# NetScaler 1000V Release Notes

**NetScaler 11.1**

Last Updated: April 2017

## Contents

# 11.1-53.11

Updated: April 27, 2017 | Release notes version: 1.0

This release notes document describes the enhancements and changes and specifies the issues that exist, for the NetScaler release 11.1 Build 53.11. See Release History.

# What’s New

The enhancements and changes that are available in Build 53.11.

**AAA-TM**

- Maximum relayState size accepted by NetScaler SAMLIdP is increased to 2500 bytes

  With this enhancement, maximum relayState size accepted by NetScaler SAMLIdP is increased to 2500 bytes as compared to existing 1024 bytes.

  [# 678694]

**Admin Partitions**

- Support for sending SNMP traps of all partitions through NetScaler GUI

  On a partitioned NetScaler appliance, you can now use the NetScaler GUI to enable sending SNMP trap messages of all partitions to the configured trap destination. In the default partition, enable the allPartitions option for the traps that you want to send. Previously, you had to use the NetScaler command line to enable this option.

  Navigate to System > SNMP > Traps, select a trap, click Edit, and select or clear the Send Traps of All Partitions check box.

  [# 677551]

**DNS**

- Flushing Negative Records

  Negative records (NXDOMAIN records and NODATA records) cannot be deleted from the NetScaler

  appliance's DNS cache. With this fix, you can use the flush dns proxyrecords command to flush

  negative DNS records from the DNS cache.

  [# 665527]

- Dropping a DNS Query When the Query is Split into Multiple Packets

  New option '-splitPktQueryProcessing' is added to 'dns parameter' list. This option can be disabled to prevent processing of requests split across multiple packets.

  [# 665067]

- Restricting TTL of Negative Records

  You cannot set an appropriate time to live (TTL) value for the negative records.With this fix, you can use the new maxnegcacheTTL option in the set DNS parameter list to set a TTL for negative records.

  [# 665528]

- Collecting Statistics of the DNS Responses Served from the Cache

  You can now collect statistics of the DNS responses served from cache and use these statistics to create a threshold beyond which additional DNS traffic is dropped. You can enforce the threshold with a bandwidth based policy. Previously, bandwidth calculation for a DNS load balancing virtual server was not accurate, because the number of cache hits was not reported. In proxy mode, the statistics for Request bytes, Response bytes, Total Packets rcvd, and Total Packets sent statistics are continuously updated. Previously, these statistics were not always updated, particularly for a DNS load balancing virtual server.

  [# 665081]

**Load Balancing**

- Enabling or Disabling Persistence Session on Services that are in TROFS State

  You can set the new trofsPersistence flag to specify whether or not to honour persistence sessions for services that are in the TROFS state. Persistence sessions are honored If this flag is set to ENABLED. They are not honored if it is set to DISABLED. Previously, persistence sessions were honored if a service was in the TROFS state.

  [# 657750]

# Fixed Issues

The issues that are addressed in Build 53.11.

**AAA-TM**

- Inflate of the data fails intermittently when NetScaler IdP receives authentication request from external Service Provider (SP) incase "Redirect Binding" is used as the transport mechanism in SAML flow.

  When "Redirect Binding" is used as the transport mechanism in SAML flow, when NetScaler IdP receives authentication request from external Service Provider (SP), occasionally inflate of the data fails. This is highly intermittent. Current enhancement offers different variants for doing inflate as the issue in intermittent and disappears on a reboot.

  [# 680064]

**AppFlow**

- If you enable AppFlow on a SQL virtual server, the NetScaler appliance might become unresponsive.

[# 671462, 672362, 668129, 670343]

- Memory usage on a NetScaler appliance might increase over time if AppFlow Client-side Measurements is enabled.

  [# 666358, 672859]

- If AppFlow Client-Side Measurements action is enabled on any of the AppFlow policies, connection to backend server might get terminated intermittently thereby sending only partial response to the client.

  [# 672863, 673532, 677954, 677576]

**Application Firewall**

- A NetScaler AppFirewall appliance displays the following error message when you try to deploy learned rules with the WAF learning mode enabled: Error in retrieving Application Firewall learning data.

  [# 674023]

- The IP reputation feature does not get enabled when the application firewall add-on license is added.

  [# 675202]

- The NetScaler appliance crashes during a field-consistency check if processing a large number of form-select fields.

  [# 668627, 664482]

- A NetScaler appliance might fail to start after URL transformation, because of low memory allocation.

  [# 674415, 675793, 679479, 678765, 677990]

- The NetScaler appliance fails to upload files for a policy profile with signatures when the NetScaler AppFirewall signature function is enabled.

  [# 660112]

- The NetScaler application firewall blocks web-service URLs and displays the following error message: No_Service_URL.

  [# 673630]

- The NetScaler appliance fails to restart after an upgrade to software release 11.1 build 51.21. The failure, caused by memory corruption in the AppSecure module, occurs while evaluating an invalid session cookie.

  [# 674361]

- On a NetScaler Application Firewall appliance in a high availability configuration, learning mode does not work after an upgrade to release 11.0 build 68.10.

[# 679411]

**DNS**

- The set lb vserver command allows you to assign the same IP address to the DNS name server and the DNS virtual server. With this fix, neither the set lb vserver nor the add dns nameServer command, nor the NetScaler GUI, allows you to assign the same address to both virtual servers.

  [# 665651]

**GSLB**

- The NetScaler GUI displays an error message when you autosync for a non-default partition.

  [# 648396]

- A NetScaler appliance might go DOWN while unbinding a GSLB domain from a GSLB virtual server. This issue occurs rarely, but can occur if GSLB site persistence is configured.

  [# 666105]

**Load Balancing**

- An error message appears when the Test Connectivity button is clicked for FQDN based LDAP/RADIUS servers,

  [# 671494]

- Admin partition packets originating from FreeBSD and destined to a virtual server's VIP address are not forwarded to FreeBSD in the return path.

  [# 671789, 621010]

**NetScaler ICA**

- When session reliability is enabled for the high availability feature, memory usage by the NetScaler appliance spikes and causes a failover.

  [# 671918, 673784, 656996, 672949, 676413]

- Debug prints that are enabled by default cause unwanted logging in the NetScaler console.

  [# 681196]

**NetScaler VPX Appliance**

- A NetScaler VPX instance might stop responding and dump core memory if you allocate a large disk size for log messages. The higher the rate of log messages, the more quickly the instance runs out of memory and fails.

  [# 646674]

- Certificate-based authentication (SSH key pair) does not work on a NetScaler VPX appliance running on Azure. This happens due to internal logic that uses different keys to encrypt and decrypt certificate data.

  [# 668007]

**SSL**

- In a cluster setup, you cannot make any change to a service or service group if you have associated a common name with the service or the service group and enabled or disabled server name indication (SNI).

  [# 665340]

- Twenty-five days after a NetScaler appliance is restarted, memory utilization continuously increases. As a result, the appliance might stop processing traffic or dump core memory and restart if the memory is exhausted.

  [# 670731, 669812, 675045, 677777, 677322, 680221]

- If both OCSP stapling and session ticket are enabled on an SSL virtual server, and a client sends a session reuse request that contains an OCSP stapling status extension, the appliance dumps core memory and restarts.

  [# 678743, 678740]

- The NetScaler appliance might dump core memory and restart if AppFlow and SSL features are enabled and the appliance receives SSL traffic.

  [# 673897, 674543, 674479, 674128, 676165, 680784]

- The wrong counter increments when alerts are received from a client counter. Instead of the ssl_tot_sslError_FatalAlertRecdCount counter, the ssl_tot_sslError_FatalAlertSentCount counter increments.

  [# 659782, 662587, 675640, 676150, 674138, 673277, 677793, 676317, 679944, 681715]

- "Client Cert Required" appears in the CLI output of SSL services, even if requirements for a client certificate have been met. This is only a display issue.

  Example:

  > sh ssl service svc1

  ...

  Server Auth: DISABLED Client Cert Required:

  ...

  [# 668085]

- If an OCSP responder URL incorrectly resolves to a NetScaler reserved IP address, the appliance dumps core memory and restarts.

[# 675887]

**System**

- After an upgrade, a NetScaler Weblogging (NSWL) HTTP record size is miscalculated if the HTTP header size is greater than 16 kilobytes and it is not a multiple of the word boundary.

  [# 671996, 672244, 678903]

- If the HTTP/2 window update frame is not handled properly and when integrated caching and application firewall is enabled, the queued packets are transmitted and they fail to update the HTTP/2 window. This results in an appliance crash

  [# 634356, 660867, 668809, 670748, 674245]

- Memory usage on a NetScaler appliance might increase over time if Multipath TCP (MPTCP) is enabled and MPTCP to Subflow sequence number mapping fails because of a split packet error in the lower client-side MSS. The appliance becomes unresponsive after the memory is exhausted.

  [# 672009, 670102]

- If an HTTP WebSocket upgrade connection request contains a Content-Length header field, WebSocket applications malfunction.

  [# 673826]

- When NetScaler Web Logging (NSWL) is trying to send log data and client connection does not exist, it causes the appliance to crash.

  [# 671383]

- If Multipath TCP (MPTCP) is enabled, the NetScaler appliance might dump core memory and restart because a protocol control block (PCB) is freed twice.

  [# 673228]

- The default Rx ring size is set as 512. However, you can use the nsif command through the NetScaler command line to change the Rx size to 1024 or 2048 at run time.

  [# 623977, 649735, 665707, 676636]

# Known Issues

The issues that exist in Build 53.11.

**AAA-TM**

- The NetScaler appliance exhibits some inconsistency in the way expired cookies (TEMP) are handled:

- On an existing TCP connection, access to backend resources is allowed.

- On a new TCP connection, the request is denied.

[# 610091]

- If a user name containing special characters is prefilled in the login forms, the RfWeb user interface fails to render the form.

  Workaround: Escape the angular brackets.

  Example:

  Username is prefilled in the login forms on the basis of the value of the InitialValue tag in the authentication schema file.

  Change

  <InitialValue>${http.req.user.name}</InitialValue>

  To

  <InitialValue><![CDATA[${http.req.user.name}]]></InitialValue>

  [# 646139]

- If you log on to the NetScaler Traffic Management (TM) virtual server using "401 Basic" authentication, you might observe authentication failures if your username or password contains special characters. This is because only UTF-8 characters below ASCII 128 (for example, A-Z, a-z, 0-9, and ~ ! @ # $ % ^ & * ( ) _ + - = [ { ] } \ | ; : ' " / ? . > , < special characters) are allowed.

  [# 620845, 589509, 650263, 672340]

**Admin Partitions**

- In a non-default partition, if the network traffic exceeds the partition bandwidth limit, the FTP control connection fails but the data connection remains established.

  [# 620673]

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

  [# 493668, 516396]

**AppFlow**

- ICA parsing uses a lot of memory, so the NetScaler appliance reaches its memory limit with a lower than expected number of connections.

  [# 459458]

**Application Firewall**

- In high availability (HA) mode, high memory consumption failover occurs when the IP reputation feature is enabled.

  [# 668205]

- If you upgrade NetScaler appliance in a high availability (HA) setup from version 10.5.56.15 to version 11.1.51.1901 and skip 250 rules with active traffic, the GUI or CLI displays a "failed to skip some rules" error message and an operation time-out error message.

  Workaround: Turn off the Learning feature when skipping learned rules.

  [# 671807]

- The output of the appfw learningdata command does not include a caret and dollar sign (^$) at the beginning and end of a URL string. Therefore, the URLs are not in proper regex format. If you do not enclose a URL in ^$ characters when you specify a learned rule to be deleted, all the rules are deleted.

  [# 668255]

- In the Visualizer, if you use Mozilla Firefox or Internet Explorer, some buttons in the Visualizer might not work.

  Workaround: Use a different web browser, such as Google Chrome.

  [# 648272]

- 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]

- If you use the NetScaler GUI to access the application firewall security check violation log messages from a profile, the syslog viewer cannot display the logs if they are not in the CEF log format. You can enable CEF logging from the application firewall settings pane in GUI the or use the following command from CLI:

  > set appfw settings CEFLogging ON

[# 630056]

- A NetScaler AppFirewall appliance with the compression feature enabled sometimes puts blank lines in HTTP response headers, resulting in garbled page rendering by the browser.

  [# 629128]

**Cache Redirection**

- In a cluster deployment, if a request is received by a node other than the node on which the client request is received, a packet loop delays the response to the request.

  [# 591265]

**Clustering**

- In a cluster setup, if you use an interface on one node to create an LACP channel on another node, the channel is created and runs smoothly, but the system reports a configuration error.

  [# 644080]

- In a cluster setup, after a reboot, tagged VLAN configuration is lost on the vlan 1 interface.

  [# 642947]

**Command Line Interface**

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

  [# 512526, 527066, 545578, 631658, 635938, 643466, 652771, 667794]

**DNS**

- A NetScaler appliance configured for DNSSEC offloading might fail because of a race condition that can occur when the appliance receives a DNS query for a type A record for a domain that also has a CNAME record, and the canonical name identifies a domain that is in the zone offloaded for DNSSEC processing.

  [# 599741]

- In a high-availability setup, a failover might occur because of memory exhaustion on the appliance if all of the following conditions are met:

  - There is a large DNS-related configuration.

  - Heavy traffic is observed for a long duration.

- CPU utilization for one CPU is very high.

[# 621661]

**High Availability**

- If you upgrade a NetScaler appliance in a high availability (HA) setup to 11.1-51.21 from an older build of the same release, HA synchronization and command propagation are disabled during the upgrade process. However, after both the appliances are upgraded to the same NetScaler software version, HA synchronization and command propagation are enabled automatically.

  [# 670784]

**Load Balancing**

- After a high availability failover, Web Interface on NetScaler displays "State Error" if you try to launch an application.

  [# 630435]

- The NetScaler appliance is unable to reuse an existing probe connection if an HTTP wildcard load balancing virtual server is configured in MAC mode with use source IP (USIP) mode enabled and the Use Proxy Port option turned off. As a result, the connection fails and client the receives a TCP reset.

  [# 632872]

**NITRO**

- A NetScaler appliance returns error code 0 if the showtechsupport script fails while uploading the collector bundle to the Citrix server.

  To identify the failure, search the script's response data for the following string pattern:

  Upload of collector archive [] failed

  [# 629572]

**NetScaler CLI**

- When you use the Net::SSH::Perl library to connect to the NetScaler appliance, and run a command with an argument that has an @ character, an error message reports that the argument does not exist.

  For example, an error message appears if you use the @ character in the tacacsSecret parameter of the following command:

  > set authentication tacacsAction TACACS-0101 -tacacsSecret Sl4make5f0rd@enc5

Workaround: Use one of the following alternate approaches:

- If you use the Net::SSH::Perl library, include double quotes around the command when calling $ssh->cmd().

- Use the Net::Telnet library.

- Use the Net::SSH::Expect library.

[# 346066]

**NetScaler CPX**

- In NetScaler CPX, the newnslog file is rotated once in two days or when the file size reaches 600MB (whichever comes first). The file rotation can go up to 200 files.

  [# 644009]

  NetScaler GUI

- Certificate bundles are not supported in cluster setups.

  [# 644199]

- If the feature "Force password change for nsroot user when default nsroot password is being used" is enabled and the nsroot password is changed at the first logon to the NetScaler appliance, the nsroot password change is not propagated to non-CCO nodes. Therefore, when an nsroot user logs on to non-CCO nodes, the appliance asks for password change again.

  [# 658132]

- If the feature "Force password change for nsroot user when default nsroot password is being used" is enabled, and you log on as nsroot user, an extra session is created.

  [# 657924]

- In older versions of Internet Explorer version 7, the browser incompatibility message does not appear for NetScaler build 11.1. The logon page directly appears, and you can log on successfully.

  [# 649052]

- The service group members do not appear in the output of the "show lb vserver" command if it is run on a cluster IP address.

  [# 642802, 668935]

- LDAP configuration failed if the virtual server name started with an underscore ("_").

  [# 646751]

- The Upgrade Wizard sometimes does not display a message when the appliance is rebooting. However, the NetScaler appliance reboots and the upgrade is successful.

  [# 557379, 585649, 609615, 617161, 646039]

- In an ESX environment, the Interface HAMON Configuration option is not available in the NetScaler GUI.

  [# 641498]

- Compatibility issues between Linux-KVM and the Intel XL710 interface might cause a NetScaler virtual appliance configured with a PCI passthrough to become unresponsive during startup.

  Workaround: Restart the Linux-KVM host.

  [# 660139]

- For IPv6 or LACP support, promiscuous mode must be enabled for VMXNET3 interfaces at the ESX Hypervisor.

  [# 641748]

- The physical link status of a PCI passthrough interface of a NetScaler VPX appliance is not updated when the state of the link is changed (for example, when the link is enabled, disabled or reset) because of a limitation in the Intel XL710 NIC. As a result, any active traffic over the PCI passthrough interface fails during this time.

  [# 660159]

- In ESX-5.5.0 (Patch-2456374), you cannot restart or shut down the NetScaler VPX instance from the VPX console.

  [# 617922]

- Untagged packets are allowed to pass through an SRIOV VF interface (Intel 82599 NIC) if the VMWare vCenter 6.0 Distributed Virtual Switch( DVS) is used to configure the VLAN trunk mode.

  [# 616044]

- When you add custom DNS name server in the NetScaler VPX appliance through NetScaler CLI, DNS lookup fails. This happens due to a default Azure DNS server entry present in /etc/resolv.conf.

  Workaround: Follow these steps:

  1. Edit the /nsconfig/.AZURE/resolv.conf file to remove the entry "nameserver 168.63.129.16", and save the file.

2. Add the required DNS name server through NetScaler CLI, by using the command "add dns nameserver <dns server IP>".

3. Save the ns config file.

4. Restart the NetScaler VPX appliance.

[# 672344]

- In an ESX environment, file transfer from a NetScaler instance to an external connection stalls if the MTU is changed during the file transfer.

  [# 630639]

- The VLAN Trunk mode of operation does not work for SRIOV VF interfaces (Intel 82599 NIC) with ixgbe PF driver 3.21.6 or later. This is a known limitation reported by Intel.

  Workaround: Use ixgbe PF driver 3.21.4.3.

  [# 636360]

- Enabling trunk mode with tagged VLAN settings on an SR-IOV interface fails with the following error message:

  "ERROR: Maximum number of tagged VLANs bound to the interface exceeded or the binding of this VLAN is not allowed on the interface."

  However, trunk mode with tagged VLAN settings is shown as enabled in the output of the following command:

  show interface summary

  [# 657462]

- If you use the following command to remove an allowed-VLAN list from an SR-10V interface, the list is not removed, and therefore you cannot configure new VLAN settings for the interface.

  unset int -trunkallowedVlan

  Workaround: Restart the NetScaler virtual appliance.

  [# 657468]

- Due to a limitation in Linux-KVM and VMware ESX platforms, if you add new PCI passthrough interfaces to an existing NetScaler virtual appliance configured with SR-IOV interface, the PCI passthrough interfaces might take precedence over the existing SR-IOV interfaces.

[# 660000]

- Traffic might not pass through an SRIOV interface if you use the VMWare vCenter 6.0 Distributed Virtual Switch (DVS) to reconfigure a VLAN trunk policy.

  This is a known issue with VMWare vCenter 6.0. Please contact VMWare support for possible workarounds.

  [# 622392]

- The NetScaler virtual appliance might fail to start if you have configured 15 or more SR-IOV and PCI passthrough interfaces.

  [# 657492]

**Networking**

- VLAN trunk mode and allowed VLAN list configurations are not supported on Link Aggregation (LA) channels and redundant interface sets.

  [# 590805]

- If a VLAN specified in the allowed VLAN list of a trunk interface overlaps with the native VLAN of another interface, both the interfaces participate in packet processing on that VLAN.

  [# 631589]

- In a high availability setup, allowed VLAN list is not propagated or synchronized. Therefore, you have to configure allowed VLAN list on both the nodes.

  [# 631592]

- When a NetScaler appliance processes traffic at line rate, management CPU spike is observed on the appliance while configuring allowed VLAN list.

  [# 638915]

- The NetScaler appliance might become unresponsive while processing a route dependency check for multiple recursive BGP routes if the next hop for any of the routes changes or goes down.

  [# 625841]

- In a high-availability setup, NSVLAN is synchronized to the secondary node as a regular VLAN if the same NSVLAN is not configured on the secondary node.

  [# 629102]

- If an interface and an IP address are bound to a VLAN, binding them to another VLAN fails with the following error message: "ERROR: Either the subnet is not directly connected or subnet already bound to another VLAN." The interface is unbound from its current VLAN and gets bound to the native VLAN.

  [# 643341]

**Policies**

- If a policy expression name is same as any function name, subsequent use of the expression results in an error. In addition, if you restart the appliance and use the policy expression in a running configuration, the policy expression receives errors, which results in a configuration loss.

  Workaround: Do not name a policy expression with the same name as any function. The simplest way to rename a policy expression is to add a prefix or suffix to the expression name (for example, myco_func or func_myco).

  [# 637060]

**SSL**

- In a high availability (HA) setup, if the primary node supports a SafeNet HSM, the HSM configuration is propagated to the secondary node even though the secondary node is not configured to support the SafeNet HSM. For information about configuring an HA setup with SafeNet network HSMs, see the NetScaler documentation for SafeNet network HSM.

  [# 628082]

- ECDHE support with SSLv3 protocol on the NetScaler appliance is not compatible with RFC 4492, because SSLv3 does not support extensions and ECDHE needs extension support.

  [# 610588, 657755]

- If you restart the SafeNet network HSM, you must also restart the SafeNet gateway daemon.

  [# 628067]

- If you run the command "sh ssl service group" on the cluster IP (CLIP) and nodes on a cluster setup, ECC curves are displayed as unbound from the CLIP.

  [# 660257]

- The number of SSL cards that are UP is not displayed for non-default partitions. Because SSL cards are shared between the default partition and the non-default partitions, the total number of SSL cards that are UP in all the non-default partitions is equal to the number of cards that are UP in the default partition.

  [# 628914]

- An incorrect error message is displayed in both the following cases:

  1. Client authentication is enabled, root CA certificate is not bound to the SSL virtual server, and a request with a valid client certificate is sent to the virtual server.

  2. Client authentication is enabled, root CA certificate is bound to the SSL virtual server, and a request with a wrong certificate is sent to the virtual server.

  The error message that appears is "Handshake failure-Internal Error" instead of "No client certificate received."

  [# 664574]

- If you create a custom cipher group and bind it to an SSL entity, the profile name "SSL_EMBEDDED_PROFILE" incorrectly appears in the output of the "show ciphergroup" command. This error does not occur if you enable the Default profile before creating the custom cipher group and binding it to the SSL entity.

  [# 637230]

- In a cluster setup, if a client certificate is bound to a back-end SSL service or service group, it appears as a "Server Certificate" instead of a "Client Certificate" when you run the "show ssl service" or the "show ssl servicegroup" command on the CLIP address.

  [# 667389]

- All SSL-based policy expressions evaluate to FALSE in HTTP/2 connections.

  [# 660674]

- If you use the add crl command in release 9.3 to add a certificate revocation list (CRL) with refresh enabled, and you don't specify a method, the add crl command returns an error after an upgrade to a later release. Unlike 9.3, later releases do not have a default method.

  [# 604061]

- After you upgrade to this build, the priority of the cipher groups changes in the default profile.

  [# 579059]

- If you have configured two SafeNet HSMs in a high availability setup on a standalone NetScaler appliance, and the primary HSM goes down, the secondary HSM does not serve traffic after a failover.

  [# 628075]

- The SSL entities to which a policy is bound do not appear in the output of the "show ssl policy" command if it is run on the cluster IP address.

[# 668520]

- An HTTPS monitor bound to an SSL service fails to send ECDHE ciphers in Client Hello messages.

  [# 658154]

**System**

- Connection failover might fail if it is enabled on virtual servers that have the same IP address and port but different listen policies.

  [# 582087, 587620]

- No Error or Warning is announced if a user tries to set trunk mode on the loopback interface.

  [# 643131]

- A NetScaler appliance does not open a new connection to the back-end server if the following set of conditions is met:

  - The global maxconn parameter is set to 1.

  - The appliance is unable to reuse the connection for probing.

  As a result, the transaction fails.

  [# 636416]

- If you enable AppQoe and AppFlow features with client-side-measurements, more memory is allocated for storing the URL and host header without freeing the first allocation. As a result, a memory leak occurs in the HI memory pool.

  [# 640545]

- If a NetScaler appliance sends a large number of packets on a TCP connection, and the network randomly drops a few of the packets, multiple sets of continuous packet loss ("holes") are created. When the appliance retransmits the packets, the network interface card (NIC) drops packets.

  [# 643929]

- A NetScaler appliance might not honor persistence for a load balancing virtual server with a wildcard configuration if information about the back-end server is not available.

  [# 556385]

- In a high availability environment, if you add Network Time Protocol (NTP) to a primary node by specifying the NTP server's DNS name, the command is not propagated to the secondary node.

  Workaround: Specify the NTP server's IP address.

  [# 639529]

- The HTML page rendering might fail if you insert a prebody script before the header tag. The HTML specification requires the character-encoding declaration to be serialized within the first 1024 bytes of the document, and the script might push the meta tag past the 1024-byte limit.

  [# 305196, 393696]

- When transmitting a TCP packet, a NetScaler appliance reuses the same IP-ID for packet retransmission. This impacts the customer if a firewall, Intrusion Prevention System (IPS) or Intrusion Detection System (IDS) drops the packet during retransmission.

  [# 670056]

- Data might be dropped when a client requests a small window size. When client sends a small window size (less than 8190 bytes) in its request packet to a NetScaler appliance, the appliance advertises a window size of 8190 bytes to the back-end server. Upon receiving this information, the server sends up to 8190 bytes of data to the appliance, and in turn the appliance, in transparent mode, sends the same amount of data to the client, even if the actual window size is less than the window size advertised by the client. If a device between the appliance and client checks the window size before accepting the data, that device might drop the data that does not fit in the client's window size.

  Workaround: Enable the end point processing features on NetScaler to control the complete TCP stack independently. Such features are TCP Buffering, SSL Offload etc

  [# 622573]

**Telco**

- In a high availability setup, forcing synchronization does not synchronize Port Control Protocol (PCP) mappings to the secondary node.

  [# 647630]

# What's New in Previous NetScaler 11.1 Releases

The enhancements and changes that were available in NetScaler 11.1 releases prior to Build 53.11. The build number provided below the issue description indicates the build in which this enhancement or change was provided.

**AAA-TM**

- OAuth Support for Multi-Factor Authentication

  The NetScaler appliance now supports OAuth in a multifactor deployment and for cascading authentication. That is, OAuth can be now be used anywhere in a cascade, in first factor or in any of the factors, and as a fallback authentication policy. In earlier releases, OAuth could be used only for the first factor. Note: To use OAuth in a factor other than the first, you must register an authentication FQDN with the application because OAuth must start and end on the same virtual server.

  [From Build 41.26] [# 611735]

- OAuth Support for Multi-Factor Authentication

  The NetScaler appliance now supports OAuth in a multifactor deployment and for cascading authentication. That is, OAuth can be now be used anywhere in a cascade, in first factor or in any of the factors, and as a fallback authentication policy.

  In earlier releases, OAuth could be used only for the first factor.

  Note: To use OAuth in a factor other than the first, you must register an authentication FQDN with the application because OAuth must start and end on the same virtual server.

  [From Build 47.14] [# 611735, 572701, 572705]

- This enhancement allows the user to Preview the custom portal themes by binding it to an Authentication virtual server. Earlier this support was only present for Gateway virtual servers.

  [From Build 47.14] [# 620908]

- At present, customization of the Portal pages is only offered for Gateway virtual server. Admins often have the same branding requirements for the Login page that is presented on the Authentication virtual server page - for example, tmindex.html. This enhancement supports Portal Theme binding for the Authentication virtual server.

  [From Build 47.14] [# 581544, 475585, 552072, 606858, 619869]

- You can now change the credential default behavior by defining the loginschema so that the desired credentials (username and password) are used for SSO. To use the first factor for the SSO, you configure the loginschema to store the first factor credential at the specified indexes and use attribute expressions for the traffic policies.

  Previously, multiple sets of login credentials were required for nFactor authentication. By default, the credentials used for the final factor were the default single sign-on (SSO) user name and password. If the first factor was LDAP (Lightweight Directory Access Protocol) but the second factor OTP (One Time Password) on a non-Active Directory password, the default credentials became OTP. This procedure was complex and affected usability.

Configuration:

> set authentication loginSchema ls1 -SSOCredentials YES Done

> set authentication loginSchema ls1 -SSOCredentials NO Done

[From Build 49.16] [# 647382]

**Admin Partitions**

- On a partitioned NetScaler appliance, you can now bind a VLAN as a dedicated VLAN for a particular partition or as a shared VLAN across multiple partitions.

  [From Build 41.26] [# 581671]

- Shared VLAN Support
  On a partitioned NetScaler appliance, you can now bind a VLAN as a dedicated VLAN for a particular partition or as a shared VLAN across multiple partitions.

  [From Build 47.14] [# 581671]

- Binding System Group to Administrative Partition

  In a partitioned NetScaler appliance, you can now bind a system group to a specific administrative partition by using the bind system group <grpname>-partitionname <partitionname> command.

  [From Build 51.21] [# 629434]

- Role-based access (RBA) for System Groups

  Admin partitions now provide role based authentication for system groups. With this access control mechanism, a NetScaler appliance supports the following actions:

  1. Bind an existing partition or all partitions to a system group.

  2. Authenticate a user (bound to a system group), using local or external authentication, and allow the user to switch to a partition that is bound to the system group.

  3. Bind the system group to a custom command policy.

  [From Build 51.21] [# 627888]

- Instant Visibility of the HA Status of Partitions

On a partitioned NetScaler appliance in a high availability configuration, the top pane of the NetScaler GUI displays the high availability status of the partitions. This instant visibility helps you monitor the HA configuration efficiently.

[From Build 51.21] [# 628478]

- A group user associated with a superuser command policy is unable to switch partitions through the NetScaler GUI.

  [From Build 51.21] [# 627770]

- Role-based Access in an Administrative Partition

  As the root administrator of a partitioned NetScaler appliance, you can now designate partition administrators to control user access to entities within specific partitions. A partition administrator can provide granular, role-based access for a partition user and specify a set of permissions and allowed operations. The authorization is specific to the partition. The partition administrator and the users authorized by the partition administrator access the partition through a SNIP address.

  [From Build 51.21] [# 594425]

- Binding System Group to Administrative Partition

  In a partitioned NetScaler appliance, you can now bind a system group to a specific administrative partition by using the bind system group <grpname>-partitioname <partitioname> command.

  [From Build 51.26] [# 629434]

- A group user associated with a superuser command policy is unable to switch partitions through the NetScaler GUI.

  [From Build 51.26] [# 627770]

- Role-based access (RBA) for System Groups

  Admin partitions now provide role based authentication for system groups. With this access control mechanism, a NetScaler appliance supports the following actions:

  1. Bind an existing partition or all partitions to a system group.

  2. Authenticate a user (bound to a system group), using local or external authentication, and allow the user to switch to a partition that is bound to the system group.

  3. Bind the system group to a custom command policy.

  [From Build 51.26] [# 627888]

- Instant Visibility of the HA Status of Partitions

On a partitioned NetScaler appliance in a high availability configuration, the top pane of the NetScaler GUI displays the high availability status of the partitions. This instant visibility helps you monitor the HA configuration efficiently.

[From Build 51.26] [# 628478]

- Role-based Access in an Administrative Partition

  As the root administrator of a partitioned NetScaler appliance, you can now designate partition administrators to control user access to entities within specific partitions. A partition administrator can provide granular, role-based access for a partition user and specify a set of permissions and allowed operations. The authorization is specific to the partition. The partition administrator and the users authorized by the partition administrator access the partition through a SNIP address.

  [From Build 51.26] [# 594425]

**AppExpert**

- Rate Limiting at the Packet Level

  You can configure a stream selector and a responder policy to collect statistics at the packet level and identify defective or attack-prone packets flowing through all the connections identified by the selector. If, at any point, the percentage of defective or attack-prone packets exceeds the configured threshold, the policy applies a corrective action (RESET or DROP).

  [From Build 51.21] [# 615910]

- Rate Limiting at the Packet Level

  You can configure a stream selector and a responder policy to collect statistics at the packet level and identify defective or attack-prone packets flowing through all the connections identified by the selector. If, at any point, the percentage of defective or attack-prone packets exceeds the configured threshold, the policy applies a corrective action (RESET or DROP).

  [From Build 51.26] [# 615910]

**Cluster**

- PBR Support for Cluster

  Partially striped and spotted policy based routes (PBR) are now supported on a Layer 3 NetScaler cluster.

  [From Build 41.26] [# 611938]

- PBR Support for Cluster

Partially striped and spotted policy based routes (PBR) are now supported on a Layer 3 NetScaler cluster.

[From Build 47.14] [# 611938]

- SNMP MIB Support for Cluster Nodes

  In a cluster setup, you can now configure the SNMP MIB in any node by including the ownerNode parameter in the set snmp mib command. Without this parameter, the set snmp mib command applies only to the cluster coordinator node.

  To display the MIB configuration for an individual node other than the cluster coordinator node, include the ownerNode parameter in the show snmp mib command.

  [From Build 49.16] [# 628136, 623888]

- You can now avoid closing a node's connections when you add the node to or remove it from a cluster. Before adding or removing a node, log on to the cluster IP (CLIP) address and set the "retain connections on cluster" option. Then log on to the node's NSIP address and specify a timeout interval for graceful shutdown.

  [From Build 51.21] [# 635529, 634785]

- LLDP Support in a Cluster Setup

  LLDP is a layer 2 protocol that enables a NetScaler appliance to advertise its identity and capabilities to the directly connected (neighbor) devices, and to learn the identity and capabilities of these neighbor devices. In a cluster setup, the NetScaler GUI and NetScaler CLI now display the LLDP neighbour configuration of all or specific cluster nodes when the GUI or CLI is accessed through the Cluster IP address (CLIP). Any change made to the global level LLDP mode is applied to the global level LLDP mode on each of the cluster nodes.

  [From Build 51.21] [# 470187]

- You can now avoid closing a node's connections when you add the node to or remove it from a cluster. Before adding or removing a node, log on to the cluster IP (CLIP) address and set the "retain connections on cluster" option. Then log on to the node's NSIP address and specify a timeout interval for graceful shutdown.

  [From Build 51.26] [# 635529, 634785]

- LLDP Support in a Cluster Setup

  LLDP is a layer 2 protocol that enables a NetScaler appliance to advertise its identity and capabilities to the directly connected (neighbor) devices, and to learn the identity and capabilities of these neighbor devices. In a cluster setup, the NetScaler GUI and NetScaler CLI now display the LLDP neighbour configuration of all or specific cluster nodes when the GUI or CLI is accessed through the Cluster IP address (CLIP). Any change made to the global level LLDP mode is applied to the global level LLDP mode on each of the cluster nodes.

[From Build 51.26] [# 470187]

**DNS**

- Preserving NetScaler Memory by Limiting the Memory Consumed by DNS Cache

  You can now limit the amount of memory consumed by the DNS cache. You can specify the maximum cache size (in MB), and also the cache size (in MB) for storing negative responses. When either limit is reached, no more entries are added to the cache. Also, SNMP traps are generated and syslog messages are logged.

  The maximum cache size is set using the maxCacheSize parameter and the cache size for negative responses is set using the maxNegativeCacheSize parameter.

  This limitation is added per packet engine. For example, if the maxCacheSize is set to 5 MB and the appliance has 3 packet engines, then the total configured cache size is 15 MB.

  [From Build 52.13] [# 665068]

- Enabling DNS Cache Bypass

  You can now configure the cacheHitBypass parameter so that the cache is built but not used. When this parameter is enabled, the requests bypass the DNS cache and are sent to the back-end servers. When this parameter is disabled, the NetScaler appliance starts responding from the cache that has been built so far.

  [From Build 52.13] [# 665073]

- Retaining DNS Records in Cache

  You can now retain the cache that is built so far and prevent it from being aged out. You can do so by using the cacheNoExpire parameter. When this parameter is enabled, the entries in the DNS cache is retained. When this parameter is disabled, the records are flushed out when the TTL expires.

  This option can be used only when the maximum cache size (maxCacheSize parameter) is specified.

  [From Build 52.13] [# 665070]

- Generating SNMP Traps and Syslog Messages When the Memory Consumed by DNS Cache Reaches the Limits Set for Caching DNS Records and Negative Responses

  You can now limit the amount of memory consumed by the DNS cache. You can specify the maximum cache size (in MB), and also the cache size (in MB) for storing negative responses. When either limit is reached, no more entries are added to the cache. Also, SNMP traps are generated and syslog messages are logged.

  [From Build 52.13] [# 665532]

**GSLB**

- Support for EDNS0 Client Subnet

  The NetScaler appliance now supports the EDNS0 client subnet (ECS) option in deployments that include the NetScaler appliance configured as an ADNS server authoritative for a GSLB domain. In the deployment, if you use static proximity as the load balancing method, you can now use the IP subnet in the ECS option, instead of using the LDNS IP address, to determine the geographical proximity of the client. In the case of proxy mode deployment, the appliance forwards a DNS query with the ECS option as-is to the back-end servers and does not cache DNS responses that include the ECS option.

  Note: The EDNS0 client subnet (ECS) option is not applicable for some other deployment modes, such as ADNS mode for non-GSLB domains, resolver mode, and forwarder mode. In such modes, the ECS option is ignored by the NetScaler appliance.

  [From Build 41.26] [# 457159]

- Support for EDNS0 Client Subnet

  The NetScaler appliance now supports the EDNS0 client subnet (ECS) option in deployments that include the NetScaler appliance configured as an ADNS server authoritative for a GSLB domain. In the deployment, if you use static proximity as the load balancing method, you can now use the IP subnet in the ECS option, instead of using the LDNS IP address, to determine the geographical proximity of the client. In the case of proxy mode deployment, the appliance forwards a DNS query with the ECS option as-is to the back-end servers and does not cache the DNS responses that include ECS option.

  Note: The EDNS0 client subnet (ECS) option is not applicable for some other deployment modes, such as ADNS mode for non-GSLB domains, resolver mode, and forwarder mode. In such modes, the ECS option is ignored by the NetScaler appliance.

  [From Build 47.14] [# 457159]

- Configuring GSLB by Using a Wizard in the NetScaler GUI

  You can now use a wizard to configure the GSLB deployment types (active-active and active-passive) and parent-child topologies. In the NetScaler GUI, navigate to Configuration > Traffic Management > GSLB, and click Get Started.

  You can also start the GSLB configuration wizard from the dashboard. The dashboard provides the overall status of the GSLB sites participating in GSLB. You can also synchronize the sites and test the GSLB setup from the dashboard. To access the GSLB dashboard, navigate to Configuration > Traffic Management > GSLB > Dashboard.

  [From Build 51.21] [# 664467]

- Backing UP a Parent Site in a Parent-Child Deployment

The backup parent site topology is useful in scenarios wherein a large number of child sites are associated with a parent site. If this parent site goes DOWN, all of its child sites become unavailable. To prevent this, you can now configure a backup parent site to which the child sites can connect if the original parent site is DOWN.

[From Build 51.21] [# 605605]

- Real-time Synchronization of the GSLB Configuration

  When you create or change the GSLB configuration on a master site, you can use the new AutomaticConfigSync option to automatically synchronize the slave sites.

  When AutomaticConfigSync option is enabled, you do not have to manually trigger the AutoSync option.

  [From Build 51.21] [# 605595]

- Time Delay for Setting a Site as DOWN When Metrics Exchange Protocol Connection to a Remote Site is DOWN

  In a GSLB high availability setup, if the status of a Metrics Exchange Protocol (MEP) connection to a remote site changes to DOWN, you can set a delay to allow some time for reestablishment of the MEP connection before the site is marked as DOWN. If the MEP connection is back UP before the delay expires, the services are not affected.

  [From Build 51.21] [# 621435]

- Time Delay for Setting a Site as DOWN When Metrics Exchange Protocol Connection to a Remote Site is DOWN

  In a GSLB high availability setup, if the status of a Metrics Exchange Protocol (MEP) connection to a remote site changes to DOWN, you can set a delay to allow some time for reestablishment of the MEP connection before the site is marked as DOWN. If the MEP connection is back UP before the delay expires, the services are not affected.

  [From Build 51.26] [# 621435]

- Backing UP a Parent Site in a Parent-Child Deployment

  The backup parent site topology is useful in scenarios wherein a large number of child sites are associated with a parent site. If this parent site goes DOWN, all of its child sites become unavailable. To prevent this, you can now configure a backup parent site to which the child sites can connect if the original parent site is DOWN.

  [From Build 51.26] [# 605605]

- Testing the GSLB Setup

  You can test the GSLB setup to make sure that the ADNS services or the DNS servers are responding with the correct IP address for the domain name that is configured in the GSLB setup.

  This is supported in NetScaler GUI only.

[From Build 51.26] [# 664467]

- Real-time Synchronization of the GSLB Configuration

  When you create or change the GSLB configuration on a master site, you can use the new AutomaticConfigSync option to automatically synchronize the slave sites.

  When AutomaticConfigSync option is enabled, you do not have to manually trigger the AutoSync option.

  [From Build 51.26] [# 605595]

**Load Balancing**

- Closing Monitor Connections at Service and Service Group Level

  A parameter named monConnectionClose has been added at the service and service group levels. If this parameter is not set, the monitor connection is closed by using the value set in the global load balancing parameters. If this parameter is set at the service or service group level, the monitor connection is closed by sending a connection termination message, with the FIN or RESET bit set, to the service or service group.

  [From Build 41.26] [# 607661]

- NetScaler appliances now support load balancing virtual servers of type SSL_FIX, which can load balance FIX-protocol requests at the FIX message level and allow FIX-specific session persistence.

  [From Build 41.26] [# 634096]

- Secure FTP Monitoring Support

  The NetScaler appliance now supports secure FTP monitoring. That is, you can now configure the appliance to send secure FTP probes to your FTP services.

  [From Build 47.14] [# 237766]

- Support for Load Balancing Profile

  A load balancing configuration has a large number of parameters, so setting the same parameters on a number of virtual servers can become tedious. You can now set load balancing parameters in a profile and associate this profile with virtual servers, instead of setting these parameters on each virtual server.

  [From Build 47.14] [# 353669]

- Closing Monitor Connections at the Service Group Level

A parameter named monConnectionClose has been added at the service group level. If this parameter is not set, the monitor connection is closed by using the value set in the global load balancing parameters. If this parameter is set at the service group level, the monitor connection is closed by sending a connection termination message, with the FIN or RESET bit set, to the service group.

[From Build 47.14] [# 628111]

- Improved Support for Persistency

  In certain cases, cores of a NetScaler appliance might not be synchronized, because a core-to-core monitoring or service update has not reached one of the cores. For example, if the core that owns persistency has not received notification that a service is DOWN, that service remains in the persistency table. If a traffic-owner core that has been notified that the service is DOWN finds it in the persistency table, it requests a different service from the persistency-owner core, so that it can redirect the request. Before this enhancement, if the persistency owner returned the same service, the traffic-owner core dropped the user's request. Now, instead of immediately dropping the request, the traffic owner queries the persistency owner a second time. Sending the second query usually gives the persistency owner enough time to have received the update, in which case it returns a different service.

  [From Build 47.14] [# 571771]

- Closing Monitor Connections at the Service Level

  A parameter named monConnectionClose has been added at the service level. If this parameter is not set, the monitor connection is closed by using the value set in the global load balancing parameters. If this parameter is set at the service level, the monitor connection is closed by sending a connection termination message, with the FIN or RESET bit set, to the service.

  [From Build 47.14] [# 607661]

- Configuring an HTTPS Virtual Server to accept HTTP Traffic

  You can now configure an HTTPS virtual server to also process all HTTP traffic. That is, if HTTP traffic is received on the HTTPS virtual server, the appliance internally prepends "https://" to the incoming URL or redirects the traffic to another HTTPS URL, depending on the option configured.

  [From Build 47.14] [# 570157]

- Setting SSL Parameters on a Secure Monitor

  A monitor inherits either the global settings or the settings of the service to which it is bound. If a monitor is bound to a non-SSL or non-SSL_TCP service, such as SSL_BRIDGE, you cannot configure it with SSL settings such as the protocol version or the ciphers to be used. Therefore, in such deployments, SSL-based monitoring of the back-end servers is ineffective.

This enhancement gives you more control over SSL-based monitoring of back-end servers, by enabling you to bind an SSL profile to a monitor. An SSL profile contains SSL parameters, cipher bindings, and ECC bindings. For example, you can set server authentication, ciphers, and protocol version in an SSL profile and bind the profile to a monitor. Note that to perform server authentication, you must also bind a CA certificate to a monitor. To perform client authentication, you must bind a client certificate to the monitor. New parameters for the "bind lb monitor" command enable you to do so.

Note: The SSL settings take effect only if you add a secure monitor. Also, the SSL profile type must be BackEnd.

SSL profiles can be bound to the following monitor types:

- HTTP

- HTTP-ECV

- TCP

- TCP-ECV

- HTTP-INLINE

To specify an SSL profile while adding a monitor by using the command line

At the command prompt, type:

add lb monitor <monitorName> <type> -secure YES -sslprofile <string>

set lb monitor <monitorName> <type> -secure YES -sslprofile <string>

Example:

add ssl profile prof1 -sslProfileType BackEnd

add lb monitor mon1 HTTP -secure YES -sslprofile prof1

To bind a certificate-key pair to a monitor by using the command line

At the command prompt, type:

bind monitor <monitor name> -certkeyName <string> [(-CA [-crlCheck ( Mandatory | Optional ) | -ocspCheck ( Mandatory | Optional )]

[From Build 47.14] [# 506771]

- Required Unbind Operation Prevents Accidentally Disabling a Virtual Server

Accidentally deleting a service or service group that is bound to a virtual server can result in the virtual server going DOWN. With this release, you cannot delete a service or service group that is bound to a virtual server until you first unbind it from the virtual server.

[From Build 47.14] [# 258327]

- FIX Protocol Support

  NetScaler appliances now support load balancing virtual servers of type SSL_FIX, which can load balance FIX-protocol requests at the FIX message level and allow FIX-specific session persistence.

  [From Build 47.14] [# 634096]

- Support for Reverse TCP Monitors

  The NetScaler appliance now supports reverse TCP monitors. A reverse monitor marks the service as DOWN if the probe criteria are satisfied and UP if they are not satisfied.

  A direct TCP monitor marks the service as DOWN if it receives a RESET in response to the monitor probe. However, a reverse TCP monitor treats RESET as a successful response and marks the service as UP.

  To configure a reverse TCP monitor by using the NetScaler command line

  At the command prompt, type:

  add lb monitor <monitor-name> tcp -reverse yes -destip <primary-service ip> -destport <primary-service port>

  bind service <svc-name> -monitorname <monitor-name>

  To configure a reverse TCP monitor by using the NetScaler GUI

  1. Navigate to Traffic Management > Load Balancing > Monitors.

  2. Create a TCP monitor and select Reverse.

  [From Build 49.16] [# 630159]

**NITRO**

- Automate NetScaler Upgrade and Downgrade with a Single API

  A new API, install, can be used to upgrade or downgrade a NetScaler appliance. You can specify a local or remote location for the build file used to upgrade or downgrade the appliance.

  [From Build 41.26] [# 598557]

- Handle Multiple NITRO Calls in a Single Request

  A new API, macroapi, can be used to configure a set of homogeneous or heterogeneous objects in a single API request. The query parameter "onerror" specifies the action to be taken if an error is encountered. Possible values for this parameter are exit, continue, and rollback.

  [From Build 41.26] [# 598559]

- Simplify Management Operations with an idempotent API

  You can add or update resources seamlessly, with a single API, by using the new "idempotent" query parameter. Previously, an attempt to add a resource that was already configured, or to update a resource that was not yet configured, caused an error.

  Now, if you include "idempotent=yes" in a POST request, NITRO executes the request in an idempotent manner.

  [From Build 41.26] [# 601351]

- Retrieve Bindings in Bulk

  You can use a bulk GET API to fetch bindings of all the entities of a given entity type.

  For example, you can fetch bindings of all the load balancing virtual servers in one call instead of by using multiple GET by "name" calls.

  [From Build 41.26] [# 600350]

- Support for ping and traceroute commands

  You can now direct ping and traceroute operations to any host, by using the NITRO API through the NetScaler appliance.

  [From Build 49.16] [# 406603]

**NetScaler CLI**

- 1) A new system parameter was added. "totalAuthTimeout" - the default value is 20 seconds, minimum value 5 seconds and maximum 120 seconds. set system parameter - totalAuthTimeout <positive_integer>

  2) A new aaa radius param was added. authservRetry - the default value is 3 retries, minimum 1 and maximum 10 retries can be configured.

  set aaa radiusParams - authservRetry <positive_integer>

  [From Build 47.14] [# 492179]

- Force Password Change

  The default root credentials for a NetScaler appliance is "nsroot". However, for security reasons, you might enforce a password change to ensure the credentials are changed to a new value other than the default value. To implement this, a new parameter, "forcePasswordChange" is introduced.

  If you, as a root administrator log on with default credentials and set forcePasswordChange to ENABLED, on your next subsequent logon attempt, you will be prompted to change the password, and will not be allowed to log on without doing so. After the password is changed, the prompt no longer appears.

  Note: You are prompted to change the current password to a new one only if the ForcePasswordChange parameter is enabled. Otherwise, you can access the appliance with the default login credentials (user name: NSROOT, password: NSROOT).

  [From Build 51.21] [# 490116, 638504]

- Force Password Change

  The default root credentials for a NetScaler appliance is "nsroot". However, for security reasons, you might enforce a password change to ensure the credentials are changed to a new value other than the default value. To implement this, a new parameter, "forcePasswordChange" is introduced.

  If you, as a root administrator log on with default credentials and set forcePasswordChange to ENABLED, on your next subsequent logon attempt, you will be prompted to change the password, and will not be allowed to log on without doing so. After the password is changed, the prompt no longer appears.

  Note: You are prompted to change the current password to a new one only if the ForcePasswordChange parameter is enabled. Otherwise, you can access the appliance with the default login credentials (user name: NSROOT, password: NSROOT).

  [From Build 51.26] [# 490116, 638504]

**NetScaler CPX**

- Container-Based Application Delivery Controller

  Citrix NetScaler CPX is a container-based application delivery controller that can be provisioned on a Docker host. NetScaler CPX enables customers to leverage Docker engine capabilities and use NetScaler load balancing and traffic management features for container-based applications. You can deploy one or more NetScaler CPX instances as standalone instances on a Docker host. For more information, see About NetScaler CPX: http://docs.citrix.com/en-us/netscaler-cpx/11-1/about-netscaler-cpx.html.

  [From Build 47.14] [# 627953, 632576]

- Open Source packages are now available in NetScaler CPX

All the open source packages that are used in NetScaler CPX are available in the contrib/cpx/ folder.

[From Build 49.16] [# 652842]

- New End User License Agreement (EULA) for NetScaler CPX Express

  You now need to accept an End User License Agreement (EULA) to install and use the NetScaler CPX Express.

  The End User Licensing Agreement is available at: https://www.microloadbalancer.com/eula.

  [From Build 51.21] [# 656632]

- New End User License Agreement (EULA) for NetScaler CPX Express

  You now need to accept an End User License Agreement (EULA) to install and use the NetScaler CPX Express.

  The End User Licensing Agreement is available at: https://www.microloadbalancer.com/eula.

  [From Build 51.26] [# 656632]

**NetScaler GUI**

- Support for High Availability Configuration for a Secure Access Only Remote node

  The NetScaler GUI now supports configuration of a node in High Availability (HA) mode even if the Secure Access Only option is enabled for the NetScaler IP (NSIP) address of the other node in the HA pair.

  [From Build 47.14] [# 624858]

- High Availability Status Information in the Top Pane

  The top pane of the NetScaler GUI now displays the High Availability status of the node. This instant visibility of HA status helps you monitor the HA configuration efficiently.

  [From Build 47.14] [# 423777, 466239, 582803]

- Usability Support to Upload Technical Support Collector Archive

  You can now automatically upload the technical support collector archive to Citrix Support servers.

  Navigate to System > Diagnostics > Technical Support Tools > Generate support file, and select Upload the Collector Archive. Type your user credentials and click Run.

  [From Build 47.14] [# 614285, 620953]

- Diagnostic of Start New Trace and Support Stop Running Trace

  Starting and stopping nstrace are now separate options in the NetScaler GUI. As a result, it is easier to stop a running trace and download the results.

  Navigate to System > Diagnostics and select "Start New Trace" or "Stop Running Trace."

  [From Build 47.14] [# 564499, 565594]

- In the load balancing visualizer, you can now seamlessly migrate the configuration of a service to all the services bound to the virtual server. To copy the settings of one service to all the other services, in the visualizer, click "Configuration Sets," select a service, and then click "Migrate Config."

  [From Build 47.14] [# 619498]

- Tabular, One-page Application Firewall Wizard

  The new, tabular, Application Firewall wizard improves flexibility and accelerates the completion of tasks. You can go back to any page and edit any details about profiles, policies, and signatures, and skip screens that are not mandatory. In addition, all resource-consuming tasks, such as submission and binding, are completed after you click Finish.

  [From Build 47.14] [# 587433, 557185, 619712]

- Icons for Action and Information Menus

  Two new icons in the NetScaler GUI display action menus and information menus. If you are in a window with detail-view rows, and the rows have actions, you can now display the action menu by clicking the action icon in that row, rather than right-clicking the row.

  Similarly, you can display the info menu by clicking the info icon.

  [From Build 47.14] [# 614868]

- SSL Certificate Management GUI Enhancements and Changes

  1) Links to the following pages have been removed from the SSL overview page:

  - Create RSA Key

  - Create DSA Key

  - Create CSR

  - Create Certificate

To access these pages, navigate to Traffic Management > SSL > SSL Files.

2) Server, client, and CA certificates are now segregated. When you bind a certificate to an SSL end point, only the list of appropriate certificates appears. For example, when you bind a server certificate to an SSL virtual server, only the server certificates are listed. In earlier releases, all the certificates, including client and CA certificates, were listed.

3) You can configure an SNMP trap from the "Install Certificate" page to send a notification when the certificate is about to expire. For a valid certificate in the notification period, status changes to yellow. For an expired certificate, status changes to red.

4) "Certificate format" field has been removed, because the format (PEM/DER/PFX/Bundle) is automatically detected by the software during certificate installation. Also, if the file is not password protected, you are not prompted for a password.

5) The key files, CSR files, and certificate files are segregated onto different tabs for ease of use.

6) The SSL certificate overview page now explains the end-to-end flow of managing certificates on your appliance.

[From Build 47.14] [# 612894]

- To test connectivity from a subnet IP (SNIP) address to another IP address, you can now select the source address from a list of SNIP addresses instead of typing the SNIP address. If the SNIP address is not in the list, you can add it. To use this feature, navigate to System > Diagnostics. In Utilities, select ping or ping6, and then select "SNIP."

  [From Build 47.14] [# 597501]

- The NetScaler appliance was enhanced so Negotiate Authentication is available for VPN Virtual Servers. The GUI reflects this under the NSG > Policies > Authentication node.

  [From Build 47.14] [# 600708]

- Improved IPv4 Address Fields

  IPv4 address fields now do not have dot separators, which improve the usability of these fields.

  [From Build 47.14] [# 610522]

**NetScaler Insight Center**

- You can now assign IPv4, IPv6, or both IP addresses to your NetScaler Insight Center server. To assign a new IP address, navigate to Configuration > System > Network Configuration, and select IPv4, IPv6, and/or both and specify the network parameters.

  If you specify both IPv4 and IPv6 addresses, you can access the NetScaler Insight Center by using any one of the IP addresses.

[From Build 47.14] [# 582943]

- You can now view USB event reports of a user's active sessions on HDX Insight. You can view details such as USB Status, Number of USB Instances Accepted, Number of USB Instances Rejected, and Number of USB Instances Stopped.

  [From Build 47.14] [# 549746]

- You can now search for a specific application, client, or server by using the Search option in Web Insight.

  [From Build 47.14] [# 590782]

- In HDX Insight, you can view a diagrammatic representation of a client's current session details.

  Navigate to HDX insight > Users, and in the Current Sessions section, click the Diagram button to display details such as Client IP address, NetScaler IP address, Origin Server IP address, Country, Region, Session ID, and Client Version.

  [From Build 47.14] [# 606189]

- 1)You can now use NetScaler Insight Center to monitor and manage your incoming traffic's IP Reputation.

  2)You can now enable/disable AppFlow for Security insight separately from the Enable AppFlow option. To Enable or Disable AppFlow, navigate to Configuration > Inventory and open your NetScaler Instance. Select the virtual servers and click Enable AppFlow and select the Security Insight option.

  [From Build 47.14] [# 635528]

- You can now enable, edit, or clear AppFlow on multiple virtual servers simultaneously. To perform these actions, navigate to Configuration > Inventory and open your NetScaler Instance. Select the virtual servers and click Enable AppFlow, Edit AppFlow Settings, or Clear AppFlow Configuration, respectively.

  [From Build 47.14] [# 534805]

- You can now view the client's machine name for any user session in HDX Insight.

  [From Build 47.14] [# 606187]

- NetScaler Insight Center can now use the X-Forwarded-For header to display the actual client IP address instead of the IP address of the proxy that forwarded the request.

  [From Build 47.14] [# 541439]

- You can now use NetScaler Insight Center to monitor NetScaler integrated caching. Cache Insight enables you to see and monitor the various actions performed by the NetScaler cache.

[From Build 47.14] [# 498439]

- You can now view the current session details from the Geomaps section in HDX Insight.

  [From Build 47.14] [# 606188]

- The following thin clients now support HDX Insight:

  -WYSE Windows based thin clients

  -WYSE Linux based thin clients

  -WYSE ThinOS based thin clients

  -10Zig Ubuntu based thin clients

  [From Build 47.14] [# 614892, 550997, 604388, 620422, 632370]

- You can only enable or disable the X-Forwarded-For feature using the NetScaler appliance's CLI. To enable this feature, at the command prompt, type: "set appflow param httpXForwardedFor ENABLED".

  [From Build 51.21] [# 643724]

- You can only enable or disable the X-Forwarded-For feature using the NetScaler appliance's CLI. To enable this feature, at the command prompt, type: "set appflow param httpXForwardedFor ENABLED".

  [From Build 51.26] [# 643724]

**NetScaler VPX Appliance**

- New license for NetScaler VPX on ESX and KVM platforms

  40G license is now available for NetScaler VPX appliance on ESX and KVM platforms

  For more information about recommended interfaces and performance details, refer to the latest VPX datasheet.

  [From Build 47.14] [# 623179]

- New license for NetScaler VPX on KVM platforms

  A 100G license is now available for NetScaler VPX appliances on a KVM platform.

  For more information about recommended interfaces and performance details, see the latest VPX datasheet.

  [From Build 49.16] [# 660256]

- The number of unique IPv6 addresses that you can add to a NetScaler virtual appliance configured with SR-IOV interfaces is limited to 30 on the following platforms:

  * Linux-KVM

  * VMware ESX

  [From Build 49.16] [# 639229]

- Support for PCI Passthrough Interfaces on NetScaler VPX Appliances Installed on VMware ESX Server

  You can now configure a NetScaler VPX instance deployed on VMware ESX Server to use PCI passthrough interfaces.

  For performance information about PCI passthrough interfaces on ESX Server, see the latest VPX datasheet.

  [From Build 51.21] [# 661840]

- Support for PCI Passthrough Interfaces on NetScaler VPX Appliances Installed on VMware ESX Server

  You can now configure a NetScaler VPX instance deployed on VMware ESX Server to use PCI passthrough interfaces.

  For performance information about PCI passthrough interfaces on ESX Server, see the latest VPX datasheet.

  [From Build 51.26] [# 661840]

**Networking**

- NetScaler Support for Microsoft Direct Access Deployment

  Microsoft Direct Access is a technology that enables remote users to seamlessly and securely connect to enterprise's internal networks, without the need to establish a separate VPN connection. Unlike VPN connections, which require user intervention to start and close connections, a Direct Access-enabled client connects automatically to the enterprise's internal networks whenever the client connects to the Internet.

  Manage-Out is a Microsoft Direct Access feature that allows administrators inside the enterprise network to connect to Direct Access clients outside the network and manage them (for example, performing administration tasks, such as scheduling service updates, and providing remote support.

  In a Direct Access deployment, NetScaler appliances provide high availability, scalability, high performance, and security. NetScaler load balancing functionality sends client traffic through the most appropriate server. The appliances can also forward the Manage-Out traffic through the right path to reach the client.

  [From Build 41.26] [# 612455]

- IPv6 Support in Active-Active Mode using VRRP

NetScaler Appliances Support VIP6 Addresses in Active-Active Deployments.

An active-active deployment, in addition to preventing downtime, makes efficient use of all the NetScaler appliances in the deployment. In an IPv6 active-active deployment mode, the same VIP6 address is assigned to every NetScaler appliance in the configuration, but with different priorities, so that a given VIP6 can be active on only one appliance at a time.

The active VIP6 address is called the master VIP6, and the corresponding VIP6s on the other NetScaler appliances are called the backup VIP6s. If a master VIP6 fails, the backup VIP6 with the highest priority takes over and becomes the master VIP6. All the NetScaler appliances in an active-active deployment use the Virtual Router Redundancy Protocol (VRRP) to advertise their VIP6s and the corresponding priorities at regular intervals.

NetScaler appliances in active-active mode can be configured so that no appliance is idle. In this configuration, different sets of VIPs are active on each appliance.

The following features of IPv4 active-active configuration are also supported for IPv6 active-active configuration:

* Preemption

* Delaying preemption

* Sharing

* Changing VIP address priority automatically

[From Build 41.26] [# 553570]

- Graceful Restart for Dynamic Routing Protocols

  In a non-INC high availability (HA) setup in which a routing protocol is configured, after a failover, routing protocol is converged and routes between the new primary node and the adjacent neighbor routers are learned. Route learning take some time to complete. During this time, forwarding of packets is delayed, network performance might get disrupted, and packets might get dropped.

  Graceful restart enables an HA setup during a failover to direct its adjacent routers to not remove the old primary node's learned routes from their routing databases. Using the old primary node's routing information, the new primary node and the adjacent routers immediately start forwarding packets, without disrupting network performance.

  The following routing protocols support graceful restart in a non-INC high availability setup:

  - Border Gateway Protocol (BGP)

  - IPv6 Border Gateway protocol (IPv6 BGP)

- Open Shortest Path First (OSPF)

- IPv6 Open Shortest Path First (OSPFv3)

[From Build 41.26] [# 571033]

- Adding Default Route for the changed NSIP address Before a Restart

  If you change the NSIP address of a NetScaler appliance, you can now add a default route to the new address's subnet before restarting the NetScaler appliance. This change makes the new NSIP address accessible from other networks after the appliance is restarted.

  In previous releases, if the subnet address of the new NSIP address is different from the previous one, you cannot add a default route for this new subnet until you restart the appliance. Because of this restriction, the new NSIP address is unreachable from other networks after a restart.

  [From Build 47.14] [# 551505]

- Configuring Source IP Persistency for Backend Communication

  By default, for a load balancing configuration with the USIP option disabled and a net profile bound to a virtual server or services or service groups, the NetScaler appliance uses the round-robin algorithm to select an IP address from the net profile for communicating with the servers. Because of this selection method, the IP address selected can be different for different sessions of a specific client.

  Some situations require that the NetScaler appliance sends all of a specific client's traffic from the same IP address when sending the traffic to servers. The servers can then, for example, identify traffic belonging to a specific set for logging and monitoring purposes.

  The source IP persistency option of a net profile enables the NetScaler appliance to use the same address, specified in the net profile, to communicate with servers for all sessions initiated from a specific client to a virtual server

  [From Build 47.14] [# 530670]

- Stateful Connection Failover Support for RNAT

  Connection failover helps prevent disruption of access to applications deployed in a distributed environment. The NetScaler appliance now supports stateful connection failover for connections related to RNAT rules in a NetScaler High Availability (HA) setup.

  In an HA setup, connection failover (or connection mirroring) refers to the process of keeping an established TCP or UDP connection active when a failover occurs. The primary appliance sends messages to the secondary appliance to synchronize current information about the RNAT connections. The secondary appliance uses this connection information only in the event of a failover. When a failover occurs, the new primary NetScaler appliance has information about the connections established before the failover and hence continues to serve those connections

even after the failover. From the client's perspective this failover is transparent. During the transition period, the client and server may experience a brief disruption and retransmissions.

Connection failover can be enabled per RNAT rule. For enabling connection failover on an RNAT rule, you enable the connFailover (Connection Failover) parameter of that specific RNAT rule by using either NetScaler command line or configuration utility. Also, you must disable the tcpproxy (TCP Proxy) parameter globally for all RNAT rules in order for connection failover to work properly for TCP connections.

[From Build 47.14] [# 457167]

- NITRO API Support for Dynamic Routing

  NetScaler appliances now support NITRO API for configuring dynamic routing protocols.

  [From Build 47.14] [# 626083]

- IPv6 Support in Active-Active Mode using VRRP

  NetScaler Appliances Support VIP6 Addresses in Active-Active Deployments.

  An active-active deployment, in addition to preventing downtime, makes efficient use of all the NetScaler appliances in the deployment. In an IPv6 active-active deployment mode, the same VIP6 address is assigned to every NetScaler appliance in the configuration, but with different priorities, so that a given VIP6 can be active on only one appliance at a time.

  The active VIP6 address is called the master VIP6, and the corresponding VIP6s on the other NetScaler appliances are called the backup VIP6s. If a master VIP6 fails, the backup VIP6 with the highest priority takes over and becomes the master VIP6. All the NetScaler appliances in an active-active deployment use the Virtual Router Redundancy Protocol (VRRP) to advertise their VIP6s and the corresponding priorities at regular intervals.

  NetScaler appliances in active-active mode can be configured so that no appliance is idle. In this configuration, different sets of VIPs are active on each appliance.

  The following features of IPv4 active-active configuration are also supported for IPv6 active-active configuration:

  * Preemption

  * Delaying preemption

  * Sharing

  * Changing VIP address priority automatically

  [From Build 47.14] [# 553570]

- Configuring Allowed VLAN List

  NetScaler accepts and sends tagged packets of a VLAN on an interface if the VLAN is explicitly configured on the NetScaler appliance and the interface is bound to the VLAN. Some deployments (for example, Bump in the wire) require the NetScaler appliance to function as a transparent device to accept and forward tagged packets related to a large number of VLANs. For this requirement, configuring and managing a large number of VLANs is not a feasible solution.

  Allowed VLAN list on an interface specifies a list of VLANs. The interface transparently accepts and sends tagged packets related to the specified VLANs without the need for explicitly configuring these VLANs on the appliance.

  [From Build 47.14] [# 495219]

- Managing High Availability Heartbeat Messages on a NetScaler Appliance

  The two nodes in a high availability configuration send and receive heartbeat messages to and from each other on all interfaces that are enabled. The heartbeat messages flow regardless of the HA MON setting on these interfaces. If NSVLAN or SYNCVLAN or both are configured on an appliance, the heartbeat messages flow only through the enabled interfaces that are part of the NSVLAN and SYNCVLAN.

  If a node does not receive the heartbeat messages on an enabled interface, it sends critical alerts to the specified Command Center and SNMP managers. These critical alerts give false alarms and draw unnecessary attention from the administrators for interfaces that are not configured as part of the connections to the peer node.

  To resolve this issue, the HAHeartBeat option for interfaces and channels is used for enabling or disabling HA heartbeat-message flow on them.

  [From Build 47.14] [# 477162, 575447, 604578]

- Using a Source Port from a Specified Port Range for Backend Communication

  By default, for configurations with USIP option disabled or with USIP and use proxy port options enabled, the NetScaler appliance communicates to the servers from a random source port (greater than 1024).

  The NetScaler supports using a source port from a specified port range for communicating to the servers. One of the use case of this feature is for servers that are configured to identify received traffic belonging to a specific set on the basis of source port for logging and monitoring purposes. For example, identifying internal and external traffic for logging purpose.

  [From Build 47.14] [# 420067, 420039]

- Setting the MTU on the NSVLAN

  By default, the MTU of the NSVLAN is set to 1500 bytes. You can now modify this setting to optimize throughput and network performance. For example, you can configure the NSVLAN to process jumbo frames.

[From Build 47.14] [# 425950]

- Graceful Restart for Dynamic Routing Protocols

  In a non-INC high availability (HA) setup in which a routing protocol is configured, after a failover, routing protocol is converged and routes between the new primary node and the adjacent neighbor routers are learned. Route learning take some time to complete. During this time, forwarding of packets is delayed, network performance might get disrupted, and packets might get dropped.

  Graceful restart enables an HA setup during a failover to direct its adjacent routers to not remove the old primary node's learned routes from their routing databases. Using the old primary node's routing information, the new primary node and the adjacent routers immediately start forwarding packets, without disrupting network performance.

  The following routing protocols support graceful restart in a non-INC high availability setup:

  - Border Gateway Protocol (BGP)

  - IPv6 Border Gateway protocol (IPv6 BGP)

  - Open Shortest Path First (OSPF)

  - IPv6 Open Shortest Path First (OSPFv3)

  [From Build 47.14] [# 571033]

- Dynamic Routing support for Link-Local Subnet IPv6 addresses

  NetScaler appliances now support dynamic routing on a link-local Subnet IPv6 (SNIP6) address for a VLAN. In a default admin partition, link-local SNIP6 address takes precedence over the link-local NSIP6 address for running dynamic routing on a VLAN. In a non-default partition, the NetScaler appliance does not support dynamic routing on link-local NSIP6 address for a VLAN. Link-local SNIP6 address can now be used for running dynamic routing on the VLAN.

  [From Build 47.14] [# 553544]

- Logging Start Time and Connection Closure Reasons in RNAT Log Entries

  For diagnosing or troubleshooting problems related to RNAT connections, the NetScaler appliance now logs the following additional information:

  - Start time of the RNAT session.

  - Reason for closure of the RNAT session. The NetScaler appliance logs closure reason for TCP RNAT sessions that do not use the TCP proxy (TCP proxy disabled) of the appliance. The following are the type of closure reasons that are logged for TCP RNAT sessions:

-- TCP FIN. The RNAT session was closed because of a TCP FIN sent by either the source or destination device.

-- TCP RST. The RNAT session was closed because of a TCP Reset that was sent by either the source or destination device.

-- TIMEOUT. The RNAT session timed out.

[From Build 47.14] [# 609410]

- NetScaler Support for Microsoft Direct Access Deployment

  Microsoft Direct Access is a technology that enables remote users to seamlessly and securely connect to enterprise's internal networks, without the need to establish a separate VPN connection. Unlike VPN connections, which require user intervention to start and close connections, a Direct Access-enabled client connects automatically to the enterprise's internal networks whenever the client connects to the Internet.

  Manage-Out is a Microsoft Direct Access feature that allows administrators inside the enterprise network to connect to Direct Access clients outside the network and manage them (for example, performing administration tasks, such as scheduling service updates, and providing remote support).

  In a Direct Access deployment, NetScaler appliances provide high availability, scalability, high performance, and security. NetScaler load balancing functionality sends client traffic through the most appropriate server. The appliances can also forward the Manage-Out traffic through the right path to reach the client.

  [From Build 47.14] [# 612455]

- Using NULL Policy Based Routes to Drop Outgoing Packets

  Some situations might demand that the NetScaler appliance drops specific outgoing packets instead of routing them, for example, in testing cases and during deployment migration. NULL policy based routes can be used to drop specific outgoing packets. A NULL PBR is a type of PBR that has the nexthop parameter set to NULL. The NetScaler appliance drops outgoing packets that match a NULL PBR.

  [From Build 47.14] [# 451632]

- Network Service Header support for Service Function

  Network Services Header (NSH) is a new standard that enables the Service Function Chaining (SFC) architecture. NSH enables you to define the service chain paths and forward the data-plane traffic through multiple service nodes in a dynamic and fail-proof manner.

  A NetScaler appliance can now play the service-function role in a SFC architecture. The NetScaler appliance receives packets with Network Service headers and, upon performing the service, modifies the NSH bits in the response packet to indicate that the service has been performed. In that role, the appliance supports symmetric service chaining with

features (for example, INAT, TCP and UDP load balancing services, and routing). The NetScaler appliance as service-function does not support IPv6 and Reclassification.

[From Build 47.14] [# 593459]

- Wildcard TOS Monitors

  In a load balancing configuration in DSR mode using TOS field, monitoring its services requires a TOS monitor to be created and bound to these services. A separate TOS monitor is required for each load balancing configuration in DSR mode using TOS field, because a TOS monitor requires the VIP address and the TOS ID to create an encoded value of the VIP address. The monitor creates probe packets in which the TOS field is set to the encoded value of the VIP address. It then sends the probe packets to the servers represented by the services of a load balancing configuration. With a large number of load balancing configurations, creating a separate custom TOS monitor for each configuration is a big, cumbersome task. Managing these TOS monitors is also a big task. Now, you can create wildcard TOS monitors. You need to create only one wildcard TOS monitor for all load balancing configurations that use the same protocol (for example, TCP or UDP).

  A wildcard TOS monitor has the following mandatory settings:

  -Type = <protocol>

  -TOS = Yes

  The following parameters can be set to a value or can be left blank:

  -Destination IP

  -Destination Port

  -TOS ID

  A wildcard TOS monitor (with destination IP, Destination port, and TOS ID not set) bound to a DSR service automatically learns the TOS ID and the VIP address of the load balancing virtual server. The monitor creates probe packets with TOS field set to the encoded VIP address and then sends the probe packets to the server represented by the DSR service.

  [From Build 49.16] [# 615975]

- Support of Automatic ARP Resolution to Special MAC address

  In a cluster deployment, when the client-side or server side-link to a node goes down, traffic is steered to this node through the peer nodes for processing. Previously, the steering of traffic was implemented on all nodes by configuring dynamic routing and adding static ARP entries pointing to the special MAC address of each node. If there are a large number of nodes in a cluster deployment, adding and managing static ARP entries with special MAC addresses on all

the nodes is a cumbersome task. Now, nodes implicitly use special MAC addresses for steering packets. Therefore, static ARP entries pointing to special MAC addresses no longer have to be added to the cluster nodes.

[From Build 49.16] [# 635235]

- Monitoring Command Propagation Failures in a Cluster Deployment

  In a cluster deployment of NetScaler appliances, you can use the new command "show prop status" for faster monitoring and troubleshooting of issues related to command-propagation failure on non-CCO nodes. This command displays up to 20 of the most recent command propagation failures on all non-CCO nodes. You can use either the NetScaler command line or the NetScaler GUI to perform this operation after accessing them through the CLIP address or through the NSIP address of any node in the cluster deployment.

  [From Build 49.16] [# 623707]

- Automatic TCP-Connection Reset for Inactive Nodes

  Previously, a cluster node did not reset its existing TCP connections (to clients and servers) when its state became Inactive. As a result, the states of the client and server connections became undefined. Now, a node resets all its TCP connections before entering the Inactive state.

  [From Build 49.16] [# 635826]

- Support for Sending Response Traffic Through an IP-IP tunnel

  You can now configure a NetScaler appliance to send response traffic through an IP-IP tunnel instead of routing it back to the source. Previously, when the appliance received a request from another NetScaler or a third-party device through an IP-IP tunnel, it had to route the response traffic instead of sending it through the tunnel. You can now use policy based routes (PBRs) or enable MAC-Based Forwarding (MBF) to send the response through the tunnel.

  In a PBR rule, specify the subnets at both end points whose traffic is to traverse the tunnel. Also set the next hop as the tunnel name. When response traffic matches the PBR rule, the NetScaler appliance sends the traffic through the tunnel.

  Alternatively, you can enable MBF to meet this requirement, but the functionality is limited to traffic for which the NetScaler appliance stores session information (for example, traffic related to load balancing or RNAT configurations). The appliance uses the session information to send the response traffic through the tunnel.

  [From Build 49.16] [# 632279]

- Loop Prevention Mechanism based on VLAN ID

  For a MAC-mode based load balancing configuration, the NetScaler appliance maintains a source MAC table. This table maps the virtual server to the MAC addresses of all the bound services. The appliance uses this table to prevent (loop prevention mechanism) the server traffic from reaching the virtual server.

For a trunk link that is shared by the VLANs of the servers and the VLANs of clients, the appliance also prevents traffic from these clients from reaching the virtual server. To solve this issue, the NetScaler loop prevention mechanism now considers the VLAN ID along with the MAC address, so that the client traffic in a trunk link reaches the virtual server.

[From Build 51.21] [# 663400]

- Advertisement of SNIP and VIP Routes to Selective Areas

  In a cluster setup, for a requirement to advertise spotted SNIP addresses to only the server-side routers, enabling DRADV mode or redistribute connect ZebOS operations cannot be used. This is because these operations send all the connected routes to ZebOS. Also, adding dummy static routes in ZebOS for the required subnets, or adding ACLs in ZebOS to filter unwanted connected routes is a cumbersome and tedious task.

  A new option, Network Route, addresses this issue. You can enable this option for only one SNIP address per subnet. The connected route for that SNIP address is sent as a kernel route to ZebOS.

  For VIP and SNIP addresses, another new option, Tag, can be assigned an integer from 1 to 4294967295. This parameter can be set only when Host Route or Network Route is enabled for VIP or SNIP addresses. The tag value associated with VIP and SNIP addresses are also sent along with their routes to ZebOS. Tags with different values can be set for VIP and SNIP routes. These tag values can then be matched in routemaps in ZebOS and advertised to selective areas.

  [From Build 51.21] [# 633418]

- Stateful Connection Failover Support for RNAT configurations with TCP Proxy On

  Connection failover helps prevent disruption of access to applications deployed in a distributed environment. In a High Availability (HA) setup, stateful connection failover for RNAT is now supported with TCP proxy.

  Connection failover can be enabled per RNAT rule. For enabling connection failover on an RNAT rule, you enable the "connFailover" ("Connection Failover") parameter of that specific RNAT rule. To enable TCP proxy for RNAT, you must enable "tcpproxy" parameter by using the "set rnatparam" command in the NetScaler CLI or select "Enable RNAT Source IP Persistency" (System > Setting > Change Global System Settings) in the NetScaler GUI.

  [From Build 51.21] [# 439206]

- Stateful Connection Failover Support for RNAT configurations with TCP Proxy On

  Connection failover helps prevent disruption of access to applications deployed in a distributed environment. In a High Availability (HA) setup, stateful connection failover for RNAT is now supported with TCP proxy.

  Connection failover can be enabled per RNAT rule. For enabling connection failover on an RNAT rule, you enable the "connFailover" ("Connection Failover") parameter of that specific RNAT rule. To enable TCP proxy for RNAT, you must enable "tcpproxy" parameter by using the "set rnatparam" command in the NetScaler CLI or select "Enable RNAT Source IP Persistency" (System > Setting > Change Global System Settings) in the NetScaler GUI.

[From Build 51.26] [# 439206]

- Advertisement of SNIP and VIP Routes to Selective Areas

  In a cluster setup, for a requirement to advertise spotted SNIP addresses to only the server-side routers, enabling DRADV mode or redistribute connect ZebOS operations cannot be used. This is because these operations send all the connected routes to ZebOS. Also, adding dummy static routes in ZebOS for the required subnets, or adding ACLs in ZebOS to filter unwanted connected routes is a cumbersome and tedious task.

  A new option, Network Route, addresses this issue. You can enable this option for only one SNIP address per subnet. The connected route for that SNIP address is sent as a kernel route to ZebOS.

  For VIP and SNIP addresses, another new option, Tag, can be assigned an integer from 1 to 4294967295. This parameter can be set only when Host Route or Network Route is enabled for VIP or SNIP addresses. The tag value associated with VIP and SNIP addresses are also sent along with their routes to ZebOS. Tags with different values can be set for VIP and SNIP routes. These tag values can then be matched in routemaps in ZebOS and advertised to selective areas.

  [From Build 51.26] [# 633418]

- Loop Prevention Mechanism based on VLAN ID

  For a MAC-mode based load balancing configuration, the NetScaler appliance maintains a source MAC table. This table maps the virtual server to the MAC addresses of all the bound services. The appliance uses this table to prevent (loop prevention mechanism) the server traffic from reaching the virtual server.

  For a trunk link that is shared by the VLANs of the servers and the VLANs of clients, the appliance also prevents traffic from these clients from reaching the virtual server. To solve this issue, the NetScaler loop prevention mechanism now considers the VLAN ID along with the MAC address, so that the client traffic in a trunk link reaches the virtual server.

  [From Build 51.26] [# 663400]

- Disabled ACL logging for Loopback Traffic

  By default, the NetScaler appliance bypasses ACL processing for loopback traffic, but it logs the loopback traffic for ACL rules for which the ACL logging option is enabled. These log entries for loopback traffic create a false impression that the NetScaler appliance has processed loopback traffic for ACL rules.

  Now, the NetScaler appliance does not log loopback traffic for ACL rules.

  [From Build 52.13] [# 671305]

**SSL**

- Support for SNI on the Back-End Service

The NetScaler appliance now supports Server Name Indication (SNI) at the back end. That is, the common name is sent as the server name in the client hello to the back-end server for successful completion of the handshake. In addition to helping meet federal system integrator customer security requirements, this enhancement provides the advantage of using only one port instead of opening hundreds of different IP addresses and ports on a firewall.

Federal system integrator customer security requirements include support for Active Directory Federation Services (ADFS) 3.0 in 2012R2 and WAP servers. This requires supporting SNI at the back end on a NetScaler appliance.

[From Build 41.26] [# 471431, 559271, 595785]

- Segregation of Certificates According to Type

  To facilitate certificate selection, certificates are now segregated according to type, such as server certificate, client certificate, and CA certificate.

  To view the certificates in the GUI, navigate to Traffic Management > SSL > Certificates.

  To view the certificates in the CLI, type "show ssl certkey"

  [From Build 47.14] [# 620923, 623890]

- Support for ECC curves in Service Groups

  You can now bind ECC curves to back-end service groups by using the NetScaler command line.

  At the command prompt, type:

  bind ssl serviceGroup <serviceGroupName> -eccCurveName <eccCurveName>

  [From Build 47.14] [# 592418]

- Support to create a Certificate Signing Request signed with the SHA256 Digest Algorithm

  The NetScaler appliance supports creating a CSR signed with the SHA256 digest algorithm. The encryption hash algorithm used in SHA256 makes it stronger than SHA1.

  [From Build 47.14] [# 606874, 595902]

- Support for AES-GCM/SHA2 ciphers on the front-end of VPX appliances

  The NetScaler VPX appliance now supports AES-GCM/SHA2 ciphers on the front end.

  [From Build 47.14] [# 498207]

- New Counters at the SSL Virtual Server Level and at the Global Level

Six counters have been added to the output of the "stat ssl vserver" command, as follows:

1. ssl_ctx_tot_enc_bytes: Tracks the number of encrypted bytes.

2. ssl_ctx_tot_dec_bytes: Tracks the number of decrypted bytes.

3. ssl_ctx_tot_hw_enc_bytes: Tracks the number of hardware encrypted bytes.

4. ssl_ctx_tot_hw_dec_bytes: Tracks the number of hardware decrypted bytes.

5. ssl_ctx_tot_session_new: Tracks the number of new sessions created.

6. ssl_ctx_tot_session_hits: Tracks the number of session hits.

Five counters have been added to the output of the "stat ssl -detail" command, as follows:

1. ssl_tot_sslServerInRecords: Tracks the number of SSL records processed by the appliance.

2. ssl_cur_sslInfo_SPCBInUseCount: Tracks the number of SSL protocol control blocks (SPCBs) used at any given point.

2. ssl_cur_session_inuse: Tracks the number of active SSL sessions.

4. ssl_cur_sslInfo_cardinBlkQ: Tracks the number of bulk encryption and decryption operations that are pending for card.

5. ssl_cur_sslInfo_cardinKeyQ: Tracks the number of handshake-related operations that are pending for card.

[From Build 47.14] [# 597279, 582601]

Removing RC4-MD5 cipher from the default cipher list

The RC4-MD5 cipher is removed from the list of default ciphers that are supported on a NetScaler appliance.

[From Build 47.14] [# 258311]

- Support for SNI on the Back-End Service

  The NetScaler appliance now supports Server Name Indication (SNI) at the back end. That is, the common name is sent as the server name in the client hello to the back-end server for successful completion of the handshake. In addition to helping meet federal system integrator customer security requirements, this enhancement provides the advantage of using only one port instead of opening hundreds of different IP addresses and ports on a firewall.

  Federal system integrator customer security requirements include support for Active Directory Federation Services (ADFS) 3.0 in 2012R2 and WAP servers. This requires supporting SNI at the back end on a NetScaler appliance.

[From Build 47.14] [# 471431, 559271, 595785]

- Optimizing ECDHE Computation

  ECDHE-RSA computation has been optimized by using a combination of software and hardware offload capabilities.

  [From Build 50.10] [# 643480]

- Support for TLS Session Ticket Extension

  An SSL handshake is a CPU-intensive operation. If session reuse is enabled, the server/client key exchange operation is skipped for existing clients. They are allowed to resume their sessions. This improves the response time and increases the number of SSL transactions per second that a server can support. However, the server must store details of each session state, which consumes memory and is difficult to share among multiple servers if requests are load balanced across servers.

  NetScaler appliances now support the SessionTicket TLS extension. Use of this extension indicates that the session details are stored on the client instead of on the server. The client must indicate that it supports this mechanism by including the session ticket TLS extension in the client Hello message. For new clients, this extension is empty. The server sends a new session ticket in the NewSessionTicket handshake message. The session ticket is encrypted with a key known only to the server. If a server cannot issue a new ticket at this time, it completes a regular handshake.

  To resume a session, the client must include the session ticket in the request. If, for any reason, the server does not honor the ticket, it attempts to initiate a full handshake with the client.

  [From Build 51.21] [# 416800, 577122, 648240]

- Providing the Revocation Status of a Server Certificate to a Client

  To avoid unnecessary congestion when each client requests the revocation status of a server certificate during an SSL handshake, the NetScaler appliance now supports OCSP stapling. That is, the appliance can now send the revocation status of a server certificate to a client, at the time of the SSL handshake, after validating the certificate status from an OCSP responder. The revocation status of a server certificate is "stapled" to the response the appliance sends to the client as part of the SSL handshake. To use the OCSP stapling feature, you must enable it on an SSL virtual server and add an OCSP responder on the appliance.

  Note: NetScaler appliances support OCSP stapling as defined in RFC 6066.

  Important: NetScaler support for OCSP stapling is limited to handshakes using TLS protocol version 1.0 or higher. This feature is not supported in a cluster setup.

  [From Build 51.21] [# 367538]

- Support for Client Certificate Thumbprint

NetScaler appliances now support inserting the thumbprint (also called a fingerprint) of a certificate into the header of a request sent to a back-end server. If client authentication is enabled, the appliance computes the thumbprint of the certificate, and uses an SSL policy action to insert the thumbprint into the request. The server searches for the thumbprint, and grants secure access if there is a match.

[From Build 51.21] [# 537629, 632507]

- Providing the Revocation Status of a Server Certificate to a Client

  To avoid unnecessary congestion when each client requests the revocation status of a server certificate during an SSL handshake, the NetScaler appliance now supports OCSP stapling. That is, the appliance can now send the revocation status of a server certificate to a client, at the time of the SSL handshake, after validating the certificate status from an OCSP responder. The revocation status of a server certificate is "stapled" to the response the appliance sends to the client as part of the SSL handshake. To use the OCSP stapling feature, you must enable it on an SSL virtual server and add an OCSP responder on the appliance.

  Note: NetScaler appliances support OCSP stapling as defined in RFC 6066.

  Important: NetScaler support for OCSP stapling is limited to handshakes using TLS protocol version 1.0 or higher. This feature is not supported in a cluster setup.

  [From Build 51.26] [# 367538]

- Support for Client Certificate Thumbprint

  NetScaler appliances now support inserting the thumbprint (also called a fingerprint) of a certificate into the header of a request sent to a back-end server. If client authentication is enabled, the appliance computes the thumbprint of the certificate, and uses an SSL policy action to insert the thumbprint into the request. The server searches for the thumbprint, and grants secure access if there is a match.

  [From Build 51.26] [# 537629, 632507]

- Support for TLS Session Ticket Extension

  An SSL handshake is a CPU-intensive operation. If session reuse is enabled, the server/client key exchange operation is skipped for existing clients. They are allowed to resume their sessions. This improves the response time and increases the number of SSL transactions per second that a server can support. However, the server must store details of each session state, which consumes memory and is difficult to share among multiple servers if requests are load balanced across servers.

  NetScaler appliances now support the SessionTicket TLS extension. Use of this extension indicates that the session details are stored on the client instead of on the server. The client must indicate that it supports this mechanism by including the session ticket TLS extension in the client Hello message. For new clients, this extension is empty. The server sends a new session ticket in the NewSessionTicket handshake message. The session ticket is encrypted with a key known only to the server. If a server cannot issue a new ticket at this time, it completes a regular handshake.

To resume a session, the client must include the session ticket in the request. If, for any reason, the server does not honor the ticket, it attempts to initiate a full handshake with the client.

[From Build 51.26] [# 416800, 577122, 648240]

- Send Certificates to Back-End Servers as Strings, without Spaces

  A CLI command has been added to send a certificate to a back-end server as a string without any spaces, instead of in its original format (with spaces). Previously, you had to use an nsapimgr option to do this.

  At the NetScaler CLI, type:

  set ssl parameter -insertCertSpace ( YES | NO )

  [From Build 52.13] [# 661342]

**Security**

- Configuring DNS Security Options from the Add DNS Security Profile Page in the NetScaler GUI

  You can now configure the DNS security options from the Add DNS Security Profile page in the NetScaler GUI. This page provides a user-friendly graphical user interface for configuring DNS security settings. The Cache Poisoning Protection option is always enabled. The other security options can be applied to all DNS endpoints or to specific DNS virtual server(s) in your deployment.

  Two of the security options, Bypass the Cache and Provide root details in the DNS response, can be applied to all DNS endpoints. The following security options can be applied either to all DNS endpoints or to specific DNS virtual servers:

  DNS DDoS protection

  Manage exceptions - whitelist/blacklist servers

  Prevent random subdomain attacks

  Enforce DNS transactions over TCP

  [From Build 51.21] [# 617479]

- Configuring DNS Security Options from the Add DNS Security Profile Page in the NetScaler GUI

  You can now configure the DNS security options from the Add DNS Security Profile page in the NetScaler GUI. This page provides a user-friendly graphical user interface for configuring DNS security settings. The Cache Poisoning Protection option is always enabled. The other security options can be applied to all DNS endpoints or to specific DNS virtual server(s) in your deployment.

Two of the security options, Bypass the Cache and Provide root details in the DNS response, can be applied to all DNS endpoints. The following security options can be applied either to all DNS endpoints or to specific DNS virtual servers:

DNS DDoS protection

Manage exceptions - whitelist/blacklist servers

Prevent random subdomain attacks

Enforce DNS transactions over TCP

[From Build 51.26] [# 617479]

**System**

- TCP Fast Open (TFO) is a TCP mechanism that enables speedy and safe data exchange between a client and a server during TCP's initial handshake. This feature is available as a TCP option in the TCP profile bound to a virtual server of a NetScaler appliance. TFO uses a TCP Fast Open Cookie (a security cookie) that the NetScaler appliance generates to validate and authenticate the client initiating a TFO connection to the virtual server. By using the TFO mechanism, you can reduce an application's network latency and the delay experienced in short TCP transfers.

  [From Build 41.26] [# 358990]

- A new slow-start algorithm, Hybrid Start (Hystart) is configured as a TCP option in the relevant TCP profile bound to a virtual server. This algorithm dynamically determines a safe point at which to terminate (ssthresh) and enables a transition to avoid congestion with heavy packet losses. This option is disabled by default.

  [From Build 41.26] [# 603099]

- The "start nstrace" command has a new parameter, -capsslkeys, with which you can capture the SSL master keys for all SSL sessions. If the capsslkeys option is enabled, a file named nstrace.sslkeys is generated along with the packet trace and imported into Wireshark to decrypt the SSL traffic in the trace file.

  [From Build 41.26] [# 603225]

- MAC Address is tied to the IP Address in case of an IP Conflict

  An SNMP trap that is sent as a result of an IP address conflict now contains the MAC address of the device. You can therefore identify the device by its MAC address. Previously, identifying the device was not possible, because the conflict lasts for only a short time.

  [From Build 47.14] [# 570372, 524621]

- Capturing SSL Keys during NetScaler Trace

The "start nstrace" command has a new parameter, -capsslkeys, with which you can capture the SSL master keys for all SSL sessions. If the capsslkeys option is enabled, a file named nstrace.sslkeys is generated along with the packet trace and imported into Wireshark to decrypt the SSL traffic in the trace file.

[From Build 47.14] [# 603225]

- TCP Hystart Algorithm

  A new slow-start algorithm, Hybrid Start (Hystart) is configured as a TCP option in the relevant TCP profile bound to a virtual server. This algorithm dynamically determines a safe point at which to terminate (ssthresh) and enables a transition to avoid congestion with heavy packet losses. This option is disabled by default.

  [From Build 47.14] [# 603099]

- Dynamic TCP Buffer Management

  When you enable the Dynamic Receive Buffer option in a TCP profile, the NetScaler appliance can dynamically adjust the TCP receive buffer size for optimized memory usage based on the congestion window.

  [From Build 47.14] [# 628115]

- TCP Fast Open Mechanism

  TCP Fast Open (TFO) is a TCP mechanism that enables speedy and safe data exchange between a client and a server during TCP's initial handshake. This feature is available as a TCP option in the TCP profile bound to a virtual server of a NetScaler appliance. TFO uses a TCP Fast Open Cookie (a cryptographic cookie) that the NetScaler appliance generates to validate the client initiating a TFO connection to the virtual server. By using the TFO mechanism, you can reduce an application's network latency and the delay experienced in short TCP transfers.

  [From Build 47.14] [# 358990]

- Proportional Rate Recovery Algorithm

  The Proportional Rate Recovery (PRR) algorithm is a fast recovery algorithm that evaluates TCP data during a loss recovery. It is patterned after Rate-Halving, by using the fraction that is appropriate for the target window chosen by the congestion control algorithm. It minimizes window adjustment, so that the actual window size at the end of recovery is close to the Slow-Start threshold (ssthresh).

  [From Build 47.14] [# 473777]

- Warning about an Unsaved NetScaler Configuration

  The NetScaler GUI displays a Save icon with a red dot when a running configuration is not saved. A unsaved configuration could be lost if a power outage or restart occurs.

To save the configuration(s), you can click the Save icon and then click Yes at the configuration prompt. When you return to the main screen by clicking OK, the icon is white.

Note: In some cases, the red dot might appear even though there is no unsaved configuration. In that case, if you click the Save icon, the following message appears: "The running configuration has not changed."

[From Build 47.14] [# 626225]

- Bridge Group Support for Cluster

  Bridge Group functionality is now supported on a Layer 3 NetScaler cluster.

  [From Build 47.14] [# 587548]

- Configuring SNMP Audit Log Levels

  After you enable the SNMP trap logging option, a NetScaler appliance on which at least one trap listener is configured can log SNMP trap messages (for SNMP alarms in which logging capability is enabled). Now, you can specify the audit log level of trap messages sent to an external log server. The default log level is Informational. Possible values are Emergency, Alert, Critical, Error, Warning, Debug, and Notice.

  For example, you can set the audit log level to Critical for an SNMP trap message generated by a logon failure. That information is then available on the NSLOG or SYSLOG server for troubleshooting.

  [From Build 47.14] [# 569317]

- RDX Error Management

  In the NetScaler GUI, if you skip a mandatory field or make an invalid entry, an error message appears beside the field or in the page header, depending on the type of error, and remains until you enter a valid value. For example, on the Add Virtual Server page, if you enter an invalid server IP address or port number, an error message appears beside the IP Address or Port field, and you cannot submit the page until you correct the error.

  [From Build 47.14] [# 552575]

- Proactive Support for Hardware Errors

  Policy Infrastructure (PI) for Auditlog Framework

  Audit log actions now support advance policies and expressions. Advance policy expressions are very powerful and provide endless use cases to work with. Previously, the audit module supported only classic policies. You can now bind advanced audit-log policies to the syslog and nslog global entities.

  [From Build 49.16] [# 522692, 607221]

- In a NetScaler appliance, if the Ring Receive buffer is full, the appliance starts to discard data packets at the Network Interface Card (NIC). As a result, the appliance drops packets leading to a probe failure.

  [From Build 49.16] [# 623977, 649735]

- Specifying a domain name for a logging server

  When configuring an auditlog action, you can specify the domain name of a syslog or nslog server instead of its IP address. Then, if the server's IP address changes, you do not have to change it on the NetScaler appliance.

  [From Build 49.16] [# 314438]

- TCP Burst Rate Control

  A NetScaler appliance now uses a technique called "TCP Burst Rate Control" for burst management in a high speed mobile network. This technique evenly spaces the flow of data into the network, avoiding bursts by waiting for a period of time before sending the next group of packets. By using this technique, you can achieve better throughput and lower packet drop rates. This feature is available as a TCP option in the TCP profile bound to a virtual server on a NetScaler appliance.

  [From Build 49.16] [# 628114]

- Half-closed or established TCP connections, between clients and a NetScaler appliance, cleaned up by the NetScaler zombie process can now be dropped silently, that is, without sending RST packets to the clients.

  To configure this feature, run the following commands at the NetScaler shell prompt:

  - nsapimgr_wr.sh -ys tcp_hc_zombie_silent_drop=1

  - nsapimgr_wr.sh -ys tcp_est_zombie_silent_drop=1

  [From Build 50.10] [# 656135]

- The TCP timestamp is now an interoperable parameter for TCP and Multipath TCP (MPTCP) data transmission.

  [From Build 50.10] [# 646496]

- By default, a NetScaler appliance ignores the non-standard and obsolete "Proxy-Connection" HTTP header. To change this behavior, use the nsamimgr command to set the proxyConnection parameter to 1. This setting prioritizes the Proxy-Connection header over the Connection header.

  For example, nsapimgr -ys proxyconnection=1

  [From Build 51.21] [# 654560]

- Changes in NetScaler Telco Software Licensing Editions

  The software licensing editions for NetScaler Telco platforms (NetScaler T1000 series and NetScaler VPX-T) have changed as follows:

  Basic edition

  * Features added: Content Filtering

  * Features removed: None

  Advanced edition

  * Features added: AAA, Content Optimization, RDP Proxy, RISE, and Internet On Hold (IOH)

  * Features removed: None

  [From Build 51.21] [# 656361]

- Changes in NetScaler Telco Software Licensing Editions

  The software licensing editions for NetScaler Telco platforms (NetScaler T1000 series and NetScaler VPX-T) have changed as follows:

  Basic edition

  * Features added: Content Filtering

  * Features removed: None

  Advanced edition

  * Features added: AAA, Content Optimization, RDP Proxy, RISE, and Internet On Hold (IOH)

  * Features removed: None

  [From Build 51.26] [# 656361]

- By default, a NetScaler appliance ignores the non-standard and obsolete "Proxy-Connection" HTTP header. To change this behavior, use the nsamimgr command to set the proxyConnection parameter to 1. This setting prioritizes the Proxy-Connection header over the Connection header.

  For example, nsapimgr -ys proxyconnection=1

  [From Build 51.26] [# 654560]

- New Hardware-Script Option Removes Media Errors

  In a hardware script, the new -d option extracts CF, SSD, and HDD media errors from the log files.

  >ns_hw_err.bash -d

  [From Build 52.13] [# 628137]

**Telco**

- Large Scale NAT64

  Because of the imminent exhaustion of IPv4 addresses, ISPs have started transitioning to IPv6 infrastructure. But during the transition, ISPs must continue to support IPv4 along with IPv6, because most of the public Internet still uses IPv4. Large scale NAT64 is an IPv6 transition solution for ISPs with IPv6 infrastructure to connect their IPv6-only subscribers to the IPv4 Internet. DNS64 is a solution for enabling discovery of IPv4-only domains by IPv6-only clients. DNS64 is used with large scale NAT64 to enable seamless communication between IPv6-only clients and IPv4-only servers.

  A NetScaler appliance implements large scale NAT64 and DNS64 and is compliant with RFCs 6145, 6146, 6147, 6052, 3022, 2373, 2765, and 2464.

  The following lists some of the large scale NAT64 features supported on NetScaler appliance:

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

  - Deterministic/Fixed NAT. Support for pre-allocation of blocks of ports to subscribers to minimize logging.

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

  - Filtering. Support of Endpoint-Independent Filtering (EIF), Address-Dependent Filtering (ADF), and Address-Port-Dependent Filtering (APDF).

  - Quotas. Configurable limits on number of ports, sessions per subscriber, and sessions per LSN group.

  - Static Mapping. Support for manually defining a large scale NAT64 mapping.

  - Hairpin Flow. Support for communication between subscribers or internal hosts using NAT IP addresses.

  - 464XLAT connections. Support for communication between IPv4-only aware applications on IPv6 subscriber hosts and IPv4 hosts on the Internet through IPv6 network.

  - Variable length NAT64 and DNS64 prefixes. The NetScaler appliance supports defining NAT64 and DNS64 prefixes of lengths of 32, 40, 48, 56, 64, and 96.

- Multiple NAT64 and DNS64 prefix. The NetScaler appliance supports multiple NAT64 and DNS64 prefixes.

- LSN Clients. Support for specifying or identifying subscribers for large scale NAT64 by using IPv6 prefixes and extended ACL6 rules.

- Logging. Support for logging NAT64 sessions for law enforcement. In addition, the following are also supported for logging.

-- Reliable SYSLOG. Support for 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 reduce the large scale NAT64 log volume.

-- Logging MSISDN information. Support for including subscribers' MSISDN information in large scale NAT64 logs to identify and track subscriber activity over the Internet.

[From Build 41.26] [# 496866]

- Compact Logging for Large Scale NAT

  Logging LSN information is one of the important functions needed by ISPs to meet legal requirements and be able to identify the source of traffic at any given time. This eventually results in a huge volume of log data, requiring the ISPs to make large investments to maintain the logging infrastructure.

  Compact logging is a technique for reducing the log size by using a notational change involving short codes for event and protocol names. For example, C for client, SC for session created, and T for TCP. Compact logging results in an average of 40 percent reduction in log size.

  Compact logging is supported for NAT44, DS-Lite, and NAT64.

  [From Build 41.26] [# 496812]

- NAT44 Wildcards Static Maps

  A static mapping entry is usually a one-to-one LSN mapping between a subscriber IP address:port and a NAT IP address:port. A one-to-one static LSN mapping entry exposes only one port of the subscriber to the Internet.

  Some situations might require exposing all ports (64K) of a subscriber to the Internet (for example, a server hosted on an internal network and running a different service on each port). To make these internal services accessible through the Internet, you have to expose all the ports of the server to the Internet.

One way to meet this requirement is to add 64K one-to-one static mapping entries, one mapping entry for each port. Creating 64K entries is very cumbersome and a big task. Also, this large number of configuration entries might lead to performance issues in the NetScaler appliance.

Another simple method is to use wildcard ports in a static mapping entry. You just need to create one static mapping entry with NAT-port and subscriber-port parameters set to the wildcard character (*), and the protocol parameter set to ALL, to expose all the ports of a subscriber to the Internet. For a subscriber's inbound or outbound connections matching a wildcard static mapping entry, the subscriber's port does not change after the NAT operation.

[From Build 41.26] [# 614784]

- Subscriber Aware LSN Session Termination

  Currently, if a subscriber session is deleted when a RADIUS Accounting STOP or a PCRF-RAR message is received, or as a result of any other event, such as TTL expiry or flush, the corresponding LSN sessions of the subscriber are removed only after the configured LSN timeout period. LSN sessions that are kept open until this timeout expires continue to consume resources on the appliance.

  This enhancement adds a new parameter (subscrSessionRemoval). If this parameter is enabled, and the subscriber information is deleted from the subscriber database, LSN sessions corresponding to that subscriber are also removed. If this parameter is disabled, the subscriber sessions are timed out as specified by the LSN timeout settings.

  [From Build 41.26] [# 578275]

- Port Control Protocol for Large Scale NAT

  NetScaler appliances now support Port Control Protocol (PCP) for large scale NAT (LSN). Many of an ISP's subscriber applications must be accessible from Internet (for example, Internet of Things (IOT) devices, such as an IP camera that provides surveillance over the Internet). One way to meet this requirement is to create static large scale NAT (LSN) maps. But for a very large number of subscribers, creating static LSN NAT maps is not a feasible solution.

  Port Control Protocol (PCP) enables a subscriber to request specific LSN NAT mappings for itself and/or for other 3rd party devices. The large scale NAT device creates an LSN map and sends it to the subscriber. The subscriber sends the remote devices on the Internet the NAT IP address:NAT port at which they can connect to the subscriber.

  Applications usually send frequent keep-alive messages to the large scale NAT device so that their LSN mappings do not time out. PCP helps reduce the frequency of such keep-alive messages by enabling the applications to learn the timeout settings of the LSN mappings. This helps reduce bandwidth consumption on the ISP's access network and battery consumption on mobile devices.

  PCP is a client-server model and runs over the UDP transport protocol. A NetScaler appliance implements the PCP server component and is compliant with RFC 6887. Port Control Protocol is supported for NAT44, DS-Lite and NAT64 on the NetScaler appliance.

  [From Build 41.26] [# 496807]

- Large Scale NAT64

  Because of the imminent exhaustion of IPv4 addresses, ISPs have started transitioning to IPv6 infrastructure. But during the transition, ISPs must continue to support IPv4 along with IPv6, because most of the public Internet still uses IPv4. Large scale NAT64 is an IPv6 transition solution for ISPs with IPv6 infrastructure to connect their IPv6-only subscribers to the IPv4 Internet. DNS64 is a solution for enabling discovery of IPv4-only domains by IPv6-only clients. DNS64 is used with large scale NAT64 to enable seamless communication between IPv6-only clients and IPv4-only servers.

  A NetScaler appliance implements large scale NAT64 and DNS64 and is compliant with RFCs 6145, 6146, 6147, 6052, 3022, 2373, 2765, and 2464.

  The following lists some of the large scale NAT64 features supported on NetScaler appliance:

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

  - Deterministic/Fixed NAT: Support for pre-allocation of blocks of ports to subscribers to minimize logging.

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

  - Filtering: Support of Endpoint-Independent Filtering (EIF), Address-Dependent Filtering (ADF), and Address-Port-Dependent Filtering (APDF).

  - Quotas: Configurable limits on number of ports, sessions per subscriber, and sessions per LSN group.

  - Static Mapping: Support for manually defining a large scale NAT64 mapping.

  - Hairpinning Flow: Support for communication between subscribers or internal hosts using NAT IP addresses.

  - 464XLAT connections: Support for communication between IPv4-only aware applications on IPv6 subscriber hosts and IPv4 hosts on the Internet through IPv6 network.

  - Variable length NAT64 and DNS64 prefixes: The NetScaler appliance supports defining NAT64 and DNS64 prefixes of lengths of 32, 40, 48, 56, 64, and 96.

  - Multiple NAT64 and DNS64 prefix: The NetScaler appliance supports multiple NAT64 and DNS64 prefixes.

  - LSN Clients: Support for specifying or identifying subscribers for large scale NAT64 by using IPv6 prefixes and extended ACL6 rules.

  - Logging: Support for logging NAT64 sessions for law enforcement. In addition, the following are also supported for logging.

-- Reliable SYSLOG: Support for 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 reduce the large scale NAT64 log volume.

-- Logging MSISDN information: Support for including subscribers' MSISDN information in large scale NAT64 logs to identify and track subscriber activity over the Internet.

[From Build 47.14] [# 496866]

- Support for SIP and RTSP ALGs for DS-Lite

  The NetScaler appliance now supports SIP and RTSP application layer gateways (ALGs) for DS-Lite.

  [From Build 47.14] [# 604029]

- Subscriber Aware LSN Session Termination

  Currently, if a subscriber session is deleted when a RADIUS Accounting STOP or a PCRF-RAR message is received, or as a result of any other event, such as TTL expiry or flush, the corresponding LSN sessions of the subscriber are removed only after the configured LSN timeout period. LSN sessions that are kept open until this timeout expires continue to consume resources on the appliance.

  This enhancement adds a new parameter (subscrSessionRemoval). If this parameter is enabled, and the subscriber information is deleted from the subscriber database, LSN sessions corresponding to that subscriber are also removed. If this parameter is disabled, the subscriber sessions are timed out as specified by the LSN timeout settings.

  [From Build 47.14] [# 578275]

- HTTP Header Logging Support for DS-Lite

  The NetScaler appliance can now log request header information of an HTTP connection that is using the NetScaler's DS-Lite functionality. The HTTP header logs can be used by ISPs to see the trends related to the HTTP protocol among a set of subscribers. For example, an ISP can use this feature to find out the most popular website among a set of subscribers.

  [From Build 47.14] [# 558159, 559227]

- NAT44 Wildcards Static Maps

A static mapping entry is usually a one-to-one LSN mapping between a subscriber IP address:port and a NAT IP address:port. A one-to-one static LSN mapping entry exposes only one port of the subscriber to the Internet.

Some situations might require exposing all ports (64K) of a subscriber to the Internet (for example, a server hosted on an internal network and running a different service on each port). To make these internal services accessible through the Internet, you have to expose all the ports of the server to the Internet.

One way to meet this requirement is to add 64K one-to-one static mapping entries, one mapping entry for each port. Creating 64K entries is very cumbersome and a big task. Also, this large number of configuration entries might lead to performance issues in the NetScaler appliance.

Another simple method is to use wildcard ports in a static mapping entry. You just need to create one static mapping entry with NAT-port and subscriber-port parameters set to the wildcard character (*), and the protocol parameter set to ALL, to expose all the ports of a subscriber to the Internet. For a subscriber's inbound or outbound connections matching a wildcard static mapping entry, the subscriber's port does not change after the NAT operation.

[From Build 47.14] [# 614784]

- Policy-based TCP Profile

  You can now configure the NetScaler appliance to perform TCP optimization based on subscriber attributes. For example, the appliance can now select different TCP profiles at run time, based on the network to which the user equipment (UE) is connected. As a result, you can improve a mobile user's experience by setting some parameters in the TCP profiles and then using policies to select the appropriate profile.

  [From Build 47.14] [# 622947]

- Port Control Protocol for Large Scale NAT

  NetScaler appliances now support Port Control Protocol (PCP) for large scale NAT (LSN). Many of an ISP's subscriber applications must be accessible from Internet (for example, Internet of Things (IOT) devices, such as an IP camera that provides surveillance over the Internet). One way to meet this requirement is to create static large scale NAT (LSN) maps. But for a very large number of subscribers, creating static LSN NAT maps is not a feasible solution.

  Port Control Protocol (PCP) enables a subscriber to request specific LSN NAT mappings for itself and/or for other 3rd party devices. The large scale NAT device creates an LSN map and sends it to the subscriber. The subscriber sends the remote devices on the Internet the NAT IP address:NAT port at which they can connect to the subscriber.

  Applications usually send frequent keep-alive messages to the large scale NAT device so that their LSN mappings do not time out. PCP helps reduce the frequency of such keep-alive messages by enabling the applications to learn the timeout settings of the LSN mappings. This helps reduce bandwidth consumption on the ISP's access network and battery consumption on mobile devices.

PCP is a client-server model and runs over the UDP transport protocol. A NetScaler appliance implements the PCP server component and is compliant with RFC 6887. Port Control Protocol is supported for NAT44, DS-Lite and NAT64 on the NetScaler appliance.

[From Build 47.14] [# 496807]

- Global override LSN parameter removed from L3 parameters

  The global override LSN parameter has been removed from L3 parameters. To override LSN, you must now create a net profile with the overrideLsn parameter enabled and bind this profile to all the load balancing virtual servers that are configured for value added services.

  [From Build 47.14] [# 642585]

- Compact Logging for Large Scale NAT

  Logging LSN information is one of the important functions needed by ISPs to meet legal requirements and be able to identify the source of traffic at any given time. This eventually results in a huge volume of log data, requiring the ISPs to make large investments to maintain the logging infrastructure.

  Compact logging is a technique for reducing the log size by using a notational change involving short codes for event and protocol names. For example, C for client, SC for session created, and T for TCP. Compact logging results in an average of 40 percent reduction in log size. Compact logging is supported for NAT44, DS-Lite, and NAT64.

  [From Build 47.14] [# 496812]

- Wildcard Port Static Large Scale NAT64 Maps

  A static large scale NAT64 mapping entry is usually a one-to-one mapping between a subscriber IPv6 address:port and a NAT IPv4 address:port. A one-to-one static large scale NAT64 mapping entry exposes only one port of the subscriber IP address to the Internet.

  Some situations might require exposing all ports (64K - limited to the maximum number of ports of a NAT IPv4 address) of a subscriber IP address to the Internet (for example, a server hosted on an internal network and running a different service on each port). To make these internal services accessible through the Internet, you have to expose all the ports of the server to the Internet.

  One way to meet this requirement is to add 64 thousand one-to-one static mapping entries, one mapping entry for each port. Creating that entries is very cumbersome and a big task. Also, this large number of configuration entries might lead to performance issues in the NetScaler appliance.

  A simpler method is to use wildcard ports in a static mapping entry. You just need to create one static mapping entry with NAT-port and subscriber-port parameters set to the wildcard character (*), and the protocol parameter set to ALL, to expose all the ports of a subscriber IP address for all protocols to the Internet.

For a subscriber’s inbound or outbound connections matching a wildcard static mapping entry, the subscriber’s port does not change after the NAT operation. When a subscriber-initiated connection to the Internet matches a wildcard static mapping entry, the NetScaler appliance assigns a NAT port that has the same number as the subscriber port from which the connection is initiated. Similarly, an Internet host gets connected to a subscriber's port by connecting to the NAT port that has the same number as the subscriber's port.

[From Build 51.21] [# 651078]

- Wildcard Port Static Large Scale NAT64 Maps

  A static large scale NAT64 mapping entry is usually a one-to-one mapping between a subscriber IPv6 address:port and a NAT IPv4 address:port. A one-to-one static large scale NAT64 mapping entry exposes only one port of the subscriber IP address to the Internet.

  Some situations might require exposing all ports (64K - limited to the maximum number of ports of a NAT IPv4 address) of a subscriber IP address to the Internet (for example, a server hosted on an internal network and running a different service on each port). To make these internal services accessible through the Internet, you have to expose all the ports of the server to the Internet.

  One way to meet this requirement is to add 64 thousand one-to-one static mapping entries, one mapping entry for each port. Creating those entries is very cumbersome and a big task. Also, this large number of configuration entries might lead to performance issues in the NetScaler appliance.

  A simpler method is to use wildcard ports in a static mapping entry. You just need to create one static mapping entry with NAT-port and subscriber-port parameters set to the wildcard character (*), and the protocol parameter set to ALL, to expose all the ports of a subscriber IP address for all protocols to the Internet.

  For a subscriber's inbound or outbound connections matching a wildcard static mapping entry, the subscriber's port does not change after the NAT operation. When a subscriber-initiated connection to the Internet matches a wildcard static mapping entry, the NetScaler appliance assigns a NAT port that has the same number as the subscriber port from which the connection is initiated. Similarly, an Internet host gets connected to a subscriber's port by connecting to the NAT port that has the same number as the subscriber's port.

  [From Build 51.26] [# 651078]

# Fixed Issues in Previous NetScaler 11.1 Releases

The issues that were addressed in NetScaler 11.1 releases prior to Build 53.11. The build number provided below the issue description indicates the build in which this issue was addressed.

**AAA-TM**

- In a high availability setup, a session does not time out even if a force timeout is configured on a traffic action that is bound to a load balancing or content switching virtual server and a force fail over is performed.

  [From Build 50.10] [# 623053]

- The StoreFront FQDN is not accepted as valid when a user uses it for the Test Connection function in the XA/XD Wizard. After the StoreFront FQDN is entered, the XA/XD Wizard displays an error when the user clicks Continue.

  [From Build 50.10] [# 612276, 621861, 639203, 650065, 651022]

- In a multi-core NetScaler environment, user sessions sometimes do not get terminated if the decision to terminate is based on a force timeout value that is configured on a TM traffic action.

  [From Build 50.10] [# 610604, 618760]

- If you configure "CLI Accounting" on the NetScaler appliance, the RADIUS server does not send accounting message with Session ID.

  [From Build 51.21] [# 538997]

- In a multifactor SAML IdP configuration, if a SAML request is resent from the service provider during authentication, the NetScaler appliance sends an assertion before authentication is complete.

  [From Build 51.21] [# 666161]

- The NetScaler appliance fails if all of the following conditions are met:

  - The appliance is used as a SAML service provider.

  - Multiple load balancing and content switching virtual servers are configured for the same external identity provider (IdP) but with different FQDN.

  - SAML login happens on a virtual server with an existing SAML session from the same IdP.

  [From Build 51.21] [# 664171, 670657]

- The NetScaler appliance might restart if role-based access is enabled in admin partitions.

  [From Build 51.21] [# 653702]

- If you configure "CLI Accounting" on the NetScaler appliance, the RADIUS server does not send accounting message with Session ID.

  [From Build 51.26] [# 538997]

- The NetScaler appliance fails if all of the following conditions are met:

  - The appliance is used as a SAML service provider.

- Multiple load balancing and content switching virtual servers are configured for the same external identity provider (IdP) but with different FQDN.

- SAML login happens on a virtual server with an existing SAML session from the same IdP.

[From Build 51.26] [# 664171, 670657]

- The NetScaler appliance might restart if role-based access is enabled in admin partitions.

  [From Build 51.26] [# 653702]

- In a multifactor SAML IdP configuration, if a SAML request is resent from the service provider during authentication, the NetScaler appliance sends an assertion before authentication is complete.

  [From Build 51.26] [# 666161]

**Admin Partitions**

- SNMP profiles have been modified to avoid dropping SNMP responses intended for non-default partitions. An SNMP agent can now track each SNMP request and send a response to a non-default partition. Previously, if a non-default partition received an SNMP request through a subnet IP address, the SNMP agent on the partition responded to the default partition, because the SNIP address was defined on the default partition.

  [From Build 49.16] [# 609367]

**AppFlow**

- ICA parsing uses a lot of memory, so the NetScaler appliance reaches its memory limit with a lower than expected number of connections.

  [From Build 49.16] [# 459458]

- If HDX Insight is enabled on a NetScaler appliances in high-availability mode, and if the nodes are set to STAY PRIMARY or STAY SECONDARY, session reliability fails when a failover happens.

  [From Build 49.16] [# 653438]

- A NetScaler load balanced server responds with a 411 error code for a corrupted HTTP request.

  [From Build 50.10] [# 629223]

- If AppFlow clientside measurements and AppFirewall are enabled, due to incomplete and incorrect order of the restore/cleanup of AppFlow and AppFirewall feature, NetScaler might become unresponsive.

  [From Build 50.10] [# 655309, 658547]

- If AppFlow clientside measurements are enabled, NetScaler instance does not buffer the response packets even though it acknowledges the packet to the server. This will cause page load issues if the packets are lost.

[From Build 50.10] [# 670464]

- Applications do not launch when AppFlow is enabled and connection chaining is disabled. This is because when a full sized packet is received, the connection chain ID is added to the packet resulting in the size of packet going beyond the maximum transmission unit (MTU). So, the packet gets dropped and the application fails to launch.

  [From Build 50.10] [# 650618, 653126, 661587, 664792]

- When Web Insight is enabled, and if the configuration has wild card virtual servers, the NetScaler appliance might become unresponsive when writing the appflow records.

  [From Build 50.10] [# 658624, 660103]

- Memory usage on a NetScaler appliance might increase over time if the AppFlow feature is enabled and HTTP pipelined requests are sent to the HTTP or SSL virtual servers that match at least one of the AppFlow policies.

  [From Build 51.26] [# 672102, 670990, 671906]

- Memory usage on a NetScaler appliance might increase over time if the AppFlow feature is enabled and HTTP pipelined requests are sent to the HTTP or SSL virtual servers that match at least one of the AppFlow policies.

  [From Build 52.13] [# 672102, 670990, 671906]

- If you enable AppFlow on a SQL virtual server, the NetScaler appliance might become unresponsive.

  [From Build 52.13] [# 671462, 672362]

- If AppFlow Client-Side Measurements action is enabled on any of the AppFlow policies, connection to backend server might get terminated intermittently thereby sending only partial response to the client.

  [From Build 52.13] [# 672863, 673532]

- Memory usage on a NetScaler appliance might increase over time if AppFlow Client-side Measurements is enabled.

  [From Build 52.13] [# 666358, 672859]

- If numerous GET requests are sent to a NetScaler appliance on which AppFlow is enabled, at some point the requests begin to time out.

  [From Build 52.13] [# 671993, 674438]

**Application Firewall**

- The NetScaler appliance fails if the signature match function accesses invalid memory while matching signature rules.

  [From Build 48.10] [# 643854]

- The exported, learned data for field formats does not match the output of the following command: sh appfw learning data.

  [From Build 48.10] [# 329025, 303481]

- Sites that use the NetScaler application firewall have excessive high availability failovers because of a faulty error-handling routine related to memory allocation.

  [From Build 48.10] [# 647309]

- Applications might not load properly when the memory_max_allowed value for the AppFW pool is low. This low memory condition can also cause memory allocation errors that result in numerous connection resets.

  [From Build 48.10] [# 649031, 651536]

- If the HTML response page contains a pair of hyphens (--) in the comment tag, the NetScaler appliance might parse the response page incorrectly and not add the URLs to starturl closure. This could result in some starturl violations.

  [From Build 48.10] [# 648104]

- The name of a user defined signature objects must not contain a hash character (#), even though the feedback message inaccurately lists it as an allowed character.

  [From Build 48.10] [# 648010]

- If the NetScaler appliance sends AppFlow data with application firewall records to the Security Insight collector, the appliance might fail. This might occur if the built-in NOPOLICY policy, which does not have any specified action, is configured as a global policy.

  [From Build 49.16] [# 656771]

- The NetScaler appliance might fail if both of the following conditions are met:

  - The application firewall and compression modules are both active for a connection.

  - The connection is aborted for any reason, such as connection failure on the client or server, or invalid HTTP content is received from the client or server.

  Typically, the application firewall and compression modules free the resources, including references to the connection. However, in rare cases, freeing a connection results in a dangling connection structure pointer or duplicate freeing of the structure pointer. In either of these cases, the appliance might fail.

  [From Build 49.16] [# 648981, 648996, 653492, 654739]

- A NetScaler AppFirewall appliance might run out of memory, because firewall sessions might not get cleaned up in a high availability environment if sync or propagation is disabled or the software versions running on a pair of nodes do not match. This is due to DHT not being able to clean up entries properly.

  [From Build 49.16] [# 646293, 645547, 658502]

- If the NetScaler appliance sends AppFlow data with application firewall records to the Security Insight collector, the appliance might fail. This might occur if the built-in NOPOLICY policy, which does not have any specified action, is configured as a global policy.

  [From Build 50.10] [# 656771]

- A NetScaler appliance in a high availability configuration might fail when the Application Firewall HTTP request is chunked or the chunk-header information is split across the packet and the content-type is "application/x-www-form-urlencoded" and "multipart/form-data".

  [From Build 50.10] [# 642238, 646749, 650320]

- A NetScaler appliance fails under the following set of conditions:

  - The appliance is configured to log for parsing errors in XML responses, and the configuration includes a confidential field. Webform fields can be designated as confidential fields to protect the information that users type into them.

  - The appliance receives a request in which query parameters are set.

  - A parsing error occurs during processing of the XML response.

  [From Build 50.10] [# 658561, 639647]

- On a NetScaler Application Firewall appliance in a high availability configuration, learning mode does not work after an upgrade to release 11.0 build 68.10.

  [From Build 51.21] [# 662734]

- Application Firewall uses master-slave communication for processing security checks and retrieves connection information through Protocol Control Block (PCB). In a high availability mode, the NetScaler appliance might fail, if factory reset occurs when PCB variables are cleared before freeing Application Firewall context data when accessing null pointer during processing.

  [From Build 51.21] [# 664159, 665334]

- A log message is not generated when the FormFieldConsistency protection is enabled on an Application Firewall profile and the generated hidden field "as_fid" is modified.

  With this fix, the NetScaler Application Firewall now generates a log message when the "FormFieldConsistency" protection is enabled and the hidden field "as_fid" is modified in the NetScaler Application Firewall profile.

  [From Build 51.21] [# 664211]

- On a NetScaler Application Firewall appliance in a high availability configuration, learning mode does not work after an upgrade to release 11.0 build 68.10.

  [From Build 51.21] [# 662359, 670726]

- A NetScaler appliance might fail when Application Firewall processes a request for SQL injection inspection, if the request has the SQLInjectiontype field set to "SQL Special Char or Keyword" and SQL comment handling is set to "ANSI/Nested".

  [From Build 51.21] [# 665631, 669524]

- Executing force sync operation using the nssync -s command from the shell triggers NetScaler appliance reboot and crash. The nsnetsvc crash occurs when the import filename length exceeds MAX_FILE_PATH_LEN.

  [From Build 51.21] [# 657920]

- In a high availability setup, after successful deployment of the Application Firewall learned StartURL rule from the GUI, the rule remains in the learned database and is not removed. Deploying the same startURL rule results in the following error message: "The StartURL check is already in use."

  [From Build 51.21] [# 661111]

- NetScaler release 11.0 build 47 or later logs error messages when you enable the Application Firewall feature on a NetScaler appliance in high availability mode.

  [From Build 51.21] [# 660528]

- If the NetScaler Application Firewall learning feature is enabled, Form Field Consistency violations result in blocking URL requests that end with a question mark (?), with no query parameters.

  [From Build 51.21] [# 666019]

- The Onhover pattern has been added to the default list of cross-site scripting (XSS) denied patterns that the Application Firewall looks for when scanning traffic.

  [From Build 51.21] [# 665595]

- CPU utilization becomes high if you upgrade the NetScaler appliance to release 11.0 build 65 and enable Application Firewall Starturl Closure protection.

  [From Build 51.21] [# 656708, 656061, 658404, 670134]

- CPU utilization becomes high if you upgrade the NetScaler appliance to release 11.0 build 65 and enable Application Firewall Starturl Closure protection.

  [From Build 51.26] [# 656708, 656061, 658404, 670134]

- If the NetScaler Application Firewall learning feature is enabled, Form Field Consistency violations result in blocking URL requests that end with a question mark (?), with no query parameters.

  [From Build 51.26] [# 666019]

- The Onhover pattern has been added to the default list of cross-site scripting (XSS) denied patterns that the Application Firewall looks for when scanning traffic.

  [From Build 51.26] [# 665595]

- A log message is not generated when the FormFieldConsistency protection is enabled on an Application Firewall profile and the generated hidden field "as_fid" is modified.

  With this fix, the NetScaler Application Firewall now generates a log message when the "FormFieldConsistency" protection is enabled and the hidden field "as_fid" is modified in the NetScaler Application Firewall profile.

  [From Build 51.26] [# 664211]

- A NetScaler appliance might fail when Application Firewall processes a request for SQL injection inspection, if the request has the SQLInjectiontype field set to "SQL Special Char or Keyword" and SQL comment handling is set to "ANSI/Nested".

  [From Build 51.26] [# 665631, 669524]

- Executing force sync operation using the nssync -s command from the shell triggers NetScaler appliance reboot and crash. The nsnetsvc crash occurs when the import filename length exceeds MAX_FILE_PATH_LEN.

  [From Build 51.26] [# 657920]

- Application Firewall uses master-slave communication for processing security checks and retrieves connection information through Protocol Control Block (PCB). In a high availability mode, the NetScaler appliance might fail, if factory reset occurs when PCB variables are cleared before freeing Application Firewall context data when accessing null pointer during processing.

  [From Build 51.26] [# 664159, 665334]

- On a NetScaler Application Firewall appliance in a high availability configuration, learning mode does not work after an upgrade to release 11.0 build 68.10.

  [From Build 51.26] [# 662734]

- On a NetScaler Application Firewall appliance in a high availability configuration, learning mode does not work after an upgrade to release 11.0 build 68.10.

  [From Build 51.26] [# 662359, 670726]

- NetScaler release 11.0 build 47 or later logs error messages when you enable the Application Firewall feature on a NetScaler appliance in high availability mode.

  [From Build 51.26] [# 660528]

- In a high availability setup, after successful deployment of the Application Firewall learned StartURL rule from the GUI, the rule remains in the learned database and is not removed. Deploying the same startURL rule results in the following error message: "The StartURL check is already in use."

  [From Build 51.26] [# 661111]

- During the downgrade process, the NetScaler appliance becomes unresponsive and generates an aslearn core file if the application firewall schema profile of the learned database files is not installed properly.

  [From Build 52.13] [# 670752]

- The IP reputation feature does not get enabled when the application firewall add-on license is added.

  [From Build 52.13] [# 675202]

- If you upgrade NetScaler appliance in a high availability (HA) setup from version 10.5.56.15 to version 11.1.51.1901 and skip 250 rules with active traffic, the GUI or CLI displays a "failed to skip some rules" error message and an operation time-out error message.

  [From Build 52.13] [# 661111]

- The NetScaler appliance crashes during a field-consistency check if processing a large number of form-select fields.

  [From Build 52.13] [# 668627, 664482]

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

  [From Build 52.13] [# 674031]

- NetScaler release 11.0 build 47 or later logs error messages when you enable the Application Firewall feature on a NetScaler appliance in high availability mode.

  [From Build 52.13] [# 660528]

- On a NetScaler Application Firewall appliance in a high availability configuration, learning mode does not work after an upgrade to release 11.0 build 68.10.

  [From Build 52.13] [# 662359, 670726]

- Application firewall signature rule #14990 has a PCRE expression pattern to detect the presence of a violation string in the Accept-Charset header. This expression is computationally intensive and results in generation of log message "PCRE match limit exceeded with regex..." With this fix, rule #14990 is deprecated and replaced by signature rule #999972, which has an optimized PCRE expression. The new rule shows the source as Snort and the Snort ID as 14990.

  [From Build 52.13] [# 669824]

- On a NetScaler Application Firewall appliance in a high availability configuration, learning mode does not work after an upgrade to release 11.0 build 68.10.

  [From Build 52.13] [# 662734]

- An archive error can occur when application firewall profiles are exported and archived, because the export file is not removed from the /var/archive/appfw/ and /var/tmp directories after profile export is successful. The problem is caused by an uppercase profile name when the archived export file is saved with the same case as the profile name.

  [From Build 52.13] [# 670744]

**Cache**

- A NetScaler VPX instance becomes unresponsive if a range request is greater than the cached response size. This issue happens if you enable the media classification mode on a NetScaler appliance. While parsing range header and creating range records table, the value for parameter object size is set incorrectly. So when a range request is received, the incorrect value of the stored response causes failure.

  [From Build 51.21] [# 657823, 659374, 661940, 662460, 667599]

- A NetScaler VPX instance becomes unresponsive if a range request is greater than the cached response size. This issue happens if you enable the media classification mode on a NetScaler appliance. While parsing range header and creating range records table, the value for parameter object size is set incorrectly. So when a range request is received, the incorrect value of the stored response causes failure.

  [From Build 51.26] [# 657823, 659374, 661940, 662460, 667599]

**Clustering**

- If a load balancing server is trying to synchronize its states, occasionally one or more cluster nodes might get stuck in a Service state. As a result, the other nodes in the cluster might be unavailable, which leads to an improper cluster formation.

  [From Build 50.10] [# 651828]

**Content Switching**

- The NetScaler appliance might fail if you change the target of a content switching policy action from virtual server based to expression based.

  [From Build 50.10] [# 657325, 653722, 659696, 661214]

**DNS**

- A clear config operation in a Cluster deployment does not set non-CCO nodes to the default value for the "max pipeline" parameter.

[From Build 48.10] [# 648087]

- A NetScaler appliance configured as an DNS end resolver sometimes fails to respond to DNS queries. When the appliance is configured as an end resolver, it generates iterative DNS queries to name servers on behalf of the client and returns the final responses. If a DNS zone has multiple NS records, the appliance queries the first name server in the NS record. If this resolution fails, the appliance does not retry with other name servers in the NS records, and it does not send any response to the client.

  [From Build 49.16] [# 645836]

- If the DNS server from which the cached DNS records are being served goes DOWN, the proactive DNS update queries are redirected to the back-end server.

  [From Build 52.13] [# 660562]

- When a NetScaler appliance on which DNSSEC is configured is an authoritative DNS server for two domain zones, the appliance might send the same RRSIG responses to both zones instead of responding to only the appropriate zone.

  [From Build 52.13] [# 671880]

- When the NetScaler appliance receives a DNS TCP packet that has dnspayloadlen as zero, the appliance might dump core memory.

  [From Build 52.13] [# 666803]

- If the load balancing feature is disabled and DNS name servers are being used, DNS resolution uses the most recently configured name server. If that name server is disabled, one of the name servers that is UP is used for DNS resolution.

  [From Build 52.13] [# 670588]

**DataStream**

- The DataStream feature does not work if you use a MySQL database at the back end.

  [From Build 51.21] [# 629504]

- The DataStream feature does not work if you use a MySQL database at the back end.

  [From Build 51.26] [# 629504]

**GSLB**

- In a GSLB setup, if you have configured static proximity as the primary load balancing method and RTT as the backup load balancing method, the NetScaler appliance might intermittently send an empty response to a DNS query requesting the GSLB domain.

  [From Build 50.10] [# 616321]

- The MEP connection for site metrics goes DOWN if the dynamic RTT and GSLB server persistence features are unused for more than 249 days. In some cases, however, the MEP connection for site metrics remains UP, but the MEP connection for network metrics goes DOWN.

  [From Build 51.21] [# 658890]

- In a GSLB high availability setup, if a node stays in secondary state for more than 249 days, the service state might not be updated on this node after it becomes the primary node.

  [From Build 51.21] [# 658093]

- In a GSLB high availability setup, if a node stays in secondary state for more than 249 days, the service state might not be updated on this node after it becomes the primary node.

  [From Build 51.26] [# 658093]

- The MEP connection for site metrics goes DOWN if the dynamic RTT and GSLB server persistence features are unused for more than 249 days. In some cases, however, the MEP connection for site metrics remains UP, but the MEP connection for network metrics goes DOWN.

  [From Build 51.26] [# 658890]

**Integrated Caching**

- A NetScaler appliance fails if a Page Tracking session is enabled on the appliance by Appflow or AppQoE modules for partial content responses. This happens only for partial content responses served from Integrated Cache.

  [From Build 51.21] [# 656556]

- A NetScaler appliance fails if a Page Tracking session is enabled on the appliance by Appflow or AppQoE modules for partial content responses. This happens only for partial content responses served from Integrated Cache.

  [From Build 51.26] [# 656556]

- A NetScaler appliance fails if a Page Tracking session is enabled on the appliance by Appflow or AppQoE modules for partial content responses. This happens only for partial content responses served from Integrated Cache.

  [From Build 52.13] [# 656556]

**Load Balancing**

- A secure HTTP-ECV monitor might time out if the back-end server sends a large certificate.

  [From Build 48.10] [# 638148]

- The NetScaler appliance fails to send an assertion back to the service provider when the SAML request comes without an ID field. When behaving as a samlidp, the ID field from the authnReq is remembered, so it can be sent back in the

assertion. If service providers don't send IDs, we fail due to logic error. The logic was revised so if we don not get an ID, we don't send it back.

[From Build 48.10] [# 648489]

- In the SAML response, the RelayState field is truncated. When the samlidp feature is processed, the URL decodes the entire content before parsing for individual elements. The customer's service provider sends the RelayState that was encoded. When the service provider posts the assertion back, the RelayState is truncated resulting in an SP failure.

  [From Build 48.10] [# 648337]

- In a high availability (HA) setup, after a forced HA synchronization, the configuration is first cleared and then reapplied on the secondary node. As part of the synchronization operation, the service state changes are logged in the ns.log file. Repeated forced synchronizations can flood the ns.log file. However, the service state messages are applicable only to the primary node and not relevant to the secondary node. Therefore, these messages are not logged in the ns.log file on the secondary node.

  [From Build 50.10] [# 645197]

- If the same IP address is assigned to both a GSLB service and a load balancing virtual server, the NetScaler appliance dumps core and restarts, because the internal service weight is set to zero.

  [From Build 50.10] [# 628937]

- If a GSLB service goes DOWN and then returns to the UP state, the configured hash-based load balancing methods might produce incorrect load balancing decisions, because the cache maintained for hash-based load balancing algorithms is not cleared when the GSLB service state is updated through MEP.

  [From Build 51.21] [# 658463, 658940]

- If a GSLB service goes DOWN and then returns to the UP state, the configured hash-based load balancing methods might produce incorrect load balancing decisions, because the cache maintained for hash-based load balancing algorithms is not cleared when the GSLB service state is updated through MEP.

  [From Build 51.26] [# 658463, 658940]

- Redirection does not work properly if you initially configure an HTTPS redirect URL without a slash at the end, then change the URL by adding a slash, and then remove the slash.

  [From Build 52.13] [# 662640]

- The NetScaler appliance might become unresponsive because of an internal issue related to CRC check if custom monitors are configured for a load balancing configuration of type TFTP.

  [From Build 52.13] [# 658860]

- In a cluster setup, the configuration of a service might be lost if you restart the appliance after you have configured a request timeout action (-reqTimeoutAction) in an HTTP profile and attached the profile to the service.

  [From Build 52.13] [# 649994, 649940]

- The NetScaler appliance dumps core and restarts if an autoscale service group is configured with SSL as the service type.

  [From Build 52.13] [# 656734]

**Log Streaming**

- If you enable AppFlow feature and ULFD mode on a NetScaler appliance, memory usage on the NetScaler appliance might increase.

  [From Build 52.13] [# 663260]

**NetScaler GUI**

- In Security > AAA > Virtual Servers, you can now bind an SSL profile to a virtual server.

  [From Build 48.10] [# 651031]

- If you have configured static proximity as the load balancing method on a load balancing virtual server, you cannot set a backup method by using the GUI.

  [From Build 48.10] [# 648408]

- When creating a cluster node group, you no longer have to specify a node state. The "Add Node Group" page in the NetScaler GUI displays "state" as optional, not as a required field.

  Page Navigation: Configuration > System >Cluster > NodeGroup > Add Node Group

  [From Build 48.10] [# 650357]

- The field value for X-Forwarded-For HTTP header is not displayed as client IP in NetScaler Security Insight violation logs.

  [From Build 49.16] [# 645284, 636390]

- SSL GSLB services are configured on port 443. However, if you try to edit the service by using the NetScaler GUI, port 80 appears instead of 443. This was a display issue and is fixed.

  [From Build 49.16] [# 654239]

- If you try to bind a default load balancing virtual server to a content switching virtual server in an admin partition, the following error message appears:

Operation not permitted.

[From Build 50.10] [# 653058]

- You cannot unbind a transform policy from a virtual server by using the GUI.

  [From Build 51.21] [# 652579]

- If the features "Force password change for nsroot user when default nsroot password is being used" and "strong password" are enabled, any password is accepted when you change the nsroot password.

  [From Build 51.21] [# 656825]

- When a partition admin tries to perform the Download, Create, or Create Directory operation on the "Manage Certificate" screen, an "operation not permitted" error appears. The expected behavior is that the buttons must be disabled.

  [From Build 51.21] [# 491353]

- If the name of a load balancing virtual server contains a space, the virtual server is not listed by the reporting tool. (Reporting > Counters > System entities statistics > Entities)

  [From Build 51.21] [# 642269]

- If you use "clear ns configuration" command to clear the NetScaler configuration and reset it to factory default, the command policies are restored to the default values.

  [From Build 51.21] [# 643546, 200969]

- If the features "Force password change for nsroot user when default nsroot password is being used" and "strong password" are enabled, any password is accepted when you change the nsroot password.

  [From Build 51.26] [# 656825]

- If you use "clear ns configuration" command to clear the NetScaler configuration and reset it to factory default, the command policies are restored to the default values.

  [From Build 51.26] [# 643546, 200969]

- If the name of a load balancing virtual server contains a space, the virtual server is not listed by the reporting tool. (Reporting > Counters > System entities statistics > Entities)

  [From Build 51.26] [# 642269]

- You cannot unbind a transform policy from a virtual server by using the GUI.

  [From Build 51.26] [# 652579]

- When a partition admin tries to perform the Download, Create, or Create Directory operation on the "Manage Certificate" screen, an "operation not permitted" error appears. The expected behavior is that the buttons must be disabled.

  [From Build 51.26] [# 491353]

- You cannot bind a cipher or cipher group to an SSL entity by using the NetScaler GUI.

  [From Build 52.13] [# 648293, 638254]

- You cannot unbind a transform policy from a virtual server by using the GUI.

  [From Build 52.13] [# 652579]

**NetScaler ICA**

- The NetScaler appliance fails because it attempts to process a large unexpected value for an Expander variable. This fix adds checks to prevent this condition.

  [From Build 51.21] [# 660894, 662489, 668651]

- When Session Reliability is disabled on Storefront and Session Reliability on HA Failover is enabled on a NetScaler high availability pair running on 11.1 build 49.16, the passive NetScaler instance might reboot when you launch an application or a virtual desktop.

  You can avoid the issue by performing the following steps:

  * Disable "Session Reliability on HA" feature on the NetScaler instance.

  [From Build 51.21] [# 664372, 667057]

- The NetScaler appliance fails because it attempts to process a large unexpected value for an Expander variable. This fix adds checks to prevent this condition.

  [From Build 51.26] [# 660894, 662489, 668651]

- When Session Reliability is disabled on Storefront and Session Reliability on HA Failover is enabled on a NetScaler high availability pair running on 11.1 build 49.16, the passive NetScaler instance might reboot when you launch an application or a virtual desktop.

  You can avoid the issue by performing the following step:

  * Disable "Session Reliability on HA" feature on the NetScaler instance.

  [From Build 51.26] [# 664372, 667057]

**NetScaler Insight Center**

- AppFlow configuration fails if you use the NetScaler Insight Center FQDN instead of the NetScaler Insight Center IP address.

  [From Build 48.10] [# 652425]

- System groups cannot be created in the NetScaler Insight Center GUI.

  [From Build 48.10] [# 650657]

- For a NetScaler appliance in multicore setup, reports from all cores were not getting generated except "0" core.

  [From Build 50.10] [# 656225]

- When you use LDAP for external authentication, you will receive a "Error: Resource does not exist" error message when you click Configuration tab.

  [From Build 50.10] [# 658344]

- The whitelist of Citrix Receiver versions used by HDX Insight now includes version 13.0.2.265571 of Citrix Receiver for Linux.

  [From Build 51.21] [# 614558, 606817]

- Ability to Export Gateway Insight Reports in Other Formats

  You can now export the Gateway Insight reports with all the details shown in the GUI in excel and csv format from NetScaler Insight Center.

  [From Build 51.26] [# 631822]

- The whitelist of Citrix Receiver versions used by HDX Insight now includes version 13.0.2.265571 of Citrix Receiver for Linux.

  [From Build 51.26] [# 614558, 606817]

**NetScaler SAMLIdP**

- If the RelayState value in a SAML Authentication request is more than 512 bytes but less than 1024 bytes, the SAML IdP server causes buffer overrun when sending an assertion after successful authentication.

  [From Build 50.10] [# 656779, 664051, 664765]

**NetScaler VPX Appliance**

- In a KVM environment, a NetScaler VPX instance fails to start if you have configured more than 11 vCPUs.

  [From Build 49.16] [# 647348]

- If you deploy NetScaler VPX on Azure in HA mode, the VPN virtual servers on the secondary node are not reachable after a failover. This is because, during a synchronization operation, the NSIP address of the primary node is used to create the virtual server on the secondary node. After a failover, when the secondary node becomes the new primary, the VPN virtual server has the NSIP address of the old primary.

  [From Build 49.16] [# 651670]

- A NetScaler VPX instance might stop responding and dump core memory if you allocate a large disk size for log messages. The higher the rate of log messages, the more quickly the instance runs out of memory and fails.

  [From Build 50.10] [# 646674]

- A NetScaler VPX instance might stop responding and dump core memory if you allocate a large disk size for log messages. The higher the rate of log messages, the more quickly the instance runs out of memory and fails.

  [From Build 51.21] [# 646674]

- If you add additional SR-IOV or PCI passthrough interfaces to an existing NetScaler virtual appliance configured with SR-IOV or PCI passthrough interfaces, the existing interface names might get corrupted.

  [From Build 51.21] [# 659827, 662429]

- In an ESX environment, a CLAG channel that includes a VMXNET3 interface might continue to send LACPDUs to its partner even when it is in DETACHED state.

  [From Build 51.21] [# 642389]

- In an ESX environment, a CLAG channel that includes a VMXNET3 interface might continue to send LACPDUs to its partner even when it is in DETACHED state.

  [From Build 51.26] [# 642389]

- A NetScaler VPX instance might stop responding and dump core memory if you allocate a large disk size for log messages. The higher the rate of log messages, the more quickly the instance runs out of memory and fails.

  [From Build 51.26] [# 646674]

- If you add additional SR-IOV or PCI passthrough interfaces to an existing NetScaler virtual appliance configured with SR-IOV or PCI passthrough interfaces, the existing interface names might get corrupted.

  [From Build 51.26] [# 659827, 662429]

- If you add additional SR-IOV or PCI passthrough interfaces to an existing NetScaler virtual appliance configured with SR-IOV or PCI passthrough interfaces, the existing interface names might get corrupted.

  [From Build 52.13] [# 659827, 662429]

- A NetScaler VPX appliance running on a VMware ESX server and configured with a VMXNET3 network interface stops responding and restarts if any traffic is sent to a tagged interface. Also, in the log message, the VLAN ID of the tagged interface is incorrect.

  [From Build 52.13] [# 671581, 676316]

- The NetScaler VPX GUI incorrectly shows Moscow's time zone as GMT+4 instead of GMT+3.

  [From Build 52.13] [# 662630]

- A NetScaler VPX instance might stop responding and dump core memory if you allocate a large disk size for log messages. The higher the rate of log messages, the more quickly the instance runs out of memory and fails.

  [From Build 52.13] [# 646674]

- When a remote tagged IP address is accessed through a NetScaler VPX appliance hosted on Linux KVM, the checksum value in each sent packet is incorrect.

  [From Build 52.13] [# 655067, 668302]

**Networking**

- A NetScaler appliance with OSPFv3 dynamic routing protocol configured might measure the length of OSPFv3 LSA packets in Network Byte Order instead of Host Byte Order for comparison with the minimum required packet length. As a result, the NetScaler appliance becomes unresponsive.

  [From Build 48.10] [# 652131]

- During a "force sync" operation in a cluster deployment, performing a "save config" operation on a node might lead to a full or partial configuration loss on that node. With this fix, the "save config" operation is not permitted during a "force sync" operation.

  [From Build 49.16] [# 642375, 658619]

- Restarting a NetScaler appliance that has a VLAN bound to a traffic domain and is configured as a SYNC VLAN or NSVLAN might cause configuration loss of binding between the VLAN and the traffic domain.

  [From Build 51.21] [# 648839]

- In a high availability (HA) setup, after an HA force failover operation, the NetScaler appliance removes (but not properly) static default route6s of all non-default traffic domains from its memory.

  Though the "show route6 operation" does not display these route6s but adding them again fails with the following error message: "ERROR: Resource already exist". This is because these route6s were not completely removed from memory.

  This issue also happens on a standalone NetScaler appliance when a traffic domain that has default route6s is removed.

[From Build 51.21] [# 644265]

- In a high availability setup, after a failover, the new primary node does not set the R bit and F bit in BGP open messages that are used to inform the upstream router that the node has restarted gracefully.

  [From Build 51.21] [# 665774]

- For extended ACL rules that are associated with NAT configurations (for example, RNAT rules and Large Scale NAT configurations), the NetScaler GUI displays the TCP established parameter as enabled even though the parameter is disabled.

  [From Build 51.21] [# 597458]

- In a high availability (HA) setup, after an HA force failover operation, the NetScaler appliance removes (but not properly) static default route6s of all non-default traffic domains from its memory.

  Though the "show route6 operation" does not display these route6s but adding them again fails with the following error message: "ERROR: Resource already exist". This is because these route6s were not completely removed from memory.

  This issue also happens on a standalone NetScaler appliance when a traffic domain that has default route6s is removed.

  [From Build 51.26] [# 644265]

- Restarting a NetScaler appliance that has a VLAN bound to a traffic domain and is configured as a SYNC VLAN or NSVLAN might cause configuration loss of binding between the VLAN and the traffic domain.

  [From Build 51.26] [# 648839]

- In a high availability setup, after a failover, the new primary node does not set the R bit and F bit in BGP open messages that are used to inform the upstream router that the node has restarted gracefully.

  [From Build 51.26] [# 665774]

- For extended ACL rules that are associated with NAT configurations (for example, RNAT rules and Large Scale NAT configurations), the NetScaler GUI displays the TCP established parameter as enabled even though the parameter is disabled.

  [From Build 51.26] [# 597458]

- On a NetScaler appliance, when a routing daemon (for example, BGP routing daemon) is restarted multiple times over a short period of time, the corresponding routing configuration (for example, BGP routing configuration) might get removed from the appliance.

  [From Build 52.13] [# 669005]

- In a high availability setup, after a failover, the new primary node does not set the R bit and F bit in BGP open messages that are used to inform the upstream router that the node has restarted gracefully.

  [From Build 52.13] [# 665774]

**SSL**

- You can bind ECDSA ciphers to an SSL virtual server on a platform that does not have N3 chips even though ECDSA ciphers are supported only on platforms with N3 chips.

  [From Build 48.10] [# 635234]

- Adding a certificate revocation list (CRL) on the NetScaler appliance fails with the error message "Certificate Issuer Mismatch" for a DER certificate, and with the error message "Invalid CRL" for a PEM certificate. This issue occurs because the attribute type of the common name field is different for the CA certificate than for the CRL.

  [From Build 48.10] [# 623058, 634017]

- A certificate-key pair bound to a secure monitor is not saved in the configuration file (ns.conf). As a result, the binding is lost after you restart the appliance.

  [From Build 49.16] [# 654722]

- A NetScaler virtual appliance sometimes fails because of a memory leak if you use GCM-based ciphers on a VPX appliance. The ciphers can eventually exhaust memory, causing the appliance to fail if the memory exhaustion error is not gracefully handled.

  [From Build 49.16] [# 652477, 654559, 656035, 657343]

- Client authentication causes memory leak if a client sends a certificate that includes its intermediate CA certificates. This exhausts memory on the NetScaler appliance.

  [From Build 49.16] [# 656671]

- In a cluster setup, if you rename a load balancing virtual server of type SSL, the local database table that is used for all GET operations is not updated.

  [From Build 50.10] [# 620964, 576828, 641041]

- TLS handshake fails if client authentication is set to mandatory.

  [From Build 50.10] [# 656490]

- A NetScaler appliance might dump core and restart repeatedly if the SSL3-EDH-RSA-DES-CBC3-SHA cipher is selected when heavy traffic has exhausted the appliance's memory.

  [From Build 51.21] [# 661818]

- If you upgrade to release 10.5, SSL client authentication fails if it uses a 4096-bit client certificate.

  [From Build 51.21] [# 600815, 343395]

- If you try to load large certificate files (> 256kB), the NetScaler appliance might dump core and restart, because of insufficient memory.

  [From Build 51.21] [# 643614, 624364, 646510, 667980]

- If a profile is bound to an SSL virtual server, the NITRO API displays incorrect SSL virtual server settings. The correct settings are displayed in the profile.

  [From Build 51.21] [# 628135]

- SSL processing is delayed if the server sends a DES cipher with TLS1.2 protocol in the server_hello message to the NetScaler appliance. Although this combination is deprecated, the appliance tries to process it. The operation fails at the SSL card and blocks the card for a few seconds, causing latency in processing any new requests on the same card.

  [From Build 51.21] [# 661628]

- The NetScaler appliance dumps core memory and restarts if all of the following conditions are met:

  - SNI feature is enabled.

  - Exact server certificate match is unsuccessful.

  - The common name field is greater than 253 characters.

  [From Build 51.21] [# 664338, 670653]

- A NetScaler appliance might dump core and restart repeatedly if the SSL3-EDH-RSA-DES-CBC3-SHA cipher is selected when heavy traffic has exhausted the appliance's memory.

  [From Build 51.26] [# 661818]

- If you upgrade to release 11.1, SSL client authentication fails if a 4096-bit client certificate is used.

  [From Build 51.26] [# 600815, 343395]

- The "set ssl vserver" or the "unset ssl vserver" command fails with the following error:

  Internal error

  [From Build 51.26] [# 670927, 673889]

- If you try to load large certificate files (> 256kB), the NetScaler appliance might dump core and restart, because of insufficient memory.

  [From Build 51.26] [# 643614, 624364, 646510, 667980]

- The NetScaler appliance might dump core and restart if it receives SSL traffic while AppFlow is enabled.

  [From Build 51.26] [# 668689, 672376, 672526, 673897, 674128]

- The SSL handshake fails if a server certificate is linked to its issuer certificate, OCSP stapling is enabled, and an SSL client requests the server-certificate status.

  [From Build 51.26] [# 671777]

- SSL processing is delayed if the server sends a DES cipher with TLS1.2 protocol in the server_hello message to the NetScaler appliance. Although this combination is deprecated, the appliance tries to process it. The operation fails at the SSL card and blocks the card for a few seconds, causing latency in processing any new requests on the same card.

  [From Build 51.26] [# 661628]

- If a profile is bound to an SSL virtual server, the NITRO API displays incorrect SSL virtual server settings. The correct settings are displayed in the profile.

  [From Build 51.26] [# 628135]

- The NetScaler appliance might dump core and restart if it receives SSL traffic while AppFlow is enabled.

  [From Build 52.13] [# 668689, 672376, 672526]

- The NetScaler appliance might dump core memory and restart if AppFlow and SSL features are enabled and the appliance receives SSL traffic.

  [From Build 52.13] [# 673897, 674543, 674479, 674128, 676165]

- In a cluster setup, you cannot make any change to a service or service group if you have associated a common name with the service or the service group and enabled or disabled server name indication (SNI).

  [From Build 52.13] [# 665340]

- You cannot enable server-name indication (SNI) in a back-end profile.

  [From Build 52.13] [# 670267]

- The "stat ssl vserver" command for a content switching, cache redirection, or VPN virtual server fails, and the following error message appears:

  No such resource [vServerName, <vservername>]

  [From Build 52.13] [# 644731, 671337]

- A DTLS configuration fails if it uses MAC-Based forwarding or a VLAN.

  [From Build 52.13] [# 615454, 629512]

- An SSL handshake fails if a client hello includes an ECC extension but the NetScaler appliance does not support any of the ECDHE ciphers in the cipher list sent by the client. The handshake fails even if the list contains some non-ECDHE ciphers that are supported.

  [From Build 52.13] [# 668239]

- The SSL handshake fails if a server certificate is linked to its issuer certificate, OCSP stapling is enabled, and an SSL client requests the server-certificate status.

  [From Build 52.13] [# 671777]

- If memory allocation fails during a TLS1.2 protocol handshake, the handshake is not terminated. As a result, the appliance might dump memory core and restart.

  [From Build 52.13] [# 630547, 639222, 639465, 646023, 647371, 649201, 658037, 662933, 663160, 665797, 668460, 676013, 679036]

- The "set ssl vserver" or the "unset ssl vserver" command fails and the following error message appears:

  Internal error

  [From Build 52.13] [# 670927, 673889, 673829, 671270]

- The SSL parameter "deny SSL renegotiation" is now set to ALL by default in all admin partitions. Previously, it was set to NO in the non-default partitions.

  [From Build 52.13] [# 663601]

- The NetScaler appliance might dump core memory and restart if you bind a secure monitor to a domain based service.

  [From Build 52.13] [# 661808, 662002, 672103, 672532, 674664, 671558, 674758]

**System**

- The CPU parameter value on the LCD panel does not match the value reported by the NetScaler CLI or GUI.

  [From Build 48.10] [# 643237]

- Heavy traffic through a NetScaler appliance can result in a web log buffer overrun, causing a NetScaler Web logging (NSWL) client to reconnect. When the client reconnects, the use of surplus connections results in omission of the PCB's user-name information (part of connection related information) during cloning. This leads to a loss of log data.

  [From Build 48.10] [# 633308, 646753, 648657]

- The Configd daemon fails if the number of session IDs exceeds the preset limit and existing client sessions are renumbered.

[From Build 49.16] [# 639380, 657168, 657781]

- If NetScaler appliance is setup with Web Log feature and weblog clients are connected then under traffic stress, a buffer overrun can cause the weblog client to reconnect. When the clients reconnect, we lose part of the data on connections where reconnect was triggered and hence log data is not complete.

  [From Build 49.16] [# 633308, 646753, 648657, 656502]

- Memory allocation failures occur, because the NetScaler appliance does not allocate sufficient memory for packet engines.

  [From Build 49.16] [# 647072, 643407, 650630]

- On a NetScaler appliance, if a FIN packet is held back by the forwarding interface and in the meantime, if Selective Acknowledgement (SACK) blocks are generated for the previous packet, the appliance fails.

  [From Build 49.16] [# 648446]

- NetScaler appliance crashes when a large host-name header is received and AppFlow logging for host-name and domain-name is enabled.

  [From Build 50.10] [# 660075, 664886]

- On a NetScaler appliance, if a FIN packet is held back by the forwarding interface and in the meantime, if Selective Acknowledgement (SACK) blocks are generated for the previous packet, the appliance fails.

  [From Build 50.10] [# 648446]

- When page tracking is enabled on an AppFlow, the NS_ESNS cookie is inserted into the response being served from the cache. The extra bytes added to the response are not accounted internally and so, when the ACK is received for those extra bytes, NetScaler crashes.

  [From Build 50.10] [# 649334, 653370, 656768, 662177]

- A NetScaler appliance fails if a TCP/IP session is simultaneously reused for TCP and Multipath TCP (MPTCP) operation and not mutually exclusive with TCP KeepAlive enabled for MPTCP subflows.

  [From Build 50.10] [# 654080]

- If a FASTCLOSE packet from a NetScaler appliance to a client is lost, the multipath TCP (MPTCP) session does not notify the application about the abrupt connection closure and close the socket. As a result, the appliance does not retransmit the lost packet.

  [From Build 50.10] [# 649968]

- The initial probe connection that a NetScaler appliance makes with the back-end internet server to check for server availability is now reusable for actual server connection with internet server.

[From Build 50.10] [# 654087]

- Syslog analysis is affected if the date/month format in a syslog message is not a user configured timestamp. This issue occurs if the Syslogaction uses the default date format (MM/DD/YYYY) instead of a user defined data format.

  [From Build 50.10] [# 659197, 656437]

- A NetScaler appliance might become unresponsive if it has a TCP profile with the TCP keepalive option enabled and is bound to a load balancing virtual server. The cause is an interoperability issue between the TCP keepalive and TCP packet retransmission functionalities.

  [From Build 50.10] [# 619349, 626027]

- In a NetScaler appliance, if there is an incoming TCP traffic from Wireless VLANs, the appliance routes the data packets to an IP router but now the appliance performs Policy Based Routing (PBR) to route the data packets based on incoming packet parameters, such as VLAN, MAC address, Interface, SRCIP, SRCPort, destination IP address, and destination port, to different routers through configured VLANs.

  [From Build 50.10] [# 649180]

- In a MPTCP connection, if a client negotiates a Maximum Segment Size (MSS) value of more than 1460 bytes, and the NetScaler appliance receives an ICMP protocol error message after fragmenting and sending a Data Security Standard (DSS) packet, the appliance fails. This happens because of incorrect handling of DSS packets with a segment sizes.

  [From Build 51.21] [# 648275]

- A NetScaler appliance might crash at a random location and dump a core file unrelated to the actual cause of the problem.

  [From Build 51.21] [# 660574]

- A NetScaler appliance constantly fails and dumps core memory, filling the Var directory with core files.

  [From Build 51.21] [# 647955]

- In an MPTCP connection, a NetScaler appliance sets the TCP PSH flag during retransmission of FastClose and DataFIN packets.

  [From Build 51.21] [# 667765]

- A NetScaler appliance constantly fails and dumps core memory, filling the Var directory with core files.

  [From Build 51.26] [# 647955]

- In an MPTCP connection, a NetScaler appliance sets the TCP PSH flag during retransmission of FastClose and DataFIN packets.

[From Build 51.26] [# 667765]

- A NetScaler appliance might crash at a random location and dump a core file unrelated to the actual cause of the problem.

  [From Build 51.26] [# 660574]

- In a MPTCP connection, if a client negotiates a Maximum Segment Size (MSS) value of more than 1460 bytes, and the NetScaler appliance receives an ICMP protocol error message after fragmenting and sending a Data Security Standard (DSS) packet, the appliance fails. This happens because of incorrect handling of DSS packets with a segment sizes.

  [From Build 51.26] [# 648275]

- A NetScaler appliance in the process of dumping core memory might malfunction and generate a misleading core file.

  [From Build 52.13] [# 660574]

- Processing audit-log leads to a memory-buffer overflow and disrupts other modules on the NetScaler appliance if the audit-log message size and log levels were not validated properly before audit-log processing began.

  [From Build 52.13] [# 670496]

- An overflow of integers updates the NetScaler memory statistics with a false value. This results in SNMP memory traps not reaching the configured threshold.

  [From Build 52.13] [# 663720, 612313]

- The NetScaler appliance might stop functioning and report a segmentation violation if your configuration includes policies or actions that use the following functions and one of them fails to obtain the memory that it needs:

  XPATH()

  XPATH_WITH_MARKUP()

  XPATH_JSON()

  XPATH_JSON_WITH_MARKUP()

  XPATH_HTML()

  XPATH_HTML_WITH_MARKUP()

  [From Build 52.13] [# 656646]

- In a TACACS authentication configuration, if you clear the system global TACACS policy, the NetScaler appliance displays a warning error message: "Config NodeGroup changed, force cluster sync should be fired on the newly added node to be in sync."

  [From Build 52.13] [# 666392]

- After an upgrade, a NetScaler Weblogging (NSWL) HTTP record size is miscalculated if the HTTP header size is greater than 16 kilobytes and it is not a multiple of the word boundary.

  [From Build 52.13] [# 671996, 672244]

- The NetScaler command line does not come out of the execution logic and does not display the command prompt when multiple grep with pipe operations are performed.

  [From Build 52.13] [# 667214]

- In CUBIC or BIC algorithms, the clamping congestion window is set maximum and sender congestion (snd_cwnd) is constantly increasing causing a NetScaler appliance unable to push data more than 2 GB.

  [From Build 52.13] [# 663551, 656192]

# Release History

For details of a specific release, see the corresponding release notes.

- Build 53.11 (2017-04-12) (Current build)
- Build 52.13 (2017-02-28)
- Build 51.21 (2017-02-02)
- Build 51.26 (2016-12-23) Replaces: 51.21
- Build 50.10 (2016-10-28)
- Build 49.16 (2016-09-28)
- Build 48.10 (2016-08-04)
- Build 47.14 (2016-06-30)
- Build 41.26 (2016-05-19)