Citrix NetScaler 1000V Release Notes

Citrix NetScaler 11.1-47.14
First Published: 2016-08-02

Cisco Systems, Inc.
www.cisco.com

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

Contents

11.1-47.14

Updated: July 28, 2016 | Release notes version: 1.0

This release notes document describes the enhancements and changes and specifies the issues that exist, for the NetScaler 11.1 Build 47.14 release.

What's New?

The enhancements and changes that are available in Build 47.14.

AAA-TM

- **OAuth Support for Multi-Factor Authentication**

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

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

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

  [# 611735, 572701, 572705]

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

  [# 620908]

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

  [# 581544, 475585, 552072, 606858, 619869]

  With this enhancement, NetScaler now supports GET requests from SAML SPs.

[# 564947, 590768]

Admin Partitions

- Shared VLAN Support

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

[# 581671]

Clustering

- PBR Support for Cluster

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

  [# 611938]

GSLB

- Support for EDNS0 Client Subnet

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

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

  [# 457159]

Load Balancing

- Secure FTP Monitoring Support

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

  [# 237766]

- Closing Monitor Connections at the Service Level

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

[# 607661]

- FIX Protocol Support

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

  [# 634096]

- Required Unbind Operation Prevents Accidentally Disabling a Virtual Server

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

  [# 258327]

- Support for Load Balancing Profile

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

  [# 353669]

- Setting SSL Parameters on a Secure Monitor

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

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

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

  SSL profiles can be bound to the following monitor types:

  - HTTP

- HTTP-ECV

- TCP

- TCP-ECV

- HTTP-INLINE

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

At the command prompt, type:

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

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

Example:

add ssl profile prof1 -sslProfileType BackEnd

add lb monitor mon1 HTTP -secure YES -sslprofile prof1

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

At the command prompt, type:

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

[# 506771]

- Improved Support for Persistency

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

  [# 571771]

- Closing Monitor Connections at the Service Group Level

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

  [# 628111]

- Configuring an HTTPS Virtual Server to accept HTTP Traffic

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

  [# 570157]

NITRO

- Handle Multiple NITRO Calls in a Single Request

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

  [# 598559]

- Simplify Management Operations with an idempotent API

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

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

  [# 601351]

- Automate NetScaler Upgrade and Downgrade with a Single API

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

  [# 598557]

- Retrieve Bindings in Bulk

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

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

  [# 600350]

NetScaler CLI

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

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

  set aaa radiusParams - authservRetry <positive_integer>

  [# 492179]

NetScaler GUI

- Tabular, one-page Application Firewall Wizard

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

  [# 587433, 557185, 619712]

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

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

  [# 624858]

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

  [# 619498]

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

  [# 597501]

- Icons for Action and Information Menus

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

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

  [# 614868]

- 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.

  [# 614285, 620953]

- SSL Certificate Management GUI Enhancements and Changes

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

  - Create RSA Key

  - Create DSA Key

  - Create CSR

  - Create Certificate

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

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

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

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

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

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

[# 612894]

- Improved IPv4 Address Fields

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

  [# 610522]

  - Diagnostic of Start New Trace and Support Stop Running Trace

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

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

  [# 564499, 565594]

- High Availability Status Information in the Top Pane

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

  [# 423777, 466239, 582803]

- Non-Blocking LDAP SSL/TLS authentication support added. This enhancement reduces the authentication bottleneck to prevent delayed/denied user logons.

  [# 609519]

- This enhancement provides NetScaler, which is acting as SAML IDP, the capability to sign the entire SAML response along with the assertion.

[# 620844]

- This enhancement provides the ability to a clear config basic command so it will not erase the TACACS related configuration.

  [# 515227]

- This enhancement allows an admin to view the different LoginSchemas present in NetScaler. This is done from the NetScaler admin GUI under LoginSchema Profiles configuration. Use the following path to see the LoginSchemas: Configuration > Security > AAA - Application Traffic > Login Schema > Profiles.

  [# 617921]

- This enhancement introduces support for a new key transport algorithm (RSA-V1_5). The RSA-V1_5 can used to encrypt SAML assertions along with RSA-OAEP.

  [# 580078]

- The NetScaler appliance inserts an NS_ESNS cookie for page tracking (for showing a waterfall chart) when AppFlow is enabled. Cookie insertion was controlled by the clientSideMeasurements option in the appflow action in release 10.5, but in release 11.0 the default became to always insert the cookie when appflow is enabled. Android receiver (HTTP client) was not able to handle this cookie. This fix adds the Enable/Disable page tracking (cookie insertion) option to the appflow action.

  You can now enable or disable Page tracking feature from NetScaler Insight Center.

  To perform this action, navigate to Configuration > System > Appflow > Actions. Edit an AppFlow action name, and select the Page Tracking check box.

  [# 613351, 598478, 608448]

NetScaler Insight Center

- 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

  [# 614892, 550997, 604388, 620422, 632370]

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

  [# 590782]

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

  [# 498439]

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

  [# 534805]

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

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

  [# 606189]

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

  [# 606187]

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

  [# 549746]

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

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

  [# 635528]

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

  [# 541439]

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

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

  [# 582943]

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

  [# 606188]

- NetScaler VPX Appliance

  New license for NetScaler VPX on ESX and KVM platforms

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

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

  [# 623179]

Networking

- Network Service Header support for Service Function

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

  A NetScaler appliance can now play the service-function role in a SFC architecture. The NetScaler appliance receives packets with Network Service headers and, upon performing the service, modifies the NSH bits in the response packet to indicate that the service has been performed. In that role, the appliance supports symmetric service chaining with features (for example, INAT, TCP and UDP load balancing services, and routing). The NetScaler appliance as service-function does not support IPv6 and Reclassification.

  [# 593459]

- NITRO API Support for Dynamic Routing

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

[# 626083]

- Setting the MTU on the NSVLAN

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

  [# 425950]

- NetScaler Support for Microsoft Direct Access Deployment

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

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

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

  [# 612455]

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

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

  - Start time of the RNAT session.

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

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

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

  -- TIMEOUT. The RNAT session timed out.

[# 609410]

- Using NULL Policy Based Routes to Drop Outgoing Packets

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

  [# 451632]

- Managing High Availability Heartbeat Messages on a NetScaler Appliance

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

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

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

  [# 477162, 575447, 604578]

- Dynamic Routing support for Link-Local Subnet IPv6 addresses

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

  [# 553544]

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

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

  The NetScaler supports using a source port from a specified port range for communicating to the servers. One of the use case of this feature is for servers that are configured to identify received traffic belonging to a

specific set on the basis of source port for logging and monitoring purposes. For example, identifying internal and external traffic for logging purpose.

[# 420067, 420039]

- Configuring Allowed VLAN List

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

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

  [# 495219]

- 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.

  [# 457167]

- 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]

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

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

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

  [# 551505]

- IPv6 Support in Active-Active Mode using VRRP

  NetScaler Appliances Support VIP6 Addresses in Active-Active Deployments.

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

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

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

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

* Preemption

* Delaying preemption

* Sharing

* Changing VIP address priority automatically

[# 553570]

- Graceful Restart for Dynamic Routing Protocols

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

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

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

  - Border Gateway Protocol (BGP)

  - IPv6 Border Gateway protocol (IPv6 BGP)

  - Open Shortest Path First (OSPF)

  - IPv6 Open Shortest Path First (OSPFv3)

  [# 571033]

NS1000V

- Deprecated vPath feature

  vPath feature is not supported from 11.1 Release.

  [#TSK0627231]

SSL

- Support for ECC curves in Service Groups

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

At the command prompt, type:

bind ssl serviceGroup <serviceGroupName> -eccCurveName <eccCurveName>

[# 592418]

- Removing RC4-MD5 cipher from the default cipher list

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

  For the updated list of ciphers supported by the NetScaler appliance, see http://docs.citrix.com/en-us/netscaler/11-1/ssl/supported-ciphers-list-release-11.html.

  [# 258311]

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

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

  1. ssl_ctx_tot_enc_bytes: Tracks the number of encrypted bytes.

  2. ssl_ctx_tot_dec_bytes: Tracks the number of decrypted bytes.

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

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

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

  6. ssl_ctx_tot_session_hits: Tracks the number of session hits.

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

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

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

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

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

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

[# 597279, 582601]

- Support for SNI on the Back-End Service

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

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

  [# 471431, 559271, 595785]

- Support for SafeNet Network HSM

  All NetScaler MPX, SDX, and VPX appliances except the MPX 9700/10500/12500/15500 FIPS appliances now support the SafeNet network hardware security module (HSM). A NetScaler ADC used with a SafeNet Network HSM provides FIPS 140-2 Level 2 and FIPS 140-2 Level 3 protection, depending on which SafeNet HSM is being used.

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

  [# 450699]

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

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

  For the updated cipher/protocol support matrix, see http://docs.citrix.com/en-us/netscaler/11-1/ssl/cipher_protocl_support_matrix.html.

  [# 498207]

- Segregation of Certificates According to Type

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

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

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

[# 620923, 623890]

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

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

  [# 606874, 595902]

System

- Capturing SSL Keys during NetScaler Trace

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

  [# 603225]

- Warning about an Unsaved NetScaler Configuration

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

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

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

  [# 626225]

- TCP Hystart Algorithm

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

  [# 603099]

- Dynamic TCP Buffer Management

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

[# 628115]

- Bridge Group Support for Cluster

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

  [# 587548]

- Proactive Support for Hardware Errors

  The Citrix Call Home service automatically generates a support case and uploads the system data to the Technical Support server if a critical hardware error, such as failure of a hard disk drive (HDD), Compact Flash (CF) device, SSL card, or power supply unit (PSU), occurs in a NetScaler appliance on which the Call Home feature is enabled.

  [# 639336, 599891]

- TCP Fast Open Mechanism

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

  [# 358990]

- RDX Error Management

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

  [# 552575]

- Proportional Rate Recovery Algorithm

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

[# 473777]

- Configuring SNMP Audit Log Levels

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

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

  [# 569317]

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

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

  [# 570372, 524621]

Telco

- Port Control Protocol for Large Scale NAT

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

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

  Applications usually send frequent keep-alive messages to the large scale NAT device so that their LSN mappings do not time out. PCP helps reduce the frequency of such keep-alive messages by enabling the

applications to learn the timeout settings of the LSN mappings. This helps reduce bandwidth consumption on the ISP's access network and battery consumption on mobile devices.

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

[# 496807]

- NAT44 Wildcards Static Maps

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

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

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

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

  [# 614784]

- Compact Logging for Large Scale NAT

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

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

  [# 496812]

- Large Scale NAT64

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

-- Load balancing of log servers: Support for load balancing of external log servers for preventing storage of redundant log messages.

-- Minimal Logging: Deterministic LSN configurations or Dynamic LSN configurations with port block significantly reduce the large scale NAT64 log volume.

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

[# 496866]

- Global override LSN parameter removed from L3 parameters

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

  [# 642585]

- Subscriber Aware LSN Session Termination

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

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

  [# 578275]

- Support for SIP and RTSP ALGs for DS-Lite

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

  [# 604029]

- HTTP Header Logging Support for DS-Lite

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

[# 558159, 559227]

- Policy-based TCP Profile

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

  [# 622947]

Known Issues

The issues that exist in Build 47.14.

AAA-TM

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

  [# 621010]

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

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

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

  [# 610091]

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

  Workaround: Escape the angular brackets.

  Example:

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

  Change

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

  To

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

[# 646139]

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

  [# 620845]

Admin Partitions

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

  [# 609367]

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

  [# 620673]

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

  [# 493668, 516396]

Application Firewall

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

  [# 547282]

- The user defined signature objects must not contain "#" character in the name even though the feedback message inaccurately lists it as an allowed character.

  [# 648010]

- 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]

Cache Redirection

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

  [# 591265]

Clustering

- In a Cluster setup, if you create an LACP channel on a secondary node with the interface present on the first node, the channel gets created and runs smoothly but the system displays a configuration error.

  [# 644080]

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

  [# 642947]

DNS

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

  [# 599741]

GSLB

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

  Workaround: You must autosync a non-default partition by using the command-line interface.

  [# 648396]

- On a NetScaler appliance, the default memory allocation is 10 MB per partition. In certain use cases, the allotted memory might not be sufficient for adding the maximum number of entities. You can increase the memory allocation by running the following command:

set partition <partition_name> -MaxMemLimit <limit>

For example, To increase the partition memory allocation to 50MB, at the NetScaler command prompt, type:

set partition p1 -MaxMemLimit 50

[# 614357]

- In the NetScaler GUI, GSLB services related to NAPTR are not listed when you bind the services to the GSLB virtual server.

  Workaround: Bind the GSLB services to the GSLB virtual servers by using the command-line interface.

  [# 648388]

Integrated Caching

- A NetScaler appliance fails multiple times if a cache parameter is enabled during an HA persistency test.

  [# 610085]

Load Balancing

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

  [# 630435]

Monitoring

- If the IP address and port of a dispatcher for a user monitor are set to the IP address and port of a service, and this service is later deleted, the appliance fails if you try to add a service with the same IP address and port.

  [# 618052]

NITRO

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

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

  Upload of collector archive [] failed

  [# 629572]

NS-Orchestration

- During service insertion, one VLAN is acquired from the VLAN range configured in the deployment settings. This VLAN is used for L2 bridge creation for every runtime NIC. Acquired VLANS cannot be reused for other runtime NICs.

  [# 648726]

NetScaler CLI

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

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

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

  Workaround: Use one of the following alternate approaches:

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

  - Use the Net::Telnet library.

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

  [# 346066]

NetScaler GUI

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

  [# 648408]

- Certificate bundles are not supported in cluster setups.

  [# 644199]

- If a policy is bound to or unbound from system global or the priority of the policy is modified, the changes are not reflected automatically. To see the current status, click the Refresh icon at the top right corner of the policy view. After you refresh the view, the policies display their bound status as well as their priorities.

  [# 452669, 502720, 479434, 453597, 391434, 478131, 622724, 573976, 481397, 453555]

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

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

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

  Workaround: Use the NetScaler CLI.

  [# 648293]

NetScaler Insight Center

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

  [# 630276]

- In Security Insight, IP reputation is not displayed in the Violation Category drop-down list in the Application Summary page.

  [# 643629]

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

  [# 543632, 570264, 567628, 565706]

- The period for which data is displayed on the dashboard might not match the selected period. If no data is available for part of the selected period, NetScaler Insight Center shows data from the date on which it started receiving the AppFlow data.

  [# 601474]

- 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]

- In Security Insight, clicking on the Geomap does not show location (Country/Region/City) details.

  [# 645613]

- Upgrading the NetScaler Insight Center as "read only" user throws an "Invalid Session ID".

[# 643011]

- In Security Insight, time slider for custom time duration setting, which is displayed on the dashboard, might not work intermittently.

  [# 630524]

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

  [# 643724]

- NetScaler Insight Center does not report an application-launch failure caused by a user trying to launch an application or desktop to which the user does not have access.

  [# 609604]

NetScaler VPX Appliance

- In NetScaler VPX appliance configured with VMXNET3 network interface, you cannot perform suspend or resume operation.

  [# 644785]

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

  Workaround: Use ixgbe PF driver 3.21.4.3.

  [# 636360]

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

  [# 642495]

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

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

  [# 622392]

- In ESX environment, file transfer from NetScaler to an external connection is stalled when the MTU is changed during the file transfer.

  [# 630639]

- Promiscuous Mode needs to be enabled for VMXNET3 interfaces at the ESX Hypervisor for IPv6 or LACP support.

  [# 641748]

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

  [# 642389]

- The following features are not supported for on SRIOV interface with an Intel 82599 10G NIC on ESX VPX:

  - L2 mode switching

  - Static Link Aggregation and LACP

  - Clustering

  - Admin partitioning [Shared VLAN mode]

  - High Availability [Active-Active mode]

  - Jumbo frames

  - IPv6

  The following features are not supported for on SRIOV interface with an Intel 82599 10G NIC on KVM VPX:

  - L2 mode switching

  - Clustering

  - Admin partitioning [Shared VLAN mode]

  - High Availability [Active-Active mode]

  - Jumbo frames

  - IPv6

- VLAN configuration on Hypervisor for SRIOV VF interface through ip link command is not supported

[# 605846]

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

[# 641498]

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

[# 616044]

- In KVM environment, the NetScaler VPX appliance fails to boot, if you have configured more than 11 vCPUs.

[# 647348]

- In ESX environment, the NetScaler VPX appliance configured with VMXNET3 network interface does not support Auto Negotiation feature. However, NetScaler GUI displays the Auto Negotiation feature as ENABLED for the VMXNET3 network interface.

[# 641256]

Networking

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

[# 643341]

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

[# 638915]

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

[# 625841]

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

[# 629102]

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

  [# 631589]

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

  [# 631592]

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

  [# 590805]

Policies

- If a named expression has the same name of its function, the subsequent use of the function will result in an error. In addition, if you reboot the appliance and use the function in a running configuration, the named expression displays errors thereby leading to a configuration loss.

  The workaround is to rename a named expression by using a prefix or suffix to the expression name, for example myco_func or func_myco.

  [# 637060]

- The command for configuring a content filtering action is being saved in a wrong order in the ns.conf file. Service is a mandatory parameter for adding a add content filtering action, but the add content filter action command is saved before the command that adds the service. As a result, when the build is upgraded, the content filtering action is not configured as required.

  [# 603551]

SSL

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

  [# 628075]

- The number of SSL cards that are UP is not displayed in the 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]

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

  [# 623058, 634017]

- The output of the "stat ssl vserver" command includes the statistics for non-SSL virtual servers.

  [# 627650]

- You can bind an ECDSA cipher or cipher group to a back-end SSL service even though ECDSA ciphers are not supported on back-end SSL services.

  [# 635953]

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

  [# 637230]

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

  [# 604061]

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

  [# 628082]

- Network latency might cause the "./vtl verify" command to fail. This command is used to verify network trust links (NTLs) connectivity between the NetScaler appliance and the SafeNet network HSM.

  [# 623200]

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

[# 628067]

System

- 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 default setting for auto-negotiation is 'OFF', which causes an error if you configure the interface from the SVM.

  [# 598688]

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

  [# 617922]

- If a wildcard virtual server's redirection mode is set to IP (-m IP), the NetScaler appliance cannot forward a TCP connection request to a service bound to that virtual server if the back-end server is down.

  [# 331889]

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

  [# 582087, 587620]

- 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 device (for example, TCP Buffer) in the NetScaler appliance.

  [# 622573]

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

[# 556385]

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

  [# 305196, 393696]

- If an LACP channel is bound to nine or more interfaces and is a member of a tagged VLAN, deleting the channel from a service VM can cause the NetScaler appliance to fail intermittently.

  [# 524320, 630772]

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

  - global maxconn parameter is set to 1.

  - unable to reuse the connection for probing.

  As a result, the transaction fails.

  [# 636416]

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

  [# 632872]

Telco

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

  [# 647630]