic installation fails, you can install the plug-in separately before you attempt to log on to the configuration utility or Dashboard.

The system requirements for the workstation running the GUI are as follows:

- For Windows-based workstations, a Pentium 166 MHz or faster processor with at least 48 MB of RAM is recommended for applets running in a browser using a Java plug-in product. You should have 40 MB free disk space before installing the plug-in.
- For Linux-based workstations, a Pentium platform running Linux kernel v2.2.12 or above, and glibc version 2.12-11 or later. A minimum of 32 MB RAM is required, and 48 MB RAM is recommended. The workstation should support 16-bit color mode, KDE and KWM window managers used in conjunction, with displays set to local hosts.
- For Solaris-based workstations, a Sun running either Solaris 2.6, Solaris 7, or Solaris 8, and the Java 2 Runtime Environment, Standard Edition, version 1.6 or later.

Your workstation must have a supported web browser and version 1.6 or above of the Java applet plug-in installed to access the configuration utility and Dashboard.

The following browsers are supported.

Operating System	Browser	Versions
Windows 7	Internet Explorer	8, 9, and 10
	Mozilla Firefox	3.6.25 and above
	Google Chrome	Latest
Windows 64 bit	Internet Explorer	8 and 9
	Google Chrome	Latest
MAC	Mozilla Firefox	12 and above
	Safari	5.1.3
	Google Chrome	Latest

Using the Configuration Utility

Once you log on to the configuration utility, you can configure the appliance through a graphical interface that includes context-sensitive help.

If your computer does not have a supported Java plug-in installed, the first time you log on to the appliance, the configuration utility will prompt you to download and install the plug-in.

Note: Prior to installing the Java 2 Runtime Environment, ensure that you have installed the full set of required operating system patches needed for the current Java release.

To log on to the configuration utility

1. Open your web browser and enter the NetScaler IP (NSIP) as an HTTP address. If you have not yet set up the initial configuration, enter the default NSIP (<http://192.168.100.1>). The Citrix Logon page appears.

Note: If you have two NetScaler appliances in a high availability setup, make sure that you do not access the GUI by entering the IP address of the secondary NetScaler. If you do so and use the GUI to configure the secondary NetScaler, your configuration changes will not be applied to the primary NetScaler.

2. In the User Name text box, type `nsroot`.
3. In the Password text box, type the administrative password you assigned to the nsroot account during initial configuration and click Login. The Configuration Utility page appears.

Note: If your workstation does not already have a supported version of the Java runtime plug-in installed, the NetScaler prompts you to download the Java Plug-in. After the download is complete, the configuration utility page appears.

If you need to access the online help, select Help from the Help menu at the top right corner.

Using the Statistical Utility

Dashboard, the statistical utility, is a browser-based application that displays charts and tables on which you can monitor the performance of a NetScaler.

To log on to Dashboard

1. Open your web browser and enter the NSIP as an HTTP address (<http://<NSIP>>). The Citrix Logon page appears.
2. In the User Name text box, type `nsroot`.
3. In the Password text box, type the administrative password you assigned to the nsroot account during initial configuration.

Installing the Java Runtime Plug-in

If automatic installation of the Java plug-in fails, you can install the plug-in separately before you attempt to log on to the configuration utility.

Note: Before installing the Java 2 Runtime Environment, make sure that you have installed the full set of required operating system patches needed for the current Java release.

To install the Java runtime plug-in on your workstation

1. In your web browser, enter the NSIP and port number of your appliance: <http://<NSIP>:80> The Java plug-in icon appears.
2. Click the Java plug-in icon and follow the screen prompts to copy the plug-in installer to your workstation hard disk. The Java plug-in setup icon (for example, **j2re-1.6.0**) appears on your computer at the location you specified.
3. Double-click the plug-in setup icon, and follow the screen prompts to install the plug-in.
4. Return to your web browser and click the Java plug-in icon a second time to display the GUI logon screen.

Configuring a NetScaler for the First Time

Initial configuration is the same for the multifunction Citrix NetScaler, the dedicated NetScaler Gateway Enterprise Edition, and the dedicated Citrix NetScaler Application Firewall appliances. You can use any of the following interfaces for initial configuration of your appliance:

- First-time use wizard—If you use a web browser to connect to the appliance, you are prompted to enter the network configuration and licensing information, if it is not already specified.
- LCD keypad—You can specify the network settings, but you must use a different interface to upload your licenses.
- Serial console—After connecting to the serial console, you can use the command line interface to specify the network settings and upload your licenses.
- NITRO API—You can use the NITRO API suite to configure the NetScaler appliance.

For initial configuration, use nsroot as both the administrative user name and the password. For subsequent access, use the password assigned during initial configuration.

If you are setting up two NetScaler appliances as a high availability pair, you configure one as primary and the other as secondary.

The configuration procedure for a FIPS appliance is slightly different from the procedure for a NetScaler MPX appliance or a NetScaler virtual appliance.

Using the First-time Setup Wizard

To configure a NetScaler appliance (or NetScaler virtual appliance) for the first time, you need an administrative computer configured on the same network as the appliance.

You must assign a NetScaler IP (NSIP) address as the management IP address of your NetScaler appliance. This is the address at which you access the NetScaler for configuration, monitoring, and other management tasks. Assign a subnet IP (SNIP) address for your NetScaler to communicate with the backend servers. Specify a host name to identify your NetScaler, an IP address for a DNS server to resolve domain names, and the time zone in which your NetScaler is located.

The wizard automatically appears if any of the following conditions are met:









- The appliance is configured with the default IP address (192.168.100.1).
- A subnet IP address is not configured.
- Licenses are not present on the appliance.

To perform first-time configuration of your appliance

1. In a web browser, type: <http://192.168.100.1>
Note: The NetScaler software is preconfigured with a default IP address. If you have already assigned as NSIP address, type that address in a web browser.
2. In User Name and Password, type the administrator credentials. The following screen appears.

Welcome!

Use this wizard for initial configuration of your NetScaler virtual appliance. To configure or to change a previously configured setting, click each of the sections below. If a parameter has already been configured, a check mark appears within a green circle. An orange circle containing a dash indicates that you have chosen to skip this section.

	NetScaler IP Address IP address at which you access the NetScaler for configuration, monitoring, and other management tasks. NetScaler IP Address: 10.102.29.165 Netmask: 255.255.255.0	
	Subnet IP Address Specify an IP address for your NetScaler to communicate with the backend servers. Subnet IP Address: Not configured	
	Host Name, DNS IP Address, and Time Zone Specify a host name to identify your NetScaler, an IP address for a DNS server to resolve domain names, and the time zone in which your NetScaler is located. Host Name: ns DNS IP Address: Not configured Time Zone: GMT-11:00-10:00-Pacific/Midway	
	Licenses Upload licenses from your local computer or allocate licenses from the Citrix licensing portal. There are 3 license file(s) present on this NetScaler.	

[Continue](#)

3. To configure or to change a previously configured setting, click inside each section. When done, click Continue.
4. When prompted, select Reboot.

Using the LCD Keypad

When you first install the appliance, you can configure the initial settings by using the LCD keypad on the front panel of the appliance. The keypad interacts with the LCD display module, which is also on the front panel of these appliances.

Note: You can use the LCD keypad for initial configuration on a new appliance with the default configuration. The configuration file (ns.conf) should contain the following command and default values.

```
set ns config -IPAddress 192.168.100.1 -netmask 255.255.0.0
```

The functions of the different keys are explained in the following table.

Table 1. LCD Key Functions

Key	Function
<	Moves the cursor one digit to the left.
>	Moves the cursor one digit to the right.
^	Increments the digit under the cursor.
v	Decrements the digit under the cursor.
.	Processes the information, or terminates the configuration, if none of the values are changed. This key is also known as the ENTER key.

To perform the initial configuration by using the LCD keypad press the "<" key.

You are prompted to enter the subnet mask, NetScaler IP address (NSIP), and gateway in that order respectively. The subnet mask is associated with both the NSIP and default gateway IP address. The NSIP is the IPv4 address of the NetScaler appliance. The default gateway is the IPv4 address for the router, which will handle external IP traffic that the NetScaler cannot otherwise route. The NSIP and the default gateway should be on the same subnet.

If you enter a valid value for the subnet mask, such as 255.255.255.224, you are prompted to enter the IP address. Similarly, if you enter a valid value for the IP address, you are prompted to enter the gateway address. If the value you entered is invalid, the following error message appears for three seconds, where xxx.xxx.xxx.xxx is the IP address you entered, followed by a request to re-enter the value.

```
Invalid addr!  
xxx.xxx.xxx.xxx
```

If you press the ENTER (.) key without changing any of the digits, the software interprets this as a user exit request. The following message will be displayed for three seconds.

```
Exiting menu...  
xxx.xxx.xxx.xxx
```

If all the values entered are valid, when you press the ENTER key, the following message appears.

```
Values accepted,  
Rebooting...
```

The subnet mask, NSIP, and gateway values are saved in the configuration file.

Note: For information about deploying a high availability (HA) pair, see "<http://support.citrix.com/proddocs/topic/ns-system-10-5-map/ns-nw-ha-cnfgng-ha-con.html>".

Using the NetScaler Serial Console

When you first install the appliance, you can configure the initial settings by using the serial console. With the serial console, you can change the system IP address, create a subnet or mapped IP address, configure advanced network settings, and change the time zone.

Note: To locate the serial console port on your appliance, see "RS232 Serial Console Port" in "Ports."

To configure initial settings by using a serial console

1. Connect the console cable into your appliance. For more information, see "Connecting the Console Cable" in "Connecting the Cables."
2. Run the vt100 terminal emulation program of your choice on your computer to connect to the appliance and configure the following settings: 9600 baud, 8 data bits, 1 stop bit, parity, and flow control set to NONE.
3. Press ENTER. The terminal screen displays the Logon prompt.
Note: You might have to press ENTER two or three times, depending on which terminal program you are using.
4. Log on to the appliance with the administrator credentials. Your sales representative or Citrix Customer Service can provide you with the administrator credentials.
5. At the prompt, type config ns to run the NetScaler configuration script.

6. To complete the initial configuration of your appliance, follow the prompts.

Note: To prevent an attacker from breaching your ability to send packets to the appliance, choose a non-routable IP address on your organization's LAN as your appliance IP address.

You can replace steps 5 and 6 with the following NetScaler commands. At the NetScaler command prompt, type:

```
set ns config -ipaddress<IPAddress> -netmask<subnetMask>
```

```
add ns ip<IPAddress> <subnetMask> -type<type>
```

```
add route<network> <netmask> <gateway>
```

```
set system user <userName> -password
```

```
save ns config
```

```
reboot
```

Example

```
set ns config -ipaddress 10.102.29.60 -netmask 255.255.255.0
add ns ip 10.102.29.61 255.255.255.0 -type snip
add route 0.0.0.0 0.0.0.0 10.102.29.1
set system user nsroot -password
Enter password: *****
Confirm password: *****
save ns config
reboot
```

You have now completed initial configuration of your appliance. To continue configuring the appliance, choose one of the following options:

Citrix NetScaler.

If you are configuring your appliance as a standard NetScaler with other licensed features, see "[Load Balancing](#)."

Citrix NetScaler Application Firewall.

If you are configuring your appliance as a standalone application firewall, see "[Application Firewall](#)."

NetScaler Gateway.

If you are configuring your appliance as a NetScaler Gateway, see "[NetScaler Gateway 10.5](#)."

Note: For information about deploying a high availability (HA) pair, see "[Configuring High Availability](#)."

Configuring a NetScaler by Using the NITRO API

You can use the NITRO API to configure the NetScaler appliance. NITRO exposes its functionality through Representational State Transfer (REST) interfaces. Therefore, NITRO applications can be developed in any programming language. Additionally, for applications that must be developed in Java or .NET or Python, NITRO APIs are exposed through relevant libraries that are packaged as separate Software Development Kits (SDKs). For more information, see [NITRO API](#).

Configuring a High Availability Pair for the First Time

You can deploy two NetScaler appliances in a high availability configuration, where one unit actively accepts connections and manages servers while the secondary unit monitors the first. The NetScaler that is actively accepting connections and managing the servers is called a primary unit and the other one is called a secondary unit in a high availability configuration. If there is a failure in the primary unit, the secondary unit becomes the primary and begins actively accepting connections.

Each NetScaler in a high availability pair monitors the other by sending periodic messages, called heartbeat messages or health checks, to determine the health or state of the peer node. If a health check for a primary unit fails, the secondary unit retries the connection for a specific time period. For more information about high availability, see "[High Availability](#)." If a retry does not succeed by the end of the specified time period, the secondary unit takes over for the primary unit in a process called failover. The following figure shows two high availability configurations, one in one-arm mode

Figure 1. High availability in one-arm mode





Figure 2. High availability in two-arm mode



In one-arm configuration, both NS1 and NS2 and servers S1, S2, and S3 are connected to the switch.

In two-arm configuration, both NS1 and NS2 are connected to two switches. The servers S1, S2, and S3 are connected to the second switch. The traffic between client and the servers passes through either NS1 or NS2.

To set up a high availability environment, configure one NetScaler as primary and another as secondary. Perform the following tasks on each of the NetScalers:

- Add a node.
- Disable high availability monitoring for unused interfaces.

Adding a Node

Updated: 2013-06-24

A node is a logical representation of a peer NetScaler appliance. It identifies the peer unit by ID and NSIP. An appliance uses these parameters to communicate with the peer and track its state. When you add a node, the primary and secondary units exchange heartbeat messages asynchronously. The node ID is an integer that must not be greater than 64.

To add a node by using the command line interface

At the command prompt, type the following commands to add a node and verify that the node has been added:

- `add HA node <id> <IPAddress>`
- `show HA node <id>`

Example

```

add HA node 0 10.102.29.170
Done
> show HA node 0
1)      Node ID:      0
        IP:    10.102.29.200 (NS200)
        Node State: UP
        Master State: Primary
        SSL Card Status: UP
        Hello Interval: 200 msecs
        Dead Interval: 3 secs
        Node in this Master State for: 1:0:41:50 (days:hrs:min:sec)
  
```

To add a node by using the configuration utility

1. In the navigation pane, expand System and click High Availability. The High Availability page appears.
2. On the High Availability page, select the Nodes tab.
3. Click Add. The High Availability Setup dialog box appears.
4. In the High Availability Setup dialog box, in the Remote Node IP Address text box, type an IP Address (for example, 10.102.29.170).
5. Ensure that the Configure remote system to participate in High Availability setup check box is selected. By default, this check box is selected.
6. Select the Turn off HA monitor on interfaces/channels that are down check box to disable the HA monitor on interfaces that are down. By default, this check box is selected.
7. Verify that the node you added appears in the list of nodes under the Nodes tab.

Disabling High Availability Monitoring for Unused Interfaces

Updated: 2013-06-24

The high availability monitor is a virtual entity that monitors an interface. You must disable the monitor for interfaces that are not connected or being used for traffic. When the monitor is enabled on an interface whose status is DOWN, the state of the node becomes NOT UP. In a high availability configuration, a primary node entering a NOT UP state might cause a high availability failover. An interface is marked DOWN under the following conditions:

- o The interface is not connected
- o The interface is not working properly
- o The cable connecting the interface is not working properly

To disable the high availability monitor for an unused interface by using the command line interface

At the command prompt, type the following commands to disable the high availability monitor for an unused interface and verify that it is disabled:

- o `set interface <id> -haMonitor OFF`
- o `show interface <id>`

Example

```
> set interface 1/8 -haMonitor OFF
Done
> show interface 1/8
Interface 1/8 (Gig Ethernet 10/100/1000 Mbits) #2
flags=0x4000 <ENABLED, DOWN, down, autoneg, 802.1q>
MTU=1514, native vlan=1, MAC=00:d0:68:15:fd:3d, downtime 238h55m44s
Requested: media AUTO, speed AUTO, duplex AUTO, 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)
Bandwidth thresholds are not set.
```

When the high availability monitor is disabled for an unused interface, the output of the show interface command for that interface does not include "HAMON."

To disable the high availability monitor for unused interfaces by using the configuration utility

1. Navigate to System > Network > Interfaces.
2. Select the interface for which the monitor must be disabled.
3. Click Open. The Modify Interface dialog box appears.
4. In HA Monitoring, select the OFF option.
5. Click OK.
6. Verify that, when the interface is selected, "HA Monitoring: OFF" appears in the details at the bottom of the page.

Configuring a FIPS Appliance for the First Time

A certificate-key pair is required for HTTPS access to the configuration utility and for secure remote procedure calls. RPC nodes are internal system entities used for system-to-system communication of configuration and session information. One RPC node exists on each appliance. This node stores the password, which is checked against the one provided by the contacting appliance. To communicate with other NetScaler appliances, each appliance requires knowledge of the other appliances, including how to authenticate on the other appliance. RPC nodes maintain this information, which includes the IP addresses of the other NetScaler appliances and the passwords used to authenticate on each.

On a NetScaler MPX appliance virtual appliance, a certificate-key pair is automatically bound to the internal services. On a FIPS appliance, a certificate-key pair must be imported into the hardware security module (HSM) of a FIPS card. To do so, you must configure the FIPS card, create a certificate-key pair, and bind it to the internal services.

To configure secure HTTPS by using the command line interface

1. Initialize the hardware security module (HSM) on the FIPS card of the appliance. For information about initializing the HSM, see "Configuring the HSM."
2. If the appliance is part of a high availability setup, enable the SIM. For information about enabling the SIM on the primary and secondary appliances, see "Configuring FIPS Appliances in a High Availability Setup."
3. Import the FIPS key into the HSM of the FIPS card of the appliance. At the command prompt, type:

```
import ssl fipskey serverkey -key ns-server.key -inform PEM
```

4. Add a certificate-key pair. At the command prompt, type:

```
add certkey server -cert ns-server.cert -fipskey serverkey
```

5. Bind the certificate-key created in the previous step to the following internal services. At the command prompt, type:

```
bind ssl service nshttps-127.0.0.1-443 -certkeyname server
```

```
bind ssl service nshttps-:::11-443 -certkeyname server
```

To configure secure HTTPS by using the configuration utility

1. Initialize the hardware security module (HSM) on the FIPS card of the appliance. For information about initializing the HSM, see "Configuring the HSM."
2. If the appliance is part of a high availability setup, enable the secure information system (SIM). For information about enabling the SIM on the primary and secondary appliances, see "Configuring FIPS Appliances in a High Availability Setup."
3. Import the FIPS key into the HSM of the FIPS card of the appliance. For more information about importing a FIPS key see "Importing an Existing FIPS Key."
4. Navigate to Traffic Management > SSL > Certificates.
5. In the details pane, click Install.
6. In the Install Certificate dialog box, type the certificate details.
7. Click Create, and then click Close.
8. Navigate to Traffic Management > Load Balancing > Services.
9. In the details pane, on the Action tab, click Internal Services.
10. Select nshttps-127.0.0.1-443 from the list, and then click Open.
11. On the SSL Settings tab, in the Available pane, select the certificate created in step 7, click Add, and then click OK.
12. Select nshttps-:::11-443 from the list, and then click Open.
13. On the SSL Settings tab, in the Available pane, select the certificate created in step 7, click Add, and then click OK.
14. Click OK.

To configure secure RPC by using the command line interface

1. Initialize the hardware security module (HSM) on the FIPS card of the appliance. For information about initializing the HSM, see "Configuring the HSM."
2. Enable the secure information system (SIM). For information about enabling the SIM on the primary and secondary appliances, see "Configuring FIPS Appliances in a High Availability Setup."
3. Import the FIPS key into the HSM of the FIPS card of the appliance. At the command prompt, type:

```
import ssl fipskey serverkey -key ns-server.key -inform PEM
```

4. Add a certificate-key pair. At the command prompt, type:

```
add certkey server -cert ns-server.cert -fipskey serverkey
```

5. Bind the certificate-key pair to the following internal services. At the command prompt, type:


```
bind ssl service nsrpcs-127.0.0.1-3008 -certkeyname server
```

```
bind ssl service nskrpcs-127.0.0.1-3009 -certkeyname server
```

```
bind ssl service nsrpcs-::11-3008 -certkeyname server
```

6. Enable secure RPC mode. At the command prompt, type:

```
set ns rpcnode <IP address> -secure YES
```

To configure secure RPC by using the configuration utility

1. Initialize the hardware security module (HSM) on the FIPS card of the appliance. For information about initializing the HSM, see "Configuring the HSM."
2. Enable the secure information system (SIM). For information about enabling the SIM on the primary and secondary appliances, see "Configuring FIPS Appliances in a High Availability Setup."
3. Import the FIPS key into the HSM of the FIPS card of the appliance. For more information about importing a FIPS key see "Importing an Existing FIPS Key."
4. Navigate to Traffic Management > SSL > Certificates.
5. In the details pane, click Install.
6. In the Install Certificate dialog box, type the certificate details.
7. Click Create, and then click Close.
8. Navigate to Traffic Management > Load Balancing > Services.
9. In the details pane, on the Action tab, click Internal Services.
10. Select nsrpcs-127.0.0.1-3008 from the list, and then click Open.
11. On the SSL Settings tab, in the Available pane, select the certificate created in step 7, click Add, and then click OK.
12. Select nskrpcs-127.0.0.1-3009 from the list, and then click Open.
13. On the SSL Settings tab, in the Available pane, select the certificate created in step 7, click Add, and then click OK.
14. Select nsrpcs-::11-3008 from the list, and then click Open.
15. On the SSL Settings tab, in the Available pane, select the certificate created in step 7, click Add, and then click OK.
16. Click OK.
17. Navigate to System > Network > RPC
18. In the details pane, select the IP address, and click Open.
19. In the Configure RPC Node dialog box, select Secure.
20. Click OK.

Understanding Common Network Topologies

As described in "Physical Deployment Modes," you can deploy the Citrix NetScaler appliance either inline between the clients and servers or in one-arm mode. Inline mode uses a two-arm topology, which is the most common type of deployment.

This document includes the following:

- [Setting Up Common Two-Arm Topologies](#)
- [Setting Up Common One-Arm Topologies](#)

Setting Up Common Two-Arm Topologies

In a two-arm topology, one network interface is connected to the client network and another network interface is connected to the server network, ensuring that all traffic flows through the appliance. This topology might require you to reconnect your hardware and also might result in a momentary downtime. The basic variations of two-arm topology are multiple subnets, typically with the appliance on a public subnet and the servers on a private subnet, and transparent mode, with both the appliance and the servers on the public network.

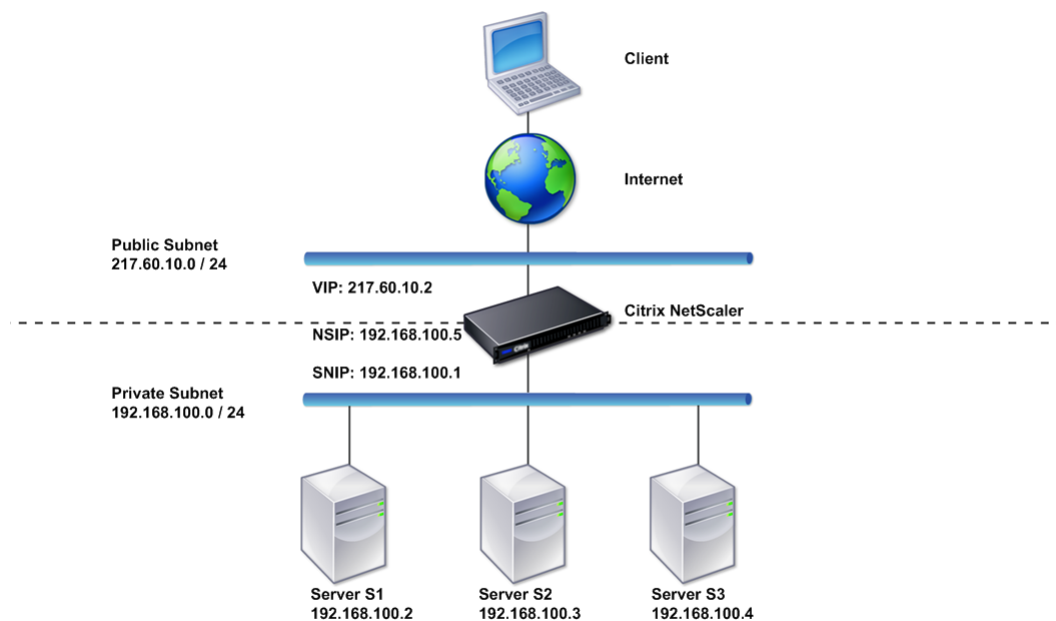
Setting Up a Simple Two-Arm Multiple Subnet Topology

One of the most commonly used topologies has the NetScaler appliance inline between the clients and the servers, with a virtual server configured to handle the client requests. This configuration is used when the clients and servers reside on different subnets. In most cases, the clients and servers reside on public and private subnets, respectively.

For example, consider an appliance deployed in two-arm mode for managing servers S1, S2, and S3, with a virtual server of type HTTP configured on the appliance, and with HTTP services running on the servers. The servers are on a private subnet and a SNIP is configured on the appliance to communicate with the servers. The Use SNIP (USNIP) option must be enabled on the appliance so that it uses the SNIP instead of the MIP.

As shown in the following figure, the VIP is on public subnet 217.60.10.0, and the NSIP, the servers, and the SNIP are on private subnet 192.168.100.0/24.

Figure 1. Topology Diagram for Two-Arm Mode, Multiple Subnets



Task overview: To deploy a NetScaler appliance in two-arm mode with multiple subnets


1. Configure the NSIP and default gateway, as described in "[Configuring the NetScaler IP Address \(NSIP\)](#)."
2. Configure the SNIP, as described in "[Configuring Subnet IP Addresses](#)."
3. Enable the USNIP option, as described in "[To enable or disable USNIP mode](#)."
4. Configure the virtual server and the services, as described in "[Creating a Virtual Server](#)" and "[Configuring Services](#)."

5. Connect one of the network interfaces to a private subnet and the other interface to a public subnet.

Setting Up a Simple Two-Arm Transparent Topology

Use transparent mode if the clients need to access the servers directly, with no intervening virtual server. The server IP addresses must be public because the clients need to be able to access them. In the example shown in the following figure, a NetScaler appliance is placed between the client and the server, so the traffic must pass through the appliance. You must enable L2 mode for bridging the packets. The NSIP and MIP are on the same public subnet, 217.60.10.0/24.

Figure 2. Topology Diagram for Two-Arm, Transparent Mode



Task overview: To deploy a NetScaler in two-arm, transparent mode

1. Configure the NSIP, MIP, and default gateway, as described in "Configuring a NetScaler by Using the Command Line Interface."
2. Enable L2 mode, as described in "Enabling and Disabling Layer 2 Mode."
3. Configure the default gateway of the managed servers as the MIP.
4. Connect the network interfaces to the appropriate ports on the switch.


Setting Up Common One-Arm Topologies

The two basic variations of one-arm topology are with a single subnet and with multiple subnets.

Setting Up a Simple One-Arm Single Subnet Topology

You can use a one-arm topology with a single subnet when the clients and servers reside on the same subnet. For example, consider a NetScaler deployed in one-arm mode for managing servers S1, S2, and S3. A virtual server of type HTTP is configured on a NetScaler, and HTTP services are running on the servers. As shown in the following figure, the NetScaler IP address (NSIP), the Mapped IP address (MIP), and the server IP addresses are on the same public subnet, 217.60.10.0/24.

Figure 3. Topology Diagram for One-Arm Mode, Single Subnet




Task overview: To deploy a NetScaler in one-arm mode with a single subnet

1. Configure the NSIP, MIP, and the default gateway, as described in "Configuring the NetScaler IP Address (NSIP)".
2. Configure the virtual server and the services, as described in "Creating a Virtual Server" and "Configuring Services".
3. Connect one of the network interfaces to the switch.

Setting Up a Simple One-Arm Multiple Subnet Topology

You can use a one-arm topology with multiple subnets when the clients and servers reside on the different subnets. For example, consider a NetScaler appliance deployed in one-arm mode for managing servers S1, S2, and S3, with the servers connected to switch SW1 on the network. A virtual server of type HTTP is configured on the appliance, and HTTP services are running on the servers. These three servers are on the private subnet, so a subnet IP address (SNIP) is configured to communicate with them. The Use Subnet IP address (USNIP) option must be enabled so that the appliance uses the SNIP instead of a MIP. As shown in the following figure, the virtual IP address (VIP) is on public subnet 217.60.10.0/24; the NSIP, SNIP, and the server IP addresses are on private subnet 192.168.100.0/24.

Figure 4. Topology Diagram for One-Arm Mode, Multiple Subnets



Task overview: To deploy a NetScaler appliance in one-arm mode with multiple subnets

1. Configure the NSIP and the default gateway, as described in "[Configuring the NetScaler IP Address \(NSIP\)](#)".
2. Configure the SNIP and enable the USNIP option, as described in "[Configuring Subnet IP Addresses](#)".
3. Configure the virtual server and the services, as described in "[Creating a Virtual Server](#)" and "[Configuring Services](#)".
4. Connect one of the network interfaces to the switch.

Configuring System Management Settings

Once your initial configuration is in place, you can configure settings to define the behavior of the Citrix NetScaler appliance and facilitate connection management. You have a number of options for handling HTTP requests and responses. Routing, bridging, and MAC based forwarding modes are available for handling packets not addressed to the NetScaler. You can define the characteristics of your network interfaces and can aggregate the interfaces. To prevent timing problems, you can synchronize the NetScaler clock with a Network Time Protocol (NTP) server. The NetScaler can operate in various DNS modes, including as an authoritative domain name server (ADNS). You can set up SNMP for system management and customize syslog logging of system events. Before deployment, verify that your configuration is complete and correct.

This document includes the following:

- [Configuring System Settings](#)
- [Configuring Modes of Packet Forwarding](#)
- [Configuring Network Interfaces](#)
- [Configuring Clock Synchronization](#)
- [Configuring DNS](#)
- [Configuring SNMP](#)
- [Verifying the Configuration](#)

Note: In addition to the tasks listed above, you can configure Syslog logging. For instructions, see [“Audit Logging.”](#)

Configuring System Settings

Configuration of system settings includes basic tasks such as configuring HTTP ports to enable connection keep-alive and server offload, setting the maximum number of connections for each server, and setting the maximum number of requests per connection. You can enable client IP address insertion for situations in which a proxy IP address is not suitable, and you can change the HTTP cookie version.

You can also configure a NetScaler appliance to open FTP connections on a controlled range of ports instead of ephemeral ports for data connections. This improves security, because opening all ports on the firewall is insecure. You can set the range anywhere from 1,024 to 64,000.

Before deployment, go through the verification checklists to verify your configuration. To configure HTTP parameters and the FTP port range, use the NetScaler configuration utility.

You can modify the types of HTTP parameters described in the following table.

Table 1. HTTP Parameters

Parameter Type	Specifies
HTTP Port Information	<p>The web server HTTP ports used by your managed servers. If you specify the ports, the appliance performs request switching for any client request that has a destination port matching a specified port.</p> <p>Note: If an incoming client request is not destined for a service or a virtual server that is specifically configured on the appliance, the destination port in the request must match one of the globally configured HTTP ports. This allows the appliance to perform connection keep-alive and server off-load.</p>
Limits	<p>The maximum number of connections to each managed server, and the maximum number of requests sent over each connection. For example, if you set Max Connections to 500, and the appliance is managing three servers, it can open a maximum of 500 connections to each of the three servers. By default, the appliance can create an unlimited number of connections to any of the servers it manages. To specify an unlimited number of requests per connection, set Max Requests to 0.</p> <p>Note: If you are using the Apache HTTP server, you must set Max Connections equal to the value of the MaxClients parameter in the Apache httpd.conf file. Setting this parameter is optional for other web servers.</p>
Client IP Insertion	<p>Enable/disable insertion of the client's IP address into the HTTP request header. You can specify a name for the header field in the adjacent text box. When a web server managed by an appliance receives a mapped IP address or a subnet IP address, the server identifies it as the client's IP address. Some applications need the client's IP address for logging purposes or to dynamically determine the content to be served by the web server.</p> <p>You can enable insertion of the actual client IP address into the HTTP header request sent from the client to one, some, or all servers managed by the appliance. You can then access the inserted address through a minor modification to the server (using an Apache module, ISAPI interface, or NSAPI interface).</p>
Cookie Version	<p>The HTTP cookie version to use when COOKIEINSERT persistence is configured on a virtual server. The default, version 0, is the most common type on the Internet. Alternatively, you can specify version 1.</p>
Requests/Responses	<p>Options for handling certain types of requests, and enable/disable logging of HTTP error responses.</p>
Server Header Insertion	<p>Insert a server header in NetScaler-generated HTTP responses.</p>

To configure HTTP parameters by using the configuration utility

1. In the navigation pane, expand System, and then click Settings.
2. In the details pane, under Settings, click Change HTTP parameters.

3. In the Configure HTTP parameters dialog box, specify values for some or all of the parameters that appear under the headings listed in the table above.
4. Click OK.

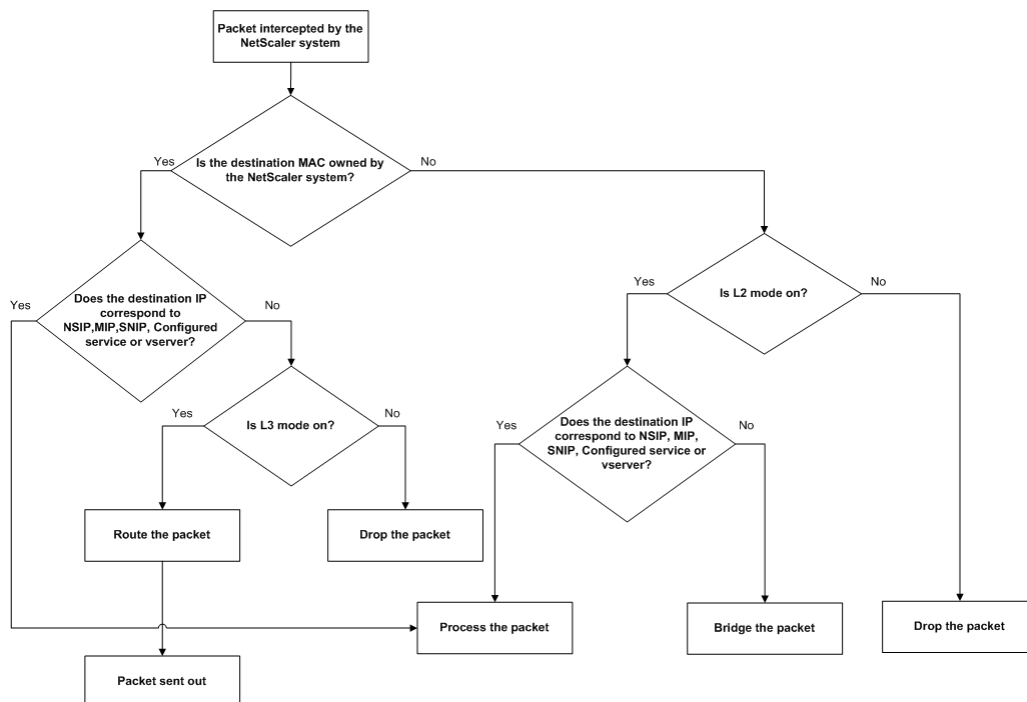
To set the FTP port range by using the configuration utility

1. In the navigation pane, expand System, and then click Settings.
2. In the details pane, under Settings, click Change global system settings.
3. Under FTP Port Range, in the Start Port and End Port text boxes, type the lowest and highest port numbers, respectively, for the range you want to specify (for example, 5000 and 6000).
4. Click OK.

Configuring Modes of Packet Forwarding

The NetScaler appliance can either route or bridge packets that are not destined for an IP address owned by the appliance (that is, the IP address is not the NSIP, a MIP, a SNIP, a configured service, or a configured virtual server). By default, L3 mode (routing) is enabled and L2 mode (bridging) is disabled, but you can change the configuration. The following flow chart shows how the appliance evaluates packets and either processes, routes, bridges, or drops them.

Figure 1. Interaction between Layer 2 and Layer 3 Modes



An appliance can use the following modes to forward the packets it receives:

- Layer 2 (L2) Mode
- Layer 3 (L3) Mode
- MAC-Based Forwarding Mode

Enabling and Disabling Layer 2 Mode

Updated: 2013-09-13

Layer 2 mode controls the Layer 2 forwarding (bridging) function. You can use this mode to configure a NetScaler appliance to behave as a Layer 2 device and bridge the packets that are not destined for it. When this mode is enabled, packets are not forwarded to any of the MAC addresses, because the packets can arrive on any interface of the appliance and each interface has its own MAC address.

With Layer 2 mode disabled (which is the default), the appliance drops packets that are not destined for one of its MAC address. If another Layer 2 device is installed in parallel with the appliance, Layer 2 mode must be disabled to prevent bridging (Layer 2) loops. You can use the configuration utility or the command line to enable Layer 2 mode.

Note: The appliance does not support spanning tree protocol. To avoid loops, if you enable L2 mode, do not connect two interfaces on the appliance to the same broadcast domain.

To enable or disable Layer 2 mode by using the command line interface

At the command prompt, type the following commands to enable/disable Layer 2 mode and verify that it has been enabled/disabled:

- `enable ns mode <Mode>`
- `disable ns mode <Mode>`
- `show ns mode`

Examples

```

> enable ns mode l2
Done
> show ns mode

      Mode                      Acronym      Status
      -----
1)  Fast Ramp                  FR        ON
2)  Layer 2 mode               L2        ON
.
.
.
Done
>

> disable ns mode l2
Done
> show ns mode

      Mode                      Acronym      Status
      -----
1)  Fast Ramp                  FR        ON
2)  Layer 2 mode               L2        OFF
.
.
.
Done
>

```

To enable or disable Layer 2 mode by using the configuration utility

1. In the navigation pane, expand System, and then click Settings.
2. In the details pane, under Modes and Features, click Configure modes.
3. In the Configure Modes dialog box, to enable Layer 2 mode, select the Layer 2 Mode check box. To disable Layer 2 mode, clear the check box.
4. Click OK. The Enable/Disable Mode(s)? message appears in the details pane.
5. Click Yes.

Enabling and Disabling Layer 3 Mode

Updated: 2013-09-13

Layer 3 mode controls the Layer 3 forwarding function. You can use this mode to configure a NetScaler appliance to look at its routing table and forward packets that are not destined for it. With Layer 3 mode enabled (which is the default), the appliance performs route table lookups and forwards all packets that are not destined for any appliance-owned IP address. If you disable Layer 3 mode, the appliance drops these packets.

To enable or disable Layer 3 mode by using the command line interface

At the command prompt, type the following commands to enable/disable Layer 3 mode and verify that it has been enabled/disabled:

- enable ns mode <Mode>
- disable ns mode <Mode>
- show ns mode

Examples

```

> enable ns mode l3
Done
> show ns mode

      Mode                      Acronym      Status
      -----
1)  Fast Ramp                  FR        ON
2)  Layer 2 mode               L2        OFF
.
.
.

```

```

.
9) Layer 3 mode (ip forwarding) L3 ON
.
.
.
Done
>

> disable ns mode l3
Done
> show ns mode

      Mode                      Acronym      Status
      -----
1) Fast Ramp                    FR         ON
2) Layer 2 mode                  L2         OFF
.
.
.
9) Layer 3 mode (ip forwarding) L3         OFF
.
.
.
Done
>

```

To enable or disable Layer 3 mode by using the configuration utility

1. In the navigation pane, expand System, and then click Settings.
2. In the details pane, under Modes and Features, click Configure modes.
3. In the Configure Modes dialog box, to enable Layer 3 mode, select the Layer 3 Mode (IP Forwarding) check box. To disable Layer 3 mode, clear the check box.
4. Click OK. The Enable/Disable Mode(s)? message appears in the details pane.
5. Click Yes.

Enabling and Disabling MAC-Based Forwarding Mode

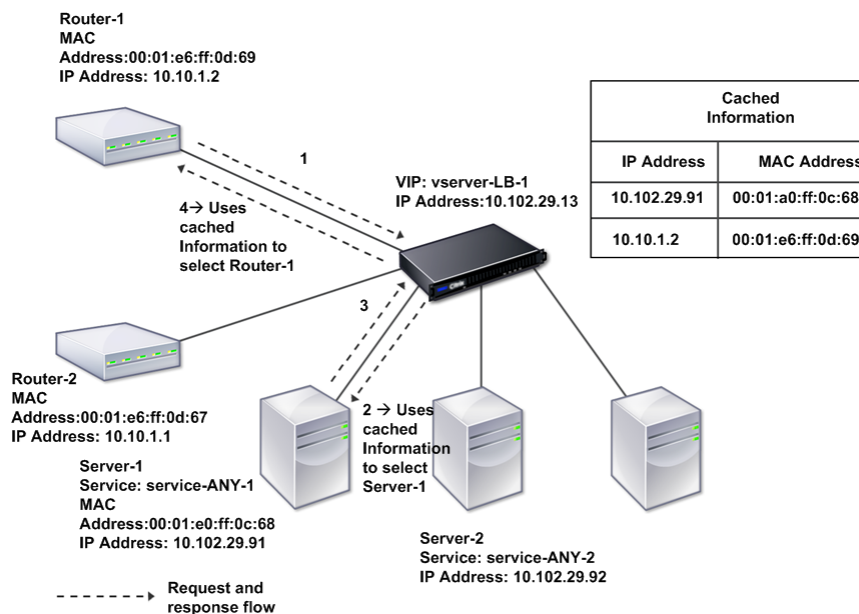
Updated: 2013-09-13

You can use MAC-based forwarding to process traffic more efficiently and avoid multiple-route or ARP lookups when forwarding packets, because the NetScaler appliance remembers the MAC address of the source. To avoid multiple lookups, the appliance caches the source MAC address of every connection for which it performs an ARP lookup, and it returns the data to the same MAC address.

MAC-based forwarding is useful when you use VPN devices because the appliance ensures that all traffic flowing through a particular VPN passes through the same VPN device.

The following figure shows the process of MAC-based forwarding.

Figure 2. MAC-Based Forwarding Process



When MAC-based forwarding is enabled, the appliance 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 responds through an 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 an appliance initiates a connection, it uses the route and ARP tables for the lookup function. To enable MAC-based forwarding, use the configuration utility or the command line.

Some deployments require the incoming and outgoing paths to flow through different routers. In these situations, MAC-based forwarding breaks the topology design. For a global server load balancing (GSLB) site that requires the incoming and outgoing paths to flow through different routers, you must disable MAC-based forwarding and use the appliance's default router as the outgoing router.

With MAC-based forwarding disabled and Layer 2 or Layer 3 connectivity enabled, a route table can specify separate routers for outgoing and incoming connections. To disable MAC-based forwarding, use the configuration utility or the command line.

To enable or disable MAC-based forwarding by using the command line interface

At the command prompt, type the following commands to enable/disable MAC-based forwarding mode and verify that it has been enabled/disabled:

- `enable ns mode <Mode>`
- `disable ns mode <Mode>`
- `show ns mode`

Example

```
> enable ns mode mbf
Done
> show ns mode
```

Mode	Acronym	Status
-----	-----	-----
1) Fast Ramp	FR	ON
2) Layer 2 mode	L2	OFF
.		
.		

```

.
6)  MAC-based forwarding          MBF          ON
.
.
.
Done
>

> disable ns mode mbf
Done
> show ns mode

      Mode                      Acronym          Status
      -----
1)  Fast Ramp                  FR           ON
2)  Layer 2 mode              L2           OFF
.
.
.
6)  MAC-based forwarding        MBF          OFF
.
.
.
Done
>

```

To enable or disable MAC-based forwarding by using the configuration utility

1. In the navigation pane, expand System, and then click Settings.
2. In the details pane, under Modes and Features group, click Configure modes.
3. In the Configure Modes dialog box, to enable MAC-based forwarding mode, select the MAC Based Forwarding check box. To disable MAC-based forwarding mode, clear the check box.
4. Click OK. The Enable/Disable Mode(s)? message appears in the details pane.
5. Click Yes.

Configuring Network Interfaces

NetScaler interfaces are numbered in slot/port notation. In addition to modifying the characteristics of individual interfaces, you can configure virtual LANs to restrict traffic to specific groups of hosts. You can also aggregate links into high-speed channels.

Virtual LANs

The NetScaler supports (Layer 2) port and IEEE802.1Q tagged virtual LANs (VLANs). VLAN configurations are useful when you need to restrict traffic to certain groups of stations. You can configure a network interface to belong to multiple VLANs by using IEEE 802.1q tagging.

You can bind your configured VLANs to IP subnets. The NetScaler (if it is configured as the default router for the hosts on the subnets) then performs IP forwarding between these VLANs. A NetScaler supports the following types of VLANs.

Default VLAN

By default, the network interfaces on a NetScaler are included in a single, port-based VLAN as untagged network interfaces. This default VLAN has a VID of 1 and exists permanently. It cannot be deleted, and its VID cannot be changed.

Port-Based VLANs

A set of network interfaces that share a common, exclusive, Layer 2 broadcast domain define the membership of a port-based VLAN. You can configure multiple port-based VLANs. When you add an interface to a new VLAN as an untagged member, it is automatically removed from the default VLAN.

Tagged VLAN

A network interface can be a tagged or untagged member of a VLAN. Each network interface is an untagged member of only one VLAN (its native VLAN). The untagged network interface forwards the frames for the native VLAN as untagged frames. A tagged network interface can be a part of more than one VLAN. When you configure tagging, be sure that both ends of the link have matching VLAN settings. You can use the configuration utility to define a tagged VLAN (nsvlan) that can have any ports bound as tagged members of the VLAN. Configuring this VLAN requires a reboot of the NetScaler and therefore must be done during initial network configuration.

Link Aggregate Channels

Link aggregation combines incoming data from multiple ports into a single high speed link. Configuring the link aggregate channel increases the capacity and availability of the communication channel between a NetScaler and other connected devices. An aggregated link is also referred to as a channel.

When a network interface is bound to a channel, the channel parameters have precedence over the network interface parameters. A network interface can be bound to only one channel. Binding a network interface to a link aggregate channel changes the VLAN configuration. That is, binding network interfaces to a channel removes them from the VLANs that they originally belonged to and adds them to the default VLAN. However, you can bind the channel back to the old VLAN, or to a new one. For example, if you have bound network interfaces 1/2 and 1/3 to a VLAN with ID 2, and then you bind them to link aggregate channel LA/1, the network interfaces are moved to the default VLAN, but you can bind them to VLAN 2.

Note: You can also use Link Aggregation Control Protocol (LACP) to configure link aggregation. For more information, see "Configuring Link Aggregation by Using the Link Aggregation Control Protocol."

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. NTP uses User Datagram Protocol (UDP) port 123 as its transport layer. You have to add NTP servers in the NTP configuration file so that the appliance periodically gets updates from these servers.

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 at <http://www.ntp.org>.

To configure clock synchronization on your appliance

1. Log on to the command line and enter the `shell` command.
2. At the shell prompt, copy the `ntp.conf` file from the `/etc` directory to the `/nsconfig` directory. If the file already exists in the `/nsconfig` directory, make sure that you remove the following entries from the `ntp.conf` file:

```
restrict localhost  
  
restrict 127.0.0.2
```

These entries are required only if you want to run the device as a time server. However, this feature is not supported on the NetScaler.

3. Edit `/nsconfig/ntp.conf` by typing the IP address for the desired NTP server under the file's server and restrict entries.
4. Create a file named `rc.netscaler` in the `/nsconfig` directory, if the file does not already exist in the directory.
5. Edit `/nsconfig/rc.netscaler` by adding the following entry: `/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.

Note: If the time difference between the NetScaler and the time server is more than 1000 sec, the `ntpd` service terminates with a message to the NetScaler log. To avoid this, you need to start `ntpd` with the `-g` option, which forcibly syncs the time. Add the following entry in `/nsconfig/rc.netscaler`:

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

If you do not want to forcibly sync the time when there is a large difference, you can set the date manually and then start `ntpd` again. You can check the time difference between the appliance and the time server by running the following command in the shell:

```
ntpdate -q <IP address or domain name of the NTP server>
```

6. Reboot the appliance to enable clock synchronization.
Note: If you want to start time synchronization before you restart the appliance, enter the following command (which you added to the `rc.netscaler` file in step 5) at the shell prompt:

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

Configuring DNS

You can configure a NetScaler appliance to function as an Authoritative Domain Name Server (ADNS), DNS proxy server, End Resolver, or Forwarder. You can add DNS resource records such as SRV Records, AAAA Records, A Records, MX Records, NS Records, CNAME Records, PTR Records, and SOA Records. Also, the appliance can balance the load on external DNS servers.

A common practice is to configure an appliance as a forwarder. For this configuration, you need to add external name servers. After you have added the external servers, you should verify that your configuration is correct.

You can add, remove, enable, and disable external name servers. You can create a name server by specifying its IP address, or you can configure an existing virtual server as the name server.

When adding name servers, you can specify IP addresses or virtual IP addresses (VIPs). If you use IP addresses, the appliance load balances requests to the configured name servers in a round robin manner. If you use VIPs, you can specify any load balancing method.

To add a name server by using the command line interface

At the command prompt, type the following commands to add a name server and verify the configuration:

- `add dns nameServer <IP>`
- `show dns nameServer <IP>`

Example

```
> add dns nameServer 10.102.29.10
Done
> show dns nameServer 10.102.29.10
1)      10.102.29.10 - State: DOWN
Done
>
```

To add a name server by using the configuration utility

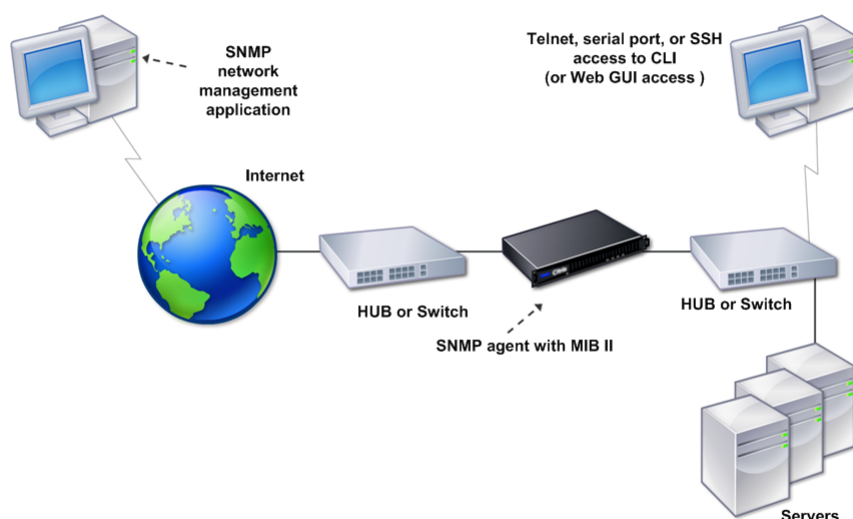
1. Navigate to Traffic Management > DNS > Name Servers.
2. In the details pane, click Add.
3. In the Create Name Server dialog box, select IP Address.
4. In the IP Address text box, type the IP address of the name server (for example, 10.102.29.10). If you are adding an external name server, clear the Local check box.
5. Click Create, and then click Close.
6. Verify that the name server you added appears in the Name Servers pane.

Configuring SNMP

The Simple Network Management Protocol (SNMP) network management application, running on an external computer, queries the SNMP agent on the NetScaler. The agent searches the management information base (MIB) for data requested by the network management application and sends the data to the application.

SNMP monitoring uses traps messages and alarms. SNMP traps messages are asynchronous events that the agent generates to signal abnormal conditions, which are indicated by alarms. For example, if you want to be informed when CPU utilization is above 90 percent, you can set up an alarm for that condition. The following figure shows a network with a NetScaler that has SNMP enabled and configured.

Figure 1. SNMP on the NetScaler



The SNMP agent on a NetScaler supports SNMP version 1 (SNMPv1), SNMP version 2 (SNMPv2), and SNMP version 3 (SNMPv3). Because it operates in bilingual mode, the agent can handle SNMPv2 queries, such as Get-Bulk, and SNMPv1 queries. The SNMP agent also sends traps compliant with SNMPv2 and supports SNMPv2 data types, such as counter64. SNMPv1 managers (programs on other servers that request SNMP information from the NetScaler) use the NS-MIB-smiv1.mib file when processing SNMP queries. SNMPv2 managers use the NS-MIB-smiv2.mib file.

The NetScaler supports the following enterprise-specific MIBs:

- A subset of standard MIB-2 groups
 - Provides MIB-2 groups SYSTEM, IF, ICMP, UDP, and SNMP.
- A system enterprise MIB
 - Provides system-specific configuration and statistics.

To configure SNMP, you specify which managers can query the SNMP agent, add SNMP trap listeners that will receive the SNMP trap messages, and configure SNMP Alarms.

Adding SNMP Managers

Updated: 2013-06-05

You can configure a workstation running a management application that complies with SNMP version 1, 2, or 3 to access an appliance. Such a workstation is called an SNMP manager. If you do not specify an SNMP manager on the appliance, 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 to SNMP queries from only those specific IP addresses. When specifying the IP address of an SNMP manager, you can use the netmask parameter to grant access from entire subnets. You can add a maximum of 100 SNMP managers or networks.

To add an SNMP manager by using the command line interface

At the command prompt, type the following commands to add an SNMP manager and verify the configuration:

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

Example

```
> add snmp manager 10.102.29.5 -netmask 255.255.255.255
Done
> show snmp manager 10.102.29.5
1)      10.102.29.5      255.255.255.255
Done
>
```

To add an SNMP manager by using the configuration utility

1. In the navigation pane, expand System, expand SNMP, and then click Managers.
2. In the details pane, click Add.
3. In the Add SNMP Manager dialog box, in the IP Address text box, type the IP address of the workstation running the management application (for example, 10.102.29.5).
4. Click Create, and then click Close.
5. Verify that the SNMP manager you added appears in the Details section at the bottom of the pane.

Adding SNMP Traps Listeners

Updated: 2013-09-13

After configuring the alarms, you need to specify the trap listener to which the appliance will send the trap messages. Apart from specifying parameters like IP 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.

To add an SNMP trap listener by using the command line interface

At the command prompt, type the following command to add an SNMP trap and verify that it has been added:

- o add snmp trap specific <IP>
- o show snmp trap

Example

```
> add snmp trap specific 10.102.29.3
Done
> show snmp trap
Type      DestinationIP  DestinationPort  Version  SourceIP      Min-Severity  C
----      -
generic   10.102.29.9    162             V2       NetScaler IP  N/A           I
generic   10.102.29.5    162             V2       NetScaler IP  N/A           I
generic   10.102.120.101 162             V2       NetScaler IP  N/A           I
.
.
.
specific  10.102.29.3    162             V2       NetScaler IP  -             I
Done
>
```

To add an SNMP trap listener by using the configuration utility

1. In the navigation pane, expand System, expand SNMP, and then click Traps.
2. In the details pane, click Add.
3. In the Create SNMP Trap Destination dialog box, in the Destination IP Address text box, type the IP address (for example, 10.102.29.3).
4. Click Create and then click Close.
5. Verify that the SNMP trap you added appears in the Details section at the bottom of the pane.

Configuring SNMP Alarms

You configure alarms so that the appliance generates a trap message when an event corresponding to one of the alarms occurs. Configuring an alarm consists of enabling the alarm and setting the severity level at which a trap is generated. There are five severity levels: Critical, Major, Minor, Warning, and Informational. A trap is sent only when the severity of the alarm matches the severity specified for the trap.

Some alarms are enabled by default. If you disable an SNMP alarm, the appliance will not generate trap messages when corresponding events occur. For example, if you disable the Login-Failure SNMP alarm, the appliance will not generate a trap message when a login failure occurs.

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

At the command prompt, type the following commands to enable or disable an alarm and verify that it has been enabled or disabled:

- `set snmp alarm <trapName> [-state ENABLED | DISABLED]`
- `show snmp alarm <trapName>`

Example

```
> set snmp alarm LOGIN-FAILURE -state ENABLED
Done
> show snmp alarm LOGIN-FAILURE
  Alarm          Alarm Threshold  Normal Threshold  Time  State      Severity
  -----          -
1) LOGIN-FAILURE  N/A                        N/A                N/A   ENABLED    -
Done
>
```

To set the severity of the alarm by using the command line interface

At the command prompt, type the following commands to set the severity of the alarm and verify that the severity has been set correctly:

- `set snmp alarm <trapName> [-severity <severity>]`
- `show snmp alarm <trapName>`

Example

```
> set snmp alarm LOGIN-FAILURE -severity Major
Done
> show snmp alarm LOGIN-FAILURE
  Alarm          Alarm Threshold  Normal Threshold  Time  State      Severity
  -----          -
1) LOGIN-FAILURE  N/A                        N/A                N/A   ENABLED    Major
Done
>
```

To configure alarms by using the configuration utility

1. In the navigation pane, expand System, expand SNMP, and then click 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, to enable the alarm, select Enabled in the State drop-down list. To disable the alarm, select Disabled.
4. In the Severity drop-down list, select a severity option (for example, Major).
5. Click OK, and then click Close.
6. Verify that the parameters for the SNMP alarm you configured are correctly configured by viewing the Details section at the bottom of the pane.

Verifying the Configuration

After you finish configuring your system, complete the following checklists to verify your configuration.

Configuration Checklist

- o The build running is:
- o There are no incompatibility issues. (Incompatibility issues are documented in the build's release notes.)
- o The port settings (speed, duplex, flow control, monitoring) are the same as the switch's port.
- o Enough mapped IP addresses have been configured to support all server-side connections during peak times.

The number of configured mapped IP addresses is: _____

The expected number of simultaneous server connections is:

☐ 62,000 ☐ 124,000 ☐ Other _____

Topology Configuration Checklist

- o The routes have been used to resolve servers on other subnets.

The routes entered are:

- o If the NetScaler is in a public-private topology, reverse NAT has been configured.
- o The failover (high availability) settings configured on the NetScaler resolve in a one arm or two-arm configuration. All unused network interfaces have been disabled:

- o If the NetScaler is placed behind an external load balancer, then the load balancing policy on the external load balancer is not "least connection."

The load balancing policy configured on the external load balancer is:

- o If the NetScaler is placed in front of a firewall, the session time-out on the firewall is set to a value greater than or equal to 300 seconds.

Note: The TCP idle connection timeout on a NetScaler appliance is 360 seconds. If the timeout on the firewall is also set to 300 seconds or more, then the appliance can perform TCP connection multiplexing effectively because connections will not be closed earlier.

The value configured for the session time-out is: _____

Server Configuration Checklist

- o "Keep-alive" has been enabled on all the servers.

The value configured for the keep-alive time-out is: _____

- o The default gateway has been set to the correct value. (The default gateway should either be a NetScaler or upstream router.) The default gateway is:

- o The server port settings (speed, duplex, flow control, monitoring) are the same as the switch port settings.

- o If the Microsoft® Internet Information Server is used, buffering is enabled on the server.
- o If an Apache Server is used, the MaxConn (maximum number of connections) parameter is configured on the server and on the NetScaler.

The MaxConn (maximum number of connections) value that has been set is:

- o If a Netscape® Enterprise Server® is used, the maximum requests per connection parameter is set on the NetScaler. The maximum requests per connection value that has been set is:
-

Software Features Configuration Checklist

- o Does the Layer 2 mode feature need to be disabled? (Disable if another Layer 2 device is working in parallel with a NetScaler.)

Reason for enabling or disabling:

- o Does the MAC-based forwarding feature need to be disabled? (If the MAC address used by return traffic is different, it should be disabled.)

Reason for enabling or disabling:

- o Does host-based reuse need to be disabled? (Is there virtual hosting on the servers?)

Reason for enabling or disabling:

- o Do the default settings of the surge protection feature need to be changed?

Reason for changing or not changing:

Access Checklist

- o The system IPs can be pinged from the client-side network.
- o The system IPs can be pinged from the server-side network.
- o The managed server(s) can be pinged through the NetScaler.
- o Internet hosts can be pinged from the managed servers.
- o The managed server(s) can be accessed through the browser.
- o The Internet can be accessed from managed server(s) using the browser.
- o The system can be accessed using SSH.
- o Admin access to all managed server(s) is working.

Note: When you are using the ping utility, ensure that the pinged server has ICMP ECHO enabled, or your ping will not succeed.

Firewall Checklist

The following firewall requirements have been met:

- o UDP 161 (SNMP)
- o UDP 162 (SNMP trap)
- o TCP/UDP 3010 (GUI)
- o HTTP 80 (GUI)
- o TCP 22 (SSH)

Load Balancing Traffic on a NetScaler Appliance

The load balancing feature distributes client requests across multiple servers to optimize resource utilization. In a real-world scenario with a limited number of servers providing service to a large number of clients, a server can become overloaded and degrade the performance of the server farm. A Citrix NetScaler appliance uses load balancing criteria to prevent bottlenecks by forwarding each client request to the server best suited to handle the request when it arrives.


To configure load balancing, you define a virtual server to proxy multiple servers in a server farm and balance the load among them.

When a client initiates a connection to the server, a virtual server terminates the client connection and initiates a new connection with the selected server, or reuses an existing connection with the server, to perform load balancing. The load balancing feature provides traffic management from Layer 4 (TCP and UDP) through Layer 7 (FTP, HTTP, and HTTPS).

The NetScaler appliance uses a number of algorithms, called load balancing methods, to determine how to distribute the load among the servers. The default load balancing method is the Least Connections method.

A typical load balancing deployment consists of the entities described in the following figure.

Figure 1. Load Balancing Architecture



The entities function as follows:

- o **Virtual server.** An entity that is represented by an IP address, a port, and a protocol. The virtual server IP address (VIP) is usually a public IP address. The client sends connection requests to this IP address. The virtual server represents a bank of servers.
- o **Service.** A logical representation of a server or an application running on a server. Identifies the server's IP address, a port, and a protocol. The services are bound to the virtual servers.
- o **Server object.** An entity that is represented by an IP address. The server object is created when you create a service. The IP address of the service is taken as the name of the server object. You can also create a server object and then create services by using the server object.
- o **Monitor.** An entity that tracks the health of the services. The appliance periodically probes the servers using the monitor bound to each service. If a server does not respond within a specified response timeout, and the specified number of probes fails, the service is marked DOWN. The appliance then performs load balancing among the remaining services.

Configuring Load Balancing

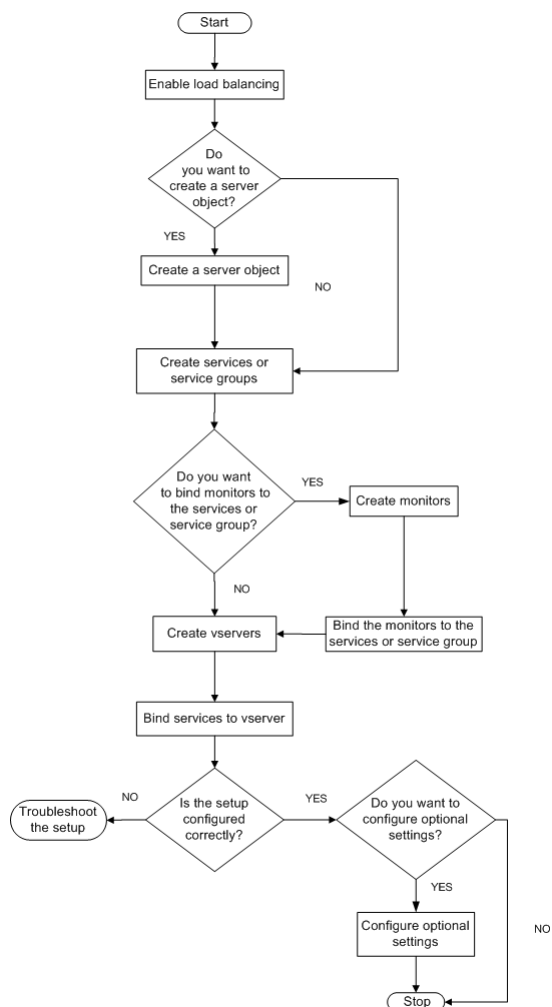
To configure load balancing, you must first create services. Then, you create virtual servers and bind the services to the virtual servers. By default, the NetScaler appliance binds a monitor to each service. After binding the services, verify your configuration by making sure that all of the settings are correct.

Note: After you deploy the configuration, you can display statistics that show how the entities in the configuration are performing. Use the statistical utility or the `stat lb vserver <vserverName>` command.

Optionally, you can assign weights to a service. The load balancing method then uses the assigned weight to select a service. For getting started, however, you can limit optional tasks to configuring some basic persistence settings, for sessions that must maintain a connection to a particular server, and some basic configuration-protection settings.

The following flow chart illustrates the sequence of the configuration tasks.

Figure 1. Sequence of Tasks to Configure Load Balancing



Enabling Load Balancing

Updated: 2013-06-05

Before configuring load balancing, make sure that the load balancing feature is enabled.

To enable load balancing by using the command line interface

At the command prompt, type the following commands to enable load balancing and verify that it is enabled:

- `enable feature lb`
- `show feature`

Example

```

> enable feature lb
Done
> show feature

      Feature                      Acronym      Status
      -----                      -
1)  Web Logging                    WL          OFF
2)  Surge Protection               SP          OFF
3)  Load Balancing                LB          ON
.
.
.
9)  SSL Offloading                 SSL         ON
.
.
.
Done

```

To enable load balancing by using the configuration utility

1. In the navigation pane, expand System, and then click Settings.
2. In the details pane, under Modes and Features, click Change basic features.
3. In the Configure Basic Features dialog box, select the Load Balancing check box, and then click OK.
4. In the Enable/Disable Feature(s)? message, click Yes.

Configuring Services and a Virtual Server

Updated: 2013-06-24

When you have identified the services you want to load balance, you can implement your initial load balancing configuration by creating the service objects, creating a load balancing virtual server, and binding the service objects to the virtual server.

To implement the initial load balancing configuration by using the command line interface

At the command prompt, type the following commands to implement and verify the initial configuration:

- o add service <name> <IPAddress> <serviceType> <port>
- o add lb vserver <vServerName> <serviceType> [<IPAddress> <port>]
- o bind lb vserver <name> <serviceName>
- o show service bindings <serviceName>

Example

```

> add service service-HTTP-1 10.102.29.5 HTTP 80
Done
> add lb vserver vserver-LB-1 HTTP 10.102.29.60 80
Done
> bind lb vserver vserver-LB-1 service-HTTP-1
Done
> show service bindings service-HTTP-1
    service-HTTP-1 (10.102.29.5:80) - State : DOWN

    1)          vserver-LB-1 (10.102.29.60:80) - State : DOWN
Done

```

To implement the initial load balancing configuration by using the configuration utility

1. Navigate to Traffic Management > Load Balancing.
2. In the details pane, under Getting Started, click Load Balancing wizard, and follow the instructions to create a basic load balancing setup.
3. Return to the navigation pane, expand Load Balancing, and then click Virtual Servers.
4. Select the virtual server that you configured and verify that the parameters displayed at the bottom of the page are correctly configured.
5. Click Open.

6. Verify that each service is bound to the virtual server by confirming that the Active check box is selected for each service on the Services tab.

Choosing and Configuring Persistence Settings

You must configure persistence on a virtual server if you want to maintain the states of connections on the servers represented by that virtual server (for example, connections used in e-commerce). The appliance then uses the configured load balancing method for the initial selection of a server, but forwards to that same server all subsequent requests from the same client.

If persistence is configured, it overrides the load balancing methods once the server has been selected. If the configured persistence applies to a service that is down, the appliance uses the load balancing methods to select a new service, and the new service becomes persistent for subsequent requests from the client. If the selected service is in an Out Of Service state, it continues to serve the outstanding requests but does not accept new requests or connections. After the shutdown period elapses, the existing connections are closed. The following table lists the types of persistence that you can configure.

Table 1. Limitations on Number of Simultaneous Persistent Connections

Persistence Type	Persistent Connections
Source IP, SSL Session ID, Rule, DESTIP, SRCIPDESTIP	250K
CookieInsert, URL passive, Custom Server ID	Memory limit. In case of CookieInsert, if time out is not 0, any number of connections is allowed until limited by memory.

If the configured persistence cannot be maintained because of a lack of resources on an appliance, the load balancing methods are used for server selection. Persistence is maintained for a configured period of time, depending on the persistence type. Some persistence types are specific to certain virtual servers. The following table shows the relationship.

Table 2. Persistence Types Available for Each Type of Virtual Server

Persistence TypeHeader 1	HTTP	HTTPS	TCP	UDP/IP	SSL_Bridge
Source IP	YES	YES	YES	YES	YES
CookieInsert	YES	YES	NO	NO	NO
SSL Session ID	NO	YES	NO	NO	YES
URL Passive	YES	YES	NO	NO	NO
Custom Server ID	YES	YES	NO	NO	NO
Rule	YES	YES	NO	NO	NO
SRCIPDESTIP	N/A	N/A	YES	YES	N/A
DESTIP	N/A	N/A	YES	YES	N/A

You can also specify persistence for a group of virtual servers. When you enable persistence on the group, the client requests are directed to the same selected server regardless of which virtual server in the group receives the client request. When the configured time for persistence elapses, any virtual server in the group can be selected for incoming client requests.

Two commonly used persistence types are persistence based on cookies and persistence based on server IDs in URLs.

Configuring Persistence Based on Cookies

Updated: 2013-08-23

When you enable persistence based on cookies, the NetScaler adds an HTTP cookie into the Set-Cookie header field of the HTTP response. The cookie contains information about the service to which the HTTP requests must be sent. The client stores the cookie and includes it in all subsequent requests, and the NetScaler uses it to select the service for those requests. You can use this type of persistence on virtual servers of type HTTP or HTTPS.

The NetScaler inserts the cookie <NSC_XXXX>= <ServiceIP> <ServicePort>

where:

- <NSC_XXXX> is the virtual server ID that is derived from the virtual server name.
- <ServiceIP> is the hexadecimal value of the IP address of the service.
- <ServicePort> is the hexadecimal value of the port of the service.

The NetScaler encrypts ServiceIP and ServicePort when it inserts a cookie, and decrypts them when it receives a cookie.

Note: If the client is not allowed to store the HTTP cookie, the subsequent requests do not have the HTTP cookie, and persistence is not honored.

By default, the NetScaler sends HTTP cookie version 0, in compliance with the Netscape specification. It can also send version 1, in compliance with RFC 2109.

You can configure a timeout value for persistence that is based on HTTP cookies. Note the following:

- If HTTP cookie version 0 is used, the NetScaler inserts the absolute Coordinated Universal Time (GMT) of the cookie's expiration (the expires attribute of the HTTP cookie), calculated as the sum of the current GMT time on a NetScaler, and the timeout value.
- If an HTTP cookie version 1 is used, the NetScaler inserts a relative expiration time (Max-Age attribute of the HTTP cookie). In this case, the client software calculates the actual expiration time.

Note: Most client software currently installed (Microsoft Internet Explorer and Netscape browsers) understand HTTP cookie version 0; however, some HTTP proxies understand HTTP cookie version 1.

If you set the timeout value to 0, the NetScaler does not specify the expiration time, regardless of the HTTP cookie version used. The expiration time then depends on the client software, and such cookies are not valid if that software is shut down. This persistence type does not consume any system resources. Therefore, it can accommodate an unlimited number of persistent clients.

An administrator can use the procedure in the following table to change the HTTP cookie version.

To change the HTTP cookie version by using the configuration utility

1. Navigate to System > Settings.
2. In the details pane, click Change HTTP Parameters.
3. In the Configure HTTP Parameters dialog box, under Cookie, select Version 0 or Version 1.

Note: For information about the parameters, see "[Configuring Persistence Based on Cookies](#)."

To configure persistence based on cookies by using the command line interface

At the command prompt, type the following commands to configure persistence based on cookies and verify the configuration:

- set lb vserver <name> -persistenceType COOKIEINSERT
- show lb vserver <name>

Example

```
> set lb vserver vserver-LB-1 -persistenceType COOKIEINSERT
Done
> show lb vserver vserver-LB-1
vserver-LB-1 (10.102.29.60:80) - HTTP      Type: ADDRESS
.
.
.
Persistence: COOKIEINSERT (version 0)    Persistence Timeout: 2 min
.
.
.
Done
>
```

To configure persistence based on cookies by using the configuration utility

1. Navigate to Traffic Management > Load Balancing > Virtual Servers.
2. In the details pane, select the virtual server for which you want to configure persistence (for example, vserver-LB-1), and then click Open.
3. In the Configure Virtual Server (Load Balancing) dialog box, on the Method and Persistence tab, in the Persistence list, select COOKIEINSERT.
4. In the Time-out (min) text box, type the time-out value (for example, 2).
5. Click OK.

6. Verify that the virtual server for which you configured persistence is correctly configured by selecting the virtual server and viewing the Details section at the bottom of the pane.

Configuring Persistence Based on Server IDs in URLs

Updated: 2013-08-23

The NetScaler can maintain persistence based on the server IDs in the URLs. In a technique called URL passive persistence, the NetScaler extracts the server ID from the server response and embeds it in the URL query of the client request. The server ID is an IP address and port specified as a hexadecimal number. The NetScaler extracts the server ID from subsequent client requests and uses it to select the server.

URL passive persistence requires configuring either a payload expression or a policy infrastructure expression specifying the location of the server ID in the client requests. For more information about expressions, see "[Policy Configuration and Reference](#)."

Note: If the server ID cannot be extracted from the client requests, server selection is based on the load balancing method.

Example: Payload Expression

The expression, URLQUERY contains sid= configures the system to extract the server ID from the URL query of a client request, after matching token sid=. Thus, a request with the URL `http://www.citrix.com/index.asp?&sid;=c0a864100050` is directed to the server with the IP address 10.102.29.10 and port 80.

The timeout value does not affect this type of persistence, which is maintained as long as the server ID can be extracted from the client requests. This persistence type does not consume any system resources, so it can accommodate an unlimited number of persistent clients.

Note: For information about the parameters, see "[Load Balancing](#)."

To configure persistence based on server IDs in URLs by using the command line interface

At the command prompt, type the following commands to configure persistence based on server IDs in URLs and verify the configuration:

- `set lb vserver <name> -persistenceType URLPASSIVE`
- `show lb vserver <name>`

Example

```
> set lb vserver vserver-LB-1 -persistenceType URLPASSIVE
Done
> show lb vserver vserver-LB-1
vserver-LB-1 (10.102.29.60:80) - HTTP      Type: ADDRESS
.
.
.
Persistence: URLPASSIVE Persistence Timeout: 2 min
.
.
.
Done
>
```

To configure persistence based on server IDs in URLs by using the configuration utility

1. Navigate to Traffic Management > Load Balancing > Virtual Servers.
2. In the details pane, select the virtual server for which you want to configure persistence (for example, vserver-LB-1), and then click Open.
3. In the Configure Virtual Server (Load Balancing) dialog box, on the Method and Persistence tab, in the Persistence list, select URLPASSIVE.
4. In the Time-out (min) text box, type the time-out value (for example, 2).
5. In the Rule text box, enter a valid expression. Alternatively, click Configure next to the Rule text box and use the Create Expression dialog box to create an expression.
6. Click OK.
7. Verify that the virtual server for which you configured persistence is correctly configured by selecting the virtual server and viewing the Details section at the bottom of the pane.

Configuring Features to Protect the Load Balancing Configuration

You can configure URL redirection to provide notifications of virtual server malfunctions, and you can configure backup virtual servers to take over if a primary virtual server becomes unavailable.

Configuring URL Redirection

Updated: 2013-06-24

You can configure a redirect URL to communicate the status of the appliance in the event that a virtual server of type HTTP or HTTPS is down or disabled. This URL can be a local or remote link. The appliance uses HTTP 302 redirect.

Redirects can be absolute URLs or relative URLs. If the configured redirect URL contains an absolute URL, the HTTP redirect is sent to the configured location, regardless of the URL specified in the incoming HTTP request. If the configured redirect URL contains only the domain name (relative URL), the HTTP redirect is sent to a location after appending the incoming URL to the domain configured in the redirect URL.

Note: If a load balancing virtual server is configured with both a backup virtual server and a redirect URL, the backup virtual server takes precedence over the redirect URL. In this case, a redirect is used when both the primary and backup virtual servers are down.

To configure a virtual server to redirect client requests to a URL by using the command line interface

At the command prompt, type the following commands to configure a virtual server to redirect client requests to a URL and verify the configuration:

- o set lb vserver <name> -redirectURL <URL>
- o show lb vserver <name>

Example

```
> set lb vserver vserver-LB-1 -redirectURL http://www.newdomain.com/mysite/maint
Done
> show lb vserver vserver-LB-1
vserver-LB-1 (10.102.29.60:80) - HTTP      Type: ADDRESS
State: DOWN
Last state change was at Wed Jun 17 08:56:34 2009 (+666 ms)
.
.
.
Redirect URL: http://www.newdomain.com/mysite/maintenance
.
.
.
Done
>
```

To configure a virtual server to redirect client requests to a URL by using the configuration utility

1. Navigate to Traffic Management > Load Balancing > Virtual Servers.
2. In the details pane, select the virtual server for which you want to configure URL redirection (for example, vserver-LB-1), and then click Open.
3. In the Configure Virtual Server (Load Balancing) dialog box, on the Advanced tab, in the Redirect URL text box, type the URL (for example, http://www.newdomain.com/mysite/maintenance), and then click OK.
4. Verify that the redirect URL you configured for the server appears in the Details section at the bottom of the pane.

Configuring Backup Virtual Servers

Updated: 2013-06-24

If the primary virtual server is down or disabled, the appliance can direct the connections or client requests to a backup virtual server that forwards the client traffic to the services. The appliance can also send a notification message to the client regarding the site outage or maintenance. The backup virtual server is a proxy and is transparent to the client.

You can configure a backup virtual server when you create a virtual server or when you change the optional parameters of an existing virtual server. You can also configure a backup virtual server for an existing backup virtual server, thus

creating a cascaded backup virtual server. The maximum depth of cascading backup virtual servers is 10. The appliance searches for a backup virtual server that is up and accesses that virtual server to deliver the content.

You can configure URL redirection on the primary for use when the primary and the backup virtual servers are down or have reached their thresholds for handling requests.

Note: If no backup virtual server exists, an error message appears, unless the virtual server is configured with a redirect URL. If both a backup virtual server and a redirect URL are configured, the backup virtual server takes precedence.

To configure a backup virtual server by using the command line interface

At the command prompt, type the following commands to configure a backup server and verify the configuration:

- o set lb vserver <name> [-backupVserver <string>]
- o show lb vserver <name>

Example

```
> set lb vserver vserver-LB-1 -backupVserver vserver-LB-2
Done
> show lb vserver vserver-LB-1
vserver-LB-1 (10.102.29.60:80) - HTTP    Type: ADDRESS
State: DOWN
Last state change was at Wed Jun 17 08:56:34 2009 (+661 ms)
.
.
.
Backup: vserver-LB-2
.
.
.
Done
>
```

To set up a backup virtual server by using the configuration utility


1. Navigate to Traffic Management > Load Balancing > Virtual Servers.
2. In the details pane, select the virtual server for which you want to configure the backup virtual server (for example, vserver-LB-1), and then click Open.
3. In the Configure Virtual Server (Load Balancing) dialog box, on the Advanced tab, in the Backup Virtual Server list, select the backup virtual server (for example, vserver-LB-2, and then click OK.
4. Verify that the backup virtual server you configured appears in the Details section at the bottom of the pane.
Note: If the primary server goes down and then comes back up, and you want the backup virtual server to function as the primary server until you explicitly reestablish the primary virtual server, select the Disable Primary When Down check box.

A Typical Load Balancing Scenario

In a load balancing setup, the NetScaler appliances are logically located between the client and the server farm, and they manage traffic flow to the servers.

The following figure shows the topology of a basic load balancing configuration.

Figure 1. Basic Load Balancing Topology




The virtual server selects the service and assigns it to serve client requests. Consider the scenario in the preceding figure, where the services service-HTTP-1 and service-HTTP-2 are created and bound to the virtual server named virtual server-LB-1. Virtual server-LB-1 forwards the client request to either service-HTTP-1 or service-HTTP-2. The system selects the service for each request by using the Least Connections load balancing method. The following table lists the names and values of the basic entities that must be configured on the system.

Table 1. LB Configuration Parameter Values

Entity Type	Required parameters and sample values			
	Name	IP Address	Port	Protocol
Virtual Server	vserver-LB-1	10.102.29.60	80	HTTP
Services	service-HTTP-1	10.102.29.5	8083	HTTP
	service-HTTP-2	10.102.29.6	80	HTTP
Monitors	Default	None	None	None

The following figure shows the load balancing sample values and required parameters that are described in the preceding table.

Figure 2. Load Balancing Entity Model



The following tables list the commands used to configure this load balancing setup by using the command line interface.

Table 2. Initial Configuration Tasks

Task	Command
To enable load balancing	enable feature lb
To create a service named service-HTTP-1	add service service-HTTP-1 10.102.29.5 HTTP 80
To create a service named service-HTTP-2	add service service-HTTP-2 10.102.29.6 HTTP 80
To create a virtual server named vserver-LB-1	add lb vserver vserver-LB-1 HTTP 10.102.29.60 80
To bind a service named service-HTTP-1 to a virtual server named vserver-LB-1	bind lb vserver vserver-LB-1 service-HTTP-1
To bind a service named service-HTTP-2 to a virtual server named vserver-LB-1	bind lb vserver vserver-LB-1 service-HTTP-2

For more information about the initial configuration tasks, see "Enabling Load Balancing" and "Configuring Services and a Vserver."

Table 3. Verification Tasks

Task	Command
To view the properties of a virtual server named vserver-LB-1	show lb vserver vserver-LB-1
To view the statistics of a virtual server named vserver-LB-1	stat lb vserver vserver-LB-1
To view the properties of a service named service-HTTP-1	show service service-HTTP-1
To view the statistics of a service named service-HTTP-1	stat service service-HTTP-1
To view the bindings of a service named service-HTTP-1	show service bindings service-HTTP-1

Table 4. Customization Tasks

Task	Command
------	---------

To configure persistence on a virtual server named vserver-LB-1	set lb vserver vserver-LB-1 -persistenceType SOURCEIP -persistenceMask 255.255.255.255 -timeout 2
To configure COOKIEINSERT persistence on a virtual server named vserver-LB-1	set lb vserver vserver-LB-1 -persistenceType COOKIEINSERT
To configure URLPassive persistence on a virtual server named vserver-LB-1	set lb vserver vserver-LB-1 -persistenceType URLPASSIVE
To configure a virtual server to redirect the client request to a URL on a virtual server named vserver-LB-1	set lb vserver vserver-LB-1 -redirectURL http://www.newdomain.com/mysite/maintenance
To set a backup virtual server on a virtual server named vserver-LB-1	set lb vserver vserver-LB-1 -backupVserver vserver-LB-2

For more information about configuring persistence, see ["Choosing and Configuring Persistence Settings."](#) For information about configuring a virtual server to redirect a client request to a URL and setting up a backup virtual server, see ["Configuring Features to Protect the Load Balancing Configuration."](#)

Accelerating Load Balanced Traffic by Using Compression

Compression is a popular means of optimizing bandwidth usage, and most web browsers support compressed data. If you enable the compression feature, the NetScaler appliance intercepts requests from clients and determines whether the client can accept compressed content. After receiving the HTTP response from the server, the appliance examines the content to determine whether it is compressible. If the content is compressible, the appliance compresses it, modifies the response header to indicate the type of compression performed, and forwards the compressed content to the client.

NetScaler compression is a policy-based feature. A policy filters requests and responses to identify responses to be compressed, and specifies the type of compression to apply to each response. The appliance provides several built-in policies to compress common MIME types such as text/html, text/plain, text/xml, text/css, text/rtf, application/msword, application/vnd.ms-excel, and application/vnd.ms-powerpoint. You can also create custom policies. The appliance does not compress compressed MIME types such as application/octet-stream, binary, bytes, and compressed image formats such as GIF and JPEG.

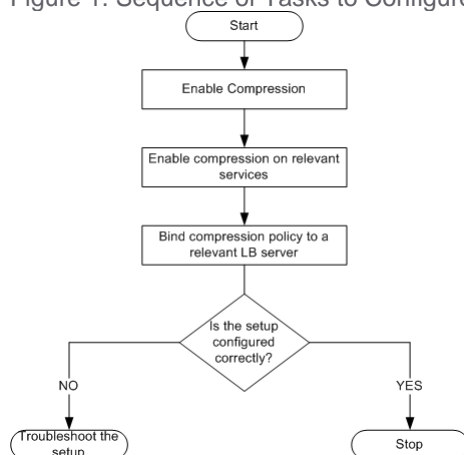
To configure compression, you must enable it globally and on each service that will provide responses that you want compressed. If you have configured virtual servers for load balancing or content switching, you should bind the policies to the virtual servers. Otherwise, the policies apply to all traffic that passes through the appliance.

Compression Configuration Task Sequence

Updated: 2013-08-22

The following flow chart shows the sequence of tasks for configuring basic compression in a load balancing setup.

Figure 1. Sequence of Tasks to Configure Compression



Note: The steps in the above figure assume that load balancing has already been configured.

Enabling Compression

Updated: 2013-06-07

By default, compression is not enabled. You must enable the compression feature to allow compression of HTTP responses that are sent to the client.

To enable compression by using the command line interface

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

- `enable ns feature CMP`
- `show ns feature`

Example

```
> enable ns feature CMP
Done
> show ns feature
```

Feature

Acronym

Status

1)	Web Logging	WL	ON
2)	Surge Protection	SP	OFF
.			
7)	Compression Control	CMP	ON
8)	Priority Queuing	PQ	OFF
.			

Done

To enable compression by using the configuration utility

1. In the navigation pane, expand System, and then click Settings.
2. In the details pane, under Modes and Features, click Change basic features.
3. In the Configure Basic Features dialog box, select the Compression check box, and then click OK.
4. In the Enable/Disable Feature(s)? dialog box, click Yes.

Configuring Services to Compress Data

Updated: 2013-08-22

In addition to enabling compression globally, you must enable it on each service that will deliver files to be compressed.

To enable compression on a service by using the command line

At the command prompt, type the following commands to enable compression on a service and verify the configuration:

- o set service <name> -CMP YES
- o show service <name>

Example

```
> show service SVC_HTTP1
SVC_HTTP1 (10.102.29.18:80) - HTTP
State: UP
Last state change was at Tue Jun 16 06:19:14 2009 (+737 ms)
Time since last state change: 0 days, 03:03:37.200
Server Name: 10.102.29.18
Server ID : 0    Monitor Threshold : 0
Max Conn: 0     Max Req: 0         Max Bandwidth: 0 kbits
Use Source IP: NO
Client Keepalive(CKA): NO
Access Down Service: NO
TCP Buffering(TCPB): NO
HTTP Compression(CMP): YES
Idle timeout: Client: 180 sec   Server: 360 sec
Client IP: DISABLED
Cacheable: NO
SC: OFF
SP: OFF
Down state flush: ENABLED

1)      Monitor Name: tcp-default
State: DOWN    Weight: 1
Probes: 1095   Failed [Total: 1095 Current: 1095]
Last response: Failure - TCP syn sent, reset received.
Response Time: N/A
Done
```

To enable compression on 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 configure compression (for example, service-HTTP-1), and then click Open.
3. On the Advanced tab, under Settings, select the Compression check box, and then click OK.
4. Verify that, when the service is selected, HTTP Compression(CMP): ON appears in the **Details** section at the bottom of the pane.

Binding a Compression Policy to a Virtual Server

Updated: 2013-09-04

If you bind a policy to a virtual server, the policy is evaluated only by the services associated with that virtual server. You can bind compression policies to a virtual server either from the Configure Virtual Server (Load Balancing) dialog box or from the Compression Policy Manager dialog box. This topic includes instructions to bind compression policies to a load balancing virtual server by using the Configure Virtual Server (Load Balancing) dialog box. For information about how you can bind a compression policy to a load balancing virtual server by using the Compression Policy Manager dialog box, see ["Configuring and Binding Policies with the Policy Manager."](#)

To bind or unbind a compression policy to a virtual server by using the command line

At the command prompt, type the following commands to bind or unbind a compression policy to a load balancing virtual server and verify the configuration:

- o (bind|unbind) lb vsrver <name> -policyName <string>
- o show lb vsrver <name>

Example

```
> bind lb vsrver lbvip -policyName ns_cmp_msapp
Done
> show lb vsrver lbvip
lbvip (8.7.6.6:80) - HTTP          Type: ADDRESS
State: UP
Last state change was at Thu May 28 05:37:21 2009 (+685 ms)
Time since last state change: 19 days, 04:26:50.470
Effective State: UP
Client Idle Timeout: 180 sec
Down state flush: ENABLED
Disable Primary Vserver On Down : DISABLED
Port Rewrite : DISABLED
No. of Bound Services : 1 (Total)      1 (Active)
Configured Method: LEASTCONNECTION
Current Method: Round Robin, Reason: Bound service's state changed to UP
Mode: IP
Persistence: NONE
Vserver IP and Port insertion: OFF
Push: DISABLED  Push VServer:
Push Multi Clients: NO
Push Label Rule:

Bound Service Groups:
1)      Group Name: Service-Group-1

1) Service-Group-1 (10.102.29.252: 80) - HTTP State: UP Weight:

1)      Policy : ns_cmp_msapp Priority:0
Done
```

To bind or unbind a compression policy to a load balancing virtual server by using the configuration utility

1. Navigate to Traffic Management > Load Balancing > Virtual Servers.
2. In the details pane, select the virtual server to which you want to bind or unbind a compression policy (for example, Vserver-LB-1), and then click Open.
3. In the Configure Virtual Server (Load Balancing) dialog box, on the Policies tab, click Compression.
4. Do one of the following:
 - o To bind a compression policy, click Insert Policy, and then select the policy you want to bind to the virtual server.
 - o To unbind a compression policy, click the name of the policy you want to unbind from the virtual server, and then click Unbind Policy.
5. Click OK.

Securing Load Balanced Traffic by Using SSL

The Citrix NetScaler SSL offload feature transparently improves the performance of web sites that conduct SSL transactions. By offloading CPU-intensive SSL encryption and decryption tasks from the local web server to the appliance, SSL offloading ensures secure delivery of web applications without the performance penalty incurred when the server processes the SSL data. Once the SSL traffic is decrypted, it can be processed by all standard services. The SSL protocol works seamlessly with various types of HTTP and TCP data and provides a secure channel for transactions using such data.

To configure SSL, you must first enable it. Then, you configure HTTP or TCP services and an SSL virtual server on the appliance, and bind the services to the virtual server. You must also add a certificate-key pair and bind it to the SSL virtual server. If you use Outlook Web Access servers, you must create an action to enable SSL support and a policy to apply the action. An SSL virtual server intercepts incoming encrypted traffic and decrypts it by using a negotiated algorithm. The SSL virtual server then forwards the decrypted data to the other entities on the appliance for appropriate processing.

This document includes the following:

- [SSL Configuration Task Sequence](#)
- [Enabling SSL Offload](#)
- [Creating HTTP Services](#)
- [Adding an SSL-Based Virtual Server](#)
- [Binding Services to the SSL Virtual Server](#)
- [Adding a Certificate Key Pair](#)
- [Binding an SSL Certificate Key Pair to the Virtual Server](#)
- [Configuring Support for Outlook Web Access](#)

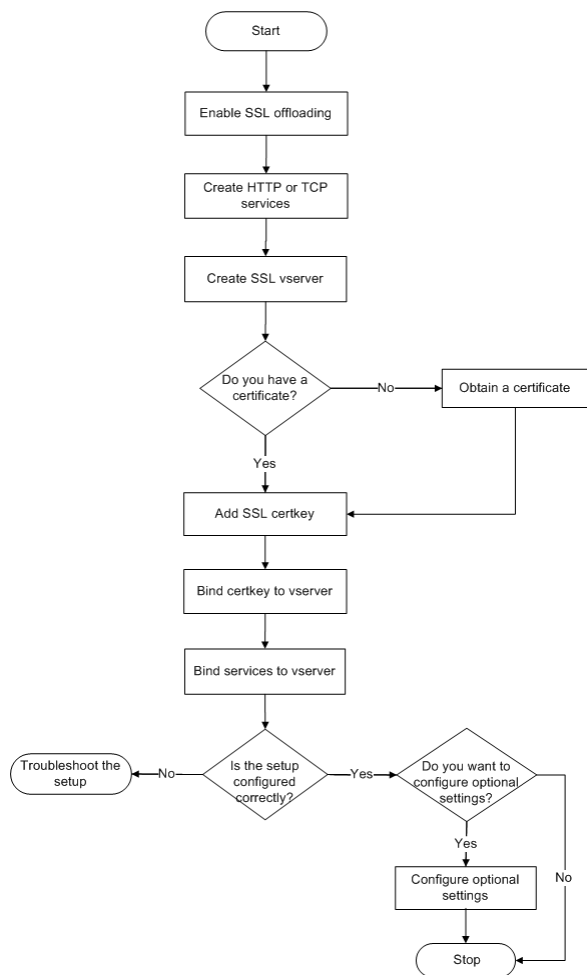
SSL Configuration Task Sequence

To configure SSL, you must first enable it. Then, you must create an SSL virtual server and HTTP or TCP services on the NetScaler. Finally, you must bind a valid SSL certificate and the configured services to the SSL virtual server.

An SSL virtual server intercepts incoming encrypted traffic and decrypts it using a negotiated algorithm. The SSL virtual server then forwards the decrypted data to the other entities on the NetScaler for appropriate processing.

The following flow chart shows the sequence of tasks for configuring a basic SSL offload setup.

Figure 1. Sequence of Tasks to Configure SSL Offloading



Enabling SSL Offload

Updated: 2013-06-05

You should enable the SSL feature before configuring SSL offload. You can configure SSL-based entities on the appliance without enabling the SSL feature, but they will not work until you enable SSL.

To enable SSL by using the command line interface

At the command prompt, type the following commands to enable SSL Offload and verify the configuration:

- o `enable ns feature SSL`
- o `show ns feature`

Example

```

> enable ns feature ssl
Done
> show ns feature
Feature Acronym Status
-----
1) Web Logging WL ON
2) SurgeProtection SP OFF
3) Load Balancing LB ON . . .
  9) SSL Offloading SSL ON
10) Global Server Load Balancing GSLB ON . .
Done >
  
```

To enable SSL by using the configuration utility

1. In the navigation pane, expand System, and then click Settings.
2. In the details pane, under Modes and Features, click Change basic features.
3. Select the SSL Offloading check box, and then click OK.

4. In the Enable/Disable Feature(s)? message box, click Yes.

Creating HTTP Services

Updated: 2013-08-23

A service on the appliance represents an application on a server. Once configured, services are in the disabled state until the appliance can reach the server on the network and monitor its status. This topic covers the steps to create an HTTP service.

Note: For TCP traffic, perform the procedures in this and the following topics, but create TCP services instead of HTTP services.

To add an HTTP service by using the command line interface

At the command prompt, type the following commands to add a HTTP service and verify the configuration:

- `add service <name> (<IP> | <serverName>) <serviceType> <port>`
- `show service <name>`

```
> add service SVC_HTTP1 10.102.29.18 HTTP 80
Done
> show service SVC_HTTP1
SVC_HTTP1 (10.102.29.18:80) - HTTP
State: UP
Last state change was at Wed Jul 15 06:13:05 2009
Time since last state change: 0 days, 00:00:15.350
Server Name: 10.102.29.18
Server ID : 0    Monitor Threshold : 0
Max Conn: 0     Max Req: 0         Max Bandwidth: 0 kbits
Use Source IP: NO
Client Keepalive(CKA): NO
Access Down Service: NO
TCP Buffering(TCPB): NO
HTTP Compression(CMP): YES
Idle timeout: Client: 180 sec    Server: 360 sec
Client IP: DISABLED
Cacheable: NO
SC: OFF
SP: OFF
Down state flush: ENABLED
```

```
1)    Monitor Name: tcp-default
      State: UP           Weight: 1
      Probes: 4           Failed [Total: 0 Current: 0]
      Last response: Success - TCP syn+ack received.
      Response Time: N/A
```

Done

To add an HTTP service by using the configuration utility

1. Navigate to Traffic Management > SSL Offload > Services.
2. In details pane, click Add.
3. In the Create Service dialog box, in the Service Name, Server, and Port text boxes, type the name of the service, IP address, and port (for example, SVC_HTTP1, 10.102.29.18, and 80).
4. In the Protocol list, select the type of the service (for example, HTTP).
5. Click Create, and then click Close. The HTTP service you configured appears in the Services page.
6. Verify that the parameters you configured are correctly configured by selecting the service and viewing the Details section at the bottom of the pane.

Adding an SSL-Based Virtual Server

Updated: 2013-06-05

In a basic SSL offloading setup, the SSL virtual server intercepts encrypted traffic, decrypts it, and sends the clear text messages to the services that are bound to the virtual server. Offloading CPU-intensive SSL processing to the appliance allows the back-end servers to process a greater number of requests.

To add an SSL-based virtual server by using the command line interface

At the command prompt, type the following commands to create an SSL-based virtual server and verify the configuration:

- `add lb vserver <name> <serviceType> [<IPAddress> <port>]`
- `show lb vserver <name>`

Example

```
> add lb vserver vserver-SSL-1 SSL 10.102.29.50 443
Done
> show lb vserver vserver-SSL-1
vserver-SSL-1 (10.102.29.50:443) - SSL Type: ADDRESS
State: DOWN[Certkey not bound] Last state change was at Tue Jun 16 06:33:08 20
Time since last state change: 0 days, 00:03:44.120
Effective State: DOWN Client Idle Timeout: 180 sec
Down state flush: ENABLED
Disable Primary Vserver On Down : DISABLED
No. of Bound Services : 0 (Total) 0 (Active)
Configured Method: LEASTCONNECTION Mode: IP
Persistence: NONE
Vserver IP and Port insertion: OFF
Push: DISABLED Push VServer: Push Multi Clients: NO Push Label Rule: Done
```

Caution: To ensure secure connections, you must bind a valid SSL certificate to the SSL-based virtual server before you enable it.

To add an SSL-based virtual server by using the configuration utility

1. Navigate to Traffic Management > SSL Offload > Virtual Servers.
2. In the details pane, click Add.
3. In the Create Virtual Server (SSL Offload) dialog box, in the Name, IP Address, and Port text boxes, type the name of the virtual server, IP address, and port (for example, `Vserver-SSL-1`, `10.102.29.50`, and `443`).
4. In the Protocol list, select the type of the virtual server, for example, `SSL`.
5. Click Create, and then click Close.
6. Verify that the parameters you configured are correctly configured by selecting the virtual server and viewing the Details section at the bottom of the pane. The virtual server is marked as `DOWN` because a certificate-key pair and services have not been bound to it.

Caution: To ensure secure connections, you must bind a valid SSL certificate to the SSL-based virtual server before you enable it.

Binding Services to the SSL Virtual Server

Updated: 2013-08-23

After decrypting the incoming data, the SSL virtual server forwards the data to the services that you have bound to the virtual server.

Data transfer between the appliance and the servers can be encrypted or in clear text. If the data transfer between the appliance and the servers is encrypted, the entire transaction is secure from end to end. For more information about configuring the system for end-to-end security, see "[SSL Offload and Acceleration](#)."

To bind a service to a virtual server by using the command line interface

At the command prompt, type the following commands to bind service to the SSL virtual server and verify the configuration:

- `bind lb vserver <name> <serviceName>`
- `show lb vserver <name>`

Example

```
> bind lb vserver vserver-SSL-1 SVC_HTTP1
Done
> show lb vserver vserver-SSL-1 vserver-SSL-1 (10.102.29.50:443) - SSL Type:
ADDRESS State: DOWN[Certkey not bound]
Last state change was at Tue Jun 16 06:33:08 2009 (+174 ms)
Time since last state change: 0 days, 00:31:53.70
Effective State: DOWN Client Idle
Timeout: 180 sec
Down state flush: ENABLED Disable Primary Vserver On Down :
DISABLED No. of Bound Services : 1 (Total) 0 (Active)
Configured Method: LEASTCONNECTION Mode: IP Persistence: NONE Vserver IP and
Port insertion: OFF Push: DISABLED Push VServer: Push Multi Clients: NO Push I

1) SVC_HTTP1 (10.102.29.18: 80) - HTTP
State: DOWN Weight: 1
Done
```

To bind a service to a virtual server by using the configuration utility

1. Navigate to Traffic Management > SSL Offload > Virtual Servers.
2. In the details pane, select a virtual server, and then click Open.
3. On the Services tab, in the Active column, select the check boxes next to the services that you want to bind to the selected virtual server.
4. Click OK.
5. Verify that the Number of Bound Services counter in the Details section at the bottom of the pane is incremented by the number of services that you bound to the virtual server.

Adding a Certificate Key Pair

Updated: 2013-06-24

An SSL certificate is an integral element of the SSL Key-Exchange and encryption/decryption process. The certificate is used during SSL handshake to establish the identity of the SSL server. You can use a valid, existing SSL certificate that you have on the NetScaler appliance, or you can create your own SSL certificate. The appliance supports RSA/DSA certificates of up to 4096 bits.

Note: Citrix recommends that you use a valid SSL certificate that has been issued by a trusted certificate authority. Invalid certificates and self-created certificates are not compatible with all SSL clients.

Before a certificate can be used for SSL processing, you must pair it with its corresponding key. The certificate key pair is then bound to the virtual server and used for SSL processing.

To add a certificate key pair by using the command line interface

At the command prompt, type the following commands to create a certificate key pair and verify the configuration:

- o add ssl certKey <certkeyName> -cert <string> [-key <string>]
- o show sslcertkey <name>

Example

```
> add ssl certKey CertKey-SSL-1 -cert ns-root.cert -key ns-root.key
Done
> show sslcertkey CertKey-SSL-1
Name: CertKey-SSL-1 Status: Valid,
Days to expiration:4811 Version: 3
Serial Number: 00 Signature Algorithm: md5WithRSAEncryption Issuer: C=US,ST=California,OU=Citrix ANG,OU=NS Internal,CN=default
Validity Not Before: Oct 6 06:52:07 2006 GMT Not After : Aug 17 21:26:47 2022
Subject: C=US,ST=California,L=San Jose,O=Citrix ANG,OU=NS Internal,CN=default
Algorithm: rsaEncryption Public Key
size: 1024
Done
```

To add a certificate key pair by using the configuration utility

1. Navigate to Traffic Management > SSL > Certificates.
2. In the details pane, click Add.
3. In the Install Certificate dialog box, in the Certificate-Key Pair Name text box, type a name for the certificate key pair you want to add, for example, Certkey-SSL-1.
4. Under Details, in Certificate File Name, click Browse (Appliance) to locate the certificate. Both the certificate and the key are stored in the /nsconfig/ssl/ folder on the appliance. To use a certificate present on the local system, select Local.
5. Select the certificate you want to use, and then click Select.
6. In Private Key File Name, click Browse (Appliance) to locate the private key file. To use a private key present on the local system, select Local.
7. Select the key you want to use and click Select. To encrypt the key used in the certificate key pair, type the password to be used for encryption in the Password text box.
8. Click Install.
9. Double-click the certificate key pair and, in the Certificate Details window, verify that the parameters have been configured correctly and saved.

Binding an SSL Certificate Key Pair to the Virtual Server

Updated: 2013-06-24

After you have paired an SSL certificate with its corresponding key, you must bind the certificate key pair to the SSL virtual server so that it can be used for SSL processing. Secure sessions require establishing a connection between the client computer and an SSL-based virtual server on the appliance. SSL processing is then carried out on the incoming traffic at the virtual server. Therefore, before enabling the SSL virtual server on the appliance, you need to bind a valid SSL certificate to the SSL virtual server.

To bind an SSL certificate key pair to a virtual server by using the command line interface

At the command prompt, type the following commands to bind an SSL certificate key pair to a virtual server and verify the configuration:

- o bind ssl vserver <vServerName> -certkeyName <string>
- o show ssl vserver <name>

Example

```
> bind ssl vserver Vserver-SSL-1 -certkeyName CertKey-SSL-1
Done
> show ssl vserver Vserver-SSL-1
```

```
Advanced SSL configuration for VServer Vserver-SSL-1:
DH: DISABLED
Ephemeral RSA: ENABLED Refresh Count: 0
Session Reuse: ENABLED Timeout: 120 seconds
Cipher Redirect: ENABLED
SSLv2 Redirect: ENABLED
ClearText Port: 0
Client Auth: DISABLED
SSL Redirect: DISABLED
Non FIPS Ciphers: DISABLED
SSLv2: DISABLED SSLv3: ENABLED TLSv1: ENABLED
```

```
1) CertKey Name: CertKey-SSL-1 Server Certificate
1) Cipher Name: DEFAULT
   Description: Predefined Cipher Alias
Done
```

To bind an SSL certificate key pair to a virtual server by using the configuration utility

1. Navigate to Traffic Management > SSL Offload > Virtual Servers.
2. Select the virtual server to which you want to bind the certificate key pair, for example, Vserver-SSL-1, and click Open.
3. In the Configure Virtual Server (SSL Offload) dialog box, on the SSL Settings tab, under Available, select the certificate key pair that you want to bind to the virtual server (for example, Certkey-SSL-1), and then click Add.
4. Click OK.
5. Verify that the certificate key pair that you selected appears in the Configured area.

Configuring Support for Outlook Web Access

If you use Outlook Web Access (OWA) servers on your NetScaler appliance, you must configure the appliance to insert a special header field, FRONT-END-HTTPS: ON, in HTTP requests directed to the OWA servers, so that the servers generate URL links as https:// instead of http://.

Note: You can enable OWA support for HTTP-based SSL virtual servers and services only. You cannot apply it for TCP-based SSL virtual servers and services.

To configure OWA support, do the following:

- Create an SSL action to enable OWA support.
- Create an SSL policy.
- Bind the policy to the SSL virtual server.

Creating an SSL Action to Enable OWA Support

Updated: 2013-06-24

Before you can enable Outlook Web Access (OWA) support, you must create an SSL action. SSL actions are bound to SSL policies and triggered when incoming data matches the rule specified by the policy.

To create an SSL action to enable OWA support by using the command line interface

At the command prompt, type the following commands to create an SSL action to enable OWA support and verify the configuration:

- add ssl action <name> -OWASupport ENABLED
 - show SSL action <name>
- ```
> add ssl action Action-SSL-OWA -OWASupport enabled
Done
> show SSL action Action-SSL-OWA
 Name: Action-SSL-OWA
 Data Insertion Action: OWA
 Support: ENABLED
Done
```

#### To create an SSL action to enable OWA support by using the configuration utility

1. Navigate to Traffic Management > SSL > Policies.
2. In the details pane, on the Actions tab, click Add.
3. In the Create SSL Action dialog box, in the Name text box, type Action-SSL-OWA.
4. Under Outlook Web Access, select Enabled.
5. Click Create, and then click Close.
6. Verify that Action-SSL-OWA appears in the **SSL Actions** page.

### Creating SSL Policies

Updated: 2013-09-04

SSL policies are created by using the policy infrastructure. Each SSL policy has an SSL action bound to it, and the action is carried out when incoming traffic matches the rule that has been configured in the policy.

#### To create an SSL policy by using the command line interface

At the command prompt, type the following commands to configure an SSL policy and verify the configuration:

- add ssl policy <name> -rule <expression> -reqAction <string>
- show ssl policy <name>

#### Example

```
> add ssl policy Policy-SSL-1 -rule ns_true -reqaction Action-SSL-OWA
Done
> show ssl policy Policy-SSL-1
Name: Policy-SSL-1 Rule: ns_true
```

```
Action: Action-SSL-OWA Hits: 0
Policy is bound to following entities
1) PRIORITY : 0
Done
```

### To create an SSL policy by using the configuration utility

1. Navigate to Traffic Management > SSL > Policies.
2. In the details pane, click Add.
3. In the Create SSL Policy dialog box, in the Name text box, type the name of the SSL Policy (for example, Policy-SSL-1).
4. In Request Action, select the configured SSL action that you want to associate with this policy (for example, Action-SSL-OWA). The `ns_true` general expression applies the policy to all successful SSL handshake traffic. However, if you need to filter specific responses, you can create policies with a higher level of detail. For more information about configuring granular policy expressions, see "[Understanding Policies and Expressions](#)."
5. In Named Expressions, choose the built-in general expression `ns_true` and click Add Expression. The expression `ns_true` now appears in the Expression text box.
6. Click Create, and then click Close.
7. Verify that the policy is correctly configured by selecting the policy and viewing the Details section at the bottom of the pane.

### Binding the SSL Policy to an SSL Virtual Server

Updated: 2013-06-24

After you configure an SSL policy for Outlook Web Access, bind the policy to a virtual server that will intercept incoming Outlook traffic. If the incoming data matches any of the rules configured in the SSL policy, the policy is triggered and the action associated with it is carried out.

### To bind an SSL policy to an SSL virtual server by using the command line interface

At the command prompt, type the following commands to bind an SSL policy to an SSL virtual server and verify the configuration:

- o `bind ssl vserver <vServerName> -policyName <string>`
- o `show ssl vserver <name>`

#### Example

```
> bind ssl vserver Vserver-SSL-1 -policyName Policy-SSL-1
Done
> show ssl vserver Vserver-SSL-1
Advanced SSL configuration for VServer Vserver-SSL-1:
DH: DISABLED
Ephemeral RSA: ENABLED Refresh Count: 0
Session Reuse: ENABLED Timeout: 120 seconds
Cipher Redirect: ENABLED
SSLv2 Redirect: ENABLED
ClearText Port: 0
Client Auth: DISABLED
SSL Redirect: DISABLED
Non FIPS Ciphers: DISABLED
SSLv2: DISABLED SSLv3: ENABLED TLSv1: ENABLED

1) CertKey Name: CertKey-SSL-1 Server Certificate

1) Policy Name: Policy-SSL-1
 Priority: 0

1) Cipher Name: DEFAULT
 Description: Predefined Cipher Alias

Done
>
```

### To bind an SSL policy to an SSL virtual server by using the configuration utility

1. Navigate to Traffic Management > SSL Offload > Virtual Servers.
2. In the details pane, select the virtual server (for example, Vserver-SSL-1), and then click Open.

3. In the Configure Virtual Server (SSL Offload) dialog box, click Insert Policy, and then select the policy that you want to bind to the SSL virtual server. Optionally, you can double-click the Priority field and type a new priority level.
4. Click OK.

## Features at a Glance

Citrix NetScaler features can be configured independently or in combinations to address specific needs. Although some features fit more than one category, the numerous NetScaler features can generally be categorized as application switching and traffic management features, application acceleration features, and application security and firewall features, and an application visibility feature.

To understand the order in which the features perform their processing, see ["Processing Order of Features."](#)

This document includes the following:

- [Application Switching and Traffic Management Features](#)
- [Application Acceleration Features](#)
- [Application Security and Firewall Features](#)
- [Application Visibility Feature](#)
- [Cloud Integration Feature](#)

# Application Switching and Traffic Management Features

## SSL Offloading

Transparently offloads SSL encryption and decryption from web servers, freeing server resources to service content requests. SSL places a heavy burden on an application's performance and can render many optimization measures ineffective. SSL offload and acceleration allow all the benefits of Citrix Request Switching technology to be applied to SSL traffic, ensuring secure delivery of web applications without degrading end-user performance.

For more information, see ["SSL Offload and Acceleration."](#)

## Access Control Lists

Compares incoming packets to Access Control Lists (ACLs). If a packet matches an ACL rule, the action specified in the rule is applied to the packet. Otherwise, the default action (ALLOW) is applied and the packet is processed normally. For the appliance to compare incoming packets to the ACLs, you have to apply the ACLs. All ACLs are enabled by default, but you have to apply them in order for the NetScaler to compare incoming packets against them. If an ACL is not required to be a part of the lookup table, but still needs to be retained in the configuration, it should be disabled before the ACLs are applied. A NetScaler does not compare incoming packets to disabled ACLs.

For more information, see ["Access Control List."](#)

## Load Balancing

Load balancing decisions are based on a variety of algorithms, including round robin, least connections, weighted least bandwidth, weighted least packets, minimum response time, and hashing based on URL, domain source IP, or destination IP. Both the TCP and UDP protocols are supported, so the NetScaler can load balance all traffic that uses those protocols as the underlying carrier (for example, HTTP, HTTPS, UDP, DNS, NNTP, and general firewall traffic). In addition, the NetScaler can maintain session persistence based on source IP, cookie, server, group, or SSL session. It allows users to apply custom Extended Content Verification (ECV) to servers, caches, firewalls and other infrastructure devices to ensure that these systems are functioning properly and are providing the right content to users. It can also perform health checks using ping, TCP, or HTTP URL, and the user can create monitors based on Perl scripts. To provide high-scale WAN optimization, the CloudBridge appliances deployed at data centers can be load balanced through NetScaler appliances. The bandwidth and number of concurrent sessions can be improved significantly.

For more information, see ["Load Balancing."](#)

## Traffic Domains

Traffic domains provide a way to create logical ADC partitions within a single NetScaler appliance. They enable you to segment network traffic for different applications. You can use traffic domains to create multiple isolated environments whose resources do not interact with each other. An application belonging to a specific traffic domain communicates only with entities, and processes traffic, within that domain. Traffic belonging to one traffic domain cannot cross the boundary of another traffic domain. Therefore, you can use duplicate IP addresses on the appliance as long as an addresses is not duplicated within the same domain.

For more information, see ["Traffic Domains."](#)

## 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 network's source IP addresses when data passes through the NetScaler.

The NetScaler appliance supports the following types of network address translation:

**INAT**—In Inbound NAT (INAT), an IP address (usually public) configured on the NetScaler appliance listens to connection requests on behalf of a server. For a request packet received by the appliance on a public IP address, the NetScaler replaces the destination IP address with the private IP address of the server. In other words, the appliance acts as a proxy between clients and the server. INAT configuration involves INAT rules, which define a 1:1 relationship between the IP address on the NetScaler appliance and the IP address of the server.

**RNAT**—In Reverse Network Address Translation (RNAT), for a session initiated by a server, the NetScaler appliance replaces the source IP address in the packets generated by the server with an IP address (type SNIP) configured on the appliance. The appliance thereby prevents exposure of the server's IP address in any of the packets generated by the server. An RNAT configuration involves an RNAT rule, which specifies a condition. The appliance performs RNAT processing on those packets that match the condition.

**Stateless NAT46 Translation**—Stateless NAT46 enables communication between IPv4 and IPv6 networks, by way of IPv4 to IPv6 packet translation and vice versa, without maintaining any session information on the NetScaler appliance. A stateless NAT46 configuration involves an IPv4-IPv6 INAT rule and an NAT46 IPv6 prefix.



**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 involves an NAT64 rule and an NAT64 IPv6 prefix.

For more information, see ["Configuring Network Address Translation."](#)

#### Multipath TCP Support

NetScaler appliances support Multipath TCP (MPTCP). MPTCP is a TCP/IP protocol extension that identifies and uses multiple paths available between hosts to maintain the TCP session. You must enable MPTCP on a TCP profile and bind it to a virtual server. When MPTCP is enabled, the virtual server functions as an MPTCP gateway and converts MPTCP connections with the clients to TCP connections that it maintains with the servers.

For more information, see ["MPTCP \(Multi-Path TCP\)."](#)

#### Content Switching

Determines the server to which to send the request on the basis of configured content switching policies. Policy rules can be based on the IP address, URL, and HTTP headers. This allows switching decisions to be based on user and device characteristics such as who the user is, what type of agent is being used, and what content the user requested.

For more information, see ["Content Switching."](#)

#### Global Server Load Balancing (GSLB)

Extends the traffic management capabilities of a NetScaler to include distributed Internet sites and global enterprises. Whether installations are spread across multiple network locations or multiple clusters in a single location, the NetScaler maintains availability and distributes traffic across them. It makes intelligent DNS decisions to prevent users from being sent to a site that is down or overloaded. When the proximity-based GSLB method is enabled, the NetScaler can make load balancing decisions based on the proximity of the client's local DNS server (LDNS) in relation to different sites. The main benefit of the proximity-based GSLB method is faster response time resulting from the selection of the closest available site.

For more information, see ["Global Server Load Balancing."](#)

#### Dynamic Routing

Enables routers to obtain topology information, routes, and IP addresses from neighboring routers automatically. When dynamic routing is enabled, the corresponding routing process listens to route updates and advertises routes. The routing processes can also be placed in passive mode. Routing protocols enable an upstream router to load balance traffic to identical virtual servers hosted on two standalone NetScaler units using the Equal Cost Multipath technique.

For more information, see ["Configuring Dynamic Routes."](#)

#### Link Load Balancing

Load balances multiple WAN links and provides link failover, further optimizing network performance and ensuring business continuity. Ensures that network connections remain highly available, by applying intelligent traffic control and health checks to distribute traffic efficiently across upstream routers. Identifies the best WAN link to route both incoming and outbound traffic based on policies and network conditions, and protects applications against WAN or Internet link failure by providing rapid fault detection and failover.

For more information, see ["Link Load Balancing."](#)

#### TCP Optimization

You can use TCP profiles to optimize TCP traffic. TCP profiles define the way that NetScaler virtual servers process TCP traffic. Administrators can use the built-in TCP profiles or configure custom profiles. After defining a TCP profile, you can bind it to a single virtual server or to multiple virtual servers.

Some of the key optimization features that can be enabled by TCP profiles are:

- **TCP keep-alive**—Checks the operational status of the peers at specified time intervals to prevent the link from being broken.
- **Selective Acknowledgment (SACK)**—Improves the performance of data transmission, especially in long fat networks (LFNs).
- **TCP window scaling**—Allows efficient transfer of data over long fat networks (LFNs).

For more information on TCP Profiles, see ["Configuring TCP Profiles."](#)

#### Web Interface on NetScaler

Provides access to XenApp and XenDesktop resources, which include applications, content, and desktops. 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 6.0.26 on the NetScaler appliance.

Note: Web Interface is supported only on NetScaler nCore releases.

For more information, see ["Web Interface."](#)

#### CloudBridge Connector

The Citrix NetScaler CloudBridge Connector feature, a fundamental part of the Citrix OpenCloud framework, is a tool used to build a cloud-extended data center. The OpenCloud Bridge enables you to connect one or more NetScaler



appliances or NetScaler virtual appliances on the cloud-to your network without reconfiguring your network. Cloud hosted applications appear as though they are running on one contiguous enterprise network. The primary purpose of the OpenCloud Bridge is to enable companies to move their applications to the cloud while reducing costs and the risk of application failure. In addition, the OpenCloud Bridge increases network security in cloud environments. An OpenCloud Bridge is a Layer-2 network bridge that connects a NetScaler appliance or NetScaler virtual appliance on a cloud instance to a NetScaler appliance or NetScaler virtual appliance on your LAN. The connection is made through a tunnel that uses the Generic Routing Encapsulation (GRE) protocol. The GRE protocol provides a mechanism for encapsulating packets from a wide variety of network protocols to be forwarded over another protocol. Then Internet Protocol security (IPsec) protocol suite is used to secure the communication between the peers in the OpenCloud Bridge.

For more information, see "[CloudBridge](#)."

## DataStream

The NetScaler DataStream feature provides an intelligent mechanism for request switching at the database layer by distributing requests on the basis of the SQL query being sent.

When deployed in front of database servers, a NetScaler ensures optimal distribution of traffic from the application servers and Web servers. Administrators can segment traffic according to information in the SQL query and on the basis of database names, user names, character sets, and packet size.

You can configure load balancing to switch requests according to load balancing algorithms, or you can elaborate the switching criteria by configuring content switching to make a decision based on SQL query parameters, such as user name, database names, and command parameters. You can further configure monitors to track the states of database servers.

The advanced policy infrastructure on the NetScaler appliance includes expressions that you can use to evaluate and process the requests. The advanced expressions evaluate traffic associated with MySQL database servers. You can use request-based expressions (expressions that begin with `MYSQL.CLIENT` and `MYSQL.REQ`) in advanced policies to make request switching decisions at the content switching virtual server bind point and response-based expressions (expressions that begin with `MYSQL.RES`) to evaluate server responses to user-configured health monitors.

Note: DataStream is supported for MySQL and MS SQL databases.

For more information, see "[DataStream](#)."

## Application Acceleration Features

### AppCompress

Uses the gzip compression protocol to provide transparent compression for HTML and text files. The typical 4:1 compression ratio yields up to 50% reduction in bandwidth requirements out of the data center. It also results in significantly improved end-user response time, because it reduces the amount of data that must be delivered to the user's browser.

For more information, see "[Compression](#)."

### Cache Redirection

Manages the flow of traffic to a reverse proxy, transparent proxy, or forward proxy cache farm. Inspects all requests, and identifies non-cacheable requests and sends them directly to the origin servers over persistent connections. By intelligently redirecting non-cacheable requests back to the origin web servers, the NetScaler appliance frees cache resources and increases cache hit rates while reducing overall bandwidth consumption and response delays for these requests.

For more information, see "[Cache Redirection](#)."

### AppCache

Helps optimize web content and application data delivery by providing a fast in-memory HTTP/1.1 and HTTP/1.0 compliant web caching for both static and dynamic content. This on-board cache stores the results of incoming application requests even when an incoming request is secured or the data compressed, and then reuses the data to fulfill subsequent requests for the same information. By serving data directly from the on-board cache, the appliance can reduce page regeneration times by eliminating the need to funnel static and dynamic content requests to the server.

For more information, see "[Integrated Caching](#)."

### TCP Buffering

Buffers the server's response and delivers it to the client at the client's speed, thus offloading the server faster and thereby improving the performance of web sites.

For more information, see "[TCP Buffering](#)."

# Application Security and Firewall Features

## Denial of Service Attack (DoS) Defense

Detects and stops malicious distributed denial-of-service (DDoS) attacks and other types of malicious attacks before they reach your servers, preventing them from affecting network and application performance. The NetScaler appliance identifies legitimate clients and elevates their priority, leaving suspect clients unable to consume a disproportionate percentage of resources and cripple your site. The appliance provides application-level protection from the following types of malicious attacks:

- SYN flood attacks
- Pipeline attacks
- Teardrop attacks
- Land attacks
- Fraggle attacks
- Zombie connection attacks

The appliance aggressively defends against these types of attacks by preventing the allocation of server resources for these connections. This insulates servers from the overwhelming flood of packets associated with these events.

The appliance also protects network resources from ICMP based attacks by using ICMP rate limiting and aggressive ICMP packet inspection. It performs strong IP reassembly, drops a variety of suspicious and malformed packets, and applies Access Control Lists (ACLs) to site traffic for further protection.

For more information, see "[HTTP Denial-of-Service Protection](#)."

## Content Filtering

Provides protection from malicious attacks for web sites at the Layer 7 level. The appliance inspects each incoming request according to user-configured rules based on HTTP headers, and performs the action the user configured. Actions can include resetting the connection, dropping the request, or sending an error message to the user's browser. This allows the appliance to screen unwanted requests and reduces your servers' exposure to attacks.

This feature can also analyze HTTP GET and POST requests and filter out known bad signatures, allowing it to defend your servers against HTTP-based attacks.

For more information, see "[Content Filtering](#)."

## Responder

Functions like an advanced filter and can be used to generate responses from the appliance to the client. Some common uses of this feature are generation of redirect responses, user defined responses, and resets.

For more information, see "[Responder](#)."

## Rewrite

Modifies HTTP headers and body text. You can use the rewrite feature to add HTTP headers to an HTTP request or response, make modifications to individual HTTP headers, or delete HTTP headers. It also enables you to modify the HTTP body in requests and responses.

When the appliance receives a request or sends a response, it checks for rewrite rules, and if applicable rules exist, it applies them to the request or response before passing it on to the web server or client computer.

For more information, see "[Rewrite](#)."

## Priority Queuing

Prioritizes user requests to ensure that the most important traffic is serviced first during surges in request volume. You can establish priority based on request URLs, cookies, or a variety of other factors. The appliance places requests in a three-tier queue based on their configured priority, enabling business-critical transactions to flow smoothly even during surges or site attacks.

For more information, see "[Priority Queuing](#)."

## Surge Protection

Regulates the flow of user requests to servers and controls the number of users that can simultaneously access the resources on the servers, queuing any additional requests once your servers have reached their capacity. By controlling the rate at which connections can be established, the appliance blocks surges in requests from being passed on to your servers, thus preventing site overload.

For more information, see "[Surge Protection](#)."

## NetScaler Gateway

NetScaler Gateway is a secure application access solution that provides administrators granular application-level policy and action controls to secure access to applications and data while allowing users to work from anywhere. It gives IT administrators a single point of control and tools to help ensure compliance with regulations and the highest levels of information security across and outside the enterprise. At the same time, it empowers users with a single point of accessâ€”optimized for roles, devices, and networksâ€”to the enterprise applications and data they need. This unique combination of capabilities helps maximize the productivity of today's mobile workforce.

For more information, see "[NetScaler Gateway](#)."

#### Application Firewall

Protects applications from misuse by hackers and malware, such as cross site scripting attacks, buffer overflow attacks, SQL injection attacks, and forceful browsing, by filtering traffic between each protected web server and users that connect to any web site on that web server. The application firewall examines all traffic for evidence of attacks on web server security or misuse of web server resources, and takes the appropriate action to prevent these attacks from succeeding.

For more information, see "[Application Firewall](#)."

## Application Visibility Feature

### NetScaler Insight Center

NetScaler Insight Center is a high performance collector that provides end-to-end user experience visibility across Web and HDX (ICA) traffic. It collects HTTP and ICA AppFlow records generated by NetScaler ADC appliances and populates analytical reports covering Layer 3 to Layer 7 statistics. NetScaler Insight Center provides in-depth analysis for the last five minutes of real-time data, and for historical data collected for the last one hour, one day, one week, and one month.

HDX (ICA) analytic dashboard enables you to drill down from HDX Users, Applications, Desktops, and even from gateway-level information. Similarly, HTTP analytics provide a bird's eye view of Web Applications, URLs Accessed, Client IP Addresses and Server IP Addresses, and other dashboards. The administrator can drill down and identify the pain points from any of these dashboards, as appropriate for the use case.

### EdgeSight for NetScaler

Support for application performance monitoring based on end user experience. This solution leverages the HTML injection feature to obtain various time values, which are used by EdgeSight server for analysis and report generation. EdgeSight for NetScaler provides a way to monitor the performance benefits of a NetScaler and determine potential bottlenecks in a network.

For more information, see "[EdgeSight Monitoring for NetScaler](#)."

### Enhanced Application Visibility Using AppFlow

The Citrix NetScaler appliance is a central point of control for all application traffic in the data center. It collects flow and user-session level information valuable for application performance monitoring, analytics, and business intelligence applications. AppFlow transmits this 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.

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.

To limit the types of flows to monitor, by sampling and filtering the application traffic, 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.

For more information, see "[AppFlow](#)."

### Stream Analytics

The performance of your web site or application depends on how well you optimize the delivery of the most frequently requested content. Techniques such as caching and compression help accelerate the delivery of services to clients, but you need to be able to identify the resources that are requested most frequently, and then cache or compress those resources. You can identify the most frequently used resources by aggregating real-time statistics about website or application traffic. Statistics such as how frequently a resource is accessed relative to other resources and how much bandwidth is consumed by those resources help you determine whether those resources need to be cached or compressed to improve server performance and network utilization. Statistics such as response times and the number of concurrent connections to the application help you determine whether you must enhance server-side resources.

If the web site or application does not change frequently, you can use products that collect statistical data, and then manually analyze the statistics and optimize the delivery of content. However, if you do not want to perform manual optimizations, or if your web site or application is dynamic in nature, you need infrastructure that can not only collect statistical data but can also automatically optimize the delivery of resources on the basis of the statistics. On the NetScaler appliance, this functionality is provided by the Stream Analytics feature. The feature operates on a single NetScaler appliance and collects run-time statistics on the basis of criteria that you define. When used with NetScaler policies, the feature also provides you with the infrastructure that you need for automatic, real-time traffic optimization.

For more information, see "[Stream Analytics](#)."

# Cloud Integration Feature

## AutoScale

All applications have specific usage patterns that comprise peaks and troughs. These load variations can be dynamic in nature and difficult to predict, given that they depend on several factors that are intrinsic to the use case. Cloud users have to constantly monitor the load on their application fleet and make sure that these variations have minimum impact on end users. During periods of peak usage, when the application fleet is overloaded and end users experience significant latency, they have to deploy additional application instances. During trough periods, the expanded fleet is underutilized. So they might have to remove additional instances or bear unnecessary cost overheads. In most cases, they have to perform these tasks manually.

If your organization uses Citrix CloudPlatform to deploy and manage the cloud environment, users can use the *AutoScale* feature in CloudPlatform, in conjunction with a Citrix NetScaler appliance, to automatically scale their applications as needed. The AutoScale feature is part of the elastic load balancing feature in CloudPlatform. A CloudPlatform user can use the AutoScale feature to specify thresholds for various conditions for automatically scaling the application fleet upward and downward. CloudPlatform, in turn, configures the NetScaler appliance (by using the NetScaler NITRO API) 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.

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.

For more information about AutoScale, see ["AutoScale: Automatic Scaling in the Citrix CloudPlatform Environment."](#)

