Citrix NetScaler 1000V Release Notes

Citrix NetScaler 11.0-62.10
2015-09-04

Cisco Systems, Inc.
www.cisco.com

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

Contents

11.0-62.10

The release notes provides the changes or enhancements, issues that are fixed, and known issues that exist in Build 62.10. The list of known issues is cumulative, that is, it includes known issues that existed in previous builds and issues that are newly found in this build.

Release history:

- Build 62.10 (2015-08-12) (Current build)
- Build 55.20 (2015-06-30)

What's New?

The enhancement and changes that are available in Build 62.10.

AAA-TM

- **Multi-Factor (nFactor) Authentication**

  The NetScaler appliance now supports a new approach to configuring multi-factor authentication. With this approach, you can configure any number of authentication factors. You can also customize the login form as required.

  In NetScaler terminology, this feature is called "nFactor Authentication." For more information, see http://docs.citrix.com/en-us/netscaler/11/security/ns-aaa-app-trafc-wrapper-con-10/multi-factor-nfactor-authentication.html.

  [# 482250, 451913, 549966]

Admin Partitions

- **Partition Specific Load Balancing Parameters**

  When you update load balancing parameters in an admin partition, the updates now apply to that partition only. You can have different load balancing parameter settings in different partitions.

  Note:

  - In previous releases, any updates to these parameters were applied across all partitions, regardless of the partition in which the changes were made.

  - These parameters are set in the CLI by using the â€œset lb parameter commandâ€• or in the GUI by navigating to Traffic Management > Load Balancing.

  [# 563004]

Cluster

- **FTP Load Balancing Support on a Cluster**

FTP load balancing is now supported in a NetScaler cluster deployment.

[# 513612]

- **Web Interface on NetScaler (WIonNS) Support on a Cluster**

  WIonNS can now be configured on a NetScaler cluster deployment. To use WIonNS on a cluster, you must do the following:

  1. Make sure that the Java package and the WI package are installed in the same directory on all the cluster nodes.

  2. Create a load balancing virtual server that has persistency configured.

  3. Create services with IP addresses as the NSIP address of each of the cluster nodes that you want to serve WI traffic.

  4. Bind the services to the load balancing virtual server.

  Note: If you are using WIonNS over a VPN connection, make sure that the load balancing virtual server is set as WIHOME.

  [# 498295, 489463]

HDX Insight

- HDX Insight now supports displaying of Appflow records from Netscaler cluster.

  [# 525758]

SSL

- Support for Thales nShield(R) HSM

  All NetScaler MPX, and VPX appliances except the MPX 9700/10500/12500/15500 appliances now support the Thales nShield(R) Connect external Hardware Security Module (HSM). With a Thales HSM, the keys are securely stored as application key tokens on a remote file server and can be reconstituted only inside the Thales HSM. Thales HSMs comply with FIPS 140-2 Level 3 specifications.

  Note: Thales integration with the ADC is currently supported for TLS version 1.0.

  [# 440351, 477544]

Telco

- **Subscriber-Aware Traffic Steering**

  Traffic steering is directing subscriber traffic from one point to another based on subscriber information. When a subscriber connects to the network, the packet gateway associates an IP address with the subscriber and forwards the data packet to the NetScaler appliance. The appliance communicates with the PCRF server over the Gx interface to get the policy information. Based on the policy information, the appliance performs one of the following actions:

  - Forwards the data packet to another set of services

- Drops the packet

- Performs LSN if configured on the appliance

[# 402473]

- Provide Internet Access to a Large Number of Private IPv4 Subscribers of a Telecom Service Provider

  The Internet’s phenomenal growth has resulted in a shortage of public IPv4 addresses. Large Scale NAT (LSN/CGNAT) provides a solution to this issue, maximizing the use of available public IPv4 addresses by sharing a few public IPv4 addresses among a large pool of Internet users. LSN translates private IPv4 addresses into public IPv4 addresses. It includes network address and port translation methods to aggregate many private IP addresses into fewer public IPv4 addresses. LSN is designed to handle NAT on a large scale.

  The NetScaler supports LSN and is compliant with RFC 6888, 5382, 5508, and 4787. The NetScaler LSN feature is very useful for Internet Service Providers (ISPs) and carriers providing millions of translations to support a large number of users (subscribers) and at very high throughput. The LSN architecture of an ISP using Citrix products consists of subscribers (Internet users) in private address spaces accessing the Internet through a NetScaler appliance deployed in ISP’s core network.

  The following lists some of the LSN features supported on a NetScaler appliance:

  * ALGs: Support of application Layer Gateway (ALG) for SIP, PPTP, RTSP, FTP, ICMP, and TFTP protocols.

  * Deterministic/ Fixed NAT: Support for pre-allocation of block of ports to subscribers for minimizing logging.

  * Mapping: Support of Endpoint-independent mapping (EIM), Address-dependent mapping (ADM), and Address-Port dependent mapping.

  * Filtering: Support of Endpoint-independent filtering (EIF), Address-dependent filtering, and Address-Port-dependent filtering.

  * Quotas: Configurable limits on number of ports and sessions per subscriber.

  * Static Mapping: Support of manually defining an LSN mapping.

  * Hairpin Flow: Support for communication between subscribers or internal hosts using public IP addresses.

  * LSN Clients: Support for specifying or identifying subscribers for LSN NAT by using IPv4 addresses and extended ACL rules.

  * Logging: Support for logging LSN session for law enforcement. In addition, the following are also supported for logging:

  ** Reliable SYSLOG: Support of sending SYSLOG messages over TCP to external log servers for a more reliable transport mechanism.

** Load balancing of Log Servers. Support for load balancing of external log servers for preventing storage of redundant log messages.

** Minimal Logging: Deterministic LSN configurations or Dynamic LSN configurations with port block significantly reduces the LSN log volume.

[# 316909]

- **Provide Internet Access to IPv4 Subscribers Through the IPv6 Core Network of a Telecom Service Provider**

  Because of the shortage of IPv4 addresses, and the advantages of IPv6 over IPv4, many ISPs have started transitioning to IPv6 infrastructure. But during this transitioning, ISPs must continue to support IPv4 along with IPv6 because most of the public Internet still uses only IPv4, and many subscribers do not support IPv6.

  Dual Stack Lite (DS Lite) is an IPv6 transition solution for ISPs with IPv6 infrastructure to connect their IPv4 subscribers to the Internet. DS Lite uses IPv6 tunneling to send a subscriberâ€™s IPv4 packet over a tunnel on the IPv6 access network to the ISP. The IPv6 packet is de capsulated to recover the subscriberâ€™s IPv4 packet and is then sent to the Internet after NAT address and port translation other LSN related processing. The response packets traverse through the same path to the subscriber.

  [# 407162]

- **Subscriber-Aware Service Chaining**

  Service chaining is determining the set of services through which the outbound traffic from a subscriber must pass before going to the Internet. Multiple services, such as antivirus services, parental control services, firewalls, and web filter, are running in a Telco network. Different subscribers have different plans and each plan has specific services associated with it. The decision to direct a subscriberâ€™s request to a service is based on the subscriber information. Instead of sending all the traffic to all the services, the NetScaler appliance intelligently routes all requests from a subscriber to a specific set of services on the basis of the policy defined for that subscriber. The appliance receives the subscriber information from the PCRF over a Gx interface.

  [# 561747]

- **Support for Gx Interface**

  The NetScaler appliance can now dynamically receive the subscriber information over a Gx interface. The appliance communicates with the PCRF server over the Gx interface, receives the subscriber information, and uses this information to direct the flow of traffic. The PCRF server can send updates over this interface at any point during the subscriber session.

  [# 402469]

- **Support for RADIUS Accounting Message**

  The NetScaler appliance can now dynamically receive the subscriber information through a RADIUS accounting message. It receives the subscriber IP address and MSISDN and uses this information to retrieve the subscriber rules from the PCRF server.

[# 526981]

- **Provide Visibility into SLA Reports**

  An ISP often purchases international bandwidth from upstream ISPs, who then become layer 2 ISPs. To provide the redundancy required for reliable service to its customers, the purchasing ISP negotiates Service Level Agreements with multiple layer 2 ISPs. The SLAs stipulate a penalty in the event that the layer 2 ISP fails to maintain a specified level of service.

  NetScaler Insight Center and the NetScaler cache redirection feature can now be used to monitor the traffic flowing through the NetScaler appliances and calculate SLA breaches. The NetScaler cache redirection feature helps save bandwidth over international links. NetScaler Insight Center works with the NetScaler cache redirection feature to calculate, and provide visibility into, the percentage of bandwidth saved and any breaches of the SLA. ISP administrators are alerted whenever there is a breach for response time, hit rate/sec, or bandwidth.

  For a specific domain, NetScaler calculates the following SLA breaches and forwards the data to NetScaler Insight Center:

  * SLA Breach. A breach that occurs when a metric (response time, hits, or bandwidth) crosses the defined threshold value. For example, SLA breach is considered if the response time for a specific domain crosses 100 ms.

  * SLA Breach Duration. Time period in which a SLA breach lasted. For example, SLA Breach Duration is considered 5 mins, if the response time for a domain is greater than 100 ms consistently for 5 mins.

  * Breached Request Percentage. Percentage of requests whose response time is not within the minimum response time and maximum response time range. For example, if you configure this value as 10%, then among 100 requests, the response time of 10 requests are not within the minimum response time and maximum response time.

  NetScaler Insight Center then calculates the following SLA breaches:

  * SLA Breach Frequency- SLA Breach Frequency is the defined as the number of times the SLA breach occurs for the SLA Breach Duration. For example, SLA Breach Frequency is considered 1, if the response time for a domain is greater than 100 ms consistently for 5 mins.

  All of these metrics are calculated for a SLA group, which contains a list of domains defined by the ISP administrator.

  [# 495288, 501269, 501277, 501278, 501279, 501280]

Bug Fixes

The issues that are addressed in Build 62.10.

Application Firewall

- After processing a request that consists of multiple headers of the same type, a subsequent request might invoke a 302 response due to the way the application firewall stores the information regarding the parsed headers. With this fix, the variable which stores the information regarding the headers is reinitialized accurately prior to processing the next request.

  [# 580564]

- When processing a form for response side security check inspection, if the application firewall resets a connection, the partially parsed form is not freed leading to memory leak. With this fix, the memory allocated to the partially parsed forms is freed when a connection is reset.

  [# 572637, 581520]

Cluster

- You cannot add LB routes in a link load balancing setup that is deployed on a cluster.

  [# 574717]

- In a NetScaler cluster, a "sh nslogaction" command that is issued from the NSIP address of a cluster node, goes into an infinite loop. The issue is not observed when the command is issued from the cluster IP address.

  [# 574333, 573645]

- In a cluster setup, for active FTP, the server cannot initiate a data connection from a random port.

  [# 559230, 571042]

Configuration Utility

- The operation to download the nstrace file from the configuration utility fails.

  [# 571814, 581955]

NetScaler Insight Center

- The NetScaler Insight Center appliance might fail and not respond, when you add, update, or delete the private IP address block that is used for geo location.

  [# 576477, 581927]

- If there are more than 25 records to display in the skip flow window, then only 25 records are displayed as the window does not provide support for pagination.

  [# 576471]

Netscaler Insight Center

- **Media Classification Support for Insight Center**

  Web Insight supports content and media type classification reports. Viewing these features are optional similar to the existing HTTP header fields User Agents, Operating Systems, Request Methods etc. You can enable or disable these features from the Configuration

section. For media classification and httpContentType Appflow parameter, you must first enable Appflow on virtual server from Insight center configuration.

Insight Center's Web Insight dash board reports the following Media types:

1) Uncategorized

2) FLV F4V Audio

3) FLV F4V Video

4) MP4 M4V Audio

5) MP4 M4V Video

6) GP 3G2 Video

7) ADTS Audio

8) APPLE Video

9) MICROSOFT Video

10) AAC Audio

11) MICROSOFT PLAYLIST Video

12) APPLE PLAYLIST Video

13) MP3 Audio

14) Unknown

[# 558890]

Networking

- An ACL6 rule might not get evaluated if you set the operator option to NEQ (!=) for source and destination IPv6 addresses.

  [# 573516]

- High availability (HA) synchronization fails if the NetScaler IP (NSIP) addresses of the nodes in the HA configuration are IPv6 addresses.

  [# 573935]

- Duplicate address detection might fail for a global IPv6 address.

  [# 560243]

- ICMPv6 requests with a payload greater than 1232 bytes (fragmented ICMPv6 requests) from a nondefault NetScaler admin partition might not succeed.

  [# 506332]

- A PBR6 rule might not get evaluated if you set the operator option to NEQ (!=) for source and destination IPv6 addresses.

  [# 575906]

SSL

- On a NetScaler MPX appliance, AES-GCM/SHA2 ciphers are supported only on the front end SSL entities.

  [# 575001]

System

- The option to set the transport type has been removed from the SET and UNSET operations. You can specify the transport type while adding a Syslog action. In a Syslog action, by default the transport type is set as UDP.

  Note: Once you have set the transport type in a Syslog action, you cannot change the transport type.

  [# 580890]

- The NetScaler appliance fails intermittently when trace is started in 'RX' mode.

  [# 576067]

Web Interface on NetScaler (WIonNS)

- After upgrading to nswi-1.8.tgz, existing WI sites are not accessible till you remove the sites and then add them back.

  [# 576883]

Known Issues

The issues that exist in Build 62.10.

AAA-TM

- When IBM Tivoli IdP is used for SAML authentication with NetScaler appliance as the service provider, there could be an issue with SAML assertion verification.

  [# 540396]

- The NetScaler implementation of Kerberos does not fully implement the ktutil functionality. While this does not affect Kerberos authentication, it restricts some administrative tasks, such as the ability to merge keytab files.

  [# 551091]

- The status of a LDAP server on the authentication dashboard of the NetScaler GUI, will be shown as UP, regardless of the actual status of the LDAP server, for the following combinations:

  - Security type is SSL and port is 389.

- Security type is TLS or PLAINTEXT and port is 636.

[# 567376, 567379]

Admin Partitions

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

  [# 489512]

- RPCSVR services cannot be configured in admin partitions.

  [# 498477]

- Admin partitions are not supported on FIPS appliances. However, owing to this issue, you can create admin partitions on FIPS appliances. You are advised against creating such partitions as they will not function properly.

  [# 517145]

- The IC memory once set for an admin partition, cannot be reduced. An appropriate error message is displayed.

  For example, if the IC memory of admin partition is 10 GB, you cannot reduce it to 8 GB. The memory limit can however be increased to a required value.

  [# 568106, 570578]

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

  [# 493668, 516396]

AppFlow Insight

- Hiding or displaying a URL, and some configuration changes might take longer than expected.

  [# 570896, 574278]

Application Firewall

- If a user request triggers an application firewall policy that is bound to the APPFW_BYPASS profile, the application firewall might fail to generate an SNMP alarm.

  [# 489691]

- If the server sends less data than the amount specified in the Content-length header, the NetScaler application firewall might send a 9845 response and reset the connection.

  [# 506653]

- In NetScaler 9.3, if there is a standalone application firewall license, the user is able to bind a classic application firewall policy to the load balancing virtual server. However, in NetScaler 10.1, the design is changed. If the load balancing feature is not licensed, binding a classic

application firewall policy to the load balancing virtual server now results in an error message.

[# 510509]

- The customer's application does not work when the application firewall is deployed to inspect the request for security check violations. When the application firewall forwards the request to the backend server, the server responds with a 403 HTTP error code, indicating that it cannot properly validate the CORBA session, and sends the page without the expected data in the form fields. The root cause is under investigation.

  Workaround: Turn off form field tagging and credit card checks.

  [# 511254]

- On a NetScaler appliance that has standalone application firewall license, when you bind a classic application firewall policy to a load balancing virtual server, an error message is displayed in the graphical user interface. The binding operation is successful. The error message is harmless and can be safely ignored.

  [# 522712]

- The Skip operation for the application firewall learned rules might take longer than expected.

  [# 547978]

- The application firewall learning engine is not able to connect to the packet engine in certain circumstances. When this happens, the aslearn process does not start and the application firewall learning functionality stops working.

  [# 576713, 582879]

- "Operation timed out" error is displayed in the CLI and the configuration utility while viewing learned rules. This error is only seen intermittently.

  [# 527190]

- A POST request with an attached word document is silently blocked by the application firewall for a customized application.

  [# 530277]

- The Graphical User Interface (GUI) for the NetScaler application firewall has significantly changed to provide enhanced user experience and remove browser plugin dependencies. The GUI steps in the current application firewall documents are in need of revision. Some of them do not match the new GUI display.

  [# 548432]

- The application firewall Graphical User Interface might display a warning when the Qualys signature file is uploaded to the NetScaler appliance. The transformation program that reads the input file is treating a warning message as an error.

  [# 547282]

- In the configuration utility (GUI), selecting the "Remove All Learned Data" action in the application firewall Learned Rules section might not remove the learned data for some of the security checks for the profile.

  [# 549255]

- When a user-defined application firewall signature object is updated by using the configuration utility, the enabled rules might get disabled and the configured actions in some signature rules might not be preserved.

  [# 561567]

- Application firewall memory allocation failures might occur, when the integrated cache is also enabled and the memory usage limit for the cache parameter is set to a high value.

  [# 567119, 568260]

- The application firewall learning engine stops recommending new rules when the learning database grows to approximately 20-22 megabytes in size. The database size limit is applied on a per profile basis.

  [# 554591]

- When a NetScaler appliance is upgraded from a 10.1 build to a 10.5 build, the application firewall signature names are converted to all lowercase characters. If the name of the signature contains any uppercase character, the conversion affects the binding between profile and signature. Any attempt to modify either the profile or the signature object displays an error message in the configuration utility.

  [# 568705]

- During an upgrade of a NetScaler appliance from version 10.0 to version 10.1 (build 121.1 or subsequent), the default JSON content type is not automatically configured. The default JSON content type is configured when version 10.1 (build 121.1) is installed on new hardware or in a new VPX instance. To check whether your appliance or instance has the correct default setting, log onto the NetScaler command line and type the following command:

  show appfw JSONContentType

  If the default content type is configured, the command output is similar to the following example:

  > show appfw JSONContentType

  1) JSONContenttypevalue: "^application/json$" IsRegex: REGEX

  Done

  If it is not, the screen shows only the following:

  > show appfw JSONContentType

  Done

To add the default content type to the configuration, after upgrading to 10.1 (121.1), log onto the NetScaler command line, and then type the following commands to configure the default content type and verify the configuration:

add appfw JSONContentType ^application/json$ -isRegex REGEX

show appfw JSONContentType

[# 430014]

Cisco RISE Integration

- Cisco RISE now supports the following commands:

  - show rise param

  - set rise param

  Following is the usage of the set rise param command:

  set rise param [-directMode (ENABLED | DISABLED)] [-indirectMode (ENABLED | DISABLED)]

  The show rise param command displays the current setting. For example,

  RISE-MPX-194-80> show rise param

  DirectMode: ENABLED IndirectMode: ENABLED

  Done

  [# 497410]

Cluster

- When WIonNS is deployed in a cluster setup, an error is thrown if you change the IP address of the WI service to point to the IP address of the cluster configuration coordinator.

  [# 582801]

- When L2 mode and MBF is enabled in a cluster deployment, access to * 80 services can fail intermittently.

  [# 479899]

- In a cluster setup, a command that is executed on the cluster configuration coordinator is propagated to the other cluster nodes. Therefore, a command that takes a long time to complete (such as "save ns config"), can take a little extra time to complete on all the cluster nodes. During this time, if you execute another command on the cluster (through another session), that command will fail because the previous command is not yet complete.

  [# 551607, 495270, 562651]

- When a cluster is connected to more than one upstream router:

- When AS OVERRIDE is not configured on the upstream router, spare nodes will learn VIP routes from one of the routers, but they will be dropped as the path contains its own AS to prevent loop formation.

  - When AS OVERRIDE is configured on any upstream router for cluster neighbors, upstream router will change AS path in VIP to its own AS while sending updates to cluster neighbors. Spare nodes will not detect any loop and learnt VIP routes are advertised to other routers.

  Spare nodes will not advertise their configured VIP routes but there is no such restriction on BGP learnt routes.

  [# 547749]

- When WIonNS is deployed in a cluster setup, if the service IP address is modified using the "set" command, the "show" command continues to display the previous IP address.

  [# 582805]

- When WIonNS is deployed in a cluster setup, an error is thrown when you rename a service that points to the IP address of the cluster configuration coordinator.

  [# 583424]

- When a node is removed from a L3 cluster, IPv6 SNIP addresses and routes are being erroneously cleared from the appliance. This behavior is seen only for IPv6 entities. IPv4 SNIPs and routes are not being removed from the appliance.

  [# 542693]

- When WIonNS is deployed in a cluster setup, if you add a service that points to the NSIP of a newly joined node, the command fails on the newly joined node but succeeds on the other cluster nodes.

  [# 584699]

Command Line Interface

- The NetScaler command line interface exits abruptly upon executing the "show dns addRec -format old" command.

  [# 512526, 527066, 545578]

Configuration Utility

- An interface does not appear as tagged or untagged in the network visualizer.

  [# 540980]

- In the NetScaler configuration utility, the page at System> Network > IPs does not display the Type for LSN NATIPs, and the value shown for Traffic Domain is incorrect.

  Workaround: Run the sh nsip command to display the values in the command line interface.

  [# 505121]

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

  [# 498004]

- If you click a VLAN in the network visualizer, details such as VLAN ID and bound interfaces are not displayed in a separate pane.

  [# 540943]

- You cannot upgrade to NetScaler 11 from the following builds by using the Upgrade Wizard of the NetScaler GUI:

  - All builds of NetScaler 9.3

  - All builds of NetScaler 10.1

  - Any build before Build 57.x of NetScaler 10.5

  Workaround: Use the command line interface to upgrade the NetScaler appliance.

  [# 563410]

- The subnet mask does not appear after an IPv4 address in the network visualizer.

  [# 540927]

- In the network visualizer, if you click a tagged interface that is part of two or more VLANs, only the VLAN at the top of the list of bound VLANs is highlighted.

  [# 541011]

- The bridge group and VLAN association is not displayed in the network visualizer.

  [# 542214]

GSLB

- If you rename a server associated with a GSLB service and then run the sync gslb command, the GSLB configuration might not synchronize with the other GSLB sites.

  Workaround: Manually update the server name in the other GSLB sites.

  [# 511994]

Load Balancing

- The appliance fails if non-reachable autoscale entities that are part of a service group later become reachable and, in the interim, the service group name has changed.

  [# 583647]

- IPV6 addresses are trimmed when data is retrieved from the packet engine because the prefix length variable is unset during the GET operation.

[# 573463]

- If the state of the IPv6 service on which a client's persistent session is running changes to out-of-service, the session might lose persistence before the client's transaction is completed.

  [# 571771]

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

  [# 511222]

- If a NetScaler appliance sending a DNSSEC negative response over UDP is not able to include the required records (for example, SOA, NSECs, and RRSIG records) in the Authority section, the appliance might send a truncated response in the wrong packet format.

  [# 540965]

NetScaler Insight Center

- Any port other than 1494 and 2598, that needs to be considered as an ICA or CGP port, needs to be explicitly configured as a global ICA port to get the HDX Insight LAN user configuration working.

  [# 530702]

- If you use the refresh button, it does not have any effect on the slider. Refresh operation does not have any affect on the time shown in the slider. Also, when you change tabs, it does not impact the slider. You can change the time by changing the time duration.

  [# 576469]

- If you have configured the ICA session timeout value to a high value, say 10 minutes or more, and there is no traffic flow from the NetScaler appliances, neither the timeline chart nor the tabular chart displays any data. However, the Active sessions and Active Desktops columns display the data until the ICA session timeout occurs.

  [# 536056]

- If the ICA Rtt column is the column in extreme left of the session details table, the pop-up box gets cropped in display.

  [# 573089]

- Adding a new data node is now driven by Auto Registration. When a kernel is imported, it requests for input from user and does an auto registration with the Insight Server. This allows the Insight Deployment Manager GUI to display the same. Removing a datanode is not presently supported.

  [# 543632, 565706, 567628, 570264]

- Insight Agent should only be added after configuring and deploying Insight DB Cluster.

  [# 570619]

- Geo report is only available for daily, weekly, and monthly reports for Web Insight.

  [# 556534]

NetScaler VPX Appliance

- A NetScaler that is deployed on the Hyper-V may crash or unexpectedly reboot if it uses three or more virtual interfaces in the VPX instance.

  [# 467734, 469552, 471601, 476833, 484210, 489880]

Networking

- In an active-active high availability configuration using Virtual Router Redundancy Protocol (VRRP) protocol, a ping to a virtual IP address (VIP) might fail from a node that is a backup node for this VIP address.

  [# 485260]

- For an RNAT connection, the NetScaler appliance drops the first packet that the server sends to the client.

  [# 543171]

- After the clear config operation, reconfiguring a VXLAN entity fails to retrieve the VXLAN SNMP counters.

  [# 572525, 574734]

- If you configure an INAT rule with the useproxyport parameter disabled, connections to the server fail if the source port is in the reserve port range (0-1023).

  [# 550488]

- A clear config operation does not remove VXLANS. The configuration utility and the CLI continue to show the VXLANs, but with incorrect IDs.

  [# 574734]

- A TCP connection involved in INAT times out at 120 seconds, regardless of what global timeout value you set for TCP client and server connections. For example, the connection times out at 120 seconds even after you run the following command:

  set ns timeout -anyTcpClient 50 -anyTcpServer 50

  [# 569874]

- RNAT source IP persistency is not supported on a virtual server configured for link load balancing.

  [# 546066]

- If you add an NTP time server by specifying the server name (host name), and the ns.conf file is very large, the result is a race condition in which the NTP daemon (NTPD) is started before host name services are ready.

  Workaround: Do one of the following:

-Restart the NTP daemon after starting the NetScaler appliance.

-Add the NTP server by specifying the IP address of the server instead of specifying the host name.

[# 573306]

Policy

- Some IP based expressions on the NetScaler appliance may not work for the IP addresses starting from octet 128 or greater (128.x.x.x - 254.x.x.x).

  [# 534244]

SSL

- If importing a certificate-key file fails because of a wrong file, and you run the command again with the correct file, the operation fails and the following error message appears:

  "ERROR: Import failed. Another resource with the same name being processed"

  Workaround: Import the file with a different name.

  [# 526433]

- If you disable SSLv3 on the "nskrpcs-127.0.0.1-3009" service, an "ERROR: Operation not permitted" message appears even though SSLv3 has been successfully disabled on the service.

  [# 521569]

- Secure renegotiation using SSLv3 protocol fails on MPX-FIPS appliances running firmware version 2.2.

  [# 550788]

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

  [# 512334]

- Server Name Indication (SNI) is not supported on a DTLS virtual server. However, if you enable SNI on a DTLS virtual server, an appropriate error message does not appear.

  [# 572429]

- If you try to add a certificate bundle with the complete path to a certificate-bundle file, an error message appears. For example,

  > add ssl certkey bundle -cert /nsconfig/ssl/bundle3.pem -key /nsconfig/ssl/bundle3.pem -bundle YES

  ERROR: Processing of certificate bundle file failed.

  Workaround: Specify only the file name. For example,

> add ssl certkey bundle -cert bundle3.pem -key /nsconfig/ssl/bundle3.pem -bundle YES

[# 481878, 521933]

- Even though the clientAuthUseBoundCAChain parameter can be enabled and disabled in the backend profile, it is supported only on the front end profile.

  [# 554782]

- Even though TLS protocol versions 1.1 and 1.2 are not supported by firmware version 1.1, the protocols incorrectly appear as enabled by default on an SSL virtual server.

  Workaround: Disable TLS1.1/1.2 explicitly on the virtual server.

  [# 576274]

- FIPS keys that are created on firmware version 2.2 are lost after you downgrade to firmware version 1.1.

  Workaround: Export the FIPS keys before you downgrade the firmware. Import the FIPS keys after the downgrade.

  [# 559796]

System

- In rare circumstances, the VPX instance can dump kernel core after a warm restart.

  [# 559176]

- The initial client connection on the NetScaler appliance might fail if a wildcard virtual server is configured and the useProxyPort option is disabled globally on the appliance.

  [# 542776, 571357]

- FTP connections through a TCP wildcard virtual server on the NetScaler appliance might fail for one of the following reasons:

  - A mismatch in TCP parameters is preventing the appliance from reusing the probe connection.

  - The server is sending data before the client-side TCP connection is established.

  [# 545858]

- If the HTML injection feature is enabled, the NetScaler appliance injects JavaScript into responses sent to clients. If a subsequent request from one of the clients is generated from the JavaScript, the appliance responds with a 404 error.

  [# 518272]

- In a high availability setup, if stateful connection failover is configured on a virtual server that is serving traffic for some time, running the "clear config extended" command results in a warm restart on both the primary and secondary appliances. Unsetting connection failover on the virtual server results is a warm restart on only the secondary appliance.

[# 575108]

Telco

- With a large number of active subscribers, and a high traffic rate for SIP over TCP, the NetScaler appliance can fail during ALG processing.

  [# 582464]

- Where there are over 140K SIP calls over UDP, the NetScaler appliance can fail during ALG processing.

  [# 574303, 582451]

- In the output of the â€œshow lsn sipalgcall -callidâ€• command, the port value of the SIP control channel is incorrect.

  [# 574257]

- In a Large Scale NAT deployment, the NetScaler appliance does not generate and send an ICMP error message to the subscriber in the event of a port allocation failure.

  [# 540162]

- If the provisional response to a SIP REGISTER message does not contain an expiry value, the NetScaler appliance drops the message.

  [# 574725]

- SIP registration might fail, if authentication is enabled in the SIP proxy server.

  [# 579797]

- An RTSP request might be logged on two different Syslog servers.

  [# 581086]

- After a failover occurs in a high availability configuration, some LSN static maps might become inactive on the new secondary node.

  Workaround: Delete the LSN static maps on the primary node and then add them again.

  [# 487318]

WAN Insight

- On the NetScaler Insight Center dashboard, the latency values displayed on the graph and the network topology diagram might not match due to time synchronization issues.

  [# 533063]

- NetScaler Insight Center takes two minutes to display the current connection details on the dashboard.

  [# 536696]

- CSV report exports elements that are present in the GUI. Additional elements like Client IP and Branch IP in the application node are denoted as 0.0.0.0 or " " as these are not present in GUI.

  [# 547380]

- If NetScaler Insight Center does not get a connection closure update for a particular connection ID, and the ID is reused, the IP data of the previous connection may be displayed.

  [# 549679]

- NetScaler Insight Center displays the latency value between two hops as 0 ms, though the minimum latency value is 1 ms.

  [# 553536]

- If you upgrade NetScaler Insight Center appliance to release 10.5 build 55.8xxx.e, the compression ratio values will be displayed as -NA-.

  [# 554960]

WIoNS

- Since the install wi package command takes more than usual time to complete, it is not possible to return the status from other nodes. Hence it is required that all the WI related packages, that is, JRE+WI be present on system on the same path for all the nodes.

  [# 507753]

Web Interface

- **OpenJDK version for Web Interface on NetScaler (WIonNS)**

  For NetScaler 10.5 and later releases, Web Interface on NetScaler (WIonNS) must use the OpenJDK7 package since NetScaler now uses FreeBSD 8.x/amd64. You can download the package from either one of the following links:

  *
  http://ftp.freebsd.org/pub/FreeBSD/releases/amd64/amd64/8.4-RELEASE/packages/java/openjdk-7.17.02_2.tbz

  *
  ftp://mirror.is.co.za/FreeBSD/ports/amd64/packages-8.4-release/devel/openjdk-7.17.02_2.tbz

  Background: When the NetScaler is upgraded to version 10.5, it still has OpenJDK1.6 instead of OpenJDK1.7 which is required for NetScaler version 10.5. Therefore, when the configurations are saved (after upgrading), the Web Interface sites become inaccessible.

  Workaround: Before you save the configurations on the upgraded appliance, make sure you reinstall the Web Interface on NetScaler version 10.5 by using OpenJDK1.7.

  [# 464854]

Web Interface on NetScaler (WIonNS)

- WIonNS v1.7 does not work when WebFront is installed.

Workaround: Upgrade to WIonNS v1.8.

[# 577988]

vPath

- In a cluster environment, vPath encapsulation may fail when MAC based forwarding is enabled.

  [# 580137]

- The first packet which is off-loaded from NS1000V is dropped by VEM.

  [#549254]

What's New in Previous NetScaler 11.0 Releases

The enhancements and changes that were available in NetScaler 11.0 releases prior to Build 62.10. The build number provided below the issue description indicates the build in which this enhancement or change was provided.

AAA-TM

- **Fallback to NTLM Authentication**

  When the NetScaler appliance is configured for Negotiate authentication and sends a 401 Negotiate response to client, if client is not able to reach domain controller or is not domain joined, then it automatically falls back to NTLM authentication and the client starts NTLM handshake. The NetScaler appliance is able to verify the credentials presented as part of NTLM authentication.

  This feature allows user logins locally or remotely.

  [From Build 55.20] [# 509829]

- **Support for Redirect Binding for SAML SP**

  When used as a SAML SP (service provider), in addition to POST bindings, the NetScaler appliance now supports redirect bindings. In redirect bindings, SAML assertions are in the URL, as against POST bindings where the assertions are in the POST body.

  Using the CLI:

  > add authentication samlAction <name> . . . [-samlBinding (REDIRECT | POST)]

  [From Build 55.20] [# 493220, 462777, 493224]

- **Encrypting SAML IdP Assertion**

  When used as a SAML IdP (identity provider), the NetScaler appliance can now be configured to encrypt the assertions by using the public key of the SAML SP (service provider).

Note:

- Make sure the SAML SP certificate is specified.

- For enhanced security, it is recommended that you encrypt assertions that contain sensitive information.

This configuration must be specified on the SAML IdP profile as follows:

On the CLI:

> set authentication samlIdPProfile <name> [-encryptAssertion (ON | OFF)] [-encryptionAlgorithm <encryptionAlgorithm>]

On the GUI:

Navigate to the screen where you configure the SAML IdP profile and specify the corresponding parameters.

[From Build 55.20] [# 482185]

- The NetScaler appliance now supports the SiteMinder SAML SP.

  [From Build 55.20] [# 488077]

- **Logging Errors in NetScaler Log Files**

  The NetScaler appliance now stores AAA authentication logs.

  - Errors and warnings are logged in the /var/nslog/ns.log file

  - Information and debug level logs are logged in the /var/log/nsvpn.log file.

  [From Build 55.20] [# 482228, 479557]

- **Using 401-based Authentication to Log on to a SAML IdP**

  When used as a SAML IdP (identity provider), the NetScaler appliance now allows logon using the following 401-based authentication mechanisms: Negotiate, NTLM, and Certificate.

  [From Build 55.20] [# 496725, 508689]

- **OAuth/OpenID-Connect Mechanisms for AAA-TM**

  The NetScaler AAA-TM feature now supports OAuth and OpenID-Connect mechanisms for authenticating and authorizing users to applications that are hosted on applications such as Google, Facebook, and Twitter.

  Note: OAuth on NetScaler is currently qualified only for Google applications.

  A major advantage is that user's information is not sent to the hosted applications and therefore the risk of identity theft is considerably reduced.

  In the NetScaler implementation, the application to be accessed is represented by the AAA-TM virtual server. So, to configure OAuth, an action must be configured and

associated with a AAA-TM policy which is then associated with a AAA-TM virtual server. The configuration to define a OAuth action is as follows:

> add authentication OAuthAction <name> -authorizationEndpoint <URL> -tokenEndpoint <URL> [-idtokenDecryptEndpoint <URL>] -clientID <string> -clientSecret <string> [-defaultAuthenticationGroup <string>] [-Attribute1 <string>] [-Attribute2 <string>] [-Attribute3 <string>] ....

Note:

- Refer to the man page for information on the parameters.

- Attributes (1 to 16) are attributes that can be extracted in OAuth response. Currently, these are not evaluated. They are added for future reference.

[From Build 55.20] [# 491920]

- **Using the SHA256 Algorithm to Sign SAML IdP Assertions**

  When used as a SAML IdP (identity provider), the NetScaler appliance can now be configured to digitally sign assertions by using the SHA256 algorithm. Additionally, you can configure the appliance to accept only digitally signed requests from the SAML SP (service provider).

  These configurations must be specified in the SAML IdP profile as follows:

  From the CLI:

  > set authentication samlIdPProfile <name> [-rejectUnsignedRequests (ON | OFF)] [-signatureAlg (RSA-SHA1 | RSA-SHA256)] [-digestMethod (SHA1 | SHA256)]

  From the GUI:

  Navigate to the screen where you configure the SAML IdP profile, and specify the corresponding parameters.

  [From Build 55.20] [# 474977]

- **Fallback from Certificate to Other Authentication Mechanisms**

  When authentication is configured to be done by using certificates and then followed by LDAP or other authentication mechanisms, the following behavior holds true:

  - In previous releases: If certificate authentication fails (or was skipped), the other authentication mechanism is not processed.

  - From this release onwards: Even if certificate authentication is not done, the other authentication mechanism is processed.

  [From Build 55.20] [# 550946]

- **Using Cookies to Track SAML Sessions**

  In a deployment where a NetScaler appliance is configured as a SAML IdP (identity provider) for multiple SAML SPs (service provider), the appliance allows a user to access

multiple SPs without explicitly authenticating every time.The appliance creates a session cookie for the first authentication and every subsequent request uses this cookie for authentication.

[From Build 55.20] [# 503882]

- **Using Certificates to Log on to a SAML IdP**

  When used as a SAML IdP (identity provider), the NetScaler appliance now allows logon using certificates.

  [From Build 55.20] [# 512125]

- **Including Additional Attributes in SAML IdP Assertion**

  When used as a SAML IdP (identity provider), the NetScaler appliance can now be configured to send 16 additional attributes in addition to the NameId attribute. These attributes must be extracted from the appropriate authentication server. For each of them, you can specify the name, the expression, the format, and a friendly name.

  These attributes must be specified in the SAML IdP profile as follows:

  From the CLI:

  > set authentication samlIdPProfile <name> [-Attribute1 <string> -Attribute1Expr <string> [-Attribute1FriendlyName <string>] [-Attribute1Format (URI | Basic)]] [-Attribute2 <string> -Attribute2Expr <string> [-Attribute2FriendlyName <string>] [-Attribute2Format (URI | Basic)]]

  For example, the following command adds the attribute "MyName":

  > add authentication samlIdPProfile ns-saml-idp -samlSPCertName nssp -samlIdPCertName nssp -assertionConsumerServiceURL "http://nssp.nsi-test.com/cgi/samlauth" -Attribute1 MyName -Attribute1Expr http.req.user.name -Attribute1FriendlyName Username -Attribute1Format URI

  From the GUI:

  Navigate to the screen where you configure the SAML IdP profile, and specify the additional attributes as required.

  [From Build 55.20] [# 460680, 504703]

- **Supporting Encrypted Assertions on SAML SP**

  When used as a SAML SP (service provider), the NetScaler appliance can now decrypt the encrypted tokens that it receives from the a SAML IdP. No configuration is required on the NetScaler.

  [From Build 55.20] [# 291693]

- The configuration of a AAA-TM virtual server in the NetScaler GUI is simplified for ease of configuring the required authentication mechanism.

  [From Build 55.20] [# 524386]

- The output of "show ns ip" now also includes the aaadnatIp address.

  [From Build 55.20] [# 472912]

Admin Partitions

- **Getting Web Logs for Specific Partitions/Users**

  Using the NetScaler Web Logging (NSWL) client, the NetScaler can now retrieve the web logs for all the partitions with which the logged in user is associated. To view the partition for each log entry, customize the log format to include the %P option. You can then filter the logs to view the logs for a specific partition.

  [From Build 55.20] [# 534986]

- **Getting NetScaler Trace for Specific Partitions**

  You can now generate the NetScaler trace for a specific admin partition. To do so, you must access that admin partition and run the "nstrace" operation. The trace files for the admin partition will be stored in the /var/partitions/<partitionName>/nstrace/ directory.

  [From Build 55.20] [# 496937, 515294]

- Scriptable monitors can now be configured on the admin partitions that are available on a NetScaler appliance.

  [From Build 55.20] [# 535494]

- **Setting L2 and L3 parameters in Admin Partitions**

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

  - For L2 parameters that are set by using the "set L2Param" command, the following parameters can be updated only from the default partition, and their values are applicable to all the admin partitions: maxBridgeCollision, bdgSetting, garpOnVridIntf, garpReply, proxyArp, resetInterfaceOnHAfailover, and skip_proxying_bsd_traffic. The other L2 parameters can be updated in specific admin partitions, and their values are local to those partitions.

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

  [From Build 55.20] [# 513564]

- **Supporting Dynamic Routing in Admin Partitions**

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

  > set L3Param -dynamicRouting ENABLED

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

[From Build 55.20] [# 514848]

- **Configuring Integrated Caching on a Partitioned NetScaler**

  Integrated caching (IC) can now be configured for admin partitions. After defining the IC memory on the default partition, the superuser can configure the IC memory on each admin partition such that the total IC memory allocated to all admin partitions does not exceed the IC memory defined on the default partition. The memory that is not configured for the admin partitions remains available for the default partition.

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

  - Partition1: 4 GB

  - Partition2: 3 GB

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

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

  [From Build 55.20] [# 481444, 484618]

Application Firewall

- The field format rules specify the inputs that are allowed in the target form fields. You can also limit the minimum and the maximum allowed length for the inputs. The application firewall learning engine monitors the traffic and provides field format recommendations based on the observed values. If the initial field format learned rules are based on a small sample of data, a few non typical values might possibly result in a recommendation that is too lenient for the target field. Updates to the application firewall have now decoupled violations and learning for the field formats. The firewall learns the field formats regardless of the violations. The learning engine monitors and evaluates all the incoming new data points to recommend new rules. This allows fine tuning the configuration to specify optimal input formats with adequate min/max range values. If a rule has already been deployed for a field/URL combination, the GUI allows the user to update the field format. A dialog box asks for confirmation to replace the existing rule. If you are using the command line interface, you have to explicitly unbind the previous binding and then bind the new rule.

  [From Build 55.20] [# 450326, 483677, 513927]

- The NetScaler application firewall offers SQL/XSS security check protections to detect and block possible attacks against the applications. You now have much tighter security control when configuring SQL/XSS protections. Instead of deploying relaxation rules that completely bypass the security check inspection for a field, you now have an option to relax a specific subset of violation patterns. You can continue to inspect the relaxed field in the incoming requests to detect and block the rest of the SQL/XSS violation patterns. The commands used in relaxations and learning now have optional parameters for value type and value expression. You can specify whether the value expression is a regular expression or a literal string.

  Command Line Interface:

bind appfw profile <name> -SQLInjection <String> [isNameRegex (REGEX | NOTREGEX)] <formActionURL> [-location <location>] [-valueType (Keyword| SpecialString|Wildchar) [<valueExpression>][-isValueRegex (REGEX | NOTREGEX)]]

unbind appfw profile <name> -SQLInjection <String><formActionURL> [-location <location>][-valueType (Keyword|SpecialString|Wildchar) [<valueExpression>]]

bind appfw profile <name> -crossSiteScripting <String> [isNameRegex (REGEX | NOTREGEX)] <formActionURL> [-location <location>] [-valueType (Tag| Attribute|Pattern) [<valueExpression>][-isValueRegex (REGEX | NOTREGEX)]]

unbind appfw profile <name> -crossSiteScripting <String> <formActionURL> [-location <location>] [-valueType (Tag|Attribute|Pattern) [<valueExpression>]]

[From Build 55.20] [# 450324, 483683]

- The NetScaler application firewall module offers data leak prevention and supports credit card protection. It can examines the credit card numbers in the response and takes the specified action if a match is found. In some scenarios, it might be desirable to exclude a specific set of numbers from the credit card security check inspection. For example, server responses for some internet applications might include a string of digits that is not a credit card number but matches the pattern of a credit card number. These responses can trigger false positives and therefore get blocked by the application firewall's Credit Card security check. The application firewall now offers the ability to learn and deploy relaxations for the credit card numbers. The credit card relaxation rule provides the flexibility to exclude a specific string of numbers from the safe commerce check without compromising credit card security. These numbers are not examined in the responses even if the credit card check is ON.

  Examples of CLI Commands:

  1. Bind the credit card number to profile:

  bind appfw profile <profile-name> -creditCardNumber <any number/regex> "<url>"

  2. Unbind credit card number from profile:

  unbind appfw profile <profile-name> -creditCardNumber <credit card number> "<url>"

  3. Log: Enable Logging of credit card Numbers

  add appfw profile <profilename> - doSecureCreditCardLogging <ON/OFF>

  set appfw profile <profilename> - doSecureCreditCardLogging <ON/OFF>

  4. Learn:

  show appfw learningdata <profilename> creditCardNumber

  rm appfw learningdata <profilename> -creditcardNumber <credit card number> "<url>"

  export appfw learningdata <profilename> creditCardNumber

  [From Build 55.20] [# 383298]

- Geolocation, which identifies the geographic location from which requests originate, can help you configure the application firewall for the optimal level of security. For example, if an excessively large number of requests are received from a specific area, it is easy to determine whether they are being sent by users or a rogue machine. The application firewall offers you the convenience of using the built-in NetScaler database or any other geolocation based database to identify the source of origin of coordinated attacks launched from a country. This information can be quite useful for enforcing the optimal level of security for your application to block malicious requests originating from a specific geographical region. Geolocation logging uses the Common Event Format (CEF).

  To use Geolocation Logging

  1. Enable CEFLogging and GeoLocationLogging.

  >set appfw settings GeoLocationLogging ON CEFLogging ON

  2. Specify the database

  >add locationfile /var/netscaler/inbuilt_db/Citrix_Netscaler_InBuilt_GeoIP_DB.csv

  or

  add locationfile <path to database file>

  [From Build 55.20] [# 483703]

- The application firewall is fully supported in striped, partially striped, or spotted configurations. The two main advantages of striped and partially striped virtual server support in cluster configurations are the following:

  - Session failover support: Striped and partially striped virtual server configurations support session failover. The advanced application firewall security features, such as Start URL Closure and the Form Field Consistency check, maintain and use sessions during transaction processing. In ordinary high availability configurations, or in spotted cluster configurations, when the node that is processing the application firewall traffic fails, all the session information is lost and the user has to reestablish the session. In striped virtual server configurations, user sessions are replicated across multiple nodes. If a node goes down, a node running the replica becomes the owner. Session information is maintained without any visible impact to the user.

  - Scalability: Any node in the cluster can process the traffic. Multiple nodes of the cluster can process the incoming requests served by the striped virtual server. This improves the application firewall's ability to handle multiple simultaneous requests, thereby improving the overall performance.

  Security checks and signature protections can be deployed without the need for any additional cluster-specific application firewall configuration. You just do the usual application firewall configuration on the configuration coordinator (CCO) node for propagation to all the nodes.

  Cluster details are available at
  http://docs.citrix.com/en-us/netscaler/11/system/clustering.html.

[From Build 55.20] [# 408831, 403780]

- All application firewall graphical user interface (GUI) dialog boxes, including the ones for signatures, visualizer, and syslog viewer, are now completely free from any java dependencies and show a significant improvement in the overall performance. The HTML based GUI dialogues have been re-organized for enhanced user experience and intuitive workflow of information. Instead appearing in of pop-up dialog boxes with tabs, the information is now displayed as an in-line expansion. You can expand all the configuration sections and scroll up and down for a comprehensive view.

  [From Build 55.20] [# 506157]

- All application firewall graphical user interface (GUI) dialog boxes, including the ones for signatures, visualizer, and syslog viewer, are now completely free from any java dependencies and show a significant improvement in the overall performance. The HTML based GUI dialogues have been re-organized for enhanced user experience and intuitive workflow of information. Instead appearing in of pop-up dialog boxes with tabs, the information is now displayed as an in-line expansion. You can expand all the configuration sections and scroll up and down for a comprehensive view.

  [From Build 55.20] [# 520048]

Cache Redirection

- **Support for default syntax expressions**

  You can now use default syntax expressions in cache redirection policies. The NetScaler appliance provides built-in cache redirection policies based on default syntax expressions, or you can create custom cache redirection policies to handle typical cache requests. In addition to the same types of evaluations done by classic cache redirection policies, the default syntax 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, directing requests to either cache or origin server).

  [From Build 55.20] [# 490297, 495915, 536986, 536992, 537010, 537014, 538269]

Cluster

- **Disabling Steering on the Cluster Backplane**

  By default, a NetScaler cluster steers traffic over the cluster backplane, from the flow receiver node to the flow processor node. You can disable steering so that the process becomes local to the flow receiver and thereby ensure that the flow receiver also becomes the flow processor. Such a configuration can come in handy when you have a high latency link.

  Note: This configuration is applicable only for striped virtual servers.

  Steering can be disabled at the global NetScaler level or at the individual virtual server level. The global configuration takes precedence over the virtual server setting.

  - At the global level, steering can be disabled for all striped virtual servers. It is configured at cluster instance level. Traffic meant for any striped virtual server will not be steered on cluster backplane. The command is:

> add cluster instance <clId> -processLocal ENABLED

- At a virtual server level, you can disable steering for a specific striped virtual server. It is configured on a striped virtual server. Traffic meant for that virtual server will not be steered on cluster backplane. The command is:

> add lb vserver <name> <serviceType> -processLocal ENABLED

For more information, see http://docs.citrix.com/en-us/netscaler/11/system/clustering/cluster-managing/cluster-steering-disable.html.

[From Build 55.20] [# 539136]

- **Routing on Striped SNIP addresses**

  You can now run dynamic routing on a striped SNIP address in a NetScaler cluster. The routes advertised by the cluster have the striped SNIP as the next hop. There is just one adjacency with the cluster. Internally, the cluster picks one of the active nodes as the routing leader. When the current routing leader goes down, the routing ownership moves to an active node.

  Note:

  - Striped SNIP addresses are useful mainly for cluster LA (link aggregation) deployments. They can also be used for ECMP, but the multipath routing functionality is unavailable.

  - Striped SNIP addresses can also be used in asymmetrical topologies.

  - Routing on striped SNIPs and routing on spotted SNIPs can coexist in a cluster.

  To specify leader node configurations, in the VTYSH shell, use the "owner-node leader" command.

  [From Build 55.20] [# 329439]

- **Reduce Backplane Steering for Spotted and Partially-striped Virtual Serves when Using ECMP**

  With the Equal Cost Multiple Path (ECMP) mechanism, virtual server IP addresses are advertised by all active cluster nodes. This means that traffic can be received by any cluster node, which then steers the traffic to the node that must process the traffic. While there are no hassles in this approach, there can be a lot of redundant steering in case of spotted and partially striped virtual servers. Therefore, from NetScaler 11 onwards, spotted and partially striped virtual server IP addresses are advertised only by the owner nodes. This reduces the redundant steering.

  You can override this default behavior, by entering the following command in the VTYSH shell:

  ns(config)# ns spotted-vip-adv all-nodes

  [From Build 55.20] [# 317706]

- **Nodegroup for Datacenter Redundancy**

A cluster nodegroup can now be configured to provide datacenter redundancy. In this use case, nodegroups are created by logically grouping the cluster nodes. You must create active and spare nodegroups. When the active nodegroup goes down, the spare nodegroup which has the highest priority (the lower priority number) is made active and it starts serving traffic.

For more information, see http://docs.citrix.com/en-us/netscaler/11/system/clustering/cluster-managing/cluster-nodegroups-datacenter-redundancy.html.

[From Build 55.20] [# 495019]

- BridgeGroups are now supported in a NetScaler cluster deployment.

  [From Build 55.20] [# 494991]

- **Routing in a L3 Cluster**

  In a L3 cluster, different nodegroups can have different VLANs and subnets associated with them. This can result in a VLAN getting exposed only in some nodes. Therefore, you can now configure dynamic routing on a VLAN to expose the VLAN to ZebOS even when there are no IP addresses with dynamic routing that are bound to it. The command to configure this is:

  > add/set vlan <id> -dynamicRouting (ENABLED | DISABLED)

  Note:

  - This option is also available for VXLAN and BridgeGroups.

  - This configuration can also be used for L2 clusters.

  [From Build 55.20] [# 531868]

- **Cluster to Include Nodes from Different Networks (L3 Cluster)**

  You can now create a cluster that includes nodes from different networks. To configure a cluster over L3, you must add the nodes of different networks to different nodegroups. For more information, see http://docs.citrix.com/en-us/netscaler/11/system/clustering/cluster-setup.html.

  You can transition an existing L2 cluster to an L3 cluster. For instructions, see http://docs.citrix.com/en-us/netscaler/11/system/clustering/cluster-usage-scenarios/cluster-migrate-between-l2-l3.html.

  [From Build 55.20] [# 374289, 317257]

- **Link Redundancy based on Minimum Throughput**

  In a dynamic cluster link aggregation (LA) deployment that has link redundancy enabled, you can configure the cluster to select the partner channel or interface on the basis of its throughput. To do this, configure a threshold throughput on the channel or interface as follows:

  > set channel CLA/1 -linkRedundancy ON -lrMinThroughput <positive_integer>

The throughput of the partner channels is checked against the configured threshold throughput. The partner channel that satisfies the threshold throughput is selected in FIFO manner. If none of the partner channel meets the threshold, or if threshold throughput is not configured, the partner channel with the maximum number of links is selected.

[From Build 55.20] [# 508993]

DNS

- **Rewrite and responder support for DNS**

  The rewrite and responder features now support DNS. You can now configure rewrite and responder functionalities to modify DNS requests and responses as you would for HTTP or TCP requests and responses.

  [From Build 55.20] [# 405769]

- **Enable or disable negative caching of DNS records**

  The NetScaler appliance supports caching of negative responses for a domain. You can enable or disable negative caching from the command line, by setting cacheNegativeResponses with the set dns parameter command, or in the configuration utility, in the Configure DNS Parameters dialog box.

  Note: You can enable or disable negative caching independent of global caching. By default, negative caching is enabled.

  [From Build 55.20] [# 391254]

- **Support for DNS Logging**

  You can now configure a NetScaler appliance to log DNS requests and responses. The logs are in SYSLOG format. You can use these logs to:

  - Audit the DNS responses to the client

  - Audit DNS clients

  - Detect and prevent DNS attacks

  - Troubleshoot

  [From Build 55.20] [# 419632, 561291]

GSLB

- Support for binding a single Virtual Server as a backup for multiple GSLB Virtual servers

  In a GSLB site deployment, you can now bind a single virtual server as a backup virtual server for multiple GSLB virtual servers in the deployment.

  [From Build 55.20] [# 373061]

Load Balancing

- **IPv6 Support for HTTP based User Monitors**

  You can now use IPv6 addresses in the following monitors:

- USER

- SMTP

- NNTP

- LDAP

- SNMP

- POP3

- FTP_EXTENDED

- STOREFRONT

- APPC

- CITRIX_WI_EXTENDED

Note: The monitor for MySQL does not support IPv6 addresses.

[From Build 55.20] [# 510111]

- **Support for Secure LDAP Monitor**

  You can now monitor LDAP services over SSL. To monitor the LDAP services over SSL, use the built-in LDAP monitor or create a user monitor and enable the "secure" option.

  [From Build 55.20] [# 418061, 556530]

- **Automatic Restart of the Internal Dispatcher**

  In earlier releases, if the internal dispatcher failed, the services that used scriptable monitors also went down and the appliance had to be restarted. From release 11, if the internal dispatcher fails, the pitboss process restarts it. As a result, you no longer have to restart the appliance. For information about user monitors, see http://docs.citrix.com/en-us/netscaler/11/traffic-management/load-balancing/load-balancing-custom-monitors/understand-user-monitors.html.

  [From Build 55.20] [# 368128]

- **Setting the Maintenance State for your Server with Minimal Interruption**

  You can now set the maintenance state for your server with minimal interruption and without changing any configuration on the NetScaler appliance. In the maintenance state, the server continues to accept persistent client connections while new connections are load balanced among the active servers. On the NetScaler appliance, configure a transition out of service (TROFS)-enabled monitor and bind it to a service representing the server. Specify a trofsCode or trofsString in the monitor. Upon receipt of a matching code or string from the server in response to a monitor probe, the appliance places the service in the TROFS state. During this time, it continues to honor persistent client connections.

  To avoid disrupting established sessions, you can place a service in the TROFS state by doing one of the following:

- Adding a TROFS code or string to the monitorâ€”Configure the server to send a specific code or string in response to a monitor probe.

Note: This enhancement is available from release 10.5 build 56.16.

- Explicitly disable the service and:

- Set a delay (in seconds).

- Enable graceful shut down.

Adding a TROFS Code or String

Note: This enhancement is not applicable to GSLB services.

From release 10.5, build56.16, if you bind only one monitor to a service, and the monitor is a TROFS-enabled monitor, it can place the service in the TROFS state on the basis of the serverâ€™s response to a monitor probe. This response is compared with the value in the trofsCode parameter for an HTTP monitor or the trofsString parameter for an HTTP-ECV or TCP-ECV monitor. If the code matches, the service is placed in the TROFS state. In this state, it continues to honor the persistent connections.

If multiple monitors are bound to a service, the effective state of the service is calculated on the basis of the state of all the monitors that are bound to the service. Upon receiving a TROFS response, the state of the TROFS-enabled monitor is considered as UP for the purpose of this calculation. For more information about how a NetScaler appliance designates a service as UP, see http://docs.citrix.com/en-us/netscaler/11/traffic-management/load-balancing/load-balancing-advanced-settings/set-monitor-threshold.html.

Important!

- You can bind multiple monitors to a service, but only one monitor must be TROFS-enabled.

- You can convert a TROFS-enabled monitor to a monitor that is not TROFS-enabled, but not vice versa.

[From Build 55.20] [# 408103]

- The following global timeouts has been introduced for TCP sessions on a NetScaler appliance related to RNAT rules, forwarding sessions, or load balancing configuration of type ANY:

  * Any TCP Client. Global idle timeout, in seconds, for TCP client connections. Client timeout set for an entity overrides the global timeout setting.

  * Any TCP Server. Global idle timeout, in seconds, for TCP server connections. Server timeout set for an entity overrides the global timeout setting.

  These timeout can be set either from the NetScaler command line (set ns timeout command) or from the configuration utility (System > Settings > Change Timeout Values page).

Note: For applying these timeouts to a virtual server or service of type ANY, set these timeouts before adding the virtual server or the service.

[From Build 55.20] [# 507701]

- **New Trap for Spillover**

  If you have configured spillover on a virtual server and also configured a trap listener on the appliance, an SNMP trap is now sent to the trap listener when the virtual server experiences spillover. The trap message displays the name of the virtual server that experienced the spillover, the spillover method, the spillover threshold, and the current spillover value. If the spillover is policy based, the rule causing it appears in the Spillover Threshold field. If the virtual server is DOWN or disabled, the status message "vserver not up" appears in the trap message.

  [From Build 55.20] [# 486268, 475400]

- If you have set the persistence type to COOKIEINSERT, you can now encrypt the cookie in addition to any existing SSL encryption by using the NetScaler command line and configuration utility.

  At the NetScaler command prompt, type:

  set lb parameter -useSecuredPersistenceCookie Enabled-cookiePassphrase test

  In the configuration utility, navigate to Traffic Management > Load Balancing > Change Load Balancing Parameters and select Use Secured Persistence Cookie and Cookie Passphrase and enter a passphrase.

  [From Build 55.20] [# 347108, 323325, 348588]

- If you configure cookie persistence and custom cookie on a virtual server, and later change the name or IP address of the virtual server, persistence is not honored.

  [From Build 55.20] [# 524079, 559022]

NetScaler Insight Center

- **Exporting Reports**

  You can now save the Web Insight reports or HDX Insight reports in PDF, JPEG, PNG , or CSV format on your local computer. You can also schedule the export of the reports to specified email addresses at various intervals.

  For more information, see http://docs.citrix.com/en-us/netscaler-insight/11-0/viewing-reports/ni-export-report-con.html.

  [From Build 55.20] [# 320860]

- You can now configure NetScaler Insight Center to display the reports in your local time or GMT time.

  [From Build 55.20] [# 491073]

- You can now identify the root cause of a terminated ICA session by viewing the session termination reason on the HDX Insight node. Along with the termination reason, it also displays the session TCP metrics such as ICA RTT and WAN Latency.

  [From Build 55.20] [# 488279]

- You can configure NetScaler Insight Center to display the geo maps for a particular geographical location or LAN by specifying the private IP range (start and end IP address) for the location.

  [From Build 55.20] [# 502478]

- Multi-Hop support for NetScaler Insight Center enables Insight Center to detect which Citrix appliances a connection passes through, and in which order, for improved reporting.

  [From Build 55.20] [# 383172]

- You can now increase the storage space of NetScaler Insight Center to 512 GB.

  [From Build 55.20] [# 425761, 553254]

- **Insight Deployment Management**

  You can now improve the processing power of and increase storage space in your NetScaler Insight Center deployment by adding agents, connectors, and databases. An agent processes HTTP traffic and sends the data to the connectors that distribute this data across databases. You can add multiple agents, connectors, and databases to scale your deployment. In this deployment, you can also the decide the number of resources you have to allocate and determine the elements you need in the database architecture, on the basis of the number of HTTP requests per second, number of ICA sessions, and number of active WAN connections.

  [From Build 55.20] [# 404919]

- You can now configure a DNS server when you set up NetScaler Insight Center. Configuring a DNS server helps resolve the host name of a server into its IP address.

  For example, while creating an email server, you now have an option to specify the server name of the server rather than the IP address.

  [From Build 55.20] [# 514612]

- The NetScaler Insight Center configuration utility now displays the progress of the upgrade process.

  [From Build 55.20] [# 519788, 522021]

- **Hop Diagram Support**

  The HDX Insight reports now support hop diagrams, which provide complete details about the client, NetScaler ADC, and server in an active session.

  To display the hop diagram, on the dashboard tab, navigate to HDX Insight > Users >, click on a user name and, in the Current Application Sessions table, click on the session diagram icon.

[From Build 55.20] [# 443824]

Networking

- **OSPFv3 Authentication**

  For ensuring the integrity, data origin authentication, and data confidentiality of OSPFv3 packets, OSPFv3 authentication must be configured on OSPFv3 peers.

  The NetScaler appliance supports OSPFv3 authentication and is partially compliant with RFC 4552. OSPFv3 authentication is based on the two IPSec protocols: Authentication Header (AH) and Encapsulating Security Payload (ESP). The NetScaler supports only the AH protocol for OSPFv3 authentication.

  OSPFv3 authentication use manually defined IPSec Security Associations (SAs) between the OSPFv3 peers and does not rely on IKE protocol for forming dynamic SAs. Manual SAs define the security parameter Index (SPI) values, algorithms, and keys to be used between the peers. Manual SAs require no negotiation between the peers; therefore, same SA must be defined on both the peers.

  You can configure OSPFv3 authentication on a VLAN or for an OSPFv3 area. When you configure for a VLAN, the settings are applied to all the interfaces that are member of the VLAN. When you configure OSPFv3 authentication for an OSPF area, the settings are applied to all the VLANs in that area. The settings are in turn applied to all the interfaces that are members of these VLANs. These settings do not apply to member VLANs on which you have configured OSPFv3 authentication directly.

  [From Build 55.20] [# 471703]

- **Configuring Communication Intervals for an Active-Active Deployment**

  In an active-active deployment, all NetScaler nodes use the Virtual Router Redundancy Protocol (VRRP) to advertise their master VIP addresses and the corresponding priorities in VRRP advertisement packets (hello messages) at regular intervals.

  VRRP uses the following communication intervals:

  * Hello Intervalâ€”Interval between successive VRRP hello messages that a node sends, for all of its active (master) VIP addresses, to the other nodes of the VRRP deployment. For a VIP address, nodes on which the VIP address is in the inactive state use the hello messages as verification that the master VIP address is still UP.

  * Dead Intervalâ€”Time after which a node of a backup VIP address considers the state of the master VIP address to be DOWN if VRRP hello messages are not received from the node that has the master VIP address. After the dead interval, the backup VIP address takes over and becomes the master VIP address.

  You can change these intervals to a desired value on each node. They apply to all VIP addresses on that node.

  [From Build 55.20] [# 512843]

- The NetScaler appliance supports sending static IPv6 routes through a VXLAN. You can enable the NetScaler appliance to send an IPv6 route through either a VXLAN or a VLAN. A VXLAN parameter is added to the static IPv6 route command set.

  [From Build 55.20] [# 472443]

- **Specifying a VLAN in a Static ARP Entry**

  In a static ARP entry, you can specify the VLAN through which the destination device is accessible. This feature is useful when the interface specified in the static ARP entry is part of multiple tagged VLANs and the destination is accessible through one of the VLANs. The NetScaler appliance includes the specified VLAN ID in the outgoing packets matching the static ARP entry. If you don't specify a VLAN ID in an ARP entry, and the specified interface is part of multiple tagged VLANs, the appliance assigns the interface's native VLAN to the ARP entry.

  For example, say NetScaler interface 1/2 is part of native VLAN 2 and of tagged VLANs 3 and 4, and you add a static ARP entry for network device A, which is part of VLAN 3 and is accessible through interface 1/2. You must specify VLAN 3 in the ARP entry for network device A. The NetScaler appliance then includes tagged VLAN 3 in all the packets destined to network device A, and sends them from interface 1/2.

  If you don't specify a VLAN ID, the NetScaler appliance assigns native VLAN 2 for the ARP entry. Packets destined to device A are dropped in the network path, because they do not specify tagged VLAN 3, which is the VLAN for device A.

  [From Build 55.20] [# 520355]

- **Changing the Priority of a VIP Address Automatically in an Active-Active Deployment**

  To ensure that a backup VIP address takes over as the master VIP before the node of the current master VIP address goes down completely, you can configure a node to change the priority of a VIP address on the basis of the states of the interfaces on that node. For example, the node reduces the priority of a VIP address when the state of an interface changes to DOWN, and increases the priority when the state of the interface changes to UP. This feature is configured on each node. It applies to the specified VIP addresses on the node.

  To configure this feature on a node, you set the Reduced Priority (trackifNumPriority) parameter, and then associate the interfaces whose state is to be tracked for changing the priority of the VIP address. When any associated interface's state changes to DOWN or UP, the node reduces or increases the priority of the VIP address by the configured Reduced Priority (trackifNumPriority) value.

  [From Build 55.20] [# 512848]

- **Support of IPv6 Dynamic Routing Protocols on VXLANs**

  The NetScaler appliance supports IPv6 dynamic routing protocols for VXLANs. You can configure various IPv6 Dynamic Routing protocols (for example, OSPFv3, RIPng, BGP) on VXLANs from the VTYSH command line. An option IPv6 Dynamic Routing Protocol has been added to VXLAN command set for enabling or disabling IPv6 dynamic routing protocols on a VXLAN. After enabling IPv6 dynamic routing protocols on a VXLAN,

processes related to the IPv6 dynamic routing protocols are required to be started on the VXLAN by using the VTYSH command line.

[From Build 55.20] [# 472432]

- **Layer 2 PBR Support for Forwarding Sessions**

  In earlier releases, Layer 2 information (for example, destination MAC address, source VLAN, and Interface ID) about packets related to forwarding sessions were ignored during a PBR lookup. In other words, any packet related to a forwarding session was not considered for matching against a PBR having Layer 2 parameters as its condition.

  Now, layer 2 information about a packet related to a forwarding session is matched against layer 2 parameters in the configured PBRs.

  This feature is useful in a scenario where packets related to a forwarding session must be processed by another device before being sent to their destination.

  Following are the benefits of this support:

  - Instead of defining new PBRs that are based on Layer 3 parameters, you can use existing PBRs based on Layer 2 parameters for sending the packets related to forwarding sessions to the desired next hop device.

  - In a deployment that includes NetScaler appliances and optimization devices, PBRs based on Layer 2 parameters can be very handy compared to other, complex configuration for identifying the forwarding session related packets for PBR processing.

  - Identifying forwarding session related Ingress packets for sending them to the optimization device.

  - Identifying egress packets, which also matched a forwarding session rule, from the optimization device for sending the packets to the desired next hop device.

  [From Build 55.20] [# 484458]

- **Logging HTTP Header Information**

  The NetScaler appliance can now log header information of HTTP requests related to an LSN configuration. The following header information of an HTTP request packet can be logged:

  - URL that the HTTP request is destined to.

  - HTTP Method specified in the HTTP request.

  - HTTP version used in the HTTP request.

  - IP address of the subscriber that sent the HTTP request.

  An HTTP header log profile is a collection of HTTP header attributes (for example, URL and HTTP method) that can be enabled or disabled for logging. The HTTP header log profile is then bound to an LSN group. The NetScaler appliance then logs HTTP header attributes, which are enabled in the bound HTTP header log profile for logging, of any HTTP requests related to the LSN group.

An HTTP header log profile can be bound to multiple LSN groups but an LSN group can have only one HTTP header log profile.

[From Build 55.20] [# 496835]

- **As-Override Support in Border Gateway Protocol**

  As a part of BGP loop prevention functionality, if a router receives a BGP packet containing the router's Autonomous System Number (ASN) in the Autonomous Systems (AS) path, the router drops the packet. The assumption is that the packet originated from the router and has reached the place from where it originated.

  If an enterprise has several sites with a same ASN, BGP loop prevention causes the sites with an identical ASN to not get linked by another ASN. Routing updates (BGP packets) are dropped when another site receives them.

  To solve this issue, BGP AS-Override functionality has been added to the ZebOS BGP routing module of the NetScaler.

  With AS-Override enabled for a peer device, when the NetScaler appliance receives a BGP packet for forwarding to the peer, and the ASN of the packet matches that of the peer, the appliance replaces the ASN of the BGP packet with its own ASN number before forwarding the packet.

  [From Build 55.20] [# 503566]

- **GRE Payload Options in a GRE IP Tunnel**

  For a configured GRE IP tunnel, the NetScaler appliance encapsulates the entire Layer 2 packet, including the Ethernet header and the VLAN header (dot1q VLAN tag). IP GRE tunnels between NetScaler appliances and some 3rd party devices might not be stable, because these 3rd party devices are not programmed to process some or the Layer 2 packet headers.

  To configure a stable IP GRE tunnel between a NetScaler appliance and a 3rd party device, you can use a new parameter with the GRE IP tunnel command set. You can set the GRE payload parameter to do one of the following before the packet is sent through the GRE tunnel:

  - Carry the Ethernet header but drop the VLAN header

  - Drop the Ethernet header as well as the VLAN header

  - Carry the Ethernet header as well the VLAN header

  [From Build 55.20] [# 518397]

- **Blocking Traffic on Internal Ports**

  The NetScaler appliance does not block traffic that matches an ACL rule if the traffic is destined to the appliance's NSIP address, or one of its SNIP addresses, and a port in the 3008-3011 range.

This behavior is now specified by the default setting of the new Implicit ACL Allow (implicitACLAllow) parameter (of the L3 param command). You can disable this parameter if you want to block traffic to ports in the 3008-3011 range. An appliance in a high availability configuration makes an exception for its partner (primary or secondary) node. It does not block traffic from that node.

To disable or enable this parameter by using the command line interface

At the command prompt, type:

> set l3param -implicitACLAllow [ENABLED|DISABLED]

Note: The parameter implicitACLAllow is enabled by default.

Example

> set l3param -implicitACLAllow DISABLED

Done

[From Build 55.20] [# 529317]

- **Redundant Interface Sets**

  A redundant interface set is a set of interfaces in which one interface is active and the others are on standby. If the active interface fails, one of the standby interfaces takes over and becomes active.

  Following are the main benefits of using redundant interface sets:

  - The back-up links between the NetScaler appliance and a peer device ensure connection reliability.

  - Unlike link redundancy using LACP, no configuration is required on the peer device for a redundant interface set. To the peer device, a redundant interface set appears as individual interfaces, not as a set or collection.

  - In a high availability (HA) configuration, redundant interface sets can minimize the number the HA failovers.

  A redundant interface set is specified in LR/X notation, where X can range from 1 to 4. For example, LR/1.

  [From Build 55.20] [# 355237, 186503, 249551]

- **Jumbo Frames Support for NetScaler VPX Appliances**

  NetScaler VPX appliances now support receiving and transmitting jumbo frames containing up to 9216 bytes of IP data. Jumbo frames can transfer large files more efficiently than is possible with the standard IP MTU size of 1500 bytes.

  A NetScaler appliance can use jumbo frames in the following deployment scenarios:

  - Jumbo to Jumbo. The appliance receives data as jumbo frames and sends it as jumbo frames.

- Non-Jumbo to Jumbo. The appliance receives data as regular frames and sends it as jumbo frames.

- Jumbo to Non-Jumbo. The appliance receives data as jumbo frames and sends it as regular frames.

Jumbo Frames support is available on NetScaler VPX appliances running on the following virtualization platforms:

- VMware ESX (Note that NetScaler VPX appliances running on VMware ESX support receiving and transmitting jumbo frames containing up to only 9000 bytes of IP data.)

- Linux-KVM

For configuring Jumbo Frames on a NetScaler VPX appliance, you must:

- Set the MTU of the interface or channel of the VPX appliance to a value in the range 1501-9216. Use the NetScaler command line interface or the configuration utility of the VPX appliance to set the MTU size.

- Set the same MTU size on the corresponding physical interfaces of the virtualization host by using its management applications.

[From Build 55.20] [# 464830, 478103, 485905]

- **Client Source Port for Server Side ConnectionsÂ related to INAT and RNAT Rules**

  The NetScaler appliance, for INAT and RNAT rules, now supports using client port as the source port for server side connections. A parameter Use Proxy Port has been added to the INAT and RNAT command set. WhenÂ Use Proxy Port is disabledÂ for an INAT rule or a RNAT rule, the NetScaler appliance retains the source port of the client's requestÂ for the server side connection. When the option is enabled (default), the NetScaler appliance uses a random port as the source port for the server side connection.Â Â

  You must disable thisÂ parameterÂ for proper functioning of certain protocols that requireÂ a specificÂ source port in the request packet.

  [From Build 55.20] [# 399821]

- **MAC Address Wildcard Mask for Extended ACLs**

  A new wildcard mask parameter for extended ACLs and ACL6s can be used with the source MAC address parameter to define a range of MAC addresses to match against the source MAC address of incoming packets.

  MAC Address Wildcard Mask for PBRs

  A new wildcard mask parameter for PBRs and PBR6s can be used with the source MAC address parameter to define a range of MAC addresses to match against the source MAC address of outgoing packets.

  [From Build 55.20] [# 391630]

Optimization

- **Support for JPEG-XR image format in Front End Optimization (FEO)**

  The front end optimization feature now supports the conversion of GIF, JPEG, TIFF, and PNG images to JPEG-XR format as part of the image optimization functionality.

  [From Build 55.20] [# 504044]

- **Media classification support on the NetScaler appliance**

  You can now monitor and display the statistics of the media traffic going through the NetScaler appliance.

  [From Build 55.20] [# 493103]

- **Support for WebP image format in Front End Optimization (FEO)**

  The front end optimization feature now supports the conversion of GIF, JPEG, and PNG images to WEBP format as part of the image optimization functionality.

  [From Build 55.20] [# 509338]

Policies

- **Policy extensions support on NetScaler appliance**

  The NetScaler appliance now supports policy extensions, which you can use to add customized functions to default syntax policy expressions. An extension function can accept text, double, Boolean or number values as input, perform a computation, and produce a text, double, Boolean or number result.

  [From Build 55.20] [# 248822]

- **Transaction Scope Variables**

  Transaction scope variables are added to variables feature. You can now use transaction scope variables to specify separate instances with values for each transaction processed by the NetScaler appliance. Transaction variables are useful for passing information from one phase of the transaction to another. For example, you can use a transaction variable to pass information about the request onto the response processing.

  [From Build 55.20] [# 444109]

SSL

- **Support for SNI with a SAN Extension Certificate**

  The NetScaler appliance now supports SNI with a SAN extension certificate. During handshake initiation, the host name provided by the client is first compared to the common name and then to the subject alternative name. If the name matches, the corresponding certificate is presented to the client.

  [From Build 55.20] [# 250573]

- **Support for TLS_FALLBACK_SCSV signaling cipher suite value**

  The NetScaler appliance now supports the TLS_FALLBACK_SCSV signaling cipher suite value. The presence of this SCSV extension in the Client Hello indicates that the client is

retrying to connect to the server by using a lower SSL version, after its previous attempt to communicate with a higher version failed. Therefore, if the server finds this extension in Client Hello and also finds that the client is proposing a version that is lower than the maximum version supported by the server, it is a likely indication of a "man in the middle attack." The server drops these handshakes.

For more information, see http://docs.citrix.com/en-us/netscaler/11/traffic-management/ssl/customize-ssl-config/config-protocol-settings.html.

[From Build 55.20] [# 509666, 573528]

- **DH Key Performance Optimization**

  DH key generation is optimized on a VPX appliance by adding a new parameter dhKeyExpSizeLimit. You can set this parameter on an SSL virtual server or on an SSL profile and bind the profile to the SSL virtual server. The key generation is optimized as defined by NIST in http://csrc.nist.gov/publications/nistpubs/800-56A/SP800-56A_Revision1_Mar08-2007.pdf. Additionally, the minimum DH count is set to zero. As a result, you can now generate a DH key for each transaction as opposed to a minimum of 500 transactions earlier. This helps to achieve perfect forward secrecy (PFS).

  [From Build 55.20] [# 498162, 512637]

- **Support for Additional Ciphers on a DTLS Virtual Server**

  EDH, DHE, ADH, EXP, and ECDHE ciphers are now supported on a DTLS virtual server.

  [From Build 55.20] [# 508440, 483391]

- **Support for Auto-Detection of the Certificate-Key Pair Format**

  The NetScaler software has been enhanced to automatically detect the format of the certificate-key pair. To do so, the format of the certificate and key file should be the same. If you specify the format in the inform parameter, it is ignored by the software. Supported formats are PEM, DER, and PFX.

  [From Build 55.20] [# 209047, 432330, 481660]

- **New SNMP OIDs for SSL transactions per second**

  The following SNMP OIDs have been added to the display the SSL transactions per second:

  NS-ROOT-MIB::sslTotTransactionsRate.0 = Gauge32: 0

  NS-ROOT-MIB::sslTotSSLv2TransactionsRate.0 = Gauge32: 0

  NS-ROOT-MIB::sslTotSSLv3TransactionsRate.0 = Gauge32: 0

  NS-ROOT-MIB::sslTotTLSv1TransactionsRate.0 = Gauge32: 0

  [From Build 55.20] [# 449923]

- **Support for TLS Protocol Version 1.1 and 1.2 on the NetScaler MPX, MPX-FIPS, Appliances**

  The NetScaler MPX appliance now supports TLS protocol versions 1.1 and 1.2 on the backend. MPX-FIPS appliances running firmware version 2.2 also support TLSv1.1/1.2 on the backend.

  [From Build 55.20] [# 494082, 566364]

- **Stricter Control on Client Certificate Validation**

  You can configure the SSL virtual server to accept only client certificates that are signed by a CA certificate bound to the virtual server. To do so, enable the ClientAuthUseBoundCAChain setting in the SSL profile bound to the virtual server.

  For more information, see http://docs.citrix.com/en-us/netscaler/11/traffic-management/ssl/config-ssloffloading/ssl-profiles.html.

  [From Build 55.20] [# 533241]

- **Changes to the Default Cipher Suite**

  If user-defined ciphers or cipher groups are not bound to an SSL virtual server, the DEFAULT cipher group is used for cipher selection at the front end and the ALL cipher group is used for cipher selection at the back end. In this release, the predefined cipher suites, such as DEFAULT and ALL, are modified to give strong ciphers a higher priority. For example, earlier RC4-MD5 was given a higher priority but it is deprioritized in the new list because it is a weak cipher.

  [From Build 55.20] [# 226713, 258311, 384491]

- **Support for Displaying the Hex Code of a CIpher**

  The show ciphersuite command now displays the IETF standard hexadecimal code of the cipher. It is helpful in debugging, because a hex code is unique to a cipher but the cipher name might differ on the NetScaler appliance, OpenSSL, and Wireshark.

  At the NetScaler command line, type:

  show ciphersuite

  In the configuration utility, navigate to Traffic Management > SSL > Cipher Groups.

  [From Build 55.20] [# 491286]

- **Support for Checking the Subject Alternative Name in addition to the Common Name in s Server Certificate**

  If you configure a common name on an SSL service or service group for server certificate authentication, the subject alternative name (SAN), if specified, is matched in addition to the common name. Therefore, if the common name does not match, the name that you specify is compared to the values in the SAN field in the certificate. If it matches one of those values, the handshake is successful. Note that in the SAN field, only DNS names are matched.

[From Build 55.20] [# 439161]

- **2048-bit Default Certificates on the NetScaler Appliance**

  With this release, the default certificate on a NetScaler appliance is 2048-bits. In earlier builds, the default certificate was 512-bits or 1024-bits. After upgrading to release 11.0, you must delete all your old certificate-key pairs starting with "ns-", and then restart the appliance to automatically generate a 2048-bit default certificate.

  [From Build 55.20] [# 451441, 405363, 458905, 465280, 540467, 551603, 559154]

- **Support for TLS Protocol Version 1.1 and 1.2 on the NetScaler VPX Appliances**

  The NetScaler VPX appliance now supports TLS protocol versions 1.1 and 1.2 on the front end.

  [From Build 55.20] [# 424463, 481970]

System

- **Support for FACK on TCP profiles**

  The TCP profiles on a NetScaler appliance now support forward acknowledgement (FACK). FACK avoids TCP congestion by explicitly measuring the total number of data bytes outstanding in the network, and helping the sender (either a NetScaler ADC or a client) control the amount of data injected into the network during retransmission timeouts.

  [From Build 55.20] [# 439130]

- **User configurable congestion window for TCP profile**

  You can now set the maximum congestion window size for a TCP profile on the NetScaler appliance.

  [From Build 55.20] [# 248711]

- During the execution of the "nstrace.sh" script (from shell) or the "start nstrace" command (from CLI), when the trace file is rolled over, some packets might not be available in the trace. The number of packets that will be dropped from the trace is directly proportional to the traffic rate.

  [From Build 55.20] [# 480258, 494482, 523853]

- Support for milliseconds, microseconds, and nanoseconds in Time Format Definition table

  You can now configure NetScaler web logging clients to capture transaction times in milliseconds, microseconds, and nanoseconds for logging on the NetScaler appliance.

  [From Build 55.20] [# 505840, 505377]

- **Support for HTTP/2 on the NetScaler Appliance**

  The NetScaler appliance supports HTTP/2 connections with clients supporting HTTP/2 protocol.

  [From Build 55.20] [# 490096, 505747]

- The NetScaler introduces a new role called sysadmin. A sysadmin is lower than a superuser is terms of access allowed on the appliance. A sysadmin user can perform all NetScaler operations with the following exceptions: no access to the NetScaler shell, cannot perform user configurations, cannot perform partition configurations, and some other configurations as stated in the sysadmin command policy.

  [From Build 55.20] [# 548516]

- The NetScaler appliance fails intermittently when trace is started in 'RX' mode.

  [From Build 55.20] [# 576067]

- **NTP Version Update**

  In NetScaler release 11, the NTP version has been updated from 4.2.6p3 to 4.2.8p2.

  If you upgrade your NetScaler appliance from any earlier release to release 11, the NTP configuration is automatically upgraded with additional security policies. For more information about configuring an NTP server, see http://docs.citrix.com/en-us/netscaler/11/system/basic-operations/configuring-clock-sychronization.html.

  [From Build 55.20] [# 440375, 440591]

- **Showtechsupport utility enhancement**

  If your NetScaler appliance has Internet connectivity, you can now directly upload the newly generated collector archive to the Citrix technical support server from the appliance.

  [From Build 55.20] [# 480797]

- The NetScaler Web Logging (NSWL) client logs a hyphen (-) instead of a user name when %u is specified in the log format.

  [From Build 55.20] [# 238440, 239481, 247372, 422873]

- **Maintaining minimum number of reuse pool connections in HTTP Profiles**

  You can now specify the minimum number of reuse pool connections to be opened from the NetScaler appliance to a particular server. This setting helps in optimal memory utilization and reduces the number of idle connections to the server.

  [From Build 55.20] [# 397478]

- The NetScaler appliance generates SNMP clear alarm traps for successful cases of haVersionMismatch, haNoHeartbeats, haBadSecState, haSyncFailure, and haPropFailure error events in an HA configuration.

  [From Build 55.20] [# 368832]