

NetYCE 7.0.3 Build_20180215

Release notes

Date: 2018-02-15

Enhancement

Scheduled jobs

On NetYCE environments using multiple servers, each server runs its own scheduler to execute the jobs. The 'Operate - Jobs - Scheduled jobs' tool gives a status listing of all the servers job and allow the operator to manage them.

Since each server uses its own scheduler, the operator needed to switch to the appropriate server to manage the jobs scheduled there. This task has now been simplified by extending the 'Scheduled jobs' tool with all jobs of all schedulers and enabling the operator to manage these jobs from any server.

Service-type Add-Port(s)

Two new service-type commands were added that allow the creation of ports for a node. These are ADD - PORT(s) - TYPE_FROM_NAME and permit the operator to pass interface names that include the interface-type.

The interface-type part of the interface-name is validated to determine the required port type to create. This validation is very flexible in this determination since you can either specify the type using the NetYCE Port_class abbreviations (Fa, Gi, Te, Ma, Lo, Po, etc) or the longer generic names (FastEthernet, GigabitEthernet, Port_channel, etc), but also the vendor specific interface-names (100GE, xe, Bridge-Aggregation, etc).

The examples below demonstrates this behaviour, each format produces the same interface for a HP C5 device:

```
Gi/0/1/2
Gi00/01/02
Gigabit_ethernet0/1/2
GigabitEthernet/00/01/02
```

Existing service type commands that include the port-type (e.g. ADD - PORT - GIGABIT_ETHERNET) also support the port-type sepecification but is optional. When included in the port-name, the type

overrides the type in the command.

SUPERNET service_types

The Service_types had so far only very limited support for managing a Clients IP supernets. To allow manipulation of the supernets, the single existing service-type call was replaced for eight new calls:

Name call	Description
ADD - SUPERNET - IP_SUPERNET	add new ip-supernet range to client. Format value as "<net-address>/<prefix>". Set ip-plan using Assign-Supernet-Ip_plan or use API custom var "ip_plan"
ADD - SUPERNET - IP_PLAN	add new ip-supernet range to client from pool of free Supernets using numeric ip-plan ID. Or use API custom var "ip_supernet" to specify the supernet
ASSIGN - SUPERNET - IP_PLAN	assign the numeric ip-plan ID to the supernet-alias. Supernet and plan prefixes must match. Cannot change a supernet plan-id which has active subnets
ASSIGN - SUPERNET - DNS_DOMAIN	assign the value to the supernet Dns_domain attribute
LOCATE - SUPERNET - IP_SUPERNET	find the supernet within the client matching the "<address>/<prefix>" value
LOCATE - SUPERNET - IP_PLAN	find the (first) supernet within the client matching the Ip_plan value
DELETE - SUPERNET - FREE	remove the supernet from the client IF there are no subnets left. Place supernet in free pool
DELETE - SUPERNET - COMPLETE	remove the supernet from the client after DELETING all its subnets. Place supernet in free pool

Device file transfers

To simplify the upload and download of files between network devices and the NetYCE system, new Scenario commands have been added.

The **file_get** retrieves a file from the NetYCE system and transfers it to the device.

The **file_put** retrieves a file from the node and transfers it to the NetYCE system.

The transfer uses SFTP or TFTP depending on the available support from the vendor and the connectivity available. The implementation of these commands is at its introduction only available for the Juniper junos vendor. Others will follow shortly.

Mpls / Vrf type attribute

To assist creating 'categories' of Mpls vrfs and Node vrfs, the attributes 'Mpls_type' and 'Vrf_type' have been added to their respective object types.

The values for these attributes are defined in the 'Lookup' using the 'Mpls_type' variable of the 'Translation' class. As many entries as desired can be created using this lookup variable. These

values will be presented as drop-down menu lists from which one value can be selected.

IPv6 Extra buttons

The Create IPv6 Subnets form has been modified with four extra buttons in order to get more control over the start and end of the range of subnets you want to assign.

Customer special update: RN3 Core vlans

The custom tool “RN3 Core - Add/remove vlan” could not modify the RN3 core vlan in the migration scenario where a local node from a different location (changed ip-ranges) was to be attached to an existing EVPN connection. Changes to the evpn 'identifier' and the used relation were required.

Changes to the evpn 'identifier' and the used relation were required to visualize these extra connections and schedule their related change jobs. The embedded scenario was updated to the version 7 syntax. An updated Relation definition for “RX_core_vlan” is required.

IP Subnet Plans Layout

Tweaked the layout for the IP subnet create forms a bit so that it's more clear where every option belongs.

Service type records multiselect

Multiselect now works for the records of a service type. You can select multiple records using the shift-key. The edit form disappears, but you can delete multiple records at the same time using this.

IPv4 subnet range delete extra check

Deleting an IPv4 plan now first checks whether there are any subnets using this plan. If there are: this will be displayed to the user as a warning. This way the user won't have to go all the way back to the build menu to check. This also applies when deleting IPv4 subnet ranges.

IPv6 subnet range grid columns

Column names for IPv6 subnet ranges were rearranged to be more in line with IPv4 subnet ranges.

Customer logo

Customers that want to include their corporate logo in the NetYCE application can do so using the gui. In the "Admin - System - System status" tool, using the 'Full report' option, the section 'Current license info' now includes to buttons that allow the system manager to upload a logo image file or restore the original NetYCE logo.

It will be displayed in the upper left corner above the menu in a 118 by 30 pixel space. The image will be downscaled to fit this size and should support a transparent background for best effect. The preferred file format is '.png'.

IPv6 subnet range delete extra checks

Deleting an IPv6 plan now first checks whether there are any subnets using this plan. If there are: this will be displayed to the user as a warning. This way the user won't have to go all the way back to the build menu to check. This also applies when deleting IPv6 subnet ranges.

Change

Job tools selection simplification

Most of the job tools require the user to make a number of selections before the job can be submitted or an action is executed. This selection usually involves selecting a set of clients, then a set of nodes, each time clicking an 'add' or 'next' button.

To simplify this selection process, double-clicking is now added in many such places allowing the user to select and 'add' or select and 'next' using but a single mouse action.

Chrome mouse scroll behaviour

For those users that access the NetYCE front-end using the Google Chrome browser a new feature of that browser has been put to good use: overscroll. The normal behaviour is that when the top or bottom of a scrolling object (e.g. data grid) is reached, the parent object (e.g. the page) will then start scrolling. In many applications resulting in an annoying experience.

With this feature in place, only the object where the mouse is currently pointing at will scroll. So while scrolling a grid, the end of the grid will not result in the page scrolling. And while scrolling the page, the page keeps scrolling until the mouse enters the grid.

Currently only the latest Chrome browser version (63) supports this feature. Firefox is expected to follow in the upcoming releases.

Vendor configuration items

With a growing number of configuration items for Vendor-types in the Lookup we decided this information was more efficiently managed and retrieved using a dedicated database table.

The 'Vendor_type' and other vendor related items in the 'Vendors' class were removed from the Lookup and can now be accessed using the "Admin - Custom data" menu as the YCE table "Vendors". Access requires 'system' level privileges.

The Vendors table now includes setting for the file-transfer options for each vendor. The use of these settings will gradually be implemented in each of the vendor-modules.

Fix

Long cyphers

Recently the various password fields in the database were made significantly longer allowing strings to be stored up to 200 characters. But since password fields are generally sorted in encrypted format, some overhead is consumed limiting the unencrypted length somewhat.

The implementation of the database column encryption features of NetYCE are described in the [Encryption](#) article.

This implementation uses checksums to validate the encrypted string before attempting decryption. As a consequence of the much longer encryption strings, the checksums also became longer, which was not taken into account and caused strings of more than 63 characters to fail unencryption.

This situation has been corrected, providing transparent database encryption of sensitive information once more.

Junos indentation

Evaluating the commands in the Command job 'Evaluate' function now properly ignores comments in the indentation if they contain '{' or '}'

Copy port-setup

When migrating some existing forms to a newer technology shared files with other forms were removed along with the obsolete. Unfortunately, one of these shared files was not obsolete since is still being used in the Ports 'Copy' function where port configurations can be copied from one switch to another.

This broken functionality has been restored.

Template function arguments

The templates use over 20 functions to perform dynamic calculations, conversions or formatting. Most of these functions accept multiple arguments that are specific to the function. A problem was identified when using multiple arguments that were quoted but separated by commas as is mandatory, but also with spaces for readability.

These extra spaces were parsed incorrectly and were concatenated with the arguments resulting in incorrect (or none) output. The parsing of the function arguments has been reworked to prevent these problems.

Port_name correction

When adding ports using the web GUI, some ports were assigned incorrect (internal) Port_names.

If the slot was blank (no slot-id) and the port-type is one of the non-Sys ports, the port_name would be set using the port-id without a padding '0'. So would a GigabitEthernet port be named like 'Gi5' and not using the convention 'Gi05'. This behaviour was corrected.

Node selection in tools

An odd node selection issue with several operational tools would occur when the client-code and the node-name were the same.

In those cases the inclusion of the node-name would result in the list of select nodes to be extended with the list of nodes from the client with that same name. Some nodes would be included twice due to this issue.

By separating the node selection by node name from the node selection by client-code could this issue be resolved.

Database auto-increment validation

Due to a still unknown cause, one table could not accept additional data rows because the maximum value of the autonumber function would be exceeded. Although the highest column value that is assigned this autonumber was still at least a billion short of the maximum, the internal table administration had somehow concluded it had reached the maximum. A simple manual adjustment of the auto-increment value fixed the issue, but because the issue could not be explained, a workaround is now implemented.

The daily maintenance of log files and databases is extended with a function that validates and corrects any table using an autonumber where the auto-increment value does not match the maximum value of the column using the autonumber value.

Unauthorized redirects

Instead of being redirected to the login page when you visit a vmenu page you do not have access to, you simply see a message.

Custom attribute password

Password custom attributes used to disappear when you clicked on them. Now, when you click on them they are visible.

Scenario empty string

String replacement in scenarios now also works for variables containing an empty string.

IE11 login

Internet Explorer 11 in 'EDGE'-mode can be used with NetYCE but is not recommended. The performance using IE is very poor compared with the Firefox and Chrome browsers. Other caveats concern the improper rendering of some elements weakening the user experience.

Recent customer experiences however demonstrated that logging in using IE11 was not possible. This was caused by a slight difference in the request header. This situation is corrected.

Inverted Checkpoint's enabled / disabled state

When developing the Checkpoint vendor module there was only a account provided which defaulted to bash mode, therefore the vendor module first disabled when it logged in to get to the 'clish' (default shell on Checkpoint devices). As it turns out this is not default behaviour because the norm is to log in in clish mode. The necessary changes are incorporated in the State_actions table and in the vendor module.

checking for OS file presence in ios dir

The path we use to check for the presence of a OS file in the ios dir was incorrect, this has been fixed.

Daily database maintenance

The NetYCE database uses several tables to collect logging information on the users' actions, the API actions and the node configuration sessions. With high levels of mutations, these tables grow quickly. A daily process ('log_maint') removes the oldest of each of these tables as per the settings in the Lookup.

Some customers experienced problems with this maintenance where these tables became corrupted and consequently interrupted the master/master replication. The underlying cause proved to be dependent on the table defragmentation that is part of the daily maintenance. Without table defragmentation the removal of old records would not result in freeing disk space and so is deemed essential.

This table defragmentation is executed by a database command that proved very sensitive to a lack of disk space and is and prone to corruption or even crashes. To alleviate these issues, the daily maintenance process has been reworked to consider each of these known issues and take appropriate action. It now also repeats the entire process up to three times to ensure the integrity of the database. Should spontaneous table corruption occur (likely when disk space is low), this same process will automatically attempt the repair.

Job logs engineer permissions

Engineers now have permissions to access the job logs form.

HTML files caching

The NetYCE web front-end adopted the 'appcache' mechanism of controlling the browser cache in version 6.3. The support of the appcache did not see the expected adoption of this mechanism in other browsers, but instead its support dwindled away over time.

In an earlier release we reverted back to include versioning information with some files to force cache renewal. This behaviour has now been extended to other file types which should result in fewer cases where caching is preventing proper display or behaviour after a version update.

Node VRF Dropdown

The new vrf dropdown in the node vrf form is now properly sorted.

From:
<https://wiki.netyce.com/> - **Technical documentation**

Permanent link:
https://wiki.netyce.com/doku.php?id=maintenance:releases:7.0.3_20180215

Last update: **2024/07/03 12:31**

