

Not for resale or distribution

Implementing Citrix NetScaler 11 for Application and Desktop Solutions

Citrix Course CNS-207-31

Not for resale or distribution



Implementing Citrix NetScaler 11 for Application and Desktop Solutions

Citrix Course CNS-207-31
December 2015
Version 3.0

CITRIX[®]
Authorized
Courseware

Not for resale or distribution

Table of Contents

Module 1: Getting Started	21
Getting Started Manual	23
Overview	23
Introduction to Citrix NetScaler	23
NetScaler Functionality	23
NetScaler Overview	24
Product Features	25
Application Availability	25
Application Acceleration	26
Application Security	26
Simple Manageability	27
Front-end and TCP Protocol Optimizations	28
Lower Total Cost of Ownership	28
Discussion Question	28
Hardware Platforms	28
Hardware Components	30
Network and Serial Interfaces	30
File System	30
nCore Configuration Architecture	31
Deployment Scenarios	33
New Technology	33
Displacement	34
SDX	35
NetScaler Gateway Overview	36
NetScaler Gateway Platforms	37
MPX Appliance	37
VPX	37
Initial NetScaler Access	38
NetScaler Licenses	38
NetScaler Gateway Licensing	39
NetScaler Gateway License Types	39
Platform License	40
Universal License	40
NetScaler Gateway Licensing Considerations	40
To Install the NetScaler Gateway License	41
NetScaler Gateway Pre-Installation Checklist	41
Replacing Secure Gateway	42
Configuring NetScaler Gateway For First-Time Use	42
Settings Configuration	43
NetScaler Gateway Wizard	43
NetScaler for XenApp and XenDesktop Wizard	45
End-User Access with the FQDN	46

To View an Imported Certificate FQDN	46
Configuration Testing	46
Name Service Providers Configuration	47
Performing an Upgrade	47
Upgrading a Standalone NetScaler System	47
Save the Configuration	48
Running Configuration	48
Discussion Question	48
Module 2: Basic Networking	49
Basic Networking Manual	51
Overview	51
OSI Networking Model	51
NetScaler Architecture Overview	52
NetScaler IP Addresses	53
NetScaler IP Address	54
Configuring a NetScaler IP Address	54
Mapped IP Address	54
Configuring a Mapped IP Address	55
Subnet IP Address	55
Configuring a Subnet IP Address	56
Virtual IP Address	56
To configure a virtual IP address	56
To remove a NetScaler-owned IP address	57
Network Topology	57
One-Arm Mode	58
Two-Arm Mode	59
NetScaler Gateway Deployment	60
NetScaler Gateway Deployment in the Perimeter Network	60
NetScaler Gateway Deployment in the Secure Network	61
Discussion Question	62
Double-Hop Deployment	62
NetScaler Gateway Deployment with XenMobile, XenApp and XenDesktop	62
NetScaler Gateway Deployment with Web Interface	63
NetScaler Network Interfaces	64
Enabling or Disabling Network Interfaces Using the Configuration Utility	64
Enabling or Disabling Network Interfaces Using the Command-Line Interface	65
Virtual Local Area Networks (VLAN)	65
VLAN Configuration	66
Port-based VLANs	67
Tagged VLANs	67
NSVLAN	68
Configuring VLANs	68
IP Address Routing	69
Static Routes	69
MAC-based Forwarding Mode	69

Determining the Source IP Address	70
Sending a Client IP Address to Servers	70
Link Aggregation	71
Discussion Question	71
Module 3: High Availability	73
High Availability Manual	75
Overview	75
High-Availability Functionality	76
High Availability Process	77
High-Availability Node Configuration	78
Pre-Configuration Checklist	78
High-Availability Failover Process	79
Configuring Primary and Secondary Nodes	79
High-Availability Status	79
Verifying Status on the Appliance	80
Propagation and Synchronization	80
Disabling Command Propagation	81
Automatic Configuration Synchronization	81
Forced Synchronization	81
Performing a Forced Failover	82
Fail-Safe Mode	82
High-Availability Management	83
Upgrading a High-Availability Pair	83
Discussion Question	83
High-Availability Issues	84
Configuration Synchronization Failure	84
File Synchronization Failure	84
Unexpected Failover	85
High-Availability Pair in Different Subnets	86
Adding a Remote Node	86
To Add a Remote Node for Independent Network Computing Mode	86
Discussion Question	87

Module 4: Integrating NS Gateway with Other Resources (Unified Gateway) **89**

Integrating NetScaler with XenApp and XenDesktop Using the Unified Gateway Wizard	91
Overview	91
NetScaler Unified Gateway Key Features and Benefits	91
NetScaler Unified Gateway 11 New Features	92
NetScaler Unified Gateway Prerequisites	92
Firewall Rules	93
Unified Gateway with XenApp, XenDesktop, and XenMobile	94
Unified Gateway with Other Applications	94
Client Connection Methods	95

Unified Gateway Configuration Wizard	95
Unified Gateway Traffic Flow in Double Hop DMZ	96

Module 5: Basic Load Balancing	99
Basic Load Balancing Manual	101
Overview	101
Load-Balancing Basics	102
Entity Management	103
Server Creation	104
Server Configuration Parameters	105
To Create a Server Using the Configuration Utility	106
Service Monitoring	106
Services Configuration Overview	107
Services Creation	107
Service Configuration Parameters	109
To Create a Service Using the Configuration Utility	110
Service Groups	110
To Create a Service Group Using the Configuration Utility	111
Virtual Server Creation	111
Virtual Server Configuration Parameters	111
To Create a Virtual Server Using the Configuration Utility	112
Binding Services or Service Groups to a Virtual Server	112
To Bind a Service to a Load-Balancing Virtual Server Using the Configuration Utility	113
To Bind a Service Group to a Virtual Server Using the Configuration Utility	113
Traffic Types	113
Application Protocols	113
Session Protocols	114
General Traffic	114
Built-in Monitors	115
TCP-based Applications	115
SSL Services	115
FTP Services	115
RADIUS Services	116
DNS and DNS-TCP Services	116
LDAP Services	116
XML Broker Services	116
Monitor Parameters	116
Standard Parameters	117
Creating Monitors	118
HTTP Monitoring	119
Extended Application Verification (EAV) Monitoring	120
Extended Content Verification (ECV) Monitoring	120
HTTP-ECV and TCP-ECV Monitoring Process	120
Reverse Condition Monitoring	121
Setting Monitor Thresholds	121
Custom Monitors	122

XenDesktop Delivery Controller Monitoring	122
StoreFront Store Monitoring	122
TFTP Server Monitoring	123
Load-Balancing Methods	123
Least Connections	124
Round Robin	124
Diameter Load Balancing	124
Service Weights	125
To Configure a Virtual Server to Assign Weights to Services Using the Configuration Utility	125
Persistence and Persistence Connections	125
Session Persistence Methods	126
Cookie-Insert Persistence	126
Cookie Version	127
Source-IP-Address Persistence	127
SSL Session ID Persistence	128
Persistence Tables	128
Persistence Group Configuration	128
To View Persistence Sessions Using the Configuration Utility	129
Clearing of Persistence Sessions	129
To Clear Persistence Sessions Using the Configuration Utility	129
Load Balancing Configuration Protection	130
Disabling Services	130
To Disable a Service Using the Configuration Utility	130
Graceful Shutdown of Services	130
Removing Services	131
To Remove a Service Using the Configuration Utility	131
Configuration Verification	131
Viewing the Properties of a Server Object	131
Viewing the Properties of a Virtual Server	131
Viewing the Properties of a Service	132
Viewing the Bindings of a Service	132
The Load-Balancing Visualizer	132
Discussion Question	133
Module 6: NetScaler Security and SSL	135
SSL Offload Manual	137
Overview	137
SSL	137
SSL Session Process	137
Features and Benefits	139
Offload Performance	140
Digital Certificates	140
SSL Administration	140
SSL and Policies	141
SSL Keys	141

Public Key Encryption	142
Certificate Signing Request	142
SSL Certificates	143
Root CA Certificates	143
Intermediate CA Certificates	143
Server Certificates	143
Client Certificates	144
Certificate Generation	144
Certificate Key Pairs	144
Adding a Certificate Key Pair	145
Binding a Certificate Key Pair to a Virtual Server	146
Intermediate Certificate Not Linked	146
Certificate Revocation List (CRL)	146
Certificate Updates	147
Multiple Hosts on a Single IP Address	147
Discussion Question	148
SSL Offload Overview	148
Configuring SSL Offload	149
Securing Traffic Between a Client and the NetScaler System	150
SSL Virtual Servers	150
Creating an SSL Virtual Server	150
Binding an SSL Virtual Server	151
SSL Termination Points	151
Terminating SSL on the NetScaler System	151
Terminating SSL on the Servers	152
Deployment Scenarios	152
Front-end SSL with Back-end HTTP Requirements	152
Front-end SSL with Back-end SSL Requirements	153
Securing Traffic Between the NetScaler and the Server	154
Front-end SSL_TCP with Back-end TCP Requirements	154
Securing TCP Traffic Between a Client and the NetScaler System	155
SSL Bridge	156
SSL Bridge Requirements	156
Citrix Recommendations for SSL	156
SSL Renegotiation Attack	157
SSL Troubleshooting	157
SSL Command-Line Interface Commands	158
Troubleshooting of Encrypted SSL Connections	158
SSL Offload Troubleshooting	159
Access to SSL VIP Address Failing	160
Discussion Question	160

Module 7: Authentication and Authorization	161
Authentication and Authorization Manual	163
Overview	163
System and AAA Users Groups	163

Local Accounts	165
External Authentication	165
External Authentication for System Users	166
Authentication Actions and Policies	166
Local Authentication Configuration	166
NTLMv2 Authentication	166
Configuration of Command Policies for Delegated Administrators	167
Custom Command Policy Configuration for Delegated Administrators	167
Authentication Configuration	168
Authentication Types Supported on NetScaler	168
Default Global Authentication Types Configuration	169
Local Users Configuration	170
Discussion Question	170
Authentication Policies	170
Authentication Order	171
Authentication Policy Bindings	171
LDAP Authentication Configuration	171
Determining LDAP Attributes	171
RADIUS Authentication Configuration	172
Troubleshooting Authentication	172
Client Certificate Authentication Configuration	172
Using Certificates as Two-Factor Authentication	173
To Configure a Client Certificate Authentication Policy	173
Smart Card Authentication Configuration	173
Common Access Cards	174
To Disable Authentication	175
Authorization Configuration	175
Default Global Authorization	175
Authorization Policies Configuration	175
Authentication, Authorization, and Auditing (AAA) Issues	176
Authentication Troubleshooting	176
Discussion Question	176
Module 8: Access Policies	177
Access Policies Manual	179
Overview	179
Endpoint Analysis	179
Endpoint Analysis Plug-in	179
Endpoint Policies	180
Preauthentication Policies	181
Policy Profiles	181
Policy Expressions	182
Policy Bindings	182
Creating a Preauthentication Policy	183
Preauthentication Policies and Profiles	183
Security Preauthentication Expressions Configuration for End-User Devices	183

Operating System Policies Configuration	184
Antivirus, Firewall, Internet Security, or Anti-Spam Expressions Configuration	184
To Configure a Client Service Scan	186
Security Checks Configuration	187
To Configure Process Policies	187
File Scan Policies Configuration	187
To Configure a Session or Preauthentication Policy to Check for a File on the End-User Device	188
Registry Policies Configuration	188
To Configure a Registry Policy	189
To Create a Sample Preauthentication Scan	189
Custom Expressions Configuration	190
To Bind Preauthentication Policies	190
Preauthentication Policy Considerations	191
To Change the Priority of a Preauthentication Policy	192
To Unbind a Preauthentication Policy	192
To Remove a Preauthentication Policy	192
Post-Authentication Policies	193
To Configure a Post-Authentication Policy	193
Post-Authentication Scan Frequency Configuration	193
End-User Logon Options Evaluation	194
Quarantine Groups	195
Endpoint Analysis Troubleshooting	195
Discussion Question	195
Module 9: End User Access and Experience (includes RDP Proxy)	197
End User Access and Experience Exercises (includes RDP Proxy)	199
Overview	199
Connection Methods	199
Secure Tunnel Establishment	200
Network Firewalls and Proxies	201
Secure Tunnel Termination	201
NetScaler Gateway Plug-in Support	202
Software Firewalls	203
NetScaler Gateway Plug-in Integration with Citrix Receiver	203
Citrix Receiver Home Page	203
To Configure the Citrix Receiver Home Page	203
Application of the Citrix Receiver Theme to the Logon Page	204
NetScaler Gateway Plug-in Selection for End Users	204
NetScaler Gateway Plug-in Deployment, Upgrading and Removal from Active Directory	205
NetScaler Gateway Plug-in Installation Troubleshooting Using Active Directory	205
Client Ports	206
NetScaler Gateway Plug-in Connection Configuration	206
Internal Network Resources Connection	206
End-User Connections Proxy Support Enablement	207
Session Policies	208

Session Profiles	209
Policy Bindings	210
Creating Session Policies	211
Credential Passing	211
Credential Passing Configuration	211
Single Sign-on for File Shares	212
Configuring Global Credential Passing	212
Configuring the Default Home Page for SSL VPN and Clientless VPN Connection	212
Timeout Settings	212
Configuring Session and Client Idle Timeout Values	213
Configuring Forced Timeout Values	214
Split Tunneling	214
Configuring Local LAN Access	215
Restricting Local LAN Access to Private Network IP Addresses	215
Configuring Split Tunneling	215
Reverse Split Tunneling	216
Split Tunneling Considerations	216
Timeout Settings Configuration	217
To Configure Forced Timeouts	218
To Configure a Forced Timeout Within a Session Policy	218
Configuring Session or Idle Timeouts	219
To Configure a Session or Client Idle Timeout Globally	219
To Configure Session or Client Idle Timeout Settings Using a Session Policy	219
Client Cleanup	220
Cleanup Options	220
Cleanup Prompt	221
Cleanup Data Sets	221
Cleanup Level	222
Configuring Client Cleanup	223
Single Sign-on Configuration	223
Single Sign-on with Windows Configuration	224
Enabling Single Sign-On	224
Single Sign-on to Web Applications Configuration	225
Single Sign-on to Web Applications Using LDAP Configuration	225
Modifying LDAP Authentication policy	225
Domain Single Sign-on Configuration	226
Client Interception	227
Configuring Client Interception	228
To Configure an IP Address Range	228
Address Pools	228
Configuring Intranet IP Addresses	229
Defining Options for an Address Pool	230
Reclaiming of IP Addresses	231
To Configure Name Service Resolution	231
Access Interface Configuration	232
To Replace Access Interface with a Custom Home Page	232
Clientless Access	233

Clientless Access Enablement	234
Web Address Encoding	235
Clientless Access Policies	235
Clientless Access Policy Creation	236
Domain Access Configuration for End Users	236
SharePoint Site	237
Clientless Access Settings Using Web Interface	238
Client Choices Page Configuration	238
Access Scenario Fallback Configuration	239
Access Scenario Fallback Policy Creation	240
NetScaler Gateway Advanced Concepts	240
Discussion Question	241

Module 10: Integrating NetScaler 11 with XenApp and XenDesktop 243

Integrating NetScaler with XenApp and XenDesktop Manual	245
Overview	245
NetScaler Gateway Prerequisites	245
Firewall Rules	246
StoreFront Services Deployment	247
Web Interface Site Configuration with NetScaler Gateway	247
Web Interface or XenApp Integration with NetScaler Gateway	248
Manual Configuration of Certificate Trust for Web Interface in NetScaler	248
NetScaler Gateway Integration with Web Interface and StoreFront Services	249
Citrix Receiver for Web	250
Secure Ticket Authority Configuration on NetScaler Gateway	250
To Configure Remote Access to StoreFront	251
Optimal NetScaler Gateway Routing	252
To Configure Optimal Routing	253
Beacons	254
Enabling Access Method Fallback with Policies	255
SSL Certificate Trust	255
Session Policies	255
Session Policy and Profile Creation for Citrix Receiver	256
Session Policy Expressions	257
Session Policy Creation for the NetScaler Gateway Plug-In	257
To Create the Session Policy Rule for the NetScaler Gateway Plug-In	257
Session Policies Binding and Priority Setting	258
Session Profile Creation	258
To Create a Session Profile Using the Configuration Utility	259
To Add a Profile to a Session Policy Using the Configuration Utility	259
Custom Clientless Access Policies Configuration for StoreFront Services	259
Remote Desktop Proxy	260
XenApp and XenDesktop Addition to a Single Site	260
XenMobile Platform Overview	260
NetScaler Gateway with XenMobile Integration	261
NetScaler Gateway with XenMobile App Edition Deployment	262

Discussion Question	262
Module 11: AppExpert Policy Engines	263
AppExpert Default Policy Engine Manual	265
Overview	265
Understanding Policies	265
Policy Expressions	265
Expression Result Types	265
Understanding Packet-Processing Flow	266
Policy Process Evaluation Flow	268
Identifying Default Policy Expressions	269
Default Policy Expression Syntax	269
Default Policy Expression Objects	271
Default Policy Conversion	273
Actions	273
Action Syntax	274
Configuring Policies and Actions	274
Configuring Rewrite or Responder Policies	274
Understanding Bind Points	275
Policy gotoPriorityExpression Statement	276
Policy Binding Evaluation Process	277
Binding in the Policy Manager	277
Understanding Policy Labels	277
Configuring Policy Labels	278
Pattern Sets	278
How String Matching with a Pattern Set Works	278
Configuring a Pattern Set	278
Using Pattern Sets	279
Pattern-Set Operators	280
String Maps	281
Typecasting	282
Responder	284
Rewrite	285
URL Transformation	285
Responder Process	286
Responder Actions	286
Respond With	287
To Create a Respond With Action	288
Responder Action for Timeouts	288
Responder Policies	289
Responder HTML Page Imports	289
Responder Undefined Actions	289
Rewrite Process	291
Configuring a Rewrite Action	292
Rewrite Action Parameters	294
Rewrite Undefined Actions	295

To Configure a Rewrite Action Using the Configuration Utility	296
Configuring a Rewrite Action Using the Command-line Interface	297
Rewrite Policies	297
Binding Policies	297
Discussion Question	298
AppExpert Classic Policy Engine Manual	299
Overview	299
Policies Overview	299
Policy Basics	299
Basic Policy Components	300
Policy Priorities	300
Priority of the Policy	301
HTTP Header Settings	301
HTTP Request Headers	302
HTTP Response Headers	302
Discussion Question	305
Policies and Profiles Configuration	305
Policies	305
Bind Points	306
Citrix SmartAccess	306
Types of Policies	306
Policy Structure	307
Configuring Conditional Policies	309
Conditional Policies Configuration	309
System Expressions Configuration	310
Expression Types	310
Expressions	311
Expression Structures	312
Qualifiers, Operators, and Expression Values	312
Simple Expressions	312
Compound Expressions	313
Custom Expressions	313
Context-Sensitive Fields	314
Wildcards	314
Available Operators	314
Available Qualifiers for HTTP Traffic	315
Available Qualifiers for Non-HTTP Traffic	316
Identifying Operands, Wildcards, and Context-Sensitive Fields	317
Policies Creation	318
Accessing Policy Manager	318
Creating Policies in the Configuration Utility	318
Configuring Session Policies	318
To Create a Session Policy	319
To Bind a Session Policy by Using the Policy Manager	319
Traffic Policies	320
Traffic Policy Creation Example	320
To Configure a Traffic Policy	321

To Bind a Traffic Policy by Using the NetScaler Gateway Policy Manager	321
To Bind a Traffic Policy Globally by Using the Configuration Utility	322
To Unbind a Traffic Policy by Using the NetScaler Gateway Policy Manager	322
To Remove a Traffic Policy by Using the NetScaler Gateway Policy Manager	322
Discussion Question	322
Module 12: Multi-Tenancy and NetScaler SDX	323
Multi-Tenancy and NetScaler SDX	325
Multi-Tenancy with NetScaler	325
Multi-Tenancy Considerations	325
Introduction to the Different Kinds of Tenant Isolation	325
Tenancy Options with NetScaler	326
NetScaler Admin Partitions and Traffic Domains	326
Introduction to the NetScaler SDX Appliance	327
Product Benefits: True Multi-Tenancy	327
Product Benefits: Flexible Licensing and Configuration	328
Product Benefits: Simple Management	328
Hardware Platforms	328
Model Performance Comparison	328
Deployment Scenarios: Datacenter Consolidation	329
Deployment Scenarios: Multi-Tenancy	330
Deployment Scenarios: Tenant Isolation	330
Deployment Scenarios: Enabling Cloud Infrastructure Expansion	330
Licensing	330
Base Architecture	331
IO Virtualization	332
VLAN Filtering	332
Enabling VLAN Filtering on an Interface	333
Restricting VLANs to Specific Virtual Interfaces	333
NetScaler SDX High Availability	333
Service VM Overview	334
Device Management	334
Instance Management	334
Provisioning a NetScaler VPX Instance on an SDX Appliance	335
Administrator Profiles	335
XVA Template Repository	336
Uploading a NetScaler .xva File	336
Provisioning a NetScaler VPX Instance on a NetScaler SDX	336
NetScaler SDX Service VM Internals	337
Simple Consolidation	337
Data Plane Isolation with Shared Interfaces	339
Simple Consolidation with Delegated Administration	339
Consolidation Across Security Zones	340
SNMP	341
SNMP Trap Destinations	341
Adding an SNMP Trap Destination	341

Adding an SNMP Manager Community	342
Configuring an SNMP Manager	342
Configuring the NetScaler for SNMPv3 Queries	342
SNMP Views	343
Configuring an SNMP View	343
SNMP Users	343
Configuring an SNMP User	343
SNMP Alarms	344
Modifying a Predefined SNMP Alarm	345
System Health Monitoring	345
Third-Party Virtual Machines	345
Managing a NetScaler SDX Appliance: Overview	346
Managing a NetScaler SDX Appliance: Adding and Configuring a NetScaler VPX Instance	346
Module 13: Monitoring and Administration	349
Monitoring and Management Manual	351
Overview	351
Monitoring Needs	351
Monitoring Methods	352
SNMP	352
AppFlow	352
EdgeSight Monitoring	352
Insight Center Overview	353
Action Analytics	353
Command Center	353
Dashboard	355
System and Feature Counters	355
Dashboard Components	356
Reporting Tool	357
Working with Reports	358
Using Built-in Reports	358
AppFlow on the NetScaler System	359
How AppFlow Works	359
How Insight Center Collects AppFlow Data	359
HDX Insight	360
Configuring Insight Center	361
Enabling Data Collecting for Insight Center	361
NetScaler Log Management	362
Audit Logging	362
Configuring NetScaler for Audit Logging	363
Configuring an Auditing Server	363
Global Auditing Parameters	364
Configuring Auditing Policies	365
Binding Auditing Policy	365
Audit Messages	365

Discussion Question	366
Troubleshooting Resources	366
Citrix Knowledge Center	366
Citrix Product Documentation	366
Citrix Technical Support	366
Collected NetScaler Data	367
Additional Resources	367
Troubleshooting Tools	368
Newslog	368
Show TechSupport Command	369
Using the Visualizer	369
Display NetScaler System Information	369
Troubleshooting Hardware Issues	369
Troubleshooting License Issues	370
Non-Maskable Interrupt (NMI) Button	371
Display Software Information	371
Discussion Question	371

Not for resale or distribution

Not for resale or distribution

Module 1

Getting Started

1

Not for resale or distribution

Not for resale or distribution

Getting Started Manual

Overview

The Citrix NetScaler product line optimizes the delivery of applications over the Internet and private networks, combining application-level security, optimization, and traffic management into a single, integrated appliance. The NetScaler features that you enable and the policies you set are then applied to incoming and outgoing traffic.

After completing this module, you will be able to:

- Identify the capabilities, functionality, and features of the NetScaler.
- Identify basic NetScaler nCore configuration architecture.
- Identify basic NetScaler hardware and components.
- Identify deployment considerations and decide which deployment configurations are best suited for specific scenarios.
- Perform the initial NetScaler setup and configuration.
- Obtain, install, and manage NetScaler licenses.
- Describe how to enable, license, upgrade, and initially deploy a NetScaler Gateway configuration.

Introduction to Citrix NetScaler

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 control and manage incoming and outgoing traffic.

NetScaler Functionality

NetScaler content switching and load balancing dramatically improve the throughput and scalability of an Internet application infrastructure by decoupling each application request/response flow from the underlying transport.

Content switching and load balancing ensure the most efficient use of transport protocols and resources, even in a scenario where all the content is encrypted or compressed.

The NetScaler system manages the complete life cycle of the request/response transaction. With this management, the NetScaler system is uniquely equipped to direct and control application requests most efficiently, from the client to the server and back again.

Connection multiplexing (also known as connection reuse) enables the servers to provide all requested data while handling fewer connections than are received by the NetScaler system.

This efficient use of the HTTP specification provides a significant boost to the effective capacity of the server by reducing server CPU load. With this separation, the NetScaler system can use the TCP proxy architecture to multiplex and reuse the server-side TCP connection independently from a client-side connection. This reuse of established and idle server-side TCP connections reduces the TCP overhead on web servers.

NetScaler Overview

Citrix NetScaler is an application switch that performs application-specific traffic analysis to intelligently distribute, optimize, and secure layer-4 through layer-7 (L4-L7) network traffic for web applications. For example, a NetScaler system makes load-balancing decisions on individual HTTP requests rather than on the basis of long-lived TCP connections, so that the failure or slowdown of a server is managed much more quickly and with fewer disruptions to clients. NetScaler functions are broadly categorized into features, such as switching, security, protection, and farm optimization.

Switching

When deployed in front of application servers, a NetScaler system ensures ideal distribution of traffic. You can segment application traffic according to information in the body of an HTTP or TCP request, and on the basis of L4-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 correct servers.

Security and Protection

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

Optimization

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. Optimization improves server performance in the farm and therefore speeds up applications. A NetScaler system 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.

Product Features

The key feature sets of a NetScaler system provide:

- Application availability
- Application acceleration
- Application security
- Simple manageability
- Web 2.0 optimization

Feature accessibility depends on the hardware, licenses, and software of the specific NetScaler system.

For more information about NetScaler features, see the Citrix NetScaler Datasheet at http://citrix.com/content/dam/citrix/en_us/documents/products/netscaler-data-sheet.pdf.

Application Availability

The NetScaler system ensures that applications and services are available to all users. The following table identifies the features supported by the NetScaler system for application availability.

Feature	Platinum Edition	Enterprise Edition	Standard Edition
L4 load balancing and L7 content switching	X	X	X
Microsoft SQL, MySQL	X	X	X
AppExpert Rate Controls	X	X	X
IPv6 support	X	X	X

Feature	Platinum Edition	Enterprise Edition	Standard Edition
Traffic domains	X	X	X
Global server load balancing (GSLB)	X	X	X*
Dynamic routing protocols	X	X	
Surge protection and priority queuing	X	X	X
TriScale clustering*	X	X	X*

X*: Optional feature

Application Acceleration

The NetScaler system maximizes application performance with its set of powerful acceleration capabilities, including intelligent data compression, static and dynamic content caching, and multiple TCP optimizations that improve the efficiency of the network.

The following table identifies the licenses supported by the NetScaler system for enhancing application acceleration.

Feature	Platinum Edition	Enterprise Edition	Standard Edition
Client and server TCP optimizations	X	X	X
Citrix AppCompress for HTTP	X	X	X*
Citrix AppCache	X	X*	

X*: Optional feature

Application Security

NetScaler protects against a wide variety of threats with integrated security capabilities that protect application resources, augmenting existing network-layer security protections. The following table identifies the licenses supported by the NetScaler system for enhancing application security.

Feature	Platinum Edition	Enterprise Edition	Standard Edition
L4 DoS defenses	X	X	X
L7 rewrite and responder	X	X	X
NetScaler Gateway, SSL VPN	X	X	X
XenMobile NetScaler Connector	X	X	
SAML2 support	X	X	X
AAA for traffic management	X	X	
Citrix Application Firewall	X	X*	
Citrix XML Firewall	X	X*	NetScaler CloudBridge Connector
X			

X*: Optional feature

Simple Manageability

The NetScaler system provides key features for simple manageability:

- NetScaler Insight Center (HDX Insight not included in Standard Edition)
- AppExpert visual policy builder
- Action Analytics
- AppExpert service callouts, templates, and visualizers
- Role-based administration and Authentication, Authorization, and Auditing (AAA) for administration
- Configuration wizards
- Native Citrix Web Interface
- Citrix Command Center for manageability

Front-end and TCP Protocol Optimizations

The following table identifies the licenses supported by NetScaler for Front-end and TCP protocol optimizations.

Feature	Platinum Edition	Enterprise Edition	Standard Edition
Content Layout	X	X	
Domain Sharing	X	X	
Image Optimization	X	X	
Style Sheets and Javascript Optimization	X	X	
Multi-path TCP	X	X	X
BIC and Cubic TCP	X	X	X

Lower Total Cost of Ownership

NetScaler reduces the total cost of ownership with web cache redirection. In the Enterprise and Platinum editions, you can automatically direct requests with content not cached on NetScaler to your cache farm. In addition, these editions include N-tier multilayer load balancing support of cache servers.

Discussion Question

Which NetScaler features do you or your organization plan to implement and why? Are there any NetScaler features that you would not implement for specific reasons?

Hardware Platforms

NetScaler product is available in three different appliance options to match the broadest range of business, performance, and deployment requirements.

- NetScaler MPX contains a full portfolio of hardware-based application delivery appliances delivering 500 Mbps to 120 Gbps of performance.
- NetScaler SDX hardware-based appliances with advanced virtualization can consolidate as many as 40 independently-managed NetScaler instances with as much as 50 Gbps of overall performance.

- NetScaler VPX software-based virtual appliances run on widely deployed hypervisors and support 10 Mbps to 3 Gbps performance levels.



The full suite of NetScaler capabilities is now available on the Amazon Web Services (AWS) environment.

Platform Performance	NetScaler MPX	NetScaler SDX	NetScaler VPX
System throughput, Gbps (range)	0.5 - 120	4 - 50	Up to 3.0
HTTP requests/sec (range)	175,000 - 4,700,000	600,000 - 3,700,000	Up to 100,000
SSL transactions/sec (2K key certificates) (range)	1,500 - 560,000	8,000 - 98,000	Up to 500
SSL throughput, Gbps (range)	0.5 - 75	3.5 - 11	Up to 1.0
Compression throughput, Gbps (range)	0.5 - 14.7	2.4 - 10	Up to 0.75
SSL VPN: concurrent users	5,000 - 35,000	30,000	Up to 500


NetScaler appliances are automatically compliant with the RoHS (Restriction of Hazard Substances) directives, as well as the WEEE (Waste Electrical & Electronic Equipment) directive. The NetScaler MPX appliance is SVHC (Substance-of-Very-High-Concern) compliant.

In addition, Citrix offers NetScaler MPX appliances that are FIPS (Federal-Information-Processing-Standard) compliant and support more than 4.5 Gbps of SSL throughput. For more information about FIPS-enabled NetScaler systems, see Citrix article CTX129543 at <http://support.citrix.com>.

For more information about NetScaler platforms, see the Citrix NetScaler Datasheet at http://citrix.com/content/dam/citrix/en_us/documents/products/netscaler-data-sheet.pdf.

Hardware Components

Each NetScaler system includes specific components. Some of these components are easily accessible on the exterior of a NetScaler appliance but should not be removed. The following are images of a NetScaler MPX appliance.



Redundant power supplies are also available as an option on all NetScaler appliances. These power supplies are interchangeable while the appliance is operating, which allows you to replace one power supply without shutting down the appliance, provided the other power supply is working. For more information, see the latest NetScaler Quick Start Guides on <http://support.citrix.com>.

Network and Serial Interfaces

On all physical NetScaler appliances, the network interfaces are on the front of the appliance.

The RS232 serial console port on the front of most NetScaler appliances provides a direct connection between the appliance and a workstation or laptop, allowing direct access to the appliance for initial configuration or troubleshooting.

For VPX appliances, the console is accessible using Hypervisor management tools.

File System

Important NetScaler data systems include:

- RAM Drive
- Flash Memory
- Hard Disk

The RAM drive is mounted on the root or `/` file system. During operation, binaries are run from the RAM drive. The running configuration files for BSD-level tools are used from the `/etc` file system residing on the RAM drive.



BSD (Berkeley Software Distribution) is a Unix operating system derivative. The NetScaler software runs on a FreeBSD operating system.

The flash memory is mounted in the NetScaler system as `/flash`. During startup, the RAM drive is initially populated from a compressed build file residing on the flash drive, and then the configuration files needed for BSD-level processes are copied from the `/nsconfig` directory on the flash drive into the `/etc` directory in the RAM drive. The `/nsconfig` directory is then linked to the `/flash/nsconfig` directory as a shortcut. When a save configuration command is executed, the running (in-memory) configuration is saved to this directory. The flash drive is on the back of the appliance; do not remove the flash drive.

The NetScaler system hard disk is used to store log data and core files and is used as long-term storage for unused NetScaler firmware builds. The `/var` directory represents the physical hard disk. The hard disk is on the back of the appliance; do not remove the hard disk.



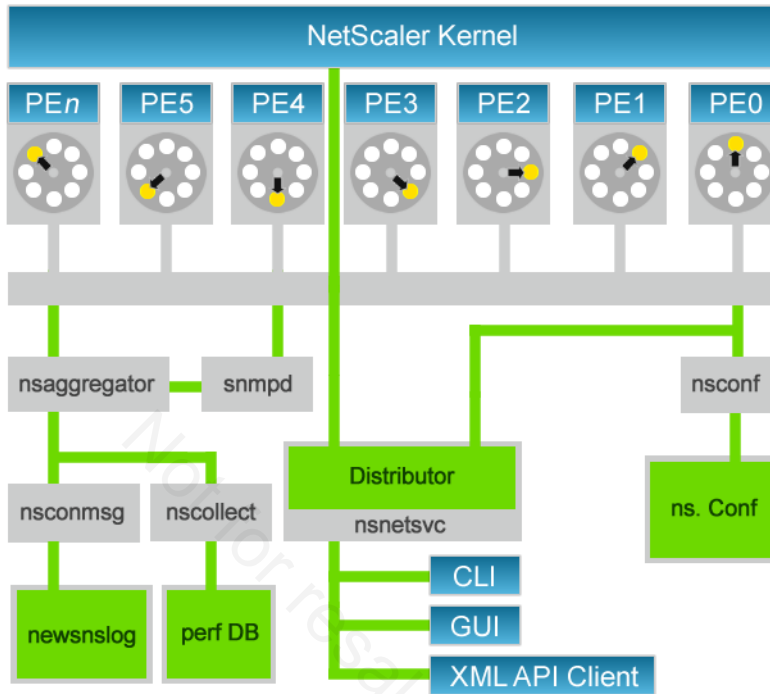
In some NetScaler modules, the hard disk has been replaced with a solid state drive (SSD).

nCore Configuration Architecture

NetScaler nCore uses multiple CPU cores for packet handling. The NetScaler nCore architecture includes the underlying NetScaler kernel and the cores, which are separate packet engines. The packet engines are designed to work independently. However, the cores communicate with each other through core-to-core messaging.

Commands that change the configuration are sent, by default, to all packet engines. Commands that do not change the configuration are sent to PE-0. `Dist_send` handles connection management, sends and receives commands and data to and from the packet engines and rolls back a single input-output control (IOCTL) command if it fails on any one packet engine. If the packet engine moves to the user space, it needs to be able to pick up packets from the user space as well.

The queues of the network interface cards (NICs) are initiated in the kernel and then sent to the packet engine in the user space without involving the kernel. This function preserves the line rate and does not introduce the kernel as an intermediary.



The above image shows the following components:

- "PE" represents a packet engine.
- "GUI" represents the configuration utility and other graphical user interfaces.
- "CLI" represents the command-line interface.
- "nsagggregator" collects and monitors data from the packet engines and communicates with various aggregator clients, such as:
 - "nscollect"
 - "nsconmsg"
 - "stat" command
 - "snmpd"
- "newsnslog" contains the performance data stored by nsconmsg.
- The distributor sends configuration commands to all of the packet engines.
- "nsnetsvc" includes the distributor module, which talks to the kernel and packet engines.
- "nsconf" contains all of the configuration information.



For more information, see the following document:

http://www.citrix.com/content/dam/citrix/en_us/documents/products-solutions/citrix-netscaler-ncore-tecnology.pdf

Deployment Scenarios


You can integrate the NetScaler system into any network, either as a complement to existing load balancers, servers, caches, and firewalls, or as a standalone solution that can provide one or more of those functions. A successful NetScaler deployment requires planning for the correct deployment type, as well as a full migration of current network functions.

NetScaler deployment options include:

- New technology
- Displacement

New Technology


The following illustrates a new technology deployment scenario in a high-availability scenario. Before starting a new technology deployment, you should map out the network architecture to maximize the placement of the NetScaler system. The mapping process should include identifying current network functions and deciding which additional NetScaler functions will be added.



Displacement

The following figure illustrates a displacement deployment where a NetScaler system (inserted in-line) replaces another traffic manager and attempts to meet the configuration of the old device. The NetScaler system often provides superior performance and more compelling functionality compared to that offered by other traffic managers.

You can deploy NetScaler systems in many scenarios to replace other products. As with a new deployment, it is critical with displacement to analyze and decide how to replicate the current functions of the network before deployment. Doing so ensures that network features are properly migrated to the new setup. You should explore new functionality only after you ensure that the existing functions are replicated.



SDX

A NetScaler SDX appliance is a true service delivery platform for enterprise and cloud datacenters. An advanced virtualized architecture supports multiple NetScaler instances on a single hardware appliance, while an advanced control plane unifies provisioning, monitoring, and management to meet the most demanding multi-tenant requirements.

With NetScaler SDX, each instance runs as a separate virtual machine with its own dedicated NetScaler kernel, CPU resources, memory resources, address space, and bandwidth allocation. Network I/O is performed in a way that not only maintains total system performance but also enables complete segregation of each tenant's data-plane and management-plane traffic.

NetScaler SDX effectively delivers multiple virtual application delivery controllers (ADCs) by enabling fully isolated and independent NetScaler instances to run on a single appliance. A significant difference compared to other implementations of the virtual devices concept, however, is

that NetScaler SDX provides complete isolation. As a result, memory, CPU cycles, and SSL cards are allocatable resources that you can move around and definitively assign to different NetScaler instances. This characteristic, namely complete isolation, is important since NetScaler processing is application centered.

NetScaler Gateway Overview

Citrix NetScaler Gateway is a secure application access solution that provides administrators with granular application-level policy and action controls to secure access to applications and data while allowing end 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 end users with a single point of optimized access to the enterprise applications and data they need. This unique combination of capabilities helps maximize the productivity of today's mobile workforce.

NetScaler Gateway includes the following functionality:

Secure Remote Access	Transmitted data is protected by industry standard SSL encryption.
Granular Control	User devices can be scanned to determine whether they meet the criteria that is deemed necessary for access to resources. If they fail the scan, a more limited set of resources can be provided instead.
Single point of access	Whether providing end users with VPN functionality or simply the ability to access published applications and desktops remotely, the NetScaler Gateway can be a single point of access for both.
Regulatory compliance	NetScaler Gateway can help organizations to provide levels of access that meet some of the highest regulatory compliance requirements.
Client choices	End users determine which connection method is suitable for the task that they need to complete. Administrators can also specify whether to present only a single option or multiple options, depending upon the connecting device.

Integration with XenApp and XenDesktop

Integration with Citrix XenApp and XenDesktop allows administrators to use SmartAccess functionality. SmartAccess allows session-specific settings to be applied to XenApp or XenDesktop sessions, based upon the outcome of a NetScaler Gateway scan. For example, SmartAccess can determine that a device is not a domain member and therefore not allow it to use clipboard redirection.

Administrators can also use technologies such as clientless access to enable end users to securely access application data such as Microsoft Outlook Web Access or SharePoint, without the need for installing any client software. File shares can also be presented to end users, enabling simplified access to the corporate environment through a single access interface.



For more information about features and platform specifications, see the NetScaler Gateway Data Sheet at <http://www.citrix.com>.

NetScaler Gateway Platforms

You can install NetScaler Gateway in any network infrastructure without making changes to the existing hardware or back-end software. NetScaler Gateway works with other networking products, such as server load balancers, cache engines, firewalls, routers, and IEEE 802.11 wireless devices.

MPX Appliance

The hardware version of NetScaler Gateway is a NetScaler MPX appliance. It supports classic and nCore NetScaler Gateway software deployments. This appliance supports Versions 9.2 - 11.0 of the NetScaler Gateway software.



NetScaler Gateway 10 or later must run on an nCore version of the appliance.

VPX

Citrix NetScaler Gateway VPX is a virtual appliance that delivers the same features and functionality as the physical appliance. You can deploy NetScaler Gateway VPX as a virtual workload on your own hardware, in addition to or as an alternative to using a physical appliance.

NetScaler Gateway VPX supports the following software versions:

- NetScaler Gateway 11.0
- NetScaler Gateway 10.5

- NetScaler Gateway 10.1
- NetScaler Gateway 10
- Access Gateway 9.3, Enterprise Edition

You can install the software on your hypervisor of choice and receive the same configuration options as with the physical appliance. End-user connections work the same as with the virtual appliance, and you can use the same settings that you configure on the physical appliance.

Initial NetScaler Access

The NetScaler administrator account is nsroot, and the default password is nsroot. The nsroot account is part of the default configuration. Citrix recommends that you change the default password during installation.

You can access the NetScaler system through the GUI-based configuration utility or the command-line interface. You can configure the start menu and timeout-session in the configuration utility logon window.

You can log on to the NetScaler configuration utility when the system and software requirements for the workstation have been met.

You can access the command-line interface directly by connecting a workstation to the NetScaler serial port or by using SSH to connect to the NetScaler IP address.



The NetScaler MPX has an initial IP address of 192.168.100.1. The NetScaler VPX must be assigned an IP address through the console on first start.

NetScaler Licenses

You must properly license a NetScaler system before you can deploy it to distribute, optimize, or secure network traffic for web applications. After you have obtained the licenses, you must install them on your appliance and then verify that you have enabled the features corresponding to the licenses. If you do not install a license on the appliance, the First-time Setup Wizard appears, which provides options for licensing, including license installation.

Possible licenses include:

- NetScaler platform license
- NetScaler upgrade license
- NetScaler option licenses
- NetScaler Gateway Universal license
- NetScaler Gateway Platform license

The NetScaler platform license is responsible for enabling all necessary features and includes five SSL VPN connections. This license is allocated by default to hostname "ANY". The rest of the

NetScaler licenses need to be allocated to the Host ID (MAC) of the appliance in order to enable the corresponding features. In the case of high availability, two licenses are required. For more information about licensing your NetScaler, see Citrix article CTX121062 at <http://support.citrix.com>.

NetScaler option licenses, provides enablement of additional features to augment the features already supported by the platform license. These option features include AppCompress, AppCache, Application Firewall, Global Server Load Balancing (GSLB), and EdgeSight for NetScaler. NetScaler options licenses are not mandatory.

By purchasing an upgrade license, you can upgrade your NetScaler from one edition to another. For example, customers with Standard Edition may purchase the Standard Edition upgrade to Enterprise or Platinum Edition.

The NetScaler Gateway Universal license allows you to increase SSL VPN concurrent usage so that you are not restricted to five SSL VPN connections. It floats across high availability pairs. You allocate the universal license to the NetScaler Licensing Hostname, which you can configure in `/nsconfig/rc.conf`.

A NetScaler Gateway Platform License is required for enabling ICA connections to Citrix XenDesktop desktops and applications. This license supports as many as 10,000 ICA connections. This license floats across high availability pairs as well. You must allocate this license to the NetScaler Licensing Hostname. For more information about configuring unlimited ICA connections, see Citrix article CTX125567 at <http://support.citrix.com>.

NetScaler Gateway Licensing

Before you can deploy NetScaler Gateway to support end-user connections, the appliance must be properly licensed.



Citrix recommends that you retain a local copy of all license files that you receive. When you save a backup copy of the configuration file, all uploaded license files are included in the backup. If you need to reinstall the NetScaler Gateway appliance software and do not have a backup of the configuration, you will need the original license files.

Before installing licenses on NetScaler Gateway, set the host name of the appliance and then restart NetScaler Gateway. You use the Setup Wizard to configure the host name. When you generate the Universal license for NetScaler Gateway, the host name is used in the license.

NetScaler Gateway License Types

NetScaler Gateway requires a Platform license. To allow connections to the network through the NetScaler Gateway plug-in, you must also add a Universal license. NetScaler Gateway VPX includes the Platform license.

Platform License

The Platform license allows end-user connections to published applications on XenApp or virtual desktops from XenDesktop. Connections through Citrix Receiver do not use a NetScaler Gateway Universal license for each connection. These connections only need the Platform license. The Platform license is delivered electronically with all new NetScaler Gateway orders, whether physical or virtual. If you already own an appliance covered by a warranty or maintenance agreement, you can obtain the Platform license through the Product Upgrades/Fulfillment toolbox on the Citrix website.

Universal License

The Universal license limits the number of concurrent user sessions to the number of licenses you purchase. The Universal license supports the following features:

- Full VPN tunnel
- Endpoint analysis
- Policy-based SmartAccess
- Clientless access to web sites and file shares

If you purchase 100 licenses, you can have 100 concurrent sessions at any time. When a user ends a session, that license is released for the next user. A user who logs on to NetScaler Gateway from more than one device occupies a license for each session.



When you add a Universal license, it is important to adjust the "Max Users" setting to your desired amount. The global setting limits overall connections to the appliance, and the virtual server setting limits connections for each NetScaler Gateway virtual server.

When you receive your NetScaler Gateway appliance, licensing occurs in the following order:

1. You receive the License Authorization Code (LAC) by email.
2. You use the Setup Wizard to configure NetScaler Gateway with the host name.
3. You allocate the NetScaler Gateway licenses from the Citrix website. Use the host name to bind the licenses to the appliance during the allocation process.
4. You install the license file on NetScaler Gateway or your license server.

An Express license is also available for NetScaler Gateway VPX. It allows as many as five concurrent user connections. For more information about NetScaler Gateway licensing, see Citrix product documentation at <http://docs.citrix.com>.

NetScaler Gateway Licensing Considerations

After you install NetScaler Gateway, you are ready to obtain your Platform or Universal license files from Citrix. You log on to the Citrix website to access your available licenses and generate a license file. After the license file is generated, you download it to a computer. When the license file is on

the computer, you then upload it to the NetScaler Gateway. For more information about obtaining your license files, see Citrix product documentation at <http://docs.citrix.com>.



You can install the same license on multiple NetScaler Gateway appliances. In addition, NetScaler Gateway appliances in different locations can share the same license file to facilitate disaster recovery deployments.

Before proceeding with your Universal license installation, verify that your Platform license is installed correctly.



Reviewing the license log file in the var/log directory is a useful troubleshooting step when licenses do not appear in the console. A common cause is a host name mismatch, which is noted in the log.

To Install the NetScaler Gateway License

To install the NetScaler Gateway or NetScaler license using the configuration utility:

1. In a web browser, type the IP address of the NetScaler Gateway or NetScaler system, such as `http://192.168.100.1`.
2. In the User Name and Password fields, use the credentials of a NetScaler administrator account.
3. Browse to **Configuration > System**, and then click **Licenses**.
4. In the details pane, click **Manage Licenses**. If the `/nsconfig/license` directory does not exist, you are prompted to create it.
5. Click **Update License** and then select **Upload License Files**.
6. In the Select License Files dialog box, browse to the location of the license files and select the file you want to upload.
7. Click **Reboot** to apply the license.
8. In the **Reboot** dialog box, click **OK** to proceed with all the changes. Click **Close** to cancel the changes.

NetScaler Gateway Pre-Installation Checklist

Citrix provides a template for a pre-installation checklist that covers many different deployment options, ensuring that all of the key information is captured and ready for the deployment. For more information about pre-installation checklists, see Citrix product documentation at <http://docs.citrix.com>.

It is highly recommended that you review the checklist before implementing NetScaler Gateway in your environment.

Replacing Secure Gateway

Citrix Secure Gateway is limited to supporting only XenApp or XenDesktop 5.6 sessions. XenDesktop 7 does not support Secure Gateway or Web Interface. However, NetScaler Gateway provides enhanced functionality for the following:

- Additional protocols
- Access to email, web applications, and file shares
- Support for Citrix Receiver
- Integration with Citrix XenMobile
- VPN functionality

Deploying NetScaler Gateway as a replacement for Secure Gateway also allows an administrator to remove the existing Windows Server that is running Secure Gateway from the perimeter network.

For additional information about replacing Secure Gateway with NetScaler Gateway, see Citrix product documentation at <http://docs.citrix.com>.

Configuring NetScaler Gateway For First-Time Use

To configure NetScaler Gateway for the first time, you need an administrative computer configured on the same network as the appliance.

You must assign a NetScaler Gateway IP (NSIP) address as the management IP address of your appliance and assign a subnet IP (SNIP) address to which your servers can connect. You assign a subnet mask that applies to both NetScaler Gateway and SNIP addresses. You must also configure a time zone. If you assign a host name, you can access the appliance by specifying its name instead of the NSIP address.

There are two sections in the First-time Setup Wizard. In the first section, you configure the basic system settings for the NetScaler Gateway appliance, including:

- NSIP address, SNIP address and subnet mask
- Appliance host name
- DNS servers
- Time zone

In the second section, you install licenses. If you specify the address of a DNS server, you can use the hardware serial number (HSN) or license activation code (LAC) to allocate your licenses, instead of uploading your licenses from a local computer to the appliance. Citrix recommends saving your licenses to your local computer.

When you finish configuring these settings, NetScaler Gateway prompts you to restart the appliance. When you log on to the appliance again, you can use other wizards and the configuration utility to configure additional settings.

Settings Configuration

NetScaler Gateway provides wizards for configuring connections to hosted resources. These wizards are designed to guide an administrator through the deployment of NetScaler Gateway and automate some tasks based on responses entered in the wizard. For example, HTTP to HTTPS redirection can be automatically configured, using a fully qualified domain name that an administrator specifies during the wizard.

The wizards can be run by selecting the NetScaler Gateway node within the configuration utility and then selecting the relevant wizard under the Getting Started menu.

NetScaler Gateway Wizard

The NetScaler Gateway Wizard provides integration with App Controller, Citrix StoreFront Services and Web Interface. When you complete the wizard, NetScaler Gateway can communicate with App Controller or StoreFront and end users can access the following resources:

- Windows-based applications
- Virtual desktops
- Web applications
- Software-as-a-Service (SaaS) applications
- Mobile applications

StoreFront is used to provide access to Windows based applications and desktops; it also replaces Citrix Web Interface.

Before running the wizard to configure the NetScaler Gateway, gather the following information:

Wizard Component	Requirement
Virtual Server Name, IP address, and port	A descriptive name for the virtual server and the IP address and port that the gateway will listen on.
Redirection from an unsecured to a secure port	Provide the fully qualified domain name of the NetScaler Gateway, which will then be used to redirect connections made on port 80 to the secure port.
LDAP Server	Domain controller location and LDAP bind account details.
RADIUS Server	Location of two-factor RADIUS server.

Wizard Component	Requirement
DNS Server	DNS Server to be used for client connections, typically only used when a client establishes a VPN connection with the NetScaler Gateway plug-in.
App Controller, StoreFront and Web Interface	Specify the FQDN of AppController, StoreFront server or Web Interface services.

Running the NetScaler Gateway Wizard allows you to configure a NetScaler Gateway virtual server and provide information that is then configured globally.

Wizard Component	Requirement
Create a virtual server	Enter an IP address, port and virtual server name.
Specify a server certificate	<p>The following options are available:</p> <ul style="list-style-type: none"> • Use a test certificate • Install a new certificate • Use an installed certificate <p>The certificate will then be bound to the gateway and presented to devices when SSL connectivity is established.</p>
Configure Name Service	Specify DNS and WINS addresses. DNS is mandatory, WINS is optional. DNS or WINS can be set as the preferred name resolution service.

Wizard Component	Requirement
Configure Authentication	<p>Provide the location of primary authentication source. This can be:</p> <ul style="list-style-type: none"> • RADIUS • LOCAL (users based on NetScaler Gateway) • LDAP (Active Directory uses this option) • TACACS • CERT • SAML • SAML IDP • DFA • Web
Configure Additional Settings	<p>Specify the default authorization action, allow or deny. Provide the fully qualified domain name to which end users should be redirected if they attempt to access the gateway through an unsecured port. This simplifies configuration by ensuring that if a user who inputs http:// instead of https://, is redirected to https://.</p>
Configure Clientless Access	<p>Specify whether end users should be able to use Clientless Access to access resources. This setting can be overridden with session policies, allowing granular levels of access. Provide the host name for any SharePoint servers that clientless access will use. This ensures that clientless access functions correctly with SharePoint.</p>

The NetScaler Gateway Wizard allows the use of a test certificate while an administrator requests a suitable SSL certificate from a CA. If this certificate is to be used for testing, the Root CA certificate must be imported into the end user's device as a trusted Root CA. This certificate is offered for download by NetScaler Gateway when selecting the option to use a test SSL certificate.

NetScaler for XenApp and XenDesktop Wizard

The NetScaler for XenApp and XenDesktop Wizard enables you to configure access to servers running Citrix XenApp or XenDesktop.

Wizard Component	Requirement
Select a Virtual Server	Choose an existing NetScaler Gateway virtual server.
Configure client connections	Provide a primary Web Interface server address and provide a backup if one is available. Provide the domain's name, so that single sign-on functions correctly. Add any STA servers that should be used for ticketing and ensure that these are configured identically on Web Interface.
Configure SmartAccess	Select any policies that should be bound to the gateway and specify the order of priority. These policies can then be configured in XenApp or XenDesktop for SmartAccess use.

End-User Access with the FQDN

End users gain access to the gateway by using the fully qualified domain name (FQDN) of a NetScaler Gateway virtual server. SSL certificates are generally assigned to fully qualified domain names, which enables a device to verify the identity of the server it is connecting to. When requesting an SSL certificate, it is important to ensure that the correct fully qualified domain name is specified.

A fully qualified domain name is provided to end users to enable access. For example, `portal.example.com` could be the fully qualified domain name that is used to gain access to a NetScaler Gateway virtual server.

To View an Imported Certificate FQDN

To view the fully qualified domain name of a certificate that has been imported to the appliance take the following steps in the configuration utility:

1. In the navigation pane, expand **Traffic Management**>**SSL**, and then click **Certificates**.
2. In the details pane, select a certificate, and then click **Action**, then **Details**.
3. In the **Certificate Details** dialog box, expand **Subject**. The FQDN of the certificate appears in the lower pane.

Configuration Testing

An important step when deploying a NetScaler Gateway appliance is to test all components.

Accessing a NetScaler Gateway virtual server by using an IP address over HTTPS results in an SSL certificate warning. To work around this, modify the device's host file to map the IP address to the host name specified on the SSL certificate for the duration of testing. This will allow full functionality to be tested using the applied SSL certificate.

Testing can easily be completed by using either the fully qualified domain name of the appliance or the IP address in a web browser-based session.



XenApp and XenDesktop applications will not launch through the Web Interface if certificate warnings are present.

Policies can also be tested during this phase. To determine the possible results of analysis scans, bind and unbind test policies with endpoint analysis expressions configured during this phase.

Name Service Providers Configuration

NetScaler Gateway allows you to define a DNS server that should be used for connections. DNS provides name resolution services to networks. Ensuring that name resolution functions correctly is an important step to ensuring a smooth experience for end users.

By using the NetScaler Gateway Wizard you can set the DNS server addresses that should be used globally, and determine whether DNS is the preferred name resolution method.

WINS servers are supported and can be used if required. You should at least specify a DNS server.

For additional information about adding DNS servers, see Citrix product documentation at <http://docs.citrix.com>.

Performing an Upgrade

To upgrade the system software on NetScaler units in a high availability (HA) pair, first upgrade the secondary node and then the primary node.

Upgrading a Standalone NetScaler System

The NetScaler system can be upgraded in the configuration utility by using the Upgrade Wizard. Citrix recommends using the CLI.

If you use the command-line interface, use the following process to upgrade a standalone NetScaler system:

1. Create a backup copy of the `ns.conf` file, using the save to file option at `System\Diagnostics\Saved Configuration`.



You should back up a copy of the configuration file on another computer.

2. Create a release number `nsinstall` subdirectory in the `/var/nsinstall` directory.
3. Change directory to `/var/nsinstall/<release number>nsinstall`, create a directory named "build_<targetbuildnumber>" and change to this directory.
4. Download and extract the contents of the installation package to the `/var/nsinstall/release number/build` folder.
5. Run the `installns` script to install the new version.
6. When prompted, restart the NetScaler system.



For more information about upgrading a standalone NetScaler, see the following CitrixTV video at <http://video.citrix.com/tv/#videos/9686>

Save the Configuration

Like many networking devices, NetScaler Gateway allows you to have both a running configuration and a saved configuration. A saved configuration is loaded at startup and contains all changes that were present in the previous saved configuration.

Use either of the following methods to save the configuration:

- Running the `save config` command from the command-line interface
- Clicking the **Save** button in the configuration utility

Running Configuration

The running configuration contains all of the completed changes, including any that have not been saved. The running configuration is often used to test changes. If any issues are encountered, the configuration is reloaded when the device restarts.

To display any differences between the running and saved configurations, do the following in the configuration utility.

1. In the navigation pane, expand **System**, and then click **Diagnostics**.
2. In the details pane, under **View Configuration > Saved v/s running configuration**.

Discussion Question

Which tools do you currently use for configuration testing, and what is your typical process for configuration testing?

Module 2

Basic Networking

2

Not for resale or distribution

Not for resale or distribution

Basic Networking Manual

Overview

A NetScaler is usually deployed in front of a farm and functions as a transparent TCP proxy between clients and servers, without requiring any client-side configuration. To facilitate efficient and secure access to server resources, a NetScaler 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.

The NetScaler system is fundamentally a TCP proxy that separates the client connections from the server connections and manages separate connection tables for client and server connections. The NetScaler responds to client connections that are targeted at servers residing behind it, hiding the network topography and enforcing traffic security by acting as a single gateway that clients use to access the network.

After completing this module, you will be able to:

- Identify basic NetScaler networking architecture.
- Identify the IP address types that can be assigned to a NetScaler and the purpose of each address.
- Explain how the NetScaler fits into your network topology.
- Identify deployment considerations, and the advantages and disadvantages of specific deployment considerations.
- Configure virtual LANs (VLANs) and determine when to split your LAN into multiple VLANs.
- Identify the IP routing methods best suited for an environment.

OSI Networking Model



The Open Systems Interconnection (OSI) model, which is the accepted standard for describing networking technology, categorizes networking functionality into seven layers.

Layer	Description
Physical Layer (1)	The physical layer specifies the hardware standards for sending and receiving informational bit streams over a communication channel. This layer defines connector and interface specifications and cable requirements.


Layer	Description
Data Link Layer (2)	The data link layer provides physical addresses for devices so that they can access a network to send and receive data. This layer furnishes transmission protocol knowledge and management.
Network Layer (3)	The network layer determines how packets are routed and forwarded from one route to another. Routes are dependent on static tables that are built within the network. However, each path is calculated dynamically at the beginning of the conversation. This layer provides an end-to-end logical addressing system.
Transport Layer (4)	The transport layer provides communication (typically TCP or UDP protocols) between devices through a network. This layer handles error checking, network segmentation and multiplexing.
Session Layer (5)	The session layer permits end users to establish sessions in which they can log on to a remote time-sharing system or transfer data between devices. This layer provides a service to manage communication control, synchronization and token management.
Presentation Layer (6)	The presentation layer determines how applications represent incoming and outgoing data. This layer includes encryption, compression, graphics formatting and content translation.
Application Layer (7)	The application layer provides the end-user interface for a device connected to a network. This layer is not synonymous with the application itself.

NetScaler Architecture Overview

The NetScaler system is fundamentally a TCP (layer-4) proxy that separates the client connections from the server connections and manages separate connection tables for client and server

connections. The NetScaler system can provide great traffic optimization as a gateway device by multiplexing client connections to web servers.

As a TCP proxy device, the NetScaler system responds to client connections that are targeted at servers residing behind it. It therefore hides the network topography. The NetScaler system can enforce traffic security by acting as a single gateway that clients use to access the network.



The NetScaler system is not a UDP proxy. However, it still proxies the IP address, sourcing from a mapped IP address or subnet IP address as normal. This behavior can be turned on and off at a granular level.



For a list of the communication ports used by NetScaler and other Citrix technologies, see http://support.citrix.com/servlet/KbServlet/download/2389-102-706706/CitrixPorts_by_Product_and_Port_v2013_d.pdf

NetScaler IP Addresses

In the above figure, the virtual IP address and MIP/SNIP address are NetScaler owned IP addresses. The NetScaler system uses the following types of NetScaler owned IP addresses for management and proxying connections to the server:

- NetScaler IP (NSIP) addresses
- Mapped IP (MIP) addresses
- Subnet IP (SNIP) addresses
- Virtual IP (VIP) addresses

NetScaler IP Address

The NetScaler IP address (NSIP) is a unique IP address and the primary address for management and general system access. When NetScaler systems participate in a high-availability configuration, the NSIP address is used for primary communication between members of the high-availability configuration, and the NSIP is the only active IP address on the secondary member in a high-availability pair. The NSIP can be accessed from any enabled interface on the NetScaler system.

An NSIP address must be configured on a new NetScaler system. The default IP address and netmask is 192.168.100.1/16 (255.255.0.0).

Configuring an initial NSIP address or changing the NSIP address or subnet mask requires a restart of the NetScaler system. When using the command-line interface to change the configuration, first, save the configuration, then change the NSIP address, and then restart the NetScaler system.

Configuring a NetScaler IP Address

You can configure NSIP addresses through the Setup Wizard or through the command-line interface.

In the configuration utility, the NSIP can be changed only by running the Setup Wizard, located in the **System > Setup Wizard** node.

In the command-line interface, type the command below to start the text-driven utility for configuring and saving the NSIP address or subnet.

```
config ns
```

Or you can configure a NetScaler IP address directly, using the `set ns config` command.

```
set ns config -Ippaddress <IP address> -netmask <Subnet mask>
```

Mapped IP Address

A mapped IP (MIP) address is used for external connections from the NetScaler system. MIP addresses are used for connectivity in the absence of an SNIP address. For example, the MIP address is the proxy IP address of last resort. A MIP address, like a SNIP address, is used as the proxy address for NetScaler system-to-server communication. MIP addresses are still used even when the USNIP mode is globally disabled.

The MIP address should be available across all subnets and should never be bound to a VLAN. It is active only on the primary unit of a high-availability pair, like every other IP address on the system other than the NSIP address and shows as passive on the secondary unit.

When both an MIP address and a SNIP address are configured on the same subnet, the NetScaler system by default (USNIP mode enabled), uses the SNIP address to communicate with servers. If USNIP mode is disabled, the MIP address is used.

If multiple MIP addresses are present on a subnet, the NetScaler uses the MIP addresses in a round-robin fashion.

Configuring a Mapped IP Address

You can configure MIP addresses by defining an IP address and a subnet mask.

In the configuration utility, expand the **Network > IPs** node.

In the command-line interface, type:

```
add ns ip <IP address> <Subnet Mask> -  
mgmtAccess [Enabled | Disabled] -type MIP
```

Subnet IP Address

A subnet IP (SNIP) address is used in connection management and server monitoring. A SNIP address gives the NetScaler system an Address Resolution Protocol (ARP) presence in subnets to which the system might not be directly connected.

A NetScaler system should have a SNIP address configured for each directly connected subnet. When a SNIP is added to a NetScaler system, a static route entry is automatically added to the NetScaler system routing table; this route identifies the SNIP address as the default gateway on the NetScaler system for the corresponding subnet.

The Use Subnet IP (USNIP) mode can affect how a SNIP address is used by the NetScaler system to communicate with servers. With USNIP mode enabled (which is its default setting), the SNIP address functions as a proxy IP address, serving as the NetScaler-side end point for connections to the servers. In this mode, the SNIP address is the source IP address in the packets that a server receives from the NetScaler system.

If USNIP mode is disabled, the SNIP address is not used to send traffic from the NetScaler system to the servers. Instead, a mapped IP address must be available. In most environments, USNIP mode is left enabled.

Individual SNIP addresses can be enabled for management access. When management access is enabled, connections to the NetScaler command-line interface over SSH, and connections to the web-based configuration utility, can be made using the SNIP address (as if it were an NSIP address). Using management-enabled SNIP addresses allows you to connect to the NetScaler system from a subnet other than the one in which the NSIP is located. It also simplifies managing NetScaler systems in a high-availability configuration, since only the primary unit will respond to the SNIP. Management access is not enabled by default. Unlike the NSIP address, but like every other type of IP address, SNIP addresses are active only on the primary unit of a high-availability pair. They are passive on the secondary unit.

If multiple SNIP addresses are present on a subnet, the NetScaler alternates between the SNIP addresses in round-robin manner when communicating with servers.

Configuring a Subnet IP Address

You can configure SNIP addresses by defining an IP address and a subnet mask.

In the configuration utility, expand the **Network > IPs** node.

In the command-line interface, type:

```
add ns ip <IP address> <Subnet Mask> -  
mgmtAccess [Enabled | Disabled]
```



Because of the default behavior of the NetScaler system, IP addresses are not associated directly with interfaces (MAC addresses). As a result, all SNIP addresses are available on all interfaces by default.

Virtual IP Address

Virtual IP (VIP) addresses are assigned to virtual servers on the NetScaler system and used for client-to-NetScaler communication. A VIP address is usually presented to the clients as a logical abstraction of a physical server behind the NetScaler system.

When the VIP address is a public IP address, it usually corresponds to the DNS entry for a domain. A VIP address is automatically created when a virtual server is added. A virtual server is identified as a unique combination of IP address and port number.

Disabling or changing the status of a VIP address will affect all virtual servers using the VIP address.

To configure a virtual IP address

A VIP is typically associated with a virtual server, and some of the attributes of the VIP are customized to meet the requirements of the virtual server. You can host the same virtual server on multiple NetScaler systems residing on the same broadcast domain by using ARP and ICMP attributes. After you add a VIP or any IP address, the NetScaler sends and responds to ARP requests. VIPs are the only NetScaler-owned IP addresses that can be disabled. When a VIP is disabled, the virtual server using it goes down and does not respond to ARP or ICMP requests.

In the configuration utility, expand the **Network > IPs** node.

In the command-line interface, type:

```
add ns ip <IP address> <Subnet Mask> -type VIP -  
arp [Enabled|Disabled] -icmpresponse  
[NONE|ONE_VSERVER|ALL_VSERVERS|VSVR_CNTRLD]
```


To remove a NetScaler-owned IP address

You can remove any IP address except the NSIP. The following table provides information about the processes you must follow in order to remove the various types of IP addresses. Before removing a VIP, remove the associated virtual server.

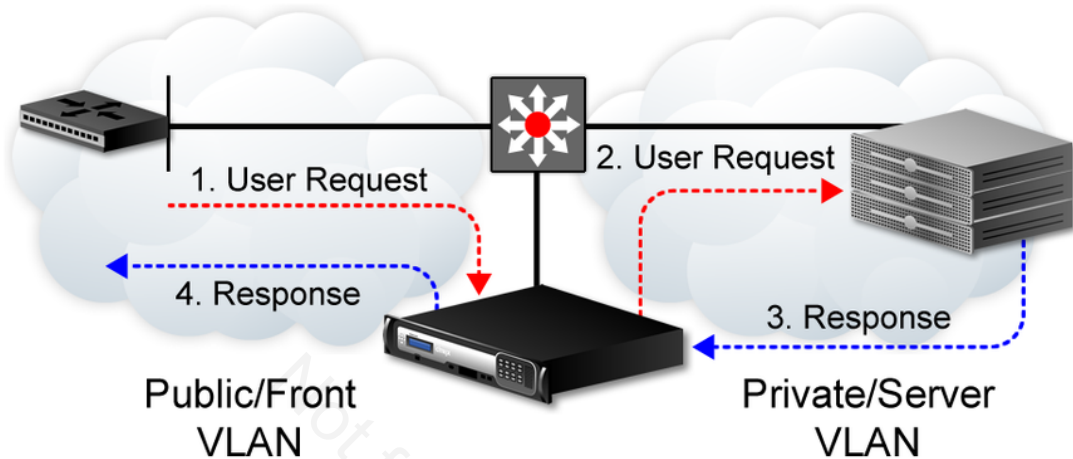
IP address type	Implications
Subnet IP address (SNIP)	If the IP address being removed is the last IP address in the subnet, the associated route is deleted from the route table. If the IP address being removed is the gateway in the corresponding route entry, the gateway for that subnet route is changed to another NetScaler-owned IP address.
Mapped IP address (MIP)	If a SNIP exists, you can remove the MIPs. The NetScaler uses the NSIP and SNIPs to communicate with the servers when the MIP is removed. Therefore, you must also enable Use SNIP (USNIP) mode.
Virtual Server IP address (VIP)	Before removing a VIP, you must first remove the virtual server associated with it. For more information about removing the virtual server, see Citrix article CTX132359 at http://support.citrix.com .

Network Topology

You can deploy the NetScaler system in the following network topologies:

- One-arm mode
- Two-arm mode

One-Arm Mode



One-arm topology requires most of the networking devices to be reconfigured to ensure that traffic passes through the NetScaler system. The two basic variations of one-arm topology are with a single subnet and with multiple subnets.

You can use a one-arm topology with a single subnet when the clients and servers reside on the same subnet. You can use a one-arm topology with multiple subnets when the clients and servers reside on different subnets.

For example, consider a NetScaler system deployed in one-arm mode for managing three servers. The servers are connected to a switch on the network. A virtual server of type HTTP is configured on the NetScaler system and HTTP services are running on the three servers. The three servers are on the private subnet, so a subnet IP address (SNIP) is configured through the NetScaler command line to communicate with them. The Use Subnet IP address (USNIP) option must be enabled so that the NetScaler system uses the SNIP instead of an MIP.

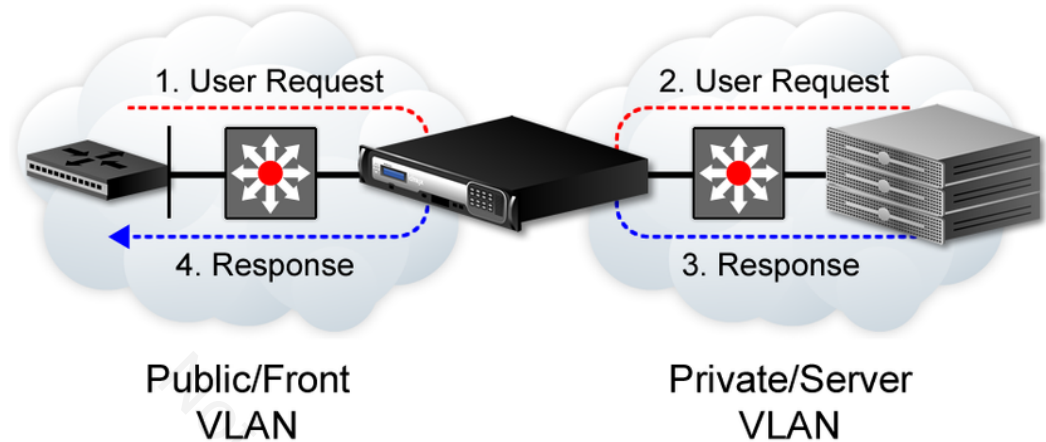
A one-arm mode configuration allows:

- A simple configuration with one physical interface and no risk of bridge loops.
- One or many VLANs with 802.1q tagging.
- Link aggregation to satisfy bandwidth requirements.

One-arm configurations are not recommended:

- When a server needs to see the real IP address of the client and the default gateway of the server cannot be changed to an IP address on the NetScaler system.
- When the servers are accessed directly and the default gateway has been changed to the NetScaler system, a situation that is sometimes addressed with the use of name-based services and explicit server definitions.

Two-Arm Mode



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 NetScaler system. The basic variations of two-arm topology are multiple subnets, typically with the NetScaler system on a public subnet and the servers on a private subnet and transparent mode, with both the NetScaler system and the servers on the public network.

One of the most commonly used topologies places the NetScaler system 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 a NetScaler deployed in two-arm mode for managing three servers, with a virtual server of type HTTP configured on the NetScaler system, and with HTTP services running on the servers. The servers are on a private subnet and a SNIP is configured on the NetScaler system to communicate with the servers. The Use SNIP (USNIP) option must be enabled on the NetScaler system so that it uses the SNIP instead of the MIP.

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 client needs them. A NetScaler system is placed between the client and the server, so that the traffic must pass through the NetScaler system. You must enable layer-2 mode for bridging the packets. The NSIP and MIP are on the same public subnet.

Two-arm mode has the following benefits:

- It allows layer-3 (routed) deployments with split subnets.
- It allows layer-2 (bridged) deployments with one subnet on each side.

Two-arm topologies work in some situations in which a one-arm configuration does not work. With two-arm mode, you can use the NetScaler system to route some traffic as well as to balance load.

NetScaler Gateway Deployment

You can deploy a NetScaler Gateway appliance at the perimeter of your organization's internal network (or intranet) to provide a secure single point of access to the servers, applications, and other network resources that reside in the internal network. All remote users must connect to the appliance before they can access any resources on the internal network.

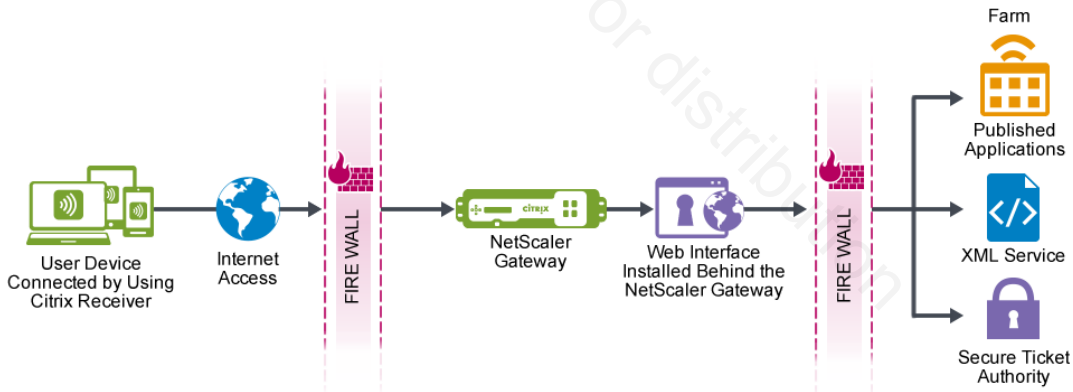
You can deploy the appliance with XenApp and XenDesktop to allow users to access their Windows, web and SaaS applications. If your deployment includes XenApp, you can deploy the appliance in a single-hop or double-hop perimeter network configuration. A double-hop deployment is not supported with XenDesktop.

You can deploy NetScaler Gateway in the following locations in your network:

- On the network perimeter
- In a secure network that does not have a perimeter network
- With additional NetScaler Gateway appliances to support load balancing and failover

NetScaler Gateway Deployment in the Perimeter Network

Many organizations protect their internal network with a perimeter network. A perimeter network is a subnet that lies between an organization's secure internal network and the Internet (or any external network). When you deploy NetScaler Gateway in the perimeter network, users connect through the Citrix NetScaler Gateway Plug-in or Citrix Receiver.



In the configuration shown above, you install NetScaler Gateway in the perimeter network and configure it to connect to both the Internet and the internal network.


When you deploy NetScaler Gateway in the perimeter network, user connections must traverse the first firewall to connect to NetScaler Gateway. By default, user connections use SSL on port 443 to establish this connection. To allow user connections to reach the internal network, you must allow SSL on port 443 through the first firewall.

NetScaler Gateway decrypts the SSL connections from the user device and establishes a connection on behalf of the user to the network resources behind the second firewall. The ports that must be open through the second firewall are dependent on the network resources that you authorize external users to access.

For example, if you authorize external users to access a Web server in the internal network and this server listens for HTTP connections on port 80, you must allow HTTP on port 80 through the second firewall. NetScaler Gateway establishes the connection through the second firewall to the HTTP server on the internal network on behalf of the external user devices.

NetScaler Gateway Deployment in the Secure Network

You can install NetScaler Gateway in the secure network. In this scenario, pictured in the following illustration, one firewall stands between the Internet and the secure network. NetScaler Gateway resides inside the firewall to control access to the network resources.



When you deploy NetScaler Gateway in the secure network, connect one interface on NetScaler Gateway to the Internet and the other interface to servers running in the secure network. Putting NetScaler Gateway in the secure network provides access for local and remote users. Because this configuration only has one firewall, the deployment is less secure for users connecting from a remote location. Although NetScaler Gateway intercepts traffic from the Internet, the traffic enters the secure network before users are authenticated. When NetScaler Gateway is deployed in a DMZ, users are authenticated before network traffic reaches the secure network.

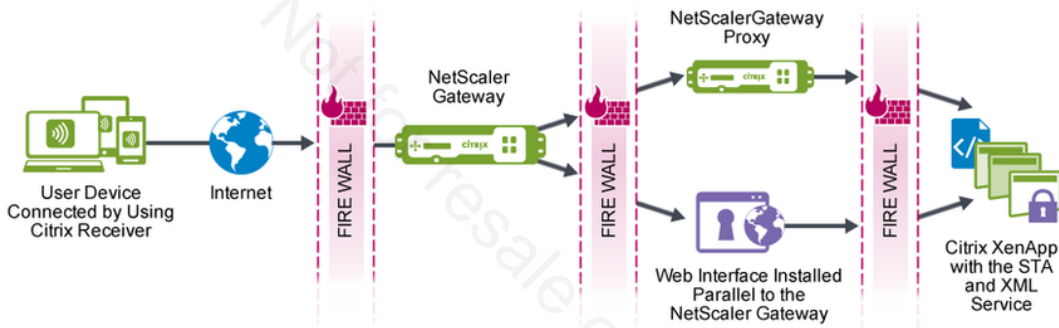
When NetScaler Gateway is deployed in the secure network, NetScaler Gateway plug-in connections must traverse the firewall to connect to NetScaler Gateway. By default, user connections use the SSL protocol on port 443 to establish this connection. To support this connectivity, you must open port 443 on the firewall.

Discussion Question

How have you deployed NetScaler or NetScaler Gateway previously in your organization? In the perimeter network or in the secure network? Which deployment do you plan to implement in the future? What are your reasons for using this type of deployment? Discuss with your classmates and instructor.

Double-Hop Deployment

Some organizations use three firewalls to protect their internal networks. The three firewalls divide the perimeter network into two stages to provide an extra layer of security for the internal network. The network configuration is called a double-hop deployment.



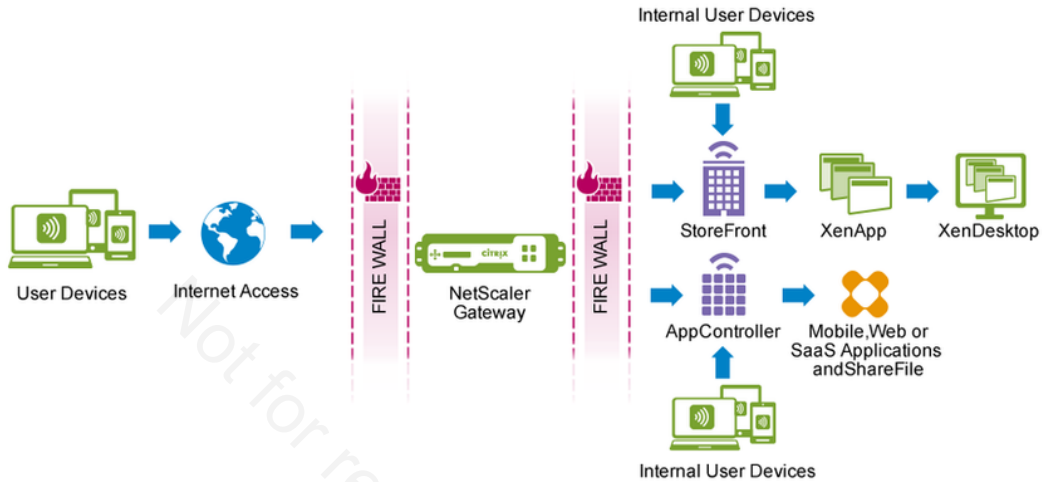
You can also have a double-hop deployment with one appliance in the perimeter network in one appliance in the secure network. If you configure a double-hop configuration with one appliance in the perimeter network and one in the secure network, you can ignore the instructions for opening ports on the third firewall.

NetScaler Gateway Deployment with XenMobile, XenApp and XenDesktop

You can provide access to your applications and desktops for remote and internal users by using NetScaler Gateway, XenMobile App Edition, XenApp and XenDesktop. NetScaler Gateway authenticates users and then allows them to access their applications by using Citrix Receiver and StoreFront.

XenMobile App Edition contains App Controller, which allows users to connect to web, SaaS and mobile applications and manage these applications for single sign-on (SSO), along with ShareFile documents. You install App Controller in the internal network. Remote users connect to App Controller through NetScaler Gateway to access their applications and ShareFile data. Remote users

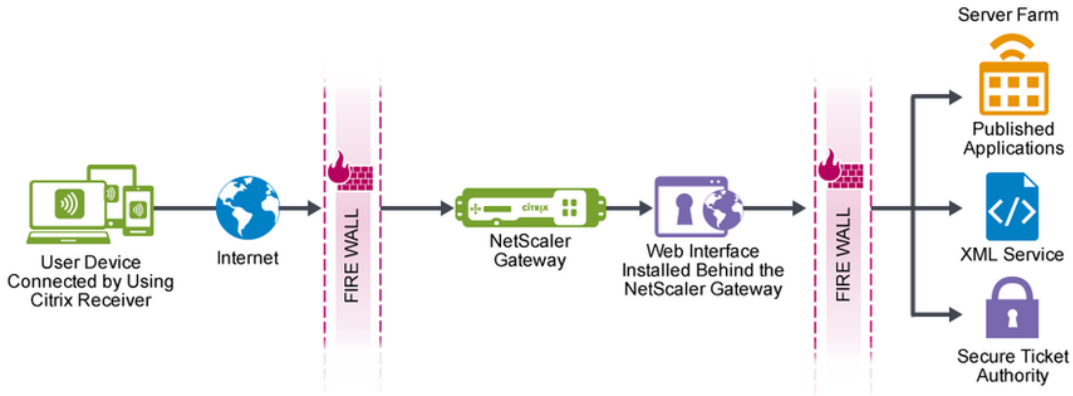
can connect with either the NetScaler Gateway plug-in, Receiver, or Worx Home to access applications and ShareFile. Users who are in the internal network can connect directly to App Controller by using Receiver. The following figure shows NetScaler Gateway deployed with App Controller and StoreFront.



With each deployment, StoreFront and App Controller must reside in the internal network and NetScaler Gateway must be in the perimeter network. For more information about deploying App Controller or StoreFront, see Citrix product documentation at <http://docs.citrix.com>.

NetScaler Gateway Deployment with Web Interface

In this configuration, both NetScaler Gateway and Web Interface are deployed in the perimeter network. When users log on with Citrix Receiver, the initial user connection goes to NetScaler Gateway and is then redirected to the Web Interface. To route all HTTPS and ICA traffic through a single external port and require the user of a single SSL certificate, NetScaler Gateway acts as a reserve web proxy for the Web Interface.



When the Web Interface is deployed behind NetScaler Gateway in the perimeter network, you can configure authentication on the appliance, but it is not required.

NetScaler Network Interfaces

Network interfaces in the NetScaler system are numbered in <slot>/<port> notation. After configuring your interfaces, you should display the interfaces and their settings to verify the configuration. You can also display this information to troubleshoot a problem in the configuration. In order to manage the network interfaces, you can enable and disable any interfaces. You can reset an interface to renegotiate its settings. You can clear the accumulated statistics for an interface. To verify the configuration, you can display the interface settings. You can also display the statistics for an interface to evaluate its health.

In addition to configuring individual interfaces, you can logically group interfaces by using VLANs to restrict data flow within a set of interfaces and you can aggregate links into channels.

Enabling or Disabling Network Interfaces Using the Configuration Utility

1. In the Navigation pane, expand the Network node, then the Interfaces node.
2. In the Interfaces pane, select the network interface that you want to enable or disable and select one of the following options:
 - Enable the network interface
 - Disable the network interface

Enabling or Disabling Network Interfaces Using the Command-Line Interface

At the NetScaler command-line interface, type one of the following pairs of commands to enable or disable an interface and verify the setting:

```
enable interface <interface_num>
show interface <interface_num>
```

or

```
disable interface <interface_num>
show interface <interface_num>
```

Virtual Local Area Networks (VLAN)

A NetScaler system supports layer-2 port and IEEE 802.1q tagged VLANs. VLAN configurations are useful when you need to restrict traffic to certain groups of stations. You can configure a network interface as a part of multiple VLANs by using IEEE 802.1q tagging.

You can configure VLANs and bind them to IP subnets. The NetScaler then performs IP forwarding between these VLANs (if it is configured as the default router for the hosts on these subnets).

The NetScaler supports the following types of VLANs:

Port-based VLAN

The membership of a port-based VLAN is defined by a set of network interfaces that share a common, exclusive, layer-2 broadcast domain. You can configure multiple port-based VLANs. By default, all network interfaces on the NetScaler are members of VLAN 1. If you apply 802.1q tagging to the port, the network interface belongs to a port-based VLAN. Layer-2 traffic is bridged within a port-based VLAN and layer-2 broadcasts are sent to all members of the VLAN if layer-2 mode is enabled. When you add an untagged network interface as a member of a new VLAN, it is removed from its current VLAN.

Tagged VLAN

802.1q tagging (defined in the IEEE 802.1q standard) allows a networking device (such as the NetScaler) to add information to a frame at layer 2 in order to identify the VLAN membership of the frame. Tagging allows network environments to have VLANs that span multiple devices. A device that receives the packet reads the tag and recognizes the VLAN to which the frame belongs. Some network devices do not support receiving both tagged and untagged packets on the same network interface, (in particular, Force10 switches). In such cases, you need to contact customer support for assistance. The network interface can be a tagged or untagged member of a VLAN. Each network interface is an untagged member of only one VLAN (its native VLAN). This network interface transmits the frames for the native VLAN as untagged frames. A network interface can be a part of more than one VLAN if the other VLANs are tagged. When you configure tagging, make sure that the VLAN configurations at the two ends of the link match one another. The port to which the NetScaler connects must be on the same VLAN as the NetScaler network interface.

NSVLAN

NSVLAN is the VLAN to which the NetScaler management IP (NSIP) address's subnet is bound. By default, NSVLAN has a VLAN ID (VID) of 1. If you unbind a network interface from its current port-based VLAN, it is added back to the default VLAN.



A Tagged VLAN configuration is neither synchronized nor propagated; you must perform the configuration independently on each unit in a high-availability pair.

VLAN Configuration

You can implement VLANs in the following environments:

- Single subnet
- Multiple subnets
- Single LAN
- VLANs (no tagging)
- VLANs (802.1q tagging)

If you configure VLANs with network interface membership only, the total number of possible VLANs is limited to the number of network interfaces available on the NetScaler. If more IP subnets are required with a VLAN configuration, 802.1q tagging must be used.

When you bind a network interface to a VLAN, the network interface is removed from the default VLAN. If the network interfaces need to be a part of more than one VLAN, you can bind the network interfaces to the VLANs as tagged members.

You can configure the NetScaler to forward traffic between VLANs at layer 3. In this case, a VLAN is associated with a single IP subnet. The hosts in a VLAN that belong to a single subnet use the same subnet mask and one or more default gateways connected to that subnet. Configuring layer 3 for a VLAN is optional. Layer 3 is used for IP forwarding (inter-VLAN routing). Each VLAN has a unique IP address and subnet mask that together define an IP subnet for the VLAN. In a high availability configuration, this IP address is shared with the other NetScaler systems. The NetScaler forwards packets between configured IP subnets (VLANs).



When you configure the NetScaler, you must not create overlapping IP subnets. Doing so impedes layer-3 functionality.

Port-based VLANs

The membership of a port-based VLAN is defined by a set of network interfaces that share a common, exclusive layer-2 broadcast domain. You can configure multiple port-based VLANs. By default, all network interfaces on the NetScaler are members of VLAN 1. This VLAN exists permanently. It cannot be deleted and its VLAN ID cannot be changed.

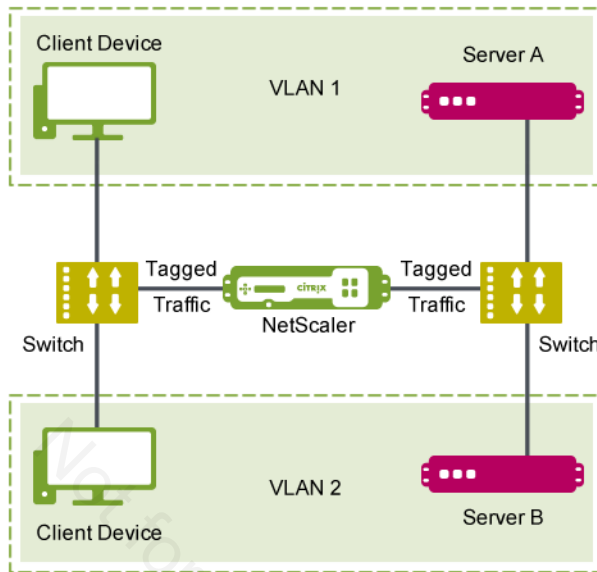
If you apply 802.1q tagging to the port, the network interface belongs to a port-based VLAN. Layer-2 traffic is bridged within a port-based VLAN and layer-2 broadcasts are sent to all members of the VLAN if layer-2 mode is enabled. When you add an untagged network interface as a member of a new VLAN, it is removed from its current VLAN. If you unbind a network interface from its current port-based VLAN, it is added back to the default VLAN.

Tagged VLANs

The figure below shows a tagged VLAN structure. IEEE 802.1Q specifications define the tagging structure of VLAN traffic. VLAN tagging inserts an additional header between the layer-2 and layer-3 headers in the packet. The additional header contains a protocol ID and a VLAN ID. The virtual network with which the packet is associated is identified by the VLAN ID. This tagging information can be used by layer-2 VLAN-aware devices to intelligently forward the data to ports associated with the same network. NetScaler tagged VLAN functionality is fully 802.1Q compliant.

IEEE 802.1Q specifies the following:

- Extra header between MAC header and IP header
- 2 byte protocol ID [TPID]
- 2 byte LAN ID [TCI]



NSVLAN

NSVLAN is a VLAN to which the NetScaler management IP (NSIP) address's subnet is bound. The NSIP subnet is available only on interfaces that are associated with NSVLAN. By default, NSVLAN is VLAN1, but you can designate a different VLAN as NSVLAN. If you do so, you must restart the NetScaler system for the change to take effect. After the restart, NSIP subnet traffic is restricted to the new NSVLAN.

The traffic from the NetScaler IP subnet can be tagged (802.1q) with the VLAN ID specified for NSVLAN. You must configure the attached switch interface to tag and allow this same VLAN ID on the connected interface.

If you remove your NSVLAN configuration, the NSIP subnet is automatically bound to VLAN1, restoring the default NSVLAN.

Configuring VLANs

In the configuration utility, expand the **Network > VLAN** node.

In the command-line interface, type:

```
show vlan <id>

bind vlan <id> -ifnum <InterfaceName>
[-tagged] [-IPaddress <IPaddress netmask>]
```

IP Address Routing

NetScaler systems support both dynamic and static routing. Because simple routing is not the primary role of a NetScaler, the main objective of running dynamic routing protocols is to enable route health injection (RHI), so that an upstream router can choose the best among multiple routes to a topographically distributed virtual server.

Most NetScaler implementations use some static routes to reduce routing overhead. You can create backup static routes, monitor the active routes to enable automatic switchover in the event that a static route goes down. You can also assign weights to facilitate load balancing among static routes, create null routes to prevent routing loops and configure IPv6 static routes. You can configure policy-based routes (PBRs), for which routing decisions are based on criteria that you specify.

Static Routes


Static routes are manually created to improve the performance of your network. You can monitor static routes to avoid service disruptions. Also, you can assign weights to Equal Cost Multi-Path Routing (ECMP) routes and you can create null routes to prevent routing loops.

The NetScaler supports monitoring of IPv4 and IPv6 static routes. You can configure the NetScaler to monitor an IPv4 static route either by creating a new ARP or PING monitor or by using existing ARP or PING monitors. You can configure the NetScaler to monitor an IPv6 static route, either by creating a new Neighbor discovery for IPv6 (ND6) or PING monitors or by using the existing ND6 or PING monitors.

MAC-based Forwarding Mode

MAC-based forwarding can be used to process traffic more efficiently by avoiding extraneous route/ARP lookups when forwarding packets. To avoid multiple lookups, the NetScaler system caches the source MAC address for every connection and returns the data to the same MAC address. With MAC-based forwarding enabled, NetScaler refers to its cache of MAC addresses before relying on its routing table.

MAC-based forwarding is useful with full VPN devices, because the NetScaler system ensures that the traffic flowing through the VPN passes through the same VPN device.



Determining the Source IP Address

When the NetScaler system communicates with the physical servers or peer devices, by default it does not use the IP address of the client. NetScaler maintains a pool of mapped IP addresses (MIPs) and subnet IP addresses (SNIPs) and selects an IP address from this pool to use as the source IP address of a connection to the physical server. Depending on the subnet in which the physical server is placed, NetScaler decides whether a MIP or SNIP should be used.



If the Use Source IP (USIP) option is enabled, NetScaler uses the IP address of the client.

Sending a Client IP Address to Servers

The NetScaler system usually functions in a transparent proxy configuration. Clients initiate connections to a VIP address on the NetScaler system. The system terminates the connection from the client, processes the packet and then initiates a connection to the appropriate server on behalf of the client.

The default behavior of the NetScaler system is to change the source and destination IP address of a packet received from a client before sending it to the server. The packet originating from the client contains the client IP address as the source IP address and the virtual server IP address (the VIP on

the NetScaler system) as the destination IP address. The NetScaler system changes the packet before sending it to the server, so that the source IP address becomes the NetScaler MIP/SNIP address and the destination IP address becomes the physical IP address of the server.

As a result of proxying the connection, the server is unable to see the IP address of the client that originated the connection; the server can see only the NetScaler MIP or SNIP address. Since some applications require the client IP address for proper logging or functionality, the NetScaler system has two ways of providing this information to the server:

- Client-IP HTTP header insertion
- Use Source IP mode

Link Aggregation

Link Aggregation combines data from multiple ports into a single high-speed link. Link aggregation allows multiple interfaces to be used in a combined manner to provide increased capacity/throughput and availability for the communication channel between the NetScaler system and other connected devices. Link aggregation is configured to address bandwidth limitations and to provide redundancy for interfaces (mitigating single points of failure).

In some environments, the speed of a single interface is not adequate for the volume of traffic that the NetScaler system manages. To address inadequate interface capability, you can combine multiple interfaces on the NetScaler system into a single, logical high bandwidth 802.3ad interface. The resulting aggregated interface is treated as a single interface for configuration purposes. The aggregate interface link throughput is the sum of the bound physical interfaces. The switch connected to the aggregate interfaces must also support 802.3ad on the NetScaler system.



Because of the default NetScaler behavior, two interfaces should never be plugged into the same VLAN or broadcast domain unless link aggregation/802.3ad is being configured.

Discussion Question

Have you used MAC-based forwarding mode in your NetScaler implementation? Would you consider using this mode in the future? Why or why not? Discuss with your classmates and instructor.

Not for resale or distribution

Module 3

High Availability

3

Not for resale or distribution

Not for resale or distribution

High Availability Manual

Overview

A high availability deployment of two Citrix NetScalers can provide uninterrupted operation in any transaction. In a high-availability pair configuration, only one system is active. This system, which is known as the primary, actively accepts connections and manages servers. All shared IP addresses are active on the primary system only.

The secondary system monitors the health of the primary system. If the secondary system is in a healthy state, it is ready to actively accept connections if the primary system is experiencing issues. This process prevents downtime and ensures that the services provided by the NetScaler system remain available even if one system ceases to function.

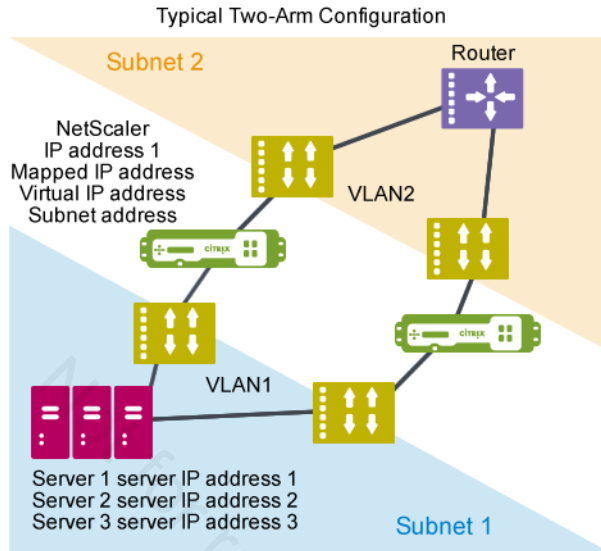


High availability packets are sent untagged by default, which can be an issue with a switch that handles tagged packets only. For more information, see Citrix Knowledge Base article CTX122921 at <http://support.citrix.com>.

After completing this module, you will be able to:

- Define high availability and the considerations for setup.
- Create a high-availability pair.
- Verify the high-availability setup.
- Configure high availability on the NetScaler system.

High-Availability Functionality



The figure shows a typical two-arm high-availability configuration. Both systems in a high-availability pair exchange UDP heartbeat messages that communicate the state of the other and ensure that only one unit is taking traffic at a time. This configuration is known as active/passive. High availability does not require feature enablement; however, high availability needs to be configured.

If a component of the primary system fails, a message to the secondary system instructs it to take over as the primary system. If the secondary system does not receive a heartbeat message from the primary system after a specific interval of time, it automatically assumes the primary role. This scenario is known as failover.



Failover occurs when three heartbeat messages are missed. High-availability heartbeat messages are sent every 200 milliseconds on UDP port 3003.

After a failover, all clients must reestablish their connections to the managed servers, but session persistence tables are synchronized between the primary and the secondary systems, honoring the previous persistence decision. For more information about TCP connection failover, refer to the Citrix NetScaler Networking Guide at <http://support.citrix.com>.

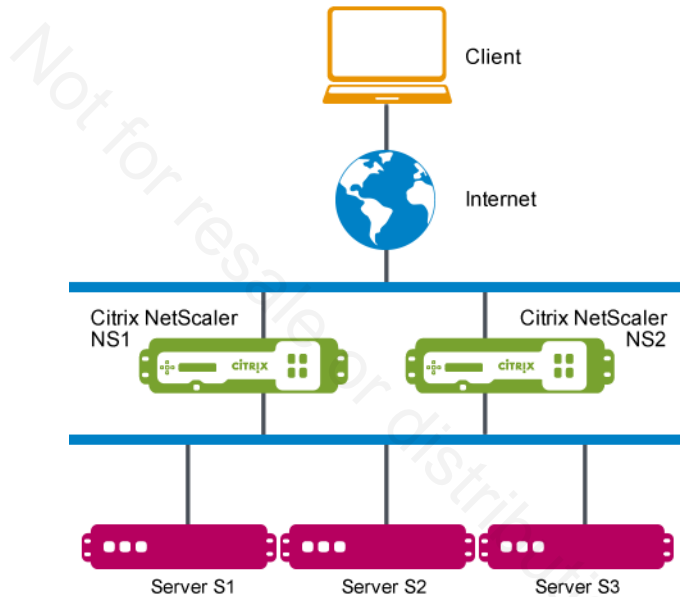


Neither system is considered primary if the primary node experiences a failure and attempts to fail over to a secondary system that is not healthy. Therefore, it is possible for both systems to be secondary at the same time and for neither system to process traffic.

High Availability Process

An HA deployment of two NetScaler Gateway appliances can provide uninterrupted operation in any transaction. When you configure one appliance as the primary node and the other as the secondary node, the primary node accepts connections and manages servers while the secondary node monitors the primary. If, for any reason, the primary node is unable to accept connections, the secondary node takes over.

The secondary node monitors the primary by sending periodic messages (often called heartbeat messages or health checks) to determine whether the primary node is accepting connections. If a health check fails, the secondary node retries the connection for a specified period, after which it determines that the primary node is not functioning normally. The secondary node then takes over for the primary (a process called failover).



With Web server logging persistence enabled, no log data is lost due to the failover. For logging persistence to be enabled, the log server configuration must carry entries for both systems in the `log.conf` file.

The basic steps to configure high availability are as follows:

1. Create a basic setup, with both nodes in the same subnet.
2. Customize the intervals at which the nodes communicate health-check information.
3. Customize the process by which nodes maintain synchronization.
4. Customize the propagation of commands from the primary to the secondary.
5. Optionally, configure fail-safe mode to prevent a situation in which neither node is primary.

6. Configure virtual MAC addresses if your environment includes devices that do not accept NetScaler Gateway gratuitous ARP messages.

High-Availability Node Configuration

A pair of NetScaler systems must be configured to become a high-availability pair. The process for configuring a high-availability pair involves first configuring the primary node then configuring the secondary node.

Citrix recommends that you set the status of the desired secondary node to stay secondary when nodes are configured. This practice ensures that an accidental failover does not occur during the configuration process, resulting in changes being made to the secondary rather than the primary node. Any changes that are made to the secondary node are not propagated to the primary node.



You should not use NetScaler VLANs when configuring high availability because it limits the NSIP address subnet. The NSIP address subnet should be available on all interfaces during high-availability configuration.

In a high-availability configuration, you can designate which interfaces to monitor for failing events. A failover occurs when any high-availability monitored interface goes down. If a particular interface is not being used, or if a failover is not required upon failure, the high-availability monitor should be disabled.



If an interface is not going to be used at all, a best practice is to disable it and turn off the associated high-availability monitor.

Pre-Configuration Checklist

Before configuring high-availability-pair nodes:

- Ensure that the NSIP addresses for the primary and the secondary nodes are unique from any other device on the network. The NSIP address can be changed using the `set ns config` command; this change requires a restart.
- Ensure that IP address conflicts can be viewed in the configuration utility from the System > Diagnostics > View console messages menu.



While the RPC node password must be identical in a high-availability configuration, the nsroot password on both systems does not have to be the same.

High-Availability Failover Process

The primary node in a high availability setup owns floating IP addresses, such as the MIP, SNIP and VIP addresses, and responds to address resolution protocol (ARP) requests with its own MAC address. Therefore, the ARP table of an external device, such as an upstream router, is updated with the floating IP address and the MAC address of the primary node.

When a failover occurs and the primary node is replaced by the secondary node, the replacement node will use the gratuitous address resolution protocol (GARP) to advertise the floating IP addresses that were learned from the original primary node. The MAC address that the new primary node advertises is the MAC address of its own network interface. Because some devices do not accept GARP messages, the external devices retain the IP address-to-MAC address mapping that the old primary node formerly advertised, which can result in a GSLB site becoming unavailable for those networks.

Configuring Primary and Secondary Nodes

You can use the following procedure to configure the primary and secondary nodes using the configuration utility or the command-line interface:

1. Disable the interfaces that are not connected or being used for traffic.
2. Disable monitoring for the interfaces whose failure should not cause a high-availability mode failover.
3. Assign a node ID to each NetScaler system.
4. Save the configuration.

Before two nodes are able to function in a high-availability pair, they must first be made aware of each other.

You join a new NetScaler system with an existing system to form a high-availability pair because Citrix recommends enabling "stay secondary" status on the new system. If "stay secondary" status is not enabled during the high-availability configuration, the new system can be elected to become the primary node, and the configuration settings on the existing system are replaced with the settings (typically a blank configuration) on the new system.

When "stay secondary" status is enabled on the new system, the existing system becomes the primary node if it passes its health checks and its configuration settings are synchronized with the new system.

High-Availability Status

After you have configured two NetScaler systems in a high-availability pair, it is important that you verify the status of each node to ensure that the system is prepared in case of failover. For each node, the node state should show as UP and the master state should show as either PRIMARY or SECONDARY, as appropriate. You should check that the sync state shows "success" on the

secondary node because it verifies the successful completion of the configuration synchronization process.

Verifying Status on the Appliance

You can check the master status of a node remotely by using either the configuration utility or the command-line interface. For example, the `show ha node` command will display the node state and master state. In addition, the LCD Configuration Display on the NetScaler appliance displays the status of the node.

Propagation and Synchronization

Command propagation, enabled by default, causes any command issued on the primary node to run on the secondary node. Consequently, you can test if high availability is working by making configuration changes to the primary node and then testing to see if the changes have been propagated to the secondary node.

Configuration synchronization occurs when a node comes up for the first time in a secondary state. The secondary node pulls the configuration from the primary node and overwrites its existing configuration.

In addition to automatic configuration synchronization, forced synchronization between two nodes is also supported. Forced synchronization is used whenever you want to ensure that changes made to a NetScaler configuration are transferred from the primary to the secondary system. For example, if a network interruption occurs and the systems are unable to communicate, a forced synchronization ensures that any changes made during the interruption are present on both systems.

In NetScaler releases 8.0 and later, if the nodes in a high-availability pair are running different versions of the NetScaler operating system, the node running the newest software version goes into listen mode. When a node is in listen mode, neither command propagation nor configuration synchronization occurs.

- Changes made using the `config ns` command in the command-line interface are not propagated. Any configurations made using this command must be performed on each node separately.
- The heartbeat messages are UDP packets sent to port 3003 of the other node in a high-availability pair.
- High-availability configuration synchronization occurs on TCP port 3010.
- Command propagation between the primary and secondary occurs on TCP port 3011.
- Secure high-availability configuration synchronization occurs on TCP port 3008.
- Secure command propagation occurs on TCP port 3009.



For more information about command propagation, see Citrix article CTX124344 at <http://support.citrix.com>.

Disabling Command Propagation

In some cases, command propagation might not be desired.

For example, if you make changes to the system configuration, and these changes cause problems, the changes are also made to the secondary node if command propagation is enabled.

When command propagation is disabled, you can first make changes on the primary node. Once the configuration is completed and is verified to be functioning properly, you can use the `force ha sync` command to ensure that the changes are propagated to the secondary node.

Automatic Configuration Synchronization

By default, configuration synchronization between the systems in a high-availability pair occurs automatically. Configuration synchronization occurs when:

- A node first comes up in the secondary state.
- A failover event occurs.
- A forced synchronization is issued.

When a save configuration is issued on the primary system, the running configuration on both systems is saved from memory to the "ns.conf" file on the flash drive. In some cases, automatic synchronization may not be desired.

Once the configuration is completed and is functioning properly, you can use the following command to make sure that the changes are synchronized into the secondary node:

```
force ha sync
```

Forced Synchronization

Forced synchronization can be performed on either the primary or the secondary node; however, if synchronization is already in progress, the command fails and a warning message is displayed.



Before performing a forced synchronization, you must save the configuration because the saved configuration is copied over but the running configuration is not.

The saved configuration file is the ns.conf file. The running configuration is stored in memory and might contain unsaved configuration changes.

Forced synchronization will also fail when it is run in the following circumstances:

- On a standalone NetScaler system
- On a NetScaler system on which high availability is disabled
- On a NetScaler system on which high-availability synchronization is disabled



The DONE message that is displayed after running the `force ha sync` command does not indicate that synchronization has been successful. To verify that synchronization was successful, run the `show node` command to verify that the configurations on both nodes are the same.

Performing a Forced Failover

When a failover occurs because the primary node fails, the primary node becomes secondary and the secondary node becomes primary. A forced failover is an administratively initiated failover on a node that produces the same results. A forced failover is desired if you want to test whether a high-availability pair configuration is performing correctly or to replace or upgrade the primary node.

Forced failover is performed in the configuration utility or with the `force ha failover` command in the command-line interface. The command can be issued on either the primary or the secondary node.

A forced failover will work only when:

- The primary node is able to determine that the status of the secondary node is UP.
- The health of the secondary node is good.
- The secondary node is not configured to stay secondary.



Performing a forced failover from the configuration utility might cause the client to lose connectivity with the NetScaler system. Citrix recommends that you perform a forced failover from the command-line interface.

Fail-Safe Mode

In a high availability configuration, fail-safe mode ensures that one node is always primary when both nodes fail the health check. This is to ensure that when a node is only partially available, backup methods are enabled to handle traffic as best as possible. The high availability fail-safe mode is configured independently on each node.

To enable fail-safe mode using the command-line interface, type:

```
set HA <node> [-failSafe (ON | OFF)]
```

For more information about fail-safe mode, see Citrix product documentation at <http://docs.citrix.com>.

High-Availability Management

While a high-availability pair can be managed using the unique NSIP addresses assigned to each node during initial configuration, Citrix recommends that you manage the pair using either an SNIP or MIP address. When a shared IP address is used to connect to either the command-line interface or the configuration utility, the connection is made to the primary node in the pair. If a failover occurs, the secondary node becomes primary, and the same MIP or SNIP address then connects to the new primary node. Using a SNIP or MIP address is also helpful when you need to perform configurations from a subnet different from that of the high-availability pair.

Every NetScaler system is assigned an MIP address or a range of MIP addresses during initial configuration; however, management access must be enabled on the MIP or SNIP address before it can be used to manage a high-availability pair.



Citrix recommends that you secure management access in the NetScaler system.

Upgrading a High-Availability Pair

When the nodes of a high-availability pair are running different versions of the system software, the node running the newest version goes into listen mode. When a node is in listen mode, neither propagation nor synchronization occurs.

Listen mode is used if you want to test the newest software version on one node before upgrading the second node. To conduct such a test, you perform a forced failover on the upgraded system, causing it to take over as the primary node. When a forced failover occurs, neither propagation nor synchronization occurs and all connections need to be re-established.

To upgrade the software of the NetScaler appliances in a high-availability setup, complete the following procedure:

1. Upgrade software of the secondary appliance.
2. Upgrade software of the primary appliance.

For more information about upgrading a high-availability appliance, see Citrix article CTX127455 at <http://support.citrix.com> and the Citrix TV video at <http://www.citrix.com/tv/#videos/9687>

Discussion Question

Have you previously updated a standalone NetScaler or high-availability pair of NetScalers? If so, what challenges did you encounter and why?

High-Availability Issues

High-availability issues include:

- Configuration synchronization failure
- File synchronization failure
- Unexpected failover

Configuration Synchronization Failure

Synchronization failure can be a result of connectivity issues, duplex mismatches, packet drops or the `/netscaler/nsnetsvc` process not running.

Perform the following tasks if synchronization between the primary and secondary node fails:

- Verify that the primary and secondary nodes can communicate with each other. Management and heartbeat messages are sent through layer-2 protocols. Layer-2 connectivity between the two high-availability nodes must allow the heartbeat to be received within 3 seconds.



One node will display a status of listening in this scenario.

- Ensure that any configured ACLs permit communication between the pair.
- Enter the following command to check the `inetd.conf` file to ensure that the `/netscaler/nsnetsvc` process is not disabled:

```
ns# more /etc/inetd.conf
```

- Ensure that the `nsnetsvc stream tcp nowait root /netscaler/nsnetsvc nsnetsvc` line contains comments.
- Enter the following command in the shell to check the `ns_com_cfg.conf` file on the secondary node:

```
ls -l
```

- Ensure that the `/tmp` directory has write permissions. For example:

```
drwxrwxrwt 4 root wheel 512 Aug 17 21:28 /tmp
```

- Verify that the two nodes are running the same version of the NetScaler operating system.

File Synchronization Failure

Both nodes in a high-availability pair might need a set of common-configuration or certificate files, depending on the needs of the deployment. If so, files might need to be manually synchronized. For example, if SSL offload is enabled, then SSL certificates must be copied to the same location

(directory) on both nodes. Similar examples include vsr.html (for SureConnect), any manually customized files or any other batch files containing configuration commands. For more information, see Citrix Knowledge Base article CTX138748.

Enter the following command in the command-line interface to manually synchronize files between nodes in a high-availability pair:

```
sync ha files mode
```

The following table lists available arguments:

Argument	Description
mode	Specifies the sync mode. Possible values include: <ul style="list-style-type: none">• all• bookmarks• ssl• htmlinjection• imports

The following table lists paths corresponding to synchronization mode:

Mode	Path
all	<ul style="list-style-type: none">• /nsconfig/ssl/• /var/vpn/bookmarks/• /nsconfig/htmlinjection/
ssl	/nsconfig/ssl/
bookmarks	/var/vpn/bookmarks/
htmlinjection	/nsconfig/htmlinjection/

Unexpected Failover

If the NetScaler systems are failing over unexpectedly, type the following command in the command-line interface to view current events that might be causing the failover.

```
show ns hardware
```

Possible causes include:

- An interface is down.

- An SSL acceleration card is down.
- The primary node has failed.

High-Availability Pair in Different Subnets

A HA deployment can also consist of two NetScaler Gateway appliances in which each appliance is located in a different network. You can also configure link redundancy and route monitors. These NetScaler Gateway functions are helpful in a cross-network high-availability configuration. The functions also cover the health check process used by each NetScaler Gateway to ensure that the partner appliance is active.

When the appliance in an HA pair reside on two different networks, the secondary NetScaler Gateway must have an independent network configuration. This means that the NetScaler Gateway appliances on different networks cannot share mapped IP addresses, virtual LANs or network routes. This type of configuration, in which the NetScaler Gateway appliances in an HA pair have different configurable parameters, is known as independent network configuration or symmetric network configuration. For more information about this configuration, see Citrix product documentation at <http://docs.citrix.com>.

Adding a Remote Node

When two nodes of an HA pair reside on different subnets, each node must have different network configurations. Therefore, to configure two independent systems to function as an HA pair, you must specify independent network computing mode during the configuration process.

When you add an HA node, you must disable the HA monitor for each interface that is not connected or being used for traffic.

To Add a Remote Node for Independent Network Computing Mode

1. In the Configuration Utility, on the **Configuration** tab, in the navigation pane, expand **System** and then click **High Availability**.
2. In the details pane, click the **Nodes** tab, then click **Add**.
3. In the **High Availability Setup** dialog box, in the **Remote Node IP Address** text box, type the NetScaler Gateway IP address of the appliance that is the remote node.
4. If you want to add the local node to the remote node automatically, select **Configure remote system to participate in High Availability setup**. If you do not select this option, you need to log on to the appliance represented by the remote node and add the node that you are currently configuring.
5. Click to enable **Turn off HA monitor on interfaces/channels that are down**.
6. Click to enable **Turn on INC (Independent Network Configuration) mode on self mode**.

7. Click **OK**. The **Nodes** page displays the local and remote nodes in your high availability configuration.

Discussion Question

Have you implemented high availability in your environment previously and what challenges or issues did you encounter? What questions do you have about configuring high availability for NetScaler?

Not for resale or distribution

Not for resale or distribution

Module 4

Integrating NS Gateway with Other Resources (Unified Gateway)

4

Not for resale or distribution

Not for resale or distribution

Integrating NetScaler with XenApp and XenDesktop Using the Unified Gateway Wizard

Overview

NetScaler Gateway communicates with Citrix services such as XenApp, XenDesktop, XenMobile. NetScaler Unified Gateway extends NetScaler Gateway connectivity allowing it to access any web, SaaS, or Cloud application through a single URL..

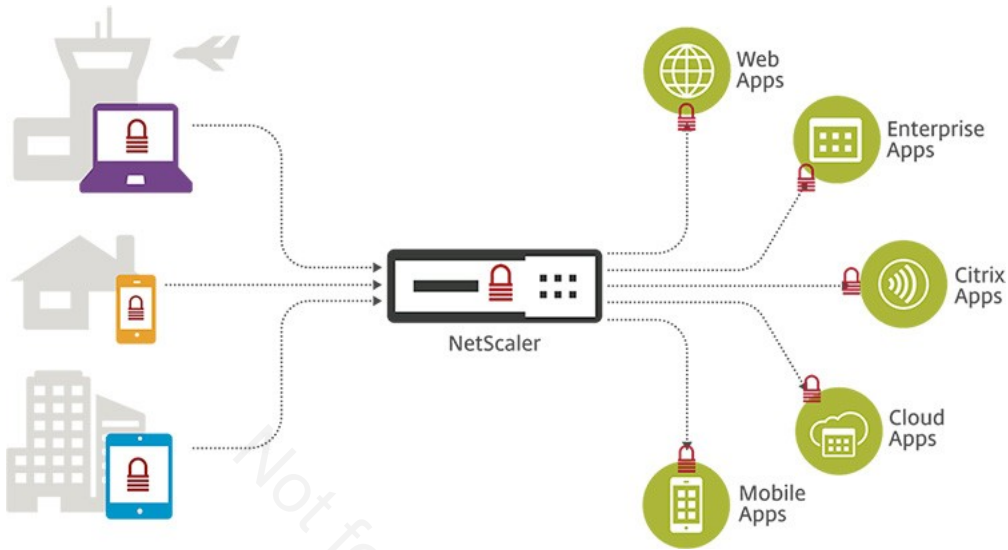
After completing this module, you will be able to:

- Configure the NetScaler Unified Gateway feature with Citrix StoreFront Services.
- Configure the NetScaler Unified Gateway feature with Citrix XenApp and Citrix XenDesktop.
- Understand how NetScaler Unified Gateway can work with other web, SaaS, and cloud services.

NetScaler Unified Gateway Key Features and Benefits

Citrix NetScaler Unified Gateway is a secure application, desktop, and data access solution that provides IT administrators granular application and device level policies to control access to corporate resources Users can connect from any device and from anywhere. Some of the features and benefits of NetScaler Unified Gateway include:

- Provides a single, secure, entry point to all resources behind the NetScaler.
- Securely proxies connections so that clients have no direct access to resources behind NetScaler
- Manages authentication at perimeter so that only authenticated traffic goes to internal network.
- Optimizes traffic which improves user experience and reduces server load.
- Can enable detailed auditing and monitoring to provide end to end network traffic analysis
- Granular control of SSL decryption and encryption.
- Granular security policy and access policy control at perimeter of network.



NetScaler Unified Gateway 11 New Features

NetScaler 11 includes many enhancements to Gateway functionality. A few of the enhancements include the following:

- NetScaler Gateway supports Windows 10
- Smart Control - Control ICA at the session level on NetScaler Gateway instead of at XenApp or XenDesktop
- Portal and EULA customization wizard.
- Gateway plug-in decoupling from Receiver - Gateway plugin behavior can be configured independently of Receiver.
- Enhanced Clustering Support - NetScaler Gateway in ICA Proxy Mode supports striped clustering.
- Granular control of SSL decryption and encryption.
- Granular security policy and access policy control at perimeter of network.

NetScaler Unified Gateway Prerequisites

Before you begin configuring NetScaler Unified Gateway, ensure that the following prerequisites are met:

- NetScaler is installed in your environment and has access to the network or networks. NetScaler is deployed in the perimeter network or internal network behind a firewall. You can also configure NetScaler in a double-hop perimeter network and for connections to a farm.

Route	Port	Details
Private network to perimeter network	80, 443, 1494, 2598	Ports 80 and 443 are used for management and administration of the NSIP through a browser. Port 443 is also used for the authentication callback URL to the NetScaler Gateway VIP from StoreFront or Web Interface. Ports 1494 and 2598 can be used for ICA or HDX traffic.
Perimeter network to private network	443, 80	Port 443 allows access to the StoreFront and port 80 allows access to XenApp or XenDesktop (if port 80 is used for secure ticketing authority traffic).
Perimeter network to private	389, 636	Ports 389 or 636 can be used for LDAP or Secure LDAP (LDAPS) authentication traffic to the LDAP servers.
Perimeter network to private	Appropriate Service Ports for application in question	Whatever ports the application uses. For example, TCP 8080 for a Tomcat application. HTTP(S) for a web application.

Unified Gateway with XenApp, XenDesktop, and XenMobile

We will cover this in greater detail in Module 10. Like NetScaler Gateway, Unified Gateway can interface with all Citrix services. Unified Gateway will provide a single access point to all Citrix services and other services.

Unified Gateway with Other Applications

Unified Gateway can proxy to almost any web application. Some of the common web applications include:

- Intranet Applications - NetScaler creates a custom URL to any internal web application. You must provide application's root relative URL and site strings. The site strings are derived from the real URL of the application in question.

- Software as a Service (SaaS) - Typically externally hosted applications such as ShareFile, SalesForce, Office 365, or SAP that require authentication. NetScaler Unified Gateway supports full VPN to these applications. Unified Gateway can also manage authentication and SSO. Can also manage SAML attributes when required. This requires configuring a SAML profile.
- A load balancing virtual server on the NetScaler - The web application can be a virtual server on the NetScaler itself.
- Clientless Access - Provides a client free access method to web services such as Outlook Web Access, or Sharepoint.

Client Connection Methods

Users can use any of the following methods to connect to network resources via NetScaler Unified Gateway:

- Citrix Receiver - Contains all Citrix plug-ins installed on the user device.
- Receiver for Web - allows user connections to applications, desktops, and ShareFile by using a Web browser.
- Worx Home - allows users to access WorxMail, WorxWeb and mobile apps from their iOS and Android devices.
- NetScaler Gateway App for iOS and Android
- NetScaler Gateway Plug-in for Windows, Mac OS X, or Linux
- NetScaler Gateway Plug-in for Java
- Clientless Access - provides basic access without installing user software

Unified Gateway Configuration Wizard

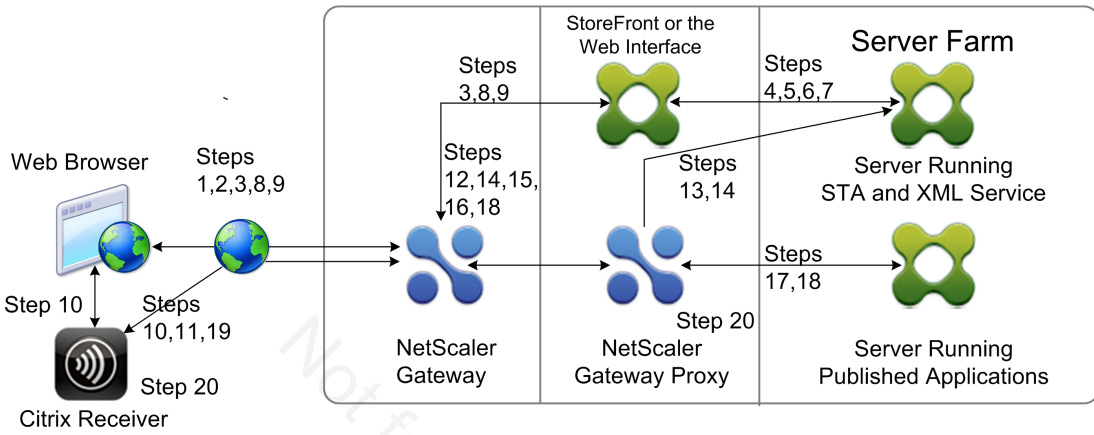
The configuration wizard manages the configuration of all the basic components to setup Unified Gateway:

- The Unified Gateway virtual server
- The SSL server certificate for the Gateway virtual server
- Authentication - Primary and Secondary
- Portal theme and optional customization
- The applications users will access via Unified Gateway such as XenApp or ShareFile.
- A non addressable VPM virtual server
- A content switching Virtual server with the Unified Gateway server as its target.

The configuration wizard does not manage basic NetScaler setup such as NetScaler owned IP addresses, VLANs, routes, licenses, and HA.:

Unified Gateway Traffic Flow in Double Hop DMZ

There are four communication stages involved:



- Authentication
- Session Ticket Creation
- Portal and EULA customization wizard.
- Initiation of Citrix Receiver
- Connection Completion

Authentication Traffic Flow

- User types the Gateway FQDN into a browser
- NetScaler Gateway in DMZ receives request
- NetScaler Gateway directs connection to the StoreFront server.
- StoreFront sends credentials to Citrix XML service
- Citrix XML service authenticates user
- Citrix XML Service creates a list of the published applications to which user is authorized to access and sends list to StoreFront server.

Session Ticket Creation

- StoreFront communicates with XML service and Secure Ticketing Authority (STA) to produce a ticket for each published application to which the user is authorized to access
- The STA then sends the requested session tickets to the Web Interface. Each session ticket includes an alias that represents the IP address of the server that hosts the published application, but not the actual IP address.
- The StoreFront server generates an ICA file for each of the published applications. The ICA file contains the ticket issued by the STA. The StoreFront server then creates and populates a web page with a list of links to the published applications and sends this web page to the web browser on the user device.

Starting Citrix Receiver

- Receiver initiates an ICA connection to NetScaler Gateway
- The NetScaler Gateway in the first DMZ communicates with the Secure Ticket Authority (STA) in the internal network to resolve the alias address in the session ticket to the real IP address of a computer running XenApp or StoreFront. This communication is proxied through the second DMZ by the NetScaler Gateway proxy
- The NetScaler Gateway proxy in the second DMZ passes the ticket validation request to the STA in the internal network. The STA validates the ticket and maps it to the computer running XenApp that hosts the published application.
- The STA sends a response to the NetScaler Gateway proxy in the second DMZ, which is passed to NetScaler Gateway in the first DMZ. This response completes the ticket validation and includes the IP address of the computer that hosts the published application.
- NetScaler Gateway in the first DMZ incorporates the address of the XenApp server into the user connection packet and sends this packet to the NetScaler Gateway proxy in the second DMZ.
- The NetScaler Gateway proxy in the second DMZ makes a connection request to the server specified in the connection packet.
- The server responds to the NetScaler Gateway proxy in the second DMZ. The NetScaler Gateway proxy in the second DMZ passes this response to NetScaler Gateway in the first DMZ to complete the connection between the server and NetScaler Gateway in the first DMZ.
- NetScaler Gateway in the first DMZ completes the SSL/TLS handshake with the user device by passing the final connection packet to the user device. The connection from the user device to the server is established.
- ICA traffic flows between the user device and the server through NetScaler Gateway in the first DMZ and the NetScaler Gateway proxy in the second DMZ.

Not for resale or distribution

Module 5

Basic Load Balancing

5

Not for resale or distribution

Not for resale or distribution

Basic Load Balancing Manual


Overview

NetScaler load balancing distributes end-user requests for web pages and other protected applications across multiple servers that host or mirror the same content. You use load balancing primarily to manage end user requests for heavily used applications, preventing poor performance and outages and ensuring that users can access your protected applications. Load balancing also provides fault tolerance; when one server that hosts a protected application becomes unavailable, the feature distributes end-user requests to the other servers that host the same application.

After completing this module, you will be able to:

- Explain basic load balancing concepts for a NetScaler system.
- Configure a basic load balancing setup.
- Explain why monitors are used in a load-balancing configuration.
- Identify the different load balancing methods.
- Identify the different session persistence methods.
- Verify load-balancing configuration.

Load-Balancing Basics




In a load-balancing configuration, the load-balancing virtual server is logically located between the client and the farm and manages traffic flow to the servers in the farm. On the NetScaler, the application servers are represented by virtual entities called services.

A load-balancing setup includes a load-balancing virtual server and multiple load-balanced application servers. The virtual server receives incoming client requests, uses the load-balancing algorithm to select an application server, and forwards the requests to the selected application server.

The load-balancing virtual server can use any of a number of algorithms, or methods, to determine how to distribute load among the load-balanced servers that it manages. The default load balancing method is the least connection method, in which the load-balancing NetScaler forwards each incoming client connection to whichever load-balanced application server currently has the fewest active user connections.

Entity Management



Entities in the NetScaler system are any configurable objects that are used with NetScaler features. Each entity is given a name when created. The entity name is used to associate the entity with a NetScaler feature. Entity names on the NetScaler system can have a maximum of 127 characters. The most commonly used entities are listed below.

Server

A server entity identifies a physical server and provides the IP address of the server. If you want to use the IP address of the server as the name of the server object, you can enter the IP address of the server when you create a service, and the server object is then created automatically. Alternatively, you can create the server object first and assign it an FQDN or other name and then specify that name instead of the IP address when you create the service.

Service	A service entity can be a logical representation of the application server itself or of an application running on a server that hosts multiple applications. A service is defined by an IP address, port, and protocol combination used to route requests to a specific load-balanced application server. The service identifies the type of traffic associated with a given server. You can configure multiple services for the same server. For example, you can configure a server to run HTTP, FTP and TCP services/applications. The NetScaler system directs traffic to the server using the appropriate service. When you create a service, you associate it with a server. For load balancing, you bind services to virtual servers. Based on these services, the virtual servers will then load-balance traffic across the available servers.
Service Group	A service group is a collection of services identified by IP address or server name. In a service group, any management changes made to the group are propagated to all members of the group.
Load-Balancing Virtual Server	A virtual server is an aggregated system entity that usually comprises multiple servers and services. Rather than traffic being routed directly to the server, it is sent to a virtual server, which then makes a decision about which server to forward the traffic to based on the services bound to the virtual server. The state of the virtual server determines whether the client requests are accepted. You need to specify the protocol, VIP and the port.
Monitor	A service monitor tracks a service and ensures that it is operating correctly. The monitor periodically performs a health check on each service that it is assigned. If the service does not respond within the time specified by the timeout, and a specified number of health checks fail, that service shows as DOWN. The NetScaler system then skips that service when performing load balancing until the issues that caused the service to stop responding are fixed.

Server Creation

Before a NetScaler system can begin load balancing for physical servers in the environment, you must first create server and service entities that represent these servers.

You can create servers explicitly as named entities before creating services. You then create the service and reference the existing server object. However, if you create the service first, the system automatically creates a server object based on the IP address used when configuring the service.

When you create a server, you must identify it by specifying its IP address or domain name. If you specify the domain name, you can later change the IP address of the physical server without having to modify the entry in the list of servers on the NetScaler. The domain name is resolved to an IP address that is specified in an address record on the DNS.

For more information about creating a server, see Citrix product documentation at <http://docs.citrix.com>.

Server Configuration Parameters

The following parameters will be used to configure a server:

name	The name assigned to the server. This alphanumeric string is required. The name must not exceed 127 characters, and the leading character must be a number or letter. The following characters are also allowed: at (@); underscore (_); dash (-); period (.); colon (:); and space ().
ipAddress	IP address of the server, in either IPv4 or IPv6 format. If the server is not reachable from the NetScaler system or is not active, the service shows as DOWN.
domain	Domain name that resolves to the IP address that represents the server.
ipv6Address	Resolve the domain name to an IPv6 address. Possible values: YES, NO. Default: NO.
state	The initial state of the server. Possible values: ENABLED, DISABLED. Default: ENABLED.
comment	A comment to help identify the server. Maximum length: 255 characters. To include spaces in a comment that you type on the NetScaler command line, enclose the entire comment inside quotation marks. The quotation marks become part of the comment. They are not required if you use the configuration utility.



To Create a Server Using the Configuration Utility

1. In the navigation pane, expand **Load Balancing**, and then click **Servers**.
2. In the details pane, click **Add**.
3. In the Create Server dialog box, specify values for the following parameters:
 - Server Name-name
 - IP Address-ipAddress
 - Domain Name-domain
 - Enable after Creating-state
 - Comment-comment
4. If you specify the domain name of the server and you want the domain name to be resolved to an IPv6 address, select the IPv6 Domain check box.
5. Click **Create** and then click **Close**.

The server you named appears in the Servers pane.

Service Monitoring

The purpose of service monitoring is to check the state of the services periodically. Monitors specify the types of requests sent to a service and the expected response from the service; this probe is known as a health check. If a service responds appropriately to the health check within the specified period of time, the state is marked as UP and requests continue to be directed to that service. If, however, a service fails a health check, the state is marked as DOWN and it no longer receives requests.

Each service defined on the NetScaler system requires some form of monitoring to ensure that requests are not sent to an unavailable service. Depending on the type of service, different kinds of monitoring may be required. Each monitor type requires its own set of parameters.

The function of a particular service must be considered when defining monitors. Is a completed three-way handshake enough to verify that the service is running? Does the service rely on other services in the enterprise to function? While the default monitors facilitate simple deployments, many environments are better served with the additional types of monitors available in the system.



The state of a service is maintained across high-availability pairs.

Services Configuration Overview

After you enable the load-balancing feature, you must create at least one service for each application server that is to be included in your load-balancing setup. The services that you configure provide the connections between the NetScaler and the load-balanced servers. Each service has a name and specifies an IP address, a port, and the type of data that is served. If you prefer to identify servers by name rather than IP address, you can create server objects and then specify the name of the server instead of its IP address when you create a service.

Services Creation

A service entity represents an application running on a server entity and is identified by a unique IP address/port combination. The service defines the traffic type between the NetScaler system and a server. Services must be bound to virtual servers before the NetScaler system is able to load balance incoming traffic between servers.

Before you create a service, you need to understand the different service types and how each is used. The following list describes the most common types of services supported on the NetScaler system. For more information about service types, see Citrix article CTX132359 on <http://support.citrix.com>.

ANY	Used for servers that accept any type of TCP, UDP, or ICMP traffic. The "ANY" parameter is used primarily with firewall load balancing and link load balancing.
DNS	Used for servers that accept DNS traffic, typically nameservers.
HTTP	Used for load-balanced servers that accept HTTP traffic, such as standard websites and web applications.
FTP	Used for servers that accept FTP traffic. You can also use TCP or "ANY" service types for FTP servers.
SSL	Used for servers that accept HTTPS traffic, such as e-commerce web sites and shopping cart applications. The SSL service type enables the NetScaler system to encrypt and decrypt SSL traffic for your secure web applications.

TCP	Used for servers that accept many different types of TCP traffic, or that accept a type of TCP traffic for which a more specific type of service is not available. You can also use the ANY service type for these servers.
SSL_TCP	Used for servers that accept non-HTTP-based SSL traffic to support SSL offloading.
SSL_BRIDGE	Used for servers that accept SSL traffic when you do not want the NetScaler system to perform SSL offloading.
UDP	Used for servers that accept UDP traffic. You can also use the ANY service type.
NNTP	NNTP (Network News Transfer Protocol) is the predominant protocol used by computer clients and servers for managing the notes posted on Usenet newsgroups.
SIP-UDP	Used for servers that accept UDP-based Session Initiation Protocol (SIP) traffic. SIP initiates, manages and terminates multimedia communications sessions and has emerged as the standard for Internet telephony (VoIP).
DNS-TCP	Used for servers that accept DNS traffic, where the NetScaler system acts as a proxy for TCP traffic sent to DNS servers.
RTSP	Used for servers that accept Real-Time-Streaming-Protocol (RTSP) traffic. RTSP provides delivery of multimedia and other streaming data.
DHCPRA	Used for servers that accept DHCP traffic. The DHCPRA service type can be used to relay DHCP requests and responses between VLANs.

DIAMETER	Used for load balancing Diameter traffic among multiple Diameter servers. Diameter uses message-based load balancing.
SSL_DIAMETER	Used for load balancing Diameter traffic over SSL.
MYSQL	Used for traffic associated with MySQL database servers.
MSSQL	Used for traffic associated with Microsoft SQL servers.
RDP	Used for remote desktop traffic.
RADIUS	Used for traffic associated with Remote Authentication Dial-In User Service (RADIUS). RADIUS supports combined authentication, authorization and auditing services for network management.

When the service is created, the service type is specified by identifying the protocol/traffic type in use. The different service types available on the NetScaler system provide the ability to handle the traffic appropriately.

The use of both server entities and service entities allows the application running on a server to be abstracted from the physical entity. This abstraction allows each service running on a single server to be treated independently of each other.

Services are designated as DISABLED until the NetScaler system connects to the associated load-balanced server and verifies that it is operational. At that point, the service is designated as ENABLED. Thereafter, the NetScaler system periodically monitors the status of the servers and places any that fail to respond to monitoring probes, called health checks, back in the DISABLED state until they respond.

Service Configuration Parameters

The following parameters will be used to configure services.

name	Name of the service. This alphanumeric string is required and cannot be changed after the service is created. The name must not exceed 127 characters and the leading character must be a number or letter. The following characters are also allowed: at (@); underscore (_); dash (-); period (.) colon (:); and space ().
-------------	---

serverName	Either the name of a previously created server object, or the IP address of the load-balanced server that hosts this service, in either IPv4 or IPv6 format. When you provide the IP address of the service, a server object is created with this IP address as its name. You can also create a server object manually and then select the server name instead of an IP address from the drop-down list box that is associated with this field. If the server is not reachable from the NetScaler or is not active, the service is designated as DOWN.
serviceType	The type of connections that the service will handle. Possible values: HTTP, SSL, FTP, TCP, SSL_TCP, UDP, SSL_BRIDGE, NNTP, DNS, "ANY", SIP-UDP, DNS-TCP and RTSP. Default: HTTP.
port	Port on which the service listens. The port number must be a positive number not greater than 65535.

To Create a Service Using the Configuration Utility

1. In the navigation pane, expand **Load Balancing**, and then click **Services**.
2. In the details pane, click **Add**.
3. In the Create Service dialog box, specify values for the following parameters:
 - Service Name-name
 - Server-serverName
 - Protocol-serviceType
 - Port-port
4. Click **Create** and then click **Close**.

The service you created appears in the Services pane.

Service Groups

Configuring a service group enables you to manage a group of services as easily as a single service. For example, if you enable an option such as compression, health monitoring, or graceful shutdown for a service group, the option is enabled for all the members of the service group.

After creating a service group, you can bind it to a virtual server and you can add services to the group. You can also bind monitors to service groups.

The members of a service group can be identified by IP address or server name. Using domain-name based (DBS) group members is advantageous because you need not reconfigure the member on the NetScaler system if the IP address of the member changes. The system automatically senses

such changes through the configured name server. This feature is particularly useful in cloud scenarios where the service provider can change a physical server or change the IP address for a service. If you specify a DBS group member, the NetScaler system learns the IP address dynamically.

You can bind both IP-based and DBS members to the same service group.



If you use DBS service group members, ensure that either a name server is specified or a DNS server is configured on the NetScaler system. A domain name will be resolved into an IP address only if the corresponding address record is present on the NetScaler system or the name server.

For more information about service groups, see Citrix article CTX132359 at <http://support.citrix.com>.

To Create a Service Group Using the Configuration Utility

1. In the navigation pane, expand **Load Balancing**, and then click **Service Groups**.
2. In the details pane, click **Add**.
3. In the Create Service Group dialog box, specify a name for the new service group.
4. In the **Protocol** list, select the protocol type.
5. Click **Create** and then click **Close**.

The service you created appears in the Service Groups pane.

Virtual Server Creation

A virtual server provides the client with access to the actual servers behind the NetScaler system. The client connects to a virtual server which consists of an IP address, port, and protocol combination that accepts incoming traffic. The virtual server is bound to a number of services running on servers. These services consist of the IP address and port of the server. The client request is routed from the virtual server to the server based on the service that handles the inbound request. DNS is not required to make an arbitrary round-robin load-balancing decision; instead, DNS returns the VIP address of the load-balancing virtual servers to the client.

Multiple load-balancing virtual servers can be configured on the same VIP address provided that the combination of the VIP address, port, and service are unique.

The virtual server is designated as DOWN until you bind the services that you created to it and until the NetScaler system connects to those services and verifies that they are operational. Only then is the virtual server designated as UP.

Virtual Server Configuration Parameters

The following parameters will be used to configure virtual servers.

name	Name of the virtual server. This alphanumeric string is required and cannot be changed after the virtual server is created. The name must not exceed 127 characters, and the leading character must be a number or letter. The following characters are also allowed: @ _ - period (.) colon (:) # and space ().
IPAddress	IP address of the virtual server. This IP address can be an IPv4 or IPv6 address and is usually a public IP address. Clients send connection requests to this IP address.
serviceType	The type of services to which the virtual server distributes requests. Possible values: HTTP, SSL, FTP, TCP, SSL_TCP, UDP, SSL_BRIDGE, NNTP, DNS, ANY, SIP-UDP, DNS-TCP, RDP and RTSP. Default: HTTP.
port	Port on which the virtual server listens for client connections. The port number must be between 0 and 65535.

To Create a Virtual Server Using the Configuration Utility

1. In the navigation pane, expand **Load Balancing**, and then click **Virtual Servers**.
2. In the details pane, click **Add**.
3. In the Create Virtual Server (Load Balancing) dialog box, specify values for the following parameters:
 - Name-name
 - IP Address-IPAddress
 - Protocol-serviceType
 - Port-port
4. Click **Create** and then click **Close**.

The virtual server you created appears in the Load Balancing Virtual Servers pane.

Binding Services or Service Groups to a Virtual Server

You can create most entities in the NetScaler system independently but they will not serve a function until they are bound to another related entity.

You can bind a service to a virtual server or service group on the NetScaler system by specifying the virtual server or service group and activating the service.

The state of the services bound to a virtual server determines the state of the virtual server: if all of the bound services are DOWN, the virtual server is marked DOWN and if any of the bound services is UP, the state of the virtual server is marked as UP.

To Bind a Service to a Load-Balancing Virtual Server Using the Configuration Utility

1. In the navigation pane, expand **Load Balancing**, and then click **Virtual Servers**.
2. In the details pane, select the virtual server for which you want to bind the service.
3. Click **Open**.
4. In the Configure Virtual Server (Load Balancing) dialog box, click the **Services** tab, then select the **Active** check box next to the service that you want to bind to the virtual server.
5. Click **OK**.

You can bind a service to multiple virtual servers.

To Bind a Service Group to a Virtual Server Using the Configuration Utility

1. In the navigation pane, expand **Load Balancing**, and then click **Virtual Servers**.
2. In the details pane, select the virtual server to bind the service group, and then click **Open**.
3. In the Configure Virtual Server (Load Balancing) dialog box, click the **Services Groups** tab.
4. In the Active column, select the check box next to the service group to bind to the virtual server and click **OK**.

Traffic Types

You can configure a load-balancing virtual server to support any number of traffic types. You can also configure multiple load-balancing servers to support different services on the same virtual IP address.

Application Protocols

HTTP, FTP, DNS, NNTP, HTTPS, MSSQL, MYSQL, RDP, DIAMETER and RADIUS are all application-layer protocols for which switching is performed at the request level rather than at the connection level. Request-level switching allows greater granularity in the load-balancing process. The SSL setting represents HTTPS and instructs the virtual server to terminate and decrypt the HTTPS connection. You can configure a load-balancing server as type SSL and associate it with a service of type HTTP. The traffic from the end user is encrypted to the NetScaler system, and plain text HTTP traffic is forwarded from the NetScaler system to the server.

Session Protocols

You can associate the load-balancing server with a session protocol, either TCP or UDP. You can also secure TCP-based protocols other than HTTP using SSL. If the incoming traffic is SSL encrypted but not HTTP traffic, a virtual server of type SSL_TCP is created. This server decrypts the traffic on arrival and forwards it based on the protocols defined on the services bound to it.

General Traffic

You use the "ANY" traffic type for any TCP, UDP and Internet Control Message Protocol (ICMP) service. The configuration of ANY is used with some firewall load balancing and link load-balancing configurations where load balancing is time based.

The following table provides an overview of the load-balancing traffic types.

Application protocols	Session protocols	General traffic type
<ul style="list-style-type: none">• HTTP• FTP• DNS• NNTP• SSL (HTTPS)• RTSP• SSL_BRIDGE• MSSQL• MYSQL• MSSQLHTTP• RDP• DIAMETER• SSL_DIAMETER• RADIUS• DHCPRA• PUSH• SSL_PUSH	<ul style="list-style-type: none">• TCP• UDP• SSL_TCP• SIP_UDP• DNS_TCP	ANY

Built-in Monitors

The NetScaler system contains a number of built-in monitors for monitoring services. These built-in monitors handle most of the common protocols. You cannot modify or remove the built-in monitors; you can only bind a built-in monitor to a service and unbind it from the service.

For a complete list of built-in monitors and more information about them, see Citrix article CTX132359 at <http://support.citrix.com>.

TCP-based Applications

The NetScaler system has two built-in monitors that monitor TCP-based applications: tcp-default and ping-default. The tcp-default monitor is bound to all TCP services; the ping-default monitor is bound to all non-TCP services.

SSL Services

The NetScaler system has both a TCPS and HTTPS default secure monitor. You can use the secure monitors to monitor HTTP as well as non-HTTP traffic. The secure monitors work as described below:

- **TCPS:** The NetScaler system establishes a TCP connection. After the connection is established, the system performs an SSL handshake with the server. After the handshake is complete, the system closes the connection.
- **HTTPS:** The NetScaler system establishes a TCP connection. After the connection is established, the system performs an SSL handshake with the server. When the SSL connection is established, the system sends HTTP requests over the encrypted channel and checks the response codes.

FTP Services

To monitor FTP services, the NetScaler system opens two connections to the FTP server. It first connects to the control port, which is used to transfer commands between a client and an FTP server. After it receives the expected response, it connects to the data port, which is used to transfer files between a client and an FTP server. When the FTP server responds as expected on both connections, the status shows as UP.

The NetScaler system has two default monitors for FTP services: the FTP monitor and the FTP-EXTENDED monitor. The FTP monitor checks basic functionality; the FTPEXTENDED monitor also verifies that the FTP server is able to transmit a file correctly.

RADIUS Services

The NetScaler system RADIUS monitor periodically checks the state of the RADIUS service to which it is bound by sending an authentication request to the service. The RADIUS server authenticates the RADIUS monitor and sends a response. By default, the monitor expects to receive a response code of 2, the default Access-Accept response, from the RADIUS server. If the monitor receives the appropriate response, it marks the status as UP.

DNS and DNS-TCP Services

The NetScaler system has two default monitors that can be used to monitor DNS services: DNS and DNS-TCP. When bound to a service, either monitor periodically checks the state of that DNS service by sending a DNS query to it. The query resolves to an IPv4 or IPv6 address. That IP address is then checked against the list of test IP addresses that you configure. The list can contain as many as five IP addresses. If the resolved IP address matches at least one IP address on the list, the DNS service is marked as UP. If the resolved IP address does not match any IP addresses on the list, the DNS service is marked as DOWN.

LDAP Services

The NetScaler system has one default monitor that can be used to monitor LDAP services: the LDAP monitor. It periodically checks the LDAP service to which it is bound by authenticating and sending a search query to it. If the search is successful, the service is shown as UP. If the LDAP server does not locate the entry, a failure message is sent to the LDAP monitor and the service is shown as DOWN.

XML Broker Services

The built-in monitor type, CITRIX-XML-SERVICE, can be used to create monitors to monitor the XML Broker services. The XML Broker services are used by Citrix XenApp. The monitor opens a connection to the service and periodically probes the XML services to which it is bound. If the server responds as expected within the configured time period, the monitor marks the service as UP. If the service does not respond, or responds incorrectly, the monitor marks the service as DOWN.

To configure a CITRIX-XML-SERVICE monitor, you need to specify the application name in addition to setting the standard parameters. The application name is the name of the application that has to be run to monitor the state of the XML Broker service. The default application is Notepad.

Monitor Parameters

Two types of parameters can be configured for monitors:

- Parameters that are standard and apply to all monitors, regardless of type
- Parameters that are specific to the type of monitor being defined

Standard Parameters

Standard parameter types apply to all load-balancing monitors, regardless of type. Each of the parameters can be set independently for individual monitors. For example, an HTTP monitor can be configured with a response timeout of 15 seconds, while another HTTP monitor can be configured with a response timeout of 10 seconds.

The following standard parameters can be applied to any monitor, regardless of type:

Interval

The interval parameter specifies the frequency at which the health-check probe is sent to a service.

Response Timeout

The response-timeout parameter specifies the interval of time for which the system waits before it marks the health check probe as FAILED. As long as the server responds within the response timeout period, the server state will remain UP, and the service will participate in load balancing. If a service fails the health check by not responding within the response timeout period, the monitor failed count is incremented. An increment of the monitor failed count constitutes one of a number of retries as specified in the retries parameter. A successful response resets the monitor failed count to zero.

Down Time

The down-time parameter specifies the duration for which the system waits before sending the next health-check probe once a service is marked as DOWN. The default down-time value is 30 seconds, except for LDNS monitors, which have a default down time value of 20 seconds.

Retries

The retries parameter specifies the number of consecutive health-check probe failures after which the system marks the service as DOWN. The default retries value is 3.

Success Retries

A monitor periodically evaluates the state of services to indicate whether it is UP or DOWN. If the monitor did not receive a response in the configured time and if the number of retries fails, the service is marked as DOWN. The monitor probes may fail intermittently. The monitor evaluates the state of the service as DOWN, making it unavailable. When you configure the success-retries parameter, the monitor marks the state of the service as UP only when the monitor probe succeeds for a consecutive number of retries.

Reverse

The reverse parameter specifies whether a monitor is configured for reverse conditions. When the reverse parameter is set to YES, a health check probe will fail if the condition of the monitor is satisfied. In normal conditions, a probe will fail when the conditions of a monitor are not met. The reverse parameter is disabled by default.

Secure

The secure parameter allows the NetScaler system to monitor SSL services. When the state of the parameter is set to YES, the NetScaler system establishes a TCP connection with the destination of the monitor and then performs an SSL handshake with the server. The secure parameter is disabled by default.



Executing the show service command for a service monitor displays:

- The total number of probes made by the monitor.
- The total number of failed responses.
- The current contiguous failed response count.
- The result of the most recent response.

Creating Monitors

You can create a monitor on the NetScaler system by specifying the monitor name, monitor type, and values for the appropriate parameters.

In the configuration utility, expand the **Load Balancing > Monitors** node.

In the command-line interface, type:

```
add lb monitor <monitorName> -type <monitorType> [-  
interval <integer [units]>]  
[-resptimeout <integer [units]>][-Retries <integer>] [-  
successRetries <integer>  
[-failureRetries <integer>] [-downTime <integer [units]>] [-  
reverse ( YES | NO ) ]  
[-transparent ( YES | NO ) ] [-secure ( YES | NO ) ]
```



The default response timeout unit is seconds. To change the unit to milliseconds, seconds, or minutes, the optional values of msec, sec, or min can be appended following the integer value.

HTTP Monitoring

HTTP monitors are used to verify the health of an HTTP service. The following process describes how HTTP monitoring occurs on the NetScaler system:

1. The NetScaler system establishes a TCP connection with the service destination specified by the monitor.
2. The NetScaler system sends an HTTP request to the service.
3. The NetScaler system compares the received response to a set of acceptable response codes that were configured in the monitor parameters.
4. If the response matches any configured responses, the probe is a success. If the response does not match any configured responses, the probe fails.



If the response code parameter is left empty, any response from the service is considered a match. The NetScaler system looks for matching response codes in the first 24 KB of data in the body of the response.

The parameters for the HTTP monitor can be configured as follows:

HTTP Request

The HTTP request parameter specifies the HTTP request to be sent to the service bound to the monitor. Default value: HEAD /.

Response Codes

The response codes parameter specifies a set of HTTP response codes expected from the service bound to the monitor. Default value: 200.

Extended Application Verification (EAV) Monitoring

You can monitor service dependencies using the following built-in monitor types for certain protocols and applications:

- RADIUS
- MYSQL
- SNMP
- LDAP
- POP3
- NNTP
- CITRIX-XML-SERVICE
- CITRIX-WEB-INTERFACE
- STOREFRONT



The CITRIX-XML-SERVICE and CITRIX-WEB-INTERFACE monitor types check Citrix XenApp services. These monitors are used in datacenters where XenApp servers are deployed.

Extended Content Verification (ECV) Monitoring

Extended content verification (ECV) monitors are used for verifying content in HTTP, TCP and UDP payload. ECV monitors are considered kernel monitors because they are basic, high-performance probes that originate from the BSD kernel. General limitations of ECV monitors include:

- One request/response
- 127 characters in the probe request
- Only the first 24 KB of the probe response parsed

For HTTP and TCP services, predefined Extended Content Verification (ECV) monitors are available. For ECV monitors, it is not enough to see that a TCP connection was accepted; some particular reply in the connection is required to mark the service as UP. For these monitors, a request string is configured along with an expected reply string. If the reply string received by the monitor matches the expected string, then the service is marked as UP.

HTTP-ECV and TCP-ECV Monitoring Process

The -secure (YES | NO) option allows the NetScaler system to monitor SSL services. TCP-ECV and HTTP-ECV probes include:

- TCP connection
- SSL handshake

- Encrypted server data (TCP-ECV probes only)
- Encrypted HTTP request (HTTP-ECV probes only)

The process describes how an HTTP-ECV or TCP-ECV monitor performs a health-check probe:

1. The NetScaler system establishes a TCP connection with the service destination specified by the monitor.
2. The NetScaler system sends data specified in the send string parameter to the service.
3. The NetScaler system compares the HTTP response or data received by the service to the expected response specified by the receive string parameter.
4. The probe is a success if the response matches the data in the receive string parameter. The probe fails if the response does not match.



If the receive string parameter is left empty, any response from the service is considered a match. The NetScaler system looks for matching responses in the first 24 KB of data in the body of the response.

Reverse Condition Monitoring

The reverse parameter specifies whether a monitor is configured for reverse conditions. When the reverse parameter is set to YES, a health-check probe will fail if the condition of the monitor is satisfied. In normal conditions, a probe will fail when the conditions of a monitor are not met.

For example, an HTTP-ECV monitor is configured with a send string of GET /FILE, a receive string of ERROR and the reverse parameter is set to YES. If the NetScaler system sends a probe that returns a response with the string ERROR, the probe will fail. If the response does not match ERROR, the probe is successful.

The reverse parameter is disabled by default.

Setting Monitor Thresholds

The NetScaler system uses monitor thresholds to determine the state of services. The threshold value is used to minimize the impact of spikes in load from shutting down a service.

The threshold value from the monitor specifies a value that determines the state of the service as UP. The NetScaler system determines the state of the service as UP only when the sum of the monitors that are UP is equal to or greater than the threshold value that you configured on the service.

In the configuration utility, expand the **Load Balancing Services** node.

In the command-line interface, type:

```
set service <ServiceName> -monThreshold <Value>
```

Custom Monitors

In addition to built-in monitors, you can create custom monitors, either based on the built-in monitors or from scratch, to check the state of your services. The NetScaler system provides several types of custom monitors based on scripts that are included with the NetScaler operating system that can be used to determine the state of services based on the load on the service or network traffic sent to the service. These are the inline monitors, end-user monitors and load monitors.

With any of these types of monitors, you can use the supplied functionality, or you can create your own scripts and use those scripts to determine the state of the service to which the monitor is bound.

For more information about custom monitors, including inline monitors, end-user monitors and load monitors, see Citrix article CTX132359 at <http://support.citrix.com>.

XenDesktop Delivery Controller Monitoring

NetScaler provides a built-in monitor, CITRIX-XD-DDC monitor, which monitors Delivery Controller servers. In addition to the health check, you can also verify whether the probe is sent by a valid user of the Delivery Controller server.

The monitor sends a probe to the Delivery Controller server in the form of an XML message. If the Delivery Controller server responds to the probe with the identity of the farm, the probe is considered to be successful and the server's status is marked as UP. If the HTTP response does not have a success code or the identity of the farm is not present in the response, the probe is considered to be a failure and the status of the server is marked as DOWN.

If you use the wizard for configuring the load balancing of XenDesktop servers, the XD-DDC monitor is automatically created and bound to the DDC services. If you do not use the wizard, add a monitor of the type XD-DDC.

To add an XD-DDC monitor with validate credentials option by using the command-line interface, type the following command:

```
add lb monitor monitorName monitorType -userName userName -  
password password  
-ddcDomain ddc_domain_name -validateCred YES
```

StoreFront Store Monitoring

You can configure an end-user monitor for a Citrix StoreFront store. The monitor determines the state of the StoreFront store by successfully probing the account service, authentication service and discovery document (in that order). If any of those services do not respond to the probe, the monitor probe fails and the StoreFront store is marked as DOWN. The monitor sends probes to the IP address and port of the bound service.

You can also bind a StoreFront monitor to a service group. A monitor is bound to each member of the service group and probes are sent to the IP address and port of the bound member (service). Also, because each member of a service group is now monitored by using the member's IP address, you can now use the StoreFront monitor to monitor StoreFront cluster nodes that are added as members of the service group.

The `secure` parameter is used to determine whether to use HTTP (default) or HTTPS to send monitor probes. To use HTTPS, set the `secure` option to YES.

To create a StoreFront monitor by using the command-line interface, type the following command:

```
add lb monitor monitorName STOREFRONT string -storeName string
[-storefrontacctservice (YES | NO)] -secure (YES | NO)
```

TFTP Server Monitoring

NetScaler contains native support for load balancing Trivial File Transfer Protocol (TFTP) servers and monitoring for a known payload. Complete the following steps to set up load balancing for TFTP servers.

1. Create the services for the TFTP servers.
2. Go to **Configuration > Traffic Management > Load Balancing > Virtual Servers** and then click **Add**.
3. Specify the parameters for the TFTP server name, protocol (TFTP), IP address, and port.
4. Click **Create** and **Close** to continue.
5. Verify the status of the virtual server. It should display as UP with two services bound. Continue with the next step.
6. Go to the DHCP administration console and modify the scope to include the VIP as part of the configuration.
7. Test the setup by inspecting the boot log and verifying that the TFTP VIP is included in the DHCP reply.

Load-Balancing Methods

Load-balancing decisions for incoming traffic are made based on the load-balancing method assigned to the service. The following sections provide available load-balancing methods. For more information about load-balancing methods, also called load-balancing algorithms, see Citrix article CTX132359 at <http://support.citrix.com>.

The following are common load-balancing methods for implementing NetScaler in a XenDesktop or XenApp environment:

- Least Connections
- Round Robin

Least Connections

Least connections is the default load-balancing method for the NetScaler system and is the most accurate method to distribute load intelligently. The least-connection calculation used to determine the service with the fewest number of connections is calculated differently depending on the service type. The following table provides an overview of the services used by the least-connections load-balancing method.

Service	Description
TCP, HTTP, HTTPS, SSL_TCP	Uses established, active connections to a service and connections to a service waiting in the Surge Queue for the least connections calculation.
UDP	Uses all sessions between the client and the physical server for the least connections calculation. These sessions are logical, time-based entities and are created on the UDP packet that arrives first. If weights are configured, they are taken into account when server selection is completed.

Round Robin

Round robin is the simplest load-balancing method, distributing traffic based on a server-rotation system, regardless of load.

An incoming request is sent to Server 1, the next request is sent to Server 2, the next request is sent to Server 3. When all servers have received a request, the cycle begins again.

This method is sufficient if all requests result in the same load on servers, but in most cases, a more robust load-balancing method based on metrics should be used.

Diameter Load Balancing

The Diameter protocol is a next generation Authentication, Authorization and Accounting (AAA) signaling protocol used mainly on mobile devices such as laptops and mobile phones. It is a peer-to-peer protocol, as opposed to the traditional client-server model used by most other protocols. However, in most Diameter deployments, the client originates the request and the server responds to the request. When Diameter messages are exchanged, the Diameter server usually does much more processing than does the Diameter client.

With the increase in control plane signaling volume, the Diameter server becomes a bottleneck. Therefore, Diameter messages must be load balanced to multiple servers. A virtual server performing load balancing of Diameter messages provides the following benefits:

- Lighter load on Diameter servers, which translates to faster response times to end users
- Server health monitoring with better failover capabilities
- Better scalability in terms of server addition without changing client configuration
- High availability
- SSL Diameter offloading

For more information about diameter load balancing, see Citrix product documentation at <http://docs.citrix.com>.

Service Weights

In a load-balancing configuration, you assign weights to services to indicate the percentage of traffic that should be sent to each service. Services with higher weights can handle more requests; services with lower weights can handle fewer requests. Assigning weights to services allows the NetScaler system to determine how much traffic each load-balanced server can handle and therefore more effectively balance load.

Service weights can be configured for the following load-balancing methods:

- Least Connections
- Round Robin
- Least Bandwidth
- Least Packets
- Least Response Time

To Configure a Virtual Server to Assign Weights to Services Using the Configuration Utility

1. In the navigation pane, expand **Load Balancing**, and then click **Virtual Servers**.
2. In the details pane, select the virtual server, and click **Open**.
3. On the **Services** tab, in the **Weights** spin box, type or select the weight to assign to the service.
4. Click **OK**.

Persistence and Persistence Connections

Unless you configure persistence, a load-balancing stateless protocol, such as HTTP, disrupts the maintenance of state information about client connections. Different transmissions from the same client might be directed to different servers even though all of the transmissions are part of the

same session. You must configure persistence on a load-balancing virtual server that handles certain types of web applications, such as shopping cart applications.

Before you can configure persistence, you need to know the difference between the different types of persistence, how they are used, and what the implications of each type is. You then need to configure the NetScaler system to provide persistent connections for those websites and web applications that require them.

You can also configure backup persistence, which takes effect in the event that the primary type of persistence configured for a load-balancing virtual server fails. You can configure persistence groups so that a client transmission to any virtual server in a group can be directed to a server that has received previous transmissions from the same client.

Session Persistence Methods

Session persistence methods are determined based on the method you assign to the service. The following sections provide information about the most commonly used persistence methods.

Cookie-Insert Persistence

Cookie-insert persistence inserts an HTTP cookie in the client responses. The cookie is then used to direct subsequent requests from the client to the same physical server that handled the initial request. The cookie contains the IP address and port number of the service to which the client request is directed.



If a client is unable to store a cookie, subsequent requests from that client will not have a cookie, and session persistence will not be honored. A load-balancing decision will be made based on the configured method.

The timeout value for cookie-insert persistence is the period of inactivity for a session. If you set the timeout value to zero, the session never expires in the NetScaler system; however, the expiration time is dependent on the client software used and cookies usually expire when the client software is closed.



If cookie-insert is configured as the primary persistence method and source-IP-address is configured as the backup, you can set different timeout values for each method.

Backup persistence is the method used when the primary persistence method fails. When you use cookie-insert persistence, you can configure source-IP-address persistence as the backup persistence method to maintain persistence sessions. For example, if a client is using a web browser that does not support cookies, persistence is based on the source IP address of the client instead.

You can also specify a name for a persistent cookie for a load-balancing virtual server. To specify the name from the command-line interface, type the following command:

```
set lb vserver name -cookieName string
```

Cookie Version

The NetScaler system uses two HTTP cookie versions:

HTTP cookie version 0 If HTTP cookie version 0 is selected, the system inserts the absolute GMT time of the cookie expiration time that is calculated as the sum of the current GMT time on the system and the timeout value. HTTP cookie version 0 is the default setting on the NetScaler system.

HTTP cookie version 1 If HTTP cookie version 1 is selected, the system inserts a relative expiration time. The client software calculates the actual expiration time.

Source-IP-Address Persistence

Source-IP-address persistence routes a client request based on the configured load-balancing method and then directs all subsequent requests from the same IP address to the physical server that handled the initial request.



Traffic passing through a NAT device will appear to originate from the same IP address. If most client requests flow through a NAT device, source-IP-address persistence might not be the most efficient load-balancing method because every request will be directed to the same physical server even though the requests originate from multiple clients.

Source-IP-address persistence sessions are stored in a persistence table, which consumes system resources. The maximum number of persistence sessions per core on an nCore NetScaler appliance is 1,000,000 (1 million). The maximum number of persistence sessions that can coexist on an nCore NetScaler appliance is equal to the product of the number of cores on the appliance and the per-core limit. For example, if the appliance has 6 CPU cores, the maximum number of persistence sessions that can coexist on the appliance is 6,000,000 ($6 * 1000000$).

The lower limit for the persistence sessions setting, per core, is 150,000. The lower limit for an nCore appliance is equal to the product of the number of cores on the appliance and the lower limit for each core. For example, the lower limit for the persistence sessions setting for an nCore NetScaler appliance with 6 cores is 900,000 ($6 * 150000$).

The timeout value for source-IP-address persistence is the period of inactivity for a session. Once the timeout value is reached, the session is discarded, and server selection resumes based on the configured load-balancing method.

SSL Session ID Persistence

SSL-session-ID persistence is based on the arriving SSL session ID in a request, which is part of the SSL handshake process. Requests containing the same SSL session ID are directed to the physical server that processed the initial request.

SSL-session-ID persistence sessions are stored in a persistence table and therefore consume system resources. A maximum of 256,000 entries can be stored in the persistence table at any one time.

The timeout value for SSL-session-ID persistence is the period of inactivity for a session. Once the timeout value is reached, the session is discarded and server selection resumes based on the configured load-balancing method.



SSL-session-ID persistence is not generally recommended as many web browsers renegotiate the SSL session as often as every two minutes. In such a case, the persistence is often broken. The SSL session is generally not long enough to use as a persistence criteria for typical web sessions.

Persistence Tables

Session persistence information for each session is stored on the NetScaler system in a persistence table. The following information is displayed in the table.

- Persistence type
- Source IP address
- Destination IP address
- Destination port
- Virtual server name
- Persistent session timeout
- Reference count
- SIP CALLID

Persistence Group Configuration

When you have load-balanced servers that handle several different types of connections (such as web servers that host multimedia), you can configure a virtual server group to handle these connections. To create a virtual server group, you bind different types of virtual servers, one for each type of connection that your load-balanced servers accept, into a single group. You then configure a persistence type for the entire group.

You can configure either source-IP-based persistence or HTTP cookie-based persistence for persistence groups. After you set persistence for the entire group, you cannot change it for individual virtual servers in the group. If you configure persistence for a group and then add a new

virtual server to the group, the persistence of the new virtual server is changed to match the persistence setting of the group.

When persistence is configured on a group of virtual servers, persistence sessions are created for initial requests and subsequent requests are directed to the same service as the initial request, regardless of the virtual server in the group that receives each client request.

For more information about creating and modifying a virtual persistency group, see Citrix article CTX132359 at <http://support.citrix.com>.

You can also view the persistence sessions that are in effect globally or for a particular virtual server.

To View Persistence Sessions Using the Configuration Utility

1. In the navigation pane, click **Load Balancing**.
2. In the details pane, under **Monitor Sessions**, click **Virtual Server persistence sessions**.

Clearing of Persistence Sessions

You might need to clear persistence sessions from the NetScaler if sessions fail to time out. You can do one of the following:

- Clear all sessions for all virtual servers at once.
- Clear all sessions for a given virtual server at once.

To Clear Persistence Sessions Using the Configuration Utility

1. In the navigation pane, click **Load Balancing**.
2. In the details pane, under **Monitor Sessions**, click **Clear persistence sessions**.
3. In the **Clear Persistence Sessions** dialog box, do one of the following:
 - If you want to clear all sessions for all virtual servers on the system, in **Virtual Server**, select **All Virtual Servers**.
 - If you want to clear all sessions for a given virtual server, in **Virtual Server**, select the virtual server.
4. Click **OK**.

Load Balancing Configuration Protection

When a load-balancing virtual server fails, or when the virtual server is unable to handle excessive traffic, the load-balancing setup can fail. You can protect your load-balancing setup against failure by configuring the NetScaler system to redirect excess traffic to an alternate URL, configuring a backup load-balancing virtual server and configuring stateful connection failover.

Disabling Services

Services are enabled by default when you create them. You can disable or enable each service individually. When disabling a service, you normally specify a wait time, in seconds, during which the service continues to handle established connections, but rejects new ones, before shutting down. If you do not specify a wait time, the service shuts down immediately. During the wait time, the service's state is OUT OF SERVICE.

To Disable a Service Using the Configuration Utility

1. In the navigation pane, expand **Load Balancing**, and then click **Services**.
2. In the details pane, select the service that you want to disable, and then click **Disable**.
3. In the **Wait Time** dialog box, type the wait time in seconds after which the service is to be disabled (for example, **30**).
4. Click **OK**.

Graceful Shutdown of Services

If you want to disable a service only when all the established connections are closed by the server or the client, you can use the graceful shutdown option.

During scheduled network outages such as system upgrades or hardware maintenance, you might have to close or disable some services. To avoid disrupting sessions that have already been established, you can specify a wait time, which places a service in the transition-out-of-service (TROFS) state until the specified wait time expires. The service then enters the out-of-service (OFS) state.

Often, however, you cannot estimate the amount of time needed for all the connections to a service to complete the existing transactions. If a transaction is unfinished when the wait time expires, shutting down the service can result in data loss. In this case, you can specify graceful shutdown for the service, so that the service is disabled only when all the established connections are closed by either the server or the client.

Persistence is maintained according to the specified method even if you enable graceful shutdown. The system continues to serve all the persistent clients, including new connections from the clients, unless the service is marked DOWN during the graceful shutdown state as a result of the checks made by a monitor.

For more information about graceful shutdown of services, see Citrix article CTX132359 at <http://support.citrix.com>.

Removing Services

You can remove a service when it is no longer used. When you remove a service, it is unbound from its virtual server and deleted from the NetScaler configuration.

To Remove a Service Using the Configuration Utility

1. In the navigation pane, expand **Load Balancing** and then click **Services**.
2. In the details pane, select the service that you want to remove and then click **Remove**.
3. In the Remove dialog box, click **Yes**.

Configuration Verification

After finishing your basic configuration, you should view the properties of each service and load-balancing virtual server in your load-balancing setup to verify that each is configured correctly. After the configuration is running, you should view the statistics for each service and load-balancing virtual server to check for possible problems.

Viewing the Properties of a Server Object

You can view properties such as the name, state and IP address of any server object in your NetScaler configuration.

- To view the properties of server objects by using the NetScaler command-line interface, type:

```
show server <serverName>
```
- To view the properties of server objects by using the configuration utility, in the navigation pane, expand **Load Balancing**, then click **Servers**. The parameter values of the available servers appear in the details pane.

Viewing the Properties of a Virtual Server

You can view properties such as the name, state, effective state, IP address, port, protocol, method and number of bound services for your virtual servers. If you have configured more than the basic load-balancing settings, you can view the persistence settings for your virtual servers, any policies that are bound to them and any cache redirection and content-switching virtual servers that have been bound to the virtual servers.

- To view the properties of a load-balancing virtual server by using the NetScaler command-line interface, type:

```
show lb vserver <name>
```

- To view the properties of a load-balancing virtual server by using the configuration utility, in the navigation pane, expand **Load Balancing** and then click **Virtual Servers**. In the details pane, click a virtual server to display its properties at the bottom of the details pane.

Viewing the Properties of a Service

You can view the name, state, IP address, port, protocol, maximum client connection, maximum requests for each connection and server type of the configured services; you can use this information to troubleshoot any mistake in the service configuration.

- To view the properties of services by using the NetScaler command-line interface, type:

```
show service <name>
```

- To view the properties of services by using the configuration utility, in the navigation pane, expand **Load Balancing** and then click **Services**. The details of the available services appear on the Services pane.

Viewing the Bindings of a Service

You can view the list of virtual servers to which the service is bound. The binding information also provides the name, IP address, port and state of the virtual servers to which the services are bound. You can use the binding information to troubleshoot any problem with binding the services to virtual servers.

- To view the bindings of a service by using the NetScaler command-line interface, type:

```
show service bindings <name>
```

- To view the bindings of a service by using the configuration utility, in the navigation pane, expand **Load Balancing** and then click **Services**. In the details pane, select the service whose binding information you want to view. Click **Show Bindings**. The bindings of the service you selected appear in the ServiceName dialog box.

The Load-Balancing Visualizer

The Load-Balancing Visualizer is a tool that you can use to view and modify the load-balancing configuration in graphical format.

You can use the visualizer to view the following:

- The services and service groups that are bound to a virtual server

- The monitors that are bound to each service
- The policies that are bound to the virtual server
- The policy labels, if configured
- Configuration details of any displayed element
- Load-balancing virtual server statistics
- Statistical information such as the number of requests received each second by the virtual server and the number of hits each second for rewrite, responder and cache policies
- A comparative list of all the parameters whose values either differ or are not defined across service containers

You can also use the Visualizer to add and bind new objects, modify existing ones, and enable or disable objects.



The Visualizer requires a graphic interface, so it is available only through the configuration utility.

For more information about the load-balancing virtualizer, see Citrix article CTX132359 at <http://support.citrix.com>

Discussion Question

Which load-balancing method do you or your organization most commonly use? Which methods do you plan to implement and why?

Not for resale or distribution

Module 6

NetScaler Security and SSL

6

Not for resale or distribution

Not for resale or distribution

SSL Offload Manual

Overview

Secure Sockets Layer (SSL) is a session-layer encryption and authentication protocol used with HTTP and other protocols to secure communications between clients and servers. SSL, which is widely used to support e-commerce websites, specifies a method for a client and a server to exchange cryptographic information and to build a secure channel through which to send requests and responses.

To establish a secure SSL connection, a digital certificate must be issued to your organization by a Certification Authority (CA). That digital certificate contains a public key which is used to establish a secure connection using the SSL protocol.

In addition, the processing of SSL transactions by a web server is CPU-intensive enough to negatively affect web application performance. Configuring SSL offload on the NetScaler can free the web servers from handling the SSL encryption and decryption tasks and maintain security without degrading web server performance. SSL offloading can also be used to optimize traffic while maintaining the security of the private data.

After completing this module, you will be able to:

- Identify requirements to secure communications with SSL certificates.
- Create and upload an SSL certificate.
- Bind an SSL certificate key.
- Identify common virtual SSL server deployments.
- Configure advanced SSL options.
- Create appropriate servers and virtual servers.

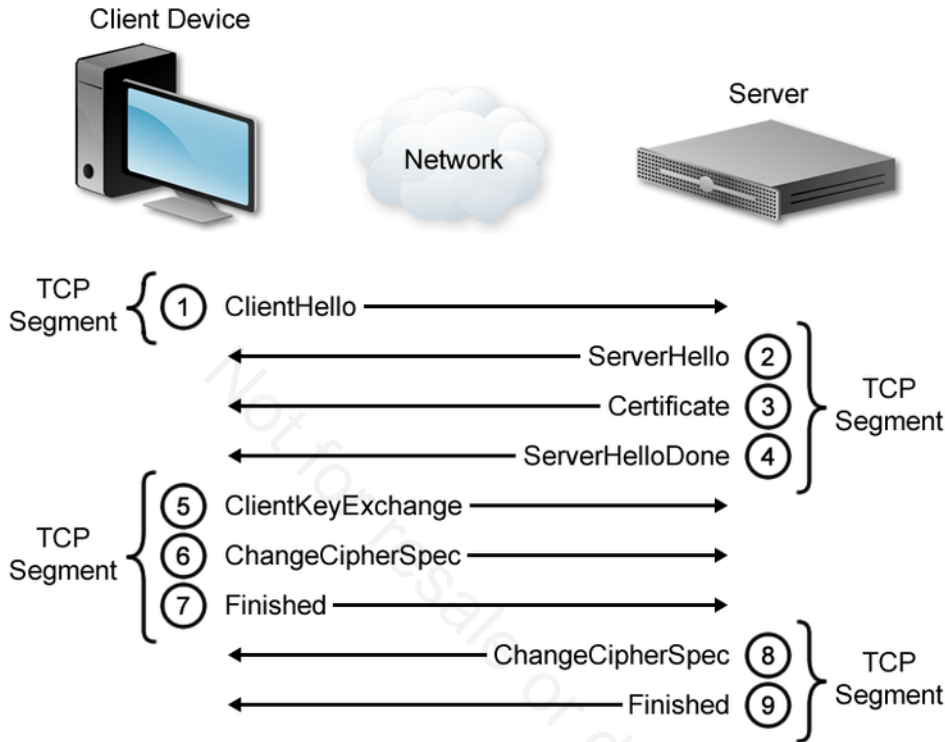
SSL

SSL is a protocol used to secure HTTP, TCP and other types of traffic; it is the industry standard security technology for establishing encrypted links between a web server and a browser.

SSL/TLS encrypts the data using a certificate that has unique credentials identifying the owner and authenticating the identity of the certificate owner.

SSL Session Process

For a client to establish a secure connection between a web browser and a server in most cases, a root certificate must be installed in the browser certificate store and on the server—which could be a web server, a service, or a NetScaler system. The following page describes the process in which a client and server initiate an SSL session.




The figure and the following process provide an overview of the process in which a client and server initiate an SSL session.

1. The client sends a Client-Hello with a set of supported cipher suites.
2. The server responds with a Server-Hello with the selected cipher suite.
3. The server sends its public key in the server certificate signed by the CA.
4. The server sends ServerHelloDone.
5. The client uses its root certificate to verify the signature of the server certificate. It then confirms that the certificate is still valid and that the subject name on the certificate matches the server name. The client generates a common key called the Premaster Secret.
6. The client transmits either the RSA-encrypted Premaster Secret or the Diffie Hellman (DH) parameters. Both options allow the client and the server to agree on the same Premaster Secret.
7. The client and the server, independently and in parallel, generate the same Master Secret using the Premaster Secret.

8. The client and the server generate session keys from the Master Secret.
9. The client and the server encrypt their communication using the session keys for the remainder of the SSL session.

Features and Benefits

The NetScaler SSL implementation supports a full feature set and is interoperable with all common SSL clients.

Feature	Benefit
Integrated or 1-arm	Scalability
SSL protocol support (SSLv2, SSLv3, TLSv1)	Rich protocol suite
Cipher suite support: <ul style="list-style-type: none"> • Key-exchange: RSA, DSS, DH [Max key: 4096 bits] • Encryption: RC4, DES, 3DES, RC2 • MAC: SHA, SHA-1, MD5 	Rich cipher suite - The level of security can be adjusted based on client needs with a full range of key exchange, encryption and authentication protocols.
Client authentication	<ul style="list-style-type: none"> • Authenticity of the end user requesting the web object • Web access on a user-by-user basis
Certificate Revocation List (CRL)/Online Certificate Status Protocol (OCSP)	Greater security, with ability to block clients with revoked certificates
<ul style="list-style-type: none"> • As many as 560,000 secure transactions every second • As many as 75 Gbps bulk encryption rate 	<ul style="list-style-type: none"> • Better performance • Quicker response time • Handling of traffic surges such as flash crowds
 Features and performance depend on your specific NetScaler platform. For more information, see the NetScaler Datasheet at http://www.citrix.com .	
Centralized certificate/key management	Ease of management
Easy drop-in installation	<ul style="list-style-type: none"> • No changes are needed to the web servers. • No additional hardware or software is needed.

Offload Performance

The NetScaler system supports extremely high-performance SSL encryption and session creation. For example, the NetScaler MPX platforms support:

- As many as 75 Gbps of bulk encryption.
- As many as 560,000 SSL handshakes every second for 2048 bit keys.

The NetScaler system can offload a significant portion of the total application load from the servers it supports as well as switch the resulting SSL traffic. This offload cannot be performed efficiently by traffic management systems. If the traffic management system is not decrypting the data, it is unable to make load-balancing decisions based on application data, such as HTTP header information or cookies. Instead, legacy traffic management systems must treat all of the traffic like raw IP address traffic and make decisions based on source and destination IP addresses. By decrypting the traffic as it enters the NetScaler system, SSL processing can be offloaded from application servers, allowing the robust traffic management functionality of the NetScaler system to come into play. SSL Offload to a NetScaler also allows for simpler certificate management. Fewer certificates are needed, thus reducing cost.

Digital Certificates

Digital certificates are small files that contain public keys and verify the identity of the holder. Certificates are issued by a Certificate Authority (CA). A CA is an entity that issues certificates after verifying the identity of a server.

Generating certificate requests and applying certificates can sometimes be a complex process depending on the CA that you use. Citrix recommends to always use a common CA such as Thawte, VeriSign, or Network Solutions. These CAs are usually trusted by all Windows and Macintosh operating systems and therefore require less administrative overhead.

Although self-signed certificates save the cost of purchasing a certificate from a commercial CA, they have limitations. When using self-signed certificates, be aware that only internal clients will trust those certificates. Any external client that is not a member of the domain of the issuing server will have to add the trusted root before a trusted connection can be established.

SSL Administration

The certificate formats that NetScaler supports are PEM and DER. An SSL certificate and key can be obtained for use on the NetScaler system using one of the following methods:

- Request certificate and key from a certificate authority (CA).
- Use an existing SSL certificate and key.
- Generate a new SSL certificate and key using the self-signing tools on the NetScaler system.

The process to create an SSL certificate is the same for the certificate authority or the NetScaler certificate tools:

1. Create a private key file.
2. Create a certificate signing request (CSR) using the private key file you created.
3. Submit the CSR to the certificate authority to create the certificate. If you are working with an external certificate authority, such as Verisign, submit the request to the CA which issued the certificate file. If you are using the NetScaler self-signing tools, the NetScaler system will use the CSR to generate the certificate.



Steps 1-3 can be skipped if a certificate and key file exist.



Use the NetScaler tools to create the key and request files when submitting a request to an external CA.

4. Install the certificate and key file on the NetScaler system.

SSL and Policies

Like other features of the NetScaler system, SSL offload can be controlled through policies.

SSL Keys

SSL Keys are strings of bits used to encrypt and decrypt messages. The NetScaler system supports RSA Algorithm keys and Digital Signature Algorithm (DSA) keys.

Keys are generated in the following situations:

- Before generating and submitting a certificate signing request to a certificate authority
- Before generating a self-signed certificate for testing purposes



The create key commands do not create objects in the NetScaler configuration; the commands create objects that are saved to the file system. Once generated, keys are saved by default to the [/nsconfig/ssl] directory in the NetScaler configuration.

When generating a key, you must designate the key size, or cipher key length, in bits. The current minimum standard for keys is 2048 bits, although the NetScaler system supports a key size as large as 4096 bits for RSA and DSA certificates. Higher key sizes result in more secure encryption algorithms.

Generate a key on the NetScaler system by specifying the key type, format and encryption in the SSL node of the configuration utility.

In the command-line interface, type one of the following:

RSA Key

```
create ssl rsakey <keyFile> <bits> [-exponents (3 | F4)]
[-keyform (DER | PEM)] [-des] [-des3] [-password <string>]
```

DSA Key

```
create ssl dsakey <keyFile> <bits> [-keyform (DER | PEM)]
[-des] [-des3] [-password <string>]
```

Public Key Encryption

SSL and TLS make use of public and private key cryptography to verify the identities of the server and client and to set up the initial secure channel for communication. Public-private key encryption is a form of asymmetric encryption, meaning that two separate but related keys are required.

Messages encrypted with the public key can only be decrypted using the private key. In turn, messages encrypted with the private key can only be decrypted using the public key. Public-private key encryption is a resource-intensive process.

Certificate Signing Request

Obtaining an SSL certificate from an authorized certificate authority requires the creation and submission of a certificate signing request (CSR).

Generate a CSR on the NetScaler system by specifying the key filename, format and encryption. In the configuration utility, expand the SSL node. In the command-line interface, type the following command:

```
create ssl certReq [-keyFile | -fipsKeyName ]
[-keyform (DER | PEM) {-PEMPassPhrase}] -countryName -stateName -
organizationName
```

The following parameters are mandatory:

```
[-keyFile] -countryName -stateName -organizationName
```

The following additional parameters are optional:

```
[organizationUnitName] [-localityName] [-commonName] [-
emailAddress]
{-challengePassword} [-companyName]
```

For more information about generating a CSR, see Citrix article CTX109260 at <http://support.citrix.com>.

SSL Certificates

The NetScaler system supports SSL certificates that are either self-signed or signed by a trusted certificate authority (CA). The NetScaler system can act as a certifying authority and provides options for generating a self-signed certificate for testing.

The NetScaler certificate tools can be used to generate the following certificate types:

- Root CA certificates
- Intermediate certificates
- Server certificates
- Client certificates

Root CA Certificates

Client devices use certificates to recognize which network entities to trust. Certificates are assigned to clients by a CA and carry an equivalent trust to that of the CA itself. A network entity requires a root_CA certificate to act as a CA and issue its own certificates to clients. This separate certificate is used to issue the client certificates. The root-CA certificate is critical, as it acts as the authority stamp on all client certificates that the CA will issue.

A NetScaler system can act as a CA and generate self-signed certificates by creating the root-CA certificate itself. In this scenario, clients must be configured to accept the NetScaler-created root-CA certificate to trust the client certificates that the NetScaler system will issue. To create a root-CA certificate, you must supply a key file and a certificate request.

Intermediate CA Certificates

Root CAs can assign certificates for an intermediate CA. An intermediate CA can then issue server certificates, client certificates or certificates for other intermediate CAs. The intermediate CA is located between your root CA and enterprise CA and requires another level of validation, creating a verification chain.

Server Certificates

The NetScaler system must have a server certificate for an initial client SSL session. The server certificate verifies the server identity to the client. The certificate authority verifies that the company using the server certificate was authorized to issue a certificate for the domain. To create a server certificate, you must specify the certificate request, the root-CA, or intermediate CA certificate that should be used to generate the server certificate, details regarding the key format, and an output file for the CA serial number. For more information about creating a server certificate, see Citrix article CTX109260 at <http://support.citrix.com>.

Client Certificates

Client certificates are issued to an end user by a certification authority. The certificate can be used only by the client for whom the certificate was issued. The difference between server and client certificates is that server certificates provide encryption and security functionality, whereas client certificates provide end-user authentication functionality.

Certificate Generation

Generate a certificate on the NetScaler system by specifying the certificate format, certificate type, CA certificate file format, CA key file name, CA key file format, CA key encryption and CA serial number file.



The CA parameters are optional.

In the configuration utility, generate a certificate under the SSL node.

In the command-line interface, type:

```
create ssl cert <certFile> <reqFile> <certType> [-  
keyFile <input_fileName>]  
[-keyForm (DER | PEM)] [-days <positive_integer>] [-  
certForm (DER | PEM)]  
[-CAcert <input_fileName>] [-CAcertForm (DER | PEM)] [-  
CAkey <input_fileName>]  
[-CAkeyForm (DER | PEM)] [-CAserial <output_fileName>]
```

For more information about replacing the default certificate of the NetScaler with a trusted CA certificate that matches the hostname of the NetScaler, see Citrix article CTX122521 at <http://support.citrix.com>.

Certificate Key Pairs

For SSL processing to occur, a certificate key entity must be bound to the virtual server. A certificate key entity is an integral element of the SSL encryption and decryption process, which is used during the SSL handshake to determine the cipher that will be used for SSL processing and also to establish the identity of the SSL server.

Before a client certificate can be used for SSL processing, it must be paired with its corresponding certificate key entity which resides on the NetScaler system. This certificate key pair is then bound to the client connection for SSL processing.

NetScaler VPX (1000) HA Status: Not Configured Info: NS11.0 62.10.nc Logout Citrix

Dashboard Configuration Reporting Documentation Downloads

← Back ?

Install Certificate

Certificate-Key Pair Name*

Certificate and Key files are stored in the folder /nsconfig/ssl/ on appliance.

Certificate File Name*
 Browse ▾ +

Key File Name
 Browse ▾ +

Certificate Format
 PEM DER

Password

Certificate Bundle
 Notify When Expires

Notification Period

Install Close

Adding a Certificate Key Pair

Add a certificate key entity on the NetScaler system by specifying the certificate file name, the private file name, the password and the certificate format:

Configuration utility location Under the SSL node

Command-line syntax Use the following command:

```
add ssl certkey <certkeyName> -cert <fileName>[(-key <fileName>
[-password]) | -fipskey <string>] [-inform (DER | PEM)] [-
expiryMonitor (ENABLED | DISABLED)]
[-notificationPeriod <positive_integer>]
```

For more information about adding a certificate key pair, see Citrix article CTX109260 at <http://support.citrix.com>.

Binding a Certificate Key Pair to a Virtual Server

Bind a certificate key pair by specifying the virtual server name and the key pair name on the NetScaler system.

Configuration utility location Under the SSL node

Command-line syntax Use the following command:

```
bind ssl vserver -certkeyname
```



Before the certificate key pair can be used, you will need to physically place the file on the NetScaler and then add it to the system for binding to the appropriate entities. If you are using the command-line interface, you will need to transfer the certificate to the NetScaler, typically by the way of SCP protocol to the directory.

Intermediate Certificate Not Linked

Certain server certificates can be issued by default by a CA that is not in the browser's trusted store. In this case, an intermediate certificate, which establishes the chain of trust, needs to be added to the NetScaler system and linked to the certkey. When the certificate is presented to the end user, the intermediate certificate is also provided.

When creating a certkey for the intermediate certificate, the process is the same as creating a certkey for a server, but you do not need the key information. To link the certificate, highlight the desired certkey and select **link**.



When linking certificates, the NetScaler system will only offer certificates that can be validly linked.

Certificate Revocation List (CRL)

A Certification Revocation List (CRL) is a list of the serial numbers of certificates that have been revoked and are no longer valid. End users should never accept this type of certificate to establish access to a secured resource.

If client or server SSL certification validation is required, the NetScaler system can perform validation against a local certificate revocation list in a local file.

This list can be updated on a specified interval automatically. Supported update methods include HTTP, LDAP and OCSP. Validation is not the default setting and must be enabled.

Certificate Updates

You might need to update or replace a certificate on the NetScaler system for a variety of reasons:

- Certificate is expired or expiring soon.
- Default certificate needs to be replaced with trusted CA certificate.
- File name or private key have changed or have been compromised.

A certificate can be updated from the configuration utility or the command line. In the configuration utility navigation pane, expand the **SSL** node, click **Certificates**, select the certificate you want to update, click **Update** and update the certificate where necessary. For notification about when a certificate expires, enable the **Notify when expires** checkbox.

SNMP monitoring can also be used to send alerts when a certificate is about to expire. Citrix recommends running reports on all certificate use and expiration dates to ensure availability of resources.

For more information about updating or replacing an SSL certificate, see Citrix article CTX109711 at <http://support.citrix.com>.

Certificates can also be updated from the Command Center. For more information about Command Center, see Citrix product documentation at <http://docs.citrix.com>.

For a demonstration about how to update certificates on the NetScaler system, see Citrix TV video 2960 at <http://www.citrix.com/tv/#videos/2960>.

Multiple Hosts on a Single IP Address

In order to implement multiple hosts on a single IP address (whether a simple or advanced implementation), you will need to present the client with a certificate for the correct FQDN. The following options can be used:

Wildcard Certificate

A wildcard certificate is a public key certificate which can be used with multiple subdomains of a domain. It allows for one single level of subdomain matching.

SubjectAltName (SAN) Extension

An SAN extension allows a certificate to contain multiple FQDNs in a single certificate.

Server Name Indication (SNI) Extension

An SNI extension allows a server to present multiple certificates on the same IP address and port number.

Discussion Question

Describe some situations in which an SSL certificate is added to the certificate revocation list.

SSL Offload Overview

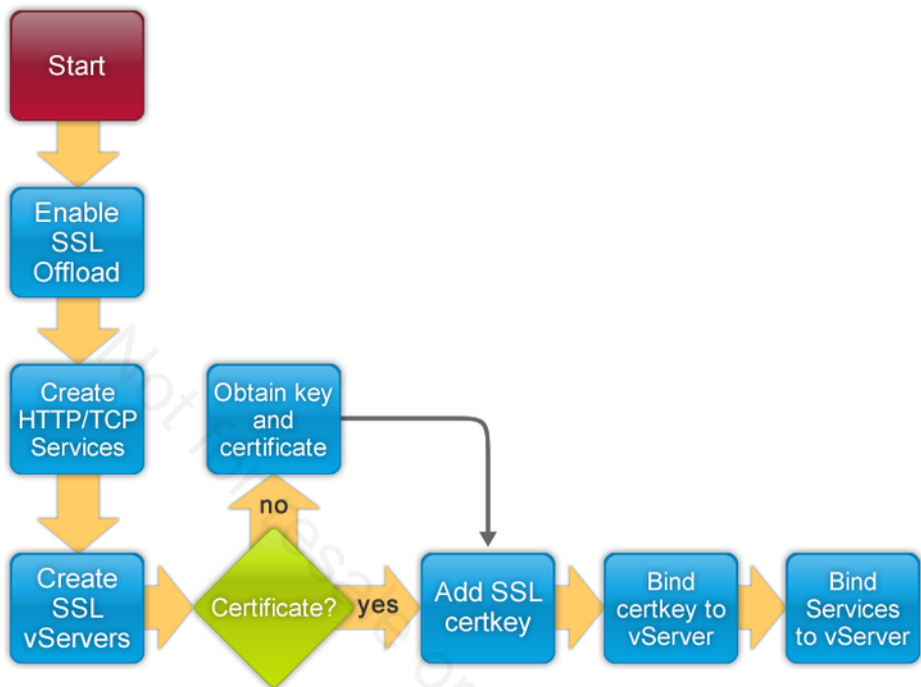
The SSL offload feature of the NetScaler system handles the CPU-intensive SSL encryption and decryption processing, allowing the web servers to dedicate more processing power to content requests. The SSL offload feature increases the performance of websites that carry out SSL transactions.

The figure provides an overview of a strict SSL offload scenario in which all SSL encrypted communication between the web servers and the client is handled by the NetScaler system. Communication between the NetScaler system and the back-end server is unencrypted, providing load reduction on the server and allowing the server to focus on performing the application role instead of on managing SSL encryption or decryption processes.



In other scenarios, the security of the data needs to be maintained end-to-end and should not be transmitted in a non-encrypted format. Encryption will occur between the client and the NetScaler system and also between the NetScaler system and the server. The server still needs to perform work to encrypt the data, but the NetScaler system is able to optimize the traffic and improve efficiencies by reusing the SSL connections.

Configuring SSL Offload



SSL must be enabled to configure the NetScaler system for SSL Offload.

The figure and the following steps provide an overview of the process used to configure SSL Offload on the NetScaler system.

1. Enable SSL Offload on the NetScaler system.
2. Create HTTP/TCP/SSL services on the NetScaler system.
3. Create SSL virtual servers on the NetScaler system.
4. Use an existing key and certificate or obtain a key and certificate.
5. Upload the key and certificate to the NetScaler system.
6. Create the certkey entity on the NetScaler system from the uploaded key and certificate.
7. Bind the certkey to the SSL or SSL_TCP virtual server.
8. Bind the services to the virtual server.

Securing Traffic Between a Client and the NetScaler System

Secure the traffic between a client and the NetScaler system using the following process:

1. Export an existing key and certificate or generate a new key, CSR and Certificate.
2. Convert the certificate file to PEM or DER using OpenSSL, if applicable.
3. Upload the certificate and key to the NetScaler system using an SCP or SFTP application.



You can also upload the certificate using the configuration utility.

4. Create a certkey entity to install the certificate and key pair on the NetScaler system.
5. Create servers that point to the back-end web servers.



If the servers have been created previously, this step is not necessary.

6. Create an HTTP service associated with the web servers.



If the service has been created previously, this step is not necessary.

7. Create a load-balancing SSL virtual server.
8. Bind the virtual server to the services.
9. Bind the certkey to the virtual server.

SSL Virtual Servers

An SSL virtual server accepts encrypted traffic, decrypts it, and sends the clear text messages to the services that are bound to the virtual server. This process allows for offloading SSL processing to the NetScaler system and the back-end servers to process a greater number of requests.

Creating an SSL Virtual Server

In the configuration utility, you can create an SSL virtual server in the NetScaler system by specifying the virtual server name, service type, IP address and port of the virtual server under the SSL Node.

In the command-line interface, type:

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

Binding an SSL Virtual Server

At least one service must be bound to a virtual server for the virtual server to accept incoming client requests.

In the configuration utility, you can bind a service to the SSL virtual server by specifying the virtual server name and the service name under the SSL node.

In the command-line interface, type:

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

SSL Termination Points

The NetScaler system can also provide protection against encrypted attacks by terminating and offloading SSL. By terminating the SSL, the NetScaler system can inspect and protect all traffic going to the servers in addition to switching it intelligently. Without this benefit, attacks can be wrapped in SSL and delivered securely to the server.

To properly configure the NetScaler system, you must determine the SSL termination points. SSL transactions can be terminated on one of the following devices:

- Citrix NetScaler
- Web server
- NetScaler Gateway

The following termination methods can be used on the NetScaler system or on the servers:

- Terminating SSL on NetScaler System
- Terminating SSL on the servers

Terminating SSL on the NetScaler System

Terminating SSL on the NetScaler system has several benefits:

- The server is freed from the extensive CPU cycles required for the handshake, encryption and decryption processes.
- Switching decisions can be made based on request content.
- Denial-of-service attack protections can be implemented.

Terminating SSL on the Servers

The primary benefit to terminating SSL on the servers is the ability to deploy a NetScaler system without any changes to the server configuration. Furthermore, if the NetScaler system is intercepting the connection, SYN attack prevention benefits can be maintained. However, the server will require significant processing power to handle the SSL transactions.

In most cases, SSL sessions will be terminated on the NetScaler system. In critical banking or financial transactions, SSL encryption must continue through to the servers to ensure that the transaction cannot be compromised under any condition.

Deployment Scenarios

The SSL requirements for a particular environment depend on how SSL will be deployed. The following scenarios are the most common:


- | | |
|---|--|
| Front-end SSL with Back-end HTTP | Front-end SSL with back-end HTTP refers to a network configuration that allows the NetScaler system to offload SSL encryption and decryption and pass traffic to web servers in plain text. |
| Front-end SSL with Back-end SSL | Front-end SSL with back-end SSL refers to a network configuration that allows the NetScaler system to pass traffic to web servers in encrypted mode. Although the NetScaler system does not fully offload SSL processing, it can provide a benefit through SSL connection multiplexing or the use of a smaller cipher. |
| Front-end TCP over SSL with Back-end TCP | Front-end TCP over SSL with back-end TCP refers to a network configuration that allows the NetScaler system to offload SSL encryption and decryption for client connections, as well as pass any TCP traffic to servers in plain text. An example of this type of deployment scenario is providing secure access for LDAP directories. |

Front-end SSL with Back-end HTTP Requirements

The figure and the following list provide an overview of the requirements necessary for a front-end SSL with back-end HTTP configuration:

- An installed certificate
- A load-balancing virtual server using the SSL protocol

- One or more HTTP services associated with back-end web servers




Front-end SSL with Back-end SSL Requirements

The figure and the following list provide an overview of the requirements necessary for a front-end SSL with back-end SSL configuration:

- An installed certkey for traffic between the NetScaler system and client
- A load-balancing virtual server set to use the SSL protocol
- An SSL service or services associated with back-end web servers

For a more secure configuration, install a certkey on the SSL server for traffic between the NetScaler and the back-end web servers. Since the service type is HTTPS, the certificate is installed on the web servers.



Securing Traffic Between the NetScaler and the Server

Secure the traffic between the NetScaler system and web servers using the following process:

1. Install certificates to secure traffic between the client and the NetScaler system, as well as between the NetScaler system and the back-end web servers.
2. Create servers that point to the back-end web servers.



If the servers have been created previously, this step is not necessary.

3. Create SSL services associated with the web servers.



If the services have been created previously, this step is not necessary.


4. Create an SSL virtual server for the web servers.
5. Bind the appropriate services or server groups to the virtual server.

Front-end SSL_TCP with Back-end TCP Requirements

The figure and the following list provide an overview of the requirements necessary for a front-end SSL_TCP with a back-end TCP configuration:

- An installed certkey

- A load-balancing virtual server using the SSL_TCP protocol
- A TCP service or services associated with back-end web servers



Securing TCP Traffic Between a Client and the NetScaler System

Secure TCP traffic between a client and the NetScaler system, as well as pass TCP traffic to server using the following process:

1. Install certificates to secure traffic between the client and the NetScaler system.
2. Create servers that point to the back-end web servers.



If the servers have been created previously, this step is not necessary.

3. Create TCP services associated with the servers.



If the services have been created previously, this step is not necessary.

4. Create an SSL_TCP virtual server.
5. Bind the appropriate services to the virtual server.
6. Bind the appropriate certkey to the SSL_TCP virtual server.

SSL Bridge

The SSL_BRIDGE functionality allows all secure traffic to be bridged transparently and directly to the back-end web server. The system does not terminate or offload this traffic. In this scenario, the web server must handle all SSL-related processing.



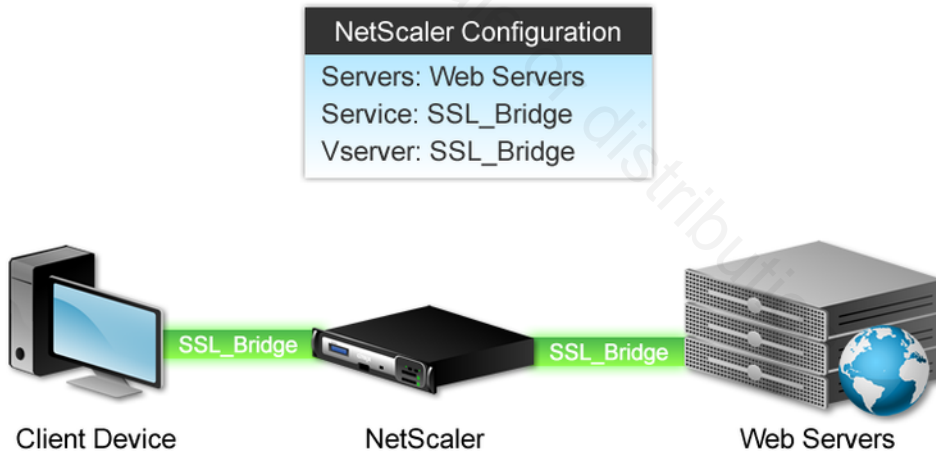
This configuration should be used only if an acceleration unit (for example, a PCI-based SSL accelerator card) is installed in the web server to handle the SSL processing overhead.

In an SSL_BRIDGE setup, the system is configured to load-balance and maintain server persistency for secure requests using the SSL Session-ID. Other features, such as content switching, Sure Connect and cache redirection do not function because the traffic passing through the NetScaler system is encrypted. Because the system does not carry out any SSL processing in an SSL_BRIDGE setup, it is not possible to bind a certkey to the SSL_BRIDGE virtual server.

SSL Bridge Requirements

The following requirements are necessary for an SSL_BRIDGE configuration:

- A load-balancing virtual server using the SSL_BRIDGE protocol
- A SSL_BRIDGE service or services associated with back-end web servers



Citrix Recommendations for SSL

Citrix recommendations for SSL include the following:

- Offload SSL processing to NetScaler.

- Be aware of which components in your infrastructure are processing SSL.
- Report on all certificate use and expiration dates.
- Document, measure and report on SSL performance.

SSL Renegotiation Attack

The SSL and TLS renegotiation process is vulnerable to an attack in which the attacker forms a TLS connection with the target server, injects content of its choice, and then splices in a new TLS connection from a client. NetScaler provides protection against these attacks by cryptographically binding renegotiation handshakes to the enclosing TLS cryptographic parameters, thus allowing the server to differentiate renegotiation from initial negotiation, as well as preventing a renegotiation from being spliced in between connections.

For more information about SSL renegotiation, see Citrix articles CTX140605 and CTX121925 at <http://support.citrix.com>.

SSL Troubleshooting

Because SSL issues involve client and server network traffic, it is useful to view HTTP headers when troubleshooting these issues. The following free tools are available for this task and can be found by performing a web search:

- Live HTTP Headers for Mozilla Firefox
- IE HTTP Headers for Microsoft Internet Explorer
- HTTP Headers Extension for Google Chrome



Header review is a well-known and effective method of troubleshooting connections for persistence failures when using cookie persistence, isolating authentication issues and back-end server response issues.

These tools provide you with the client view of the connection as it happens. As with other HTTP connections, you are able to see the server responses but will be missing the SSL details as everything is decrypted for the client before the headers are displayed. Therefore, missing portions of the SSL handshake and the certificate exchange are not displayed. These tools also allow you to save the headers for future review or sharing with others.



The Google Chrome Extension also includes the option to get client-side packet captures.

SSL Command-Line Interface Commands

When you use a show command, you typically extract a snapshot of configuration data and performance information from the NetScaler kernel. However, many commands show additional data if an exact object is selected. For example, the following command displays more specific information when you specify a virtual server:

```
show ssl vserver <vserver name>
```

More common commands include the following:

```
show ssl certkey
```

```
show ssl stats (stat ssl)
```

```
show ssl service <service name>
```

```
show ssl certlink
```

If you suspect that a global setting may be related to, or the cause of, your issue, you can list these at the command-line interface with the following command:

```
show ssl parameter
```

This command will display advanced global ssl parameters. It can be beneficial to check these global settings to determine if one of them is the cause of the issue.

Troubleshooting of Encrypted SSL Connections

With encrypted SSL connections, since the packet data is encrypted there are just a few options for troubleshooting. TCP layer and below-the-TCP layer information is viewable in the nstcpdump and nstrace files. You can also verify connections by reviewing the connection table on the NetScaler system.

In the configuration utility, this information is located in the **System > Diagnostics** menu. Then, click **Monitor Connections > TCP/IP Connections**. You can view all connections or filter by IP, port, service type and other options.

NetScaler VPX (1000) HA Status Not Configured Info NS11.0 62.10.nc Logout Citrix

Dashboard Configuration Reporting Documentation Downloads

← Back

View Sessions Help >

Expression	Link	Name	Connection Failover	Full	NNM
N/A	False	False	False	False	False

TCP/IP connections

Source IP	Source Port	Destination IP	Destination Port	Service Type	Idle Time	State	Is Link?	TD
127.0.0.1	3021	172.17.17.110	3011	RPCSVR	17	ESTABLISHED	0	
127.0.0.1	44850	127.0.0.1	5000	RPCSVR	110	ESTABLISHED	0	
127.0.0.2	38114	127.0.0.1	8008	MONITOR	34	TIME_WAIT	0	
127.0.0.2	64806	127.0.0.1	8008	MONITOR	14	TIME_WAIT	0	
127.0.0.2	12410	127.0.0.1	3335	UNKNOWN	3117	ESTABLISHED	0	
127.0.0.2	29426	127.0.0.1	8008	MONITOR	4	TIME_WAIT	0	
127.0.0.1	10760	127.0.0.1	5000	RPCSVR	3117	ESTABLISHED	0	
127.0.0.2	10565	127.0.0.1	7001	UNKNOWN	3118	ESTABLISHED	0	
127.0.0.2	3005	127.0.0.1	7776	MONITOR	36	TIME_WAIT	0	
127.0.0.2	48048	127.0.0.1	80	MONITOR	29	TIME_WAIT	0	
127.0.0.2	20487	127.0.0.1	7776	MONITOR	6	TIME_WAIT	0	
127.0.0.1	10772	127.0.0.1	5000	RPCSVR	7	ESTABLISHED	0	
127.0.0.1	10748	127.0.0.1	8777	AAA	3	ESTABLISHED	0	
127.0.0.2	24991	127.0.0.1	7776	MONITOR	1	TIME_WAIT	0	
127.0.0.1	10779	127.0.0.1	8777	AAA	9	ESTABLISHED	0	

Other options for troubleshooting encrypted SSL connections include packet-level analysis and a Wireshark capture. Nstcpdump is a standard tcpdump with a wrapper around it to allow for some NetScaler-specific information and nstrace is a tool used for creating capture files for deeper analysis. Technical support will often ask for nstrace captures, because it adds information to the trace that the nstcpdump output will not provide and provides you with the syntax to use for each case. Tcpcdumps are widely used for troubleshooting because they provide live data on the layer-3 connectivity, but tcpdumps are still limited to ensuring that the TCP connection is working or not.

A protocol analyzer such as Wireshark can help you view the packet data to see where in the connection the issue is located. Wireshark is capable of decrypting SSL traffic if the encryption key is provided. For more information about using Wireshark with Citrix Netscaler, see Citrix articles CTX122318 and CTX116557 at <http://support.citrix.com>.

SSL Offload Troubleshooting

SSL offloading issues include:

- Access to the SSL VIP address failing
- Certificate-related warnings occurring
- Intermediate certificate improperly linked
- Browser warning showing a webpage that is not secure

Access to SSL VIP Address Failing

This issue typically occurs when the certkey (certificate-key pair entity on the NetScaler system) is not bound. If the certkey is not bound, the status of the virtual server will display as DOWN.

Enter the following command in the command-line interface to show the state:

```
show ssl vserver <vserver_name>
```

The following table lists the available arguments:

Argument	Description
<i>vserver_name</i>	Specifies the name of the SSL virtual server

Enter the following command in the command-line interface to resolve the issue:

```
bind ssl certKey <vserver_name> <service_name> <certkey_name>
```

The following table lists available arguments:

Argument	Description
<i>vserver_name</i>	Specifies the name of the SSL virtual server to which the certificate-key pair needs to be bound
<i>service_name</i>	Specifies the name of the SSL service to which the certificate-key pair needs to be bound. Use the <code>add service</code> command to create this service.
<i>certkey_name</i>	Specifies the object name for the certificate-key pair

Discussion Question

What are some situations in which you would use an end-to-end SSL connection?

Module 7

Authentication and Authorization

7

Not for resale or distribution

Not for resale or distribution

Authentication and Authorization Manual

Overview

Authentication is used to allow end users to log on to the NetScaler Gateway and to connect to resources within the internal network. Authentication is an important part of security, ensuring that only authorized end users are able to log on and access resources.

Authorization is used to determine which resources end users have access to. For example, authorization can be used to permit or deny access to particular networks or IP addresses.

Authentication is often used in conjunction with LDAP to authenticate end users against Microsoft Active Directory. This allows end users to log on to the NetScaler Gateway using their domain credentials and then access resources.

Authorization can be used to ensure that once logged on, end users can only access resources that an administrator defines as allowed.

Authentication and authorization can both be configured by using policies.

After completing this module, you will be able to:

- Explain the types of authentication and authorization and determine why you would want to implement them.
- Configure and implement authentication on the NetScaler system.
- Identify key security strategies for securing NetScaler communications.

System and AAA Users Groups

System users and groups access the NetScaler system for management and configuration.

Authentication, authorization and auditing (AAA) users and groups access the NetScaler system to own and control resources, including:

- SSL VPN authentication and VPN resources (NetScaler Gateway)
- Virtual servers

The NetScaler system manages two types of users and groups:

System users and groups

These users and groups are used for logging onto the system directly for NetScaler management functions through the configuration utility or the NetScaler command-line interface. You can define command policies which determine the level of permissions each account or group contains when performing management functions. These policies are command specs or command policies, for example:

- Read-only
- Operators
- Network
- Superuser

AAA users and groups

The AAA users and groups are a set of accounts that are defined on the NetScaler system and are used to access the managed and controlled resources. Policies are defined to determine which resources the accounts can access. AAA accounts do not have management access to the NetScaler system.

Traditionally, AAA accounts were used in conjunction with the SSL VPN configuration and were used solely for logging onto and accessing SSL VPN resources allowed by the policies. Additional policies (including session, network and traffic policies) are defined to further manage the SSL VPN account behavior.

NetScaler allows for the managing of resource access that is not part of an SSL VPN configuration. NetScaler AAA for traffic management functionality enables an administrator to configure the NetScaler system to use AAA accounts to manage access to websites and web resources. AAA for traffic management allows an administrator to use the same AAA capabilities that are available with a VPN implementation to provide authentication to resources and to manage allowed and denied authorizations to those resources for non-SSL VPN scenarios.

For example, an administrator has the `www.example.com` website that is accessible through the NetScaler system. The administrator wants to control access to the site for specific accounts and limit which areas of the websites different user groups can access without modifying the site to include authentication functionality. These sites include:

- `www.example.com/techsupport`
- `www.example.com/hr`
- `www.example.com/store`

In this example, the NetScaler AAA for traffic management feature provides this functionality without configuring a VPN.

Local Accounts

For both system and AAA accounts, end users and groups are created and maintained as local accounts on the NetScaler system. End user accounts, passwords and group membership are set and managed on the NetScaler system. If local accounts are used for authentication, then the local users and groups must be created and maintained individually on each NetScaler system.

Each NetScaler system has two local system accounts that are always maintained as local accounts:

nsroot

This account is the default administrative account for the NetScaler system and cannot be disabled or removed from the system. Citrix recommends changing the default account password.

#nsinternal#

This account is used for GSLB and high availability communications through the rpc nodes. The command `set rpcnode` implicitly uses the `#nsinternal#` account.

External Authentication

The NetScaler system integrates with the following external authentication and directory services systems:

- LDAP
- RADIUS
- TACACS+

With external authentication, the NetScaler system uses the existing accounts and passwords and groups managed within the directory services infrastructure to authenticate end users to the environment.

You can define authentication policies and actions that identify which directory service resource to connect to when performing external authentication. The NetScaler system uses the external directory service to verify that the account and password supplied is valid in the directory service and to identify the group membership. The results are then returned to the NetScaler system and the appropriate authentication and authorization decisions are made.

When using external authentication, you can simplify account management by using group extraction. With this functionality, you only define the groups on the NetScaler system. The group names added to the NetScaler system, either as system groups or AAA groups, exactly match the group names and case within the external directory service. The individual user accounts do not need to be created or maintained on the NetScaler system. Instead, when a user account is supplied for authentication, the external authentication configuration relies on group extraction to identify the group to which the user account belongs. The group memberships correspond to the directory service configuration. Authentication and authorization policies and permissions are managed and assigned at the group level.

External Authentication for System Users

Local and external authentication can be configured and managed separately for system and AAA accounts. The authentication service can be different for both account types, for example:

- System authentication using local and AAA authentication using LDAP (Active Directory)
- System authentication using TACACS+ and AAA authentication using LDAP and RADIUS

Authentication Actions and Policies

Like NetScaler classic policies, authentication policies are comprised of an expression and an action. Each directory service and authentication option has different authentication policy and action types. The authentication actions include the information required to perform the authentication behavior. The authentication policy determines when the policy matches based on the defined expression. The policy is then bound to users and groups or entities on the NetScaler system.

The NetScaler system defines the policies and actions:

Policy	Action
localPolicy	no associated action
ldapPolicy	ldapAction
radiusPolicy	radiusAction
tacacsPolicy	tacacsAction
nt4Policy	nt4Action

Local Authentication Configuration

When configuring the default local authentication, you must create a system user account and group on the NetScaler system. One or more user accounts must be bound to each group. Permissions can then be managed at the group level. The appropriate command policies must be bound to either the system user accounts or the system groups. Command policies determine system management permissions levels, which are superuser, network, operator and read-only.

NTLMv2 Authentication

Credentials entered for NTLMv2 (Windows NT LAN Manager version 2) authentication can be used for single sign-on (SSO). The NetScaler system attempts to connect using NTLMv2 first and only if that fails does it then attempt to fall back to NTLMv1.

Configuration of Command Policies for Delegated Administrators

You can also delegate different administration tasks on the NetScaler by using command policies. There are four default command policies:

Read-only	Allows read-only access to show all commands except for the system command group and ns.conf show commands.
Operator	Allows read-only access and also allows access to enable and disable commands on services. This policy also allows access to set services and servers as ACCESS DOWN.
Network	Permits almost complete system access, excluding system commands and the shell command.
Superuser	Grants full system privileges, such as the privileges granted to the default administrator, nsroot.

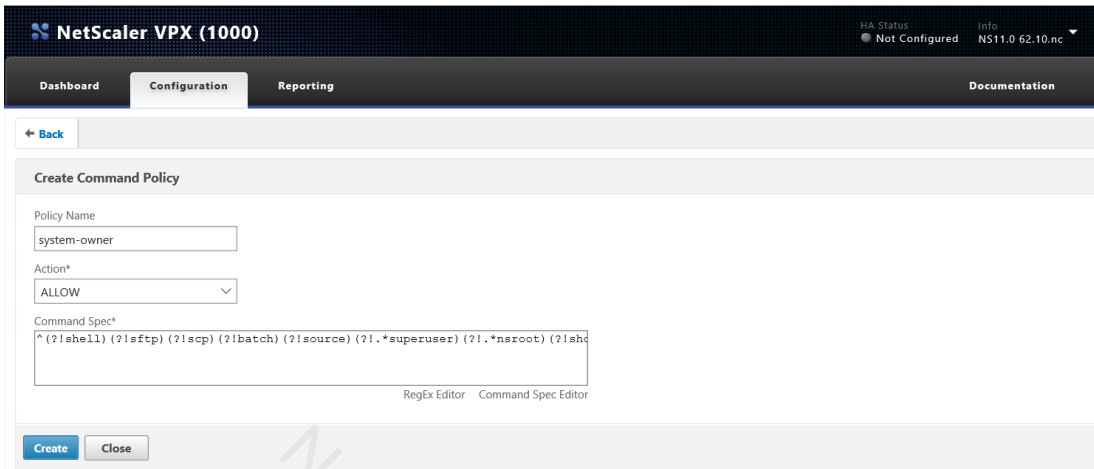
Command policies contain a list of built-in expressions that are permitted or denied for execution within the NetScaler Gateway. Command policies can be bound to system users or groups.

Command policies allow you to permit access to other users and groups with only the required permissions. This helps to prevent unscheduled changes or unnecessary downtime due to misconfiguration.

Custom Command Policy Configuration for Delegated Administrators

It is possible to configure custom command policies to ensure that delegated administrators only have a set of permissions that have been defined by an administrator. For example, you might want to create a command policy that only permits the end user to view the SSL node within the configuration utility. This can be configured within a new command policy and the policy can then be bound to either a system user or group.

A command policy contains a command specification; this determines which commands an end user can or cannot run. Command policies can have either an ALLOW or DENY action defined and can be modified or added to as necessary.



Authentication Configuration

NetScaler Gateway allows you to configure several different types of authentication, all controlled by using the NetScaler Gateway policy architecture. The policy architecture also allows for extensive customization of authentication. For example, you might choose a specific LDAP attribute to be used for single sign-on or specify the root search location of an LDAP directory, ensuring that only end users below this root are able to authenticate.

Through the use of policies you can configure multiple authentication points; for example, LDAP could be the primary authentication source with RADIUS used as a secondary authentication. The user will have to pass both authentication checks in order to be permitted to log on. Authentication policies are either bound globally or to a virtual server.

Authentication is independent of authorization; as such it can be configured without also configuring authorization.

You might want to specify certain authentication requirements based on the user's source network. For example, end users who are connecting from a partner site might be required to authenticate using a different method than other remote users.

Expressions can be used in policies to determine when a policy should apply, for example, using an expression to apply a policy if the source IP network matches a specified network.

Authentication Types Supported on NetScaler

NetScaler Gateway supports several authentication methods:

Local

Users created on the NetScaler Gateway

Lightweight Directory Access Protocol (LDAP)	Application protocol for accessing distributed directory information (for example, Microsoft Active Directory)
RADIUS	RADIUS authentication
SAML	An XML-based standard for exchanging authentication and authorization between Identity Providers (IdP) and Service Providers.
TACACS+	Similar to RADIUS authentication, however TACACS+ is used in place of RADIUS in certain organizations.
Client certificate authentication	NetScaler Gateway also allows for the configuration of authentication based on attributes of a client certificate.
KCD	NetScaler Gateway can be configured to support Kerberos Constrained Delegation (KCD) for single sign-on to Kerberos-based applications.

Proprietary authentication solutions such as RSA SecurID are typically supported by NetScaler Gateway using RADIUS for the configured authentication type.

If two-factor authentication is in place for Active Directory and RSA SecurID then two policies should be created, one policy for LDAP authentication and the other to provide RADIUS-based authentication.



The RSA SecurID administrator will need to ensure that the NetScaler NSIP is added as an object with a shared secret, otherwise the authentication request from NetScaler Gateway can be rejected.

Default Global Authentication Types Configuration

Often administrators initially configure a NetScaler Gateway by using the NetScaler Gateway Wizard. This is available from the NetScaler Gateway node within the configuration utility. You can specify the authentication method for end user logons using the wizard. Once this information has been provided the settings are bound globally.

You can add more authentication methods by creating authentication policies and then binding them either globally or to a virtual server. Ensure that if multiple policies are created then the

priority is set correctly, as otherwise the highest priority will apply and the remaining policies will not be evaluated.

When using LDAP it can be advantageous to create a load-balanced LDAP virtual server. This ensures that if a single authentication server becomes unavailable, LDAP requests can still be served.

Local Users Configuration

NetScaler Gateway allows local users to be added to the appliance. This can be useful for testing policies and profiles, while also permitting connectivity from end users who do not have an account within the authentication server.

Once end users have been created they can also be added to groups within NetScaler Gateway.

Policies can be applied to users and groups to enable specific settings to be applied. This can be by way of session policies, authorization policies, auditing policies or traffic policies.

End users can be removed from the NetScaler Gateway when they are no longer required.

An organization can use local user accounts for consultants or temporary users that do not have accounts within the authentication server.

Discussion Question

What are some of the advantages of using command policies for delegated administration?

Authentication Policies

Authentication policies are used to verify end-user logon credentials against an authentication source, for example, LDAP or RADIUS.

The default authentication type is local. This verifies users against a local user list that is located on NetScaler Gateway. To facilitate this, local users must be created on the NetScaler Gateway.

NetScaler Gateway supports the creation of multiple authentication policies, allowing you to define a detailed authentication procedure.

For example, you might want to authenticate users against Active Directory and a RADIUS service. To achieve this configuration, you can create two authentication policies: one that defines LDAP as an authentication location and a second that uses RADIUS as the authentication service. The policies can then be bound to the NetScaler Gateway virtual server as the primary and secondary authentication methods and both would share the same user name logon credentials.



Authentication traffic is generated using the NetScaler Gateway's NSIP address, not the SNIP or MIP. However, if you load balance authentication traffic, the source IP will change to the SNIP or MIP. Also, the NetScaler system has LDAP-aware health checks to verify LDAP on the backend servers.

Authentication Order

When an end user logs on to NetScaler Gateway, authentication is evaluated in the following order:

- The virtual server is checked for any bound authentication policies.
- If authentication policies are not bound to the virtual server, NetScaler Gateway checks for global authentication policies.
- If an authentication policy is not bound to a virtual server or globally, the end user is authenticated through the default authentication type.

Authentication Policy Bindings

Authentication policies can be bound either globally or to a NetScaler Gateway virtual server. In some occasions it might be preferable to bind an authentication policy globally and then use another policy to override the authentication where required.

NetScaler Gateway policy manager can be used to simplify policy configuration and bindings; this tool is available by selecting the NetScaler Gateway node within the configuration utility.

Priorities are used to determine which policy should be applied first. For example, if a global authentication policy has a priority of 10 and another authentication policy that is bound to a virtual server has a priority of 20, the global authentication policy will be applied first.

LDAP Authentication Configuration

NetScaler Gateway can use LDAP to perform authentication of end users; this provides integration with several entities:

- Microsoft Active Directory
- Novell eDirectory
- IBM Directory Server
- Lotus Domino
- Sun ONE directory

LDAP uses the following standard ports to provide services:

- 389 (unsecured LDAP) and StartTLS
- 636 (secured LDAP)

Determining LDAP Attributes

Citrix recommends using an LDAP browser such as Softerra to determine attributes within an LDAP directory. An LDAP browser allows you to connect to a LDAP server and then examine

attributes that are available within the directory; this can be useful for ensuring that the correct attributes are used when configuring authentication profiles.

RADIUS Authentication Configuration

RADIUS authentication is used to provide integration with commonly used authentication products, including:

- RSA SecurID
- SAM Express or SAS
- Gemalto Protiva

Many RADIUS-based systems require that the server sending the data has a record created within the management console. When using NetScaler Gateway, the authentication traffic uses the system's NSIP address as the source address.



Two problems are common with RADIUS authentication implementations: when an administrator does not document or loses the shared secret and when an administrator does not configure the RADIUS server to accept requests from NetScaler.

Troubleshooting Authentication

Connecting to the NSIP address allows you to perform additional command-line interface tasks. To assist in troubleshooting, you can run the following command from the shell:

```
cat /tmp/aaad.debug
```

This command will show real-time authentication attempts and the responses. This information is particularly useful when trying to isolate authentication failures.



The order of the information displayed from the `cat` command corresponds to the priority of policy usage. Group extraction is case sensitive.

Client Certificate Authentication Configuration

Some organizations use client certificates to provide an additional layer of security. Generally a certificate is created for a user and the user name is present within the subject field of the certificate. This can be used as an authentication method to confirm that the connecting user has all of the required details.

When you configure authentication using client certificates, it is possible to define which certificate field stores the user and the group data.

Using Certificates as Two-Factor Authentication

Client certificates can be used to authenticate end users first and then allow the user to authenticate using another method, for example LDAP or RADIUS.

When the end user tries to connect to the NetScaler Gateway logon page the client certificate will be examined and then authentication will be requested using either of the following options:

- Neither the user name nor the group is extracted from the certificate. The logon page appears to the user with a prompt to enter valid logon credentials. NetScaler Gateway authenticates the user credentials as in the case of normal password authentication.
- The user name and group name are extracted from the client certificate. If only the user name is extracted, a logon page appears to the end user in which the logon name is present and the end user cannot modify the name. Only the password field is blank.

To Configure a Client Certificate Authentication Policy

1. In the configuration utility, in the navigation pane, click **NetScaler Gateway > Policies > Authentication > CERT**.
2. On the **Policies** pane click **Add**.
3. In **Name** field, type a name for the policy.
4. The Authentication Type is **Cert**.
5. Next to Server, click + to go to the **Create Authentication CERT Profile** window.
6. In Name field, type a name for the profile.
7. Next to Two Factor, select **OFF**.
8. In **User Name Field and Group Name Field**, select the values and then click **Create**.
9. In the **Expression** field, click the menu for **Saved Policy Expressions** and select **ns_true**.
10. Click **Create**.



If you previously configured client certificates as the default authentication type, use the same names that you used for the policy. If you completed the User Name Field and Group Name Field for the default authentication type, use the same values for the profile.



Once the authentication policy has been created it must be bound to either a NetScaler Gateway virtual server or globally.

Smart Card Authentication Configuration

NetScaler Gateway supports the use of cryptographic smart cards to authenticate end users.

To enable support for smart card-based authentication you must complete several tasks:

1. Create an authentication policy that uses certificates as the authentication method.
2. Bind the authentication policy to a virtual server.
3. Add the root CA certificate that issued the client certificates to NetScaler Gateway, ensuring that the certificate is added as a CA certificate.
4. Bind the root certificate to the virtual server.
5. Configure NetScaler Gateway for client certificate authentication.

Common Access Cards

As an example of smart card-based authentication, the United States Department of Defense uses common access cards for identification and authentication.

To configure a common access card:

1. In the configuration utility, in the navigation pane, click **NetScaler Gateway > Policies > Authentication > CERT**.
2. On the **Policies** pane click **Add**.
3. In **Name** field, type a name for the policy.
4. The Authentication Type is **Cert**.
5. Next to Server, click + to go to the **Create Authentication CERT Profile** window.
6. In Name field, type a name for the profile.
7. Next to Two Factor, select **OFF**.
8. In the User Name Field drop down menu select `SubjectAltName:PrincipalName` and then click **Create**.
9. In the **Expression** field, click the menu for **Saved Policy Expressions** and select **ns_true**.
10. Click **Create**.
11. Bind the policy to the virtual server.

Ensure that the latest release of StoreFront is used when working with smart card pass-through.



In order to support smart card-based authentication for XenApp applications several steps must be completed within the domain, StoreFront, XenApp and NetScaler Gateway to ensure that smart card authentication works correctly.

For additional information about configuring an environment with smart card authentication, see the Citrix article CTX124603 at <http://support.citrix.com>.

For more information about smart card authentication and using Secure ICA in conjunction with smart card-based authentication, see Citrix product documentation at <http://docs.citrix.com>.

To Disable Authentication

If your NetScaler Gateway deployment does not require authentication, you can disable it. You should disable authentication for each virtual server that does not require authentication.

To disable authentication, complete the following steps:

1. In the configuration utility, in the navigation pane, expand NetScaler Gateway and then click **Virtual Servers**.
2. In the details pane, click a virtual server, and then click **Edit**.
3. In the **Authentication** field, click the existing policy and in the **Policy Binding** window, highlight the desired Policy then click **Unbind**, **Close** then **Done**.

Authorization Configuration

Authorization is used to define which network resources users have access to when they log on to NetScaler Gateway. By default NetScaler Gateway allows access to all network resources.

Authorization policies are used to determine which resources are either allowed or denied. Once the policy has been created it can be bound to users or groups.

Use the default rule of DENY and then create policies to define which resources should be accessible. This approach is often easier to implement and less prone to mistakes.

Explicit user or group policies will override the global authorization action.

Default Global Authorization

The default global authorization action specifies the default action for access to network resources. For example, the default action could be ALLOW, which then specifies that all network access is permitted unless explicitly defined in an authorization policy with a DENY action.

Authorization Policies Configuration

Authorization policies define either an ALLOW or DENY action. This action will be used if the expression specified in a policy evaluates to TRUE. For example, an expression can be configured to match traffic that is destined for 10.0.0.86; an administrator could then set the action for the policy to deny the traffic. This would deny traffic with a destination IP address of 10.0.0.86.

Policies are first applied to end users; however if no policy is applied to end users then NetScaler Gateway checks to see if an authorization policy is bound to a group to which the end user belongs. If none are found then the default global authorization rule applies.

Authorization policies can only be bound at the group or user level. The alternative is to use the default authorization setting in a session policy.



Authorization policies are often used to permit groups to access specific networks only. Administrators usually set the default authorization rule to DENY and then explicitly define which network resources should be available.

Authentication, Authorization, and Auditing (AAA) Issues

Authentication, authorization, and auditing (AAA) issues can cause content located behind a NetScaler system to become inaccessible. First, answer the following questions:

- Have configuration changes been made to servers or network devices?
- Have configuration changes been made to server, service or virtual server objects?
- Can the site be accessed directly (in other words, bypassing the NetScaler system)?
- Can the server and port be accessed using Telnet?

It is also important to gather data before troubleshooting. Useful commands include `stat aaa` (displays AAA statistics) and `show aaa` (displays information for users, groups, parameters).

For more information about troubleshooting, see Citrix articles CTX133789, CTX131684 and CTX120639 at <http://support.citrix.com>.

Authentication Troubleshooting

When troubleshooting LDAP, or any type of authentication with Citrix NetScaler, a useful tool is the `Aaad.debug` module. Using this tool, you can troubleshoot authentication issues such as general authentication errors, username or password failures, authentication policy configuration errors and group extraction discrepancies.

The `Aaad.debug` tool displays information in real time, including group extraction and authentication failures, which can be very useful for determining which authentication provider is causing the issue.

To use the `Aaad.debug` tool, begin at the command-line interface, access the shell, change to the `/tmp` directory and begin the debugging process by typing the following command:

```
cat aaad.debug
```

For more information about troubleshooting the authentication process, see Citrix article CTX114999 at <http://support.citrix.com>.

For more information about troubleshooting Active Directory group extraction for LDAP, see Citrix article CTX125797 at <http://support.citrix.com>.

Discussion Question

What are some of the common authentication issues in your environment and how do you troubleshoot them?

Module 8

Access Policies

8

Not for resale or distribution

Not for resale or distribution

Access Policies Manual

Overview

NetScaler Gateway supports the use of endpoint analysis to determine if a device meets a set of requirements as defined by an administrator. Endpoint Analysis is often referred to as a type of network access control (NAC). Based on the results of the endpoint analysis you can determine the level of access that a device should be given. Using endpoint analysis can provide a greater level of security by ensuring that devices meet the minimum corporate criteria in order to access company resources.

After completing this module, you will be able to:

- Explain how endpoint analysis is used to verify that the user device meets your requirements before allowing it to connect to your network or remain connected after end users log on.
- Explain how endpoint analysis can be used to evaluate end-user logon options.
- Configure Preauthentication policies and profiles to check for client-side security before end users are authenticated.
- Configure post-authentication policies that an end-user device must meet to keep their session active.
- Explain which security checks can be configured prior to end user logon.
- Configure file scan and registry policies to determine if end-user devices meet administrator-defined requirements.

Endpoint Analysis

Endpoint analysis allows you to define scans that routinely run once the end user is logged on. For example, you might want to ensure that the corporate firewall "CPfirewall.exe" is running on the end user's device during their session and, if not, the session should be terminated. Endpoint analysis allows the monitoring of files, processes, and registry entries on the end-user device during the session to ensure that the device continues to meet corporate requirements. Endpoint analysis includes Preauthentication as well as post-authentication scans.



Endpoint analysis policies might need to be combined with HTTP policies to determine the platform type for some operating systems.

Endpoint Analysis Plug-in

If endpoint analysis is required for a NetScaler Gateway virtual server, the end user will be presented with a plug-in for download. This plug-in facilitates the endpoint scanning. When the end user accesses the NetScaler Gateway using a web browser the plug-in will prompt the user to

confirm that the plug-in can scan the device. If the end user selects "skip" this could possibly deny them access to logon.

For more information about the endpoint analysis plug-in, see Citrix article CTX124649 at <http://support.citrix.com>.

Endpoint Policies

NetScaler Gateway allows the creation of Preauthentication policies to determine if a device is allowed to log on to the gateway. Common checks for a Preauthentication scan include:

- Antivirus
- Firewall
- Processes
- Files
- Registry Entries
- Operating Systems

A Preauthentication policy will either allow or disallow the initial logon page to be displayed based on the result of the endpoint analysis scan.

A session policy can be created to determine the configuration of the end user's session--for example, if the user is provided with clientless access or with Web Interface access. A session policy can be bound to a user or a group and can contain an endpoint analysis scan to ensure that the required items, such as whether a process is running, is still true. Session policies also integrate with the SmartAccess features of XenApp and XenDesktop, whereby you can set specific options for the end user's application or desktop session, based on the result of an endpoint analysis scan.

Client security expressions can also be specified within a session policy to determine if a device that has failed to meet the requirements of the expression is then placed into a quarantine group. Quarantine groups can be permitted to log on with a restricted access profile, for example allowing clientless access only.

Providing varying levels of access to corporate resources based on the end user's device is an effective way to enhance security. For example, devices that run a legacy version of Windows may be treated as unsecure and can only access applications and desktops by using ICA proxy. However, devices with a specified version or later version installed might be offered the ability to use the NetScaler Gateway plug-in.



Opening the NetScaler Gateway logon page in a web browser will prompt the end user to install the endpoint analysis plug-in. Once installed, the scan will be completed.

Endpoint analysis performs the following basic steps:

1. Examines an initial set of information about the end user device to determine which scans to apply.

2. Runs all applicable scans. When end users try to connect, the Endpoint Analysis Plug-in checks the end user device for the requirements specified within the Preauthentication or session policies and the result of the scan is returned. If the end-user device passes the Preauthentication policy scan, users are presented with a logon page and allowed to log on. If the user device fails the Preauthentication policy scan, a notice that the device does not meet the requirements is presented instead of a logon page and the users are not allowed to log on.



Endpoint analysis scans complete before the user session uses a license.

3. Compares property values detected on the end-user device with desired property values listed in your configured scans.
4. Produces an output verifying whether or not desired property values are found.



Some companies only allow domain-joined devices to gain remote access, as these devices are considered trusted devices. Use of the expanded features in NetScaler Gateway's registry scan function allows you to examine a machine's domain membership.

Preauthentication Policies

Preauthentication policies allow administrators to allow users to access the logon page based on client conditions. Preauthentication policies use endpoint analysis scans to determine client conditions.

Policy Profiles

A policy profile specifies the actions a policy will take if it is enforced. The following table lists the actions allowed in Preauthentication policy profiles.

Action	Description
Allow	Allow access to the logon page
Deny	Deny access to the logon page
Processes to be killed	Kill named processes
Files to be deleted	Delete specified files

Policy Expressions

A policy expression specifies the conditions under which a policy will be enforced—that is, when the profile actions will be carried out. Preauthentication policies support the following types of conditions.

Condition Type	Description
Running processes and services	Includes: <ul style="list-style-type: none">• Antivirus applications• Personal firewall applications• Internet security applications• Anti-spam applications
Existing files	Includes: <ul style="list-style-type: none">• Configuration files• License files
Existing registry items	Includes: <ul style="list-style-type: none">• Standard registry entries• Modified registry entries
Operating system parameters	Includes: <ul style="list-style-type: none">• Type• Service pack version• Hotfix version

Policy Bindings

Preauthentication policies can be bound at the following levels:

- AAA global
- Virtual server

Creating a Preauthentication Policy

Preauthentication Policies and Profiles

Preauthentication policies allow you to specify which devices should be allowed to authenticate and log on to NetScaler Gateway. Using Preauthentication policies ensures that connecting devices conform to the security standard that is required for accessing corporate resources.

You can either allow or deny access based on a certain condition. For example:

- Allow access if Symantec.exe is running
- Deny access if Keylogger.exe is running
- Allow access if C:\Corporate\device.txt is present
- Deny access if C:\Windows\system32\TypedKeys.txt is present
- Allow access if Symantec.exe is running or McAfee.exe is running

Carefully configuring Preauthentication policies can help you control which devices are accessing corporate resources, as well as help prevent infected devices from connecting to the network.

Preauthentication policies can be configured either globally or by the use of a Preauthentication policy which can then be bound to a virtual server.

The outcome of a Preauthentication policy is either to allow or deny logon.

Preauthentication policies also provide an option to terminate a process or delete a file--for example, stopping a process that is known to interfere with an application that is accessed remotely. If this setting is enabled the end user will receive a prompt which asks for confirmation of the action to be carried out.



Many companies that implement BYOD platforms allow end users to choose from one of several recommended antivirus products to which the company can provide limited assistance. In this scenario it is common for NetScaler Gateway to be configured to allow logon only if the end user's device is running the process for one of the recommended antivirus products.

Security Preauthentication Expressions Configuration for End-User Devices

Through the use of endpoint analysis you can run security checks on end user devices. The outcome of the scan determines if a user is able to log on to NetScaler Gateway and which resources users can access. Post-authentication scans can also be used to ensure that a device continues to meet the security standard that an administrator has defined.

NetScaler Gateway supports the following security checks on end user devices:

- Antivirus

- Personal firewall
- Internet security
- Anti-spam
- Service policies
- Process policies
- Operating system
- File policies
- Registry policies

Operating System Policies Configuration

NetScaler Gateway can determine the operating system version during endpoint analysis scans. The scan can also determine the current service pack level of the operating system.

The following table lists the detected operating systems and the corresponding expression values:

Operating System	Value
Mac OS X	macos
Windows 10	win10
Windows 8	win8
Windows 7	win7
Windows Server 2012	win2012
Windows Server 2008	win2008
Windows Server 2003	win2003
Windows 32 bit platform	win32
Windows 64-bit platform	win64

Antivirus, Firewall, Internet Security, or Anti-Spam Expressions Configuration

NetScaler Gateway contains built-in named expressions for several major security software providers, including:

- Symantec
- Sophos
- McAfee
- ZoneAlarm
- TrendMicro

When creating a policy expression that is going to examine antivirus, firewall, or anti-spam protection, a type of **Client Security** is used. The client security type allows for the following settings to be configured:

Component	The type of client security, such as antivirus, firewall, or registry entry
Name	The name of the application, process, file, registry entry, or operating system
Qualifier	The version or the value of the component for which the expression checks
Operator	Checks if the value exists or is equal to the value
Value	The application version for antivirus, firewall, Internet security, or anti-spam software on the user device
Frequency	How often a post-authentication scan runs, in minutes
Error	Assigns a weight to each error message contained in a nested expression when multiple expressions have different error strings. The weight determines which error message appears.
Freshness	Defines how old a virus definition can be—for example, you can configure the expression so that virus definitions are no older than three days.

Using the policy option of **Match Any Expression** allows you to specify that only a single expression needs to be matched.

To Configure a Client Service Scan

Many Windows-based applications install services that run in the background of Windows. For example, most antivirus products install a service to ensure that the software is started during device startup.

Scanning for the presence of a service is supported by NetScaler Gateway using the Client Security type, with a component type of Service.

To configure a service policy, complete the following steps:

1. In the configuration utility, in the navigation pane, click **NetScaler Gateway > Policies > Preauthentication**.
2. On the **Preauthentication Policies** tab, click **Add**.
3. In **Name**, type a name for the policy.
4. Click the **Expression Editor** to the right of the **Expression** field.
5. In the **Add Expression** dialog box, in **Select Expression Type**:, select **Client Security**.
6. Configure the settings for the following:
 - a. In **Component**, select **Service**.
 - b. In **Name**, type the name of the service—for example, SAVService (this is the name of the Sophos AntiVirus service)
 - c. In **Qualifier**, leave blank or select **Version**.
 - d. Depending on your selection in **Qualifier**, do one of the following:
 - If left blank, in **Operator**, select == (equal to) or != (not equal to).
 - If you selected **Version**, in **Operator**, in **Value**, type the value, click **Done**.



You have to configure a Request Action before completing the creation of the policy.



When vendors release updated software, they sometimes change the service name, which could cause the scan to fail.



Before creating the expression, ensure that the service name is correctly identified. This can easily be achieved by starting the Services control panel applet from within the Administrative Tools section of the control panel, then viewing the properties of the service to determine the service name.

Security Checks Configuration

When creating a session or Preauthentication policy, you can define a rule that requires all end-user devices to run a particular process when users log on. The process can be any application and can include customized applications.

For more information about configuring process policies, see Citrix product documentation at <http://docs.citrix.com>.

To Configure Process Policies

You can also configure an expression to ensure that an MD5 checksum comparison is completed. For example, if you want to confirm that SavService.exe is running and verify the MD5 checksum, you must first calculate the MD5 checksum and then use the following expression (using the appropriate checksum value):

```
CLIENT.APPLICATIONS.PROCESS(SavService.exe_md5_287811c077b90af0b013cddf47dcfd69) EXISTS
```

1. In the configuration utility, in the navigation pane, click **NetScaler Gateway > Policies > Preauthentication**.
2. On the **Preauthentication Policies** tab, click **Add**.
3. In **Name**, type a name for the policy.
4. Click the **Expression Editor** to the right of the **Expression** field.
5. In the **Add Expression** dialog box, in **Select Expression Type:**, select **Client Security**.
6. Configure the settings for the following:
 - a. In **Component**, select **Process**.
 - b. In **Name**, type the name of the application.
 - c. In **Operator**, select **EXISTS** or **NOTEXISTS**, click **Done**.

File Scan Policies Configuration

NetScaler Gateway allows you to scan an end user's device to determine if a particular file is present on the device. This functionality is also extended to allow the checking of the file's time stamp to determine the age of the file. File scan expressions can be used in either Preauthentication or session policies.

Be aware that when file scans are used, the expression will display four backslashes rather than the single backslash that is used in a regular file path. For example, "C:\Test\file.txt" would display as the following in an expression:

```
C:\\\\test\\file.txt
```



Scanning for time stamps on files can be useful in very specific circumstances, but there are other situations in which a file's time stamp can change--for example, application updates, user modification, or removal of software. When working with antivirus definitions it is preferable to use a client security antivirus expression, specifying the version of the antivirus. This allows you to determine how current antivirus software is and can often be simpler to work with than a file scan.

To Configure a Session or Preauthentication Policy to Check for a File on the End-User Device

1. In the configuration utility, in the navigation pane, click **NetScaler Gateway > Policies > Preauthentication**.
2. On the **Preauthentication Policies** tab, click **Add**.
3. In **Name**, type a name for the policy.
4. Click the **Expression Editor** to the right of the **Expression** field.
5. In the **Add Expression** dialog box, in **Expression Type**, select **Client Security**.
6. Configure the settings for the following items:
 - a. In **Component**, select **File**.
 - b. In **Name**, type the name of the application.
 - c. In **Qualifier**, leave blank or select **Time Stamp**. If **Time Stamp** is already selected, in **Value**, type the value.
 - d. In **Operator**, select the value, click **Done**.

Registry Policies Configuration

Checking the registry of an end user's device allows for granular examination of a connecting device. Registry scans can help administrators to determine the current version of software that is installed, as many applications store this data within the registry. The registry also contains domain-specific information, and many implementations use registry-based scans to determine the domain membership of a device.

Administrators often use registry-based scans to detect the domain membership of end-user devices. Administrators then use this information to provide domain-joined devices with a higher level of access, as a level of centralized control is usually in place for these devices.



When expressions are defined, the following format is used:

- Four backslashes are used to separate keys and sub-keys.
- Underscores are used to separate the sub-key and the associated value name, such as:

```
HKEY_LOCAL_MACHINE\\\\"SOFTWARE\\\\"VirusSoftware_Version
```

Registry checks can be performed for all five registry hives, including:

- HKEY_CLASSES_ROOT
- HKEY_CURRENT_USER
- HKEY_LOCAL_MACHINE
- HKEY_USERS
- HKEY_CURRENT_CONFIG

To Configure a Registry Policy

1. In the configuration utility, in the navigation pane, click **NetScaler Gateway > Policies > Preauthentication**.
2. On the **Preauthentication Policies** tab, click **Add**.
3. In **Name**, type a name for the policy.
4. Click the **Expression Editor** to the right of the **Expression** field.
5. In the **Add Expression** dialog box, in **Expression Type**, select **Client Security**.
6. Configure the settings for the following items:
 - a. In **Component**, select Registry.
 - b. In **Name**, type the name of the registry key.
 - c. In **Qualifier**, leave blank or select a value.
 - d. In **Operator**, complete one of the following tasks: If the Qualifier field is left blank, select EXISTS or NOTEXISTS. If you selected Value in the Qualifier field, select either == (equal to) or != (not equal to).
 - e. In **Value**, type the value as it appears in the registry editor, click **Done**.

To Create a Sample Preauthentication Scan

To create a test scan to detect whether the end user's device is running Notepad.exe, complete the following steps:

1. In the configuration utility, in the navigation pane, click **NetScaler Gateway > Policies > Preauthentication**.
2. On the **Preauthentication Policies** tab, click **Add**.
3. Enter a **Name** for the policy.
4. Next to **Request Action**, click +.
5. Specify a name and an Action of ALLOW.
6. Click **Create**.
7. Click the **Expression Editor** to the right of the **Expression** field.

8. In the **Add Expression** dialog box, in **Expression Type**, select **Client Security**.
9. Change the Component to **Process**.
10. Enter `Notepad.exe` in the Name field.
11. Click **Done**.
12. Click **Create**.
13. Bind the policy to the NetScaler Gateway virtual server.

Custom Expressions Configuration

NetScaler Gateway allows you to specify custom expressions for use within profiles. Creating a custom expression allows you to check for a wide variety of possible parameters, for example:

- Presence of a particular file
- Presence or value of a registry entry
- Version of antivirus
- Number of days since antivirus was updated
- Whether a process is running

Policies also allow for compound expressions to be used, allowing for increased flexibility. Commonly this is used to check for the existence of an approved antivirus software on the end user device. For example, an expression can be configured to ensure that the device is running one of these antivirus products:

- Symantec AntiVirus 12
- McAfee VirusScan 12
- Sophos Anti-Virus 9

The OR operator is used to allow NetScaler Gateway to match the expression if a single product is detected.



Depending on the antivirus product that is used, it is possible that an end user might have a newer version than the one used on corporate devices. For example, some antivirus licensing allows remote users to download and use the latest software editions that are updated from the vendor's website. In this scenario an end user might be running a newer version than those running on the internal devices.

To ensure that this does not cause issues and that the administrator does not have to constantly update expressions, it is recommended to either define the minimum version of antivirus or use a scan for the antivirus process.

To Bind Preauthentication Policies

Once a Preauthentication policy has been created, it must be bound to the relevant bind point in order to become active. Preauthentication policies can be bound either globally or to virtual servers.

1. Bind the policy to one of the following locations:
 - Under **NetScaler Gateway > NetScaler Gateway Policy Manager > AAA Global > Preauthentication Policy**.
 - Under **NetScaler Gateway Virtual Servers**, edit a virtual server and then bind the policy under the **Policies** field.
2. In the configuration utility, in the navigation pane, click **NetScaler Gateway**.
3. In the details pane, under **Policy Manager**, click **NetScaler Gateway Policy Manager**.
4. In the Policy Manager, click the + sign next to **AAA Global**, Click **Add Binding**, click in the field under **Select Policy**.
5. Check the radio button next to the desired policy then click **Select**.
6. Click **Bind**, then **Done**, then **Done** again.



Steps to Bind the Preauthentication policy to a virtual server.

- Under **NetScaler Gateway > NetScaler Gateway Policy Manager > NetScaler Gateway Virtual Server**.
- Under the **NetScaler Gateway Virtual Server** window, highlight a virtual server and then click **Edit**.
- Click the + sign next to **Policies**, select **Preauthentication** from the **Choose Policy** drop down list then click **Continue**.
- Click in the field below **Select Policy**, check the radio button next to the desired policy then click **Select**, **Bind**, **Done**, **OK** then **Done**.

Preauthentication Policy Considerations

Configure Preauthentication policies and profiles to check for client-side security before end users are authenticated.

For more information about multiple Preauthentication policies that are bound to different levels, see Citrix product documentation at <http://docs.citrix.com>.

NetScaler Gateway supports specifying multiple Preauthentication policies by binding them to different bind points. For example, a Preauthentication policy that checks for the running of a particular process could be bound globally, while another policy that checks for the presence of a text file could be bound to a virtual server.

When an end user attempts to log on to NetScaler Gateway the virtual server policy is applied first. The global policy is applied secondarily.

By using priority values for the scans it is possible to alter the order in which the scans are run. The policy with the lowest number will be evaluated first. For example, re-ordering the priority of the

global policy to 10 and the virtual server policy to 20 will instruct NetScaler Gateway to run the global policy first.

To Change the Priority of a Preauthentication Policy

1. Under **NetScaler Gateway > NetScaler Gateway Policy Manager**.
2. In the Policy Manager, expand either **AAA Global** or **NetScaler Gateway Virtual Server**.
3. If you selected **NetScaler Gateway Virtual Server**, click on the virtual servers, highlight the desired virtual server then click **Edit**.
4. Under **Policies**, click on the **AAA Preauthentication Policy**, highlight the policy then click the **Edit** button and select **Edit Binding**.
5. In the **Priority** field, type a number and then click **Bind**, then **Close**, then **Done**, then **OK**.

To Unbind a Preauthentication Policy

1. In the configuration utility, in the navigation pane, click **NetScaler Gateway**.
2. In the **details** pane, under **Policy Manager**, click **Change group settings and user permissions**.
3. In the Policy Manager, under **Configured Policies > Resources**, click the **AAA Global** or **Virtual Server** node to which the policy is bound.
4. Select the policy and under **Related Tasks**, click **Unbind Preauthentication policy** and then click **Yes**. Once the policy is no longer bound, it can be removed.

To Remove a Preauthentication Policy



Before a policy can be deleted it cannot be bound to any resources.

1. Under **NetScaler Gateway > Policies > Preauthentication**.
2. In the details pane, under **Preauthentication Policies**, highlight the desired policy then click **Delete**.
3. Click **Yes** to confirm.

Once a Preauthentication policy has been removed it must be re-created; it cannot be recovered.

Post-Authentication Policies

Through the use of a post-authentication policy you can define a set of rules for a device in order to keep the session active. If the device ceases to comply with the policy then connectivity to the NetScaler Gateway is terminated.

Client security expressions are added to session profiles to specify settings that end user devices need to maintain in order to connect to NetScaler Gateway.

Commonly, an administrator will define a quarantine group for use in conjunction with a post-authentication policy. This allows end users to be placed into a restricted access group if they fail to meet the requirements of the client security expression. For example, end users who are not running the process Symantec.exe are placed into a quarantine group which has limited access to corporate resources.

To Configure a Post-Authentication Policy

1. In the configuration utility, in the navigation pane, click **NetScaler Gateway > Policies > Session**.
2. On the **Session Policies** tab, click **Add**.
3. In **Name**, type a name for the policy.
4. Next to **Profile**, click the + sign.
5. In **Name**, type a name for the profile.
6. On the **Security** tab, click **Advanced Settings**.
7. Under **Client Security Check String**, click **Override Global** and then click **Expression Editor**.
8. Configure the Client Security expression and then click **Done**.
9. Under **Quarantine Group**, select a group.
10. In the **Error Message** field, type the message you want users to receive if the post-authentication scan fails.
11. Under **Authorization Groups**, click **Override Global**, select a group, click the right arrow to move the group to the **Configured** side, click **Create**.
12. In the **Create NetScaler Gateway Session Policy** dialog box, under **Expression**, click the **Saved Policy Expressions** drop down list, select **ns_true**, click **Create**.

Post-Authentication Scan Frequency Configuration

When post-authentication scans are implemented to ensure that client devices continue to meet access requirements, it is possible to configure how frequently the scans are run on the end user device.

The frequency of the scan is configured using the **Frequency (min)** option within expressions.



Running continuous scans on an end user device to ensure that it continues to match specified rules is only available when connections are made using the NetScaler Gateway plug-in; this is not available for any other access method. Specifying a error message if the post-authentication scan fails is helpful to end users. Often administrators instruct a user to contact an IT helpdesk or provide other information that will aid the user in re-establishing a connection to the network.

For example, in order to ensure that Sophos Antivirus auto-protect features are in use it is imperative that the Sophos executable SavService.exe is running. You might want to ensure that this process is running by checking the end user's device every 60 minutes. If the SavService.exe is no longer running then the session to the NetScaler Gateway will be terminated.

Post-authentication scans often incorporate checking the component that is used in Preauthentication policies. This helps to keep the configuration simple and consistent: for example, when checking for the presence of an antivirus vendor's process running on the user's device.



Post-authentication continuous scanning is only available for connections made using the NetScaler Gateway plug-in.

End-User Logon Options Evaluation

If a user skips the endpoint analysis scan, then this is treated as a failed scan. It is possible to configure NetScaler Gateway to provide a limited set of resources to the user based on the result of the scan. For example, a general policy can be implemented to guard against users of this kind by forcing all users to log on using ICA Proxy mode unless the process Symantec.exe is running on the end user's device.

To simplify the configuration in this type of scenario a default restrictive session policy can be configured, with an override policy that allows use of the plug-in if the user passes the endpoint analysis scan.

To complete this process, the following steps are recommended:

1. Configure restrictive settings globally.
2. Create a new session policy and profile that allows the use of the NetScaler Gateway plug-in.
3. Add an expression to the session policy to check for the presence of the Symantec.exe process:

```
CLIENT.APPLICATION.PROCESS(Symantec.exe) EXISTS
```

4. Bind the session policy to the NetScaler Gateway Virtual Server.

When a user logs on, the session policy bound to the NetScaler Gateway virtual server will be evaluated first; if the endpoint analysis scan fails then the policy will not be applied and the user will instead be given the settings that are specified globally.

Commonly, in XenApp and XenDesktop deployments ICA proxy is the primary mode of user access and the NetScaler Gateway plug-in provides a subset of users with the ability to connect. You have several options for how to restrict use of the plug-in, including:

- Assign a session policy to a specific group that enables the plug-in period.
- Use endpoint analysis to identify machines that should be enabled for use by plug-in connections.

Quarantine Groups

Quarantine groups can be specified to enable end users who have failed an endpoint scan to still gain access to a restricted set of resources. A common use of this is to allow quarantined users to access resources hosted by Web interface only, such as in a XenApp or XenDesktop environment.

For example, an endpoint analysis scan can be created to check for a particular process to be running on a device. If that process is running then the user can log on using the NetScaler Gateway plug-in. If it is not running, rather than preventing the user from accessing any resources, you can place the user into a quarantine group which can then access XenApp resources only.

To facilitate this feature a group must be created on the NetScaler Gateway. This group is then selected as a quarantine group within a session policy.



When creating the quarantine group ensure that the name does not match an Active Directory group; if the name matches, the group can receive unexpected results when logging in to the NetScaler Gateway.

Endpoint Analysis Troubleshooting

In a Bring-Your-Own Device (BYOD) model, endpoint analysis can be difficult to implement, as often an administrator cannot define a standard for the device that is requesting access. A common workaround for this is to use endpoint analysis to scan for running processes from major antivirus software vendors, ensuring that the device has active antivirus software.

Be aware that it is possible to make endpoint analysis too complex, and it is possible that a greater number of support calls can be generated. Depending on the endpoint analysis type, a user might be denied the ability to log on while being advised that their device does not meet the minimum requirement. In this scenario, a help desk call will likely be generated. To minimize the support effort, it is important to educate end users on minimum device requirements.

Discussion Question

How do Preauthentication policies behave differently than post-authentication policies? Discuss with your classmates and instructor.

Not for resale or distribution

Module 9

End User Access and Experience (includes RDP Proxy)

9

Not for resale or distribution

Not for resale or distribution

End User Access and Experience Exercises (includes RDP Proxy)

Overview

Citrix NetScaler Gateway allows for several different user connection methods, each allowing secure access to corporate-hosted resources. NetScaler Gateway incorporates SmartAccess capabilities to allow for automatic detection of end user access methods. For example, a corporate-owned laptop may be allowed to use local printing functionality, but any other device should not be allowed to have this same functionality. These types of configurations can be completed using Citrix NetScaler Gateway.

Mobile devices are fully supported for use with NetScaler Gateway, including iOS, Android, and Blackberry devices. Citrix Receiver is available for download from the relevant marketplace or application store.

At the end of this module, you will be able to:

- Explain how the NetScaler Gateway feature is used to allow end user connections, including multiple logon options.
- Explain how you can enable end users with access to published applications and virtual desktops using the NetScaler Gateway plug-in with Citrix Receiver.
- Enable clientless access to allow end users with access without requiring them to install user software.
- Enable access scenario fallback to allow an end-user device to fall back from the NetScaler Gateway plug-in to the Web Interface (using Citrix Receiver) if the end-user device does not pass the initial endpoint analysis scan.
- Configure end user device connections by defining which resources end users can access in the internal network.
- Configure end-user cleanup to remove sensitive information from the end-user device upon VPN session termination.
- Enable split tunneling to prevent the NetScaler Gateway plug-in from sending unnecessary network traffic to NetScaler Gateway.
- Configure a DNS server within a session profile.
- Configure the access level and which applications users are allowed to access in the secure network.

Connection Methods

Citrix NetScaler Gateway allows for a modular approach to end user connectivity, catering for various scenarios. For example, a customer might require external connectivity to be made available

for all users with Citrix Receiver, but the network team might also need VPN-based access so that they can connect to internally hosted servers.

NetScaler Gateway supports the following connectivity methods:

Citrix Receiver for Web	Allows connectivity using a web browser
NetScaler Gateway Plug-in for Windows	Allows the establishment of an SSL VPN session between the Windows endpoint device and the NetScaler Gateway
NetScaler Gateway Plug-in for Mac OS X	Allows the establishment of an SSL VPN session between the Mac OS X endpoint device and the NetScaler Gateway
NetScaler Gateway Plug-in for Java	Allows Macintosh, Linux, UNIX, or Windows-based devices to connect using a Java-based client
Clientless Access	Allows end users to access resources without installing a client

These connection methods are not mutually exclusive. Through the use of session profiles and policies, it is possible to present an end user with a choice screen to select whether to connect through a full VPN connection using the NetScaler Gateway plug-in, clientless access, or the Citrix Receiver for Web Interface access. Many organizations use Citrix Receiver for Web as a default setup, while also allowing a certain Active Directory group to have access to a VPN connection through the NetScaler Gateway plug-in.

Organizations can configure and restrict VPN connections to only permit certain IP addresses or networks to be reachable through the VPN tunnel. Note that this functionality is only available when use of the NetScaler Gateway plug-in is set.

Secure Tunnel Establishment

The Citrix NetScaler Gateway plug-in allows an end user device to establish a secure tunnel to the corporate network. This can be further controlled and secured by defining only certain IP addresses or subnets that should be intercepted by the plug-in.

The connection process for an end user connecting through the NetScaler Gateway plug-in involves the following steps:

1. User authenticates using one of the following methods:

- Connecting through the NetScaler Gateway plug-in: Authentication credentials are entered into the plug-in and then the details are sent to the NetScaler Gateway virtual server using HTTPS, port 443 by default.
 - Logging on to the NetScaler Gateway web portal, which can automatically invoke the NetScaler Gateway plug-in. A user can also select the option to start "Network Access" when presented with the client choices screen.
2. Configuration data is sent to the NetScaler Gateway plug-in detailing which private networks are to be secured and redirected over the secure tunnel.
 3. Any TCP, UDP or ICMP echo traffic that is destined for the secured networks is transmitted over SSL to the NetScaler Gateway.
 4. The NetScaler Gateway receives the data and then forwards the data on the internal host and relevant port.

Once the tunnel is established it will remain open until the timeout value has been reached; these values are set within session policies on the NetScaler Gateway.

Enabling the client choices option will provide a flexible approach for hosted application and desktop access. End users can use Receiver to securely start sessions using the Web Interface, or start a full VPN session if additional network access is required.

Network Firewalls and Proxies

User devices are often located within a different network than the NetScaler Gateway with NAT firewalls in between the endpoint and the NetScaler Gateway. NAT firewalls maintain a connection table that allows packets to be routed between hosts. In the case of NetScaler Gateway a reverse NAT translation table is used to allow NetScaler Gateway to match connections and send packets back over the tunnel to the end-user device with the correct port numbers so that the packets return to the correct application.

Ensure that any firewalls have granted permission for both HTTP and HTTPS protocols to the NetScaler Gateway virtual IP address. This ensures that if users mistakenly type the URL with HTTP instead of HTTPS the connection can still be made and the NetScaler Gateway can redirect the connection to HTTPS automatically without user intervention.

Secure Tunnel Termination

NetScaler Gateway SSL tunnels are used to secure traffic between an end user device and the corporate network. The tunnel is terminated on the NetScaler Gateway, then all packets that are destined for the private network are decrypted and re-generated by the NetScaler Gateway, including headers that show the source IP address as the NetScaler Gateway's address range or the address assigned to the client through an IP address pool that has been configured on the NetScaler Gateway.



When completing network tracing be aware that by default packets will appear to be coming from the NetScaler Gateway MIP or SNIP address, rather than the remote endpoint's IP address. To override this behavior a pool of IP addresses can be allocated for VPN users; this feature is called IP Pooling and can be set using the Intranet IPs tab of either a user, group or virtual server.

A company has 100 remote users who require a VPN connection. In addition, a third party contractor requires VPN access. Auditing information must be available to identify the connections from the contractor.

To achieve this, allow the default behavior of packets appearing to be from the NetScaler Gateway for the 100 users while adding the user in to the NetScaler Gateway "users" section and configuring IP address pooling to assign an IP address to the single user. The IP address will only be used by the specified user, allowing for accountability by monitoring the IP address.

NetScaler Gateway Plug-in Support

You can use one of several options for installing the NetScaler Gateway plug-in:

- Deployed through Software Distribution methods, for example SCCM or Active Directory
- Downloaded from the NetScaler Gateway when trying to connect the VPN
- Downloaded from the Citrix Downloads site

Once an end user has successfully logged on to the NetScaler Gateway and chooses to start a VPN connection using the Network Access option, client detection will begin to detect whether a suitable plug-in is installed. If the correct plug-in is present, the session will be established. If there is no plug-in or the plug-in is incorrect, the NetScaler Gateway will provide the plug-in for installation.



The plug-in installation requires the user to have administrative rights over the device; if this is not possible then consider using the Java plug-in. This restriction is only applicable to installations and some major upgrades, but not minor upgrades.

When the plug-in is installed, you might need to provide the NetScaler Gateway web address to the user; this should be provided in the format of `https://AccessGatewayFQDN`. For example, `https://portal.example.com`.



If the NetScaler Gateway is not listening on port 443 and is instead using an alternative port this must be specified in the address string. For example, if the gateway is `portal.example.com` and uses port 444, then the address would be `https://portal.example.com:444`.

Due to this requirement it is often recommended to continue the use of port 443 for the NetScaler Gateway; this ensures that the URL is easily accessible for users and the administrator will not need to remember specific ports.

If Endpoint Analysis is implemented it is important to ensure that the end-user device meets the requirements in order to establish a VPN connection. Failure to meet the criteria could mean that the session cannot be established.

Windows firewall will automatically be configured by the plug-in installation; however, third-party solutions will likely require manual configuration.

Software Firewalls

The end user's device must be able to connect to the NetScaler Gateway using the specified port (443 by default). However, be aware that software and hardware firewalls can block the connection, so you might need to add exceptions for the connection. The NetScaler Gateway plug-in installer will automatically add exceptions to the Windows firewall; any other client firewall software might require manual additions of exceptions.

NetScaler Gateway Plug-in Integration with Citrix Receiver

Citrix Receiver provides a framework for many Citrix plug-ins, including the Citrix offline plug-in, Online plug-in and the NetScaler Gateway plug-in.

When the NetScaler Gateway plug-in is installed, end users can select the option to log on to the NetScaler Gateway using Citrix Receiver. Once the logon has been completed successfully you can use single sign-on to pass the credentials through to a Receiver home page, presenting the user's applications and desktops without reporting re-authentication.

For example, when logging on to a laptop, an end user can right-click **Citrix Receiver** and select **Advanced Preferences**. This will allow the user to click the NetScaler Gateway Settings plug-in and select **logon**. A connection to the NetScaler Gateway will be initiated.

Citrix Receiver Home Page

Through the use of the Citrix Receiver home page setting it is possible to open a web browser-based session to Citrix Web Interface or StoreFront, allowing easy access to hosted applications and desktops.

The Receiver home page setting can be defined in either a session profile or set globally.



Setting items globally will ensure that they always apply unless a specific option is set within a session profile. Often base settings that must apply to all users are defined globally and then other settings are added in to session profiles.

To Configure the Citrix Receiver Home Page

To configure the Citrix Receiver Home Page setting globally, complete the following steps:

1. In the configuration utility, in the navigation pane, expand **NetScaler Gateway** and then click **Global Settings**.
2. In the **Global Settings** pane, under **Settings**, click **Change Global Settings**.

3. In the **Global NetScaler Gateway Settings** dialog box, click the **Published Applications** tab.
4. In **Citrix Receiver Home Page**, type the Web address for the Receiver and Receiver for web home page and then click **OK**.

Application of the Citrix Receiver Theme to the Logon Page

By default NetScaler Gateway virtual servers use the legacy "carbon" access interface. With the release of Citrix Receiver, a new access interface has been created. The theme files are included on the NetScaler appliance by default and can be applied using the command-line interface or using the Configuration Utility (NetScaler Gateway > Customize Access Interface > Upload the Access Interface).

For more information about applying the Citrix Receiver theme to the logon page, see Citrix product documentation at <http://docs.citrix.com>.

NetScaler Gateway Plug-in Selection for End Users

Through the use of session policies it is possible to set which plug-in type a user connection will use; this can either be the Windows/Mac OS X plug-in or the Java-based plug-in.

The Java plug-in requires JRE 1.4.2 or later in order to function on the end user device. This allows the client to function on many platforms, including Mac OS X, Linux, Internet Explorer, Safari and Firefox.

Defining the plug-in type in a session policy will ensure that the correct plug-in is used. It is also possible to specify different plug-in types for different groups of users; this can be completed by using session profiles and binding them at the relevant bind point, for example to a group. This scenario can also be completed using the following bind points:

- User
- Group
- Virtual Server

If the settings are not defined in a session policy then the global value will be used.



The NetScaler Gateway Java plug-in does not require installation by an administrator. As such this plug-in might be best suited for scenarios in which the end users do not have the ability to install software: for example, users working from Internet cafes. Alternatively some organizations use separate virtual servers for connections that originate from contractors. In this instance it might be preferable to use the Java plug-in as you might not be able to install the full plug-in on the end user device.

The NetScaler Gateway plug-in offers the most features and should be the preferred option when possible.

NetScaler Gateway Plug-in Deployment, Upgrading and Removal from Active Directory

Active Directory group policy objects are commonly used to distribute software to end user devices, as this provides a central administrative location as well as granular control over the distribution of the plug-in.

The steps to deploy the plug-in are:

1. Download the installer and then extract the MSI.
2. Create a network share to be used as a software distribution point.
3. Place the extracted MSI in to the software distribution point.
4. Create a Group Policy Object and configure software distribution of the plug-in, using the agee.msi as the software installer.
5. Link the Group Policy Object to the relevant Active Directory location.



Other software deployment technologies can also be used, for example Microsoft's SCCM product.

The plug-in is also available for download from MyCitrix.



Centrally deploying the plug-in to end user devices can greatly save on administration; however, restrictions common to BYOD scenarios can problematize central deployment. In this type of scenario the user is often an administrator of the device and as such can download and install the plug-in from the NetScaler Gateway once they have successfully authenticated.



When trying to deploy the plug-in using Active Directory you must first extract the agee.msi from the bundled installer that the NetScaler Gateway provides. To do this, run the following command on the saved installer:

```
Nsvpnc_setup /c
```

The extracted agee.msi can then be imported in to an Active Directory group policy object ready for distribution.

NetScaler Gateway Plug-in Installation Troubleshooting Using Active Directory

After a Group Policy object has been configured to deploy the NetScaler Gateway plug-in to devices, the software installation will be attempted when the device is next started. In certain circumstances this can fail to initialize, instead allowing the end user to log on when the software has not been installed.

The event logs, which are located on the end-user device, provide a good source of information regarding software installation failures; however, for more detailed information, you can enable Windows Installer Logging, as this will provide granular information regarding the software installation.



By default, Group Policy will be applied in the background once the device has started, presenting the user logon screen before Group Policy has fully completed processing. In rare instances this can cause issues with the application of certain Group Policy settings, as the user might be able to log on before new Group Policy settings have been fully applied. This behavior can be modified by setting the **Always wait for the network at computer start up and logon** option to **Enabled**. This can be found within the Group Policy Location Computer Configuration\Administrative Templates\System\Logon. Windows XP also has a similar feature known as Fast User Switching. This should be disabled if the plug-in installation is failing.

Client Ports

Administrators often expect the "client port" and "server port" to both show a port number of 443 (HTTPS); however the "client port" column will show a different number because TCP connections are assigned a random source port from the client to the server.

NetScaler Gateway Plug-in Connection Configuration

You can define end user device connection settings, ensuring that connections meet security requirements and also only permitting user access to certain resources. These options include:

- Defining the domains to which users are allowed access.
- Configuring IP addresses for users, including address pools (intranet IPs).
- Configuring time-out settings.
- Configuring single sign-on.
- Configuring client interception.
- Configuring split tunneling.
- Configuring connections through a proxy server.
- Configuring user software to connect through NetScaler Gateway.

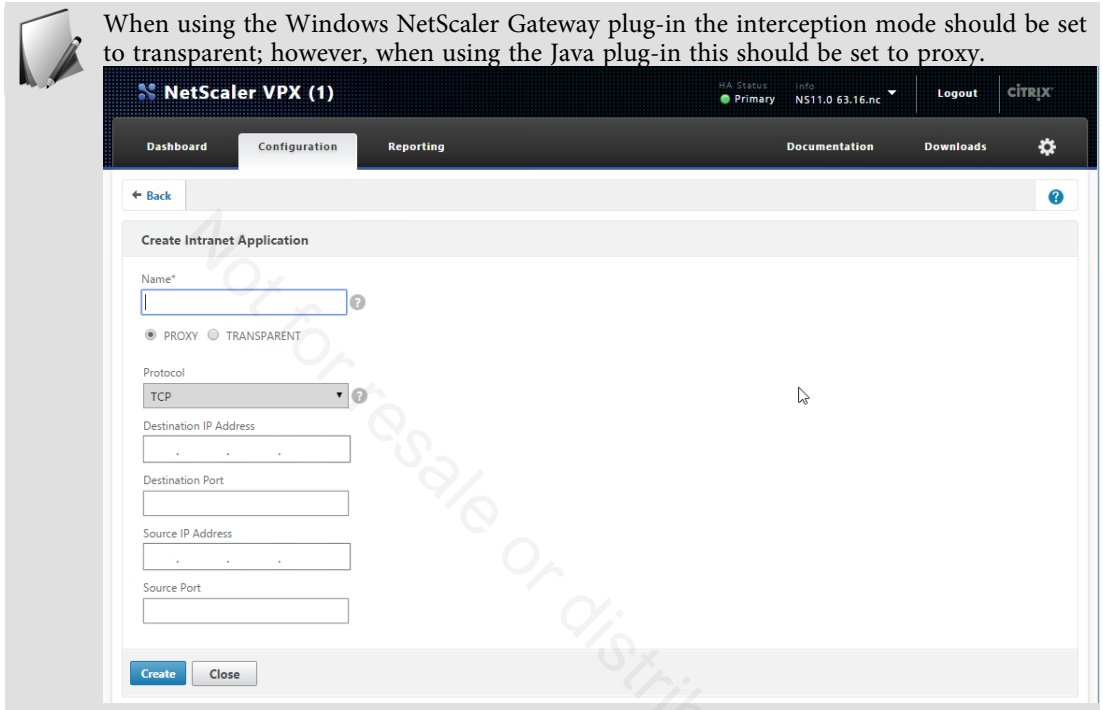
Internal Network Resources Connection

The NetScaler Gateway plug-in can be configured to direct all network traffic to the NetScaler Gateway appliance, or only traffic that is destined for internal network to be intercepted; this option is set using split tunneling.

If split tunneling is enabled then only traffic that is destined for an internal network is sent to the NetScaler Gateway. You define which IP addresses or subnets are internal by using the intranet applications configuration settings.

With split tunneling disabled, all traffic is sent to the NetScaler Gateway. The traffic is then examined and authorization policies determine if the traffic is allowed to pass through the NetScaler Gateway.

When using the Windows NetScaler Gateway plug-in the interception mode should be set to transparent; however, when using the Java plug-in this should be set to proxy.



The screenshot displays the NetScaler VPX (1) Configuration page. The main content area is titled 'Create Intranet Application'. It features a 'Name*' input field, two radio buttons for 'PROXY' (selected) and 'TRANSPARENT', a 'Protocol' dropdown menu set to 'TCP', and five input fields for 'Destination IP Address', 'Destination Port', 'Source IP Address', and 'Source Port'. At the bottom of the form are 'Create' and 'Close' buttons. The top navigation bar includes 'Dashboard', 'Configuration', and 'Reporting' tabs, along with 'Documentation' and 'Downloads' links. The Citrix logo is visible in the top right corner.

You can define an intranet application in the configuration utility by using the following steps:

1. In the navigation pane, expand **NetScaler Gateway**, expand **Resources**, then click **Intranet Applications**.
2. In the details pane, click **Add**.
3. Complete the parameters for allowing network access then click **Create**.

End-User Connections Proxy Support Enablement

Companies often deploy proxy servers within a corporate network to enhance security and to provide auditing and content filtering. The NetScaler Gateway plug-in supports setting the relevant proxy settings for an end user once a connection has been established.

You specify the proxy server address details within a session policy or globally; these settings are then incorporated into the user's web browser once the connection is established.



Companies often set proxy settings for end users through Group Policy objects when the user logs on. This ensures that traffic is automatically sent through the proxy without the user needing to configure any settings. Enabling proxy support is typically only used when split tunneling is disabled and all traffic is forwarded through the NetScaler Gateway and the corporate network.

When a device connects through VPN, it is essential that the user specifies the proxy server address in order to gain Internet access. NetScaler Gateway's proxy settings function can automate this task, in turn reducing support calls requesting assistance with setting proxy settings.

Setting the proxy is supported for HTTP, HTTPS, FTP and SOCKS. Proxy exceptions can also be added.

Not for resale or distribution

Session Policies

Session policies allow administrators to configure client-side attributes when the session is established after the end user is authenticated. Client-side attributes are settings that affect the user experience and can include timeout settings, client type and split tunneling.

Session Profiles

Session profiles specify the actions or client attributes that will be applied to the session if the policy expression conditions are met. The following settings can be configured in session profiles:

- Network configuration
 - DNS virtual server
 - WINS server IP address
 - Kill Connections
 - Usage of Mapped IP as client IP address
 - Intranet IP address (IP pooling)
 - Intranet IP DNS Suffix
 - Spoof Intranet IP address
 - HTTP ports
 - Forced timeout
 - Forced time-out warning (mins)
- Client experience
 - NetScaler Gateway home page
 - URL for Web-Based Email
 - Split tunneling
 - Session timeout and Client idle timeout
 - Clientless Access
 - Plug-in type (Windows/MAC OS X or Java)
 - Windows Plugin Upgrade
 - Linux Plugin Upgrade
 - MAC Plugin Upgrade
 - Single sign-on
 - KCD Account
 - Logon and logoff scripts
 - Client debug settings
 - Split DNS
 - Local LAN access
 - Client options and the Configuration window
 - Client cleanup behavior
 - Proxy settings
 - XenDesktop or XenApp integration
- Security

- Default authorization action
- Secure Browse
- Smartgroup
- Quarantine groups
- Error Message for quarantine.
- Authorization groups
- Groups Allowed To Login
- NetScaler Gateway and SmartAccess
 - ICA proxy mode
 - Web interface or StoreFront address
 - Web Interface Address Type (IPv4 or IPv6)
 - Web Interface Portal Mode
 - XenDesktop or XenApp Active Directory Single Sign-on domain
 - Citrix Receiver Home Page
 - Account Services Address

Policy Bindings

Session policies can be bound at the following levels:

- Global
- Virtual servers
- Group
- User



- Unless the global binding has a higher priority, settings modified by session policies override settings configured at the global and virtual server levels.
- Any attributes or parameters not configured in a session policy will be based on the settings configured at the higher order bind points.
- Any remaining attributes or parameters not configured in a session policy will be based on the settings configured at the global levels.

VPN Virtual Server Session Policy Binding			
Priority	Policy Name	Expression	Profile
100	SETPNPARAMS_POL	ns_true	SETPNPARAMS_ACT
58000	UG_VPN_SPoL_172.17.17.20	ns_true	UG_VPN_SAct_172.17.17.20

Creating Session Policies

1. You can use the following steps to create a session policy in the NetScaler Gateway Policy Manager.
 - a. Click the + sign to the right of **NetScaler Gateway Global**.
 - b. Click the drop down menu under **Bind Point** and select **Other Policies**.
 - c. In the **Connection Type** menu, select **Session Policies** then click **Continue**.
 - d. Click **Add Binding**, then click in the field below **Select Policy**.
 - e. Click **Add** then type a name for the policy.
 - f. Select an existing **Profile** or create a new one.
 - g. Select **ns_true** in the drop-down list box in **Saved Policy Expressions**.
 - h. Click **Create, Select, Bind, Done**, then click **Done**.
2. You can use the following steps to create a Session Profile with creating a Session Policy in the NetScaler Gateway Policy Manager.
 - a. When creating a new Session Policy, click the + sign next to Profile.
 - b. Type a name for the new Profile.
 - c. Specify the Profile settings.
 - d. Click **Create**, then **OK**, then **Done**.

Credential Passing

When redirecting users to a website or web application that supports non-form based authentication, such as SharePoint, credential passing enables transparent end user authentication using the NetScaler Gateway or an external authentication server.

Credential Passing Configuration

Credential passing can be configured to use either the primary or secondary user credentials. Credential passing requires the following elements:

- Double-source authentication (if secondary user credentials are used)
- Default HTTP port set to 80
- Single sign-on enabled (either globally or per session policy)

Single Sign-on for File Shares

When a user clicks a file share link on the NetScaler Gateway portal, the username and password used to log on to NetScaler Gateway are now automatically used to sign on to the file share.

Configuring Global Credential Passing

1. You can use the following steps to enable credential passing globally in the NetScaler Configuration Utility:
 - a. Select the **NetScaler Gateway** node and click **Global Settings**.
 - b. Click the **Change Global Settings** link.
 - c. Select the **Client Experience** tab and click **Single Sign-on to Web Applications**.
 - d. Select the credential information used for authentication from the **Credential Index** drop-down list box and then click **OK**.

Configuring the Default Home Page for SSL VPN and Clientless VPN Connection

1. You can use the following steps to configure an alternate landing page for web applications when credential passing is enabled.
 - a. From the **NetScaler Gateway Policy Manager**, under **NetScaler Gateway Global**, click in the Session Policy field.
 - b. Highlight the Session Policy, then click the Edit button and select **Edit Profile**. The **Configure NetScaler Gateway Session Profile** window appears.
 - c. Select the **Client Experience** tab and select **Override Global** and **Display Home Page** next to Home Page.
 - d. Type the web address for the default home page then click **OK**, **Done** and then click **Done** again.

Timeout Settings

NetScaler Gateway allows you to configure timeout settings, which can improve security by ending VPN user sessions after set time limits are reached.

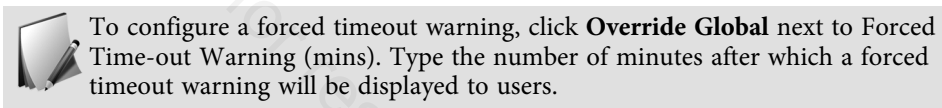
Timeout	Description
Session	The duration after which an idle VPN session is terminated. A VPN session is idle when there is no traffic traveling through the VPN tunnel.
Client Idle	The duration after which an end user's VPN session is terminated if there is no keyboard or mouse activity on the client device.
Forced	The duration after which a VPN session is terminated regardless of any activity on the VPN session or client device. NetScaler Gateway can be configured to display a warning before the session is terminated.

Configuring Session and Client Idle Timeout Values

1. You can use the following steps to configure timeout settings in the NetScaler Gateway Policy Manager.
 - a. Click the + sign to the right of **NetScaler Gateway Global**.
 - b. Click the drop down menu under **Bind Point** and select **Other Policies**.
 - c. In the **Connection Type** menu, select **Session Policies** then click **Continue**.
 - d. Click **Add Binding**, then click in the field below **Select Policy**.
 - e. Click **Add** then type a name for the policy.
 - f. Type a name for the Session Policy and then click + next to the **Profile** field to open the **Create NetScaler Gateway Session Profile** window.
 - g. Type a name for the Request Profile.
 - h. Select the **Client Experience** tab.
 - i. For a session timeout, select **Override Global** next to **Session Time-out (mins)** and type a session timeout value. For a client idle timeout, select **Override Global** next to **Client Idle Time-out (mins)** and type a client idle timeout value.
 - j. Click **Create**.
 - k. Select **ns_true** in the drop-down list box in **Saved Policy Expressions** then click **Create**.
 - l. Click **Select, Bind, Done** and then click **Done**.
 - m. Bind the policy to either the NetScaler Gateway Global, Virtual Servers, Groups or Users node under the NetScaler Gateway node or Global Bindings under the Session Node.
 - n. Set a priority for the bound policy.

Configuring Forced Timeout Values

1. You can use the following steps to configure forced timeout settings in the NetScaler Gateway Policy Manager.
 - a. Click the + sign to the right of **NetScaler Gateway Global**.
 - b. Click the drop down menu under **Bind Point** and select **Other Policies**.
 - c. In the **Connection Type** menu, select **Session Policies** then click **Continue**.
 - d. Click **Add Binding**, then click in the field below **Select Policy**.
 - e. Click **Add** then type a name for the policy.
 - f. Click the + sign next to the **Profile** field. The **Create NetScaler Gateway Session Profile** window opens.
 - g. Type a name for the Request Profile and then click **Advanced Settings** under Network Configuration.
 - h. Click **Override Global** next to Forced Timeout and type the number of minutes after which a forced session timeout will occur.



- i. Click **Create**.
- j. Select **ns_true** in the drop-down list in Saved Policy Expressions and then click **Create**.
- k. Click **Select, Bind, Done** and then click **Done**.
- l. Bind the policy to either the NetScaler Gateway Global, Virtual Servers, Groups or Users node under the NetScaler Gateway node or Global Bindings under the Session Node.

Split Tunneling

Split tunneling allows the NetScaler Gateway Plugins to distinguish between VPN session traffic and other network traffic.

Split tunneling allows end users to remotely access internal resources remotely through NetScaler Gateway while also accessing the Internet through their local network. This reduces traffic on the NetScaler Gateway, which can improve network performance.

Split tunneling works in conjunction with Intranet applications to determine which traffic is intercepted and forwarded to the NetScaler Gateway for evaluation.

When split tunneling is disabled, all network traffic on the client device is forced through the NetScaler Gateway tunnel. Disabling split tunneling may be necessary to ensure compliance with network security requirements. For example, split tunneling should be disabled when the NetScaler

Gateway is being deployed internally and visitors should be allowed access to only a limited subset of network resources.

Configuring Local LAN Access

By default, when split tunneling is disabled, end users cannot access local LAN resources. However, NetScaler Gateway allows administrators to disable split tunneling while granting end users access to their local LAN as well as the VPN tunnel. This allows users to access printers and other resources on the local LAN even though split tunneling is disabled.



This feature is only used when split tunneling is disabled. If local LAN access is enabled, users must enable the feature within the client.

Restricting Local LAN Access to Private Network IP Addresses

By default, when local LAN access is enabled, users can access any IP address with the LAN IP address space. An administrator can improve security by explicitly allowing users to access only those resources hosted on networks that are local area network ranges:

- 10.0.0.0/8
- 172.16.0.0/12
- 192.168.0.0/16

This prevents the NetScaler Gateway from treating other networks, including those added dynamically after the VPN session is established, as local area networks. This feature requires local LAN access to be enabled. An administrator enables access to private network IP addresses in the advanced client experience settings.

Configuring Split Tunneling

1. You can use the following steps to configure split tunneling in the NetScaler Gateway Policy Manager.
 - a. Click the + sign to the right of **NetScaler Gateway Global**.
 - b. Click the drop down menu under **Bind Point** and select **Other Policies**.
 - c. In the **Connection Type** menu, select **Session Policies** then click **Continue**.
 - d. Click **Add Binding**, then click in the field below **Select Policy**.
 - e. Click the + sign next to the **Profile** field. The **Create NetScaler Gateway Session Profile** window opens.
 - f. Type a name for the Profile.

- g. Select **Override Global** next to **Split Tunnel** under **Client Experience**.
- h. Configure split tunneling by selecting **ON**, **OFF** or **REVERSE** in the Split Tunnel drop-down list box.
- i. Select **ns_true** in the drop-down list in Saved Policy Expressions and then click **Create**.
- j. Click **Select**, **Bind**, **Done** and then click **Done**.
- k. Bind the policy to either the NetScaler Gateway Global, Virtual Servers, Groups or Users node under the NetScaler Gateway node or Global Bindings under the Session Node.

Reverse Split Tunneling

Reverse split tunneling allows you to define traffic that should not be intercepted. In this implementation, intranet applications are used to define which traffic the NetScaler Gateway does not intercept.



If split tunneling is disabled, all traffic is sent to the NetScaler Gateway; authorization policies are then used to control which traffic is allowed or dropped.

Split Tunneling Considerations

Before implementing split tunneling consider it carefully. Split tunneling provides clear benefits by reducing the amount of traffic sent to the Gateway as well as allowing end users to access local data. However, split tunneling can also be perceived as a security risk because end users can interact with both a secured network through the VPN, and a potentially unsecured local network, over which the administrator might not have any control.

When providing access for third-party contractors it is common practice to disable split tunneling and also only allow specific traffic using authorization policies; this helps to prevent data leakage and adds another layer of security. For example, a remote support provider might only require remote desktop access to a single server; in this case, an authorization policy could be used to ensure only that server is accessible over the VPN connection.

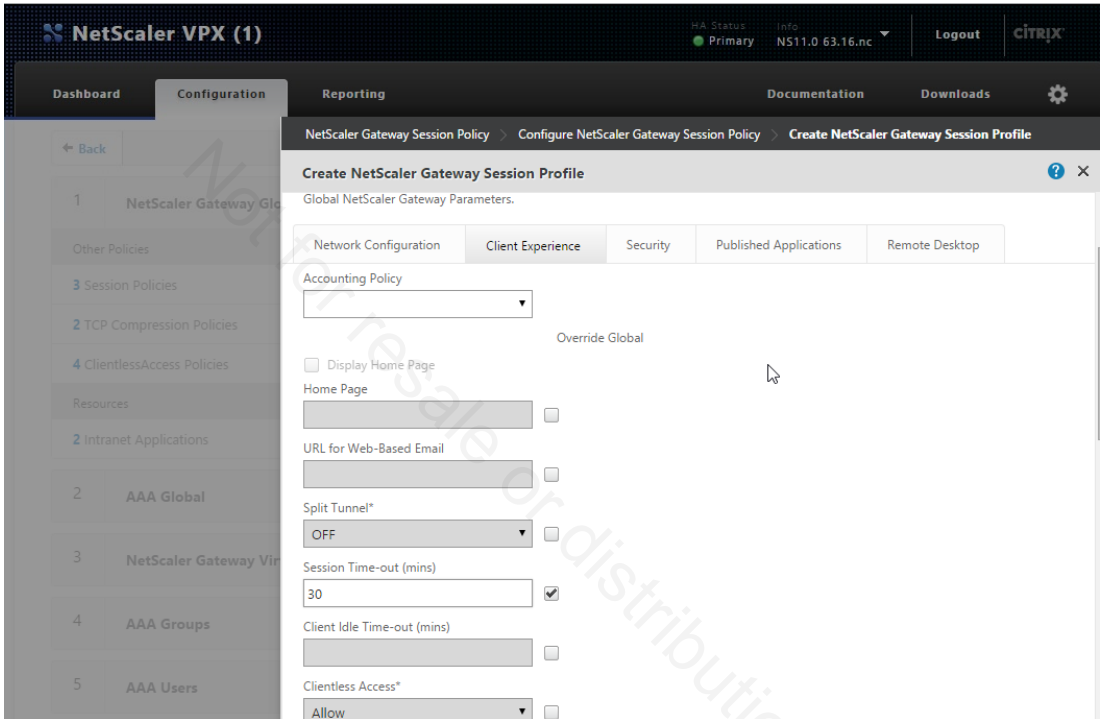
To configure split tunneling:

1. In the **NetScaler Gateway Policy Manager**, under **NetScaler Gateway Global > Other Policies**, click **x Session Policies** and then select a Session Policy.
2. Click **Edit**.
3. Click **Edit Profile**.
4. On the **Client Experience** tab, next to **Split Tunnel**, select **Global Override**, select an option and then click **OK**, then **Done**, then **Done** again.

Timeout Settings Configuration

Timeout settings allows you to end user sessions after pre-defined periods of time in one of three ways:

- Forced timeout - This ends the user's session even if it is active.
- Session timeout - A user's session will be ended if there is no traffic crossing the VPN.
- Client Idle timeout - This defines how many minutes can pass without keyboard or mouse activity. Once the timer is reached the session is disconnected.



The screenshot shows the NetScaler VPX (1) configuration interface. The main navigation bar includes Dashboard, Configuration, Reporting, Documentation, and Downloads. The current page is 'Create NetScaler Gateway Session Profile' under 'Configure NetScaler Gateway Session Policy'. The page is divided into tabs: Network Configuration, Client Experience, Security, Published Applications, and Remote Desktop. The 'Session Time-out (mins)' field is set to 30 and has a checkmark next to it. Other fields include Accounting Policy, Display Home Page, Home Page, URL for Web-Based Email, Split Tunnel* (OFF), Client Idle Time-out (mins), and Clientless Access* (Allow).



Timeouts are often applied to ensure compliance with security policies. For example, if a user has been inactive for 60 minutes then they might have left their laptop unattended; this could be a security risk.

Using timeouts also can ensure license availability for the users who need one.



Be careful using forced timeouts as this can severely impact users' work, as their VPN connection will be terminated without any option to prevent the termination.

Specifying a timeout at the global level and then overriding the timeout using session policies is a good way to simplify management of timeouts.



Setting any of the values to 0 will disable the timeout.

To Configure Forced Timeouts

Forced timeouts can be set by either a global setting or by a session policy. Often administrators will set a global setting that is used by the bulk of users, with an override timeout in a session policy for any users or groups that require a different timeout. The created session policy can also be bound to a NetScaler Gateway virtual server.

1. In the configuration utility, in the navigation pane, expand **NetScaler Gateway**, then click **Global Settings**.
2. In the details pane, under Settings, click **Change global settings**.
3. On the Network Configuration tab, click **Advanced Settings**.
4. Under Timeouts, in **Forced Timeout**, type the number of minutes that users can stay connected.
5. In **Forced Time-out Warning (mins)**, type the number of minutes before users are warned that the connection will be disconnected and then click **OK** twice.

To Configure a Forced Timeout Within a Session Policy

If you want to have further control over who receives the forced timeout, create a session policy and then apply the policy to a user or group.

1. In the configuration utility, in the navigation pane, click **NetScaler Gateway**, then click **NetScaler Gateway Policy Manager**.
2. Click the + sign to the right of **NetScaler Gateway Global**.
3. Click the drop down menu under **Bind Point** and select **Other Policies**.
4. In the **Connection Type** menu, select **Session Policies** then click **Continue**.
5. Click **Add Binding**, then click in the field below **Select Policy**.
6. Click **Add** then type a name for the Policy.
7. Click the + sign next to the **Profile** field. The **Create NetScaler Gateway Session Profile** window opens.
8. Type a name for the Profile.
9. On the **Network Configuration** tab, click **Advanced Settings**.
10. Click **Override Global** next to **Forced Timeout** type the number of minutes users can stay connected.
11. Next to **Forced Time-out Warning (mins)**, click **Override Global** and type the number of minutes users are warned that the connection is due to be disconnected. Click **Create**.

12. Select **ns_true** in the drop-down list in Saved Policy Expressions and then click **Create**.
13. Click **Select, Bind, Done** and then click **Done**.

Configuring Session or Idle Timeouts

Session timeouts and idle timeouts can be configured using either globally applied settings or session policies.

To Configure a Session or Client Idle Timeout Globally

1. In the configuration utility, in the navigation pane, expand **NetScaler Gateway** and then click **Global Settings**.
2. In the details pane, under **Settings**, click **Change global settings**.
3. Do one or both of the following: On the **Client Experience** tab, in **Session Time-out (mins)**, type the number of minutes. In **Client Idle Time-out (mins)**, type the number of minutes and then click **OK**.

To Configure Session or Client Idle Timeout Settings Using a Session Policy

1. In the configuration utility, in the navigation pane, click **NetScaler Gateway**, then click **NetScaler Gateway Policy Manager**.
2. Click the + sign to the right of **NetScaler Gateway Global**.
3. Click the drop down menu under **Bind Point** and select **Other Policies**.
4. In the **Connection Type** menu, select **Session Policies** then click **Continue**.
5. Click **Add Binding**, then click in the field below **Select Policy**.
6. Click **Add** then type a name for the Policy.
7. Click the + sign next to the **Profile** field. The **Create NetScaler Gateway Session Profile** window opens.
8. Type a name for the Profile.
9. Do one or both of the following:
 - On the **Client Experience** tab, next to **Session Time-out (mins)**, click **Override Global** and then type the number of minutes.
 - Next to **Client Idle Time-out (mins)**, click **Override Global**, type the number of minutes and then click **Create**.
10. Select **ns_true** in the drop-down list in Saved Policy Expressions and then click **Create**.
11. Click **Select, Bind, Done** and then click **Done**.

Client Cleanup

The NetScaler Gateway allows administrators to configure client cleanup for the NetScaler Gateway Plugins for Windows and ActiveX. Client cleanup removes sensitive information from the client device upon VPN session termination. This prevents the misuse of data that may be generated on a client device during a VPN session. For example, an administrator can set the VPN client to automatically remove passwords and auto-complete data stored by the web browser when the user ends the VPN session.

In addition, administrators can configure the client to present a cleanup prompt to end users. The prompt allows end users to select which data is removed from their device.

Cleanup Options

An administrator can select which types of data are removed during the client cleanup process. The following table lists the types of data that can be removed.

Data	Description
Cookies	Removes browser cookies and cache files from internal sites that were downloaded after client logon and before client logoff
Address Bar	Removes all web addresses and history data stored by Internet Explorer
Plugin	Removes the NetScaler Gateway Plugins for Windows and ActiveX
File System Application	Closes the file transfer browser
Application	Closes all applications that have accessed NetScaler Gateway services
Application Data	Removes all temporary files generated by Microsoft Word, Microsoft PowerPoint, and Microsoft Outlook
Client Certificate	Removes the SSL certificate for the NetScaler Gateway VPN virtual server used during the secure authentication process
Auto Complete	Removes all passwords and auto-complete data stored by Internet Explorer after client logon and before client logoff

Data	Description
Cache	Removes all cache data when a user closes the Session.

Cleanup Prompt

An administrator can choose to present a cleanup prompt to end users upon session termination. The cleanup prompt allows end users to control which data is removed during the cleanup process, depending on the options specified by the administrator. For example, an administrator can force the cleanup of application data stored on the end-user device, yet allow users to choose whether or not their browser history is removed.

Cleanup Data Sets

The types of data users can remove when presented with the cleanup prompt are grouped into data sets. Depending on the configuration, end users can have full, partial or no control over which data in a particular data set are removed. The following table lists and describes the data sets presented to the end user in the Citrix Windows Cleanup dialog box.

Data Set	Description
NetScaler Gateway Plugins for Windows and ActiveX	Removes the NetScaler Gateway Plugins for Windows and ActiveX
Secure certificate	Removes the SSL certificate for the NetScaler Gateway VPN virtual server used during the secure authentication process
Applications accessed from the NetScaler Gateway	Closes applications and processes that have accessed NetScaler Gateway services. This prevents the leakage of sensitive information buffered by the applications.
Application data	Removes all temporary files generated by Microsoft Word, Microsoft PowerPoint and Microsoft Outlook

Data Set	Description
Passwords and automatic completion data stored by the Web browser	Removes all passwords and auto-complete data stored by Internet Explorer. Only cookies that were saved by the default browser are removed. The default browser is the one used to log on. However, when native logon is enabled, the cookies saved by Internet Explorer are removed.
History and Web addresses typed in address bar	Removes all web addresses and history data stored by Internet Explorer. All browser windows must be closed before the cleanup.
Cookies and temporary files	Removes all files cached from the remote network or Internet, as well as temporary files and cookies. This data set works with both Internet Explorer and Firefox.

Cleanup Level

An administrator may want to allow some data to be retained on the client device, especially if the client device typically is used by only one person. For example, if all end users connect from their own client devices, an administrator may want to allow the retention of passwords, auto-complete data and history. As an alternative to selecting individual cleanup data sets, end users can select one of the three preconfigured cleanup levels from within the client cleanup prompt. The following table lists and describes the available cleanup levels.

Cleanup Level	Description
None	No data sets are deleted.
Web browser	One or more of the following data sets are deleted: <ul style="list-style-type: none"> • Passwords and auto-complete data stored by browser • History and web addresses typed in the address field • Browser cache cookies and temporary files
All items	All data sets are removed. The client also can be set to remove any temporary data generated by the client device.

Configuring Client Cleanup

1. You can use the following steps to configure client cleanup in the NetScaler Gateway Policy Manager.
 - a. In the configuration utility, in the navigation pane, click **NetScaler Gateway**, then click **NetScaler Gateway Policy Manager**.
 - b. Click the + sign to the right of **NetScaler Gateway Global**.
 - c. Click the drop down menu under **Bind Point** and select **Other Policies**.
 - d. In the **Connection Type** menu, select **Session Policies** then click **Continue**.
 - e. Click **Add Binding**, then click in the field below **Select Policy**.
 - f. Click **Add** then type a name for the Policy.
 - g. Click the + sign next to the **Profile** field. The **Create NetScaler Gateway Session Profile** window opens.
 - h. Type a name for the Profile.
 - i. Click the **Client Experience** tab.
 - j. Click **Override Global** next to Client Cleanup Prompt and select **ON**.



By default, the Client Cleanup Prompt is enabled at the global level.

- k. Click **Advanced Settings** under Client Experience.
- l. Click the **Client Cleanup** tab and select the items for cleanup.
- m. Click **Create**.
- n. Select **ns_true** in the drop-down list in Saved Policy Expressions and then click **Create**.
- o. Click **Select, Bind, Done** and then click **Done**.
- p. Bind the policy to either the NetScaler Gateway Global, Virtual Servers, Groups or Users node under the NetScaler Gateway node or Global Bindings under the Session Node.

Single Sign-on Configuration

NetScaler Gateway supports using single sign-on to ensure that end users do not have to re-authenticate multiple times when accessing applications and data.

Using single sign-on can greatly simplify the end-user logon experience as well as reduce end-user frustration. Often when end users work within a LAN environment they do not have to re-authenticate to access applications and data, but extending this behavior to remote access as well this can allow end users to work in a familiar manner.

Single sign-on natively works with web applications, file shares, and Web Interface.

Single Sign-on with Windows Configuration

When end users use the NetScaler Gateway plug-in to connect securely they are able to re-establish VPN connectivity using the plug-in rather than having to re-authenticate using a web browser-based session. This can be further simplified by instructing the NetScaler Gateway plug-in to start automatically once the end user has logged on to Windows, using single sign-on for authentication. Single sign-on will pass the user's Windows logon credentials to the NetScaler Gateway for authentication, therefore reducing the number of times that an end user must authenticate.



If the user device is not within the company domain then single sign-on to the NetScaler Gateway plug-in will fail. If single sign-on to the plug-in fails, the user will be prompted to log on to the plug-in. When a device is not on the domain, the end user will log on using cached credentials. However if the user's password is changed within the domain whilst the laptop is offline, then single sign-on will fail as the cached password is not identical to the new domain password.

Using single sign-on can greatly simplify remote working for end users. Once they have logged on to their laptop a VPN connection is established and the end user can resume the same methods of working used in the office, including local applications and shared data access.

Enabling Single Sign-On

Single sign-on can be configured either globally or within a session policy.

Global configuration can be completed using the following method:

1. In the configuration utility, in the navigation pane, expand **NetScaler Gateway** and then click **Global Settings**.
2. In the details pane, under Settings, click **Change global settings**.
3. On the Client Experience tab, click **Single Sign-on with Windows** drop down menu and select **ON**, then click **OK**.

Configuration using a Session Policy can be completed using the following steps:

1. In the Configuration utility, in the navigation pane, expand **NetScaler Gateway > Policies > Session**.
2. Click **Add** then type a name for the Policy.
3. Click the + sign next to the **Profile** field. The **Create NetScaler Gateway Session Profile** window opens.
4. Enter a name for the Profile.
5. Click the Client Experience tab, click **Single Sign-on with Windows** and then click **Create**.
6. Select **ns_true** in the drop-down list in Saved Policy Expressions and then click **Create**.

Single Sign-on to Web Applications Configuration

Single sign-on for web applications can be configured either globally or by using session policies.

To configure single sign-on to web applications globally:

1. In the configuration utility, in the navigation pane, expand **NetScaler Gateway** and then click **Global Settings**.
2. In the details pane, under **Settings**, click Change **global settings**.
3. On the Client Experience tab, click **Single Sign-on to Web Applications** and then click **OK**.

To configure single sign-on to web applications by using a session policy:

1. In the Configuration utility, in the navigation pane, expand **NetScaler Gateway > Policies > Session**.
2. Click the Session Profile tab, highlight the Profile then click **Edit**.
3. On the Client Experience tab, next to Single Sign-On to Web Applications, click **Global Override**, click **Single Sign-On to Web Applications** and then click **OK**.

Single Sign-on to Web Applications Using LDAP Configuration

LDAP-based users often log on using a format of DOMAIN\UserID. However it is also possible to log on by using the user principal name (UPN). This takes a format of `username@domain.com`. By default when an end user logs on using their UPN, single sign-on will fail and the end user must authenticate twice.

If UPN logon is to be used then the LDAP authentication policy must be amended to facilitate single sign-on.

Modifying LDAP Authentication policy

1. In the configuration utility, in the navigation pane, click **NetScaler Gateway**.
2. In the details pane, under Policy Manager, click **NetScaler Gateway policy Manager**.
3. In the Policy Manager, click the + sign next to **NetScaler Gateway Global**.
4. Click the drop down menu under **Bind Point** and select **Authentication Policies (Primary)**.
5. Click the drop down menu under **Connection Type** and select **LDAP** then click **Continue**.
6. Highlight the LDAP Policy then click **Edit**.
7. In the **Configure Authentication LDAP Policy** dialog box, next to Server, click **Edit** icon.
8. Under **Connection Settings**, in Base DN (location of users), type `DC=domainname, DC=com`.
9. In **Administrator Bind DN**, type `LDAPuser@domainname.com` where `domainname.com` is the name of your domain.

10. Check the box for **BindDN Password**.
11. In Administrator Password and Confirm Administrator Password, type the Password1.
12. In **Server Logon Name Attribute**, select UserPrincipalName.
13. In **Group Attribute**, Select memberOf.
14. In **Sub Attribute Name**, select CN.
15. In **SSO Name Attribute**, type the format by which users log on and then click **OK** twice.
16. Click **Select, Bind, Done** then **Done**.

Domain Single Sign-on Configuration

Domain single sign-on should be configured for when users are connecting to servers running Citrix XenApp by using SmartAccess. To configure single sign-on to a domain, complete the following steps:

1. In the Configuration utility, in the navigation pane, expand **NetScaler Gateway > Policies > Session**.
2. Click the Session Profile tab, highlight the Profile then click **Edit**.
3. In the Configure Session Policy dialog box, next to Request Profile, click **Modify**.
4. In the **Configure NetScaler Gateway Session Profile** dialog box, on the **Published Applications** tab, in Single Sign-on Domain, click **Override Global**, type the domain name and then click **OK**.

Client Interception

Client interception is used to define which IP address or networks should be intercepted by the NetScaler Gateway plug-in and then sent through the NetScaler Gateway.

If split tunneling is enabled, intranet applications must be defined in order for client interception to occur.

By default, when a system IP address, mapped IP or subnet IP address is configured on the appliance, subnet routes are created based on these IP addresses. Intranet applications are created automatically based on these routes and can be bound to a virtual server.



If the following is set then intranet applications do not need to be defined, as all traffic will be sent through the NetScaler Gateway:

- Interception mode is set to transparent.
- Users are connecting to NetScaler Gateway with the NetScaler Gateway plug-in for Windows.
- Split tunneling is disabled.

Configuring Client Interception

To create an Intranet Application for one IP address, complete the following steps:

1. In the configuration utility, in the navigation pane, click **NetScaler Gateway > Resources > Intranet Applications**.
2. Click **Add**.
3. Type a Name for the profile.
4. Select the **Transparent** radio button.
5. In **Protocol**, select the protocol that applies to the network resource.
6. Under **Destination Type**, select **IP Address and Netmask**.
7. In **IP Address**, type the IP address and in **Netmask**, type the subnet mask then click **Create**.

To Configure an IP Address Range

To configure an IP address range, complete the following steps:

1. In the configuration utility, in the navigation pane, click **NetScaler Gateway > Resources > Intranet Applications**.
2. Click **Add**.
3. Type a Name for the profile.
4. Select the **Transparent** radio button.
5. In **Protocol**, select the protocol that applies to the network resource.
6. Under **Destination Type**, select **IP Address Range**.
7. In **Start IP**, type the starting IP address and in **End IP**, type the ending IP address, click **Create**.

Address Pools

By default when an end user establishes a connection using the NetScaler Gateway plug-in, the traffic to any back-end servers will appear to be generated by the NetScaler Gateway appliance. In some circumstances the end-user device must have a unique IP address, for example:

- Voice over IP
- Active FTP
- Samba file sharing (mapped drives)
- Instant messaging
- SSH
- VNC
- Remote Desktop to client desktops

NetScaler Gateway can assign a static IPv4 or IPv6 address to a user connection by configuring address pools. Within an address pool a group of static IP addresses is defined for end-user connection; these are then allocated based on the binding of the policy.



Address pools can also be referred to as intranet IP addresses.

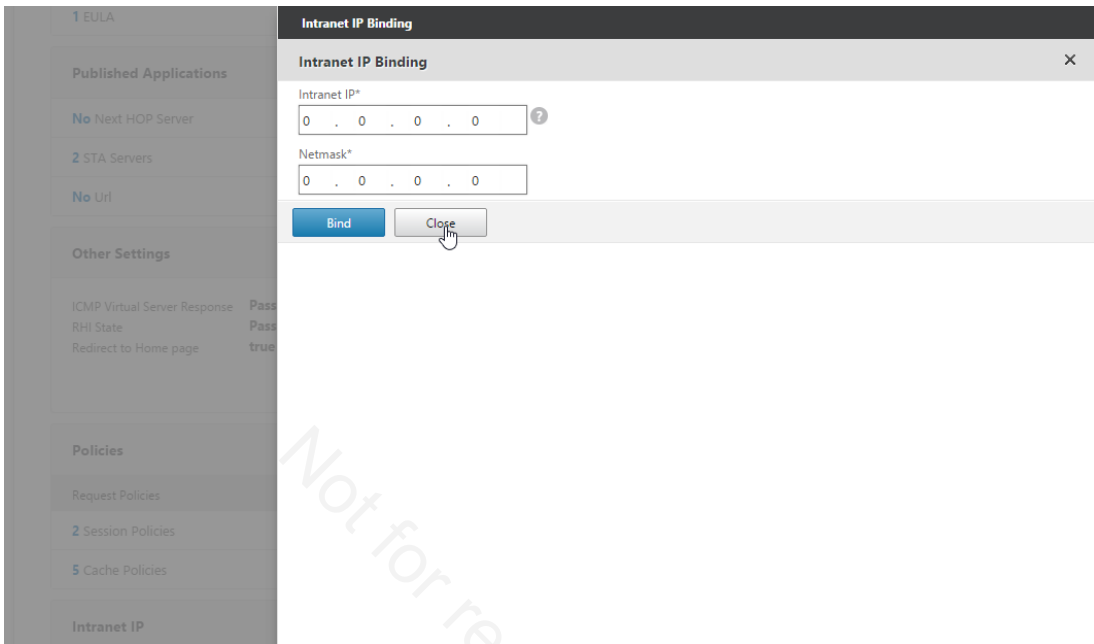
When a session is established the following hierarchy is used to determine which internal IP address should be assigned to the connection:

1. User's direct binding
2. Group assigned address pool
3. Virtual server assigned address pool
4. Global range of addresses

Configuring Intranet IP Addresses

To configure intranet IP addresses for a user, group or virtual server:

1. In the configuration utility, in the navigation pane, expand **NetScaler Gateway** and then click **Virtual Servers**.
2. Highlight the virtual server then click **Edit**.
3. In the details pane, click **Intranet IP Addresses** on the right pane.
4. On the left, click **No Intranet IP** window, click **Insert** and enter the IP address and subnet mask and then click **Bind**.
5. Repeat Step 4 for each IP address you want to add to the pool and then click **Done**.



To configure address pools globally:

1. In the configuration utility, in the navigation pane, expand **NetScaler Gateway** and then click **Global Settings**.
2. In the details pane, under **Intranet IPs**, click **To assign a unique, static IP Address or pool of IP Addresses for use by all client NetScaler Gateway sessions, configure Intranet IPs for IPv4 addresses**.
3. On the **Configure Intranet IP** window, click **Insert** and enter the IP address and subnet mask and then click **OK**.
4. Repeat Step 3 for each IP address you want to add to the pool and then click **Close**.

Defining Options for an Address Pool

It might be desirable to only use intranet IP addresses for certain connections. For example, when third-party contractors connect you can be required to audit the systems that they are accessing. Address pool options can be configured using session policies.

Configuring an address pool within a session policy can be completed using one of the following methods:

- **Nospillover** - When you configure address pools and the mapped IP address is not used, the Transfer Login page appears for end users who have used all available intranet IP addresses.
- **Spillover** - When you configure address pools and the mapped IP address is used as an intranet IP address, the mapped IP address is used when an intranet IP address cannot be assigned.

- **Off** - Address pools are not configured.

The Transfer Login page is presented if an intranet IP address is not available for the end user's connection. The page allows the end user to replace their existing NetScaler Gateway session with a new session. For example, if the user does not log off cleanly and then tries to establish a connection from another device, the end user can transfer the session to the new device.

Reclaiming of IP Addresses

When an IP address is assigned to an end user, the address is reserved for the end user's next logon until the address pool range is exhausted. When the addresses are exhausted, NetScaler Gateway reclaims the IP address from the end user who is logged off from NetScaler Gateway the longest.

If an address cannot be reclaimed and all addresses are actively in use, NetScaler Gateway does not allow the end user to log on. You can prevent this situation by allowing NetScaler Gateway to use the mapped IP address as an intranet IP address when all other IP addresses are unavailable.

To Configure Name Service Resolution

It is important to provide name resolution services for user devices when using a VPN connection. Many applications rely on suitable name resolution being in place in order to function. For example, Microsoft Outlook will look for a specific host name in order to download email. If this is not available then the end user might be unable to sync their mailbox.

NetScaler Gateway allows you to configure a DNS or WINS server for user devices using a session profile. DNS server settings can be specified globally or within a session policy.



Consider creating a load-balanced DNS virtual server that can be used for client connections. This will simplify configuration and also ensure that DNS continues to function for remote devices even if a server becomes unavailable.

Before specifying the settings, on the Network Configuration tab, do one of the following:

- To configure a DNS server, click the drop down list under **DNS Virtual Server**, select the server and then click **OK**.
- To configure a WINS server, enter the WINS Server IP address and then click **OK**.

To specify the setting with a Session Policy complete the following tasks:

1. Navigate to **NetScaler Gateway > Policies > Session**, then click the **Session Profiles** tab on the right pane.
2. Highlight the Session Profile then click **Edit**.
3. To configure a DNS server, next to DNS Virtual Server, click **Override Global**, select the server and then click **OK**.
4. To configure a WINS server, next to WINS Server IP, click **Override Global**, type the IP address and then click **OK**.

Access Interface Configuration

Once an end user has successfully logged on to NetScaler Gateway, a home page is displayed. This is known as the Access Interface. The Access Interface provides links to websites and file shares. You can customize the Access Interface by adding links to websites and file shares. Further customization is available by enabling end users to create their own bookmarks.



Often when end users attempt to open a file share from the Access Interface, a message can appear showing "Failed to make TCP connection to the server." This issue is caused by a firewall blocking traffic from the NetScaler Gateway system IP address to the file server IP address on TCP ports 445 and 139.

To Replace Access Interface with a Custom Home Page

Through the use of session policies, you can configure a custom home page to replace the default Access Interface. When a custom home page is configured the Access Interface is no longer displayed.

Custom home page settings can then be applied to users, groups, virtual servers or globally.

The screenshot shows the NetScaler VPX (1) Configuration page. The top navigation bar includes 'Dashboard', 'Configuration', 'Reporting', 'Documentation', and 'Downloads'. The main content area is titled 'Create NetScaler Gateway Session Profile'. It features a 'Name*' input field with a help icon. Below the input field, a note states: 'Unchecked Override Global check box indicates that the value is inherited from Global NetScaler Gateway Parameters.' There are five tabs: 'Network Configuration', 'Client Experience', 'Security', 'Published Applications', and 'Remote Desktop'. The 'Network Configuration' tab is active, showing 'Override Global' settings for 'DNS Virtual Server' (set to 'lb_vsrv_dns'), 'WINS Server IP', and 'Kill Connections*' (set to 'OFF'). There is an 'Advanced Settings' checkbox and 'Create' and 'Close' buttons at the bottom.

To configure a custom home page:

1. In the configuration utility, in the navigate to **NetScaler Gateway > Policies > Session**.
2. In the details pane, under Session Policies, click **Add**.
3. In **Name**, type a name for the Policy.
4. Next to **Profile**, click the + sign.
5. In **Name**, type a name for the Profile.
6. On the **Client Experience** tab, next to **Home Page**, click **Override Global**, click **Display Home Page** and then type the web address of the home page then click **Create**.
7. In the **Create NetScaler Gateway Session Policy** dialog box, click the **Saved Policy Expressions** drop down menu and select **ns_true** then click **Create**.

Clientless Access

Citrix NetScaler Gateway provides secure access to Citrix technologies and many other products, including web applications that require authentication credentials. NetScaler Gateway can provide access to certain web applications without the use of any plug-ins.

Through the use of session profiles, clientless access can be enabled or disabled in a granular fashion.

Clientless access can be used to mitigate threats to network security by securing internal systems for external use. For example, Microsoft SharePoint might not be directly addressable from the Internet, but could be secured by using NetScaler Gateway. End users once logged on to the NetScaler Gateway can be provided with access to SharePoint, ensuring that SharePoint does not have to be published externally and also ensuring that the user device does not require plug-ins.

A company can allow end users to decide which remote access method they require. NetScaler Gateway can fulfill this scenario by allowing them to choose the following options from the client choices screen:

- Clientless Access: When an employee wants to use Outlook Web Access or SharePoint only
- Network Access: When an employee requires full VPN access, for example to sync offline files or use Outlook
- Web Interface or ICA Proxy: Published applications and desktops



ICA proxy takes precedence over clientless access and SSL VPN in the absence of client choices being enabled. If ICA proxy is ON, this is how the user will connect regardless of the clientless access setting.

If ICA proxy is OFF and clientless access is ON, clientless access will be established.

If clientless access is set to DENY or ALLOW, a full VPN session will be established.

Clientless Access Enablement

Clientless access can be configured either at a global level or enabled and disabled by session policies. A session policy can then be bound to any of the following:

- User
- Group
- Virtual Server

Create NetScaler Gateway Session Profile

Name*

Unchecked Override Global check box indicates that the value is inherited from Global NetScaler Gateway Parameters.

Network Configuration | **Client Experience** | Security | Published Applications | Remote Desktop

Accounting Policy

Override Global

Display Home Page

Home Page

URL for Web-Based Email

Split Tunnel*

Session Time-out (mins)

Client Idle Time-out (mins)

Clientless Access*

Clientless Access URL Encoding*

When configuring clientless access the following options are available:

- **ON:** Enables clientless access. If you disable client choices and you do not configure or disable the Web Interface, end users log on by using clientless access.
- **ALLOW:** Clientless access is not enabled by default. If you disable client choices, and you do not configure or disable the Web Interface, end users log on with the NetScaler Gateway plug-in. If endpoint analysis fails when users log on, end users receive the choices page where clientless access is available.
- **OFF:** Clientless access is turned off. When you select this setting, end users cannot log on by using clientless access, and the icon for clientless access does not appear on the choices page.



If only a small number of end users will require clientless access, then set clientless access globally to OFF and enable access through a session policy that is bound to a user group. This keeps the configuration simple for the majority of the end users, while still allowing flexibility.



If clientless access was specified during the NetScaler Gateway Setup Wizard then clientless access set globally.

Web Address Encoding

When clientless access is used, you have the option to encode the internal URL that is being accessed. This is often used to add another layer of security, as it helps to protect internal information and also can be used to prevent people trying to access data by using bookmarks.

Three options for URL encoding are:

1. Obscure - URL is encoded to obscure the domain and protocol part of the resource
2. Clear - No encoding
3. Encrypt - Domain and protocol are encrypted using a session key. Because the key is different for each session the client cannot use bookmarks to access the page again.



URL encoding can either be set globally or by using session policies.

Many companies work with partners and contractors, often providing a platform for collaboration between the two organizations. These types of engagements are often seen as an IT security risk due to the lack of control over the user device that could be accessing the data. Encrypting the URL ensures that the data is obscured in the address bar, and that the page cannot be accessed using a web browser bookmark.



If an end user experiences issues with a web application, and the browser or session closes unexpectedly, the URL that the user was visiting can be requested; as such URL encoding might not be suitable in this scenario.

Clientless Access Policies

In order for end users to be able to access web applications using NetScaler Gateway's clientless access feature, a policy must first be created. Clientless access policies consist of a rule and a profile, and once a connection meets the rule criteria then the profile is invoked.

Pre-configured policies for SharePoint 2013, Outlook Web Access 2013, and Outlook Web App 2013 are all included with NetScaler Gateway.

Custom clientless access policies can also be created for additional web applications.



Pre-configured policies are bound globally and cannot be modified.

Clientless Access Policy Creation

Administrators might want to create additional policies to facilitate clientless access to web applications. Newly created policies can be bound to the following:

- NetScaler Gateway virtual server
- Users
- Groups

New policies can use existing settings from the default policies that are included with NetScaler Gateway or alternatively can be created as an empty policy and then expanded as required.

To use an existing policy as a template, select the policy from within the NetScaler Gateway, Policies, Clientless Access, highlight a built in policy and then click **Add**. All settings will be copied to the new profile and only a new policy name must be supplied.

Ensure that policy priorities are configured correctly, as when the first policy hit occurs, the processing will terminate and the policy's profile will be invoked for the end user's session.

Domain Access Configuration for End Users

Clientless access can offer an additional layer of security by specifically permitting or denying domains, websites and network resources.

The most common resource to be included or excluded for clientless access is domains. By including or excluding domains, you can ensure that sessions will access data only within the permitted locations.



For third-parties or partners, it is often preferred to place a strict limit on which network resources can be accessed. For example, clientless access can be deployed to allow a partner to use a custom application; however it is imperative that they cannot access any other data. In this scenario all access should be set to DENY and an inclusion list created that specifies only the web application as an allowed resource.



If you use DENY as the default action, be aware that some web applications might try to direct end users to additional network locations or resources. These resources should also be included within the inclusion list otherwise the web application might not function correctly.

SharePoint Site

NetScaler Gateway provides clientless access capabilities for Microsoft SharePoint 2013, Outlook Web Access 2013, and Outlook Web App 2013. Pre-configured clientless access policies are included and bound globally by default.

When using clientless access with SharePoint, NetScaler Gateway must have a list of host names for each SharePoint server within the network; without this list, the rewrite process will fail.

SharePoint can also be configured as the end user's homepage when clientless access is in use. For more information about configuring user homepages with clientless access, see Citrix product documentation at <http://docs.citrix.com>.

Name*

Unchecked Override Global check box indicates that the value is inherited from Global NetScaler Gateway Parameters.

Network Configuration	Client Experience	Security	Published Applications	Remote Desktop
-----------------------	--------------------------	----------	------------------------	----------------

Accounting Policy

Override Global

Display Home Page

Home Page

URL for Web-Based Email

Split Tunnel*

Session Time-out (mins)

Client Idle Time-out (mins)

Clientless Access*

Clientless Access URL Encoding*



If persistent cookies are not configured for clientless access to SharePoint some functions might fail to work correctly. For further information on setting a persistent cookie for clientless access, see Citrix product documentation at <http://docs.citrix.com>.

Clientless Access Settings Using Web Interface

Citrix Web Interface uses cookies to maintain settings between sessions. When using clientless access the cookies from previous sessions are not forwarded; as such, settings that have been modified in the previous session might not be present.

NetScaler Gateway allows specific client cookies to be captured and used between sessions. When multiple cookies are required they can be added into a pattern set; this is essentially a grouping of client cookies.

Citrix Web Interface 5.x uses three client cookies which should be added to a pattern set:

1. WIUser
2. WINGDevice
3. WINGSession

Once the cookies are placed within the pattern set, end users will be able to maintain Web Interface settings between clientless access sessions.



By default a pattern set already exists, named "ns_cvpn_default_client_cookies," the additional Web Interface cookies can be added into the existing pattern set.

Client Choices Page Configuration

NetScaler Gateway provides several different options for client connectivity; these can then be configured based on end user access requirements.

Through the use of session policies and profiles the client choices can be configured, ensuring that end users only see options for the connection methods that they are permitted to use. Session policies are used to override defined global options. The client choices screen is presented after successful authentication and can display the following options:

1. NetScaler Gateway plug-in (network access) - the NetScaler Gateway plug-in is started and a VPN connection is established.
2. Web Interface - published applications and desktops are presented.
3. Clientless access - the access interface is displayed, displaying file shares, websites, and embedded web applications such as Outlook Web Access.
4. NetScaler Gateway plug-in for Java - the java plug-in starts and allows network access.

Only the options that the end user has been permitted to use will be displayed within the client choices page. This configuration is held within a session profile and consists of the following settings:

1. Plug-in type
2. Clientless access
3. Client choices
4. Web Interface address

The configuration of these settings will determine which choices are presented to the end user.

Client choices are enabled by using the "advanced" button that is available within the client experience tab of a session profile.



Use Active Directory groups when binding session policies that contain specific settings that only certain end users require. This greatly simplifies the configuration and also allows you to designate the settings for additional end users by simply adding them to a new group.



Selecting the option to use the Java plug-in will require an intranet application to be defined and the interception mode set to proxy.

IT administrators might want to use the NetScaler Gateway plug-in to establish a VPN connection in certain circumstances but use Web Interface to launch applications for daily operations, while all other staff should only be allowed to use Web Interface. In this instance a session policy should be configured to permit the use of the client choices page and then bound to the Administrators group, ensuring that the priority number is lower than those of other session policies. Web Interface-only access should be either set globally or by using a session policy that is bound to the virtual server; this will force all other end users to use Web Interface, while allowing the administrators to choose their connection method.

Access Scenario Fallback Configuration

NetScaler Gateway's SmartAccess functionality allows for granular policy control over connecting devices. This includes scanning devices to ensure that they fulfill certain criteria and then offering connectivity options that depend on the result of the scan.

In scenarios in which devices fail a post-authentication endpoint scan, the device can be permitted to access Web Interface or use clientless access only, rather than allowing use of the NetScaler Gateway plug-in. This is achieved by creating a quarantine group.



Ensure that the name of the group being created does not match the name of an existing Active Directory group; otherwise the quarantine settings will apply to members of the Active Directory group when they log on.

Use of the NetScaler Gateway plug-in is often restricted to IT staff, whilst all other end users are allowed to use ICA Proxy mode. It may be desirable to ensure that the IT staff is only permitted to use the NetScaler Gateway plug-in when they are connecting from a trusted device: for example, checking the registry to confirm that the device is a member of the corporate domain. Using the fallback scenario it is possible to only display the NetScaler Gateway plug-in connection method if the device is a member of the domain; otherwise they will be restricted to Web Interface-only access.

Access Scenario Fallback Policy Creation

Access scenario fallback uses session policies and post-authentication scans to facilitate a flexible and secure approach to providing access to corporate resources.

For more information about the access scenario fallback procedure, see Citrix product documentation at <http://docs.citrix.com>.



The endpoint analysis plug-in is required in order to allow access scenario fallback to function. The plug-in will be offered for download from the Gateway when the end user attempts to connect. Once installed, the plug-in will prompt the end user to allow endpoint analysis to run. If the end user selects the option to "skip scan" then the end user is denied access.

NetScaler Gateway Advanced Concepts

For additional information about NetScaler Gateway advanced features, see Citrix product documentation at <http://docs.citrix.com>.

Advanced concepts include:

- **Deploying NetScaler Gateway in a double-hop perimeter network.** You can deploy two NetScaler Gateway appliances in a double-hop perimeter network in two stages to provide an extra layer of security for the internal network.
- **Configuring DNS virtual servers.** You can configure a DNS server as a virtual server and then bind the server globally or to another virtual server.
- **Resolving DNS name servers located in the secure network.** If your DNS server is located in the secure network behind a firewall and the firewall is blocking ICMP traffic, you can use a non-directly addressable DNS virtual server on NetScaler Gateway that resolves to a known fully qualified domain name (FQDN).
- **Using operators and operands.** You can use operators and operands in policy expressions.
- **Configuring server-initiated connections.** When an IP address is assigned to an end user's session, it is possible to connect to the user device from the internal network by using Remote Desktop or a virtual network client (VNC).
- **Enabling NetScaler Gateway plug-in logging.** You can configure the NetScaler Gateway plug-in to log all errors to a text file.

Discussion Question

How are you and your organization currently using the NetScaler Gateway? If you have yet to implement it, how do you plan to use NetScaler Gateway?

Not for resale or distribution

Not for resale or distribution

Module 10

Integrating NetScaler 11 with XenApp and XenDesktop

10

Not for resale or distribution

Not for resale or distribution

Integrating NetScaler with XenApp and XenDesktop Manual

Overview

NetScaler Gateway can be configured to communicate with XenApp and XenDesktop components such as StoreFront and the Web Interface, as well as the Citrix XenMobile AppController. To allow end-user connections to a farm through NetScaler Gateway, you configure settings in either StoreFront or the Web Interface.

After completing this module, you will be able to:

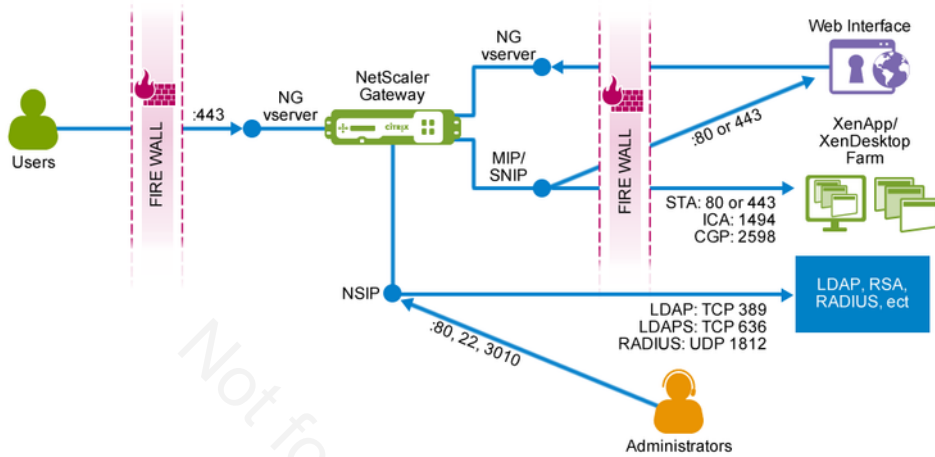
- Configure the NetScaler Gateway feature with Citrix StoreFront Services.
- Configure the NetScaler Gateway feature with Citrix XenApp and Citrix XenDesktop.
- Configure end-user device mapping to provide end users with access to XenApp applications.
- Configure SmartAccess with XenApp and XenDesktop to control user access to published applications and virtual desktops.
- Describe XenMobile and how this solution integrates with Citrix NetScaler.

NetScaler Gateway Prerequisites

Before you begin configuring settings for integration on the NetScaler Gateway, ensure that the following prerequisites are met:

- NetScaler Gateway is installed in your environment and has access to the network or networks. NetScaler Gateway is deployed in the perimeter network or internal network behind a firewall. You can also configure NetScaler Gateway in a double-hop perimeter network and for connections to a farm.
- NetScaler Gateway is configured with a default gateway or with static routes to the internal network so that end users can access resources in the network. NetScaler Gateway is configured to use static routes by default.
- The external servers used for authentication and authorization are configured and running.
- The network has a DNS or Windows Internet Naming Service (WINS) server for name resolution to provide correct NetScaler Gateway user functionality.
- The correct licenses are installed.
- NetScaler Gateway has a certificate that is signed by a trusted CA.

Firewall Rules



In order to properly configure the NetScaler Gateway, you will need to ensure that your firewall allows the proper traffic. The NetScaler Gateway can use the following ports:

Route	Port	Details
Public network to perimeter network	80, 443	Port 80 provides a redirect to port 443.
Private network to perimeter network	80, 443, 1494, 2598	Ports 80 and 443 are used for management and administration of the NSIP through a browser. Port 443 is also used for the authentication callback URL to the NetScaler Gateway VIP from StoreFront or Web Interface. Ports 1494 and 2598 can be used for ICA or HDX traffic.
Perimeter network to private network	443, 80	Port 443 allows access to the StoreFront and port 80 allows access to XenApp or XenDesktop (if port 80 is used for secure ticketing authority traffic).

Route	Port	Details
Perimeter network to private	389, 636	Ports 389 or 636 can be used for LDAP or Secure LDAP (LDAPS) authentication traffic to the LDAP servers.

StoreFront Services Deployment

Citrix StoreFront enables you to create enterprise application stores that aggregate resources from XenDesktop, XenApp, XenMobile App Controller and VDI-in-a-Box in one place. The stores you create provide your users with self-service access to their Windows desktops and applications, mobile applications, external software-as-a-service (SaaS) applications and internal web applications through a single portal from all their devices. StoreFront provides a single place to manage the provisioning of corporate desktops and applications to your users. Consolidating the delivery of resources through StoreFront eliminates the need to manage multiple delivery mechanisms for different applications or provide support for manual installations and updates.



App Controller is the unified policy controller of Citrix XenMobile App Edition that lets you securely deliver enterprise web and Software-as-a-Service (SaaS) applications, Android and iOS apps, access to public stores and much more. The ShareFile settings and application connector integrated into Citrix Receiver is also managed by XenMobile App Edition, to ensure that users have access to shared data when they need it, across devices. The App Controller virtual machine (VM) is a virtual appliance that runs on Citrix XenServer and is managed with Citrix XenCenter. You can also install App Controller on VMware ESXi and Microsoft Hyper-V.

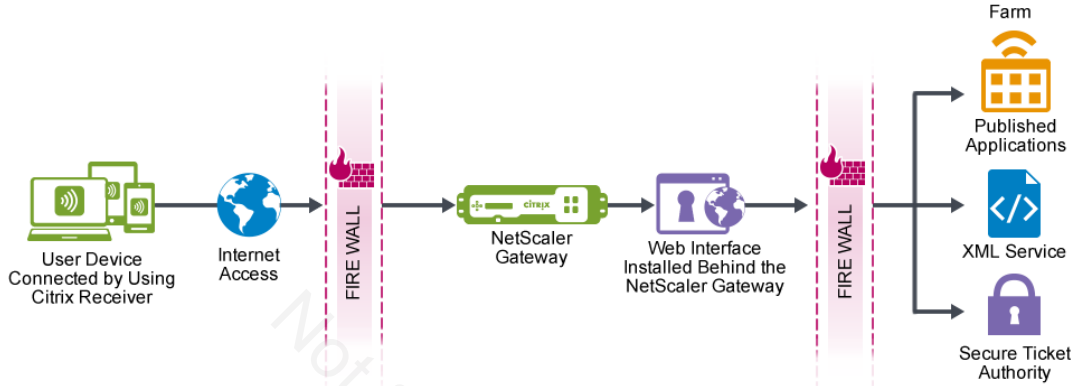
Web Interface Site Configuration with NetScaler Gateway

The Web Interface provides end users with access to XenApp applications and content as well as XenDesktop virtual desktops. End users access their published applications and desktops through a standard Web browser or through Citrix Receiver.

Use the Web Interface Management Console to create Web Interface sites for version 5.4.2 (recommended). You can install the consoles on Windows-based platforms only. To configure the Web Interface to work with NetScaler Gateway, you need to:

- Create the Web Interface site for the version you are using.
- Configure settings in the Web Interface.
- Configure session profile settings to use Web Interface on NetScaler Gateway.

Web Interface or XenApp Integration with NetScaler Gateway



The following process defines how an HDX session is established through NetScaler Gateway for integration with Web Interface or XenApp:

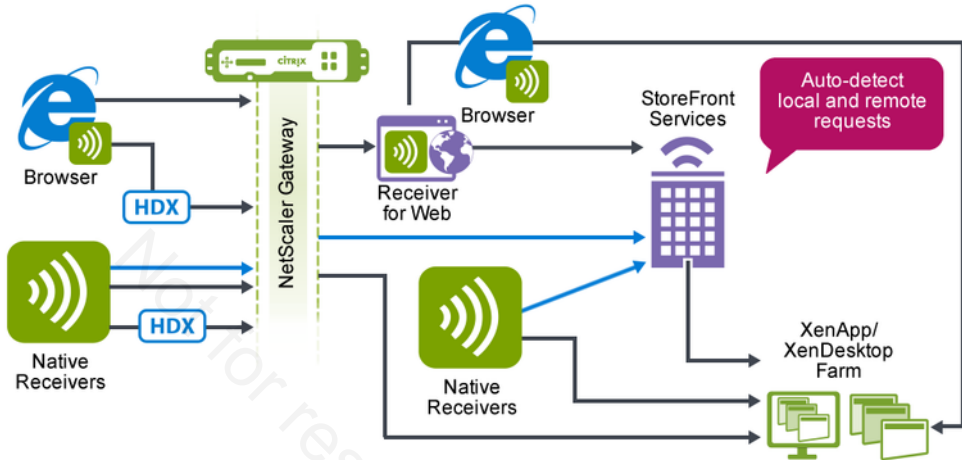
1. The end user connects to the NetScaler Gateway (for example, <https://remote.training.lab>).
2. NetScaler Gateway terminates SSL, authenticates the user, and validates the end-user device.
3. NetScaler Gateway passes user credentials and policy conditions to Web Interface.
4. Web Interface validates any SmartAccess conditions through a secure channel and then enumerates the applications.
5. The end user clicks an application icon.
6. Web Interface requests a ticket from the Secure Ticket Authority.
7. Web Interface sends the ticket to the end user in an ICA file.
8. The HDX client spawns and sends ICA traffic decrypted with SSL to the NetScaler Gateway.
9. NetScaler Gateway validates the ticket.
10. The HDX session is established.

Manual Configuration of Certificate Trust for Web Interface in NetScaler

You can import a root Certificate Authority into the NetScaler's trusted certificate store using the command-line interface. Setting up a certificate trust is necessary when deploying Web Interface for NetScaler in Gateway Direct Mode with the authentication point set to NetScaler Gateway. For more information about configuring a certificate trust for Web Interface, see Citrix article CTX127615 at <http://support.citrix.com>.

NetScaler Gateway Integration with Web Interface and StoreFront Services

NetScaler Gateway can be configured to integrate with either Web Interface or Citrix StoreFront Services as part of your implementation.



When you configure remote access and you want to provide access to resources from XenApp or XenDesktop, you can configure the Secure Ticket Authority while configuring the NetScaler Gateway.

When you configure Web Interface or StoreFront, you need to configure the FQDN for the callback URL that verifies that the request came from NetScaler Gateway. You use the same FQDN to which end users connect.



You should remember the following items when integrating with NetScaler Gateway:

- Client access to NetScaler Gateway is through the FQDN.
- Web Interface or StoreFront must call back to NetScaler Gateway using the FQDN.
- Web Interface or StoreFront must trust the SSL certificate at the machine level.
- The FQDN on the certificate must match the FQDN used by end users.
- A change in SSL certificate authorities might require a restart of the IIS web services.
- If you use domain names on NetScaler Gateway to connect to Web Interface from StoreFront, ensure that NetScaler Gateway can properly resolve the DNS.
- XenApp or XenDesktop must trust XML requests.
- Load-balance backend resources to provide resilience.

For more information about NetScaler Gateway and Receiver StoreFront integration, see Citrix product documentation at <http://docs.citrix.com>.

Citrix Receiver for Web

Receiver for Web sites enable users to access stores through a webpage. Some advanced settings can only be changed by editing the site configuration files.

Receiver for Web is a Web user interface for self-service access to IT resources. Receiver for Web contains a consistent user interface with native Receivers. Receiver for Web also contains separation between applications and desktops.

Not all of the features of Receiver for Windows are implemented with Receiver for Web, such as:

- Start Menu integration
- File Type Association
- Multi-store
- GoTo Product and streamed applications

For users who cannot install Citrix Receiver, you can enable Receiver for HTML5 on your Receiver for Web sites. Receiver for HTML5 enables users to access desktops and applications directly within HTML5-compatible web browsers without needing to install Citrix Receiver. Both internal network connections and connections through NetScaler Gateway are supported. However, for connections from the internal network, Receiver for HTML5 only enables access to resources provided by specific products. Additionally, specific versions of NetScaler Gateway are required to enable connections from outside the corporate network.

For local users on the internal network, access through Receiver for HTML5 to resources provided by XenDesktop and XenApp is disabled by default. To enable local access to desktops and applications using Receiver for HTML5, you must enable the ICA WebSockets connections policy on your XenDesktop and XenApp servers. XenDesktop and XenApp use port 8008 for Receiver for HTML5 connections. Ensure your firewalls and other network devices permit access to this port.

Receiver for HTML5 can only be used with Internet Explorer over HTTP connections. To use Receiver for HTML5 with Mozilla Firefox over HTTPS connections, users must type `about:config` in the Firefox address bar and set the "network.websocket.allowInsecureFromHTTPS" preference to true.

Secure Ticket Authority Configuration on NetScaler Gateway

The Secure Ticket Authority (STA) is responsible for issuing session tickets in response to connection requests for published applications on XenApp and published desktops on XenDesktop. These session tickets form the basis of authentication and authorization for access to published resources.

You can use any of the following methods to configure the STA on NetScaler Gateway:

- Global settings in the configuration utility
- XenDesktop and XenApp wizard
- NetScaler Gateway Policy Manager

You can bind the STA globally or to virtual servers. You can also add multiple servers running the STA when you configure a virtual server. If you are securing communications between the NetScaler Gateway and the STA, make sure a server certificate is installed on the server running the STA. For detailed steps instructing how to bind the STA globally or to virtual servers, see Citrix product documentation at <http://docs.citrix.com>.



The STA settings should be identical on StoreFront (or Web Interface) and the NetScaler, otherwise errors can occur when launching connections.

To Configure Remote Access to StoreFront

The following steps to configure remote access to StoreFront assume that you have created one or more stores in StoreFront.

1. In the Citrix StoreFront management console, click **Citrix StoreFront > Stores**.
2. In the center pane, select a store.
3. Under Actions, in the right pane, under store, click **Enable Remote Access**.
4. In the Enable Remote Access dialog box, select one of the following:
 - None



Select **None** to make the store unavailable to users on public networks. Only local users will be able to access the store.

- No VPN tunnel



Select **No VPN tunnel** to make only resources delivered through the store available through NetScaler Gateway. Users do not need to use the NetScaler Gateway plug-in.

- Full VPN tunnel



Select **Full VPN tunnel** to make the store and other resources delivered on the internal network available through an SSL VPN tunnel. Users require the NetScaler Gateway plug-in to establish the VPN tunnel.

5. In NetScaler Gateway appliances, click **Add**.
6. In the **General Settings** page, type the NetScaler Gateway name.
7. Enter the URL for NetScaler Gateway. This is the fully qualified domain name (FQDN) that is located in the secure server certificate installed on NetScaler Gateway and bound to the virtual server. Using the same FQDN for StoreFront and the NetScaler Gateway virtual server is not supported.
8. In Subnet IP address, enter the mapped IP address or subnet IP address that you configured in NetScaler Gateway.

9. In Logon type, select one of the following:
 - Domain
 - Security token
 - Domain and security token
 - SMS authentication
 - Smart card
10. Click **Next**.
11. Complete the NetScaler Gateway authentication service URL in the **Callback URL** box. StoreFront automatically appends the standard portion of the URL. Click **Next**.
12. On the Enable Silent Authentication page, list URLs for the authentication service running on the Access Controller servers. Add URLs for multiple servers to enable fault tolerance, listing the servers in order of priority to set the failover sequence. Click **Next**.
13. Configure the settings for the Secure Ticket Authority (STA).



You can configure more than one STA server for redundancy.

14. Click **Create** to add your NetScaler Gateway deployment to the list in the Enable Remote Access dialog box.

Optimal NetScaler Gateway Routing

If you want to use HDX Insight with Citrix StoreFront, you must route all traffic through the NetScaler appliance. Optimal routing allows you to facilitate this type of deployment. To configure optimal routing for your deployment, editing the store configuration files is necessary.



In multiple server deployments, use only one server at a time to make changes to the configuration of the server group. Ensure that the Citrix StoreFront management console is not running on any of the other servers in the deployment. Once complete, propagate your configuration changes to the server group so that the other servers in the deployment are updated.

If you have configured separate NetScaler Gateway appliances for your deployments, StoreFront enables you to define the optimal appliance for users to access each of the deployments providing resources for a store. For example, if you create a store that aggregates resources from two geographical locations, each with a NetScaler Gateway appliance, users connecting through an appliance in one location can start a desktop or application in the other location. However, by default, the connection to the resource is then routed through the appliance to which the user originally connected and must therefore traverse the corporate WAN.

To improve the user experience and reduce network traffic over the WAN, you can specify the optimal NetScaler Gateway appliance for each of your deployments. With this configuration, user connections to resources are automatically routed through the appliance local to the deployment

providing the resources, regardless of the location of the appliance through which the user accesses the store.

Optimal NetScaler Gateway routing can also be used in the special case where local users on the internal network are required to log on to NetScaler Gateway for endpoint analysis. With this configuration, users connect to the store through the NetScaler Gateway appliance, but there is no need to route the connection to the resource through the appliance as the user is on the internal network. In this case, you enable optimal routing, but do not specify an appliance for the deployment, so user connections to desktops and applications are routed directly and not through NetScaler Gateway. Note that you must also configure a specific internal virtual server IP address for the NetScaler Gateway appliance. Additionally, specify an inaccessible internal beacon point so that Citrix Receiver is always prompted to connect to NetScaler Gateway, regardless of the user's network location.

For more information, see the "StoreFront high availability and multi-site configuration" and "Optimal NetScaler Gateway routing example" topics located at <http://docs.citrix.com>.

To Configure Optimal Routing

1. Use a text editor to open the web.config file for the store, which is typically located in the C:\inetpub\wwwroot\Citrix\storename\ directory, where *storename* is the name specified for the store when it was created.
2. Locate the following element in the file:

```
<optimalGatewayForFarmsCollection />
```

3. Specify the optimal NetScaler Gateway routing for your deployments.



Use the following elements to define your configuration:

- "optimalGatewayForFarms": Specifies groups of deployments and defines the optimal NetScaler Gateway appliances for users to access resources provided by these deployments. Set the value of the "enabledOnDirectAccess" attribute to "true" to ensure that when local users on the internal network log on StoreFront directly, connections to their resources are still routed through the optimal appliance for the deployment. For example, the optimal appliance for the deployment could be a NetScaler Gateway appliance running HDX Insight.
- "farms": Specifies a set of XenDesktop, XenApp, App Controller and VDI-in-a-Box deployments that share a common optimal NetScaler Gateway appliance. Enter the names of deployments that you have already added to the store. The names of the deployments you specify must match exactly the names you entered when you added deployments to the store.
- "optimal gateway": Specifies details of the optimal NetScaler Gateway appliance for users to access resources provided by the listed deployments.
- "hostnames": Specifies the FQDN and port of the optimal NetScaler Gateway appliance.
- "staUrls": Specifies the URLs for XenDesktop, XenApp and VDI-in-a-Box servers running the STA.



For more information about editing the store configuration file for optimal routing, see the "To Configure optimal NetScaler Gateway routing for a store" topic at <http://docs.citrix.com>.

Beacons

The beacons "container" stores objects used to help automatically determine whether an end user is operating within the corporate network or externally. This distinction is needed to indicate whether the ICA file generated in response to a start request should include gateway information.

Native Receivers use the provided beacons. If Receiver can ping an internal beacon, a server with an address that is only accessible from within the organization, then it knows that it is operating within the corporate network.

If Receiver cannot reach the internal beacon, but can reach an external beacon, then it knows to signal StoreFront that it is operating externally.

For browser-based end users connecting with Receiver for Web, the internal or external distinction is established by a "remote" flag in the HTTP header of an end user's initial connection that indicates whether the traffic is coming in through NetScaler Gateway.

Enabling Access Method Fallback with Policies

You can use policies to determine the access method granted to an end-user device. For example, you might want to offer full SSL VPN access to a trusted corporate laptop but limit an unsecured laptop access to Web Interface or StoreFront. The following process describes an access method fallback scenario.

1. Create session profiles for:
 - SSL_VPN
 - Clientless
 - ICA_Proxy
2. Create corresponding rules for each session policy:
 - SSL_VPN - a policy to test for hidden registry entries or files, Windows version, and Symantec antivirus
 - Clientless - a policy to test for the Windows operating system and any antivirus
 - ICA_Proxy - a policy that is applied when EPA scans fail (using ns_true)
3. Bind each of the three policies to a chosen bind point. Use the following priorities:
 - Priority 10 - SSL_VPN
 - Priority 20 - Clientless
 - Priority 30 - ICA_PROXY

SSL Certificate Trust

When using public SSL certificates, the client and Web Interface or StoreFront will have the CA root certificate installed, but might not have the intermediate certificate.



You should also install the intermediate certificate on NetScaler Gateway and link the server certificate to it.

When using private SSL certificates, there is no trusted CA certificate. Both the client and Web Interface or StoreFront will receive certificate errors.



You should either import the private certificate or the private CA root certificate. The SSL certificate must be trusted at the machine level.

Session Policies

Session policies define a group of settings (a session profile) that should be applied to an end user's session. The profile for the sessions is invoked when the policy expression is met.

Session policies are applied after an end user has successfully authenticated to NetScaler Gateway. These policies can be combined with SmartAccess to permit the use of session policies as filters for XenApp or XenDesktop policies.

Using endpoint analysis expressions to control which session policy is invoked is a powerful method of ensuring compliance with corporate standards, as well as BYOD scenarios. While a company might use a BYOD model, it might be desirable that non-corporate devices are given a more restrictive connection. For example, SmartAccess policies can be used to block any local client drive mappings. By creating a sessions policy and configuring an expression for the required criteria, you can then use the policy as a filter within XenApp or XenDesktop to restrict use.



Administrators often set the primary NetScaler Gateway settings globally, then use session policies to override any specific settings. For example, consider a scenario where ICA proxy should be configured for all users, but the users of the NetworkTeam Active Directory group should also be offered the option to use the NetScaler Gateway plug-in. This scenario can be achieved by creating a session policy with the required settings and binding it to the NetworkTeam group.

Session Policy and Profile Creation for Citrix Receiver

To allow connections through NetScaler Gateway from the different versions of Receiver, you need to create session policies and profiles with specific rules to enable the connections to work. You can create separate session policies and profiles for several clients, including:

- Receiver for any supported operating system
- Receiver for Web

You can use the same policy expressions for each deployment, one to identify a Receiver of any type and one to identify browser-based access. The session policy expressions you configure depend on the version of Receiver you are using. Some versions of Receiver do not fully support the StoreFront services protocols that allow direct connections through NetScaler Gateway to stores in StoreFront. Some earlier Receiver versions do not support these protocols, such as:

- Receiver for Windows 3.0 and earlier versions
- Receiver for Mac 11.4 and earlier versions
- Receiver for Android 3.0 and earlier versions
- Receiver for iOS 5.5 and earlier version

For more information about creating session policies for Receiver and Receiver StoreFront, see Citrix product documentation at <http://docs.citrix.com>.

When you configure session profiles for use with a session policy, you need to configure parameters that are specific for the type of connection the profile supports. When you finish configuring the policy and profile, you then bind the session policy to the virtual server. The priority of the policy should be set to check for the most explicit condition first, so the expression to match any Receiver should have the lowest priority. For more information about configuring session profiles with a session policy, see Citrix product documentation at <http://docs.citrix.com>.

Session Policy Expressions

The following table shows the policy expression to configure based on the version of Receiver and the NetScaler Gateway plug-in you are using:

Version	Expression
Receiver version does not support StoreFront Services protocols	<code>REQ.HTTP.HEADER User-Agent CONTAINS CitrixReceiver && REQ.HTTP.HEADER X-Citrix-Gateway NOTEXISTS</code>
Receiver version supports StoreFront Services protocols	<code>REQ.HTTP.Header User-Agent CONTAINS CitrixReceiver && REQ.HTTP.HEADER X-Citrix-Gateway EXISTS</code>
NetScaler Gateway plug-in for Windows or Mac	<code>REQ.HTTP.HEADER User-Agent NOTCONTAINS CitrixReceiver && REQ.HTTP.HEADER Referer NOTEXISTS</code>
Receiver for Web	<code>REQ.HTTP.HEADER User-Agent NOTCONTAINS CitrixReceiver && REQ.HTTP.HEADER Referer EXISTS</code>

Session Policy Creation for the NetScaler Gateway Plug-In

You can configure NetScaler Gateway to provide end user access to published applications and virtual desktops with the NetScaler Gateway plug-in instead of with Receiver or a web browser. You can configure a session profile by using the instructions in *Configuring Access to Applications and Virtual Desktops* in the Web Interface. Then, in the session policy, you add an HTTP header rule for the NetScaler Gateway plug-in.

To Create the Session Policy Rule for the NetScaler Gateway Plug-In

1. In the Create NetScaler Gateway Session Policy dialog box, to the right of the **Expression** field, click the **Expression Editor**.
2. In the Add Expression dialog box, do the following:
 - a. In Select Expression Type, click **General**.
 - b. In Flow Type, select **REQ**.
 - c. In Protocol, select **HTTP**.

- d. In Qualifier, select **Header**.
 - e. In Operator, select **NOTEXISTS**.
 - f. In Header Name, type `Referer` and then click **Done**.
3. Click **Create**.

Session Policies Binding and Priority Setting

After you configure session policies for Citrix Receiver, NetScaler Gateway plug-in, or StoreFront Services, you can bind the policy to a user, group, virtual server or globally. Session policies are applied in the following order:

- Users
- Groups
- Virtual servers
- Globally

If you configure two or more session policies for Receiver for Windows and Receiver for Mac, Receiver for Web and the NetScaler Gateway plug-in, you bind the policies and then set the priority number for each policy.

Numerical priority takes precedence regardless of the level at which the policy is bound. If a policy that is bound globally has a priority number of one and another policy bound to a user has a priority number of two, the global policy takes precedence. A lower priority number gives the policy a higher precedence.

Session Profile Creation

Session profiles contain a group of settings that configure an end user's connection after successful authentication.

Session profiles are commonly used to set things like:

- DNS Server
- Plug-in type to be used
- Web interface site
- Enable or Disable option for split tunneling
- Proxy settings
- Session timeouts
- Quarantine groups
- Single sign-on

A session policy can only have one session profile; however, one session profile can be used in multiple session policies.

Session profiles allow you to specify key settings for a user's connection to NetScaler Gateway. By combining a profile with a session policy, you can then determine when these settings should be run.

Typically, an administrator defines access levels by creating two or more session profiles, then creating policies that contain expressions to determine when each profile should be used.

To Create a Session Profile Using the Configuration Utility

1. In the navigation pane, expand **NetScaler Gateway > Policies** and then click **Session**.
2. In the details pane, click the Profiles tab and then click **Add**.
3. Configure the settings for the profile, click **Create**.

After you create a profile, you can bind it to a session policy.

To Add a Profile to a Session Policy Using the Configuration Utility

1. In the configuration utility, in the navigation pane, expand **NetScaler Gateway > Policies** and then click **Session**.
2. On the Policies tab, do one of the following:
 - Click **Add** to create a new session policy.
 - Select a policy and then click **Edit**.
3. In Request Profile, select a profile from the list.
4. Finish configuring the Session Policy and then close the dialog box.

Custom Clientless Access Policies Configuration for StoreFront Services

You can configure a clientless access policy for Citrix Receiver and StoreFront Services.

If you create clientless access policies for Receiver and Receiver for Web, you must bind the Receiver policy to the virtual server before you bind the Receiver for Web policy. When you bind the Receiver for Web policy, choose the lower priority number to make sure that the policy takes precedence.

For detailed steps to configure a clientless access policy for Receiver, see Citrix product documentation at <http://docs.citrix.com>.

Remote Desktop Proxy

Remote Desktop Proxy is a new feature included with NetScaler 11 which provided Remote Desktop Gateway functionality via the Gateway virtual server. Remote Desktop Proxy is configured in the Session Profile and also in RDP profile.

1. Requirements for Remote Desktop Proxy Include:
 - A Gateway Universal license for each concurrent connection
 - An Enterprise or Platinum Platform License (can add on feature to standard license.)
 - ICA_Proxy set to off in the session profile.
 - Appropriate routes and allowed ports must be in place.
2. Accessing machines using Remote Desktop Proxy
 - Supports Single Sign On (SSO) so there is no need to authenticate to the RDP machine after logging in to Gateway.
 - You can create hyperlink bookmarks to the machines so users need only click on the link to launch the RDP session.
 - You can append /rdpproxy/(hostname or IP> to the Gateway URL to directly access a machine.

XenApp and XenDesktop Addition to a Single Site

If you are running XenApp and XenDesktop, you can add both applications to a single Web Interface or StoreFront site. The configuration allows you to use the same Secure Ticket Authority (STA) server from either XenApp or your XenDesktop Delivery Controller. To review detailed steps for adding XenApp or XenDesktop to a single site, see Citrix product documentation at <http://docs.citrix.com>.

XenMobile Platform Overview

With the XenMobile platform, you can offer end users connections to Windows, Web and SaaS applications and virtual desktops hosted in your network. NetScaler Gateway and XenMobile can provide internal and remote users with access to your network's applications and desktops. NetScaler Gateway authenticates users and then allows them to access their applications using Citrix Receiver.

XenMobile MDM Edition XenMobile MDM offers market-leading mobile device management capabilities that provision security to applications and data to corporate-and-employee-owned mobile devices. In this edition, ShareFile provides read access to files on network drives and SharePoint.

XenMobile App Edition

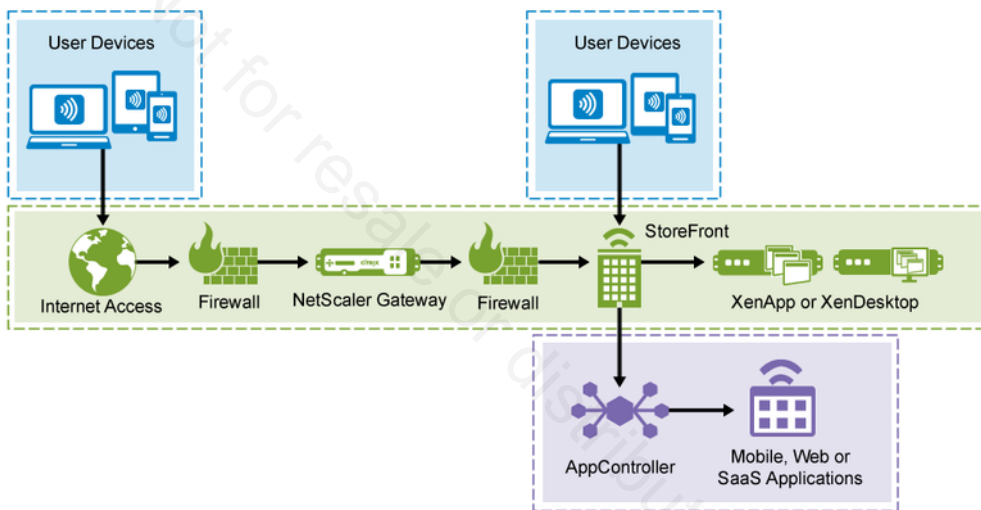
XenMobile App Edition adds advanced mobile application and data management to any MDM solution. In this edition, ShareFile provides read access to files on network drives and SharePoint.

XenMobile Enterprise Edition

Comprehensive solution to manage mobile applications, data and devices. This edition includes ShareFile Enterprise.

NetScaler Gateway with XenMobile Integration

You can configure NetScaler Gateway to work with XenMobile. In this deployment, NetScaler Gateway resides in the perimeter network. AppController and StoreFront reside in the secure network. NetScaler Gateway must have access to the same forest in which StoreFront resides.



When you configure user connections through NetScaler Gateway to AppController or StoreFront, end users can connect in the following ways:

- Receiver
- NetScaler Gateway through a web browser and Receiver for Web
- Receiver for Android or Receiver for iOS

Users can connect through NetScaler Gateway to XenMobile using the following methods:

- **Connect to Receiver for Web using the NetScaler Gateway web address in a web browser.** When end users connect with clientless access and Receiver for Web, they can start their applications from within the web browser. When you configure NetScaler Gateway to support Receiver for Web, other clientless access policies that are bound to the virtual server, such as SharePoint, are not supported.

- **Connect to XenMobile using Receiver for Windows through native protocols.** When end users connect with clientless access to AppController or StoreFront, they download a provisioning file from the Receiver for Web site and install the file on the device. Receiver uses settings within the provisioning file to determine if the end-user device is inside or outside the secure network. Users connect with the NetScaler Gateway web address, such as <https://<NetScalerGatewayFQDN>>. When logon is successful, users can start or subscribe to their Web, SaaS or mobile apps.
- **Connect to AppController using the NetScaler Gateway plug-in.** You can use the NetScaler Gateway plug-in for Windows or NetScaler Gateway plug-in for Mac to connect to web applications hosted by XenMobile.

End users can connect to StoreFront only through the following connection methods:

- **Connect to StoreFront using email-based discovery.** NetScaler Gateway supports Account Services that allows users to connect using an email address or the NetScaler Gateway FQDN. When users log on, Receiver instructs end users about how to configure access.
- **Connect to StoreFront using PNA services.** If end users connect with legacy versions of Receiver for Mac, Receiver for Android, or Receiver for iOS, they must manually configure a store within Receiver using the NetScaler Gateway web address. When end users successfully log on, they can start their published applications and virtual desktops. End users cannot connect with Receiver for Windows with PNA services.

NetScaler Gateway with XenMobile App Edition Deployment

Once you have completed any pre-configuration requirements, you can use the following process to deploy a NetScaler Gateway with XenMobile App Edition:

1. Enable the SSL Offloading and NetScaler Gateway features within the NetScaler system settings.
2. Configure authentication. Create and bind a policy.
3. Install the SSL certificate for the FQDN of the NetScaler Gateway virtual server and the intermediate certificate of the CA (if applicable).
4. Create a session profile for StoreFront-aware receivers. Create a second session profile for browser-based access.
5. Create a session policy to test the User-Agent for any Receiver. Create a second session policy to test the User-Agent for Receiver for Web. Bind the policies to the associated session profiles.
6. Create the NetScaler Gateway virtual server.
7. Optional: Create an HTTP redirect that links to a load-balancing virtual server.

Discussion Question

What are the minimum requirements necessary to set up integration on the NetScaler Gateway?

Module 11

AppExpert Policy Engines

11

Not for resale or distribution

Not for resale or distribution

AppExpert Default Policy Engine Manual

Overview

The default policy engine is an expression language used with several features, such as Content Switching, Responder, Rewrite and URL transformation.

After completing this module, you will be able to:

- Identify the syntax and uses for default policy expressions.
- Explain the policy-binding evaluation process and determine appropriate bind points for policies within the default policy engine.
- Configure and invoke pattern sets with string matching in the default policy engine.
- Extract and transform data from one type to another with typecasting in the default policy engine.
- Identify examples of how Responder, Rewrite and URL transformation are used for hosting modifications or redirection for external Web Interface for Authentication and Authorization.
- Describe content switching and the configuration process.
- Describe rule precedence in content switching.

Understanding Policies

The NetScaler system uses the policies to evaluate specified conditions and to define the actions to be taken if conditions are met. The actions defined are specific to the feature for which the policy is created. The order and flow of policy evaluation depends on the feature set and policy-expression type.

Policy Expressions

Policy expressions are used to determine if a connection matches the certain conditions. The end result of any complete policy expression must produce a Boolean value or, with default policy expressions, an UNDEF value, to tell the NetScaler system explicitly whether a particular connection matches the policy. When a policy expression is bound or activated, the feature set that is bound to or activated evaluates whether the policy expression is true for that specific feature.

Expression Result Types

NetScaler expressions can produce or return data in the form of Boolean, numeric and string values. An expression can include multiple types at the same time. For example it can match a string and a numeric type simultaneously. The following table describes these data types.

Data Type	Description
Boolean values	A Boolean expression is an expression that, when evaluated, returns one of two values: the FALSE value or the TRUE value. You can use a Boolean expression to return an unambiguous response, usually as the main expression in a policy. For example, <code>HTTP.REQ.URL.CONTAINS()</code> is a typical Boolean expression.
Integer (or Number)	An integer or number expression is an expression that, when evaluated, returns a number. You can use this expression to substitute for a literal number in the <code>gotopriorityExpr</code> portion of a rewrite rule or as a component of a Boolean expression used as a policy rule. For example, <code>HTTP.REQ.URL.LENGTH</code> is a typical number expression.
String values	A string expression is an expression that, when evaluated, returns a text string. You can use a string expression to substitute for a literal string, usually as a component of a Boolean expression used as a policy rule, or to specify an HTTP header or body string in a rewrite action. An example of a string expression is <code>HTTP.REQ.HEADER("myHeader")</code> .



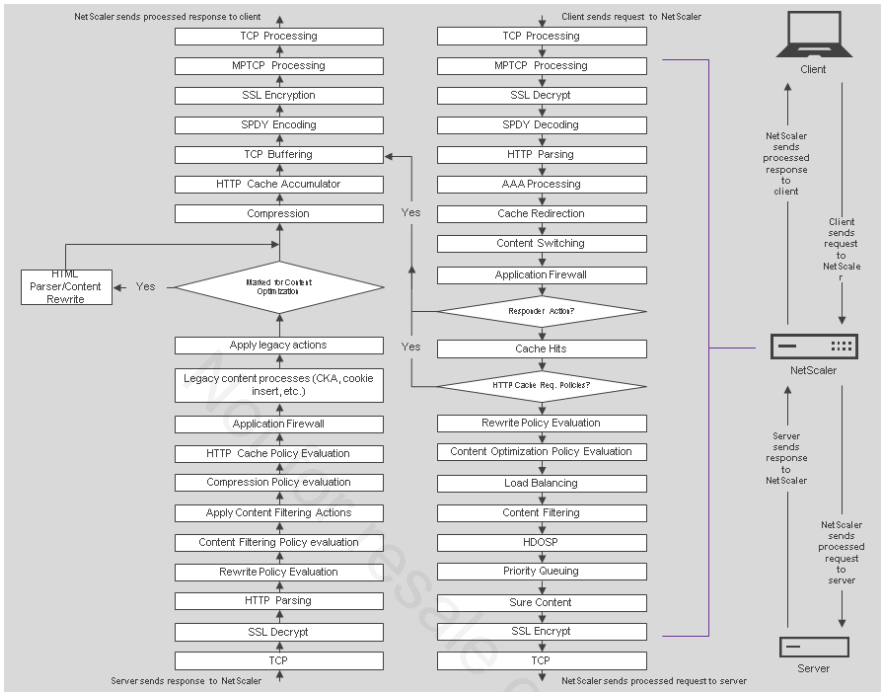
Additional data types, such as IPv6, are available internally within expressions. Most data types can be used in string contexts. Non-string results (except for Boolean results) are converted to strings if the feature requires a string result for an expression.

Understanding Packet-Processing Flow

As traffic through the NetScaler system is evaluated by the different features, each feature performs policy evaluation. Whenever a policy matches the traffic, the NetScaler system stores the action and continues processing. The NetScaler system typically applies all matching actions after processing is completed.

The following figure illustrates the overall packet-processing flow for evaluation order. For the majority of the features, the packet evaluation order does not matter because the events are logged. Note that some of the features, such as Application Firewall, Responder and caching, are first in the packet-processing flow. If there is a match on these, evaluation might not continue depending on

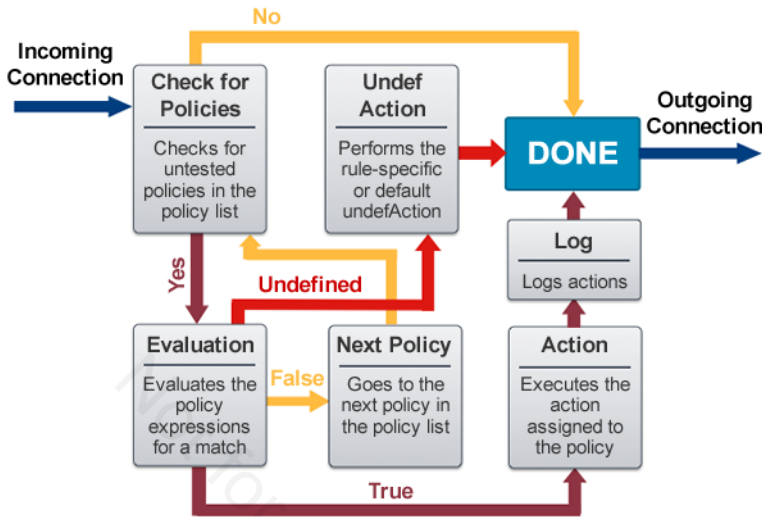
the configured action. These specific feature evaluations are performed more frequently on the incoming request, not the transformed data.



Legend for initialisms used in diagram:

Initialism	Phrase
CF	Content filtering
HDOSP	HTTP denial of service protection
SC	Sureconnect
PQ	Priority queuing
CMP	Compression
CKA	Client keep-alive mode

Policy Process Evaluation Flow



The figure shows how the evaluation of a policy produces three possible outcomes:

True A true outcome means that the connection matches the policy expression.

False A false outcome means that the connection did not match the policy expression.

Undefined An undefined outcome occurs when a policy cannot be evaluated.

The exact packet processing evaluation flow depends on the feature set performing the evaluation; the evaluation is for a single feature or module, not for all the processing that occurs on a message. The "next policy" in the figure is partially under control of the policy binding. If a policy rule is false, then the next policy evaluated is the one with the next highest priority. If the policy rule is true, then one of the following results is experienced:

- Another policy is not evaluated ("END," which is the default).
- The policy with the net highest priority is evaluated ("NEXT").
- The specific policy can be specified (example: a number).

For each feature, the policy engine evaluates all the policies that occur in the order that "next" policy specifies. For any "true" rule, the policy engine queues the action and log action. When the

list of "next" policies is completed, the feature then runs one or more actions (assuming any were queued).

Most features only perform one action (and log action) - the last one that was queued. If "END" is set as the next policy for all policies, then only one action will be queued. For some features, notably rewrite, all queued actions are started. Each feature has two phases: policy evaluation and action execution.

Identifying Default Policy Expressions

The default-policy-expression language is an object-oriented, general purpose expression language capable of extremely precise data manipulation.

Default-policy-expression types include:

- Non-compound default policy expressions
- Compound default policy expressions

In a policy, default policy expressions can be combined using Boolean operators, allowing the NetScaler system to test multiple connection elements in a structured manner. Default expressions are supported in several features, including DNS, GSLB, Content Switching, Responder and Rewrite policies.

Default Policy Expression Syntax

"Dotted function" chains in NetScaler default policy expressions read from left to right. The left-most term designates which part of the connection the expression is analyzing. The values described in the following list can appear in this position.

HTTP

Designates a complete HTTP connection, including both the request and the response. Use this term to extract information from an HTTP connection.

SYS


Designates the NetScaler system and NetScaler operating system. Use this term to operate on the date and time of the connection or to insert a classic expression in a larger default policy expression.



Many more functions are possible with the SYS value; the two mentioned are just examples.

CLIENT

Designates the client portion of the connection. Use this term to extract information from the client, such as the client IP address.

SERVER	Designates the server portion of the connection. Use this term to operate on information from the server, such as the hostname.
ANALYTICS	Designates traffic in order to gather run-time statistics. Use this term in conjunction with the Stream Analytics functionality of the NetScaler system.
SIP	Designates session initiation protocol (SIP). SIP expressions are used primarily when performing SIP load balancing for VOIP and related applications.
TEXT	Designates any piece of text. Use this term to operate on any text values associated with the connection.
	 This value is used only for specialized cases such as CVPN or Application Firewall.
MYSQL (MySQL)	Designates protocol-specific methods and properties for SQL requests. Use this term in conjunction with the DataStream functionality of the NetScaler system.
MSSQL (Microsoft SQL)	Designates protocol-specific methods and properties for SQL requests. Use this term in conjunction with the DataStream functionality of the NetScaler system.
DNS	Designates policies based on DNS requests and responses. DNS expressions are used in conjunction with DNS virtual servers as part of a DNS load- balancing or DNS proxy configuration.
CONNECTION	Designates a particular connection.

URL

Used for typecasting TEXT as a URL (example: TEXT.TYPECAST_HTTP_URL_T). It has the same restrictions in usage as "TEXT * TARGET". When used in a rewrite action for the value (string-builder_) expression, it represents the entire text selected by a target expression of the same action. When used in a rewrite action in the "refineSearch" option, it extends the current match by the specified number of bytes to the left and right.

Successive terms to the right refine the first term to define specific attributes of the connection. Each term is separated from the preceding or the following terms with a period. Arguments appear in parentheses following the term to which they apply.

Default Policy Expression Objects

The following table lists the default policy expression objects:

Prefix	Flow Type	Qualifier	Operator
HTTP	REQ	BODY	CONTAINS
SYS	RES	CACHE_CONTROL	EQ
CLIENT	IP	COOKIE	HOSTNAME
SERVER	TIME	HEADER	PATH
ANALYTICS		IS_VALID	PROTOCOL
SIP		STATUS	QUERY
TEXT		TRACKING	SUFFIX
MYSQL/MSSQL		URL	TYPECAST
DNS		VERSION	XPATH
CONNECTION		More	More



Consider the following functionality and characteristics of default policy expression objects:

- Every function exists in a class (for example, a data type). In other words, a class is a collection of functions; the prefixes are in a special prefix class. Any function returns a result of a specified class.
- A class can inherit from a parent super-class. In that case, the child class inherits all the functions of its parent and adds its own.
- You can read a dotted function expression left to right, starting with the prefix. At each point you remember the current class, starting with the prefix class. As you examine each function, you look it up in its class (if not directly in that class, it could be defined in an ancestor class) and note the resulting class of that function. This is used to set up the current class for the function to its right.
- Functions can be "overloaded." A function with the same name in the same class but with a different number or types of parameters performs a different function and can return a different result type.
- Different classes can have a function of the same name and parameter types--again these can return a different result type.

The following table lists example scenarios and expressions to demonstrate the different types of expressions:

Scenario	Expression Example
Check if the IP address of the client is in the 10.60.1.0/24 range. In this scenario, the evaluation returns a TRUE if the client is in the defined subnet and applies the action configured and bound to the policy.	<code>CLIENT.IP.SRC.IN_SUBNET (10.60.1.0/24)</code>
Check for the content of the HTTP request in the "Referer" header before the double slash. This request returns http: or https: values.	<code>HTTP.REQ.HEADER (Referer).BEFORE_STR</code>
Check for a string in the following URL query: <code>http://ads.example.com/ads/adjs.php?n=829983570&what=zone:ne:399&block=1&blockcampaign=1&exclude=</code>	<code>HTTP.REQ.URL.QUERY.CONTAINS (what=zone:)</code>
Check if the system time is between value1 and value2.	<code>SYS.TIME.BETWEEN (GMT 2006 Feb 01 12h 35m 18s, GMT 2008 Mar 01 12h 00m 00s)</code>
Check if the system time is equal to or greater than value3.	<code>SYS.TIME.GE (GMT 2007 Aug 01 11h 59m 59s)</code>



Certain characters need to be preceded by backslashes when entered in the command-line interface. The entire expression needs to be enclosed in double quotation marks.

Default Policy Conversion

You can convert a classic expression to the default expression syntax by using the `nspepi` conversion tool. You can also use the tool to convert all the classic expressions in the NetScaler configuration to the default syntax (with the exception of NetScaler entities that currently support only classic expressions).

The conversion tool does not convert policies configured for the following features, because the features currently support only classic policies:

- Authentication, Pre-authentication
- Cache redirection
- VPN (session, traffic, and tunnel traffic)
- Content filtering (The Responder feature not only provides you with functionality that is equivalent to that provided by the content-filtering feature but also surpasses the content-filtering feature in the use cases that it supports. Additionally, Responder supports the more powerful default syntax for policy expressions.)

The following NetScaler features support both classic and default syntax expressions and, therefore, support the conversion of classic expressions to default syntax expressions:

- Application Firewall policies
- Authorization policies
- Named expressions
- Compression policies
- Content-switching policies
- SSL
- User-defined, rule-based tokens and persistency (the `rule` parameter that is specified for a load-balancing virtual server)



For more information about how to convert features from classic to default policies, see Citrix article CTX131024 at <http://support.citrix.com>.

Actions

An action:

- Is bound to or activated by policies.

- Cannot depend on results of other actions.
- Is applied at the end of the policy evaluation process.
- Is owned by individual NetScaler features.

For example, actions configured in the Responder module are different from those configured in the Rewrite module. The individual feature ensures that the respective actions are applied.

Action Syntax

Each feature has feature-specific parameters for its actions. For example, at the NetScaler command prompt, you can type the following command to add the rewrite action command:

```
add rewrite action Replace_HTTP_to_HTTPS INSERT_AFTER
"HTTP.RES.HEADER(\"Location\").Value(0).Prefix(4) " "\"s\""
```

The following command is an example of action syntax:

```
add rewrite action ClientIP INSERT_HTTP_HEADER Client-
ip CLIENT.IP.SRC
```

This example illustrates that the add rewrite action command consists of different parts when used in the command-line interface. The following list defines the components of the command:

- ClientIP: The action name
- INSERT_HTTP_HEADER: The rewrite action
- Client-IP: The HTTP header name
- CLIENT.IP.SRC: The expression that produces the value for the individual client IP address

Configuring Policies and Actions

Rewrite and responder policies are configured using the default policy engine. These policies must have an action and a rule and must be bound to a bind point. After the policies are bound to a bind point, the NetScaler system performs rewrite actions.

Configuring Rewrite or Responder Policies

Use the following steps to configure a rewrite or responder action and policy:

1. Define an action to be performed.
2. Create a policy.
3. Define the rule, or expression, which determines when to apply the action.
4. Attach the action for the outcome of the evaluation.

5. Bind the policy (rule + action) to a bind point to perform rewrite.
6. Attach a priority to each policy, which determines the sequence of policy execution.
7. Optional: Refer to the next policy to be evaluated using a goto expression.
8. Optional: Invoke a policy label.

Understanding Bind Points

A bind point is the entity to which the policy is bound. Available bind points vary by feature set. For example, Responder policies can be bound either globally, to policy labels, or to virtual servers on the NetScaler system.

With default policies, policies for some features can be bound to a new policy bind point, referred to as a policy label, or sometimes as a policy bank. A policy label is a group of policies. Each policy in a policy label must be given a priority. In addition, policies have two distinct characteristics:

- Every policy is evaluated in priority order.
- No two policies can have the same priority.

In order to invoke a policy, it must be bound to an entity. You can bind a classic policy either at the global level or at the virtual server (content-switching or load-balancing virtual server) level. The default policy engine also allows you bind policies in this manner, but it offers more flexibility on how policies are bound and evaluated.

The default policy engine defines the following bind points:

- | | |
|---------------------------------|--|
| request-time bind point | Request-time bind points are evaluated and applied to requests only. |
| response-time bind point | Response-time bind points are evaluated and applied to responses only. |

You can bind a default policy to any of the following entities:

- Override global
- Load-balancing virtual server
- Content-switching virtual server
- Default global



Global bind points are used to compare policies against all traffic passing through a NetScaler system, while virtual server bind points affect only traffic specific to a single virtual server.

When binding a policy in the default policy engine, you must assign a priority. Priorities control the order of processing policies within the bind point, while the bind points have a precedence that determines which bind points are evaluated first. If multiple policies affect a single response or

request, the highest precedent policy determines the action. Bind points are processed in the following order:

1. Override global
2. Content-switching virtual server
3. Load-balancing virtual server
4. Default global



Default global policies only apply if no higher-precedent policies affect the traffic at the override global or virtual server bind points.

Policies should only be bound to a global object if the setting is needed globally or needed to apply to a large range of traffic. Otherwise, policies that are specific to an application should be bound to the appropriate virtual server.

Policy gotoPriorityExpression Statement

Processing of policies ends after all valid policies have been processed or an END is encountered on a policy match. The following behavior applies when a policy has a gotoPriorityExpression statement:

- If the condition set in the policy is TRUE, go to the NEXT specified in the goto statement priority order.
- If no gotoPriorityExpression is set, the NEXT is assumed until the final policy is reached or END is specified, at which point processing ends.

The following list contains valid values for a gotoPriorityExpression statement.

integer	This value is equal to an existing policy, but the priority must be higher numbered than the current policy to avoid loops.
NEXT	This value signals to go to the next policy in the list.
END	This value signals that the policy evaluation has ended and the actions residing in the resulting set are applied.

Policy priority is set to disallow goto statements to create loops; the gotoPriorityExpression cannot go back to the lower priority. For example, a NetScaler system has three policies of priority 10, 20 and 30. The administrator cannot create a goto statement sending a match to the policy with a priority of 30 back to the policy with a priority of 10.



A gotoPriorityExpression statement is applied only if the policy evaluated has an outcome that is TRUE.

Policy Binding Evaluation Process

In the policy evaluation process, a request is being evaluated by the NetScaler system. The following explains the evaluation process:

1. The bind point selector selects a bank based on the configuration.
2. The NetScaler system evaluates the request for matches in the selected bank.
3. The NetScaler system finds a match and the policy evaluation process ends. The specified actions are performed.
4. The NetScaler system invokes the next bank of policies and the request is evaluated against the next policy bank if no match is found.
5. The NetScaler system performs the identified action if a match is found.

If no policy evaluates to TRUE in any of the policy banks that were selected, the specific feature determines the outcome.

Binding in the Policy Manager

The policy manager provides an easy interface for managing bind points, priorities, and policy labels in the configuration utility. The policy manager makes configuring priorities within bind points much easier because the logic is built in and policy manager will not let you configure faulty logic.

Understanding Policy Labels

The logic for evaluating user-defined policy labels is the same as for evaluating other policies in banks and labels. You can create the same type of policy labels for policies that are attached to virtual servers and globally (both override and default) on the request and response side, for each valid feature and protocol combination.

Policy labels are invoked from other policies. When a policy label is invoked, all the policies bound to it are evaluated in the order of priority. When a policy is matched and the expression is evaluated to be true, the appropriate action is logged and the control is returned to the policy that invoked the policy label.

The policy defines the next policy to be evaluated using gotoPriority, or the action ends and goes to the next priority policy or the next policy bank. Policy labels are specific to certain features, such as the Rewrite and Responder features.

Configuring Policy Labels

In the command-line interface, you can type the following command to add a Rewrite policy label and Responder policy label:

```
add rewrite policylabel <labelName> <transform>
```

Pattern Sets

A pattern set is an array of indexed patterns that you configure on the NetScaler system. Pattern sets are used for string matching during default policy evaluation. The NetScaler system provides you with a set of default expression operators that you can use to compare a string in a packet with the patterns that are indexed and stored in a pattern set.

You can configure the system to compare the string that is identified in a packet with one or more patterns in the pattern set by using a simple default expression combined with an operator. Additionally, after you create a pattern set, you can use the pattern set in multiple default policies. Therefore, pattern sets eliminate the need for you to configure compound default expressions that perform string matching with multiple OR operations (one expression for each comparison). This also reduces the consumption of system resources in terms of memory and the number of expressions that the system has to evaluate.

First, you create a pattern set and bind patterns to it. Then, when you configure a policy for comparing a string in a packet with the pattern set, you use an appropriate operator and pass the name of the pattern set as an argument.

How String Matching with a Pattern Set Works

A pattern set contains a set of patterns, and each pattern is assigned a unique index. During policy evaluation, the operator compares the string that is identified in the packet with the patterns defined in the pattern set until a match is found. Then, depending on its function, the operator returns either a Boolean value that indicates whether a matching pattern was found or the index of the pattern that matches the string.

Configuring a Pattern Set

You configure a pattern set by specifying the strings that are to serve as patterns and binding them to the pattern set. Each pattern in the set has a unique index value. If you specify an index for the first pattern that you bind to a pattern set, you must specify an index for all the other patterns in the set, and you can modify the values at any time. If you do not specify an index for the first pattern, the NetScaler system assigns the pattern an index of value 1. Thereafter, the system assigns an index value to all the patterns that you bind to the set, and you cannot change them. After you have configured a pattern set for which the system has generated index values automatically, if you want to assign index values of your choice, you must create a new pattern set.

Index values are not regenerated automatically if one or more patterns are deleted or modified. For example, if the set contains five patterns, with indexes from 1 through 5 and if the pattern with an index of 3 is deleted, the other index values in the pattern set are not automatically regenerated to produce values from 1 through 4.

In the NetScaler command-line interface, you first create a pattern set by using the `add policy patset` command. Then, you create patterns and bind them to the pattern set, one pattern at a time, by using the `bind policy patset` command. In the NetScaler configuration utility, you perform all these tasks in a single dialog box.



Pattern sets are case sensitive. Therefore, the string pattern "product1," for example, is not the same as the string pattern "Product1." Case sensitivity can be circumvented by using the "ignorecase" operator.

For more information about configuring a pattern set, see Citrix product documentation at <http://docs.citrix.com>.

Using Pattern Sets

After you configure a pattern set, you can use it in a default expression that passes the pattern set as an argument to an appropriate operator.

NetScaler VPX (1000) HA Status: Not Configured Info: NS11.0 62.10.nc

Dashboard Configuration Reporting Documentation

← Back

Create String Map

Name*
sm_search_engines

Insert Delete

Key	Value
search	yahoo

Comments

Create Close

The following is the format of the default syntax expression that compares a string with a pattern set:

```
<text>.<pattern set operator>("<name>")
```

In the preceding expression format, <text> represents any default syntax expression that identifies a string in a packet and <name> is the name of the pattern set.

If you want the system to determine whether the value of the host header in a client request contains any of the patterns that are configured in a pattern set called "Patternset1", you can use the following default syntax expression:

```
HTTP.REQ.HEADER("Host").CONTAINS_ANY("Patternset1")
```

Pattern-Set Operators

The following table describes the operators that you can use with pattern sets. When you use an operator, replace <text> with the default syntax expression that identifies the string with which you want to perform string matching, and replace <pattern_set_name> with the name of the pattern set.

Operator	Description
<text>.CONTAINS_ANY(<pattern_set_name>)	Evaluates whether the target text contains any of the patterns that are bound to <pattern_set_name> and returns a Boolean TRUE if one or more matching patterns are found
<text>.EQUALS_ANY(<pattern_set_name>)	Evaluates whether the target text exactly matches any of the patterns that are bound to <pattern_set_name> and returns a Boolean TRUE or FALSE to indicate the result of the evaluation
<text>.ENDSWITH_ANY(<pattern_set_name>)	Evaluates whether the target text ends with any of the patterns that are bound to <pattern_set_name> and returns a Boolean TRUE or FALSE to indicate the result of the evaluation
<text>.STARTSWITH_ANY(<pattern_set_name >)	Evaluates whether the target text starts with any of the patterns that are bound to <pattern_set_name> and returns a Boolean TRUE or FALSE to indicate the result of the evaluation

Operator	Description
<code><text>.STARTSWITH_INDEX(<pattern_set_name>)</code>	Evaluates whether the target text starts with any of the patterns that are bound to <code><pattern_set_name></code> and, if a match is found, returns the numerical index of the matching pattern
<code><text>.ENDSWITH_INDEX(<pattern_set_name>)</code>	Evaluates whether the target text ends with any of the patterns that are bound to <code><pattern_set_name></code> and, if a match is found, returns the numerical index of the matching pattern
<code><text>.CONTAINS_INDEX(<pattern_set_name>)</code>	Evaluates whether the target text contains any of the patterns that are bound to <code><pattern_set_name></code> and, if a match is found, returns the numerical index of the matching pattern
<code><text>.EQUALS_INDEX(<pattern_set_name>)</code>	Evaluates whether the target text exactly matches any of the patterns that are bound to <code><pattern_set_name></code> and, if an exact match is found, returns the numerical index of the pattern
<code><text>.SUBSTR_ANY(<pattern_set_name>)</code>	Selects the first string that matches any pattern in the given pattern set

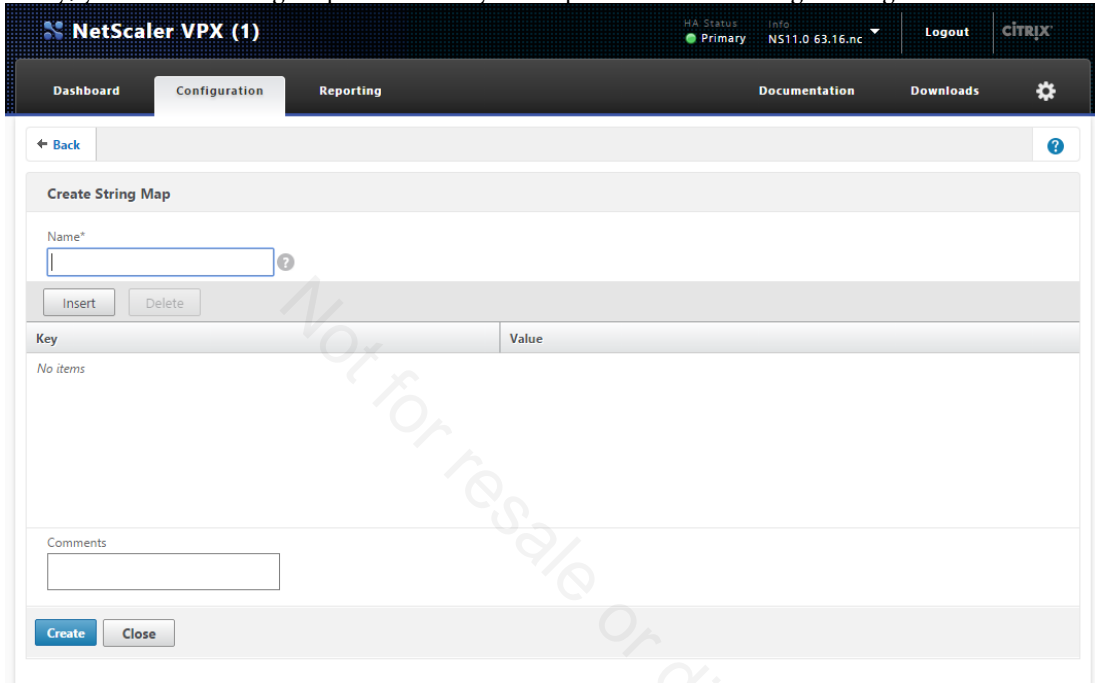
String Maps

You can use string maps to perform pattern matching in all NetScaler features that use the default policy syntax. A string map is a NetScaler entity that consists of key value pairs. The keys and values are strings in either ASCII or UTF-8 format. String comparison uses two new functions, `MAP_STRING (<string_map_name>)` and `IS_STRINGMAP_KEY (<string_map_name>)`.

A policy configuration that uses string maps performs better than one that does string matching through policy expressions and you need fewer policies to perform string matching with a large number of key value pairs. String maps are also intuitive, simple to configure and result in a smaller configuration.

String maps are similar in structure to pattern sets (a pattern set defines a mapping of index values to strings; a string map defines a mapping of strings to strings) and the configuration commands for string maps (commands such as `add`, `bind`, `unbind`, `remove` and `show`) are syntactically similar to configuration commands for pattern sets. Also, as with index values in a pattern set, each key in a string map must be unique.

You first create a string map and then bind key value pairs to it. You can create a string map from the command-line interface or the configuration utility. In the command-line interface, you first use the `add policy stringmap` command to create a string map. You then use the `bind policy stringmap` command to bind key value pairs, one pair at a time. In the configuration utility, you create a string map and bind key-value pairs to it from a single dialog box.



The screenshot shows the NetScaler VPX (1) configuration utility interface. The top navigation bar includes 'Dashboard', 'Configuration' (selected), 'Reporting', 'Documentation', and 'Downloads'. The 'Create String Map' dialog box is open, featuring a 'Name*' input field, 'Insert' and 'Delete' buttons, a table with 'Key' and 'Value' columns (currently empty), a 'Comments' input field, and 'Create' and 'Close' buttons at the bottom.

To create a string map by using the NetScaler command line, type the following command:

```
add policy stringmap <name> [-comment <string>]
```



For more information about string maps, see Citrix product documentation at <http://docs.citrix.com>.

Typecasting

Typecasting allows you to extract data of one type (such as text or an integer) from requests and responses and transform it to data of another type. For example, you can extract a string and transform the string to time format. You can also extract a string from an HTTP request body and treat it like an HTTP header or extract a value from one type of request header and insert it in a response header of a different type.

After typecasting the data, you can apply any operation that is appropriate for the new data type. For example, if you typecast text to an HTTP header, you can apply any operation that is applicable to HTTP headers to the returned value.

Typecasting is only available with the default policy engine.

The following table describes some of the values for controlling typecast results. For more information about typecasting data including more typecasting functions and examples, see Citrix product documentation at <http://docs.citrix.com>.

Function	Description
<code><text>.TYPECAST_LIST_T(<separator>)</code>	Treats the text in an HTTP request or response body as a list whose elements are delimited by the character in the <code><separator></code> argument. Index values in the list that is created start with 0. Text mode settings have no effect on the separator. For example, even if you set the text mode to IGNORECASE and the separator is the letter "p," an uppercase "P" is not treated as a separator.
<code><text>.TYPECAST_TIME_T</code>	Treats the designated text as a date string. The following formats are supported: <ul style="list-style-type: none">• RFC822: Sun, 06 Nov 1994 08:49:37 GMT• RFC850: Sunday, 06-Nov-94 08:49:37 GMT• HTTP Set-Cookie Expiry date: Sun, 06-Nov-1994 08:49:37 GMT
<code><text>.TYPECAST_HTTP_URL_T</code>	Treats the designated text as the URL in the first line of an HTTP request header. The supported format is <code>[<protocol>://<hostname>]<path>?<query></code> and the text mode is set to URLENCODED by default.

Function	Description
<prefix>.TYPECAST_NUM_T(<type>)	<p>Casts numeric string data to a signed 32-bit number. The argument <type> can be one of the following:</p> <ul style="list-style-type: none"> • DECIMAL. Treat the string as a decimal number and cast to a signed 32-bit number. • HEX. Treat the string as a hexadecimal number and cast to a signed 32-bit number. • DECIMAL_PREFIX. Consider the part of the string up to the first occurrence of a character that is not a valid decimal character and cast to a signed 32-bit number. • HEX_PREFIX. Consider the part of the string up to the first occurrence of a character that is not a valid hexadecimal character and cast to a signed 32-bit number.

The following example illustrates that typecasting expressions identify specific actions to perform on a part of a URL string. The expression takes the string after what=zone and converts it into an integer value. The expression checks if the value is greater than or equal to 399.

```
HTTP.REQ.URL.QUERY.AFTER_STR("\"what=zone\"").BEFORE_STR("&block\"")
.TYPECAST_NUM_T(DECIMAL).GE(399)
```

The following example illustrates that a typecasting expression puts the query string of the URL in a Name Value list. The expression also extracts the value of the eighth element of the string and converts it to a number for evaluation.

```
HTTP.REQ.URL.QUERY.VALUE(7).TYPECAST_NUM_T(DECIMAL)
```

Responder

With the Responder feature, responses to HTTP requests can be customized according to who sends the request, where it is sent from and other criteria with security and system management implications. For example, you can designate appropriate home pages for a given web site home page request based on the end user's location, their browser, language, or the order of preference. You might want to break the connection immediately if the request is coming from an IP address range that has been generating DDoS attacks or initiating hacking attempts.

Responder can be used for the following scenarios depending on client parameters:

- Redirecting an HTTP request to new websites or webpages
- Responding with a custom response
- Dropping or resetting a connection at request level

For more information about the responder feature, see Citrix product documentation at <http://docs.citrix.com>.

Rewrite

Rewrite refers to the rewriting of some information in the requests or responses handled by the NetScaler system. Rewriting can help in providing access to the requested content without exposing unnecessary details about the website's actual configuration. A few situations in which the rewrite feature is useful are described below:

- To improve security, the NetScaler can rewrite all the `http://` links to `https://` in the response body.
- In the SSL offload deployment, the non-secure links in the response have to be converted into secure links. Using the rewrite option, you can rewrite all the `http://` links to `https://` to ensure that the outgoing responses from NetScaler to the client have the secured links.
- If a website has to show an error page, you can show a custom error page instead of the default 404 Error page.
- If you want to launch a new website but use the old URL, you can use the rewrite option.
- When a topic in a site has a complicated URL, you can rewrite it with a simple, easy-to-remember URL.
- You can append the default page name to the URL of a website.

When you enable the rewrite feature, NetScaler can modify the headers and body of HTTP requests and responses.

For more information about the rewrite feature, including rewrite action and policy examples, see Citrix product documentation at <http://docs.citrix.com>.

URL Transformation

The URL transformation feature provides a method for modifying all URLs in designated requests from an external version seen by outside users to an internal version seen only by your web servers and IT staff. You can redirect end-user requests seamlessly without exposing your network structure to them. You can also modify complex internal URLs that users might find difficult to remember into simpler, more easily remembered external URLs.



Before you can use the URL transformation feature, you must enable the rewrite feature.

For more information about URL transformation, see Citrix product documentation at <http://docs.citrix.com>.

Responder Process



The following steps illustrate the Responder process:

1. The client browser sends a request to the web server through the NetScaler system.
2. The NetScaler system checks the request time policy bank for applicable policies.
3. The NetScaler system builds a set of actions to apply after evaluating the list of prioritized policies.
4. The NetScaler system responds to the client request with either a "redirect" or a "respondwith" action.

Responder Actions

After enabling the Responder feature, you must configure one or more actions for handling requests.

The Responder supports the following types of actions:

Respond with	Sends the response defined by the Target expression without forwarding the request to a web server. (The NetScaler system substitutes for and acts as a web server.) Use this type of action to manually define a simple HTML-based response. Normally the text for a "Respond with" action consists of a web server error code and a brief HTML page.
Respond with SQL OK	Sends the designated SQL OK response defined by the Target expression. Use this type of action to send an SQL OK response to an SQL query.
Respond with SQL Error	Sends the designated SQL Error response defined by the Target expression. Use this type of action to send an SQL Error response to an SQL query.
Respond with HTML page	Sends the designated HTML page as the response. You can choose from a list of HTML pages that were previously uploaded, or upload a new HTML page. Use this type of action to send an imported HTML page as the response.
Redirect	Redirects the request to a different webpage or web server. A redirect action can redirect requests originally sent to a "dummy" website that exists in DNS, but for which there is no actual web server, to an actual website. It can also redirect search requests to an appropriate URL. Normally, the redirection target for a redirect action consists of a complete URL.

For more information about configuring a Responder action, see Citrix product documentation at <http://docs.citrix.com>.

Respond With

A "Respond with" action is used to send an HTTP string as a response to a client.

The following table describes the parameters you must configure in order to create a "Respond with" action.

Parameter	Description
Name	The name of the action. This is a mandatory parameter and cannot be changed. The action name must not exceed 31 characters.
Type	The type of Responder action being configured. This parameter determines the behavior of the Responder action. You can choose to respond to the client or redirect the client elsewhere. Select the "Respond with" responder action type to send the configured target string to the client.
Target	This parameter contains the HTTP string to be sent as a response to the client.

To Create a Respond With Action

1. In the left pane, click **AppExpert > Responder > Actions**. The Responder Actions page appears in the right pane.
2. Click **Add**. The Create Responder Action dialog box appears.
3. In the Name text box, type the name of the responder action.
4. Under Type, select **Respond with**.
5. In the Target text area, type "HTTP/1.1 200 OK\r\n\r\n" + "Client: " + CLIENT.IP.SRC + " is not authorized to access URL:" + HTTP.REQ.URL.HTTP_URL_SAFE.
6. Click **Create**. The Responder action you created appears in the Responder Actions page.

Responder Action for Timeouts

You can invoke a Responder action when an HTTP request times out. To configure this feature, first create the Responder action that you want to invoke. Then, configure the global HTTP timeout action.

To configure the global HTTP timeout action to invoke a Responder action by using the command-line interface, type:

```
set ns httpProfile -reqTimeoutAction <responder action name>
```


Responder Policies

Responder policies are configured in the configuration utility and in the command-line interface. The following arguments are identified when adding a Responder policy:

Rule	Expression used by the Responder policy that returns a Boolean result
Action	Responder action name, NOOP, RESET or DROP
UndefAction	Responder action to be taken in the case of an undefined event during policy evaluation, with value of NOOP, RESET or DROP

Responder HTML Page Imports

The Responder feature can respond to designated requests by sending the client an HTML-based webpage that you upload to the NetScaler system. You have the option of redirecting the request or responding with a response code and answer configured on the NetScaler itself.

To use this feature, first upload an HTML-based webpage to the NetScaler by using either the NetScaler command-line interface or the configuration utility. Next, configure a Responder action with type set to RespondWithHTMLPage and the name of the HTML page. Finally, create a Responder policy and bind it to the action.

To upload an HTML page to the responder feature, at the NetScaler command prompt type the following commands:

```
import responder htmlpage [src] name [-comment comment] [-  
overwrite]
```

For more information about configuring a Responder action, see Citrix product documentation at <http://docs.citrix.com>.

Responder Undefined Actions

The NetScaler system generates an undefined event (UNDEF event) when a request does not match a responder policy, and then carries out the default action assigned to undefined events. By default, that action is to forward the request to the next feature without changing it. This default behavior is normally what you want; it ensures that requests that do not require special handling by a specific responder action are sent to your web servers and that clients receive access to the content that they requested.

If the websites that your NetScaler system protects receive a significant number of invalid or malicious requests, however, you may want to change the default action to either reset the client connection or drop the request. In this type of configuration, you would write one or more Responder policies that would match any legitimate requests and simply redirect those requests to their original destinations. Your NetScaler system would then block any other requests as specified by the default action you configured.

You can assign any one of the following actions to an undefined event:

NOOP

The NOOP action cancels Responder processing but does not alter the packet flow. This means that the system continues to process requests that do not match any Responder policy and eventually forwards them to the requested URL unless another feature intervenes and blocks or redirects the request. This action is appropriate for normal requests to your web servers and is the default setting.

RESET

If the undefined action is set to RESET, the system resets the client connection, informing the client that it must re-establish its session with the web server. This action is appropriate for repeat requests for webpages that do not exist, or for connections that might be attempts to hack or probe your protected websites.

DROP

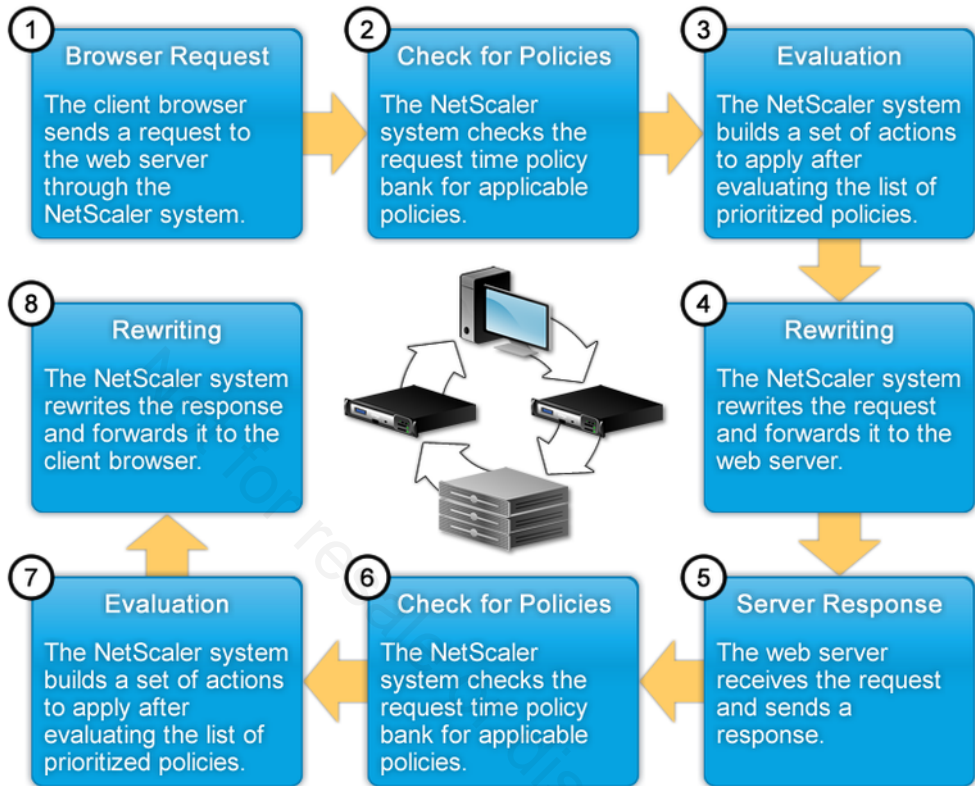
If the undefined action is set to DROP, the system silently drops the request without responding to the client in any way. This action is appropriate for requests that appear to be part of a DDoS attack or other sustained attack on your servers.



UNDEF events are triggered only for client requests. No UNDEF events are triggered for responses.

For more information about setting the responder default action, see Citrix product documentation at <http://docs.citrix.com>.

Rewrite Process



The following steps describe the rewrite process:

1. The client browser sends a request to the web server through the NetScaler system.
2. The NetScaler system checks the request time policy bank for applicable policies.
3. The NetScaler system builds a set of actions to apply after evaluating the list of prioritized policies.
4. The NetScaler system rewrites the request and forwards it to the web server.
5. The web server receives the request and sends a response.
6. The NetScaler system checks the response time policy bank for applicable policies.
7. The NetScaler system builds a set of actions to apply after evaluating the list of prioritized policies.
8. The NetScaler system rewrites the response and forwards it to the client browser.

Once rewrite rules and policies are configured, enabled, and bound to the proper policy bank, the NetScaler system begins to apply those policies. When an end user's browser sends a request to the

web server, the NetScaler system checks the request time policy bank. If the NetScaler system finds rewrite policies, the system evaluates each policy in order of priority, starting with the lowest number and proceeding to the highest number.

Configuring a Rewrite Action

After enabling the rewrite feature, you need to configure one or more actions unless a built-in rewrite action is sufficient. To configure a rewrite action, you assign it a name, specify an action type and add one or more arguments specifying additional data. The following table describes the more common action types and the arguments you use with them.



Action types that can be used only for HTTP rewrite are identified in the Rewrite Action Type column.

Rewrite Action Type	Argument 1	Argument 2
INSERT_HTTP_HEADER: Inserts the HTTP header you specify into the HTTP request or response. This is the default choice. This action type can be used only with HTTP requests and responses.	The HTTP header you want to insert. For example, if you want to insert the client IP address from which a request is sent, type <code>Client-IP</code> .	A string expression that describes the contents of the header you want to insert. For example, if you want to insert the Client IP from which a request is sent, type <code>CLIENT.IP.SRC</code> .
INSERT_BEFORE: Inserts a new string before the designated string.	A string expression that describes the string before which you want to insert a new string. For example, if you want to find the hostname <code>www.example.com</code> and insert a string before the <code>example.com</code> portion, type the following: <code>HTTP.REQ.HOSTNAME.BEFORE_STR ("example.com")</code> .	A string expression that describes the new string you want to insert. For example, if you want to insert the new string <code>en</code> before the string <code>example</code> in the hostname, type <code>en</code> followed by a period (<code>.</code>).

Rewrite Action Type	Argument 1	Argument 2
<p>INSERT_AFTER: Inserts a new string after the designated string.</p>	<p>A string expression that describes the string after which you want to insert a new string. For example, if you want to find the hostname <code>www.example.com</code> and insert a string after the "www" portion, type the following:</p> <pre>HTTP.REQ.HOSTNAME.AFTER_STR ("www. ")</pre>	<p>A string expression that describes the new string you want to insert. For example, if you want to insert the new string "en". after the string "www". in the hostname, type <code>en</code> followed by a period (.).</p>
<p>REPLACE: Replaces the designated string with a different string.</p>	<p>A string expression that describes the string you want to replace with a new string. For example, if you want to replace the entire hostname in the host header, type</p> <pre>HTTP.REQ.HOSTNAME.SERVER.</pre>	<p>A string expression that describes the new string you want to insert. For example, if you want to replace the current host header with the string <code>web01.example.net</code>, type <code>web01.example.net</code>.</p>
<p>REPLACE_HTTP_RES: Replace the HTTP response with the value specified in the target field. This action type can be used only with HTTP requests and responses.</p>	<p>A string expression that describes the string with which you want to replace the HTTP response. For example, type <code>HTTP 200 OK You are not authorized to view this page</code> to replace the entire HTTP response with this warning.</p>	
<p>DELETE: Deletes the designated string.</p>	<p>A string expression that describes the string you want to delete. For example, if you want to find and delete the string <code>.en</code> in the hostname of HTTP response headers, type the following:</p> <pre>HTTP.RES.HEADER("Host").SUBSTR(".en.")</pre>	

Rewrite Action Type	Argument 1	Argument 2
DELETE_HTTP_HEADER: Deletes the designated HTTP header, including all header contents. This action type can be used only with HTTP requests and responses.	The name of the HTTP header you want to delete. For example, if you want to delete the cache-control header from HTTP responses, type <code>HTTP.RES.HEADER("Cache-Control")</code> .	
DELETE_ALL: Delete every occurrence of the pattern specified in the target text reference.	The part of either the HTTP request or response where you want the deletion to occur.	A string pattern after which the deletion should occur.

In the command-line interface, type the following commands to create a new rewrite action and verify the configuration:

```
add rewrite action insertact INSERT_HTTP_HEADER "client-IP" CLIENT.IP.SRC
show rewrite action insertact
```

Rewrite Action Parameters

The following parameters will be used to configure a rewrite action:

- name** A name for your new action or the name of the existing action you want to modify or remove. The name can begin with a letter, number, or the underscore symbol and can comprise as many as 127 letters and numbers and the hyphen (-), period (.), pound (#), space (), at sign (@), equals (=), colon (:), and underscore (_) symbols. You should choose a name that will make it easy for others to tell what this action is supposed to do. (This cannot be changed for an existing action.)
- pattern or patset** An expression that describes the rewrite operation itself.

bypassSafetyCheck	Whether to bypass the built-in safety checks when adding or modifying this action. Values: YES, NO. Default: NO. You can bypass the safety check for actions that would otherwise generate error messages.
target	A NetScaler default syntax expression that describes the text to be rewritten by the rewrite action. For TCP rewrite actions, the target expression must begin with either CLIENT.TCP.PAYLOAD or SERVER.TCP.PAYLOAD.
stringBuilderExpr	Expression specifying the new value of the rewritten packet. Maximum length of a string literal that can be used inside the expression is 255 characters. A string literal that contains more than 255 characters can be split into smaller chunks of 255 characters each. The chunks can then be concatenated with the + operator. Maximum length of the input expression is 8191.
search	Searches for the designated string or expression in the HTTP header or body. You can use a search expression with actions of the following types: delete_all, insert_after_all, insert_before_all and replace_all.
refineSearch (-refineSearch "extend(#,#).<expression>)	A means of more efficiently searching and rewriting a lengthy response than with a simple search command. Instead of using the search command to search for all occurrences of a regular expression and then rewriting the matched strings, refineSearch enables you to search for a simple string that appears within the text you want to rewrite, refine that search by including surrounding context and then search only the selected text and context for a regular expression match.

For more information about parameters for configuring a rewrite action, see Citrix product documentation at <http://docs.citrix.com>.

Rewrite Undefined Actions

An undefined event is triggered when the NetScaler cannot evaluate a policy, usually because it detects a logical or other error in the policy or an error condition on the NetScaler. When the rewrite policy evaluation results in an error, the specified undefined action is carried out.

Undefined actions configured at the rewrite policy level are carried out before a globally configured undefined action.

The NetScaler supports the following three types of undefined actions:

undefAction NOREWRITE Cancels rewrite processing, but does not alter the packet flow. This means that the NetScaler continues to process requests and responses that do not match any rewrite policy and eventually forwards them to the requested URL unless another feature intervenes and blocks or redirects the request. This action is appropriate for normal requests to your web servers and is the default setting.

undefAction RESET Resets the client connection. This means that the NetScaler tells the client that it must re-establish its session with the web server. This action is appropriate for repeat requests for webpages that do not exist, or for connections that might be attempts to hack or probe your protected websites.

undefAction DROP Silently drops the request without responding to the client in any way. This means that the NetScaler simply discards the connection without responding to the client. This action is appropriate for requests that appear to be part of a DDoS attack or another sustained attack on your servers.



Undefined events can be triggered for both request and response flow specific policies.

For more information about configuring the default rewrite actions, see Citrix product documentation at <http://docs.citrix.com>.

To Configure a Rewrite Action Using the Configuration Utility

1. In the navigation pane, expand **AppExpert**, **Rewrite** and then click **Actions**.
2. In the details pane, do one of the following:
 - a. To create a new action, click **Add**.
 - b. To modify an existing action, select the action and then click **Edit**.
3. In the **Add Rewrite Action** or **Configure Rewrite Action** dialog box, specify values for the parameters.

4. Click **Create** or **OK**. A message appears in the status bar stating that the action has been configured successfully.
5. Repeat steps 2 through 4 to create or modify as many rewrite actions as necessary.
6. Click **Close**.

Configuring a Rewrite Action Using the Command-line Interface

In the command-line interface, type:

```
add rewrite action <name> <type> <target> [<stringBuilderExpr>]
[(-pattern <expression> | -patset string)] [-
bypassSafetyCheck (YES|NO)]

show rewrite action <name>
```

The following commands are example usages:

```
add rewrite action insertact INSERT_HTTP_HEADER "client-
IP" CLIENT.IP.SRC

show rewrite action insertact
```

Rewrite Policies

A rewrite policy consists of a rule and an action. The rule determines the traffic on which rewrite is applied and the action determines the action to be taken by the NetScaler. You can define multiple rewrite policies. For each policy, specify the bind point and priority.

A bind point refers to a point in the traffic flow at which the NetScaler examines the traffic to verify whether any rewrite policy can be applied to it. You can bind a policy to a specific load-balancing or content-switching virtual server, or make the policy global if you want the policy to be applied to the entire traffic handled by the NetScaler. These policies are referred to as global policies.

In addition to the user-defined policies, the NetScaler has some default policies. You cannot modify or delete a default policy.

Binding Policies

The NetScaler system only processes Responder or rewrite policies that are bound. Binding can be done using the configuration utility or the command-line interface.

Each policy needs a priority assigned to it that is a positive integer constant. A lower value priority number is interpreted as a higher priority. For example, a priority value set to 10 is a higher priority than a priority value set to 30. Duplicate priorities are not allowed within each bind point.

Possible GotoPriorityExpression values include:

END	Terminates policy evaluation and proceeds to apply the action
NEXT	Proceeds to the next policy in the priority ranking
Positive integer	Proceeds to the policy with the priority ranking as the specified type, indicating the type of global bind point

Discussion Question

In which situations would it be appropriate to use responder? Rewrite? URL transformation?

AppExpert Classic Policy Engine Manual

Overview

For many NetScaler features, policies control how a feature evaluates data, which ultimately determines what the feature does with the data. A policy uses a logical expression, also called a rule, to evaluate requests, responses, or other data, and applies one or more actions determined by the outcome of the evaluation. Alternatively, a policy can apply a profile, which defines a complex action.

After completing this module, you will be able to:

- Identify the classic policy expression structure and components.
- Identify the difference between classic and default policy structures.
- Explain the types of policies available with NetScaler Gateway and describe policy bind points.
- Implement content-filtering policies to manage traffic.

Policies Overview

Some NetScaler features use default syntax policies, which provide greater capabilities than classic policies. You might have to manually migrate policies from classic to default (advanced) syntax if the classic policy features are now using default policies in the latest NetScaler release.

Policy Basics

Policies provide the foundation for the behavior of most NetScaler features, enabling the features to interpret the data, such as SSL requests, HTTP requests, TCP requests, and HTTP responses that pass through it. Typically, if a policy matches the data being evaluated, the NetScaler system takes one or more actions that are associated with the policy. For example, if an HTTP request matches a rewrite policy, the NetScaler can take an action to change (rewrite) the information in the request.

Most NetScaler features have built-in policies. The number and type of user-defined policies that a feature requires differs for each feature based on implementation requirements.

A NetScaler feature can use one of two policy types:

Classic policies

Classic policies evaluate basic characteristics of traffic and other data. For example, classic policies can identify whether an HTTP request or response contains a particular type of header or URL.

Default policies

Default policies can perform the same type of evaluations as classic policies. In addition, default policies enable you to analyze more data (for example, the body of an HTTP request) and to configure more operations in the policy rule (for example, transforming data in the body of a request into an HTTP header).

Basic Policy Components

Classic and default policy components include:

Name	Each policy has a unique name.
Rule	The rule is a logical expression that enables the NetScaler feature to evaluate a piece of traffic or another object. For example, a rule can enable the NetScaler system to determine whether an HTTP request originated from a particular IP address, or whether a Cache-Control header in an HTTP request has the value No-Cache. Default policies can use all of the expressions that are available in a classic policy, with the exception of classic expressions for the SSL VPN client. In addition, default policies enable you to configure more complex expressions.
Bindings	Bindings ensure that the NetScaler system can invoke a policy when it is needed. You can bind the policy to one or more bind points.
Action	Policy evaluation ultimately results in the NetScaler system performing an action. For example, an integrated cache policy can identify HTTP requests for GIF or JPEG files. An action that you associate with this policy determines that the responses to these types of requests are served from the cache.

Policy Priorities

Lower-numbered policies have higher priority and are evaluated sooner. For example, if you have three policies with priorities of 10, 100, and 1000, the policy assigned a priority of 10 is performed first. Setting the appropriate priority is key to a successful policy; the rewrite policy evaluates multiple policies in order of priority.



Citrix recommends that you allow enough space to add policies by setting priorities with intervals of 50 (or 100) between each policy.

Priority of the Policy

NetScaler Gateway policies use a numeric order to determine which policy should be evaluated first. The policy priority is determined using the following order:

- Bind point of the policy
 - User (highest priority)
 - Group
 - Virtual Server
 - Global (lowest priority)
- Numeric ordering of the priority

Policies assigned a lower number have a higher priority than those with higher numbers. For example, a policy with a priority of 10 has a higher priority than a policy with the value of 100.

If a global policy has a priority of 20 and a user policy has a priority of 30 then the global policy is determined to be the higher priority policy.



When setting priorities for policies, ensure that the numbering is consistent where possible. If a single global policy is to be bound it can be set with a priority of 100, then any other policies should be set lower than 100 to be invoked.

HTTP Header Settings

The NetScaler system can alter the HTTP header settings as it mediates communication with the client, allowing for optimization of the TCP connection.

For example, a NetScaler system can suppress the connection, close the connection, and maintain a connection to the client while multiplexing many client connections to the server on the back-end. This behavior improves performance for both the client and the server. For example, you can modify HTTP data to redirect a request to a new home page or a new server, based on the address of the incoming request. You can also modify the data to mask server information in a response for security purposes. The URL Transformer function identifies URLs in HTTP transactions and text files for the purpose of evaluating whether a URL should be transformed.

The NetScaler system can also perform intelligent caching of data based on specific traffic conditions and parameters.

HTTP Request Headers

HTTP request header items include:

GET/HTTP/1.1	Defines the client method and the version of HTTP that is being used
Host	Defines the server being contacted
User-Agent	Defines information about the client software making the request
Accept, Accept-Language, Accept-Encoding, Accept-Charset	Define the types of file content, languages, compression algorithms and character sets that the client will accept from the server
Keep-Alive	Defines the time the session should be kept open
Connection	Defines a request to keep the session open even after the data is sent

The following is an example of an HTTP request header:

```
GET / HTTP/1.1
Host: www.msn.com
User-
Agent: Mozilla/5.0 (Windows NT 6.1; rv:20.0) Gecko/20100101
Firefox/20.0
Accept:
text/html,application/xhtml+xml,application/xml;q=0.9,*/*;q=0.8
Accept-Language: en-US,en;q=0.5
Accept-Encoding: gzip, deflate
Connection: keep-alive
```

HTTP Response Headers

HTTP response header items include:

HTTP/1.1 200 OK	Is the status of the request. Common response codes include 200 = OK; 302 = content moved; and 404 = file not found.
Connection	Is the reply of the server to the client request to keep the connection open
Set-Cookie	Sets the cookie values on the client so that the connection state can be maintained
Cache-Control	Instructs the client and any systems passing data to the client not to cache the page data
Pragma	Instructs the client and any systems passing data to the client not to cache the page data
Expires	Gives the date and time after which the response is expired

The following is an example of an HTTP response header:

```
HTTP/1.1 200 OK
Cache-Control: no-cache, no-store
Pragma: no-cache
Content-Type: text/html; charset=utf-8
Content-Encoding: gzip
Vary: Accept-Encoding
P3P: CP="NON UNI COM NAV STA LOC CURa DEVa PSAa PSDa OUR IND"
Set-
Cookie: MC1=V=3&GUID=a11d88f204844638b00c493eb2d2bc1f;
domain=.msn.com;
expires=Sat, 16-May-2015 14:41:56 GMT; path=/
Set-
Cookie: MC1=V=3&GUID=a11d88f204844638b00c493eb2d2bc1f;
domain=.msn.com;
expires=Sat, 16-May-2015 14:41:56 GMT; path=/
Set-Cookie: brdSample=0; domain=.msn.com; expires=Sat, 16-May-
2015 14:41:56 GMT; path=/
Set-Cookie: mh=MSFT; domain=.msn.com; expires=Sat, 16-May-
2015 14:41:56 GMT; path=/
Set-Cookie: CC=US; domain=.msn.com; expires=Sat, 16-May-
2015 14:41:56 GMT; path=/
Set-Cookie: CULTURE=EN-US; domain=.msn.com; expires=Sat, 16-May-
2015 14:41:56 GMT; path=/
Set-Cookie: _FS=NU=1; domain=.msn.com; path=/
Set-
Cookie: _SS=SID=587531AADAF4425EBA4265247AAA7ED8;
domain=.msn.com; path=/
Set-Cookie: MUIDB=3C453DB039BE6C570D51397F38A16C37;
expires=Sat, 16-May-2015 14:41:56 GMT; path=/
Set-Cookie: SRCHD=D=2826161&MS=2826161&AF=NOFORM;
expires=Sat, 16-May-2015 14:41:56 GMT; domain=.msn.com; path=/
Set-Cookie: SRCHUID=V=2&GUID=72BC915B4D7D43CE84EF7255BFF82EF9;
expires=Sat, 16-May-2015 14:41:56 GMT; path=/
Set-Cookie: SRCHUSR=AUTOREDİR=0&GEOVAR=&DOB=20130516;
expires=Sat, 16-May-2015 14:41:56 GMT; domain=.msn.com; path=/
errorCodeCount: [0:0]
X-AspNet-Version: 4.0.30319
S: CH1SCH060101130
Edge-control: no-store
Date: Thu, 16 May 2013 14:41:56 GMT
Content-Length: 42984
```



The Pragma field is backward compatible to HTTP/1.0, where the Cache-Control header is defined in HTTP/1.1.

Discussion Question

You created a policy but it does not seem to be working. You used the policy evaluator and it seems to be valid. What other things should you look for?

Policies and Profiles Configuration

NetScaler Gateway allows the use of policies to configure many aspects of the gateway configuration. By using policies, an administrator can choose where and when the policy should apply. For example, you might want users of a particular NetScaler Gateway vServer to be restricted to only being able to log on if their device is running Sophos antivirus version 6; however, you do not want any other NetScaler Gateway vServer impacted by this requirement. This can be achieved through the use of pre-authentication policies and policy bindings.

A policy is used to implement or manage settings when a specified set of conditions are met. For example, it might be preferable to permit use of the NetScaler Gateway plug-in use only for a certain Active Directory group. By creating a session policy and binding it to the relevant group, this configuration can be achieved.

A policy consists of the following components:

- **Expression**
 - One or more conditions
- **Profile**
 - A collection of settings
- **Priority**
 - An order of evaluation and enforcement



Although several types of policies are available to an administrator, priority and expressions are common to all policy types.

Authentication is commonly configured using authentication policies, then bound to the appropriate NetScaler Gateway virtual server. In a scenario where two-factor authentication is required, you can create two authentication policies and bind them to the virtual server.

Policies

Policies use expressions to determine if the policy should apply. If the policy is to be applied, then a collection of settings is invoked. The collection of settings is called a profile. A profile can contain many settings that should be used if the policy is matched; however a policy can only contain one profile.

Bind Points

Policies for NetScaler Gateway be bound in multiple places:

- Users
- Groups
- Virtual Servers
- Globally

By carefully selecting the bind point, you can create a flexible, secure, and low-maintenance NetScaler Gateway environment. For example, you might bind common settings globally and then provide a specific setting in an additional policy that is bound to either a user, group, or virtual server.

Citrix SmartAccess

SmartAccess allows you to create XenApp or XenDesktop policies that will apply based upon the result of a NetScaler Gateway scan. This provides increased security and can help an administrator to provide restricted access where appropriate, without impacting all users.

For example, you might want to disable local printer mapping for all end users; however, end users who have invoked a certain session policy will have mapped printers. When using NetScaler Gateway Enterprise Edition, the following are substituted:

- NetScaler Gateway farm name replaces the name of the NetScaler Gateway virtual server
- Filter name replaces the name of the pre-authentication or session policy



SmartAccess requires a XenApp or XenDesktop administrator to provide the name of the NetScaler Gateway farm and the name of the filter.

It is highly recommended to use a simple naming format for policies; this helps to eliminate common configuration mistakes with SmartAccess.

Types of Policies

NetScaler Gateway provides several types of policies. A brief overview of these is provided in the following table:

Type of policy	Explanation
Pre-authentication	Determines if a device is allowed to attempt to log on

Type of policy	Explanation
Authorization	Specifies the network resources to which end users have access; for example, permitting traffic to the 10.0.0.0 subnet
Session	Defines settings that apply to a user's NetScaler Gateway session; for example, allowing the use of the NetScaler Gateway plug-in or enforcing ICA-proxy-only access or defining the amount of time a user can stay logged on.
Traffic	Configures TCP, HTTP, and File Type Association settings
TCP Compression	Enables compression of TCP data where possible
Clientless Access	Defines settings for Clientless Access
Authentication	Provides the ability to integrate with authentication services to authenticate user logons
Auditing	Specifies how auditing of events is handled



Session policies allow you to provide session-based settings easily while allowing different uses based on the access scenario. Through the use of bind points, you can choose when to run a session policy. For example, you might want all users to have access to use the NetScaler Gateway plug-in; however a group of contractors should only be given access to clientless access. You can achieve this by creating session policies and binding them at different levels.

Policy Structure

Policies evaluate the traffic that flows through the NetScaler Gateway and then take some action based on the result of an evaluation. The evaluation is a boolean condition, and the action defines the configuration that goes into effect when the condition is met.

Every policy comprises three components:

- Expression
- Profile
- Priority

Policy Expressions

Policy expressions evaluate a condition of the connection. Some common NetScaler Gateway evaluation criteria are:

- Client operating system
- Antivirus software version and virus definition
- Connection protocol
- URL string contents
- Source and destination IP address and port
- Client SSL certificates

Additionally, the `ns_true` expression may be used to apply the policy every time NetScaler Gateway evaluates it.

Policy Profiles

Policy profiles specify the actions that the NetScaler Gateway takes if the expression is met. Different type of policies have different possible profile settings. Profile settings include:

- Authentication type
- Authentication server
- Client type
- Home page
- Time-out settings
- Network access
- Protocol
- Single sign-on
- File type association
- Compression type

The same profile may be used as part of multiple policies.

Policy Priority

Policy priority specifies which policy the NetScaler Gateway enforces if multiple conflicting policies are in place. The priority of a policy depends on both:

- The bind point
- The priority, which is set after binding

Configuring Conditional Policies

When configuring policies, you can use any Boolean expression to express the condition for when the policy applies. When you configure conditional policies, you can use any of the available system expressions, such as the following:

- Client security strings
- Network information
- HTTP headers and cookies
- Time of day
- Client certificate values

You can also create policies to apply only when the end-user device or endpoint meets specific criteria, such as a session policy for SmartAccess.

Another example of configuring a conditional policy is varying the authentication requirements for end users. As an example, you can create a policy that requires LDAP and RADIUS authentication for end users who are connecting with the NetScaler Gateway plug-in from outside the internal network, such as from their home computer, and another policy that requires LDAP authentication for end users who are connecting through a wide area network (WAN).



You cannot use policy conditions based on endpoint analysis results if the policy rule is configured as part of security settings in a session profile.

Conditional Policies Configuration

You can specify when a policy match occurs through the use of policy expressions. NetScaler Gateway contains several built-in expressions but also has flexibility to allow for custom conditions.

Some examples of possible expressions include:

- Checking to see if a certain process is running
- Scanning for the presence of a particular file
- Determining if a service is running
- Examining HTTP Request headers
- Checking for client certificates
- Ensuring that antivirus has recently updated definitions

For example, you may want to only allow logon if an end user's device or endpoint has the Sophos Antivirus process SavService.exe running, otherwise the device should not be allowed to log on. This can be configured using the following expression in a pre-authentication policy:

```
CLIENT.APPLICATION.PROCESS(SavService.exe) EXISTS
```



Conditional policies are useful to ensure that profiles are applied only when certain criteria are met. This functionality can also be greatly increased by using SmartAccess capabilities.

System Expressions Configuration

Through the use of system expressions you can specify the conditions under which a policy is enforced. If an expression is matched then the selected profile can be invoked for the connection.

For example, you might want to ensure that connecting users can only use the NetScaler Gateway plug-in if they are running the Sophos.exe process, otherwise users are forced to use ICA proxy mode. To achieve this setup, a possible configuration is:

1. Configure the ICA Proxy settings globally.
2. Create a new session policy.
3. Create a new request profile.
4. Configure the use of the NetScaler Gateway Plug-in and disabled ICA Proxy.
5. Add a client security using the following settings:
 - a. Component = Process
 - b. Name = Sophos.exe
 - c. Operator = EXISTS
6. Bind the policy to the NetScaler Gateway virtual server, ensuring that the priority value is a lower number than other policies.

Expression Types

NetScaler Gateway provides several types of expressions that can be used within policies. The following table provides an overview of the expression types and possible uses for the expressions:

Expression Type	Description
General	<p>Provides analysis of different components of client connections, such as:</p> <ul style="list-style-type: none">• HTTP Requests - including headers and URL requests• Client Certificates - Certificate is present, issuer name, subject• TCP - Source or destination port• IP - Source or destination IP

Expression Type	Description
Client Security	<p>Examines the user device to determine if the device meets certain device specific criteria, for example:</p> <ul style="list-style-type: none"> • Antivirus, personal firewall or Internet security, anti-spam: Are these products installed and are their definitions up to date? • Process: Searches for the presence of a process executable • Service: Checks if a Windows service is active • File: Confirms if a file is present within the file system • Registry: Examines registry keys and values • Operating System: Confirms the operating system version and service-pack level
Network-based	<p>Performs network analysis on either requests or responses, including source or destination MAC address, VLAN ID, or throughput</p>
Date and time	<p>Allows you to specify the following, with user-friendly calendar:</p> <ul style="list-style-type: none"> • Time: Hours, minutes and seconds • Date: Day, month and year • Day of week: Any day
File System	<p>Examine a file system to determine file parameters including:</p> <ul style="list-style-type: none"> • Time created • Size • Time modified

Expressions

NetScaler Gateway supports both simple and compound expressions. Simple expressions consist of a single comparison, for example determining whether a specific process is running on a user device.

Default expressions use existing simple and compound expression combinations. The expressions use either the AND operator or the OR operator to then determine if the compound expression is true or false.

Expression Structures

An expression defines the types of requests and associated responses to which the NetScaler system applies an action.

Inline and named expressions are used in both classic and default policies. Inline expressions are part of a policy and cannot be reused by other policies. Named expressions are a common pool of logical statements applied to the content entering the NetScaler system. Named compound expressions are independent entities and can be reused by other policies.



The NetScaler system uses a proprietary language based on Perl Compatible Regular Expression (PCRE) format that adds special terms, allowing aspects of a connection to be designated at a granular level.

Qualifiers, Operators, and Expression Values

Qualifiers specify what the policy examines. Operators determine how the qualifier will be evaluated. A qualifier is compared with the expression value, which can be literal text, a substring of text, or a numerical value.

Simple Expressions

The simple expression is the basic building block of policies. A simple expression consists of a single logical comparison, such as SOURCEIP = 10.12.120.12. In this example, the qualifier is the source IP address of the traffic and the operator is the 'equal to' sign.

The following expression matches traffic with a source IP address of 65.219.20.0 and a netmask of 255.255.255.0:

```
add pol exp bad_site "REQ.IP.SOURCEIP == 65.219.20.0 -  
netmask 255.255.255.0"
```

The expression string is marked by quotation marks in the command-line interface and the ns.conf configuration file. The following expression matches requests for specific content in the URL:

```
add pol exp big_java "REQ.HTTP.URL contains big_job.jsp"
```

The following expression matches requests with URL lengths of more than 256 characters:

```
add pol exp big_url "REQ.HTTP.URLLEN > 256"
```


The following expressions match requests based on the HTTP request headers:

```
add pol exp accepts_html "REQ.HTTP.HEADER accept contains
text/html"
```

```
add pol exp has_cookie_header "REQ.HTTP.HEADER cookie exists"
```

Compound Expressions

Compound expressions can be made up of logical combinations of existing simple and compound expressions. Compound expressions use Boolean logic to compare the results of the separate expressions and determine if the compound expression evaluates to true or false.

The logical operators that can be used with compound expressions are:

- AND, represented by the double ampersand (&&)
- OR, represented by the double pipe (||)

In the NetScaler system, compound expressions are created from the same interface as simple expressions within the configuration utility.

Unlike simple expressions, compound expressions cannot be modified once they are created because of the risk of creating an expression loop.

The following compound expressions are built from simple expressions:

```
add pol exp big_bad "(bad_site && big_java) "
```

```
add pol exp big_bad_or_long "(big_bad || big_url) "
```

Compound expressions cannot contain any qualifier other than COMPOUND. The COMPOUND qualifier does not need to be explicitly used. By referencing other expressions in the add expression statement, the statement assumes the use of COMPOUND== before the expressions. For example, the big_bad expression would be listed in the expression table as:

```
add policy expression big_bad COMPOUND== (bad_boys && big_java)
```



Once a compound expression has been bound, it can not be edited. It must be deleted and completely rewritten.

Custom Expressions

If you are creating a policy, you can create a custom expression while configuring the policy to determine when the policy should apply. For example, you are creating a session profile to allow

users to log on with the NetScaler Gateway Plug-in, set a time limit for the session, and allow single sign-on with Windows. After you create the session profile, in the Create Session Policy dialog box, you can create the expression. The following example shows an expression that checks that an antivirus application is running, and is a particular version:

```
CLIENT.APPLICATION.PROCESS(ccapp.exe)EXISTS -frequent 5 &&  
CLIENT.APPLICATION.AV(Symantec).VERSION==14.20.0.29 -freshness 5
```

Context-Sensitive Fields

The configuration utility displays context-sensitive fields depending upon the qualifier and operator selected.

If the HEADER qualifier is selected, type the appropriate header in the Header Name field. If the CONTENTS operator is selected, then type the length and offset in the Length and Offset fields.

The *length* and *offset* parameters are used to determine which part of a text string is examined for matching the operator and value. The *length* parameter controls the number of characters used and the *offset* parameter defines how far from the beginning of the field the capture begins. For example, if a string of "MSAccess" is examined by an expression with a length of 6 and an offset of 2, "Access" is the parsed string sent for evaluation.

Wildcards

The * wildcard character can be used to match a string within the specified qualifier. This character can appear only once within a string. You can restrict the processing of a string with wildcard characters. For example, the string /*gif will match on the first instance of 'gif' but not on further instances of 'gif,' if there is more than one in the string. It is important for strings to be matched properly for rule-based persistence.

Available Operators

The contents of the operator and value fields change depending on the qualifier selected. The following list describes available operators:

Operators	Description
==, EQ	Tests for URLs that are case sensitive for exact value text string matches
!=, NEQ	Tests for URLs that are case sensitive for items that do not match the exact value text string

Operators	Description
>,GT	Tests for URLs and query strings with lengths that are greater than the value integer. This operator is used with qualifiers such as URLLEN and QUERYLEN.
<, LT	Tests for URLs and query strings with lengths that are less than the value integer. This operator is used with qualifiers such as URLLEN and QUERYLEN.
>=,GE	Tests for URLs and query strings with lengths that are greater than or equal to the value integer. This operator is used with qualifiers such as URLLEN and QUERYLEN.
<=, LE	Tests for URLs and query strings with lengths that are less than or equal to the value integer. This operator is used with qualifiers such as URLLEN and QUERYLEN.
CONTAINS	Tests against the specified qualifier to determine if the specified string is contained in the qualifier. This operator is not case sensitive.
NOTCONTAINS	Tests against the specified qualifier to determine if the specified string is not contained in the qualifier. This operator is not case sensitive.
EXISTS	Tests for the existence of a particular qualifier
NOTEXISTS	Tests if the particular qualifier does not exist
CONTENTS	Tests if the qualifier exists and if it has contents

Available Qualifiers for HTTP Traffic

Qualifiers for HTTP traffic include:

Qualifier	Description
METHOD	Refers to the HTTP method used in the request, usually GET or POST, although all HTTP/1.1 standard headers are accepted for expressions
URL	Refers to the contents of the URL in an HTTP header, not including the query string
URLLEN	Refers to the length of the URL header contents, including the query string
URLQUERY	Refers to the query portion of the URL header contents
URLQUERYLEN	Refers to the length of the query portion of the URL header contents
URLTOKENS	Refers to special tokens in the URL
VERSION	Refers to the HTTP request version in the form of HTTP/x.x, in which x is an integer
HEADER	Refers to the header portion of the HTTP request by name

Available Qualifiers for Non-HTTP Traffic

Qualifiers for non-HTTP traffic include:

Qualifier	Description
DESTIP	Specifies the destination IP address or network
DESTPORT	Specifies the destination service port
SOURCEIP	Specifies the source IP address or network
SOURCEPORT	Specifies the source service port
LOCATION	Specifies the location value as defined in the GSLB database

Identifying Operands, Wildcards, and Context-Sensitive Fields

Operands and other variables are used to modify expressions. The value used for the operator depends on the qualifier and operator chosen. It can be literal text, a substring of text, or a numeric value. The variables and operators include:

- * wildcard for subset string matching
- Operands
- Context sensitive fields

Wildcards

The * wildcard character can be used to match a string within the specified qualifier. This character can appear only once within a string. By using wildcard characters, the administrator can restrict processing of a string. For example, the string "/*.gif" will match on the first instance of .gif, but not further instances of .gif if there is more than one .gif in the string. This can be of particular importance when using rule-based persistence, so it is important to carefully craft the strings that are to be matched.

Values or Operands

The value field specifies the object with which the qualifier is compared. The value may be either a string or an integer.

The -length and -offset parameters are used to determine which part of a text string is examined for matching the operator and value. The -length parameter controls the number of characters used, and the -offset parameter defines how far from the beginning of the field the capture begins. For example, if a string of "MSAccess" is examined by an expression with a length of 6 and an offset of 2, "Access" is the resulting parsed string sent for evaluation.

Context-Sensitive Fields

When creating an expression in the configuration utility, the choices available in the drop-down list box are context-sensitive, which means that they change to match the protocol chosen by the administrator. If the Header qualifier is selected, type the appropriate header in the Header Name field. If the contents operator is selected, type the length and offset in the appropriate fields.

Policies Creation

Creating policies within NetScaler Gateway can be achieved using either the NetScaler Gateway policy manager or the configuration utility.

The NetScaler Gateway policy manager provides a simplified interface to access all NetScaler Gateway items and policies, providing an interface to create, bind, and remove policies.

Once a policy has been created it must then be bound in order for it to become active. Policies can be bound at the following bind points:

- Users
- Groups
- Virtual Server
- Globally

Accessing Policy Manager

To access the NetScaler Gateway Policy Manager, use the following procedure:

1. Within the configuration node, select the **NetScaler Gateway** node.
2. Under the **policy manager** pane, select **NetScaler Gateway Policy Manager**.

Creating Policies in the Configuration Utility

To create a policy within the configuration utility, the following procedure can be used:

1. Within the configuration node, select the **NetScaler Gateway** node.
2. Expand the **Policies** node.
3. Select the relevant policy that you want to configure (for example, Pre-Authentication).
4. Click the **Add** button to create a policy.



Using the Policy Manager, you can access an overview of policies and objects from a single management location.

Configuring Session Policies

A session policy is a collection of expressions and settings that are applied to end users, groups, virtual servers, and globally.

Session policies define when session profiles should be applied. Session Profiles are used to configure the settings for end user connections. You can define settings to configure the software

users log on with, such as the NetScaler Gateway Plug-in for Windows or the NetScaler Gateway Plug-in for Java. Session policies are evaluated and applied after the user is authenticated.

Session policies are applied according to the following rules:

- Session policies always override global settings in the configuration.
- Any attributes or parameters that are not set using a session policy are set on policies established for the virtual server.
- Any other attributes that are not set by a session policy or by the virtual server are set by the global configuration.

If you deploy XenMobile App Edition or StoreFront in your network, Citrix recommends using the Quick Configuration wizard to configure session policies and profiles. When you run the wizard, you define the settings for your deployment.

To Create a Session Policy

1. In the configuration utility, on the **Configuration** tab, in the navigation pane, expand **NetScaler Gateway > Policies** and then click **Session**.
2. In the details pane, under **Policy Manager**, click **Change group settings and user permissions**.
3. In the Policy Manager, under **Available Policies / Resources**, click **Session Policies**.
4. Under **Related Tasks**, click **Create new session policy**.
5. In **Name**, type a name for the policy.
6. Next to **Request Profile**, click **New**.
7. In **Name**, type a name for the profile.
8. Complete the settings for the session profile and then click **Create**.
9. In the Create Session Profile dialog box, add an expression for the policy, click **Create**, and then click **Close**.



In the expression, you can select **NS_TRUE** so that the policy is always applied to the level to which it is bound.

To Bind a Session Policy by Using the Policy Manager

1. In the configuration utility, click the **Configuration** tab and then in the navigation pane, click **NetScaler Gateway**.
2. In the details pane, under **Policy Manager**, click **NetScaler Gateway Policy Manager**.
3. In the NetScaler Gateway Policy Manager, under **Available Policies / Resources**, expand **Session Policies** and then click a policy.
4. Drag the session policy to the user, group, virtual server, or NetScaler Gateway Global session policy under **Configured Policies / Resources**.



Session policies are applied as a hierarchy in the following order:

- Users
- Groups
- Virtual server
- Globally

Traffic Policies

You can configure traffic policies to define how traffic that meets a specified criteria is controlled by NetScaler Gateway. For example, you might want to enforce a shorter timeout value for certain sensitive applications when they are accessed from an untrusted network.

Another use of traffic policies is to enable or disable file type association. This is a feature in XenApp which allows an administrator to configure certain application extensions to open with published applications rather than with local resources.

A traffic policy has a corresponding profile assigned to implement the traffic settings. The policy must also be bound to the appropriate level. This can be one of the following:

- User
- Group
- Virtual Server
- Globally

Traffic Policy Creation Example

A company may use an HR system that must have restricted access and be further secured to ensure that end users cannot stay logged on for more than 15 minutes. Using a traffic profile, you can specify the following to ensure that this is enforced:

- Protocol = HTTP
- AppTimeout = 15

Once a traffic profile is in place to determine the action, a policy must be created with an expression to determine when the profile should be used. For example, if the HR system was accessed using an IP address of 10.0.0.86, then the following expression could be used in the policy:

```
REQ.IP.DESTIP == 10.0.0.86
```


To Configure a Traffic Policy

1. In the configuration utility, in the navigation pane, click **NetScaler Gateway**.
2. In the details pane, under **Policy Manager**, click **NetScaler Gateway Policy Manager**.
3. In the Policy Manager, under **Available Policies > Resources**, click **Traffic Policies**.
4. Under **Related Tasks**, click **Create new traffic policy**.
5. In **Name**, type a name for the policy.
6. Next to **Request Profile**, click **New**.
7. In **Name**, type a name for the profile.
8. In **Protocol**, select either **HTTP** or **TCP**.



If you select TCP as the protocol, you cannot configure single sign-on and the setting is disabled in the profile dialog box.

9. To limit the time end users can stay logged on to the Web application, in **AppTimeout (minutes)**, type the number of minutes.
10. To enable single sign-on to the Web application, in **Single Sign-On**, select **ON**.



If you want to use form-based single sign-on, you can configure the settings within the traffic profile. For more information, see [Configuring Form-Based Single Sign-On](#).

11. To specify a file type association, in **File Type Association**, select **ON**.
12. To use the Repeater Plug-in to optimize network traffic, in **Branch Repeater**, select **ON**, click **Create**, and then click **Close**.
13. In the **Create Traffic Policy** dialog box, create or add an expression, click **Create**, and then click **Close**.

To Bind a Traffic Policy by Using the NetScaler Gateway Policy Manager

1. In the configuration utility, click the **Configuration** tab and then in the navigation pane, click **NetScaler Gateway**.
2. In the details pane, under Policy Manager, click **Change group settings and user permissions**.
3. Under Available Policies / Resources, expand **Traffic Policies** and then click a traffic policy.
4. Drag the policy to Traffic Policies under Configured Policies / Resources for the level at which you want the policy bound.

To Bind a Traffic Policy Globally by Using the Configuration Utility

1. In the configuration utility, on the Configuration tab, in the navigation pane, expand **NetScaler Gateway > Policies** and then click **Traffic**.
2. In the details pane, select a policy and then in Action, click **Global Bindings**.
3. In the Bind / Unbind Traffic Policies dialog box, under Details, click **Insert Policy**.
4. Under Policy Name, select the policy and then click **OK**.

To Unbind a Traffic Policy by Using the NetScaler Gateway Policy Manager

1. In the configuration utility, click the **Configuration** tab and then in the navigation pane, click **NetScaler Gateway**.
2. In the details pane, under Policy Manager, click **Change group settings and user permissions**.
3. Under Configured Policies / Resources, expand the node that has the traffic policy bound to it, expand **Traffic Policies** and then click the traffic policy.
4. Under Related Tasks, click **Unbind traffic policy**.

After the traffic policy is unbound, you can remove the policy.

To Remove a Traffic Policy by Using the NetScaler Gateway Policy Manager

1. In the configuration utility, click the **Configuration** tab and then in the navigation pane, click **NetScaler Gateway**.
2. In the details pane, under Policy Manager, click **Change group settings and user permissions**.
3. Under Available Policies / Resources, expand **Traffic Policies**, and then select the traffic policy.
4. Under Related Tasks, click **Remove traffic policy**, and then click **Yes**.

Discussion Question

What are some of the ways that traffic policies are used?

Module 12

Multi-Tenancy and NetScaler SDX

12

Not for resale or distribution

Not for resale or distribution

Multi-Tenancy and NetScaler SDX

Multi-Tenancy with NetScaler

Multi-Tenancy with NetScaler occurs when a single NetScaler, or a single HA pair of NetScalers, serve multiple tenants. In many cases, the single NetScaler appliance is a NetScaler SDX appliance running multiple instances. The individual instances may be dedicated to a single tenant, or multiple tenants may share a NetScaler instance. A tenant is an isolated group of users with common access and privileges to resources. Usually, the isolation requirements of the users or tenants will determine what model of tenancy to use.

A tenant is an isolated group of users with common access and privileges to resources. Usually a tenant represents a company, or a division within a company.

- Private tenants have dedicated resources which provide for a higher degree of separation. They are typically more expensive to implement and maintain, and are appropriate for companies with higher security or performance requirements.
- Shared tenants are served from the same resources. Separation between shared tenants may not be required. If separation is required, it is maintained at a software level.

Multi-Tenancy Considerations

Some of the advantages of multi-tenancy include:

- It is typically more cost effective than implementing and maintaining multiple single tenants.
- It is useful for companies which rely upon network trafficking or data mining.
- Maintenance and updates are usually simpler.

Some of the difficulties associated with multi-tenancy include:

- It is typically more difficult to design and implement because multi-tenancy environments are usually more complex than single tenet environments.
- Extra consideration is often required to maintain a separation of client traffic and meta-data.

Introduction to the Different Kinds of Tenant Isolation

Tenant isolation requirements often dictate the attributes of a multi-tenant environment. Given this, it is important to understand some of the ways we can isolate tenants. Types of isolation include:

- **Performance Isolation** - When a tenant's consumption of resources does not impact other tenants.
- **Traffic/Data Isolation** - When a tenant's network traffic and/or data is separate from other tenants.

- **Fault Isolation** - When a service shutdown or failure in one tenant does not impact other tenants.
- **Administrative Isolation** - The extent to which management functions for different tenants can be separated and delegated.
- **Functional Isolation** - The ability of tenants to use different programs, operating systems, and firmware versions.

Tenancy Options with NetScaler

Citrix NetScaler can attain multiple types of isolation enforced at 3 different levels. Device level separation entails using separate physical devices for tenants:

- Usually offers the best performance and most separation.
- Usually the most expensive to implement.
- Administration is often more difficult.

Instance level separation entails using a separate VPX instance running on a generic hypervisor such as Citrix XenServer, or running on a NetScaler SDX appliance:

- Achieves a high level of isolation.
- Much more cost effective than device level separation.

Software level separation entails multiple tenants running on the same NetScaler instance which may be physical or virtual. If separation is required, it maintained with mechanisms such as :

- Most cost effective.
- Usually simplest to manage.
- NetScaler provides services such as Admin Partitions, Traffic Domains, and Subject Name Indication (SNI).
- Also attained at authentication service level.

NetScaler Admin Partitions and Traffic Domains

An admin partition is a logical entity on a NetScaler VPX or MPX appliance. Each admin partition can contain a separate configuration including separate administrative delegation. For these reasons, admin partitions are suited to provide software level separation between tenants. By default, a NetScaler contains a single admin partition known as the Default Partition. Other partitions are identified by the name assigned to them. Admin partitions support most NetScaler features which has expanded greatly since their introduction. Citrix will continue to expand the supported features with admin partitions. A complete overview on admin partitions is available at <http://docs.citrix.com/en-us/netscaler/11/system/admin-partition.html>

A traffic domain is a NetScaler entity which segments the network traffic of different applications. Application entities such as servers are assigned to a traffic domain and can only communicate with other entities within their own traffic domain. Because traffic domains provide complete separation

duplicate entities and even addresses can be used. With NetScaler traffic domains, tenants can completely separate their application traffic while residing on a single NetScaler. Much like admin partitions, NetScalers are preconfigured with a traffic domain known as the default traffic domain with an ID of 0. Other traffic domains are identified by the number assigned to them. Traffic domains support most NetScaler features with more being added often. More information on traffic domains is available at <http://docs.citrix.com/en-us/netscaler/11/networking/traffic-domains.html>.

Introduction to the NetScaler SDX Appliance

The Citrix NetScaler SDX appliance is a multi-tenant platform on which you can provision and manage multiple virtual machines (instances). The SDX appliance addresses cloud computing and multi-tenancy requirements by allowing a single administrator to configure and manage the appliance and delegate the administration of each hosted instance to tenants. The SDX appliance enables you to provide each tenant the following benefits:

- One complete instance - Each instance has identical privileges.
- A completely isolated network - Traffic meant for a particular instance is sent only to that instance.

Each complete instance has the following privileges:

- Dedicated CPU and memory resources.
- A separate space for entities.
- The independence to run the release and build of the administrator's choice.
- Life-cycle independence.

The Citrix NetScaler SDX appliance provides a Management Service that is pre-provisioned on the appliance. The Management Service provides a user interface (HTTP and HTTPS modes) and an API to configure, manage and monitor the appliance, the Management Service and the instances. A Citrix self-signed certificate is pre-packaged for HTTPS support. Citrix recommends that HTTPS mode is used to access the Management Service user interface.

Product Benefits: True Multi-Tenancy

NetScaler SDX provides a platform to run multiple independent instances of key services to meet the distinct needs of individual business units, critical applications and service provider clients. Enterprise and service provider clients gain dedicated control over their delivery infrastructure, including services such as load balancing, security and application acceleration. Complete isolation of per-client traffic helps satisfy security and compliance mandates and eases operational administration through version control and life-cycle management.

Product Benefits: Flexible Licensing and Configuration

NetScaler SDX offers flexible licensing to meet both enterprise and service provider requirements. Depending on the platform, the base solution includes 2 or 5 independent NetScaler instances and accommodates a growth in scale of as many as 8 times the base number of instances. The popular Citrix Pay-As-You-Grow licensing program lets customers scale solution performance to meet future business needs while protecting their initial NetScaler investment.

Product Benefits: Simple Management

Unified provisioning, monitoring and management of multiple concurrent NetScaler instances through a single control plane streamlines multi-tenant operations. Not only is each NetScaler instance managed independently, but each NetScaler can run a different software version and support independent IP addressing schemes to preserve end-to-end isolation of application traffic between different clients.

Hardware Platforms

NetScaler SDX is available in a variety of models to suit the most demanding IT and business needs. If additional throughput is needed, Burst Pack and Pay-as-You-Grow options help protect the initial investment and make it easier to increase the scale of a network with a simple software license upgrade. Instance packs (also applied through a simple software license) further protect the investment by entitling the deployment of additional NetScaler instances with licenses immediately and without affecting production ADC services. For more details on models and specifications, view the NetScaler SDX data sheet.

Model Performance Comparison

Model	HTTP Throughput	Memory	SSL Throughput	Included Instances
SDX 8015	15 Gbps	32 GB	6 Gbps	2
SDX 8200	4 Gbps	32 GB	3.5 Gbps	2
SDX 8600	6 Gbps	32 GB	5.5 Gbps	2
SDX 11515	15 Gbps	48 GB	14 Gbps	5
SDX 11520	20 Gbps	48 GB	15 Gbps	20
SDX 11530	30 Gbps	48 GB	17 Gbps	20

Model	HTTP Throughput	Memory	SSL Throughput	Included Instances
SDX 11540	40 Gbps	48 GB	19 Gbps	20
SDX 11542	42 Gbps	48 GB	20.5 Gbps	20
SDX 14020	20 Gbps	64 GB	21 Gbps	5
SDX 14030	30 Gbps	64 GB	23 Gbps	10
SDX 14040	40 Gbps	64 GB	34 Gbps	10
SDX 14060	60 Gbps	64 GB	40 Gbps	25
SDX 14080	80 Gbps	64 GB	43 Gbps	25
SDX 14100	100 Gbps	64 GB	46 Gbps	25
SDX 17550	20 Gbps	96 GB	8 Gbps	5
SDX 19550	30 Gbps	96 GB	9 Gbps	5
SDX 20550	40 Gbps	96 GB	9 Gbps	5
SDX 21550	50 Gbps	96 GB	11 Gbps	5
SDX 22040	40 Gbps	256 GB	35 Gbps	20
SDX 22060	60 Gbps	256 GB	45 Gbps	80
SDX 22080	80 Gbps	256 GB	55 Gbps	80
SDX 22100	100 Gbps	256 GB	65 Gbps	80
SDX 22120	120 Gbps	256 GB	35 Gbps	80
SDX 24100	100 Gbps	256 GB	40 Gbps	40
SDX 24150	150 Gbps	256 GB	44 Gbps	80

Deployment Scenarios: Datacenter Consolidation

Enterprise datacenters are migrating to virtual architectures that share resources across business units and individual applications. As part of the process, organizations are consolidating networking services and collapsing multiple network appliances into a single scalable platform that

can support the datacenter requirements. NetScaler SDX is the ideal platform to enable this service consolidation with the virtual datacenter while continuing to preserve the necessary traffic isolation, performance and service levels expected by application owners.

Deployment Scenarios: Multi-Tenancy

Virtual datacenters enable key services to be delivered efficiently to individual tenants such as internal business units, individual application owners and external service provider customers. NetScaler SDX satisfies multi-tenant requirements by:

- Running completely independent NetScaler instances with separate policies to deliver all NetScaler capabilities.
- Providing complete isolation of traffic among clients to meet compliance requirements.
- Supporting different NetScaler software versions to meet the life-cycle needs of each client.
- Maintaining separate IP addressing for easy deployment into virtual datacenters.

Deployment Scenarios: Tenant Isolation

NetScaler SDX traffic and data paths from unique clients remain segregated at the hardware layer, enabling consolidation benefits to span security zones without appliance sprawl. Hardware virtualization bypasses and replaces software-based traffic classification or sorting methods, enabling you to maintain layer 2 separation without the typical latency associated with virtualizing datacenter workloads. Independent processing stacks and routing segregate traffic and network information to support multiple isolated tenants on a single platform.

Deployment Scenarios: Enabling Cloud Infrastructure Expansion

Enterprises that use cloud infrastructures need the ability to provision networking capabilities on demand to meet dynamic application needs. As applications shift from the enterprise datacenter to a cloud environment, much of the application delivery functionality supporting the application must also migrate. NetScaler SDX distributes flexible NetScaler application delivery capabilities across multiple applications running in the cloud through a common control plane while providing the full reliability and performance of NetScaler hardware.

Licensing

Licensing the NetScaler SDX is no different from licensing a NetScaler MPX or VPX. However, the maximum number of NetScaler VPX instances is defined by the license and the platform that have been purchased.

The platform license entitles one base SDX appliance and five VPX instances on certain platforms by default. Five instance add-on licenses (Instance Pay-As-You-Grow) enables adding additional VPX instances, beyond the default five. Platform upgrade licenses (Platform Pay-As-You-Grow) upgrades to a higher throughput capacity on the same hardware platform. Platform conversion license changes MPX appliances to SDX appliances (Not applicable for FIPS, 9500, 7500, or 5500).

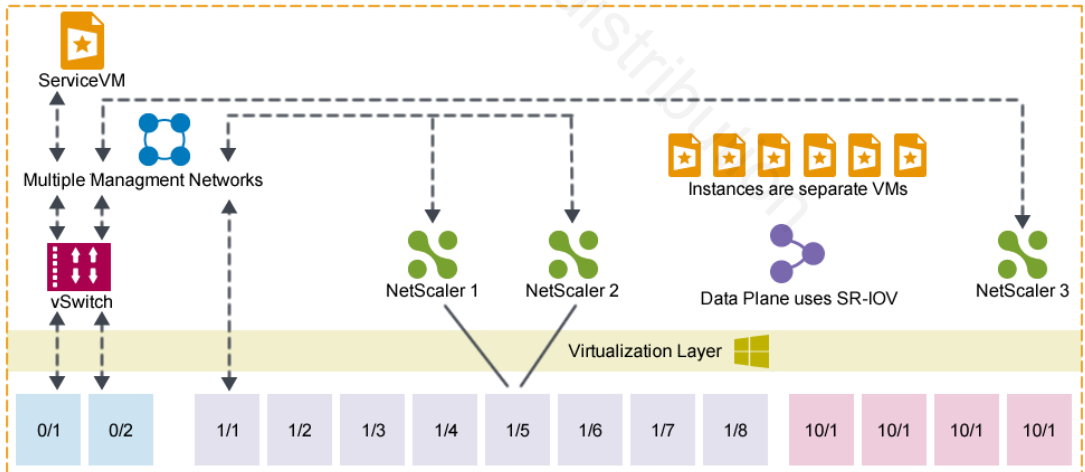
All of the VPX instances that are deployed on NetScaler SDX contain Platinum Edition licenses and functionality.

Example of a Typical Licensing Scenario:

- Purchased system: SDX 11500.
- Apply platform license.
 - As many as five VPX instances, maximum system throughput is 8 Gbps.
- Add three 5-pack instance licenses.
 - Add as many as 20 VPX instances, maximum system throughput is still 8 Gbps.
- Apply SDX-11500-to-SDX-18500 platform upgrade license.
 - Maximum system throughput increases to 36 Gbps.

Base Architecture

The NetScaler SDX has an underlying XenServer virtualization layer that provides the core functionality. The NetScaler SDX uses a vSwitch for the management plane network and that management plane network can comprise multiple networks, which is important for compliance and consolidation across security zones. However, the data plane does not allow access through the vSwitch. This is important for scalability, performance, multi-tenancy and isolation.



IO Virtualization

SR-IOV is a PCI standard that provides IO virtualization. With IO virtualization, a physical device or function, like a NIC, can be carved into virtual devices or functions. The virtual functions can be assigned to virtual machines. The virtual machine has direct access to hardware through a virtual function. IOMMU translates the physical addresses for the guest into host physical addresses. With IO virtualization, virtual machines can efficiently share the IO devices. The latest NICs from manufacturers like Intel support SR-IOV functionality.

With IO virtualization, each virtual function gets its own hardware receiving and transmitting queues and has direct access to the hardware. MAC and VLAN filters are associated with each virtual function. When the NIC receives a packet, two levels of filtering are applied:

- In the first phase, MAC filtering is applied to find the right virtual function based on the destination MAC address.
- VLAN filtering is applied to the packet later.

A packet is queued to a virtual function only if both MAC and VLAN filters pass. When a virtual function transmits a packet, it places the packet in the transmission queue and the hardware fetches the packet for actual transmission. There is no hypervisor involvement in the data path. Packet-switching is done at the hardware level, resulting in higher network performance. Hardware provides MAC and VLAN filtering capabilities to isolate the traffic across virtual machines. Using IO virtualization technologies, isolation is achieved without sacrificing performance.

VLAN Filtering

VLAN filtering provides segregation of data between NetScaler VPX instances that share a physical port. For example, if two NetScaler VPX instances are configured on two different VLANs and VLAN filtering is enabled, one instance cannot view traffic on the other instance. If VLAN filtering is disabled, all of the instances can see the tagged or untagged broadcast packets, but the packets are dropped at the software level. If VLAN filtering is enabled, each tagged broadcast packet reaches only the instance that belongs to the corresponding tagged VLAN. If none of the instances belong to the corresponding tagged VLAN, the packet is dropped at the hardware level (NIC).

If VLAN filtering is enabled on an interface, a limited number of tagged VLANs can be used on that interface (63 tagged VLANs on a 10G interface and 32 tagged VLANs on a 1G interface). A VPX instance receives only the packets that have the configured VLAN IDs. Restart the NetScaler VPX instances associated with an interface if the state of the VLAN filter is changed from DISABLED to ENABLED on that interface. VLAN filtering is enabled by default on the NetScaler SDX appliance. If you disable VLAN filtering on an interface, you can configure as many as 4096 VLANs on that particular interface.



VLAN filtering can be disabled only on a NetScaler SDX appliance running XenServer version 6.0.

Enabling VLAN Filtering on an Interface

1. On the **Configuration** tab, in the navigation pane, expand **System** and then click **Interfaces**.
2. In the Interfaces pane, click **VLAN Filter**.
3. In the Enable/Disable VLAN Filter dialog box, click **Add** to enable VLAN filtering on an interface.
4. Optionally, select **Reboot associated Instances**.
5. Click **OK**.

Restricting VLANs to Specific Virtual Interfaces

You can enforce specific 802.1Q VLANs on the virtual interfaces associated with NetScaler instances. This capability is especially helpful in restricting the usage of 802.1Q VLANs by the instance administrators. If two instances belonging to two different companies are hosted on an SDX appliance, the two companies can be restricted from using the same VLAN ID so that one company does not see the other company's traffic. If you, while provisioning or modifying a VPX instance, try to assign an interface to an 802.1Q VLAN, a validation is performed to verify that the VLAN ID specified is part of the allowed list.

By default, any VLAN ID can be used on an interface. To restrict the tagged VLANs on an interface, specify the VLAN IDs in the Network Settings at the time of provisioning a NetScaler instance, or later by modifying the instance. To specify a range, separate the IDs with a hyphen (for example, 10-12). If you initially specify some VLAN IDs but later delete all of them from the allowed list, you can use any VLAN ID on that interface. In effect, you have restored the default setting. After creating a list of allowed VLANs, you do not have to log on to an instance to create the VLANs. You can add and delete VLANs for specific instances from the Management Service.



If Layer 2 mode is enabled, the administrator must take care that the VLAN IDs on different NetScaler instances do not overlap.

NetScaler SDX High Availability

High availability is supported on the VPX instances that are provisioned on the SDX appliances. In an HA pair, failover can be configured to ensure that an instance on one appliance can fail over to another appliance without requiring a change to the entire device and every instance on the device. Embedded within this is the ability to have an active instance or active instances on both devices.

With an SDX appliance, the capability to upgrade an instance without upgrading the entire device and the capability to enable failover for an instance without failing over the entire device is available.

Service VM Overview

The NetScaler SDX service VM is a pre-provisioned FreeBSD (64-bit) virtual machine. It manages the whole appliance and as a result, the underlying XenServer never is exposed.

The SDX service VM has a management interface GUI and an accompanying API (similar to the NetScaler NITRO API).

Device Management

The SDX service VM shows system information such as number of available CPU cores, the total and available amounts of memory and the current firmware version information. The current VPX instance inventory is also displayed with the respective information for each of those systems.

Tasks that can be completed from the service VM graphical user interface include:

- Port administration (changes the interface speed and auto-negotiation settings).
- Management IP address assignment to the XenServer (only SSH is allowed) for service virtual machine failure conditions.
- File management from a local system to the service virtual machine.
- Event management.
- Task management.
- Auditing.
- Technical support for the service virtual machine and underlying XenServer.

Instance Management

The SDX Service VM also manages the individual VPX instances. From the Service VM, the following tasks can be completed:

- Start, stop, restart and remove VPX instances.
- Upgrade (single or multiple instances).
- View running vs. saved configurations.
- View current instance resource utilization.
- View audit messages.
- Managing users (without RBA).
- Save configuration on VPX instances.
- Add MIP and SNIP on VPX instances.

Provisioning a NetScaler VPX Instance on an SDX Appliance

You can provision one or more NetScaler instances on the SDX appliance by using the Management Service. The number of instances that you can install depends on the license you have purchased. If the number of instances added is equal to the number specified in the license, the Management Service does not allow provisioning more NetScaler instances.

To provision NetScaler instances on the SDX appliance:

1. Define an administrator profile to attach to the NetScaler instance. This profile specifies the user credentials that are used by the Management Service to provision the NetScaler instance and later to communicate with the instance to retrieve configuration data. You can also use the default administrator profile.
2. Upload the .xva image file to the Management Service.
3. Add NetScaler instances using the Management Service. The Management Service implicitly deploys the NetScaler instances on the SDX appliance and then downloads configuration details of the instances.



By default, an .xva image file based on the NetScaler 9.3 release is available on the SDX appliance.

Administrator Profiles

The user credentials specified in an administrator profile are also used by the client when logging on to the NetScaler instances through the command-line interface or the configuration utility.

The default administrator profile for an instance specifies a user name of nsroot and the password is also nsroot. This profile cannot be modified or deleted. However, you should override the default profile by creating a user-defined administrator profile and attaching it to the instance when you provision the instance. The Management Service administrator can delete a user-defined administrator profile if it is not attached to any NetScaler instance.



Do not change the password directly on the NetScaler VPX instance. If you do so, the instance becomes unreachable from the Management Service. To change a password, first create a new administrator profile and then modify the NetScaler instance, selecting this profile from the Administrator Profile list.

To change the password of NetScaler instances in a high-availability setup, first change the password on the instance designated as the secondary node and then change the password on the instance designated as the primary node. Remember to change the passwords only by using the Management Service.

XVA Template Repository

Multiple versions of NetScaler VPXs can be uploaded to the XVA repository on the SDX appliance from the service VM. Current supported versions are:

- NetScaler 9.3
- NetScaler 10
- NetScaler 10.1
- NetScaler 10.5
- NetScaler 11.0

You have to upload the NetScaler .xva files to the SDX appliance before provisioning the NetScaler instances. You can also download an .xva image file to a local computer as a backup. The .xva image file format is: NSVPX-XEN-ReleaseNumber-BuildNumber_nc.xva.

In the NetScaler XVA Files pane, you can view the following details:

- Name -- Name of the .xva image file. The file name contains the release and build number. For example, the file name NSVPX-XEN-11.0-63.16_nc.xva refers to release 11.0 build 63.16
- Last Modified -- Date when the .xva image file was last modified.
- Size -- Size, in megabytes (MB), of the .xva image file.

Uploading a NetScaler .xva File

1. On the **Configuration** tab, in the navigation pane, expand **NetScaler Configuration** and then click **XVA Files**.
2. In the NetScaler XVA Files pane, click **Upload**.
3. In the Upload NetScaler Instance XVA dialog box, click **Browse** and select the XVA image file that you want to upload.
4. Click **Upload**. The XVA image file appears in the NetScaler XVA Files pane after it is uploaded.

Provisioning a NetScaler VPX Instance on a NetScaler SDX

1. On the **Configuration** tab, in the navigation pane, expand **NetScaler Configuration** and then click **Instances**.
2. In the NetScaler Instances pane, click **Add**.
3. In the Provision NetScaler Wizard, follow the instructions on the screen.
4. Click **Create** and then click **Close**. The provisioning progress and any failures, such as failure to assign a virtual function to the VPX instance, are displayed.

To modify the parameters of a provisioned NetScaler instance: in the NetScaler Instances pane select the instance that you want to modify and then click **Modify**. In the Modify NetScaler Wizard, modify the parameters.



If you modify the number of SSL chips, interfaces, memory and feature license, the NetScaler instance implicitly stops and restarts to bring these parameters into effect.

You cannot modify the Image and User Name parameters.

If you want to remove a NetScaler instance provisioned on the SDX appliance, in the NetScaler Instances pane select the instance that you want to remove and then click **Delete**. In the Confirm message box, click **Yes** to remove the NetScaler instance.

NetScaler SDX Service VM Internals

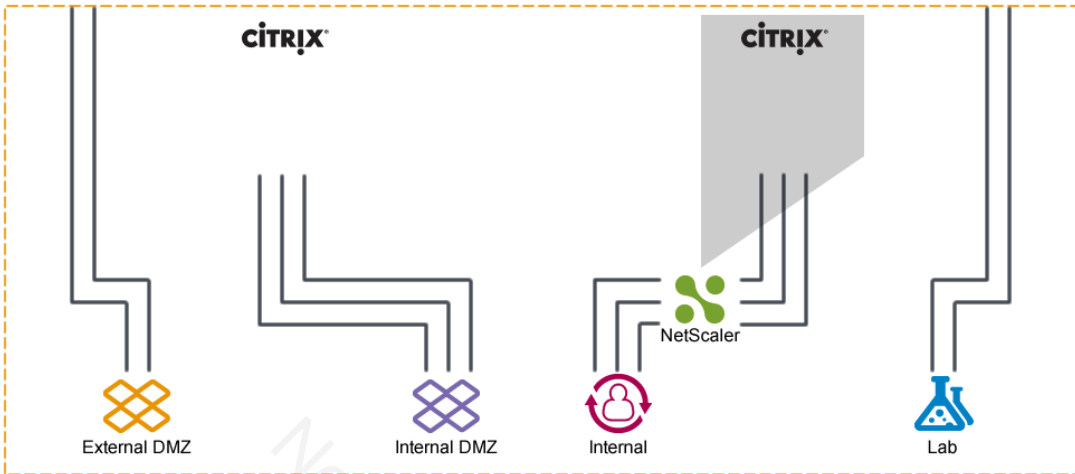
The Service VM sends API calls to the virtual machines for management tasks. There is no command-line interface for the Service VM. Memory usage is rated for each virtual machine and therefore for each dedicated core. Other monitoring screens are system-based aggregate usage.

There are some settings that must be configured before attempting to restart the SDX appliance (including XenServer):

- When the Service VM starts, it must be set to `auto_poweron` for itself through XenServer so that when XenServer is restarting, it can automatically start the service virtual machine.
- The Service VM on XenServer should contain the Service VM description to identify the service virtual machine in order to set `auto_poweron` on the virtual machine.

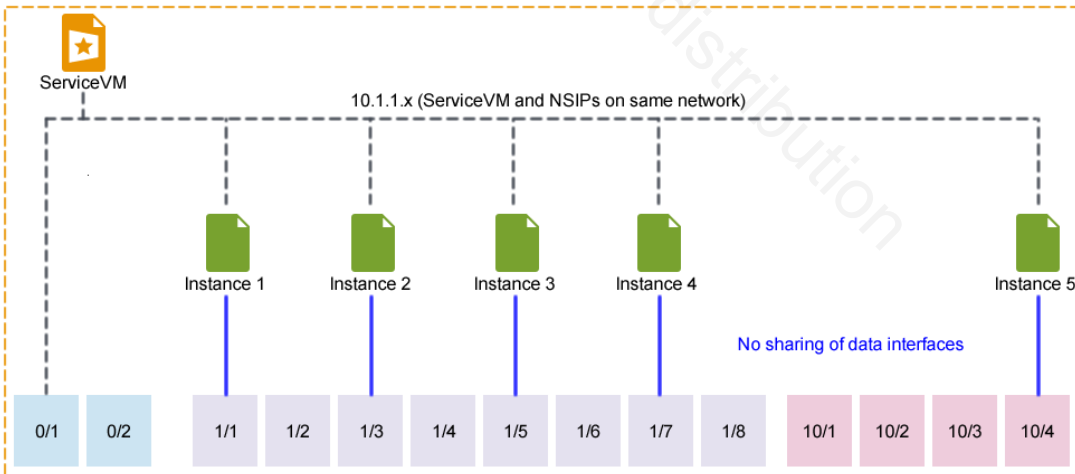
Simple Consolidation

Simple consolidation involves using a NetScaler SDX to consolidate and provide dedicated instances for a series of applications that all reside in the same security zone. The administrator for the device is also the administrator for each instance.



For example, if an administrator is supporting five different instances, with all of the instances in the same security zone and uses the same administrator for all instances, then place the Service VM and the NSIP/management interface for all instances on the same network. Therefore, a single management network on the device is acceptable.

For the data plane, one approach is to just give each instance its own dedicated physical interface or interfaces. Since the data plane traffic uses SR-IOV, this traffic does not travel through a central virtual switch; therefore the isolation is very strong. In this case, each instance can have any or all of the 4096 VLANs available (subject to how the network is configured). The data plane networks can also be completely different networks.



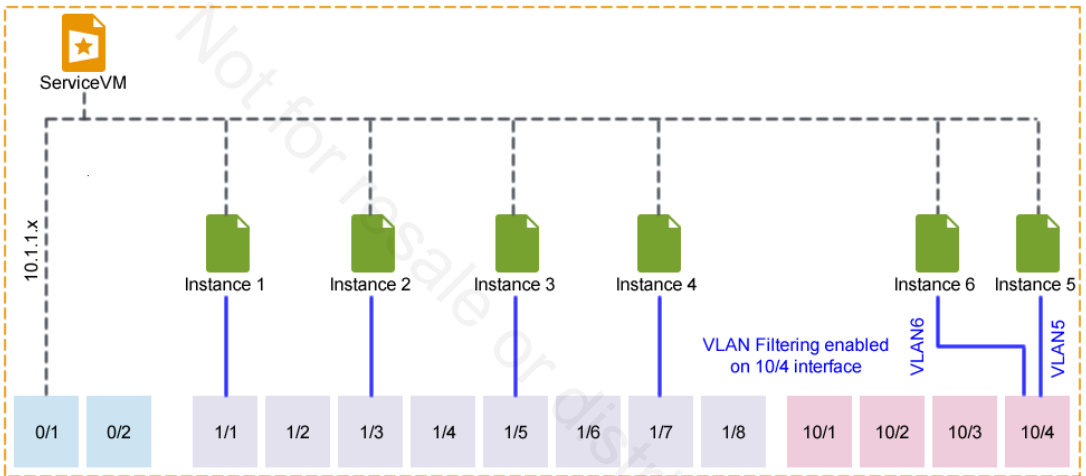
Simple consolidation deployments are used when compliance is not a concern and when all instances are located in the same security zone. Instance density is limited to the number of

physical interfaces. Data plane isolation is achieved by not sharing physical interfaces. There are 4096 VLANs per interface and instance.

Data Plane Isolation with Shared Interfaces

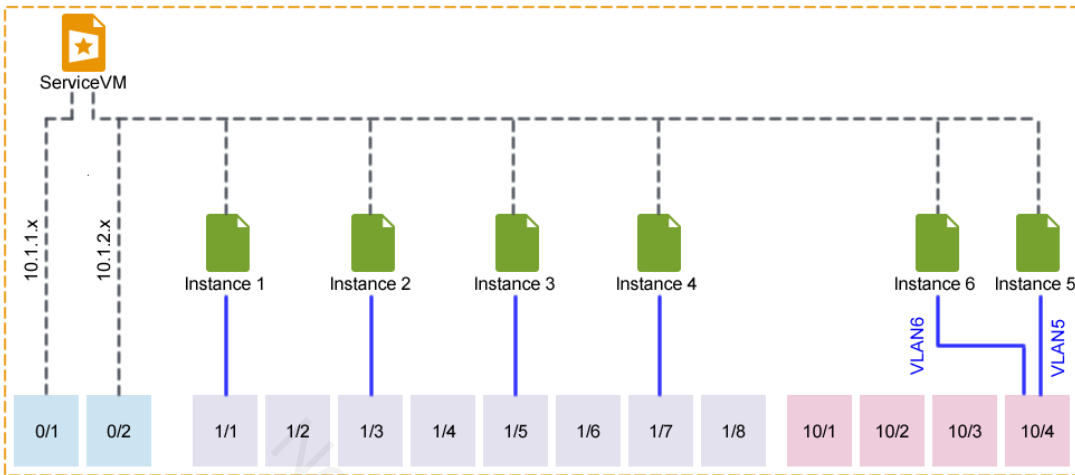
SR-IOV provides the capability to safely share an interface across instances. SR-IOV provides better performance--a side effect of its intended purpose to safely virtualize a single physical interface into multiple virtual interfaces.

SR-IOV provides the ability to isolate traffic by providing VLAN filtering at the interface level. For example, in the image below, traffic from VLAN6 is only sent to instance 6 and traffic from VLAN 5 is only sent to instance 5. You can test and validate this by doing a broadcast storm against instance 6 and instance 5 will not be impacted at all.

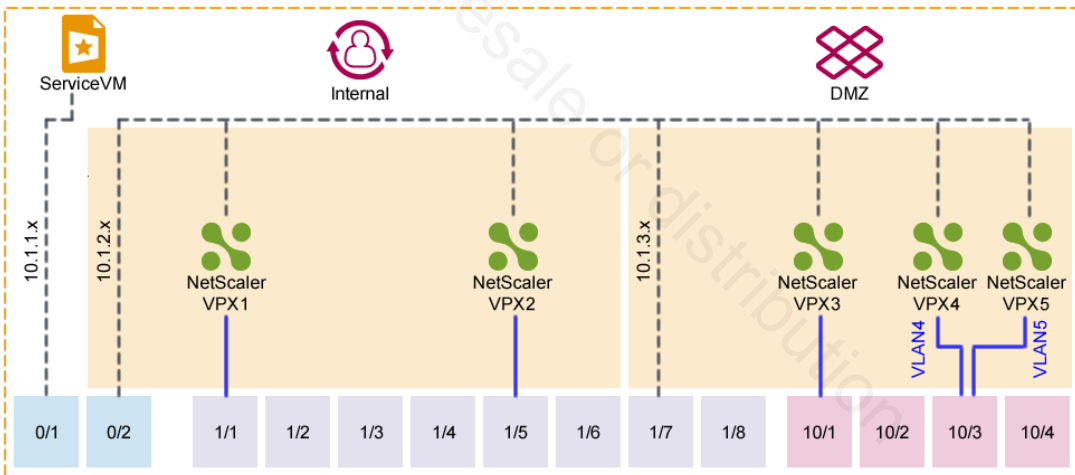


Simple Consolidation with Delegated Administration

The NetScaler SDX has the ability to keep the traffic on the device, or to force communication between the Service VM and the instances off the device and then back on. This occurs when it is important to send traffic through an intermediary such as a firewall for audit or compliance purposes.



Consolidation Across Security Zones



Data and management planes are isolated to support network segmentation use cases. There is support for multiple management networks and separate Service VM from NSIPs. Other features of this consolidation method include:

- Very strong data plane isolation options.
 - Dedicated interfaces to instances.
 - Shared interfaces with and without VLAN filtering.
- Multiple management networks.

- Support for hierarchical networking.
- Flexible data ports.
 - Dedicated interfaces for security zones.
 - Sharing of interfaces within security zones.
- Traffic isolation at the hardware level.
 - MAC and VLAN filtering.

SNMP

You can configure a Simple Network Management Protocol (SNMP) agent on the NetScaler SDX appliance to generate asynchronous events, which are called traps. The traps are generated whenever there are abnormal conditions on the NetScaler SDX appliance. The traps are then sent to a remote device called a trap listener, which signals the abnormal condition on the NetScaler SDX appliance.

Each SNMP network management application uses SNMP to communicate with the SNMP agent on the NetScaler SDX appliance. The SNMP agent searches its management information base (MIB) to collect the data requested by the SNMP Manager and provides the information to the application.

SNMP Trap Destinations

The SNMP agent on the SDX appliance generates traps that are compliant with SNMPv2 only. The supported traps can be viewed in the SDX MIB file. You can download this file from the Downloads page in the SDX user interface.

Adding an SNMP Trap Destination

1. On the configuration tab, in the navigation pane, expand **System**, **SNMP** and then click **Trap Destinations**.
2. In the **SNMP Traps** pane, click **Add**.
3. In the **Create SNMP Trap Destinations** dialog box, specify values for the following parameters:
 - Destination Server--IPv4 address of the trap listener to which you will send the SNMP trap messages.
 - Port -- UDP port at which the trap listener listens for trap messages. Must match the setting on the trap listener, or the listener drops the messages. Minimum value: 1. Default: 162.
 - Community -- Password (string) sent with the trap messages, so that the trap listener can authenticate them. This can include letters, numbers and symbolic characters: hyphen (-); period (.); hash (#); space (); at (@); &equals (=); colon (:); and underscore (_).



You must specify the same community string on the trap listener device, or the listener drops the messages. The default string is "public."

4. Click **Create**. The SNMP trap destination that you added appears in the SNMP Traps pane.
To modify the values of the parameters of an SNMP trap destination, in the SNMP Traps Destinations pane, select the trap destination that you want to modify and then click **Edit**. In the **Configure SNMP Trap Destination** dialog box, modify the parameters.
To remove an SNMP trap, in the SNMP Trap Destinations pane, select the trap destination that you want to remove and then click **Delete**. In the Confirm message box, click **Yes** to remove the SNMP trap destination.

Adding an SNMP Manager Community

You must configure the NetScaler SDX appliance to allow the appropriate SNMP managers to query it. You must also provide the SNMP manager with the required appliance-specific information. For an IPv4 SNMP manager, you can specify a host name instead of a manager IP address. If you specify a host name instead of a manager IP address, you must add a DNS name server that resolves the host name of the SNMP manager to its IP address.

You must configure at least one SNMP manager. If you do not configure an SNMP manager, the appliance does not accept or respond to SNMP queries from any IP address on the network. If you configure one or more SNMP managers, the appliance accepts and responds only to SNMP queries from those specific IP addresses.

Configuring an SNMP Manager

1. On the Configuration tab, in the navigation pane, expand **System** and then expand **SNMP**.
2. Click **Managers**.
3. In the details pane, click **Add**.
4. In the Add SNMP Manager Community dialog box, set the following parameters:
 - SNMP Manager.
 - Community.
5. Click **Add** and then click **Close**.

Configuring the NetScaler for SNMPv3 Queries

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

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

- SNMP Views.
- SNMP Users.

These entities function together to implement the SNMPv3 security features. Views are created to allow access to subtrees of the MIB.

SNMP Views

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

Configuring an SNMP View

1. On the Configuration tab, in the navigation pane, expand **System** and then expand **SNMP**.
2. Click **Views**.
3. In the details pane, click **Add**.
4. In the Add SNMP View dialog box, set the following parameters:
 - **Name:** Name for the SNMPv3 view. This can include letters, numbers and symbolic characters: hyphen (-); period (.); hash (#); space (); at (@); equals (=); colon (:); and underscore (_).
 - **Subtree:** A particular branch (subtree) of the MIB tree, which you want to associate with this SNMPv3 view. You must specify the subtree as an SNMP OID.
 - **Type:** Include or exclude the subtree, specified by the subtree parameter, in or from this view. This setting can be useful when you have included a subtree, such as A, in an SNMPv3 view and you want to exclude a specific subtree of A, such as B, from the SNMPv3 view.

SNMP Users

After you have created an SNMP view, add SNMP users. SNMP users have access to the MIBs that are required for querying the SNMP managers.

Configuring an SNMP User

1. On the Configuration tab, in the navigation pane, expand **System** and then expand **SNMP**.
2. Click **Views**.
3. In the details pane, click **Add**.

4. In the Add SNMP User dialog box, set the following parameters:
 - Name: Name for the SNMPv3 user. This can include letters, numbers and symbolic characters: hyphen (-); period (.); hash (#); space (); at (@); && equals (=); colon (:); and underscore (_).
 - Security Level: Security level required for communication between the appliance and the SNMPv3 users.
 - Authentication Protocol: Authentication algorithm used by the appliance and the SNMPv3 user for authenticating the communication between them. You must specify the same authentication algorithm when you configure the SNMPv3 user in the SNMP manager.
 - Authentication Password: Pass phrase to be used by the authentication algorithm. This can include letters, numbers and symbolic characters: hyphen (-); period (.); hash (#); space (); at (@); && equals (=); colon (:); and underscore (_).
 - Privacy Protocol: Encryption algorithm used by the appliance and the SNMPv3 user for encrypting the communication between them. You must specify the same encryption algorithm when you configure the SNMPv3 user in the SNMP manager.
 - View Name: Name of the configured SNMPv3 view that you want to bind to this SNMPv3 user. An SNMPv3 user can access the subtrees that are bound to this SNMPv3 view as type INCLUDED but cannot access the ones that are type EXCLUDED.

SNMP Alarms

The NetScaler SDX appliance provides a predefined set of condition entities called SNMP alarms. When the condition set for an SNMP alarm is met, the appliance generates SNMP trap messages that are sent to the configured trap listeners. For example, when an alarm is enabled, a trap message is generated and sent to the trap listener whenever a device (instance) is provisioned on the appliance. You can assign a severity level to an SNMP alarm, the corresponding trap messages are then assigned that severity level.

Severity levels, defined on the appliance and in decreasing order of severity are:

- Critical
- Major
- Minor
- Warning
- Informational (default)

For example, if you set a Warning severity level for the SNMP alarm named deviceAdded, the appliance generates trap messages labeled with the Warning severity level when a device is added.

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

Modifying a Predefined SNMP Alarm

1. On the Configuration tab, in the navigation pane, expand **System**.
2. Expand **SNMP** and then click **Alarms**.
3. In the SNMP Alarm Configuration details pane, select an alarm and then click **Edit**.
4. In the Configure SNMP Alarm dialog box, from the Severity list, select a severity level.
5. To enable the alarm, select **Enable Alarm**.
6. Click **OK**.

System Health Monitoring

System health monitoring detects errors in the monitored components so that you can take corrective action to avoid a failure. The following components are monitored on a NetScaler SDX appliance and can be viewed from the Monitoring tab:

- Hardware and software resources.
- Physical and virtual disks.
- Hardware sensors, such as fan, temperature, voltage and power supply sensors.
- Interfaces.

Third-Party Virtual Machines

The SDX appliance supports provisioning of the following third-party virtual machines (instances):

- SECUREMATRIX GSB.
- Websense Protector.
- BlueCat DNS/DHCP Server.

SECUREMATRIX GSB provides a highly secure password system that eliminates the need to carry any token devices.

Websense Protector provides monitoring and blocking capabilities, preventing data loss and leaks of sensitive information.

BlueCat DNS/DHCP Server delivers DNS and DHCP for your network. You can provision, monitor, manage and troubleshoot an instance from the Management Service.

Third-party instances use the SDXTools daemon to communicate with the Management Service. The daemon is preinstalled on the provisioned instance. You can upgrade the daemon when new versions become available.



The total number of instances that you can provision on an SDX appliance depends on the license installed on the appliance.



You must upgrade your XenServer version to 6.1.0 before you install any third-party instance.

Managing a NetScaler SDX Appliance: Overview

This interactive exercise demonstrates how to log on to a NetScaler SDX appliance, create a new NetScaler VPX instance, bind two NetScaler VPX instances as an HA pair and log on to the NetScaler HA pair through the SNIP address.

Managing a NetScaler SDX Appliance: Adding and Configuring a NetScaler VPX Instance

1. Log on to the NetScaler SDX Service VM.
 - a. Click the **User Name** text box and type `nsroot`, then press **Enter**.
 - b. Type `nsroot` in the **Password** field and then press **Enter**.
2. Create a new NetScaler VPX instance called CitrixEdu-2 with an IP address of 172.21.0.21 using an Platinum-level feature license.
 - a. Click on the **NetScaler** node, then **Instances**.
 - b. Click the **Add** button at the top of the **Instances** pane on the home screen.
 - c. Type `CitrixEdu-2` in the **Name** field, then press **Enter**.
 - d. Type `172.21.0.21` in the **IP Address** field, then press **Enter**.
 - e. Type `172.21.0.1` in the **Gateway** field, then press **Enter**.
 - f. Click the drop down list next to the **Browse** button under the **XVA File** option and select **Appliance**.
 - g. Select the **NSVPX-XEN-11.0-64.34_nc.xva** file then click **Open**.
 - h. Select **Platinum** from the **Feature License** drop-down list box.
3. Finish configuring the CitrixEdu-2 NetScaler VPX instance.
 - a. Under **Instance Administration**.
 - b. Type `CitrixAdmin` in the **User Name** field, then press **Enter**.
 - c. Type `Password1` in the **Password** field, then press **Enter**.
 - d. Type `Password1` in the **Confirm Password** field, then press **Enter**.
 - e. Click **Add** under the **Data Interfaces** field and verify that **1/1** is selected then click **Add**.
 - f. Scroll to the bottom of the window by clicking the bottom of the right scroll bar.
 - g. Click **Done** to complete the NetScaler VPX instance configuration.



The Username and Password are case sensitive on the NetScaler.

4. View the CitrixEdu-2 instance information dialog box in the NetScaler Instances pane.
 - a. Click the **arrow** next to the CitrixEdu-2 name in the NetScaler instances pane.



This field displays all of the configuration information for the NetScaler Instances on an SDX appliance.

5. Log on to the newly created NetScaler VPX instance using the CitrixAdmin/Password1 credentials.
 - a. Click the **new tab button** at the top of the Firefox browser window.
 - b. Type 172.21.0.21, then press **Enter**.
 - c. Select the **User Name** field, type CitrixAdmin and then press **Enter**.
 - d. Type Password1 in the **Password** field and then press **Enter**.
6. Configure the CitrixEdu-2 with a SNIP address of 172.21.0.251.
 - a. Click **Skip** on the **Citrix User Experience Improvement Program** pop-up window.
 - b. On the **Welcome** screen click in the **Subnet IP Address** field.
 - c. Type 172.21.0.251 in the IP Address field and press **Enter**, click **Done**, then click **Continue**.
 - d. Click on **System**, then **Network**, then **IPs**.
 - e. Highlight the SNIP **172.21.0.251** then click **Edit**.
 - f. Scroll to the bottom of the page and check the box for **Enable Management Access control to support the below listed applications**, then click **OK**.
7. Join the CitrixEdu-1 and CitrixEdu-2 NetScaler VPX instances in a High-Availability pair, with CitrixEdu-2 being the primary node.
 - a. Select the **System** node, select the **High Availability** sub-node and then click **Add**.
 - b. Type 172.21.0.11 in the **Remote Node IP Address** field, press **Enter**.
 - c. Click in the **User name** field and enter CitrixAdmin then press enter.
 - d. Type Password1 in the **Password** field and then press **Enter**.
 - e. Click **Create**, confirming that the High Availability pair has been set up, then click **Logout** at the top of the screen.
 - f. Connect to the **SNIP** of the HA pair.
 - g. Select the **browser address bar**, type 172.21.0.251 and press **Enter**.
 - h. Select the **User Name** field, type CitrixAdmin and press **Enter**. Type Password1 in the **Password** field and press **Enter**.
 - i. Select the **System** node, then select the **High Availability** sub-node.



The CitrixEdu-1 node is listed as Secondary and the CitrixEdu-2 node is listed as Primary.

- j. Click **Logout** at the top of the screen.

Not for resale or distribution

Module 13

Monitoring and Administration

13

Not for resale or distribution

Not for resale or distribution

Monitoring and Management Manual

Overview

A NetScaler deployment can be monitored by four methods: logging, the reporting tool, Command Center, and the Dashboard. SNMP allows you to generate traps for abnormal conditions on the NetScaler system. Citrix AppFlow is an open standards technology that transforms the data collected by existing networking devices into powerful operational and business intelligence. And EdgeSight Monitoring for NetScaler monitors the end-user experience with web applications that are served in an environment using NetScaler.

When troubleshooting a NetScaler environment, you should reference troubleshooting resources such as the Citrix Knowledge Center, product documentation, Citrix Technical Support, collected NetScaler data, and troubleshooting logs.

After completing this module, you will be able to:

- Determine which data to monitor in the environment.
- Monitor the NetScaler system by viewing statistics through the Dashboard.
- Assess the health of the NetScaler system by reviewing built-in reports.
- Monitor the NetScaler system using syslogs.
- Investigate an issue using the correct utility or information source.
- Troubleshoot traffic management issues with NetScaler.

Monitoring Needs

The NetScaler system provides integrated, real-time monitoring of application and client traffic, along with the capability to record critical application delivery information to automatically optimize the system and load-balancing policies.

The following use cases are examples of the growing monitoring and information demands:

- Mapping the end-user experience for e-commerce sales
- Ensuring improved load balancing across the datacenter
- Tracking traffic through web applications
- Web application performance
- Identifying when application response times exceed service level agreements (SLAs) for transactions
- Monitoring the end-user experience
- Monitoring abnormal client-server behavior based on packet counts, byte counts and server response times

Monitoring Methods

A NetScaler system provides several ways to monitor your environment, including:

- SNMP
- AppFlow
- EdgeSight Monitoring
- Insight Center
- Action Analytics
- Command Center
- Dashboard
- Reporting tool

SNMP

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

AppFlow

AppFlow provides visibility at the transaction level for HTTP, SSL, TCP and SSL_TCP flows. You can sample and filter the flow types that you want to monitor. AppFlow uses actions and policies to send records for a selected flow to a specific set of collectors. An AppFlow action specifies which set of collectors will receive the AppFlow records. Policies, which are based on default expressions, can be configured to select flows for which flow records will be sent to the collectors specified by the associated AppFlow action.

EdgeSight Monitoring

The EdgeSight Monitoring application provides webpage monitoring data with which you can monitor the performance of various web applications served in a NetScaler environment. You can now export this data to AppFlow collectors to get an in-depth analysis of the web applications. AppFlow, which is based on IPFIX standard, provides more specific information about web application performance than does EdgeSight Monitoring alone. You can configure both load-balancing and content-switching virtual servers to export EdgeSight Monitoring data to AppFlow collectors.

With the release of XenDesktop 7, Citrix EdgeSight and Citrix Desktop Director has been integrated into a single architecture to provide advanced analytics. EdgeSight and Director provide a detailed and intuitive overview of XenDesktop environments and enables monitoring and troubleshooting of system issues.

Insight Center Overview

NetScaler Insight Center is a virtual appliance that is installed on a hypervisor. Insight Center collects detailed information about web and virtual desktop traffic that passes through a NetScaler.

Insight Center has two components:

- Web Insight – monitors HTTP, SSL and TCP traffic passing through a load balancing and content switching virtual server defined on the NetScaler appliance. Web Insight supports HTTP, SSL and TCP virtual servers.
- HDX Insight – monitors ICA traffic passing through NetScaler Gateway virtual servers defined in the NetScaler appliance.

Action Analytics

NetScaler Action Analytics collects detailed information on administrator-selected application traffic flows and generates multidimensional views of what is happening in real time. Graphical and tabular displays enable you to instantly view critical application-delivery parameters, such as which top URLs are being accessed or which top clients are consuming application resources and then view more detailed statistics including number of overall requests, total bandwidth being consumed and the response times of the back-end infrastructure delivering the application content.

Command Center

Citrix Command Center is a management and monitoring solution for Citrix application networking products, which includes Citrix NetScaler. Command Center enables network administrators and operations teams to manage, monitor and troubleshoot the entire global application delivery-infrastructure from a single, unified console. This centralized management solution simplifies operations by providing you with real-time performance monitoring of your enterprise-wide application-delivery infrastructure and automating management tasks that need to be run across multiple devices.

The following list includes some of the management tasks that are simplified by using Command Center:

- Quickly address and resolve device and network issues and keep the network running effectively by monitoring and managing the SNMP and syslog events generated on your devices.

- Understand the traffic patterns, gather data for capacity planning, and monitor the performance of the entire application delivery infrastructure by using historical charts and performance graphs.
- Monitor and manage the states of virtual servers, services, and service groups across the NetScaler infrastructure.
- Troubleshoot configuration errors or recover unsaved configuration on sudden system shutdown by running audit policies.

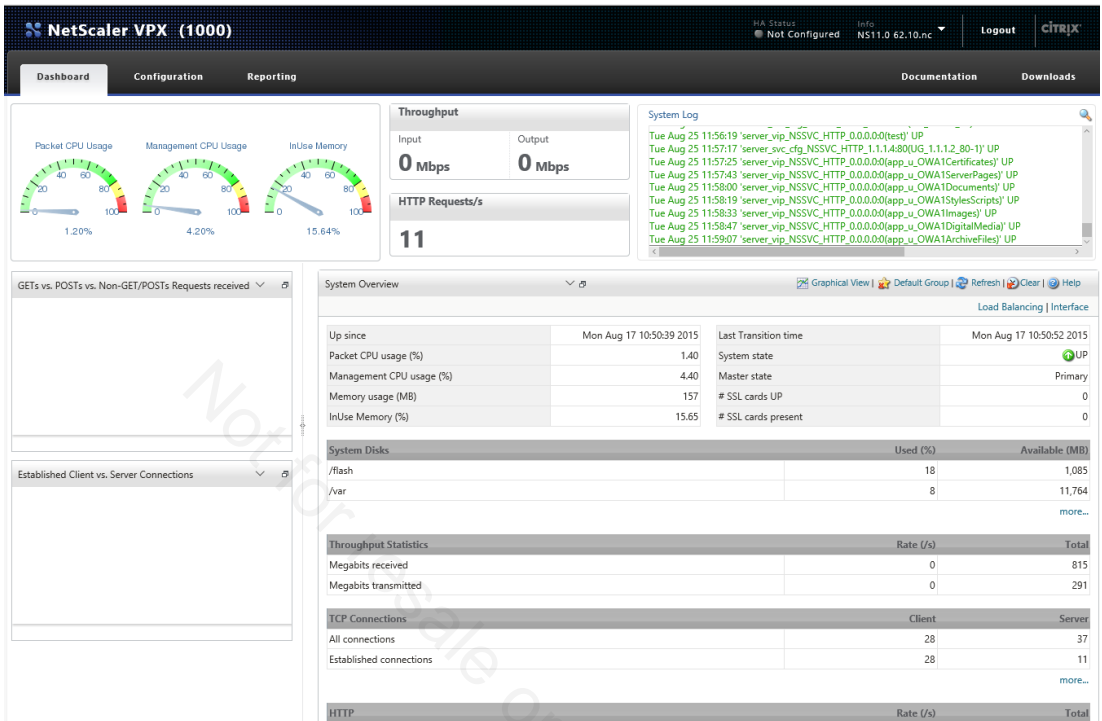
Available as a software platform or a hardware appliance, Command Center enables you to monitor and control a diversified deployment of multiple NetScaler solutions including:

- NetScaler ADC
- NetScaler SDX
- NetScaler Gateway
- NetScaler Application Firewall
- CloudBridge
- XenServer

The NetScaler Command Center hardware appliance can manage 300 Citrix networking devices, includes 500GB of built-in storage and is preloaded with a database.

To monitor and manage Citrix devices, you need to connect to the Command Center server by using the HTML Web client and then add the devices for discovery. For more information about Citrix Command Center, see Citrix product documentation at <http://docs.citrix.com>.

Dashboard



The Dashboard is the monitoring page in the configuration utility. The Dashboard is HTML-based, displays critical performance statistics and provides real-time data. The data displayed is for the last five minutes of operation and is updated every seven seconds.

System and Feature Counters

The Dashboard monitors both system and features counters.

System counters include:

- CPU utilization
- Memory utilization
- System throughput
- System statistics
- Protocol statistics
- High availability
- Bridge

- ACL
- Network entities
- SNMP
- SSL
- System health

Feature counters include:

- Audit
- AAA
- VPN
- TCP
- Load balancing
- Content switching
- Services
- HTTP compression
- Decompression
- Integrated cache



Certain features are dependent on the enabled licenses.

Dashboard Components

The Dashboard components give you information on the health and status of the NetScaler system.

CPU Utilization Pane The CPU Utilization pane displays the real-time CPU utilization of the system as a percentage of the total capacity. The color codes for the pane are:

- Green, which indicates that the CPU utilization is within safe limits
- Yellow, which indicates that the CPU utilization is high
- Red, which indicates that the CPU utilization has reached a critical limit

Memory Utilization Pane The Memory Utilization pane displays the real-time memory utilization of the system. While the dial indicates the percentage of the memory used, you can drag the mouse cursor over the pane for an absolute value. The color codes for the pane are:

- Green, which indicates that the memory utilization is within safe limits
- Yellow, which indicates that the memory utilization is high
- Red, which indicates that the memory utilization has reached a critical limit

System Throughput Pane The System Throughput pane displays throughput in terms of incoming and outgoing traffic.

HTTP Requests per Second Pane The HTTP Request per Seconds pane illustrates the HTTP requests served by the system each second.

System Log Pane The System Log pane provides a view of system events and alerts that are generated when the Dashboard is connected to the system. The data is real time and has an unlimited history; however, all data is lost when the Dashboard is restarted.

Built-in Comparative Chart Pane The Built-in Comparative Charts pane monitors a built-in set of counters using charts. The charts depict the variation of two or more counters over time. Options for viewing the charts include chart type, chart appearance, legends, independent window views and custom charts.

Group Monitoring Pane The Group Monitoring pane provides a view of all counters of a feature or protocol and the ability to plot any number of counters on custom charts. By default, this pane provides a snapshot of the system.

Reporting Tool

The NetScaler system provides the reporting tool in order to parse and display the data contained in the syslog file. The reporting tool is a web-based interface that provides built-in reports that display statistics collected by the nscollect utility.

Reports allow you to plot and monitor statistics for the various functional groups over a specified time interval, which assists in the troubleshooting and analyzing of the behavior of the NetScaler system. You can access the reporting tool by selecting **Reporting** when logged on to the NetScaler system.

Working with Reports

You can plot and monitor statistics on the NetScaler system over a specified time interval. Reports enable you to troubleshoot or analyze the behavior of your system. There are two types of reports: built-in reports and custom reports. The content for built-in or custom reports can be viewed in a graphical format or a tabular format. The graphical view consists of line, area and bar charts that can display up to 32 sets of data, also known as counters. The tabular view displays the data in columns and rows. This view is useful for debugging error counters.

The default report that is displayed in the reporting tool is CPU versus memory usage and HTTP requests rate. Reports can be generated for the last hour, last day, last week, last month, last year or a customized interval of time.

You can do the following with reports:

- Change the graphical display type, such as bar chart or line chart.
- Customize charts in a report.
- Export the chart as an Excel comma-separated value (CSV) file.
- View the charts in detail by zooming in, zooming out, or using a drag-and-drop operation (scrolling).
- Set a report as the default report for viewing whenever you log on.
- Add or remove counters.
- Refresh reports to view the latest performance data.

Using Built-in Reports

The reporting tool provides built-in reports for frequently viewed data. Built-in reports are available for the following seven functional groups:

- System
- Network
- SSL
- Compression
- Integrated Cache
- NetScaler Gateway
- Application Firewall

By default, the built-in reports are displayed for the last day. However, you can view the reports for the last hour, last week, last month, or last year.




You cannot save changes to built-in reports, but you can save any modifications to a built-in report by saving it as a custom report.

AppFlow on the NetScaler System

The NetScaler system is a central point of control for all application traffic in the datacenter. It collects information about flow and about the end user's session that is valuable for application-performance monitoring, analytics and business intelligence applications. It also collects webpage-performance data and database information. AppFlow transmits the information by using the Internet Protocol Flow Information eXport (IPFIX) format. IPFIX is the standardized version of Cisco's NetFlow and is widely used to monitor network flow 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.

AppFlow uses actions and policies to send records for a selected flow to a specific set of collectors. An AppFlow action specifies which set of collectors will receive the AppFlow records. Policies, which are based on default expressions, can be configured to select flows for which flow records will be sent to the collectors specified by the associated AppFlow action. To limit the types of flows, you can enable AppFlow for a virtual server. AppFlow can also provide statistics for the virtual server. You can enable AppFlow for a specific service, representing an application server and monitor the traffic to that application server.

How AppFlow Works



In a common deployment scenario, inbound traffic flows to a virtual IP address (VIP) on the NetScaler system and is load balanced to a server. Outbound traffic flows from the server to a mapped or subnet IP address on the NetScaler system and from the VIP to the client. A flow is a unidirectional collection of IP packets identified by the following five tuples: sourceIP, sourcePort, destIP, destPort and protocol.

How Insight Center Collects AppFlow Data

Insight Center retrieves AppFlow logging information for each monitored application, analyzes the information and presents it as visual reports. To enable data collection, you must enable AppFlow on each virtual server from which you want Insight Center to retrieve the data. Services bound to those virtual servers must also have AppFlow enabled.

The NetScaler appliance is usually the central point where traffic flows through the network making it an ideal place to collect network flow information. The AppFlow feature in a NetScaler appliance uses the Internet Protocol Flow Information Export (IPFIX) format to transmit network flow information. AppFlow provides visibility at the transaction level for HTTP, SSL, TCP and SSL_TCP flows. AppFlow uses actions and policies to send records for a selected flow to specific set of collectors. An AppFlow action specifies which set of collectors will receive the Appflow records. Policies, which are based on Advanced Expressions, can be configured to select flows for which flow records will be sent to the collectors specified by the associated AppFlow action.

Some of the metrics collected by Web Insight:

- Number of hits received by the NetScaler
- Bandwidth of the NetScaler
- Number of hits received by an application
- Bandwidth used by an application
- Response time for an application
- Number of hits received for a URL
- Number of requests sent by a client
- Latency of the client-side network
- Number of hits received by the virtual servers
- Processing time on a virtual server
- Number of hits by HTTP request method
- Number of hits by HTTP response status

Some of the metrics collected by HDX Insight:

- Latency of the client-side network
- ICA RTT
- Bandwidth used in an ICA session
- Total number of active sessions in a given time interval
- Total number of applications active during a given time period
- Average rate at which data is transferred over the ICA session
- Total number of unique end-user sessions in a given time period
- Licenses in use for SSL VPN and ICA traffic

HDX Insight

HDX Insight allows administrators of Citrix XenApp and Citrix XenDesktop environments a way to monitor the end users and performance of the applications hosted on those products. HDX Insight captures data about the ICA traffic that flows between the clients and the servers, generates AppFlow records by doing deep packet inspection of the data and presents the records as visual reports on the Insight Center dashboard.

HDX Insight currently supports collecting AppFlow data from NetScalers in single-hop mode. In single-hop mode, end users access the NetScaler through a virtual private network, such as the NetScaler Gateway.



HDX Insight does not support standalone Access Gateway Standard appliances.

Configuring Insight Center

The Insight Center console can be accessed by typing the IP address assigned to the Insight Center virtual appliance in the address bar of a browser.

When accessing the NetScaler Insight Center for the first time, a welcome screen directs you to get started by specifying the NetScaler appliance to be monitored. You must add at least one NetScaler appliance to the Insight Center inventory and provide the following information for that appliance:

- NetScaler IP address
- User name
- Password

Insight Center can be installed on the following platforms:

- Citrix XenServer 5.6 or later
- VMware ESX 4.1 or later



Insight Center is only supported on NetScaler nCore builds.

For detailed information about installing Insight Center, see the product documentation for *Installing NetScaler Insight Center* at <http://docs.citrix.com/en-us/netscaler-insight/11-0/installing-insight-center.html>.

Enabling Data Collecting for Insight Center

To enable data collection, you must enable AppFlow on each virtual server on the NetScaler from which you want to collect data. When enabling AppFlow, you specify an expression for that virtual server. There are several pre-configured expressions for web traffic. You can also use multiple expressions by using the logical operators AND (&&) or OR (||). For VPN virtual servers, the only rule available is `ns_true`.

You cannot enable AppFlow for a virtual server on a NetScaler appliance on which Appflow is already enabled for four virtual servers. Each virtual server functions as an Appflow collector and you can only put four AppFlow collectors on one NetScaler appliance.



NetScaler Insight Center does not support IPv6.

NetScaler Log Management

Newslog stores console messages, events, and performance statistics, which are written to newslog every seven seconds. The log is saved every three hours. Once saved, these newslog files roll on a periodic basis based on a specified lifetime of the nsconmsg process. When a log rolls, each older log file is compressed using GZIP and saved with a number from 0 through 100. Newslog is a binary file stored in `/var/nslog`.

The Manage Logs option in the Diagnostic pane provides the ability to view information from the log files, save viewed data, download the log files and delete the unwanted log files from the NetScaler system. The following actions can be performed:

- Viewing events
- Viewing log files duration
- Viewing events from a specific time
- Viewing console messages
- Downloading log files
- Deleting log files

Audit Logging

Audit logging enables you to log the NetScaler states and status information collected by various modules in the kernel and in the user-level daemons. The NetScaler system allows you to customize the logging of system events according to the need of a site. You can record events either to files on the NetScaler system or to external log servers. The NetScaler system supports two logging formats:

- syslog
- nslog

By default, the NetScaler system records system events in syslog format to `/var/log/ns.log` using local0 and records SSL VPN events to `/var/log/nsvpn.log` using local1. Logging of these events is enabled by default. Syslog uses UDP over port 514 and syslog traffic is sent in clear text. The NetScaler kernel controls syslog processing. Nslog is a proprietary binary format that records more detailed event information than the syslog format. You can log syslog and nslog events to either a local file on the NetScaler system or to a remote server.



`syslog.conf` should not be modified.

Configuring NetScaler for Audit Logging

Policies define the syslog or nslog protocol and the server actions define where each log is sent. For server actions, you specify the system information, which runs the syslog or the nslog server.

The NetScaler system logs the following information related to TCP connections:

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

To configure audit logging on the NetScaler system, you must:

1. Configure audit servers
2. Configure audit policies
3. Bind the audit policies globally

Configuring an Auditing Server

An external nslog audit server can run either the Linux, FreeBSD or Windows operating system.



Local logging still occurs on the NetScaler system when an external audit server is configured.

You can configure an audit server on the NetScaler system by specifying the following parameters:

- The name of the syslog server action or nslog server action
- IP address of the auditing server
- Port through which to communicate
- Severity levels of messages to be logged
- Format of the date stamp
- Facility value (RFC 3164) assigned to the log message
- Time zone for the time stamp
- TCP logging
- ACL logging
- Enable user-configurable log messages
- Enable or disable export log messages to the AppFlow collectors.

In the configuration utility, go to Configuration > System > Auditing. When using the command-line interface, type:

```
add audit logAction <name> <ServerIPAddress> -serverPort <port> -  
logLevel <LogLevel>  
-dateFormat ( MMDDYYYY | DDMMYYYY ) -logFacility <logFacility>  
-tcp (None | ALL ) -timeZone ( GMT_TIME | LOCAL_TIME )
```



Replace logAction with the type, either syslogAction for a syslog auditing server or nslogAction for an nslog auditing server.

Global Auditing Parameters

Global auditing parameters help you log all the states and status information collected by different modules in the kernel, as well as in the user-level daemons in the NetScaler system.

Configure global auditing parameters on the NetScaler system by specifying the following:

- Auditing type
- IP address of the external server used to store auditing messages
- Port number used for communication between the system and the external logging server
- Level of logging detail
- Date format
- Time zone
- Log facility value
- TCP logging

In the configuration utility, go to Configuration > System > Auditing > Policies. In the command-line interface, type:

```
add audit Params <name> <ServerIPAddress> -serverPort <port> -  
logLevel <LogLevel>  
-dateFormat MMDDYYYY | DDMMYYYY ) -logFacility <logFacility>  
-tcp (None | ALL ) -timeZone ( GMT_TIME | LOCAL_TIME )
```



Replace Params with either syslogParams for a syslog auditing server, or nslogParams for an nslog auditing server.

Configuring Auditing Policies

Auditing policies decide which messages are generated and logged during the session. These messages are logged in the syslog or nslog format. Different types of messages are logged based on the level of logging selected.

Auditing policies are used for the NetScaler Gateway feature of the NetScaler system.

You can add an auditing policy on the NetScaler system by specifying the policy name and rule action. In the configuration utility, go to Configuration > System > Auditing > Policies. Using the command-line interface, type:

```
add audit policy <Policy name> <rule> <action>
```

Binding Auditing Policy

You can configure auditing on the NetScaler system at several levels in the following order of priority:

1. Global
2. Virtual server
3. Group
4. User

At each of these levels, you can define multiple auditing policies. You can either bind auditing policies globally or to virtual servers, groups and users.

You can bind an auditing policy on the NetScaler system by specifying the policy and assigning a priority. In the configuration utility, go to Configuration > System > Auditing > Policies. When using the command-line interface, type:

```
bind system global <PolicyName> -priority <positive_integer>
```

Audit Messages

You can view recent and historical audit messages in both the configuration utility and in the command-line interface.

You can view recent audit messages on the NetScaler system by specifying the level of detail to be displayed and the number of audit messages to be shown and the style to display the messages. In the configuration utility go to Configuration > System > Auditing. When using the command-line interface, type:

```
show audit messages -logLevel <logLevel> -  
numOfMesgs <positive_integer>
```

You can also view historical audit messages on the NetScaler system by going to Configuration > System > Auditing > Syslog messages in the configuration utility.

Discussion Question

What are common logging and reporting tools? How do you use logs and reporting tools in your environment?

Troubleshooting Resources

Useful troubleshooting resources include:

- Citrix Knowledge Center
- Citrix product documentation
- Citrix Technical Support
- Collected NetScaler data
- Troubleshooting log
- Citrix Tools as a Service (TaaS)

Citrix Knowledge Center

An important tool for research, the Citrix Knowledge Center is the official resource for technical information on Citrix products, hotfixes, security bulletins and troubleshooting guides. The Knowledge Center contains forums through which external sources can be contacted. It also contains a thorough library of documentation and white papers. The forum community is also very knowledgeable and responsive.

For more information about the Citrix Knowledge Center, see Citrix articles at <http://support.citrix.com>.

Citrix Product Documentation

The most recent NetScaler product documentation is available from Citrix product documentation at <http://docs.citrix.com>. The product documentation can also be downloaded through the NetScaler graphical user interface.

Citrix Technical Support

Citrix Technical Support can be an extremely useful resource when troubleshooting a NetScaler deployment. When contacting Technical Support, compile all the previously gathered information pertaining to the issue. Be prepared to engage the Support Engineer in an exchange of ideas.

Collected NetScaler Data

Gathering all available NetScaler data is important when beginning analysis. Data resources include:

- Results of the `show techsupport` command, which provides configuration files, performance log data, system messages and other relevant system information
- User feedback, which can include screen captures
- Documented steps for reproducing the issue
- Network packet traces
- Syslogs
- Web logs
- SNMP alarms
- Network topology diagrams and other deployment documentation

Additional Resources

Additional resources for gathering data include:

- User details, which can include screen captures
- Documented steps for reproducing an issue
- Network packet traces
- Network topology diagrams and other deployment documentation
- Database status from your database administrator
- Event logs
- System logs
- Web logs
- Troubleshooting logs
- Citrix Auto Support, available at <http://taas.citrix.com>
- NetScaler Call Home feature. Call Home registers your NetScaler system with the Citrix Technical Support server (TaaS) and monitors the appliance for common error conditions. If your appliance is successfully registered with the TaaS server, Call Home automatically uploads system-debug data to that server in the event that one of the conditions occurs. Call Home is supported by any NetScaler MPX appliance running release 10 or later. For more information, see the NetScaler 11 Frequently Asked Questions (FAQs) document at <http://docs.citrix.com>.

Sometimes the cause of an issue is not readily apparent, and researching the symptoms may be required. A key external resource is the online Citrix Knowledge Center, available at <http://support.citrix.com>.

The Citrix Knowledge Center is the official resource for technical information on Citrix products, including hotfixes, security bulletins, troubleshooting guides, documentation and white papers. The Knowledge Center contains forums through which external sources can be contacted.

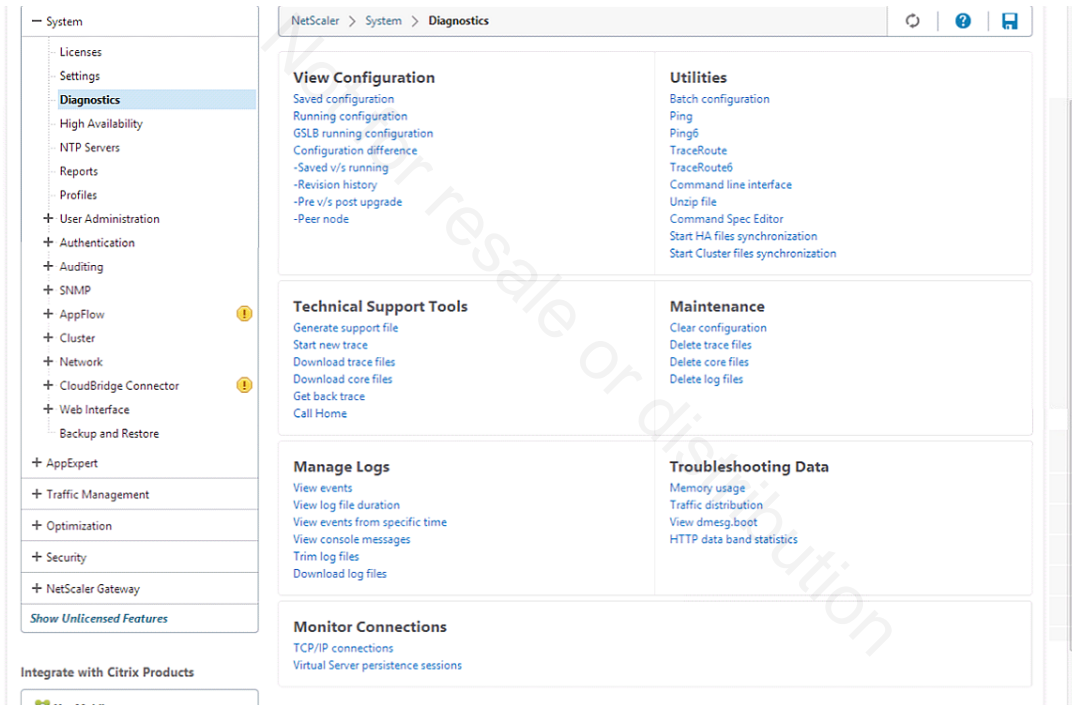
Troubleshooting Tools

NetScaler has a number of tools to help with troubleshooting. These include the following tools that are available through the BSD shell:

- nsconmsg
- nstrace

In addition, the configuration utility contains a diagnostic utility (located in System > Diagnostics) that provides access to the following tools:

- traceroute
- ping
- tcpdump



Newslog

The `/var/nslog/newslog` binary file stores console messages, events and performance statistics. This log can be accessed from the Diagnostics menu in the configuration utility.

Show TechSupport Command

When troubleshooting NetScaler, you can generate a set of files required for troubleshooting any issue related to Netscaler. A support file can also be generated from the configuration utility using the Generate support file link in the System > Diagnostics page under Technical Support Tools.

From the command-line interface, you can also run the following command:

```
show techsupport
```

Using the Visualizer

The network visualizer is a tool that you can use to view the network configuration of a NetScaler node, including the network configuration of the nodes in a high availability deployment. You can also modify the configuration of VLANs, interfaces, channels, bridge groups and perform high availability configuration tasks.

The visualizer can be a helpful tool for troubleshooting NetScaler. For example, when you are logged on to a standalone system, you can use the visualizer to view a consolidated graphical summary of key network components. You can also view the individual details of various network components, as well as view node details, node statistics and statistics for VLANs and interfaces. These capabilities can potentially help you find the source of a problem or issue in your environment.

For more information about the visualizer, see Citrix product documentation at <http://docs.citrix.com>.

Display NetScaler System Information

The model and serial numbers of a NetScaler system can be useful during troubleshooting.

To display the model and serial numbers, enter the following command in the command-line interface:

```
show hardware
```

This information can also be viewed in the system node of the configuration utility.

Troubleshooting Hardware Issues

Examples of hardware issues include:

- The interface is not detected.
- The interface is always flapping, or alternately showing as UP and DOWN.
- Continuous console messages refer to a hardware component.

- Lights do not appear on the appliance when it is turned on (assuming you have already tested your power sources and cables).
- Network interface LED lights do not appear when a network cable is inserted.
- Cannot connect to the console or get a logon prompt through a serial port (assuming you use the correct terminal emulation settings explained in the "Configuring the Application Switch Using the CLI" section of Citrix article CTX112893 at <http://support.citrix.com>).
- BIOS or POST errors
- Specific SSL card failures
- Output of the `dmesg` shell command indicates interface issues or read/write errors on the hard disk sector



Disk errors can result from file corruption, which can be corrected by running a file consistency check. Citrix Technical Support can help you run these procedures.

- You received an appliance with the wrong kind of interfaces, for example copper instead of fiber.

Additional data required to troubleshoot hardware issues include:

- The "newslog" file
- The "dmesg" shell command output
- The "show node" output from the NetScaler command-line interface if the issue is related to SSL card failures

In some cases, it might be possible to recover the unit from errors. Contact Citrix Technical Support if you are unable to resolve hardware issues and to obtain a Return Materials Authorization (RMA). For more information about NetScaler hardware issues, see Citrix articles CTX113462 and CTX109304 at <http://support.citrix.com>. In addition, refer to the Citrix Technical Support Brief Troubleshooting Guide at <http://support.citrix.com>.

Troubleshooting License Issues

You can view license issues that have been read by entering the following shell command:

```
cat /var/log/license.log
```

If a feature that you need licensed is listed, you must request a new license. If "YES" is listed next to a feature, the feature can be enabled by right-clicking the feature in the configuration utility.

An example of a licensing issue is an error message displaying a restriction such as "Certificate with key size greater than 512 bits not supported." This message displays because a license is not applied or a public NetScaler software release is not installed on the system. To resolve the issue, you can apply the appropriate license to the NetScaler system, or upgrade the software of the system to a public NetScaler software release, or do both. For more information about troubleshooting license issues, see Citrix article CTX125548 at <http://support.citrix.com>.

Non-Maskable Interrupt (NMI) Button

If the NetScaler is not responding, you can force a core dump and restart the physical appliance, using the NMI button. The core files can help you or Citrix Technical Support investigate why the NetScaler is not responding. The NMI button is the recessed red button on the back side of a physical NetScaler appliance.



The NMI button does not exist on a NetScaler VPX.

When you press this button, the appliances dumps troubleshooting data and restarts. The vmcore and kernel files are then available in the `/var/crash` directory. The process of dumping the data and restarting the appliance can take between 10 and 45 minutes, depending on the RAM of the appliance. For more information about the NMI button, see Citrix article CTX120660 at <http://support.citrix.com>.

Display Software Information

NetScaler systems are pre-loaded with software from the factory. When troubleshooting software issues, the following commands can be used to display the current software information, such as the version, feature, and license information:

```
show feature
```

```
show version
```

```
show license
```

Software information is also available in the configuration utility. The version is displayed in the top pane and the license information is located in the System > License menu.

Discussion Question

Besides the built-in NetScaler tools, which third party tools do you use in your environment?

Not for resale or distribution

Not for resale or distribution

Not for resale or distribution



851 West Cypress Creek Road Fort Lauderdale, FL 33309 USA (954) 267 3000 www.citrix.com

Rheinweg 9 8200 Schaffhausen Switzerland +41 (0) 52 63577 00 www.citrix.com

© Copyright 2016 Citrix Systems, Inc. All rights reserved.