



Citrix NetScaler 1000V Release Notes

Citrix NetScaler 11.0-69.12

First Published: 2017-03-07

Cisco Systems, Inc.

www.cisco.com

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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: This equipment has been tested and found to comply with the limits for a Class B digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation. This equipment generates, uses and can radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If the equipment causes interference to radio or television reception, which can be determined by turning the equipment off and on, users are encouraged to try to correct the interference by using one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/TV technician for help.

Modifications to this product not authorized by Cisco could void the FCC approval and negate your authority to operate the product.

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

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

CITRIX Citrix and other Citrix product names referenced herein are trademarks of Citrix Systems, Inc. and/or one of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. All other product names, company names, marks, logos, and symbols are trademarks of their respective owners.

© 2017 Cisco Systems, Inc. All rights reserved.

Contents

11.0-69.12.....	4
Fixed Issues.....	5
Known Issues.....	14
What's New in Previous NetScaler 11.0 Releases.....	25
Fixed Issues in Previous NetScaler 11.0 Releases	72
Release History.....	127

11.0-69.12

Updated: December 14, 2016 | Release notes version: 1.0

This release notes document describes the enhancements and changes and specifies the issues that exist, for the NetScaler release 11.0 Build 69.12. See [Release History](#).

Fixed Issues

The issues that are addressed in Build 69.12.

AAA-TM

- In rare cases, the NetScaler appliance dumps core and restarts if RADIUS accounting is enabled.

[# 648633, 659976]

- 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.

[# 664171]

- If SAML authentication is configured on NetScaler with artifact binding but certificates are not configured correctly in the SAML action, NetScaler fails to send the artifact resolution request to the Identity Provider.

[# 641913]

- 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.

[# 623053]

- 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.

[# 610604, 618760]

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

[# 653702]

AppFlow

- The NetScaler appliance fails to work if the built-in NOPOLICY policy is configured as a global policy and AppFlow data is included in application firewall records sent to NetScaler Insight Center.

[# 656771]

- 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.

[# 650618, 653126, 661587, 664792]

Application Firewall

- 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".

[# 642238, 646749, 650320]

- A NetScaler Application Firewall appliance in a high availability configuration might run out of memory, because firewall sessions might not get cleaned up 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 session entries properly.

[# 646293, 645547, 658502]

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

[# 656708, 656061, 658404, 670134]

- 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".

[# 665631, 669524]

- 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.

[# 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.

[# 662359]

- On a NetScaler appliance in a high availability configuration, a cookie header length of 61 or greater triggers Buffer Overflow violations and causes Application Firewall related scripts to fail.

[# 646775, 648952]

- 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.

[# 662734]

- 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."

[# 661111]

- **Enhancement: Application Firewall Log Messages**

NetScaler Application Firewall can now generate log messages when it triggers a page redirect because of being unable to parse the content-type header format properly.

[# 654228]

- 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.

[# 660528]

- 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.

[# 658561, 639647]

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.

[# 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.

[# 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.

[# 657325, 653722, 659696, 661214]

DNS

- 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.

[# 645836]

DataStream

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

[# 629504]

Domain Name System

- If you have configured a large number of HTTP profiles, occasionally the NetScaler appliance might dump core and restart.

[# 662929]

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.

[# 616321]

- 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.

[# 658093]

Load Balancing

- 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.

[# 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.

[# 628937]

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

[# 656734]

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

[# 656779, 663610, 664051, 664056, 664765]

- 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.

[# 658463, 658940]

NITRO

- When using the NITRO API to upload a file, make sure that each directory in the file path has the 755 (read, write, execute) permission.

For example, to upload a file to the `"/nsconfig/ssl/"` directory, the following directories must have the 755 permission:

- flash (because the `"/nsconfig"` folder is actually a link to `"/flash/nsconfig/"` directory)

- nsconfig

- ssl

[# 591970, 597032]

NetScaler GUI

- 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.

[# 653058]

- The download JRE link provided in the Install WebFront page is broken.

[# 654196]

NetScaler Insight Center

- Geo location does not resolve from Maxmind GeoCity.dat database file.

[# 656017]

- Geo location does not resolve IP addresses from maxmind data base file GeoCity.dat for both Web and HDX Insight.

[# 660424, 659826]

- After you update the country/region/city information for the IP Block, Geo report displays the existing information.

[# 617352]

- When XenApp/XenDesktop users launch applications/desktops that have the Advanced Encryption policy enabled, memory allocation issues cause high-availability failovers.

[# 659728]

- 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.

[# 660894, 662489, 668651]

Networking

- 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.

[# 644265]

- 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.

[# 597458]

SSL

- 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.

[# 620964, 576828, 641041]

- The version displayed in syslog is SSLv2.0 even though the session is negotiated using TLSv1.2.

[# 474417, 474413]

- 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.

[# 628135]

- For each SSL record that a NetScaler appliance receives from a server, the appliance opens a TCP window corresponding to the SSL layer's overhead. However, if SHA256 ciphers are used, the MAC size is incorrectly calculated as 20 bytes instead of 32 bytes. As a result, the advertised window shrinks by 12 bytes and, for a large file transfer, a TCP zero window is advertised and further transactions are dropped.

[#656605]

System

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

[# 646496]

- 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.

[# 654080]

- If a 10G interface is disabled, the lfnDiscard counter might increase.

[# 647082]

- 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

[# 654560]

- 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.

[# 659197, 656437]

- 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.

[# 633308, 646753, 648657, 656502]

- During a TCP transaction, when the client advertises zero window to a NetScaler appliance, the appliance periodically sends zero window probe to ascertain if the client can open the window so that the NetScaler appliance can send in new data. When sending such a probe, the appliance sends a full maximum segment size (MSS) packet during first probe and from the second probe onwards, sends a 1-byte packet. If the client does not open the window after sending such a probe, but instead sends a TCP Reset or if the connection on the NetScaler appliance gets flushed for other reasons, then it may lead to duplicate buffer free on the appliance that might cause the appliance to fail.

[# 657742, 656548, 657753, 657771, 658352, 658507, 658526, 659842, 659849, 660345, 660812, 660998, 661018, 661266, 661511, 662353, 662493, 665448]

- 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.

[# 623977, 649735]

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

[# 639380, 657168, 657781]

- 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.

[# 649968]

- 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.

[# 619349, 626027]

- If the maximum available memory of a TCP Buffer (TCPB) is unequally divided among Packet Engines (PE) running on a NetScaler appliance, the PE will install the TCPB on a TCP connection without sufficient memory buffer. This leads to a connection reset.

[# 587114]

- When a NetScaler appliance sends out full sized persist probe packet that is more than the client advertised window, firewall drops the packet causing the connection to fail.

[# 576980]

- 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 therefore when the ACK is received for those extra bytes, NetScaler crashes.

[# 649334, 653370, 656768, 662177, 663537, 669161]

- NetScaler appliance crashes when a large host-name header is received and AppFlow logging for host-name and domain-name is enabled.

[# 660075, 655521, 664886]

Known Issues

The issues that exist in Build 69.12.

AAA-TM

- If you bind a bookmark URL to a AAA user, the published URL tab displays "no URL" in the NetScaler GUI.

[# 636785]

- 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]

- You cannot load balance external AAA servers, such as LDAP, RADIUS, or TACACS servers, in a non-default partition.

[# 621010]

Admin Partitions

- The following two issues can occur if you add an external group as a system group on a NetScaler appliance and use the "set system group" command to configure the prompt string and timeout parameters at the system group level:

1. Session timeout-When a user from an external group logs on to the NetScaler command line interface (CLI), the session timeout set for the group is not applicable to sessions in the default and non-default partitions. However, if you configure the timeout parameter by using the "set system parameter" or "set cli mode -timeout <seconds>" commands, the session times out as specified.

2. Prompt string missing-When a user from an external group logs on to the NetScaler command line interface (CLI), the prompt string does not appear in the default and non-default partitions. For example, in a default partition, instead of "<pstring>" only ">" appears, and in a non-default partition, instead of "<pstring-partitionname>" only "partitionname>" appears.

However, if you set the prompt string by using the "set system parameter" or "set cli prompt" commands, the prompt string is displayed. For example, cliprompt> appears in a default partition, and cliprompt-partitionname> appears in a non-default partition.

[# 632193, 632460]

- RPCSVR services cannot be configured in admin partitions.

[# 498477]

- The IC memory allotted to an admin partition, cannot be reduced.

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

[# 568106, 570578]

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

[# 493668, 516396]

- If you change the resource allocation for any of the Admin Partitions, the NetScaler appliance displays a blank screen.

Workaround

Do one of the following:

1. Clear browser's cache and cookies.
2. Access NetScaler GUI in browser incognito mode.
3. Access NetScaler GUI through other web browsers.
4. Disable "Use software acceleration" option in browser settings and restart your browser.

[# 621722]

AppFlow

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

[# 629223]

Application Firewall

- The cookie consistency behavior changed in release 11.0. In earlier releases, the cookie consistency check invokes sessionization. The cookies are stored in the session and signed. A "wlt_" suffix is appended to transient cookies and a "wlf_" suffix is appended to the persistent cookies before they are forwarded to the client. Even if the client does not return these signed wlf/wlt cookies, the application firewall uses the cookies stored in the session to perform the cookie consistency check.

In release 11.0, the cookie consistency check is sessionless. The application firewall now adds a cookie that is a hash of all the cookies tracked by the application firewall. If this hash cookie or any other tracked cookie is missing or tampered with, the application firewall strips the cookies before forwarding the request to the back end server and

triggers a cookie-consistency violation. The server treats the request as a new request and sends new Set-Cookie header(s).

[# 571943]

- 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]

- The results of penetration (PEN) tests against the NetScaler landing pages reveal that the coding in some landing pages seems vulnerable to attempts at passing commands to back-end servers for the servers to execute; the commands enter through the username and (probably) the password fields. The PEN tests also indicated that /cgi/login is vulnerable to cross-site request forgery (CSRF).

Workaround

Currently, the available workarounds are to install additional NetScaler appliances, arrange NetScaler instances in a tiered architecture to provide AppFirewall protections to the landing page, or to use a third party application firewall in front of the NetScaler landing pages.

[# 644851]

- On a NetScaler appliance running release 11.0 or later, the web application firewall does not always function as expected if the DefaultCharset in a profile is not specified correctly. If a request does not have a content-type header, the WAF uses the DefaultCharset specified in a profile.

[# 624978]

- When editing application firewall signatures, you cannot sort on the "Enabled" column.

[# 621333]

- When XML output of an IBM AppScan report is imported into NetScaler, it shows zero entries if the report is based on version 2.2 of the IBM AppScan schema. Citrix supports the 2.0 schema, but must acquire information about version 2.2 of the IBM AppScan schema in order to solve this problem.

[# 626154]

Audit Logging

- During synchronization and saving of a system configuration, if Cache Redirection (CR) policy is configured before configuring an audit message action, it results in an improper sequence of CR policy and audit message actions.

[# 622905]

Clustering

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

[# 582801]

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

[# 582805]

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

[# 583424]

- When a node is removed from a layer 3 cluster, IPv6 SNIP addresses and routes are being erroneously cleared from the appliance. IPv4 SNIP addresses and routes are not affected.

[# 542693]

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

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

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

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

[# 547749]

- Domain Name System

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]

GSLB

- In a typical GSLB deployment, when internal user logon is disabled, GSLB auto sync uses SSH keys to synchronize the configuration. In a partitioned environment, however, GSLB auto sync cannot use SSH keys to synchronize the configuration across the GSLB sites.

Workaround: To use GSLB auto sync in partitioned environment, enable internal user logon and make sure that the partition user name is the same at the local and remote GSLB sites.

[# 625997]

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

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

[# 511994]

High Availability

- If you upgrade a NetScaler appliance in a high availability (HA) setup to the latest 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.

[# 611197]

Integrated Caching

- After an upgrade, the content acceleration feature is not supported.

[# 597415]

Load Balancing

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

[# 511222]

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

- An optimal XenApp/XenDesktop is launched by using determine_services in the policy expression.

- Static proximity is used to create a preferred list of Desktop Delivery Controllers and this information is forwarded to the StoreFront.

- Your connection is terminated or disconnected while the determine_services policy is being evaluated.

[# 668766]

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

[# 573463]

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

[# 540965]

NetScaler Documentation

- The following XM wizard names: `_XM_<VIP_IP>` should be reserved. These names should not be re-used for additional configuration entities. If the same convention is re-used, the edit/delete of the XM wizard could delete the configuration; even though, it was not pushed by the wizard.

[# 588178, 610159]

NetScaler GUI

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

[# 540980]

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

[# 541011]

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

- All builds of NetScaler 9.3

- All builds of NetScaler 10.1

- Any build before Build 57.x of NetScaler 10.5

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

[# 563410]

- In Security Insight, the Search functionality in Application Summary table does not work.

[# 630276]

NetScaler Insight Center

- When Web Insight displays URL records, the maximum size of a URL is limited to 1472 bytes.

[# 500108]

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

[# 554960]

- The current-connection details displayed on the NetScaler Insight Center dashboard have a latency of about 2 minutes.

[# 536696]

- Adding a new database node is now driven by auto-registration. When a kernel is imported, it requests input from the user and auto-registers with the NetScaler Insight Center server. Removing a database node is currently not supported.

[# 543632, 565706, 567628, 570264]

- Hiding or displaying a URL, or changing the configuration, might take longer than expected.

[# 570896, 574278]

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

[# 573089]

- Security Insight might display an incorrect total-violations count for some applications, because of a delay in receiving the safety profile configuration data.

[# 627373]

- If Expander module in NetScaler appliance fails, the NetScaler Insight Center skips monitoring a few ICA connections.

[# 631367]

- When, during a scale-out deployment of NetScaler Insight, you configure the database and connector nodes and register them with the Insight server, the console might display error code 0003 for the connectors and databases.

[# 631504, 565655]

- On the Security Insight dashboard, the time slider for custom time duration might not always work.

[# 630524]

- If you select Enable URL Data Collection in the Web Insight URL Data Collection Settings, the NetScaler Insight Center virtual appliance's available memory reduces rapidly.

[# 638324]

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

[# 570619]

- In NetScaler Insight Center, the Google geo chart sometimes does not display all regions.

[# 537007]

- If you export CSV files of WAN Insight reports, many of the fields in the CSV files might be empty.

[# 547380]

- In NetScaler Insight Center, some countries are not displayed on the Google geo chart.

[# 537003]

NetScaler VPX Appliance

- In an ESX environment, if a CLAG or Node LAG is created with one or more VMXNET3 interfaces on a NetScaler VPX instance then the NetScaler GUI might show the MAC address of the CLAG or Node LAG as 00:00:00:00:00:00.

[# 642495]

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

[# 641748]

Networking

- During a force sync operation in a cluster deployment, built-in compression policies fail on nodes because the nodes ignore the "Resource already exists" error for these built-in policies.

[# 648182]

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

```
set ns timeout -anyTcpClient 50 -anyTcpServer 50
```

[# 569874]

- **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.

[# 530670, 665196]

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]

- If your configuration includes policies or actions

that use the XPATH(), XPATH_WITH_MARKUP(), XPATH_JSON(), XPATH_JSON_WITH_MARKUP(), XPATH_HTML(), or XPATH_HTML_WITH_MARKUP() functions and one of those functions

fails to obtain memory that it needs, then the appliance may crash with a segmentation violation.

[# 656646]

SSL

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

[# 572429]

- If you bind a certificate-key pair to a DTLS virtual server, the following incorrect error message might appear:

No usable ciphers configured on the SSL vserver

You can safely ignore this message.

[# 542973]

- If you add or remove an HSM key, the following error message appears. However, the key is added or removed successfully.

ERROR: Internal error while adding HSM key.

[# 577552]

- 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]

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

[# 554782]

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

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

[# 576274]

- If you run the "show ssl servicegroup" command on the IP address of the cluster configuration coordinator, TLS protocol versions 1.1 and 1.2 appear as disabled, but they appear as enabled on the cluster nodes.

[# 669917]

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

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

Workaround: Import the file with a different name.

[# 526433]

- 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.

[# 661818]

Security Insight

- Security Insight uses late accounting for historical reporting. When you view the reports in the dashboard, you might observe the following behavior for the selected duration options:

[1] 1 hour: Data for security violations triggered in last 1 minute might not be included.

[2] 1 day: Data for security violations triggered in last 1 hour might not be included.

[3] 1 week: Data for security violations triggered in last 1 day might not be included.

[4] 1 month: Data for security violations triggered in last 1 day might not be included.

[# 619713]

System

- 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]

- If a NetScaler appliance sends a large number of packets on a TCP connection and if few packets get randomly dropped in the network, it leads to multiple sets of continuous packet loss (Holes). When the appliance retransmits the packets in these sets of continuous packet loss, it results in packet drops at the NetScaler Interface Card (NIC).

[# 643929]

- 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]

- The "tech support" API has new parameters: casenumber, description, upload, username, and password. If you configure these parameters, the API uploads the collector file to the Citrix server.

NOTE: Recommended practice is to use the POST HTTP method, with all the arguments in the request, so that your logon credentials are not logged in the apache log file.

[# 551873]

- 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-injection feature might cause dropped requests, closed connections, and possible failure of the NetScaler appliance. The HTML-injection feature generates a special request for each embedded object, for sending timestamp-related information to the EdgeSight server. The request URL contains the content type of the object. If the Content-Type field in the request contains a space, it should be percent-encoded, but the HTML-injection feature inserts the space as is. Therefore, by HTTP standards, the request is invalid. If the "drop invalid requests" option is enabled in the applicable HTTP profile, the request is dropped and the connection is closed. Also, if the URL spans multiple packets, the NetScaler appliance fails while processing the next packet after the request is marked invalid.

[# 626848]

- If you do not configure the Maximum Transmission Unit (MTU) in the VLAN or NIC interface, a NetScaler appliance does not appropriately advertise the Maximum Segment Size (MSS) option in the TCP SYN packet sent to the back-end server, which results in packet drops and a transaction failure.

Workaround: Configure the correct MSS option in the TCP profile.

[# 627394]

- A NetScaler appliance fails if the Front End Optimization (FEO) feature is enabled and Page Extend Cache is turn on in the FEO action.

[# 658045]

Telco

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

[# 540162]

- In the output of the "show lsn sipalgcalls -callid" command, the port value of the SIP control channel is incorrect.

[# 574257]

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

[# 574725]

WIoNS

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

[# 507753]

What's New in Previous NetScaler 11.0 Releases

The enhancements and changes that were available in NetScaler 11.0 releases prior to Build 69.12. The build number provided below the issue description indicates the build in which this enhancement or change was provided.

AAA-TM

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

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

[From Build 55.23] [# 512125]

Logging Errors in NetScaler Log Files

The NetScaler appliance now stores AAA authentication logs.

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

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

[From Build 55.23] [# 482228, 479557]

- **Supporting Encrypted Assertions on SAML SP**

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

[From Build 55.23] [# 291693]

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

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

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

From the CLI:

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

From the GUI:

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

[From Build 55.23] [# 474977]

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

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

[From Build 55.23] [# 496725, 508689]

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

[From Build 55.23] [# 472912]

- **Encrypting SAML IdP Assertion**

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

Note:

- Make sure the SAML SP certificate is specified.
- For enhanced security, it is recommended that you encrypt assertions that contain sensitive information.

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

On the CLI:

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

On the GUI:

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

[From Build 55.23] [# 482185]

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

[From Build 55.23] [# 524386]

- **Using Cookies to Track SAML Sessions**

In a deployment where a NetScaler appliance is configured as a SAML IdP (identity provider) for multiple SAML SPs (service provider), the appliance allows a user to access multiple SPs without explicitly authenticating every time. The appliance creates a session cookie for the first authentication and every subsequent request uses this cookie for authentication.

[From Build 55.23] [# 503882]

- **Fallback to NTLM Authentication**

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

This feature allows user logins locally or remotely.

[From Build 55.23] [# 509829]

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

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

Using the CLI:

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

[From Build 55.23] [# 493220, 462777, 493224]

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

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

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

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

In the NetScaler implementation, the application to be accessed is represented by the AAA-TM virtual server. So, to configure OAuth, an action must be configured and associated with a AAA-TM policy which is then associated with a AAA-TM virtual server. The configuration to define a OAuth action is as follows:

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

Note:

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

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

[From Build 55.23] [# 491920]

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

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

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

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

[From Build 55.23] [# 550946]

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

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

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

From the CLI:

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

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

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

From the GUI:

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

[From Build 55.23] [# 460680, 504703]

- The NetScaler appliance now supports the SiteMinder SAML SP.

[From Build 55.23] [# 488077]

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

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

In NetScaler terminology, this feature is called "nFactor Authentication."

[From Build 62.10] [# 482250, 451913, 549966]

Admin Partitions

- **Supporting Dynamic Routing in Admin Partitions**

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

```
> set L3Param -dynamicRouting ENABLED
```

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

[From Build 55.23] [# 514848]

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

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

[From Build 55.23] [# 534986]

- **Getting NetScaler Trace for Specific Partitions**

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

[From Build 55.23] [# 496937, 515294]

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

[From Build 55.23] [# 535494]

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

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

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

- Partition1: 4 GB

- Partition2: 3 GB

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

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

[From Build 55.23] [# 481444, 484618]

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

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

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

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

[From Build 55.23] [# 513564]

- **Partition Specific Load Balancing Parameters**

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

Note:

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

- These parameters are set in the CLI by using the "set lb parameter" command or in the GUI by navigating to Traffic Management > Load Balancing.

[From Build 62.10] [# 563004]

- The following load balancing features can now be configured in admin partitions:
 - DBS autoscale
 - Stateless connection mirroring
 - RDP
 - Radius
 - Graceful shutdown

[From Build 64.34] [# 588406]

The NetScaler appliance now supports FTP load balancing in admin partitions.

[From Build 64.34] [# 568811]

- **GSLB Support in Admin Partitions**

The NetScaler appliance now supports the GSLB feature in admin partitions. You can now deploy, with an admin partition, applications that need the GSLB feature to distribute traffic across globally located datacenters.

[From Build 65.35] [# 436582, 405290, 504574, 506221]

- **Stateful Connection Failover/Mirroring Support in Admin Partitions**

The NetScaler appliance now supports stateful connection mirroring in admin partitions. You can now deploy TCP-based applications in an admin partition, so that failure of one NetScaler appliance does not make the application unavailable.

Note: The application must be deployed in a NetScaler high availability (HA) setup, and connection mirroring must be configured for the application.

[From Build 65.35] [# 599629]

- **AAA-TM Support in Admin Partitions**

The NetScaler appliance now supports the AAA-TM feature in admin partitions. You can now deploy, with an admin partition, enterprise applications that require authenticated access.

[From Build 65.35] [# 481384]

Application Firewall

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

[From Build 55.23] [# 506157]

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

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

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

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

[From Build 55.23] [# 408831, 403780]

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

[From Build 55.23] [# 520048]

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

[From Build 55.23] [# 450326, 483677, 513927]

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

Command Line Interface:

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

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

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

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

[From Build 55.23] [# 450324, 483683]

- Geolocation, which identifies the geographic location from which requests originate, can help you configure the application firewall for the optimal level of security. For example, if an excessively large number of requests are received from a specific area, it is easy to determine whether they are being sent by users or a rogue machine. The application firewall offers you the convenience of using the built-in NetScaler database or any other geolocation based database to identify the source of origin of coordinated attacks launched from a country. This information can

be quite useful for enforcing the optimal level of security for your application to block malicious requests originating from a specific geographical region. Geolocation logging uses the Common Event Format (CEF).

To use Geolocation Logging

1. Enable CEFLogging and GeoLocationLogging.

```
>set appfw settings GeoLocationLogging ON CEFLogging ON
```

2. Specify the database

```
>add locationfile /var/netScaler/inbuilt_db/Citrix_NetScaler_InBuilt_GeolP_DB.csv
```

or

```
add locationfile <path to database file>
```

[From Build 55.23] [# 483703]

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

Examples of CLI Commands:

1. Bind the credit card number to profile:

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

2. Unbind credit card number from profile:

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

3. Log: Enable Logging of credit card Numbers

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

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

4. Learn:

```
show appfw learningdata <profilename> creditCardNumber
```

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

```
export appfw learningdata <profilename> creditCardNumber
```

[From Build 55.23] [# 383298]

- NetScaler now supports the IP Reputation feature, which is useful in identifying an IP address that is sending unwanted requests. You can use the IP reputation list to preemptively reject requests that are coming from an IP with a bad reputation. NetScaler uses WebRoot as the service provider for the dynamically generated malicious IP database and the metadata for those IPs. The IP Reputation feature can be configured by using PI Expressions in a policy. For example, you can configure an application firewall policy using expressions such as: CLIENT.IP.SRC.IPREP_IS_MALICIOUS.

[From Build 65.35] [# 580866]

Cache Redirection

- **Support for default syntax expressions**

You can now use default syntax expressions in cache redirection policies. The NetScaler appliance provides built-in cache redirection policies based on default syntax expressions, or you can create custom cache redirection policies to handle typical cache requests. In addition to the same types of evaluations done by classic cache redirection policies, the default syntax policies enable you to analyze more data (for example, the body of an HTTP request) and to configure more operations in the policy rule (for example, directing requests to either cache or origin server).

[From Build 55.23] [# 490297, 495915, 536986, 536992, 537010, 537014, 538269]

Cluster

- **Link Redundancy based on Minimum Throughput**

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

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

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

[From Build 55.23] [# 508993]

- **Disabling Steering on the Cluster Backplane**

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

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

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

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

```
> add cluster instance <clId> -processLocal ENABLED
```

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

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

[From Build 55.23] [# 539136]

- **BridgeGroups are now supported in a NetScaler cluster deployment**

[From Build 55.23] [# 494991]

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

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

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

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

[From Build 55.23] [# 317706]

- **Nodegroup for Datacenter Redundancy**

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

[From Build 55.23] [# 495019]

- **Routing in a L3 Cluster**

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

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

Note:

- This option is also available for VXLAN and BridgeGroups.

- This configuration can also be used for L2 clusters.

[From Build 55.23] [# 531868]

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

You can now create a cluster that includes nodes from different networks. To configure a cluster over L3, you must add the nodes of different networks to different nodegroups.

You can transition an existing L2 cluster to an L3 cluster.

[From Build 55.23] [# 374289, 317257]

- **Routing on Striped SNIP addresses**

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

Note:

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

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

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

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

[From Build 55.23] [# 329439]

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

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

1. Make sure that the Java package and the WI package are installed in the same directory on all the cluster nodes.
2. Create a load balancing virtual server that has persistency configured.
3. Create services with IP addresses as the NSIP address of each of the cluster nodes that you want to serve WI traffic.
4. Bind the services to the load balancing virtual server.

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

[From Build 62.10] [# 498295, 489463]

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

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

[From Build 62.10] [# 513612]

- **Cluster versioning**

When you are upgrading a cluster to NetScaler 11.0 build 64.x from an earlier NetScaler 11.0 build, cluster configuration propagation is disabled.

Traditionally, this issue occurred only during an upgrade of a cluster to a different NetScaler version (for example, from 10.5 to 11.0). This exception arises because the cluster version in build 64.x is different from the one in previous NetScaler 11.0 builds.

Note: Normally, the cluster version matches the NetScaler version.

Configuration propagation remains disabled until all the cluster nodes are upgraded to Build 64.x.

[From Build 64.34] [# 591877]

- **Reducing the Minimum Value for the Dead Interval**

You can now set the dead interval for a cluster instance to a minimum value of 1 second.

Note: If the dead interval value is less than 3 seconds, set the hello interval parameter to 200 ms.

[From Build 64.34] [# 573218]

Command Line Interface

- The NetScaler administrator can now specify the maximum number of concurrent sessions a user can log on to the CLI. Although logons to the configuration utility do not count against the limit, all logon attempts are denied after the limit is reached. For example, if the maximum number of concurrent sessions is set to 20, a user can log on to the CLI 19 times and can log on to the configuration utility any number of times. Once the user logs on to the CLI for the 20th time, he or she can no longer log on to the CLI or the configuration utility. Any logon attempt then results in a system error message.

[From Build 64.34] [# 491778]

DNS

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

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

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

[From Build 55.23] [# 391254]

- **Support for DNS Logging**

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

- Audit the DNS responses to the client
- Audit DNS clients
- Detect and prevent DNS attacks

- Troubleshoot

[From Build 55.23] [# 419632, 561291]

Rewrite and responder support for DNS

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

[From Build 55.23] [# 405769]

GSLB

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

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

[From Build 55.23] [# 373061]

- **GSLB Service Selection using Content Switching**

Description: You can now configure a content switching (CS) policy to customize a GSLB deployment so that you can:

* Restrict the selection of a GSLB service to a subset of GSLB services bound to a GSLB virtual server for the given domain.

* Apply different Load Balancing methods on the different subsets of GSLB services in the deployment.

* Apply spillover policies on a subset of GSLB services, and you can have a backup for a subset of GSLB services.

* Configure a subset of GSLB services to serve a specific type of content.

* Define a subset GSLB services with different priorities, and define the order in which the services in the subset are applied to a request.

For more information, see [Configuring GSLB Service Selection Using Content Switching](#).

[From Build 63.16] [# 503588]

- **Ability to specify GSLB Site IP address as source IP address for an RPC node**

You can now configure the NetScaler appliance to use GSLB Site IP address as the source IP address for an RPC node.

[From Build 64.34] [# 531395]

- **NetScaler GSLB deployments support NAPTR records**

In GSLB deployments, the NetScaler appliance now supports DNS queries with NAPTR records. You can now configure a NetScaler appliance to receive DNS queries with NAPTR records from clients (for example, Mobile Management Entity (MME)) and respond with the list of services configured for a domain. Also, the NetScaler appliance monitors the health of the services and in the response it provides only the list of services that are up.

[From Build 64.34] [# 468647]

HDX Insight

- HDX Insight now supports displaying of AppFlow records from NetScaler cluster.

[From Build 62.10] [# 525758]

Load Balancing

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

[From Build 55.23] [# 524079, 559022]

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

At the NetScaler command prompt, type:

```
set lb parameter -useSecuredPersistenceCookie Enabled-cookiePassphrase test
```

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

[From Build 55.23] [# 347108, 323325, 348588]

- **Automatic Restart of the Internal Dispatcher**

In earlier releases, if the internal dispatcher failed, the services that used scriptable monitors also went down and the appliance had to be restarted. From release 11, if the internal dispatcher fails, the pitboss process restarts it. As a result, you no longer have to restart the appliance.

[From Build 55.23] [# 368128]

- **New Trap for Spillover**

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

[From Build 55.23] [# 486268, 475400]

- The following global timeouts has been introduced for TCP sessions on a NetScaler appliance related to RNAT rules, forwarding sessions, or load balancing configuration of type ANY:
 - * Any TCP Client. Global idle timeout, in seconds, for TCP client connections. Client timeout set for an entity overrides the global timeout setting.
 - * Any TCP Server. Global idle timeout, in seconds, for TCP server connections. Server timeout set for an entity overrides the global timeout setting.

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

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

[From Build 55.23] [# 507701]

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

You can now use IPv6 addresses in the following monitors:

- USER
- SMTP
- NNTP
- LDAP
- SNMP
- POP3
- FTP_EXTENDED
- STOREFRONT

- APPC

- CITRIX_WI_EXTENDED

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

[From Build 55.23] [# 510111]

- **Support for Secure LDAP Monitor**

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

[From Build 55.23] [# 418061, 556530]

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

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

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

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

- Explicitly disable the service and:

- Set a delay (in seconds).

- Enable graceful shut down.

Adding a TROFS Code or String

Note: This enhancement is not applicable to GSLB services.

From release 11, if you bind only one monitor to a service, and the monitor is a TROFS-enabled monitor, it can place the service in the TROFS state on the basis of the server's response to a monitor probe. This response is compared with the value in the trofsCode parameter for an HTTP monitor or the trofsString parameter for an HTTP-ECV or TCP-ECV monitor. If the code matches, the service is placed in the TROFS state. In this state, it continues to honor the persistent connections.

If multiple monitors are bound to a service, the effective state of the service is calculated on the basis of the state of all the monitors that are bound to the service. Upon receiving a TROFS response, the state of the TROFS-enabled monitor is considered as UP for the purpose of this calculation.

Important!

- You can bind multiple monitors to a service, but only one monitor must be TROFS-enabled.
- You can convert a TROFS-enabled monitor to a monitor that is not TROFS-enabled, but not vice versa.

[From Build 55.23] [# 408103]

- **Support for Unauthenticated Stores**

In earlier releases, the StoreFront monitor tried to authenticate anonymous stores. As a result, a service could be marked as DOWN and you could not launch XenApp or XenDesktop by using the URL of the load balancing virtual server.

The probe order has changed. The monitor now determines the state of the StoreFront store by successively probing the account service, the discovery document, and then the authentication service, and skips authentication for anonymous stores.

[From Build 64.34] [# 575549]

- With the following new OID, you can use SNMP to learn the current number of server connections per service.

svcCurSvrConnections, 1.3.6.1.4.1.5951.4.1.2.1.1.59

[From Build 64.34] [# 548470]

- **Retaining the Original State of a Service Group Member after Disabling and Enabling a Virtual Server**

A new global option, `-retainDisableServer`, enables you to retain a service-group member's state when a server is disabled and reenabled.

Previously, a member's state would change from DISABLED to ENABLED under the following set of conditions:

- Two applications are deployed on the same port on a virtual server.
- Two service groups with a common member are bound to this virtual server, and the common member is enabled in one group and disabled in the other.
- The server is disabled and then re-enabled.

Under these conditions, disabling the server disables all the service group members, and re-enabling the server enables all the members, by default, regardless of their earlier states. To bring the members back to the original states, you must manually disable those member(s) in the service group. This is a cumbersome task and prone to errors.

[From Build 64.34] [# 493692]

- With the following new OID, you can use SNMP to learn the effective state of a virtual server.

vsvrCurEffState, 1.3.6.1.4.1.5951.4.1.3.1.1.75

[From Build 64.34] [# 538499]

NetScaler Insight Center

- **Insight Deployment Management**

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

[From Build 55.23] [# 404919]

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

[From Build 55.23] [# 519788, 522021]

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

[From Build 55.23] [# 425761, 553254]

- **Hop Diagram Support**

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

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

[From Build 55.23] [# 443824]

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

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

[From Build 55.23] [# 514612]

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

[From Build 55.23] [# 491073]

- NetScaler Insight Center now supports monitoring NetScaler appliances deployed in LAN user mode. The dashboard now displays the following user access types, depending on the NetScaler deployment:

- Transparent mode user: User connected to XenApp or XenDesktop server directly, with no intervening virtual server.

- LAN user: Internal user connected to XenApp or XenDesktop server directly, without configuring the routing rules on a NetScaler ADC.

[From Build 55.23] [# 490147, 482900]

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

[From Build 55.23] [# 502478]

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

[From Build 55.23] [# 488279]

- **Exporting Reports**

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

[From Build 55.23] [# 320860]

Security Insight

- Web and web service applications that are exposed to the Internet have become increasingly vulnerable to attacks. To protect applications from attack, you require visibility into the nature and extent of past, present, and impending threats, real-time actionable data on attacks, and recommendations on countermeasures. Security Insight provides a single-pane solution to help you assess your application security status and take corrective actions to secure your applications.

Security insight is included in NetScaler Insight Center, and it periodically generates reports based on your Application Firewall and NetScaler system security configurations. The reports include the following information for each application:

-Threat index. A single-digit rating system that indicates the criticality of attacks on the application, regardless of whether or not the application is protected by a NetScaler appliance. The more critical the attacks on an application, the higher the threat index for that application.

-Safety index. A single-digit rating system that indicates how securely you have configured the NetScaler devices to protect applications from external threats and vulnerabilities. The lower the security risks for an application, the higher the safety index.

-Actionable Information. Information that you need to lower the threat index and increase the safety index, which significantly improves application security. For example, you can review information about violations, existing and missing security configurations in Application Firewall and NetScaler security features, the rate at which the applications are being attacked, and so on.

[From Build 65.35] [# 587137]

- 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 65.35] [# 614892]

- **Viewing the GET or POST requests**

The NetScaler Insight Center now displays the GET or POST requests that are sent by the client to a domain. To view the GET or POST requests, navigate to Domains > URLs > Clients > Http Request Method, or to Domains > URLs > Http Request Method > Clients.

[From Build 65.35] [# 620323]

- **NetScaler Insight Center is now supported on KVM hypervisors.**

[From Build 66.11] [# 631295]

NetScaler VPX Appliance

- **New license for NetScaler VPX on ESX Platform**

The following licenses are now available for NetScaler VPX appliance on ESX platform:

- 25M

- 5G

- 10G

- 15G

- 25G

- 40G

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

[From Build 68.10] [# 623179]

- **New license for NetScaler VPX on ESX Platform**

The following licenses are now available for NetScaler VPX appliance on ESX platform:

- 25M

- 5G

- 10G

- 15G

- 25G

- 40G

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

[From Build 68.11] [# 623179]

Networking

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

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

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

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

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

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

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

- Set the MTU of the interface or channel of the VPX appliance to a value in the range 1501-9216. Use the NetScaler command line interface or the configuration utility of the VPX appliance to set the MTU size.
- Set the same MTU size on the corresponding physical interfaces of the virtualization host by using its management applications.

[From Build 55.23] [# 464830, 478103, 485905]

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

The NetScaler appliance, for INAT and RNAT rules, now supports using client port as the source port for server side connections. A parameter Use Proxy Port has been added to the INAT and RNAT command set. When Use Proxy Port

is disabled for an INAT rule or an RNAT rule, the NetScaler appliance retains the source port of the client's request for the server side connection. When the option is enabled (default), the NetScaler appliance uses a random port as the source port for the server side connection. You must disable this parameter for proper functioning of certain protocols that require a specific source port in the request packet.

[From Build 55.23] [# 399821]

- **Redundant Interface Sets**

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

Following are the main benefits of using redundant interface sets:

- The back-up links between the NetScaler appliance and a peer device ensure connection reliability.
- Unlike link redundancy using LACP, no configuration is required on the peer device for a redundant interface set. To the peer device, a redundant interface set appears as individual interfaces, not as a set or collection.
- In a high availability (HA) configuration, redundant interface sets can minimize the number the HA failovers.

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

[From Build 55.23] [# 355237, 186503, 249551]

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

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

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

[From Build 55.23] [# 512848]

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

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

VRRP uses the following communication intervals:

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

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

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

[From Build 55.23] [# 512843]

- **OSPFv3 Authentication**

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

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

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

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

[From Build 55.23] [# 471703]

- **Blocking Traffic on Internal Ports**

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

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

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

At the command prompt, type:

```
> set l3param -implicitACLAllow [ENABLED|DISABLED]
```

Note: The parameter implicitACLAllow is enabled by default.

Example

```
> set l3param -implicitACLAllow DISABLED
```

Done

[From Build 55.23] [# 529317]

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

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

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

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

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

[From Build 55.23] [# 503566]

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

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

MAC Address Wildcard Mask for PBRs

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

[From Build 55.23] [# 391630]

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

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

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

- Carry the Ethernet header but drop the VLAN header
- Drop the Ethernet header as well as the VLAN header
- Carry the Ethernet header as well the VLAN header

[From Build 55.23] [# 518397]

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

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

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

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

Following are the benefits of this support:

- Instead of defining new PBRs that are based on Layer 3 parameters, you can use existing PBRs based on Layer 2 parameters for sending the packets related to forwarding sessions to the desired next hop device.
- In a deployment that includes NetScaler appliances and optimization devices, PBRs based on Layer 2 parameters can be very handy compared to other, complex configuration for identifying the forwarding session related packets for PBR processing.
- Identifying forwarding session related Ingress packets for sending them to the optimization device.
- Identifying egress packets, which also matched a forwarding session rule, from the optimization device for sending the packets to the desired next hop device.

[From Build 55.23] [# 484458]

- **Logging HTTP Header Information**

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

- URL that the HTTP request is destined to.
- HTTP Method specified in the HTTP request.
- HTTP version used in the HTTP request.
- IP address of the subscriber that sent the HTTP request.

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

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

[From Build 55.23] [# 496835]

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

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

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

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

[From Build 55.23] [# 520355]

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

[From Build 55.23] [# 472443]

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

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

[From Build 55.23] [# 472432]

- **Delaying Preemption**

By default, a backup VIP address preempts the master VIP address immediately after its priority becomes higher than that of the master VIP. When configuring a backup VIP address, you can specify an amount of time by which to delay the preemption. Preemption delay time is a per-node setting for each backup VIP address.

The preemption delay setting for a backup VIP does not apply in the following conditions:

* The node of the master VIP goes down. In this case, the backup VIP takes over as the master VIP after the dead interval set on the backup VIP's node.

* The priority of the master VIP is set to zero. The backup VIP takes over as the master VIP after the dead interval set on the backup VIP's node.

[From Build 64.34] [# 553246]

- **Keeping a VIP address in the Backup State**

You can force a VIP address to always stay in backup state in a VRRP deployment. This operation is helpful in maintenance or testing of the deployment.

When a VIP address is forced to stay in backup state, it does not participate in VRRP state transitions. Also, it cannot become master even if all other nodes go down.

To force a VIP address to stay in backup state, you set the priority of the associated VMAC address to zero. To ensure that none of the VIP addresses of a node handle traffic during a maintenance process on the node, set all the priorities to zero.

[From Build 64.34] [# 553311]

- **Keeping a VIP address in the Backup State**

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 64.34] [# 420067, 420039]

- **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 65.35] [# 457167]

Optimization

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

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

[From Build 55.23] [# 504044]

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

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

[From Build 55.23] [# 509338]

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

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

[From Build 55.23] [# 493103]

- **Support for SR-IOV interfaces on NetScaler VPX Appliance in VMware ESX**

You can now configure NetScaler VPX appliance deployed on VMware ESX 6.0 to use SR-IOV network interfaces. The NetScaler VPX appliance now supports Intel 82599 10g Network Interface Card (NIC).

For performance information of SR-IOV interface on ESX, refer the latest VPX datasheet.

[From Build 68.10] [# 637341]

- **Support for VMXNET3 interfaces on NetScaler VPX Appliance in VMware ESX**

You can now configure NetScaler VPX appliance deployed on VMware ESX 6.0 to use VMXNET3 network interfaces. The NetScaler VPX appliance now supports Intel 82599 10g Network Interface Card (NIC).

For performance information of VMXNET3 interface on ESX, refer the latest VPX datasheet.

[From Build 68.10] [# 637336]

- **Support for SR-IOV interfaces on NetScaler VPX Appliance in VMware ESX**

You can now configure NetScaler VPX appliance deployed on VMware ESX 6.0 or ESX 5.5 to use SR-IOV network interfaces. The NetScaler VPX appliance now supports Intel 82599 10g Network Interface Card (NIC).

For performance information of SR-IOV interface on ESX, refer the latest VPX datasheet.

[From Build 68.11] [# 637341]

- **Support for VMXNET3 interfaces on NetScaler VPX Appliance in VMware ESX**

You can now configure NetScaler VPX appliance deployed on VMware ESX 6.0 or ESX 5.5 to use VMXNET3 network interfaces. The NetScaler VPX appliance now supports Intel 82599 10g Network Interface Card (NIC).

For performance information of VMXNET3 interface on ESX, refer the latest VPX datasheet.

[From Build 68.11] [# 637336]

Policies

- **Policy extensions support on NetScaler appliance**

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

[From Build 55.23] [# 248822]

- **Transaction Scope Variables**

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

[From Build 55.23] [# 444109]

SSL

- **Stricter Control on Client Certificate Validation**

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

[From Build 55.23] [# 533241]

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

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

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

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

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

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

[From Build 55.23] [# 449923]

- **Changes to the Default Cipher Suite**

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

[From Build 55.23] [# 226713, 258311, 384491]

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

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

[From Build 55.23] [# 439161]

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

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

[From Build 55.23] [# 451441, 405363, 458905, 465280, 540467, 551603, 559154]

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

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

At the NetScaler command line, type:

```
show ciphersuite
```

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

[From Build 55.23] [# 491286]

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

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

[From Build 55.23] [# 250573]

- **DH Key Performance Optimization**

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

[From Build 55.23] [# 498162, 512637]

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

The NetScaler appliance now supports the `TLS_FALLBACK_SCSV` signaling cipher suite value. The presence of this SCSV extension in the Client Hello indicates that the client is retrying to connect to the server by using a lower SSL version, after its previous attempt to communicate with a higher version failed. Therefore, if the server finds this extension in Client Hello and also finds that the client is proposing a version that is lower than the maximum version supported by the server, it is a likely indication of a "man in the middle attack." The server drops these handshakes.

[From Build 55.23] [# 509666, 573528]

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

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

[From Build 55.23] [# 508440, 483391]

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

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

[From Build 55.23] [# 209047, 432330, 481660]

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

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

[From Build 55.23] [# 424463, 481970]

- **Support for Thales nShield(R) HSM**

All NetScaler VPX appliances now support the Thales nShield(R) Connect external Hardware Security Module (HSM). With a Thales HSM, the keys are securely stored as application key tokens on a remote file server and can be reconstituted only inside the Thales HSM.

Thales integration with the ADC is supported for TLS versions 1.0, 1.1, and 1.2.

[From Build 62.10] [# 440351, 477544]

- **Enhanced SSL Profile**

The SSL infrastructure on the NetScaler appliance is continually updated to address the ever growing requirements for security and performance. Vulnerabilities in SSLv3 and RC4 implementation have emphasized the need to use the latest ciphers and protocols to negotiate the security settings for a network connection. Implementing any changes to the configuration, such as disabling SSLv3, across thousands of SSL end points is a cumbersome process. Therefore, settings that were part of the SSL end points configuration have been moved to the SSL profile, along with the default ciphers. To implement any change in the configuration, including cipher support, you need only modify the profile. If the profile is enabled, the change is immediately reflected in all the end points that the profile is bound to.

Important: After the upgrade, if you enable the profile, you cannot reverse the changes. That is, the profile cannot be disabled.

[From Build 64.34] [# 533640]

- **Graceful Cleanup of SSL sessions after change in any SSL entity parameter**

Some operations - for example, updating a certificate to replace a potentially exposed certificate, using a stronger key (2048-bit instead of 1024-bit), adding/removing a certificate from a certificate chain, or changing any of the SSL parameters - should clean the SSL sessions gracefully instead of abruptly terminating existing sessions. With this enhancement, existing connections continue to use the current settings but all new connections use the new certificate or settings. However, connections that are in the middle of a handshake or sessions that are renegotiating are terminated, and session reuse is not allowed. To clear the sessions immediately after a configuration change, you must disable and re-enable each entity.

[From Build 64.34] [# 529979]

- **New Counters in SSL Statistics**

Because TLS 1.1 and 1.2 are becoming the primary security protocols, the transaction and session statistics for these protocols are now included in the SSL statistics.

[From Build 64.34] [# 336395, 559165, 560353]

- If you downgrade the software on your NetScaler appliance that does not have a license to release 9.3 build 61.66 or earlier, some commands related to the default server certificate might not be saved in the running configuration. As a result, after restarting, secure access (HTTPS) to the appliance fails.

[From Build 64.34] [# 551603, 559154]

- **2048-bit Default Certificates on the NetScaler VPX Instance**

You no longer need a license on your VPX instance to generate a 2048-bit default certificate. After upgrading your VPX instance to release 11.0, if you want to replace the old internal default 512-bit certificate, delete all your old certificate-key pairs that have "ns-" as the first three characters, and then restart the instance to automatically generate a 2048-bit default certificate.

[From Build 64.34] [# 451441, 405363, 458905, 465280, 540467, 547106, 551603, 559154, 584335, 588128]

- The NetScaler appliance now supports the following "signature algorithms" extensions in the back end client hello message:

- RSA-MD5

- RSA-SHA1

- RSA-SHA256

[From Build 65.35] [# 600155, 601059]

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

[From Build 65.35] [# 498207]

- **New Client Authentication Counters for SSL Virtual Servers**

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

1. ssl_ctx_tot_clientAuth_success-Tracks the number of successful client authentications for each SSL virtual server.

2. ssl_ctx_tot_clientAuth_failures-Tracks the number of failed client authentications for each SSL virtual server.

[From Build 65.35] [# 492684]

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

[From Build 65.35] [# 543526, 579749, 619662]

System

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

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

[From Build 55.23] [# 248711]

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

[From Build 55.23] [# 480258, 494482, 523853]

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

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

[From Build 55.23] [# 505840, 505377]

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

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

[From Build 55.23] [# 490096, 505747]

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

[From Build 55.23] [# 548516]

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

[From Build 55.23] [# 576067]

- **NTP Version Update**

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

If you upgrade your NetScaler appliance from any earlier release to release 11, the NTP configuration is automatically upgraded with additional security policies.

[From Build 55.23] [# 440375, 440591]

- **Support for FACK on TCP profiles**

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

[From Build 55.23] [# 439130]

- **Showtechsupport utility enhancement**

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

[From Build 55.23] [# 480797]

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

[From Build 55.23] [# 238440, 239481, 247372, 422873]

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

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

[From Build 55.23] [# 397478]

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

[From Build 55.23] [# 368832]

- **Support for MPTCP Version Negotiation**

A client can now establish an MPTCP connection with NetScaler appliance even if the client's and the NetScaler appliance's MPTCP versions does not match. If the MPTCP version of the client is higher than the one supported on the appliance, the client falls back to a lower or equal version. If the appliance supports that version, the MPTCP session continues. Otherwise, the appliance falls back to a normal TCP session.

[From Build 64.34] [# 529883]

- **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 64.34] [# 314438]

- **Support for Configuring a Proxy Server to Install Licenses**

You no longer have to configure internet connectivity on the NetScaler appliance in order to use a hardware serial number or license activation code to allocate a NetScaler license. Instead, you can use a proxy server.

On the NetScaler GUI, navigate to Configuration > System > Licenses > Manage Licenses > Add a New License, select the Connect through Proxy Server check box, and specify the IP address and port of your proxy server.

[From Build 64.34] [# 541474]

- **Support restore operation on the NetScaler Appliance by using a remotely stored backup**

You can now use a remotely saved backup to restore a NetScaler appliance through the "add system backup <filename>", that adds the metadata to the remote backup package, so that the restore operation can successfully use the backup package.

[From Build 64.34] [# 569974]

Telco

- **Provide Visibility into SLA Reports**

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

NetScaler Insight Center and the NetScaler cache redirection feature can now be used to monitor the traffic flowing through the NetScaler appliances and calculate SLA breaches. The NetScaler cache redirection feature helps save

bandwidth over international links. NetScaler Insight Center works with the NetScaler cache redirection feature to calculate, and provide visibility into, the percentage of bandwidth saved and any breaches of the SLA. ISP administrators are alerted whenever there is a breach for response time, hit rate/sec, or bandwidth.

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

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

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

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

NetScaler Insight Center then calculates the following SLA breaches:

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

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

[From Build 62.10] [# 495288, 501269, 501277, 501278, 501279, 501280]

- **Support for RADIUS Accounting Message**

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

[From Build 62.10] [# 526981]

- **Support for Gx Interface**

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

[From Build 62.10] [# 402469]

- **Provide Internet Access to IPv4 Subscribers Through the IPv6 Core Network of a Telecom Service Provider (Dual-Stack Lite)**

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

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

The NetScaler appliance implements the AFTR component of a DS-Lite deployment and is compliant with RFC 6333.

[From Build 62.10] [# 407162]

- **Subscriber-Aware Traffic Steering**

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

- Forwards the data packet to another set of services
- Drops the packet
- Performs LSN if configured on the appliance

[From Build 62.10] [# 402473]

- **Provide Internet Access to a Large Number of Private IPv4 Subscribers of a Telecom Service Provider (Large Scale NAT)**

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

The NetScaler supports LSN and is compliant with RFC 6888, 5382, 5508, and 4787. The NetScaler LSN feature is very useful for Internet Service Providers (ISPs) and carriers providing millions of translations to support a large number of users (subscribers) and at very high throughput. The LSN architecture of an ISP using Citrix products consists of

subscribers (Internet users) in private address spaces accessing the Internet through a NetScaler appliance deployed in ISP's core network.

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

- * ALGs: Support of application Layer Gateway (ALG) for SIP, PPTP, RTSP, FTP, ICMP, and TFTP protocols.
- * Deterministic/ Fixed NAT: Support for pre-allocation of block of ports to subscribers for minimizing logging.
- * Mapping: Support of Endpoint-independent mapping (EIM), Address-dependent mapping (ADM), and Address-Port dependent mapping.
- * Filtering: Support of Endpoint-independent filtering (EIF), Address-dependent filtering, and Address-Port-dependent filtering.
- * Quotas: Configurable limits on number of ports and sessions per subscriber.
- * Static Mapping: Support of manually defining an LSN mapping.
- * Hairpin Flow: Support for communication between subscribers or internal hosts using public IP addresses.
- * LSN Clients: Support for specifying or identifying subscribers for LSN NAT by using IPv4 addresses and extended ACL rules.
- * Logging: Support for logging LSN session for law enforcement. In addition, the following are also supported for logging:
 - ** Reliable SYSLOG: Support of sending SYSLOG messages over TCP to external log servers for a more reliable transport mechanism.
 - ** Load balancing of Log Servers. Support for load balancing of external log servers for preventing storage of redundant log messages.
 - ** Minimal Logging: Deterministic LSN configurations or Dynamic LSN configurations with port block significantly reduces the LSN log volume.

[From Build 62.10] [# 316909]

- **Subscriber-Aware Service Chaining**

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

subscriber to a specific set of services on the basis of the policy defined for that subscriber. The appliance receives the subscriber information from the PCRF over a Gx interface.

[From Build 62.10] [# 561747]

- **High Availability Support for Dynamic Subscriber Sessions**

In the absence of a high availability (HA) setup, the subscriber information that is received from the RADIUS client is lost if the appliance fails. With HA support, the subscriber sessions are continually synchronized on the secondary node. In the event of a failover, the subscriber information is still available on the secondary node.

[From Build 63.16] [# 574838]

- **Subscriber Session Event Logging**

The NetScaler appliance currently maintains millions of subscriber sessions in its database (subscriber store) but does not log these messages. Telco administrators need reliable log messages to track the control plane messages specific to a subscriber. They also need historical data to analyze subscriber activities. The appliance now supports logging of RADIUS control plane accounting messages and Gx control plane logging messages. Some of the key attributes are MSISDN and time stamp. By using these logs, you can track a user by using the IP address, and the MSISDN if available.

[From Build 64.34] [# 575621, 575623]

- **IPv6 Prefix based Subscriber Sessions**

A telco user can be uniquely identified by the IPv6 prefix rather than the complete IPv6 address. The NetScaler appliance now uses the prefix instead of the complete IPv6 address (/128) to identify a subscriber in the database (subscriber store). For communicating with the PCRF server (for example, in a CCR-I message), the appliance now uses the framed-IPv6-Prefix AVP instead of the complete IPv6 address. The default prefix length is /64, but you can configure the appliance to use a different value.

[From Build 64.34] [# 574135]

- **Idle Session Management of Subscriber Sessions in a Telco Network**

Subscriber-session cleanup on the NetScaler appliance is based on control plane events, such as a RADIUS Accounting Stop message, a Diameter RAR (session release) message, or a "clear subscriber session" command. In some deployments, the messages from a RADIUS client or a PCRF server might not reach the appliance. Additionally, during heavy traffic, the messages might be lost. A subscriber session that is idle for a long time continues to consume memory and IP resources on the appliance. The idle session management feature provides configurable timers to identify idle sessions, and cleans up these sessions on the basis of the specified action.

[From Build 64.34] [# 574138]

- **Deterministic NAT Allocation for DS-Lite**

Deterministic NAT allocation for DS-Lite LSN deployments is a type of NAT resource allocation in which the NetScaler appliance pre-allocates, from the LSN NAT IP pool and on the basis of the specified port block size, an LSN NAT IP address and a block of ports to each subscriber (subscriber behind B4 device).

The appliance sequentially allocates NAT resources to these subscribers. It assigns the first block of ports on the beginning NAT IP address to the beginning subscriber IP address. The next range of ports is assigned to the next subscriber, and so on, until the NAT address does not have enough ports for the next subscriber. At that point, the first port block on the next NAT address is assigned to the subscriber, and so on.

The NetScaler appliance logs the allocated NAT IP address and the port block for a subscriber. For a connection, a subscriber can be identified by just its mapped NAT IP address and port block. For this reason, the NetScaler appliance does not log the creation or deletion of an LSN session.

[From Build 64.34] [# 582325]

- **Port Block Size in a Large Scale NAT Configuration**

Deterministic NAT and Dynamic NAT with port block allocation significantly reduce the LSN log volume. For these two types of configuration, the NetScaler appliance allocates a NAT IP address and a block of ports to a subscriber.

The minimum port block size for deterministic LSN configuration and dynamic LSN configuration with port block has been reduced from 512 ports to 256. This reduction of the minimum port block doubles the maximum number of subscribers for a NAT IP address in an LSN configuration. It also reduces the number of unused ports assigned to subscribers who do not need more than 256 ports at a time.

The port block size parameter can be set while adding or modifying an LSN group as part of an LSN configuration. A value of 256 (default) or a multiple of 256 can be set to the port block size parameter.

[From Build 64.34] [# 581285]

- **Configuring DS-Lite Static LSN Maps**

The NetScaler appliance supports manual creation of DS-Lite LSN mappings, which contain the mapping between the following information:

- * Subscriber's IP address and port, and IPv6 address of B4 device or component

- * NAT IP address and port

Static DS-Lite LSN mappings are useful in cases where you want to ensure that the connections initiated to a NAT IP address and port map to the subscriber IP address and port through the specified B4 device (for example, web servers located in the internal network).

[From Build 64.34] [# 558406]

- **IP Prefix NAT**

The NetScaler appliance supports translating a part of the source IP address instead of the complete address of packets received on the appliance. IP prefix NAT includes changing one or more octets or bits of the source IP address.

The NetScaler appliance supports IP prefix NAT for traffic related to virtual servers and services for which the NetScaler does not maintain any session information. For example, virtual servers and services of type ANY, UDP, and DNS.

IP prefix NAT is useful in a deployment of NetScaler appliances and optimization devices (for example, Citrix ByteMobile) for identifying traffic from different client networks, which share the same network address, for meeting different optimization needs for traffic from each client network.

[From Build 64.34] [# 590571]

- **Logging MSISDN Information for a Large Scale NAT configuration**

A Mobile Station Integrated Subscriber Directory Number (MSISDN) is a telephone number uniquely identifying a subscriber across multiple mobile networks. The MSISDN is associated with a country code and a national destination code identifying the subscriber's operator.

You can configure a NetScaler appliance to include MSISDNs in LSN log entries for subscribers in mobile networks. The presence of MSISDNs in the LSN logs helps the administrator in faster and accurate back tracing of a mobile subscriber who has violated a policy or law, or whose information is required by lawful interception agencies.

[From Build 64.34] [# 581315, 502083]

Fixed Issues in Previous NetScaler 11.0 Releases

The issues that were addressed in NetScaler 11.0 releases prior to Build 60.12. The build number provided below the issue description indicates the build in which this issue was addressed.

AAA-TM

- The "show aaa session" command causes a high level of CPU usage when executed with the "-username" or "-group" option.

[From Build 63.16] [# 577778, 595104, 595185]

- When kerberos token decryption fails, the NetScaler appliance responds with a 200 response with error message, instead of sending a 401 response.

[From Build 64.34] [# 567233, 593994]

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

[From Build 64.34] [# 540396]

- The NetScaler appliance might become unresponsive if the persistence cookie feature is enabled in AAA-TM deployments.

[From Build 64.34] [# 599701, 607138, 608997]

- Single sign-on to server does not succeed when native clients, such as iOS clients, connect to the NetScaler appliance using Active Sync protocol and send cookies along with authorization header.

[From Build 64.34] [# 597221]

- If the AAA virtual server is configured to a non-ActiveDirectory LDAP server, and an invalid password is used to logon, the NetScaler appliance becomes unresponsive.

[From Build 64.34] [# 599264, 610045]

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

- Security type is SSL and port is 389.

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

[From Build 64.34] [# 567376, 567379, 592941]

- Authentication fails if the server name in an LDAP action is changed from an FQDN to an IP address by using the "set ldapaction" command.

[From Build 65.35] [# 614597]

- When the NetScaler appliance is configured as SAML Service Provider (SP), the SAML Identity Provider (IdP) dishonors a logout request that is performed on the traffic management virtual server (load balancing or content switching) that uses a AAA-TM traffic policy.

This happens because the NetScaler SP sends to the SAML IdP a SAML logoutRequest that contains "Conditions" XML tag.

[From Build 65.35] [# 613700]

- 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 65.35] [# 610604, 618760, 623053]

- If a logout message from a session owner to a cached session is dropped, the NetScaler appliance might fail while trying to resend the message.

[From Build 65.35] [# 620948]

- When RADIUS is used in nFactor authentication, the NetScaler appliance fails to complete the request if user is prompted for password change.

[From Build 65.35] [# 612431]

- You cannot enter the FQDN for a RADIUS or LDAP server by using the NetScaler GUI.

[From Build 65.35] [# 596382, 618884]

- If you are using AAA-TM on an HTTP virtual server with no endpoint features enabled, the acknowledgement from the NetScaler appliance might not contain all the data that the client sent. This might cause some page elements to not load completely, or to time out.

[From Build 65.35] [# 615885]

- The NetScaler appliance intermittently fails if a user accesses a very long URL without proper AAA context.

[From Build 65.35] [# 598837]

- You cannot add an authentication virtual server as the target of a content switching virtual server in a partition.

[From Build 66.11] [# 624063]

- If you add a Negotiate Server with a Keytab file with a GUI, an error is issued: "Error in retrieving file. Directory does not exist." The error is only issued when it is executed within partition.

[From Build 66.11] [# 620774]

- When doing forms based SSO, if the back-end server sets a cookie with the login form, NetScaler does not send those cookies to the client. This behavior was observed after a successful forms SSO attempt. This applies to forms based SSO access in AAA-TM products.

[From Build 66.11] [# 624165]

- If AAA-TM logout is configured through a traffic policy on the NetScaler appliance, and the server sends a chunked response, the user encounters an error.

[From Build 66.11] [# 623005]

- If you use the Kerberos protocol for single sign-on (SSO) to access a back-end server, the NetScaler appliance might fail if heavy traffic causes allocation failures, because the appliance might detect a call to free memory that has already been freed.

[From Build 67.12] [# 637125, 649410]

- If SAML authentication is used to log on a user, and the SAML action is removed while there are active sessions, addition of a high availability node might cause occasional failures on the secondary node.

[From Build 67.12] [# 621787]

- The NetScaler appliance fails if authentication is disabled while user authentication is in progress.

[From Build 67.12] [# 617370]

Action Analytics

- A global flag that tracks stream sessions when the ICMP traffic processing begins is not initiated properly.

[From Build 64.34] [# 595915, 602701]

Admin Partitions

- In an admin partition, changes done to enable or disable a NetScaler feature or mode are not saved. Therefore, after the NetScaler appliance is rebooted, the status of the feature or mode is reset to its default value.

[From Build 64.34] [# 594845]

- Setting L2 and L3 parameters in Admin Partitions

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

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

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

[From Build 64.34] [# 513564]

- Partition administrators cannot upload scriptable monitor scripts to a partition. This can only be done by NetScaler superusers. Also, scriptable monitors for an admin partition cannot be configured by using the GUI.

[From Build 64.34] [# 583756]

- When creating an admin partition, you can now set the memory limit to a minimum value of 5 MB.

[From Build 64.34] [# 580419]

- When the NSIP password is changed (by using the "set ns rpcnode" command) on the default partition, the GSLB auto-sync function does not work in the available admin partitions.

[From Build 65.35] [# 621939]

- In a partitioned NetScaler appliance, for GSLB services that point to a local load balancing virtual server, the monitors that are bound to that GSLB service fail. Also, connection proxy between local load balancing virtual servers does not work.

[From Build 66.11] [# 613751]

- When a user is authenticated, the persistence cookie is set, and any subsequent request that includes the cookie and matches the basic pattern is granted access to the server. However, in a non-default partition, if you specify a value for the ns_aaatm_tempc_allow_patterns default pattern set, the value is ignored.

[From Build 67.12] [# 623404]

- If you remove an admin partition, the NetScaler appliance fails or corrupts an SNMPD packet queue.

[From Build 67.12] [# 618251]

AppExpert

- The order in which AppExpert evaluates application units cannot be changed. With this fix, the NetScaler GUI displays a burger icon for each application unit. After hovering over the icon, you can move an application unit up or down in the order of evaluation.

Navigation: Configuration > AppExpert > Application > Application Unit section

[From Build 64.34] [# 567425]

AppFlow

- When routes are updated after an AppFlow collector is added, the NetScaler appliance sends ARP requests for the AppFlow collector IP address, even when the collector is reachable only through a router.

[From Build 63.16] [# 574420]

- The NetScaler appliance might become unresponsive if a request generated by a client is corrupted after execution of the client-side measurement script. This issue can occur if you enable the client side measurement option for an AppFlow action.

[From Build 64.34] [# 601915, 601924, 607217]

- The NetScaler appliance might become unresponsive if you enable the client side measurement option for an AppFlow action.

[From Build 64.34] [# 595238]

- When you try to save or check-in a document on the SharePoint server that has MDS enabled and AppFlow enabled on the NetScaler appliance, you will see the following error message: "An error occurred while processing the request on the server. The status code returned from the server was: 0".

[From Build 67.12] [# 630611]

- 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 68.11] [# 650618, 653126, 657166, 661587]

Application Firewall

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

[From Build 62.10] [# 580564]

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

[From Build 62.10] [# 572637, 581520]

- If you use the default browser PDF plugin to view an application firewall report, embedded links might be inactive.

[From Build 63.16] [# 372768]

- The Citrix application firewall silently resets the connection when it receives a malformed or invalid request. With this fix, the application firewall logs such events.

[From Build 63.16] [# 577742]

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

[From Build 63.16] [# 580564]

- The NetScaler application firewall terminates the connection when the request comes with a tampered session cookie and the cookie protection is enabled.

[From Build 63.16] [# 574498, 591172]

- If learning thresholds for the application firewall security checks are set to a value greater than 1, the configuration utility displays the following error message when you try to access the learned data: "communication error with aslearn."

[From Build 63.16] [# 584621]

- The NetScaler appliance might become unresponsive when processing a request, because of an interoperability issue between the application firewall, SSL, and the responder module. The issue arises under the following set of circumstances:

The configuration includes an application firewall profile protecting an SSL virtual server.

A responder policy is configured to reset the connection, and this policy is bound either globally or to the virtual server that receives the request.

[From Build 63.16] [# 592429]

- The NetScaler appliance might fail when the application firewall is processing the cookie header(s) in an HTTP request. This occurs when the cookie transform action is enabled and all other security checks that apply to establishing a user session are disabled.

[From Build 63.16] [# 591176, 593996, 597440, 601359]

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

[From Build 63.16] [# 547978]

- During an application firewall security check inspection, a compressed response from the server might trigger a violation if the XML format check is enabled. With this fix, the Accept-Encoding request header is removed when the XML protections are enabled. If content compression is enabled on the server, the XML check inspection is bypassed when the server sends a compressed response.

[From Build 63.16] [# 580273]

- NetScaler application firewall resets the connection when the request contains tampered session cookie and the cookie protection is enabled.

[From Build 63.16] [# 591172, 574498]

- When URLTransform or CVPN policies are configured, application firewall code is invoked to carry out the validation of http packet information even if application firewall feature is disabled. When streaming code is engaged, the application firewall is not processing the conditional headers accurately and might reset connection and respond with RST code 9856. With this fix, parsing and validating the request headers is handled correctly by the application firewall module.

[From Build 64.34] [# 593960, 605920]

- The application firewall has extended external format signature support for a new scan tool called WebInspect. The WebInspect scan tool, provided by Hewlett Packard (HP), is designed to analyze the web applications and web services for security vulnerabilities. As stated in the following Data Sheet link from HP, "WebInspect provides the broadest dynamic application security testing coverage and detects new types of vulnerabilities that often go undetected by black-box security testing technologies": <http://www8.hp.com/h20195/V2/GetPDF.aspx/4AA1-5363ENW.pdf>.

[From Build 64.34] [# 588914, 609060]

- The NetScaler appliance might fail when the application firewall receives an HTTP response with an attribute value that exceeds 1 MB in length.

[From Build 64.34] [# 592018]

- When the application firewall redirects a blocked request to a customized error page, the `$(NS_APPFW_SESSION_ID);` variable on the error page might not display the session ID accurately. If the request does not contain a session cookie, the variable might display a hyphen (-) instead of the session ID.

[From Build 64.34] [# 599052]

- If the application firewall cookie proxy check is enabled and the server tries to expire and modify the same cookie in the same response, the NetScaler appliance might fail because of memory corruption.

[From Build 64.34] [# 603694, 609394]

- In a cluster deployment, accessing the application firewall learned data might display "Error in retrieving Application Firewall learning data. Communication error with aslearn". This error is triggered if buffer overflow occurs when the cluster configuration coordinator tries to get learned data from the other nodes of the cluster.

[From Build 64.34] [# 607187]

- The NetScaler appliance might become unresponsive when processing a request, because of an interoperability issue between the application firewall, SSL, and the responder module. The issue arises under the following set of circumstances:

The configuration includes an application firewall profile protecting an SSL virtual server. A responder policy is configured to reset the connection, and this policy is bound either globally or to the virtual server that receives the request.

[From Build 64.34] [# 592429, 612052]

- The NetScaler application firewall terminates the connection when the request comes with a tampered session cookie and the cookie protection is enabled.

[From Build 64.34] [# 574498, 591172]

- Signatures version may not get updated correctly if updated_signatures.xml file is present in /nsconfig folder. With this fix, this file is removed during build installation and the version of the application firewall signatures is updated accurately.

[From Build 64.34] [# 588640]

- The NetScaler appliance might fail when the application firewall is processing the cookie header(s) in an HTTP request. This occurs when the cookie transform action is enabled and all other security checks that apply to establishing a user session are disabled.

[From Build 64.34] [# 591176, 593996, 597440, 601359]

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

[From Build 64.34] [# 561567]

- NetScaler application firewall resets the connection when the request contains tampered session cookie and the cookie protection is enabled.

[From Build 64.34] [# 591172, 574498]

- Application firewall profiles that are exported and archived from one build cannot be restored to a system running a different build, because changes introduced in the newer releases can lead to compatibility issues. With this fix, the application firewall now logs an error message, in ns.log, if you attempt to restore an archived profile to a different build than the one from which it was exported.

[From Build 64.34] [# 601064]

- The application firewall buffers the entire request for security check inspections. Therefore, when the client sends the expect 100-continue header in the request, the application firewall sends the 100-continue response to get the entire request from the client. The application firewall modifies the expect 100-continue header received from the client and corrupts it before forwarding the processed request to the server. In the 11.0 release, the header was not corrupted before it was forwarded request to server. With this fix, the expect 100-continue header from the client is modified and a corrupted header is sent to the server.

[From Build 64.34] [# 598607]

- NetScaler application firewall handles memory incorrectly if XSS and "CrossSiteScriptingCheckCompleteURLs" are enabled in the application firewall profile. The errors also appear if "checkrequestHeaders" and finegrained relaxations are enabled.

[From Build 65.35] [# 606931]

- Application Firewall memory allocation errors might occur if the license on the NetScaler appliance restricts the number of packet engines.

[From Build 65.35] [# 621798]

- In certain cases, if a custom error page containing variables is served to the client, the content length in the response is incorrect. As a result, the custom error page might not be visible in the client's browser.

[From Build 65.35] [# 616947]

- You might encounter unexpected failures if form field consistency protection is enabled on the application firewall profile and you try to retrieve the form from Distributed Hash Table (DHT).

[From Build 65.35] [# 616191]

- If you use the Mozilla Firefox browser to access the NetScaler GUI, you cannot make changes to the application firewall configuration.

[From Build 65.35] [# 619978]

- In NetScaler web application firewall high availability deployments, application firewall sessions are not cleaned up on the secondary node. As a result, memory usage increases on the secondary node.

[From Build 65.35] [# 612284, 619056]

- In a cluster setup, while exporting application firewall learnt data, you might see the following error message:

"communication error with aslearn"

This message is because of a schema difference.

[From Build 65.35] [# 625807]

- The application firewall might experience a transient low-memory condition during a traffic surge if advanced security check protections (such as Form Field consistency, CSRF, form tagging and so on, which require rewriting the HTML forms in the response) are enabled for the profiles. This might result in a memory leak, and memory allocation failures might occur even after the traffic surge subsides.

[From Build 65.35] [# 598776, 597952]

- When the application firewall cookie proxy check is enabled, the NetScaler appliance might become unresponsive while updating the cookies in the distributed hash table with a set of cookies from the server.

[From Build 65.35] [# 609394, 618385]

- In release 10.5.e (enhancement builds only) as well as in the 11.0 release builds, application firewall processing of the Cookie header was changed. In those releases, every cookie is evaluated individually, and if the length of any one cookie received in the Cookie header exceeds the configured BufferOverflowMaxCookieLength, the Buffer Overflow violation is triggered. As a result of this change, requests that were blocked in 10.5 and earlier release builds might be allowed, because the length of the entire cookie header is not calculated for determining the cookie length. In some

situations, the total cookie size forwarded to the server might be larger than the accepted value, and the server might respond with "400 Bad Request".

With this fix, the change has been reverted. The behavior is now similar to that of the non-enhancement builds of release 10.5. The entire raw Cookie header is now considered when calculating the length of the cookie. Surrounding spaces and the semicolon (;) characters separating the name-value pairs are also included in determining the cookie length.

[From Build 65.35] [# 614449]

- When you use the NetScaler GUI to perform the Skip operation, the application firewall learned rules might not be deleted. This occurs because NITRO is sending wrong "Location" ("Field") data to the GUI. With this fix, the GUI converts "Field" into "FORMFIELD," and the Skip operation removes the skipped rules, as expected.

[From Build 65.35] [# 610116, 603473]

- The XSS transform for special characters in the application firewall might not work as expected if the -crossSiteScriptingTransformUnsafeHTML option or the sqlTransformAction option is set to ON in the profile.

[From Build 65.35] [# 618707]

- When you use the NetScaler GUI to perform the Skip operation, the application firewall learned rules might not be deleted. This occurs because NITRO is sending wrong "Location" ("Field") data to the GUI. With this fix, the GUI converts "Field" into "FORMFIELD," and the Skip operation removes the skipped rules, as expected.

[From Build 65.35] [# 603473]

- After an upgrade from release 10.5 to a release 11.0 build, uploading a Word document triggers false positives for application firewall SQL and XSS violations during the file upload operation. With this fix, the behavior is the same as in 10.5. The application firewall inspects text, Javascript , HTML, XML and JSON contents when a file is uploaded. It doesn't inspect any other contents.

[From Build 66.11] [# 619354]

- If a client submits a form that includes a field named "as_fid", and the application-firewall profile has signatures enabled, the signatures might block form submissions from that client.

[From Build 66.11] [# 628525]

- Starturl relaxations might not work if regex expressions use grouping for matching multiple terms. The URL might not get matched against all the terms in the group.

[From Build 66.11] [# 628789]

- The application firewall's SQL Injection special-character transform does not work properly if either of the following parameters is enabled in a profile:

-crossSiteScriptingTransformUnsafeHTML

-SQLInjectionTransformSpecialChars

[From Build 66.11] [# 617614, 624646, 624653]

- In a cluster deployment, the NetScaler web application firewall (WAF) fails if starturl closure is enabled.

[From Build 66.11] [# 617680]

- On a partitioned NetScaler appliance, connections for incoming requests might be reset if the application firewall feature is enabled.

[From Build 66.11] [# 622826]

- The NetScaler appliance fails if you enable or disable the IP Reputation feature in any partition other than the default partition.

[From Build 66.11] [# 627505, 628073]

- If referer header protection is enabled on the application firewall and the starturl action specified is either "stats" or "learn," webpages might not load correctly.

[From Build 66.11] [# 629385, 638602]

- When the application firewall signature has upper case or mixed case characters in the name, the configured profile bindings for such a signature are not displayed in the signatures pane in the configuration utility.

[From Build 66.11] [# 561845, 620915]

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

[From Build 67.12] [# 329025, 303481]

- In a high availability (HA) deployment, a memory leak can occur if auto-update of application firewall signatures is enabled or you update the signatures by using the -mergedefault option.

[From Build 67.12] [# 620878, 629043, 641457]

- The NetScaler appliance might fail when application firewall is attempting to log messages regarding the user's session but the source string is NULL due to memory corruption.

[From Build 67.12] [# 635738]

- 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 67.12] [# 647309]

- The application firewall allows configuring default field format parameters. The valid range for the maximum field format length is 1-65535. The GUI as well as CLI currently accepts zero as input even though zero is outside the allowed range.

[From Build 67.12] [# 608010, 603763, 629859]

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

[From Build 68.10] [# 643854]

- In a high availability (HA) deployment, a memory leak can occur if auto-update of application firewall signatures is enabled or you update the signatures by using the Merge Default (-mergedefault) option.

[From Build 68.10] [# 620878, 629043, 641457, 649075]

- The name of a user defined signature object must not contain a hash-mark character (#), but the feedback message lists it as an allowed character.

[From Build 68.10] [# 648010]

- 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 68.10] [# 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. This could result in a violation.

[From Build 68.10] [# 648104]

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

[From Build 68.10] [# 329025, 303481]

- If the HTML response page contains a pair of hyphens (--) in the comment tag, the NetScaler appliance might parse the response page incorrectly. This could result in a violation.

[From Build 68.11] [# 648104]

- 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 68.11] [# 649031, 651536]

- In a high availability (HA) deployment, a memory leak can occur if auto-update of application firewall signatures is enabled or you update the signatures by using the Merge Default (-mergedefault) option.

[From Build 68.11] [# 620878, 629043, 641457, 649075]

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

[From Build 68.11] [# 643854]

- The name of a user defined signature object must not contain a hash-mark character (#), but the feedback message lists it as an allowed character.

[From Build 68.11] [# 648010]

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

[From Build 68.11] [# 329025, 303481]

Audit logging

- You can now customize the log levels for logs generated for AAATM user logon or logoff, and for logs generated for executive commands by a NetScaler administrative user.

[From Build 65.35] [# 386650]

Cache Redirection

- In the GUI, the Policy drop-down list does not display the cache redirection policies.

[From Build 65.35] [# 622402]

- If a request to a cache redirection virtual server resolves to an IP address that belongs to a content switching virtual server configured on the NetScaler appliance, the appliance might fail.

[From Build 65.35] [# 621522]

- Classic cache redirection policies send CONNECT requests to the cache, as expected, if they do not match the policy rule, but default syntax cache redirection policies send them to the origin server instead. With this fix, default syntax cache redirection policies send nonmatching requests to the cache.

[From Build 68.10] [# 637826]

- Classic cache redirection policies send CONNECT requests to the cache, as expected, if they do not match the policy rule, but default syntax cache redirection policies send them to the origin server instead. With this fix, default syntax cache redirection policies send nonmatching requests to the cache.

[From Build 68.11] [# 637826]

Cisco RISE Integration

- If RISE feature is not enabled and we try to disable it, an error message is displayed for all the features.

[From Build 65.35] [# 513761]

Cluster

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

[From Build 62.10] [# 574717]

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

[From Build 62.10] [# 574333, 573645]

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

[From Build 62.10] [# 559230, 571042]

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

[From Build 63.16] [# 551607, 495270, 562651]

- A NetScaler cluster does not respond to cURL HTTP requests from outside the datacenter, because the Path MTU Discovery (PMTUD) mode gets disabled when a cluster is created.

[From Build 64.34] [# 541223]

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

[From Build 64.34] [# 584699]

- The VRRP Feature does not work in a cluster setup that includes a node with a node ID of zero (0).

[From Build 66.11] [# 618663]

- For some commands, such as "add cs policy" and "add server," the ID generated on a non-CCO node already exists for another command of same type on the cluster configuration coordinator (CCO). Therefore, command execution on the non-CCO node fails.

[From Build 66.11] [# 614718, 615459]

Clustering

- A force cluster sync operation causes the cluster's static ARP configuration to become inconsistent.

[From Build 68.10] [# 635231]

- A force cluster sync operation causes the cluster's static ARP configuration to become inconsistent.

[From Build 68.11] [# 635231]

Command Line Interface

- A customized CLI prompt is not persisted after rebooting the appliance.

[From Build 63.16] [# 583625]

- The NetScaler CLI exhibits the following issues on running the "show" and "stat" commands on a service group.

- When using the "show servicegroup -includeMembers" command: This command lists only one service per service group, although more than 1 service are bound to the service group(s).

- When using the "stat servicegroupMember <ServiceGroupName> <Service-IP-address> <port>" command: This command does not work if you specify the <Service-IP-address>. Instead, you must specify the <Service-Name>.

[From Build 66.11] [# 554652, 596571]

Configuration Utility

- You cannot configure the service path AVP by using the configuration utility.

[From Build 62.10] [# 576603]

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

[From Build 62.10] [# 571814, 581955]

- You cannot configure the service path AVP by using the configuration utility.

[From Build 63.16] [# 576603]

- The values for the parameters on the "Configure Load Balancing Parameters" page do not appear even though they have been set.

[From Build 63.16] [# 583741]

- SUBSCRIBER expressions do not appear in the list for rewrite and responder policies and action.

[From Build 63.16] [# 583751]

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

[From Build 63.16] [# 571814, 581955]

- When starting a nstrace and another instance is already running, an option to stop this is not available in the configuration utility. One has to login through the command line interface to stop the trace.

[From Build 64.34] [# 603476]

- If you are using the configuration utility to run diagnostics on the NetScaler appliance, you cannot specify a traffic domain.

[From Build 64.34] [# 609334]

- The configuration utility does not reflect the correct count of cached objects whereas this number is shown correctly through the command line interface.

[From Build 64.34] [# 607622, 608517]

- When an HTML page is imported, the content is copied to /nsconfig/ssl and then to /var/download/responder. The content is not removed from /nsconfig/ssl, although it serves no purpose there. With this fix, the content is copied directly to /var/download/responder.

[From Build 64.34] [# 590268]

- The integrated caching feature is not available on the GUI.

[From Build 64.34] [# 601429]

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

[From Build 64.34] [# 540943]

- If you create a cipher group and do not add any ciphers to it, an error message appears when you try to open the cipher group in the configuration utility.

[From Build 64.34] [# 604646]

- You cannot add user-defined values for the user name and group name fields on the Authentication CERT Profile page.

With this fix, you can specify a user-defined value by navigating to Security > AAA - Application Traffic > Policies > Authentication > Basic Policies > CERT > Profiles and selecting New in the User Name Field list and the Group Name Field list.

[From Build 64.34] [# 597708]

- If you log on to the appliance by using the GUI, the list of licenses is not retrieved.

[From Build 65.35] [# 611772]

Content Switching

- If a large number of content switching policies are bound to a content switching virtual server, using the configuration utility to bind a new policy without explicitly assigning a priority might result in the policy being assigned the priority of the first policy on the next page of the display. Since a policy is already assigned that priority, an error message stating that the priority is already used appears.

[From Build 64.34] [# 601203]

- In certain cases, if the state of a load balancing virtual server changes, the NetScaler appliance might fail while changing the state of the associated content switching virtual server.

[From Build 65.35] [# 522510, 528782, 538223, 552913, 602829]

DNS

- The query logs contain incorrect information if the UDP payload size in the OPT record is not 1280. Also, if a load balancing virtual server on the NetScaler appliance receives a request with the CD bit set, and the "RecursionAvailable" parameter is disabled on the DNS or DNS-TCP load balancing virtual server, the CD bit is not logged.

[From Build 63.16] [# 579942]

- The NetScaler appliance fails, if there is a cache miss when the backend DNS server is accessed directly through the NetScaler appliance.

[From Build 64.34] [# 609074]

- If, while a DNS-TCP client request is in surge queue, the NetScaler appliance receives a FIN from the client and responds with a FIN or ACK before the queued request is forwarded to the backend server, the appliance might fail.

[From Build 64.34] [# 581723]

- If a NetScaler appliance in DNS resolver mode is configured to resolve queries with suffixes, the appliance fails if there is no address record for the NS record associated with one of the suffixes.

[From Build 65.35] [# 605861]

- Due to a memory overwrite issue, the prev value of dns_tot_ServerQueries counter is set to zero everytime during the end of perf collection cycle, that is every 7 seconds. This results in the difference between cur and prev value get accumulated to the global counterpart even if there is no traffic.

[From Build 65.35] [# 615519, 580342]

- If, while resolving a domain name in DNS resolver mode, the NetScaler appliance does not receive a response from the first name server, it tries to resolve the domain name with the other name servers. During this process, if the address record for the associated NS record is not present, the NetScaler appliance fails.

[From Build 65.35] [# 609967, 617204]

- Non-standard query packets are altered before they are forwarded to back-end servers, which causes the server to respond with a "FORMAT error" message.

[From Build 66.11] [# 559064]

- In a deployment with heavy DNS traffic and many DNS cache entries, some entries in the cache might not get updated or deleted, even after the TTL expires.

[From Build 67.12] [# 619124, 622308]

- 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 68.10] [# 648087]

- 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 68.11] [# 648087]

DataStream

- A NetScaler client becomes unresponsive if:
 1. The NetScaler appliance receives the complete response to the client's query from the server.
 2. At the same time, the client sends an attention packet to the appliance.

The client becomes unresponsive because the appliance closes the server-side connection but does not send the client a response to the attention packet.

[From Build 65.35] [# 560401]

- If the NetScaler appliance receives a prelogin message request from a Visual Studio 2015 client, it sends an incorrect response. As a result, the client becomes unresponsive.

[From Build 66.11] [# 613239, 616404]

GSLB

- If you have configured the canonical name as the GSLB domain in NetScaler appliance, when the backend server returns the CNAME record without the requested record, NetScaler appliance changes the TTL value of the GSLB domain with the TTL value of the CNAME record.

[From Build 63.16] [# 582925]

- GSLB virtual server configured with Dynamic Proximity as LB method fails.

[From Build 63.16] [# 578969]

- GSLB Service Selection using Content Switching

Description: You can now configure a content switching (CS) policy to customize a GSLB deployment so that you can:

* Restrict the selection of a GSLB service to a subset of GSLB services bound to a GSLB virtual server for the given domain.

* Apply different Load Balancing methods on the different subsets of GSLB services in the deployment.

* Apply spillover policies on a subset of GSLB services, and you can have a backup for a subset of GSLB services.

* Configure a subset of GSLB services to serve a specific type of content.

* Define a subset GSLB services with different priorities, and define the order in which the services in the subset are applied to a request.

For more information, see [Configuring GSLB Service Selection Using Content Switching](#).

[From Build 64.34] [# 503588]

- Initiating 280k SIP sessions with 40k subscribers might cause the NetScaler appliance to fail.

[From Build 64.34] [# 582459, 591247]

- If a server entity (for example, a server IP address or server name) is associated with both a GSLB entity and a non-GSLB entity on a GSLB site, and the GSLB configuration is synced to another site that does not include this server entity, the synchronization removes the server entity and all other entities associated with that server.

[From Build 64.34] [# 590336]

- In a content switching GSLB deployment, you can bind multiple domains to a CS GSLB virtual server, but the show cs vsrver command shows only one domain bound to the CS GSLB virtual server.

[From Build 65.35] [# 612916]

- If the ACK on PUSH option is disabled in the default TCP profile, the NetScaler appliance might fail while downloading the static proximity database.

[From Build 65.35] [# 582102]

- In the GUI, on the GSLB statistics page, the local site MEP state is always displayed as DOWN instead of as a blank field.

[From Build 65.35] [# 617267]

- For GSLB deployments in a partitioned environment, the options to synchronize the GSLB configuration and view the synchronization status are provided in the GUI.

[From Build 65.35] [# 622147]

- The NetScaler appliance fails if you run the "show gslb domain" command on a non-gslb domain record.

[From Build 65.35] [# 618789]

- When using the GUI in a partitioned environment, you cannot add GSLB services.

[From Build 65.35] [# 622131]

- In a GSLB deployment, if monitors are bound to GSLB services and the trigger monitor is set to MEP_DOWN. The remote GSLB services are incorrectly marked as down when MEP goes down due to temporary network outage but the MEP connection is still active.

[From Build 66.11] [# 610065]

- When the MEP connection between two GSLB sites is reestablished after going down, the connection becomes active immediately, but the NetScaler GUI and CLI do not show it as UP for about 9 seconds.

[From Build 66.11] [# 615886]

- After being upgraded to NetScaler release 11.0, the appliance cannot access the geoIP location database.

[From Build 67.12] [# 639371]

- If the GSLB virtual server is of type RADIUS and if the rewrite or responder feature is disabled, the NetScaler appliance fails when it gets a DNS request for a GSLB domain bound to the RADIUS GSLB virtual server.

[From Build 67.12] [# 642791]

High Availability

- The HA traffic between the HA pair is abnormally high. This issue is caused by a loop that repeatedly tries to push the same sessions to the secondary appliance after failover.

[From Build 63.16] [# 560640, 566710, 576012, 576096, 579037, 582354, 590730]

- When there is a HA issue, the synchronization of persistence sessions between the primary and secondary appliances can fail. This can cause some of the persistence sessions not being replicated on the secondary appliance.

[From Build 63.16] [# 580703, 579037, 595491, 595506, 596002, 596215, 599250, 599396, 604164, 605112, 608450, 608485]

- When there is a HA issue, the synchronization of persistence sessions between the primary and secondary appliances can fail. This can cause some of the persistence sessions not being replicated on the secondary appliance.

[From Build 64.34] [# 580703, 579037, 595104, 595491, 595506, 596002, 596215, 599250, 599396, 604164, 605112, 608450, 608485, 610589]

Insight Center

- If integrated cache is enabled on a NetScaler appliance running software version 11.0-65.x or earlier, web insight reports are not available in the Insight Center.

[From Build 66.11] [# 631677]

Integrated Cache

- A VPX system can repeatedly fail if HA cache persistence is used along with HTML-injection.

[From Build 65.35] [# 581598]

- Disabling the Media Classification mode, even if the host header is missing in the GET request, does not cause a NetScaler appliance to fail.

[From Build 65.35] [# 616021, 616757, 618970, 624338]

- The NetScaler can stop responding when cache object persistency is configured in a HA setup.

[From Build 65.35] [# 589322]

- A NetScaler appliance performing integrated caching becomes unresponsive if the length of the URL is 2040 (including the hostname, query parameter, and other specific information).

[From Build 65.35] [# 605831, 612030, 612102]

- When a NetScaler appliance uses a flash cache with HTTPS traffic, only the initial client request is serviced. Subsequent client requests fail.

[From Build 65.35] [# 602984]

- When the "clear config" command is issued, the NetScaler appliance can become unresponsive if more than one CPU tries to free the same shared memory.

[From Build 65.35] [# 609928]

Integrated Caching

- If you set the PINNED option for a cache content group, caching continues in this group even if the group uses more than its allocated memory, until the integrated caching memory is exhausted. Because cached objects in these groups

cannot be removed until the appliance is restarted, there might be a situation in which no more objects can be cached and the appliance resets the connections of clients who send additional requests.

[From Build 66.11] [# 621356, 631356]

Load Balancing

- When editing a service group in the configuration utility, the cacheable option is automatically set to true.

[From Build 63.16] [# 592235]

- In a load balancing group configuration, the "sh run" command sometimes runs in a loop, which exponentially increases the size of the temporary configuration file. As a result, saving the configuration and synchronizing the nodes in a high availability setup might fail.

[From Build 63.16] [# 587812, 598499, 601918]

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

[From Build 63.16] [# 583647]

- If an SSL monitor is bound to a domain-based service that is configured with non-default SSL settings, the monitor might not show the service as UP.

[From Build 63.16] [# 575171, 576012]

- In a link load balancing (LLB) deployment, if persistence is enabled on a NetScaler appliance and a policy based routing (PBR) or LB route is configured, the appliance might fail intermittently.

[From Build 64.34] [# 574137]

- In a high availability setup, if a large number of services and service groups are configured, service state updates might fail because of a timer issue.

[From Build 64.34] [# 605596, 609999]

- In a link load balancing (LLB) deployment, if persistence is enabled on a NetScaler appliance and a policy based routing (PBR) or LB route is configured, the appliance might fail intermittently.

[From Build 64.34] [# 554841]

- A secure StoreFront monitor intermittently fails to send probes.

[From Build 64.34] [# 559164, 582153]

- In certain cases, if the name of an FTP virtual server is greater than 32 characters, the virtual server lookup fails and the request is not served.

[From Build 64.34] [# 566644]

- If the channel between the primary node and the secondary node is disrupted, the session deletion information sent from the primary node to the secondary node might get lost. As a result, while the persistent sessions are reduced to zero on the primary node, the secondary node reaches its limit.

[From Build 64.34] [# 596524, 597295]

- While probing the back-end HTTP server by using an HTTP monitor, the appliance does not send the port number in the HTTP host header. This behavior is not compliant with RFC 2616.

[From Build 64.34] [# 564295]

- The NetScaler appliance fails because of an incorrect initialization of template size in a stream analytics session info record.

[From Build 65.35] [# 598391]

- In a high availability setup in admin partition mode, the persistent sessions are not synchronized to the secondary node after performing force ha sync or force failover operation.

[From Build 65.35] [# 630344]

- The NetScaler appliance fails while trying to load balance a request that was received on a recently closed connection. This happens because the server tries to keep the connection alive by sending an RTSP request but the appliance cannot find the corresponding client side connection.

[From Build 65.35] [# 612943]

- In a high availability setup, if a large number of services and service groups are configured, service state updates might fail because of a timer issue.

[From Build 65.35] [# 605596, 609999]

- In rare cases, during a high level of CPU usage, if you disable and enable a service with zero delay, the state of the service might be inconsistent on different packet engines.

[From Build 67.12] [# 622807]

- 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 68.10] [# 648337]

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

[From Build 68.10] [# 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. When service providers do not send IDs, we used to fail due to a logic error.

[From Build 68.10] [# 648489]

- 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. When service providers do not send IDs, we used to fail due to a logic error.

[From Build 68.11] [# 648489]

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

[From Build 68.11] [# 638148]

- 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 68.11] [# 648337]

NITRO API

- For the .NET SDK, when "nitro.dll" is used along with a version later than 4.0 of the "Newtonsoft.json.dll" file, "private" properties cannot be serialized.

[From Build 63.16] [# 567162, 571309]

- The TCP connection is not persistent for NITRO requests. Therefore, the underlying TCP connection is getting closed for each NITRO request.

[From Build 63.16] [# 583395, 457969]

- The NetScaler appliance might become unresponsive when a NITRO request is fetching a large number of bound entities.

[From Build 64.34] [# 530805, 562748, 567856]

- If the NetScaler appliance receives a logon request that contains both the session token and the request payload with the logon credentials, the appliance creates a new connection without closing the previous connection. If the appliance receives multiple such requests, the following error message appears: CFE limit exceeded.

[From Build 65.35] [# 620458, 619154, 621601]

- If a large number of concurrent NITRO requests are issued, many requests time out.

[From Build 66.11] [# 616433]

NetScaler GUI

- The Goto Priority Expression field is missing for a Traffic Management session policy bound to a AAA user or group under NetScaler > Security > AAA > Policies > Session.

[From Build 66.11] [# 629828]

- The details of a custom monitor bound to a service group are not displayed correctly in the NetScaler GUI. The details appear correctly in the CLI.

[From Build 67.12] [# 640332]

- If you specify an invoke label, such as policy label or virtual server, in your policy you cannot view the details of the invoke label directly.

[From Build 67.12] [# 635940]

- The strongpassword command might fail if a strong password validation is done when the password length is Zero. And when the strongpassword command fails, there is a configuration loss.

[From Build 67.12] [# 638190, 647387]

- 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 67.12] [# 614285, 620953]

- If you disable the GSLB feature on a NetScaler appliance, the load balancing feature is automatically disabled.

[From Build 67.12] [# 643191]

NetScaler Insight Center

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

[From Build 62.10] [# 576471]

- Media Classification Support for Insight Center

Web Insight supports content and media type classification reports. Viewing these features are optional similar to the existing HTTP header fields User Agents, Operating Systems, Request Methods etc. You can enable or disable these features from the Configuration section. For media classification and httpContentType Appflow parameter, you must first enable AppFlow on virtual server from Insight center configuration.

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

- 1) Uncategorized
- 2) FLV F4V Audio
- 3) FLV F4V Video
- 4) MP4 M4V Audio
- 5) MP4 M4V Video
- 6) GP 3G2 Video
- 7) ADTS Audio
- 8) APPLE Video
- 9) MICROSOFT Video
- 10) AAC Audio
- 11) MICROSOFT PLAYLIST Video
- 12) APPLE PLAYLIST Video
- 13) MP3 Audio
- 14) Unknown

[From Build 62.10] [# 558890]

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

[From Build 62.10] [# 576477, 581927]

- The NetScaler Insight Center appliance throws an error when modifying the name of a threshold record. To fix this issue, the name field has been made read-only.

[From Build 62.10] [# 573550]

- An exported report displays the time duration as "custom" irrespective of the time duration selected in the report.

[From Build 62.10] [# 577426]

- The SNMP daemon runs on NetScaler Insight Center even though NetScaler Insight Center does not support SNMP requests.

[From Build 63.16] [# 537253]

- The NTP server configuration on NetScaler Insight Center is not propagated to the connector, agent, and database nodes.

[From Build 63.16] [# 579777]

- When you configure authentication (Configuration > System > Authentication > Authentication Configuration), the Server Name field does not display the selected server.

[From Build 64.34] [# 599322]

- If a failover occurs in a high availability configuration, an ICA connection that uses Automatic Client Reconnect (ACR) might fail to reconnect.

[From Build 64.34] [# 601318, 603208]

- If you click on a country in the Geo Maps in the XenDesktop Director GUI, the GEO maps are not displayed.

[From Build 65.35] [# 617872]

- For some elements on the dashboard, NetScaler Insight Center does not fetch records for the specified time frame.

[From Build 65.35] [# 611532, 612283]

- NetScaler Insight Center does not cache reports after you enable database caching.

[From Build 65.35] [# 611269]

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

[From Build 65.35] [# 576469]

- The network panel in the XenDesktop Director GUI does not display a graph with the session details when you select another user.

[From Build 65.35] [# 550209]

- Adding a private IP block in NetScaler Insight Center fails if you select a country name that has special characters.

[From Build 65.35] [# 609646, 620408]

- The network panel in the XenDesktop Director GUI does not display a graph with the session details for the selected user.

[From Build 65.35] [# 550227]

- NetScaler Insight Center fails to generate the technical support file, because the namedpipe file causes an error in the creation of the technical support file.

[From Build 65.35] [# 613622]

- The whitelist of Citrix Receiver versions used by HDX Insight now includes version 13.0.3.265571 of Citrix Receiver for Linux.

[From Build 65.35] [# 614558, 606817]

- The network panel in the XenDesktop Director GUI displays session details of all users, instead of for just the selected user.

[From Build 65.35] [# 607332]

- The NetScaler appliance might sometimes become unresponsive or experience intermittent HA failovers based on a particular ICA network condition.

[From Build 65.35] [# 623729, 623379]

- In NetScaler Insight Center, updating fields in a private IP block fails.

[From Build 65.35] [# 623022]

- NetScaler Insight Center might intermittently become unresponsive and not populate any reports.

[From Build 65.35] [# 618370, 622539]

- The network panel in the XenDesktop Director GUI displays the details of all of the administrative user's sessions, instead of just the details for the selected session.

[From Build 65.35] [# 594512]

- NetScaler Insight Center is now supported on KVM hypervisors.

[From Build 65.35] [# 631295]

- The NetScaler appliance might become unresponsive if you attempt to delete an AppFlow action while the traffic is flowing.

[From Build 65.35] [# 585914, 613238]

- If the two NetScaler appliances in a double hop deployment are running different NetScaler software editions (Platinum, Enterprise, or Standard), NetScaler Insight Center fails to generate reports for these appliances on the NetScaler Insight Center dashboard.

[From Build 65.35] [# 609452]

- In the network panel in the XenDesktop Director GUI, the time slider for selecting the time period for a graph is not properly displayed.

[From Build 65.35] [# 593699]

- NetScaler Insight Center restarts intermittently, and HDX insight reports might not show any data.

[From Build 65.35] [# 606455]

- In Security Insight, there might be a delay in receiving the safety profile configuration data for some applications.

[From Build 66.11] [# 628733]

- In HDX-Insight, the location of public IP addresses is not displayed on the geo map.

[From Build 66.11] [# 631633]

- Now, you can enable or disable SLA Traffic reports from NetScaler Insight Center.

To perform this action, navigate to Configuration > System > System Settings > Configure Data Record Settings, and select the Enable SLA Data Collection check box.

[From Build 67.12] [# 639810]

- The size of the graphs displayed by NetScaler Insight Center is not consistent.

[From Build 67.12] [# 583382]

- If advanced encryption is enabled for ICA in HDX Insight, the following sequence of events might cause the NetScaler appliance to become unresponsive:

1. A server sends improper public key information during shared-key negotiation.
2. Either the client or the server disconnects.

[From Build 67.12] [# 629798, 634013]

- If you enable AppFlow on Load Balancing virtual servers of TCP type, the NetScaler Insight center might become unresponsive.

[From Build 67.12] [# 645281, 645291]

- If you are concurrently using supported and unsupported versions of Citrix Receivers, the NetScaler Insight Center might become unresponsive.

[From Build 67.12] [# 641992]

NetScaler VPX Appliance

- 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 68.10] [# 651670]

- 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 68.11] [# 651670]

NetScaler VPX on Azure

- On a NetScaler instance deployed on Azure, dummy policies, actions, and load balancing virtual server configurations on an instance are lost if you restart the instance.

[From Build 67.12] [# 624685]

Networking

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

[From Build 62.10] [# 573935]

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

[From Build 62.10] [# 575906]

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

[From Build 62.10] [# 573516]

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

[From Build 62.10] [# 560243]

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

[From Build 62.10] [# 506332]

- The output of the show ACL does not display the correct hits for ICMP packets that match the ACL rules.
[From Build 63.16] [# 585265]
- You cannot configure INAT46, INAT64, or INAT66 rules by using the configuration utility.
[From Build 63.16] [# 582682]
- In a high availability configuration, when the connection between primary and secondary goes down and comes up again, the secondary node receives HA INIT request from the primary node and it terminates all BGP connections.
[From Build 63.16] [# 588509]
- The NetScaler appliance might assign the NTP module a port that is used by some other feature module. Therefore, an incoming NTP response can be processed by the feature module. This can result in the failure of the NetScaler appliance.
[From Build 63.16] [# 588477]
- The NetScaler appliance might erroneously forward DHCP broadcast packets to the default router. As a result, the broadcast packets go in loops between the appliance and the router.
[From Build 64.34] [# 591657, 595649]
- In an IPSec tunnel, the NetScaler appliance might remove sessions between client and server before encrypting (IPSec) DNS response packets, resulting in the loss of these DNS packets in the tunnel.
[From Build 64.34] [# 587718]
- If a connection matches a RNAT rule, the NetScaler appliance probes for the existence of the destination server before processing the connection based on the RNAT rule. The connection that is used for probing is sometimes left idle on the appliance and a new connection is opened once the client connection is successfully established. This probe connection stays idle for the configured idle timeout (2.5 hours) thus holding up resources on the server.
Now, these probe connections are flushed within a minute if they remain idle.
[From Build 64.34] [# 588694, 588551]
- The configuration utility does not display any route monitors configured on the NetScaler appliance.
[From Build 64.34] [# 589128]
- For extended ACL rules that are associated in NAT configurations (for example, RNAT rules, Large Scale NAT configurations), the configuration utility displays the TCP established parameter as enabled for these ACL rules.
[From Build 64.34] [# 597458]

- On a NetScaler appliance with a forwarding session rule configured and connection failover enabled, the appliance might become unresponsive when processing packets that match the forwarding session rule.

[From Build 64.34] [# 587382, 603629]

- On a NetScaler appliance with a NetScaler owned IP address configured with a VMAC address on a traffic domain, when a peer device sends an ARP request with unicast MAC for this IP address, the NetScaler appliance responds with the physical MAC address instead of the VMAC address. As a result, the NetScaler appliance drops packets forwarded by the peer device if the packets are destined to the physical MAC address for that IP address.

[From Build 64.34] [# 588912]

- A NetScaler appliance might consume a high percentage of CPU cycles, because the appliance repeatedly updates the active connections with changes in MAC addresses of servers.

[From Build 64.34] [# 579099]

- In a high availability configuration, when the connection between primary and secondary goes down and comes up again, the secondary node receives HA INIT request from the primary node and it terminates all BGP connections.

[From Build 64.34] [# 588509]

- Binding a redundant interface set (for example, LR/1) to NSVLAN might cause the NetScaler appliance to become unresponsive.

[From Build 64.34] [# 597071]

- You cannot securely access (HTTPS) the NetScaler GUI by using a subnet IP (SNIP) address that is configured on a traffic domain.

[From Build 64.34] [# 600364]

- When you remove an admin partition, the NetScaler appliance fails or corrupts an SNMPD packet queue.

[From Build 65.35] [# 613457, 614545, 617179, 621236]

- Forwarding sessions do not work as expected with bridge groups, because packets are not forwarded to the correct VLAN.

[From Build 65.35] [# 600012]

- For extended ACL rules that are associated in NAT configurations (for example, RNAT rules, Large Scale NAT configurations), the configuration utility displays the TCP established parameter as enabled for these ACL rules.

[From Build 65.35] [# 597458]

- The NetScaler appliance fails when it processes invalid UDP packets received at port 500 or port 4500.

[From Build 65.35] [# 609537, 489498]

- If an IPv6 virtual server with persistency enabled is removed from a traffic domain, the traffic domain information for the existing persistency sessions is lost, and the NetScaler appliance hosting the virtual server becomes unresponsive.

[From Build 65.35] [# 608558]

- The dynamic routing module on a NetScaler appliance might incorrectly save the command "redistribute intranet" as "redistribute trill" in the ZebOS configuration file. Because the appliance does not support the "redistribute trill" command, after a failover in a high availability setup, the new primary node treats the "redistribute trill" command as an error and does not apply the subsequent commands in the ZebOS configuration file. This results in loss of configuration.

[From Build 65.35] [# 620152]

- For backend TCP connections, a NetScaler appliance might allocate the subnet IP address and port of an active connection to a new connection. As a result, the new TCP connection fails.

[From Build 65.35] [# 613454]

- On a NetScaler appliance, connections might get reset between routing processes. As a result, the dynamic routes are occasionally deleted and added back.

[From Build 65.35] [# 599306]

- For a sessionless virtual server configuration, the NetScaler appliance might forward packets for an incoming connection without changing their source MAC address with the MAC address of one of its interfaces. As a result, the connection fails.

[From Build 65.35] [# 603477, 583499]

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

[From Build 65.35] [# 574734]

- FTP in passive mode does not work in this build.

[From Build 65.35] [# 631929]

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

[From Build 65.35] [# 572525, 574734, 614924]

- In a GSLB deployment of NetScaler appliances configured with OSPF routing protocol, the OSPF process running in one of the NetScaler appliances sources OSPF hello packets from the GSLB site IP address configured on the appliance. As a result, neighbor adjacency does not get established.

[From Build 65.35] [# 612419]

- The NetScaler appliance does not retain the entire 64 bit ID of IPv6 fragments of a session. As a result, the session might fail.

[From Build 66.11] [# 614042]

- The NetScaler appliance might fail if secure management access (HTTPS) is enabled on a SNIP6 address that is configured for a traffic domain.

[From Build 66.11] [# 618633]

- The NetScaler appliance sends GARP request for a non-addressable virtual server when the virtual server's state changes to UP or DOWN.

[From Build 66.11] [# 620697]

- In an active-active deployment using VRRP, a NetScaler appliance does not match its configured bridge ACL rules to the packets received from the inactive VIP addresses of the other NetScaler appliances.

[From Build 66.11] [# 614786]

- In a high availability setup, secondary node advertises default routes even after performing "ns block-sec-rtadv" operation in VTYSH shell.

[From Build 67.12] [# 639541]

- The NetScaler appliance does not return the value of SNMP OID "lldpRemPortId" for SNMP walk requests.

[From Build 67.12] [# 639630]

- An active FTP connection might get reset for no apparent reason, regardless of the state of the random source port.

[From Build 67.12] [# 507908, 609496, 611357, 615638]

- The existing static ARP entries might point to incorrect MAC addresses after you perform any bind or unbind, or remove VLAN operation.

[From Build 67.12] [# 635823]

- The NetScaler appliance might become unresponsive when one or both of the following conditions are met:
 - When you remove a traffic domain, which has ACLs, or ACL6s, or PBRs, or PBR6s rules, without performing apply operation for ACLs, or ACL6s, or PBRs, or PBR6s rules.
 - When you remove any ACL, or ACL6, or PBR, or PBR6 rule within a traffic domain and then remove the traffic domain before performing apply operation for ACLs, or ACL6s, or PBRs, or PBR6s rules.

[From Build 67.12] [# 636269]

- The NetScaler sends ARP requests for NSIP address on all the interfaces instead of only on the interfaces bound to the NSVLAN.

[From Build 67.12] [# 639356]

- When all ports of all IP addresses bound to a netprofile are used in different backend connections, the NetScaler appliance uses one of the SNIP addresses, which is not bound to the netprofile, for a new backend connection. Backend systems reject the connection if the SNIP address is listed in their deny ACL rules. Now, the NetScaler appliance does not initiate any new backend connection when all the ports of all IP addresses in a netprofile are being used.

[From Build 67.12] [# 627547]

- In a dual-stack lite deployment, the NetScaler appliance does not fragment the received packets if they are too large and the DF bit is not set on them.

[From Build 67.12] [# 631894]

- 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 68.10] [# 642375]

- In a high availability (HA) setup, high latency might occur during configuration synchronization, resulting in some configurations not getting synchronized to the secondary node. In this situation, an HA failover results in loss of configuration.

[From Build 68.10] [# 607929]

- 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 68.10] [# 652131]

- For SIP and RTSP Application Layer Gateways (ALGs) to work properly for a large Scale NAT (LSN) configuration on a NetScaler appliance, it is mandatory to configure all ports of the NAT IP address for FULL cone NAT. That is, Endpoint-Independent Mapping (EIM) and Endpoint Independent Filtering (EIF) must be enabled on these ports, even though they are not used by SIP and RTSP traffic.

[From Build 68.10] [# 641719]

- In a high availability (HA) setup, high latency might occur during configuration synchronization, resulting in some configurations not getting synchronized to the secondary node. In this situation, an HA failover results in loss of configuration.

[From Build 68.11] [# 607929]

- For SIP and RTSP Application Layer Gateways (ALGs) to work properly for a large Scale NAT (LSN) configuration on a NetScaler appliance, it is mandatory to configure all ports of the NAT IP address for FULL cone NAT. That is, Endpoint-Independent Mapping (EIM) and Endpoint Independent Filtering (EIF) must be enabled on these ports, even though they are not used by SIP and RTSP traffic.

[From Build 68.11] [# 641719]

- 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 68.11] [# 642375, 658619]

- 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 68.11] [# 652131]

Optimization

- Enabling the media classification feature causes the NetScaler appliance to become unresponsive.

[From Build 64.34] [# 581123, 584501, 588400, 590438, 594672, 595638, 601727, 601862, 603667, 604126, 607439, 609907, 611899]

- A NetScaler appliance crashes when Media classification mode is enabled and HTTP request of bigger URLs are received.

[From Build 64.34] [# 589825, 594694, 606589, 607919]

- The memory usage statistic shown on the LCD display of a NetScaler appliance is the allocated memory. The NetScaler configuration utility displays the currently used memory. Therefore, the two values are different.

[From Build 63.16] [# 334358, 576545]

- **Support for New Hardware Platforms**

The T1120 and T1300-40G platforms with NIC firmware 4.53 are now supported in this release.

Note: T1300-40G platform with NIC firmware 4.26 has backward compatibility.

[From Build 66.11] [# 593888]

Policies

- The NetScaler appliance fails to respond when a blocking log action is configured with a responder action.

[From Build 62.10] [# 574458, 574593]

- Some IP based expressions might not work for IP addresses starting from octet 128 or greater (128.x.x.x - 254.x.x.x).

The following expressions are not impacted:

- EQ, IN_SUBNET, IS_IPV6, GET1, GET2, GET3, GET4, MATCHES, MATCHES_LOCATION, APPEND, TYPECAST_TEXT_T, TYPECAST_IPv6_ADDRESS_AT

The following expressions do not work:

GT, GE, LT, LE, BETWEEN, NE, ADD, SUB, MUL, DIV, MOD, NEG, BITAND, BITOR, BITXOR, BITNEG, LSHIFT, RSHIFT, TYPECAST_TIME_AT, TYPECAST_IP_ADDRESS_AT, TYPECAST_DOUBLE_AT, TYPECAST_UNSIGNED_LONG_AT, WEEKDAY_STRING, WEEKDAY_STRING_SHORT, SIGNED8_STRING, UNSIGNED8_STRING, SIGNED16_STRING, UNSIGNED16_STRING, SIGNED32_STRING

[From Build 63.16] [# 534244]

- If packet tracing is configured with a default-syntax expression and non-TCP traffic is being processed, and rewrite action applied on a HTTP chunked message is occurring then the rewritten data maybe incorrect or it might crash a NetScaler appliance.

[From Build 64.34] [# 598465]

- An HTTP callout that is configured for use with a virtual server does not work with a backup virtual server (if configured).

[From Build 64.34] [# 382341, 540646, 585790]

- The default timeouts for Rewrite Processing and for Advanced Expression Regex Evaluation have changed from 1 millisecond to less than the pitboss timeout of 5 seconds. This restores the default behavior for releases prior to 11.0. In addition, an optional -timeout parameter to "set the re-write param " CLI command was added. The time is measured in milliseconds - see the man page. A "set policy param -timeout <value>" command has been added to the CLI. These ways of setting the timeout work for all partitions.

In release 11.0, the default timeouts for Rewrite Processing and for Advanced Expression Regex Evaluation have also changed from 1 millisecond to less than the pitboss timeout. This restores the default behavior from releases prior to 11.0. However, neither the CLI command, nor the GUI, nor the Nitro call is available. Instead, for 11.0, an nsapimgr command is available from the shell. This will only change the timeout on the default partition when Partitioning is used. Other partitions will only use the default. The syntax is as follows: "nsapimgr_wr.sh -ys arg1=<value> -ys call=ns_pixl_regex_set_time_limit" to set the time limit on regular expression evaluation in Advanced expressions. "nsapimgr_wr.sh -ys arg1=<value> -ys call=ns_rw_set_eval_time_limit" to set the time limit on Rewrite processing. For either of these, setting the value to 0 resets the limit to the default. These nsapimgr commands will not be supported after 11.0, and the CLI, GUI, or Nitro must be used.

[From Build 65.35] [# 577016, 578214]

- Under certain conditions, a NetScaler appliance does not insert an X-Forwarded-For field in the HTTP header for HTTP CONNECT requests that are forwarded to server.

[From Build 65.35] [# 605089]

- **Rewrite action block leading to subsequent action time outs.**

If multiple rewrite policies evaluate to TRUE for a particular protocol and direction (for example HTTP request or TCP response), and more than one associated action is selected for execution, they might not all execute. If one of the actions is suspended (blocked), the next selected action will time out, and any subsequent actions will be skipped.

The following functions in expressions can block:

- * HTTP_CALLOUT
- * MATCHES_LOCATION
- * STREAM
- * CHECK_LIMIT
- * MATCHES
- * BODY
- * PAYLOAD
- * MSSQL
- * MYSQL
- * ORACLE
- * SUBSCRIBER
- * DETERMINE_SERVICES
- * use of variables (in other words, \$<variableName>)

Note: These expressions can block, depending on specific conditions that occur at the time of execution.

[From Build 65.35] [# 628326]

- A memory leak occurs when a responder action has blocking expressions (for example, stream analytics, HTTP callout, matches_location) and body or payload based expressions.

[From Build 65.35] [# 598252, 623764, 624637, 624759, 629247, 629344]

- In a NetScaler appliance if the set memory limit is insufficient to support the configured number of VPN sessions, the appliance fails.

[From Build 67.12] [# 636579]

- A NetScaler appliance fails if you perform a Clear Configuration operation.

[From Build 67.12] [# 634124]

- While evaluating default syntax expression for local time zone, a NetScaler appliance incorrectly applies US daylight savings time (DST) rules in non-US timezone. This results in setting an offset time for an hour. For example, the default expression `!(SYS.TIME.GE (LOCAL 8h) & SYS.TIME.LE(LOCAL 17h))` returns 'False' if the local time in US timezone is between 0800 and 1700. In the UK timezone, this expression incorrectly returns 'False' if the local time is

between 0700 and 0759 and returns 'True' if the local time is between 1700 and 1759 from 8 Mar 2015 (the start of US DST) to 28 Mar 2015 (the day before the start of UK DST) and also from 25 Oct 2015 (the day after the end of UK DST) to 31 Oct (the day before the end of US DST).

[From Build 68.10] [# 556230]

- While evaluating default syntax expression for local time zone, a NetScaler appliance incorrectly applies US daylight savings time (DST) rules in non-US timezone. This results in setting an offset time for an hour. For example, the default expression `!(SYS.TIME.GE (LOCAL 8h) & SYS.TIME.LE(LOCAL 17h))` returns 'False' if the local time in US timezone is between 0800 and 1700. In the UK timezone, this expression incorrectly returns 'False' if the local time is between 0700 and 0759 and returns 'True' if the local time is between 1700 and 1759 from 8 Mar 2015 (the start of US DST) to 28 Mar 2015 (the day before the start of UK DST) and also from 25 Oct 2015 (the day after the end of UK DST) to 31 Oct (the day before the end of US DST).

[From Build 68.11] [# 556230]

Responder

- The NetScaler appliance fails if it receives a new request while an embedded expression in the responder HTML page is in blocking state.

[From Build 65.35] [# 556035]

SSL

- You cannot enable TLSv1.1/1.2 on a front end SSL service after explicitly disabling it.

[From Build 62.10] [# 574589]

- If you have configured optional client-certificate authentication and your policies target client certificate x509 extensions, such as auth keyid, a transaction with a client that doesn't have a certificate might cause the appliance to fail or to use stale values from a previous transaction.

[From Build 63.16] [# 593091]

- If TLS1.1/1.2 protocol is used with AES/3DES ciphers, the length of the TCP window at the back end shrinks to zero. As a result, after some time, the connection is terminated.

[From Build 63.16] [# 591600, 595713, 596278, 596556, 596566, 598045, 599524, 600591, 604929]

- If you update the certificate-key pair for a service group, the change is not reflected in the individual services that are bound to this service group. As a result, the old certificate-key pair continues to be used for negotiation in the SSL handshake.

[From Build 63.16] [# 554925]

- If you have a large number of SSL services (greater than 3000) in the backend, CPU usage increases exponentially and the appliance fails.

[From Build 63.16] [# 581193]

- If you downgrade the software on your NetScaler appliance that does not have a license to release 9.3 build 61.66 or earlier, some commands related to the default server certificate might not be saved in the running configuration. As a result, after restarting, secure access (HTTPS) to the appliance fails.

[From Build 63.16] [# 551603, 559154]

- If you use the "add ssl certkey" command to add an encrypted .pfx file, the password is now encrypted and saved in the configuration file (ns.conf). In earlier releases, the password was not saved, so automatic execution of the add ssl certkey command failed when the appliance was restarted.

[From Build 63.16] [# 591167]

- An incoming SSL record that spans more than 256 TCP packets and contains TCP header options causes memory corruption in the Cavium command buffer structure. As a result, the NetScaler appliance fails.

[From Build 63.16] [# 573904, 583295, 590222]

- In some cases, when client authentication is enabled, incorrect data from a client leads to a memory leak on the NetScaler appliance. If a large number of clients send incorrect data, the appliance fails.

[From Build 63.16] [# 570754]

- If TLS1.1/1.2 protocol is used with AES/3DES ciphers, the length of the TCP window at the back end shrinks to zero. As a result, after some time, the connection is terminated.

[From Build 64.34] [# 591600, 595713, 596278, 596556, 596566, 598045, 599524, 600591, 604409, 604929]

- Even though SSL renegotiation is set to deny (that is, denySSLReneg is set to ALL), the server responds with the "server renegotiation" extension in the initial SSL handshake.

[From Build 64.34] [# 559082]

- NetScaler VPX virtual appliances do not support AES-GCM/SHA2 ciphers, but in earlier builds you can bind these ciphers, incorrectly, to an SSL virtual server. From the current build, you cannot bind these ciphers to the virtual server. If you have bound AES-GCM/SHA2 ciphers to a VPX instance that you upgrade to the current build, the bind commands in the configuration return an error. In a comparison of the configurations of the old and new build, the missing bindings can be mistakenly construed as a configuration loss.

[From Build 64.34] [# 609476]

- In release 10.5 or later, TLS protocol versions 1.1 and 1.2 are enabled by default, but you can disable them for all services except SSL_BRIDGE and dynamic services, which can't otherwise be configured. In this release, you can

disable TLS1.1/1.2 on SSL_BRIDGE and dynamic services by enabling the new svctls1112disable and montls1112disable parameters, as follows:

```
> set ssl param -svctls1112disable enable -montls1112disable enable
```

After the new parameters are enabled, you cannot disable them by using the "set ssl param" command. You must edit the configuration (ns.conf) file as follows:

1. Remove these parameters from the "set ssl param" command.
2. Save the configuration.
3. Restart the appliance.

[From Build 64.34] [# 602502, 599209, 609284]

- If the passphrase for a certificate contains the "\$" character, the configuration utility becomes unresponsive.

[From Build 64.34] [# 591743]

- If you enable the DH parameter while creating an SSL profile by using the configuration utility, the following error message appears:

Error in retrieving File. Invalid args in query parameters

[From Build 64.34] [# 594922]

- If you have configured optional client-certificate authentication and your policies target client certificate x509 extensions, such as auth keyid, a transaction with a client that doesn't have a certificate might cause the appliance to fail or to use stale values from a previous transaction.

[From Build 64.34] [# 593091]

- If you bind a secure monitor to a service, such as SSL_BRIDGE, that does not allow SSL configuration, the default settings are used. The default SSL version sent in the SSL handshake record header is SSLv3.

Contact Citrix support if you want to disable SSLv3 and use the next higher protocol.

[From Build 64.34] [# 584424]

- In the OpenSSL interface in the NetScaler configuration utility, if you type a command before the OpenSSL> prompt appears, the OpenSSL> prompt might not appear at all. As a result, any commands that you type are not run in OpenSSL mode.

[From Build 64.34] [# 595413]

- When clearing the NetScaler configuration, user-defined cipher groups that are bound to an internal SSL service might get corrupted. Subsequent cipher bind or unbind operations with that service will cause the appliance to become unresponsive.

[From Build 65.35] [# 611894, 622042]

- You cannot install a certificate on the appliance if the certificate is not in the /nsconfig/ssl directory. With this fix, you can install a certificate in the appliance's default partition from any location. For other partitions, the certificate must be in the /nsconfig/partitions/ssl/ directory.

[From Build 65.35] [# 602631]

- The NetScaler appliance fails if it parses the value of an unknown certificate extension while the certificate is loading.

[From Build 65.35] [# 623996]

- After you upgrade to this build, configuring a front-end service, or creating an internal service, with default ciphers results in a cipher inconsistency between a packet engine and the cluster configDB.

[From Build 66.11] [# 625966]

- The appliance fails if a loop is created while linking the certificates. With this fix, the software checks whether a new certificate is already part of the link.

[From Build 67.12] [# 612461]

- In a cluster setup, if you have configured a front end service, or an internal service is created, with default ciphers, and then you upgrade to this build, there is a cipher inconsistency between a packet engine and the cluster configDB.

[From Build 67.12] [# 631258]

- The output of the "stat ssl -detail" command is different for back-end entities than for front-end entities. The output for back-end entities does not include statistics for sessions, handshakes, or client authentications for TLS protocol versions 1.1 and version 1.2.

At the back end, the label "Authorizations" is incorrect. It should be "Authentications."

[From Build 67.12] [# 627635]

- In a cluster setup, an error message appears if you run any of the following commands:

```
show ssl vserver [<vServerName>] [-cipherDetails]
```

```
show ssl service [<serviceName>] [-cipherDetails]
```

```
show ssl serviceGroup [<serviceGroupName>] [-cipherDetails]
```

[From Build 67.12] [# 628296]

- If you have configured a DTLS virtual server and also enabled the Default profile, clearing the configuration might cause the appliance to fail.

Note: On some STA services, if you enable DTLS, you don't have to explicitly create a DTLS virtual server, because it is automatically created. The automatically created DTLS server does not cause the configuration to fail if the Default profile is enabled.

[From Build 67.12] [# 622916, 625989, 632989, 634897, 639567]

- SSL internal services might fail if you modify any SSL parameters while the SSL feature is disabled or not licensed.

[From Build 67.12] [# 601951]

- If you configure a clear text port on an SSL-based wildcard virtual server, and run the clear config command with the extended+ option, the NetScaler appliance fails.

[From Build 67.12] [# 586199]

- In a cluster setup, the serial number and validity do not appear correctly in the output of the "sh ssl certkey" command.

[From Build 67.12] [# 635851, 504829]

- In a cluster setup, if you try to update an existing certificate by replacing the old files with new certificate and key files, the following error message appears:

ERROR: Resource already exists [certkeyName Contents, nglab-2016]

[From Build 67.12] [# 633395]

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

[From Build 68.10] [# 652477, 654559, 655657, 656035, 657343]

- The output of the "stat ssl -detail" command is different for back-end entities than for front-end entities. The output for back-end entities does not include statistics for sessions, handshakes, or client authentications for TLS protocol versions 1.1 and version 1.2.

At the back end, the label "Authorizations" is incorrect. It should be "Authentications."

[From Build 68.10] [# 627635]

- The output of the "stat ssl -detail" command is different for back-end entities than for front-end entities. The output for back-end entities does not include statistics for sessions, handshakes, or client authentications for TLS protocol versions 1.1 and version 1.2.

At the back end, the label "Authorizations" is incorrect. It should be "Authentications."

[From Build 68.11] [# 627635]

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

[From Build 68.11] [# 652477, 654559, 656035, 657343]

System

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

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

[From Build 62.10] [# 580890]

- Management CPU usage is high when you use the configuration utility's memory usage diagnostic tool (System > Diagnostics > Memory usage).

[From Build 63.16] [# 586328]

- The NetScaler appliance might become unresponsive if it receives a retransmitted TCP jumbo frame that carries the TCP FIN flag.

[From Build 63.16] [# 571176]

- In NetScaler Insight Center, NetScaler 1000V, and NetScaler VPX on ESX, the Vmtoolsd daemon fails during start up and creates a core dump in the directory /var/core. It does not affect normal VPX functionality. However, operations such as "Shut Down Guest" and "Restart Guest" from the vSphere client summary tab fail.

[From Build 63.16] [# 570166, 477094, 498384, 520519, 530951, 543554, 555689, 585809]

- The upgrade wizard in the configuration utility puts the NetScaler software in the /var directory instead of the /var/nsinstall/<build id> directory.

[From Build 63.16] [# 586721]

- When SPDY Protocol is enabled and SPDY Traffic is received on the NetScaler appliance, the TCP current clients counter goes to negative values and shows a very large value in the stat or the SNMP OID.

[From Build 63.16] [# 551562, 551786, 568554]

- After cleaning up an MPTCP session, the NetScaler appliance might not set the DATA_FIN flag in the TCP header of the data or acknowledgement packet if there is no subflow for sending the data.

[From Build 63.16] [# 553650]

- For a NetScaler appliance with extended memory configured for Large Scale NAT (LSN) feature, after warm rebooting the appliance, when the appliance is added as secondary node to an appliance that does not have the extended memory configured for LSN, the secondary appliance becomes unresponsive.

[From Build 63.16] [# 593261]

- If a server advertises a maximum segment size (MSS) greater than 1460 bytes, a TCP transaction might not generate a response after passing through the NetScaler appliance.

[From Build 63.16] [# 584079]

- On rebooting the NetScaler appliance, the timeout is not set to the value specified by the "set ns timeout" command.

[From Build 63.16] [# 587074]

- The NetScaler appliance might become unresponsive if front end optimization (FEO) is enabled with the SSL and rewrite features.

[From Build 63.16] [# 583829]

- If a NetScaler appliance that is sending auditlog messages over TCP (audit syslogaction specifies TCP as the transport protocol) has more than 200 million active sessions, the rate at which the syslogs are sent drops to 700 Kbps or lower, and the appliance consumes a high percentage of the CPU cycles.

[From Build 63.16] [# 580309]

- If you execute NTP commands, such as enable ntp sync and show ntp status, the NetScaler appliance might become unresponsive because of a memory leak.

[From Build 63.16] [# 529787, 546378, 574866, 581849]

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

[From Build 63.16] [# 575108, 581862]

- Syslog messages generated by user action are logged as error messages instead of informational messages.

[From Build 64.34] [# 538212]

- If weblog data is sent over a TCP connection that the NetScaler appliance has terminated because of buffer overflow, the appliance fails. With the fix, the connection is checked to ensure that it is not closed before the weblog data is sent.

[From Build 64.34] [# 593968, 574996]

- After you upgrade a NetScaler appliance to 10.5 build, the Client-Server Link Mapping check box is now available in the TCP Connections page

[From Build 64.34] [# 551611, 519966]

- In certain cases, the NetScaler appliance might not retransmit the lost TCP segments resulting in a transaction failure.

[From Build 64.34] [# 565938, 560394, 592227, 597160, 607864, 609068]

- A NetScaler appliance might occasionally fail when a client connects to an HTTP/SSL server and the server sends a 101 (switching protocols) response. The connection is closed before data can be sent or received from the client.

[From Build 64.34] [# 576561, 587759]

- When the NetScaler appliance receives MPTCP traffic, the number of established client connections is high, because both MPTCP sessions and subflows are treated as client connections.

With this fix, the SNMP OID of following MIBs have changed to:

mptcpCurSessWithoutSFs: 130

vsvrCurMptcpSessions: 73

vsvrCursubflowConn: 74

[From Build 64.34] [# 583292]

- When adding a syslog action for which the netProfile parameter is set, the Subnet IP (SNIP) address is used as the source IP address for sending log messages. If the netProfile parameter is not set, the NetScalerIP (NSIP) address is used as the source IP for sending the log messages.

[From Build 64.34] [# 595449]

- The appliance might fail under the following set of conditions:

1. A pipelined HTTP request is received that spans multiple TCP segments.
2. An internal HTTP response generated by NetScaler for the HTTP request in condition 1, is terminated by a TCP segment that has the TCP FIN flag set.
3. The appliance receives another HTTP request on the same connection.

[From Build 64.34] [# 587817, 587879, 589416, 594044, 595927, 601915, 610728]

- If the NetScaler appliance receives a data or an acknowledgement packet without the Data Sequence Signal (DSS) option before the MPTCP connection is established, the appliance does not seamlessly fallback to regular TCP.

[From Build 64.34] [# 588909]

- A NetScaler appliance has high memory consumption if Front End Optimization (FEO) feature is enabled.

Work around: Disable the feature or reboot the appliance.

[From Build 64.34] [# 591928]

- The NetScaler appliance does not reduce the received Maximum Segment Size (MSS) to accommodate TCP options (such as timestamps). Therefore, the NIC drops such packets.

[From Build 64.34] [# 593209]

- Some events may be logged twice if DEBUG level is enabled for syslog, by using the "set audit syslogParam" command.

[From Build 64.34] [# 594485]

- For a NetScaler appliance with extended memory configured for Large Scale NAT (LSN) feature, after warm rebooting the appliance, when the appliance is added as secondary node to an appliance that does not have the extended memory configured for LSN, the secondary appliance becomes unresponsive.

[From Build 64.34] [# 593261]

- In a HA setup, if a domain-based SNMP manager is added on the secondary appliance, the NetScaler appliance stops responding eventually. You must configure the SNMP manager on the primary appliance.

[From Build 64.34] [# 581355, 593292, 595943]

- A NetScaler appliance might crash if you attempt to start the nstrace instance with advanced filter expression.

[From Build 64.34] [# 493737, 526095, 598148]

- Failed SNMP requests were not removed properly, therefore, subsequent set requests were retained in the queue. This led to all SNMP requests getting blocked and high memory usage, due to which the SNMP module stops responding.

[From Build 64.34] [# 590289, 584527, 596242]

- When parsing a host name with no Path component, the URL parsing logic does not search for a question mark (?), so an entire string might be interpreted as the host name. This causes an error when the appliance tries to resolve the DNS name. With this fix, the parsing logic searches for question marks.

E.g.:

`http://example.com.php?&curuserid=94315577&host=wscdny203.live.changba.com&token=T59d105c1c74042e&localip=221.235.187.75&clientip=80.95.239.1&bless=1&channelsrc=market_%E7%99%BE%E5%BA%A6`

[From Build 64.34] [# 587858]

- A NetScaler appliance fails when an MPTCP subflow receives an Infinite DSS mapping in a partially retransmitted packet.

[From Build 65.35] [# 614842, 623426]

- A warning error message "Error =80000004 in nsagg_process_stat_request, closing connection" displays when a nscollect module requests counter information from a nsaggregator daemon at every 5-minute interval. The

nsagggregator daemon prints the warning message as response to the request received from nscollect module for more than 256 counters.

[From Build 65.35] [# 610809, 577474, 579560]

- If, when establishing an MPTCP connection, a NetScaler appliance receives a duplicate acknowledgment in the 3-way handshake process, the appliance reverts to a normal TCP connection.

[From Build 65.35] [# 601372]

- Client-based virtual machines are unable to access a NetScaler appliance if they are running on the same server (for example, VPX on Linux KVM). However, they are accessible if they are running on different servers.

[From Build 65.35] [# 613108]

- A configuration loss occurs when you upgrade Content Switching (CS) entities above the following limits:

Memory allotted to a packet engine 2GB 4GB

Max number of content switching virtual servers 1200 2500

Max number of content switching policies 1500 3500

Max number of content switching virtual server bindings 3000 4000

[From Build 65.35] [# 628528]

- Commands entered in the NetScaler CLI or GUI might fail because of a shortage of system resources or failure of system socket connections.

With this fix, the NetScaler appliance attempts to reestablish the socket connections. After the socket connections are established, the appliance runs the failed commands internally.

[From Build 65.35] [# 615487]

- A NetScaler appliance fails when it encounters an HTTP/2 connection level error on a TCP connection.

[From Build 65.35] [# 615395]

- In a high availability setup with stateful connection failover option enabled on a virtual sever, if a network link that is used for synchronizing connection information between the nodes becomes DOWN.

Both nodes take a lot a time to reestablish connection information synchronization through the remaining active links, as a result some connection information might not get synchronize to the secondary node.

[From Build 65.35] [# 590574]

- Support restore operation on the NetScaler Appliance by using a remotely stored backup

You can now use a remotely saved backup to restore a NetScaler appliance through the "add system backup <filename>", that adds the metadata to the remote backup package, so that the restore operation can successfully use the backup package.

[From Build 65.35] [# 569974]

- Management access to the NetScaler appliance can slow down or become unavailable when the traffic domain identifier is not initialized for jumbo frames. However, virtual servers continue to serve traffic.

[From Build 65.35] [# 583579, 594722, 626120]

- If, when processing a URL, the parser encounters a tag that has "#" as a source attribute, the URL is considered to be empty as # is a fragment identifier. This leads to corrupted values because we continue processing the empty URL.

[From Build 65.35] [# 605258]

- Due to a bug in Hard Disk Drive (HDD) monitoring logic, if a message in /var/log/messages matches "*ad* Device not configured" string pattern, it results in producing false positive errors.

[From Build 65.35] [# 611774, 598774]

- The NetScaler Weblog client intermittently fails because of incorrect indexing, leading to segmentation failure

[From Build 65.35] [# 615895, 620767]

- A NetScaler appliance fails when it receives an MP_CAPABLE final acknowledgement in a single packet with the FIN flag set.

[From Build 65.35] [# 583853, 583855, 588078, 601746, 602955]

- With the default TCP congestion control, a NetScaler appliance recovering from packet loss reduces the congestion window to half its previous length. With multiple packet loss events, the congestion window becomes small and delays transactions.

[From Build 65.35] [# 606493, 601655, 623185]

- A T1200 appliance that is used in a NetScaler deployment can become unresponsive or fail when generating the NetScaler tech support logs.

[From Build 65.35] [# 606247, 624369, 624385]

- You cannot shut down or restart the virtual machine by using the VMware vSphere tool.

[From Build 65.35] [# 607158]

- A NetScaler appliance might fail because of a segmentation fault if it receives a large HTTP/2 request Header that evicts the dynamic header table entry.

[From Build 65.35] [# 615629]

- When a NetScaler appliance is integrated with ESP or VPX devices functioning as E100 devices, it encounters buffer-allocation failure and packet-reception failure.

[From Build 65.35] [# 604971, 611176]

- A NetScaler appliance becomes unresponsive when passing an HTML response with the HTML tag exceeding 16 characters.

[From Build 65.35] [# 611723]

- In certain cases, the NetScaler appliance might not retransmit the lost TCP segments resulting in a transaction failure.

[From Build 65.35] [# 565938, 560394, 592227, 597160, 607864, 609068]

- A NetScaler appliance does not support Base64 decoded TASS cookie IDs of more than 64 characters. If Security Assertion Markup Language (SAML) or federation results in an ID longer than 64 characters, the appliance does not support the cookie ID.

[From Build 65.35] [# 594603, 607019, 615811]

- The NetScaler appliance might fail if both of the following conditions are met:
 - One or more of the following features are configured on the appliance: cache redirection, content switching, AAA-TM, Clientless VPN, full tunnel VPN, forward proxy.
 - The client connection times out while the DNS name is being resolved using the FQDN of back-end servers.

[From Build 66.11] [# 543293, 578993, 579965, 593378, 599535, 608479, 614368, 628579, 628763, 634338]

- A NetScaler appliance fails when an MPTCP subflow receives an Infinite DSS mapping in a partially retransmitted packet.

[From Build 66.11] [# 623426]

- A high availability pair fails if an HTTP response from a back-end server contains carriage return line feeds (CRLFs) after the HTTP Content Length and at the start of a new packet.

[From Build 66.11] [# 547267, 623146]

- In a wildcard virtual server configuration, a NetScaler appliance dynamically identifies an origin service by opening a probe connection. If the origin responds with a jumbo Maximum Segment Size (MSS), the appliance uses the MSS for future connections with the origin. If the jumbo frame support is disabled, it results in transactions failure.

[From Build 66.11] [# 605873]

- A NetScaler appliance fails if the Front End Optimization (FEO), Application Firewall, and SSL features are all enabled and the appliance encounters an error while parsing an HTML response.

[From Build 66.11] [# 624327]

- 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 66.11] [# 603225]

- A NetScaler appliance fails on a network interface if it receives retransmitted data for which the maximum transmission unit (MTU) is larger than 1500 bytes.

[From Build 67.12] [# 625776, 624763, 624779, 629314, 630646, 636283, 637479]

- After an upgrade, if a partition id is missing in a HTTP abort transaction log record, it results in a NetScaler Web Log (NSWL) utility failure.

[From Build 67.12] [# 623983]

- In a high availability setup, command propagation and configuration synchronization using secure RPC might fail if SSLv3 and TLS1.0 protocols are disabled for SSL internal services.

[From Build 67.12] [# 613966]

- A NetScaler appliance fails if you enable the Front End Optimization (FEO) feature and the User Agent Request header size exceeds 8 KB.

[From Build 67.12] [# 636662]

- After a reboot, if you increase the Cache memory limit with Front End Optimization (FEO) feature enabled, it leads to a NetScaler appliance failure.

[From Build 67.12] [# 626082, 628536, 633772, 642939]

- If, in a Multipath TCP (MPTCP) session with a single subflow, the client in the subflow signals a zero-window condition before the subflow connection times out, the NetScaler appliance uses small-window-protection logic to mark the subflow connection as a small-window attack from the client. The logic checks to determine whether the existing number of small window connections exceeds the threshold value (set to 100, by default) and if true, the appliance resets the subflow causing the appliance to fail.

[From Build 67.12] [# 639081]

- If a client on an IPV6 connection advertises an MSS value below 1360 (bytes), the NetScaler appliance responds with a MSS value below the (RFC) required minimum value of 1220 (bytes).

[From Build 67.12] [# 556475]

- A NetScaler Policy Infrastructure (PI) connection reset code for a non-HTTP virtual server might cause a memory leak.

[From Build 67.12] [# 626562, 632738, 634610]

- If you use the GUI to access a NetScaler appliance, the appliance sends configuration change SNMP traps, even if there has been no change in the configuration settings.

[From Build 67.12] [# 634480]

- In an HTTP/2 transaction, if a response is sent from the Integrated Caching (IC) module, the transaction might fail because the NetScaler appliance advertises zero window to the client.

[From Build 67.12] [# 635102]

- An interface based expression might be evaluated incorrectly. In previous releases, evaluation of an interface-based expression was based on the information available in the connection block as well as the information available in the individual frame. Now, only the information in the frame is considered, and this information can change during the course of a transaction.

[From Build 68.10] [# 597312]

- The TCP wait queue counter might be incorrect, because the NetScaler appliance does not update the counter properly during persistence probes.

[From Build 68.10] [# 637919]

- The NetScaler appliance might fail, because of memory corruption, if a policy uses an expression that applies the MATCHES (not MATCHES_LOCATION) function to an IPv4 or an IPv6 address and there is an issue in communicating with the DNS server.

[From Build 68.10] [# 630782, 630436, 631279, 637396, 650939, 650964]

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

[From Build 68.10] [# 643237]

- An invalid compressed header in SPDY frames causes a NetScaler appliance to restart.

[From Build 68.10] [# 637651]

- 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 68.10] [# 633308, 646753, 648657]

- During a TCP transaction, when the client advertises zero window to a NetScaler appliance, NetScaler periodically sends zero window probe to ascertain if the client can open the window so that NetScaler can send in new data. When sending such a probe, NetScaler sends a full maximum segment size (MSS) packet during first probe and from the second probe onwards, sends a 1-byte packet. If the client does not open the window after sending such a probe, but instead sends a TCP Reset or if the connection on NetScaler gets flushed for other reasons, then it may lead to duplicate buffer free on NetScaler that may lead the appliance to crash.

[From Build 68.11] [# 657742, 657753, 657771, 658352, 658507, 658526, 659842, 659849, 660345, 660812, 660998, 661018, 661266, 661511]

- 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 68.11] [# 649334, 653370, 656768]

- An interface based expression might be evaluated incorrectly. In previous releases, evaluation of an interface-based expression was based on the information available in the connection block as well as the information available in the individual frame. Now, only the information in the frame is considered, and this information can change during the course of a transaction.

[From Build 68.11] [# 597312]

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

[From Build 68.11] [# 643237]

- 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 68.11] [# 633308, 646753, 648657, 656502]

- The NetScaler appliance might fail, because of memory corruption, if a policy uses an expression that applies the MATCHES (not MATCHES_LOCATION) function to an IPv4 or IPv6 address and there is an issue in communicating with the DNS server.

[From Build 68.11] [# 630782, 630436, 631279, 637396, 650939, 650964]

- The TCP wait queue counter might be incorrect, because the NetScaler appliance does not update the counter properly during persistence probes.

[From Build 68.11] [# 637919]

- An invalid compressed header in SPDY frames causes a NetScaler appliance to restart.

[From Build 68.11] [# 637651]

- When NetScaler sends out full sized persist probe packet that is more than the client advertised window, firewall drops the packet causing the connection to fail.

[From Build 68.11] [# 576980]

Telco

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

[From Build 63.16] [# 487318]

- In a network setup that includes both dynamic and deterministic types of clients, the first request from a deterministic client is not served if a dynamic client has sent a request.

[From Build 63.16] [# 576602]

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

[From Build 64.34] [# 579797]

- In a high availability deployment with LSN and DS-Lite configuration, LSN and DS-Lite mappings for active FTP connections are not removed from the secondary node even after they time out or are flushed.

[From Build 65.35] [# 601920, 619864]

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

[From Build 65.35] [# 582464]

- In a DS-Lite configuration with a server behind the B4 device, the NetScaler appliance does not properly process FTP packets that have the following set of characteristics:

- * Are from clients on the Internet

- * Are destined to the server

- * Match DS-Lite static NAT maps configured on the NetScaler appliance

[From Build 67.12] [# 601560]

- **Web Interface on NetScaler (WlonNS)**

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

[From Build 62.10] [# 576883]

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

[From Build 63.16] [# 577988]

Release History

For details of a specific release, see the corresponding release notes.

- Build 69.12 (2016-12-05) (Current build)
- Build 68.11 (2016-10-04) Replaces: 68.10
- Build 68.10 (2016-09-01)
- Build 67.12 (2016-07-14)
- Build 66.11 (2016-05-31) Replaces: 4.0
- Build 65.35 (2016-04-21) Replaces: 65.31
- Build 64.34 (2015-12-30)
- Build 63.16 (2015-10-11)
- Build 62.10 (2015-08-12)
- Build 55.23 (2015-06-30) Replaces: 55.20