

NetYCE 7.0.4 Build_20180528

Release notes

Date: 2018-05-28

Enhancement

Custom data filter

In the custom data forms, you can now also search for strings containing a '.'-character

Infoblox integration

A number of changes have been incorporated in the Infoblox IPAM/DHCP integration.

- Infoblox 'containers' are never removed when modifying an IPAM tree in 'update' mode. This reduces the risk of inadvertently removing very large subtrees.
- When modifying an Infoblox IPAM/DHCP tree in 'update' mode and the shadow 'history' is not available, the IPAM/DHCP tree will be rebuild completely by reverting to 'renew' mode. This reduces the risk of an undetected out-of-sync IPAM tree and allows the operator an alternative method of controlling explicitly renewing trees.
- Infoblox extended-attributes may now contain spaces in their name. A change in the XML formatting was required to support this.
- Infoblox extended-attributes will now be included in the determination if a 'network' or 'range' object needs to be updated or not. This previously required a renewal of the IPAM/DHCP tree.

New port forms

Both the node and template port forms have been completely rebuilt using the new style. The same goes for the node port forms, and all sub-forms below it, including the port update forms.

The functionality remains the same, on top of the additions of port slots. These can now be used to specify different port layouts for each slot on the same device.

These changes are part of an effort to update the entire NetYCE web-interface using the Angular-Mojolicious technologies.

Service type multiselect

Multiple service type records can now be deleted at once. Use shift and/or control keys to select mutiple records.

IPv6 Subnet extra fields

The IPv6_subnet table has gotten a number of new columns to allow for point to point and four corners connections.

New vendor modules

Additional NetYCE vendor modules have been created and integrated. New is the 'Cisco_XE' to support the XE-family products that run the IOS-like 'XE' operating system.

Also the 'Cisco_Wism' module has been ported from LabYCE to support the integrated WiSM module for wireless services.

And finally, a separate Checkpoint management vendor module has been created, 'Checkpoint_Mgmt', to control the Checkpoint hypervisor that runs virtual Checkpoint firewalls. The existing 'Checkpoint' vendor module continues to control the Checkpoint firewalls using the 'Gaia' operating system.

These extensions now brings the total number of NetYCE vendor modules to 21.

Logs new form

The old Logs-forms have been rewritten to the new forms style.

Lookup

The old Lookup-form has been rewritten to the new forms style.

Jobs integration with NCCM

The NCCM will now integrate with the NetYCE jobs. When jobs make a change to the configuration, backups of the device configuration are made before and after the job change. Any differences found in these configurations with the current NCCM state will be appended to the NCCM.

The NCCM provides a history of all configuration changes detected during the preriodic configuration polling, and now with the job-integration, includes the changes made by the NetYCE tooling by each job. The configuration changes caused by NetYCE jobs will be included in the NCCM regardles of the device NCCM polling status. These NCCM entries will include the operator-id, the job-id and the job description.

This capability is as yet incomplete since 12 of the 20 vendor modules have the required modifications. The Juniper, HP, Cisco and Huawei vendor modules are now ready for testing.

NCCM Config diffs tool

The NCCM configuration differences tool has been reworked at several levels. First of all its performance. With larger number of devices, a large number of polls per device, or with very large configurations the loading times showed exponential delays.

These performance issues have been addressed in several ways. Firstly, the NCCM database structure was altered to split the actual configuration data from the NCCM administration. Secondly, the rebuilding of the configurations to be compared is now performed at request time and no longer in bulk at device selection. Thirdly, the process of detecting and hiding of security sensitive data in the configuration (passwords and keys) was redesigned to be executed primarily at collection time, not reporting time.

And finally, the front-end tool was redesigned to simplify the selection of the two configurations to compare. Now, the user selects each of the configs in two steps. First the year and month, secondly the date and time of the config change along with its cause (job-id, user, job description). Additional features include options for uncensored view (for higher user levels), side-by-side or in-line view, un-truncated view and download buttons. The views are by default truncated to 5000 lines to prevent overly long loading times due to browser limitations when rendering configs of 100.000+ lines. The 5000 line limit customer defined using a Lookup tweak.

Tasker foreach logs

Logging for a foreach statement in a scenario should now start each iteration with the variable, and the content of that variable.

Service type multiselect

Multiselect now also works for full service types, not just the records.

Custom data NMS tables are now editable

When using the 'Admin - Custom Data' tools, the key columns of the various tables could not be edited. This limitation was created to prevent inconsistencies in the database. When the key column is an autonumber, as in the case of Nms_systems, this behaviour prevented adding new records to the table.

The restriction on editing key values has not been lifted, but made less stringent so that these keys can be edited and added when using the NMS database.

Port default button color

The default port color has now been set to #d1fffc

IPv6 subnet frontend new fields

The IPv6 subnet create and edit forms now have support for four corners and point to point topology. The frontend fully supports these fields, the backend does not yet do this everywhere.

IPv6 Subnet Description

IPv6 subnets now also have a description field. This functions the same as for IPv4 subnets.

Tasker grep function

Scenarios have been extended with the '**grep**'-command. Similar to the 'like' command, but instead it compares each item of a list with the given regex string. It returns a list of all matching items.

Restore NCCM config

The scenario command 'config_restore' is now integrated with the NCCM. Using various options to specify what NCCM config to select, the configuration of a device can be restored by uploading it to the device and making it the startup configuration.

When this configuration file is to be activated by a reboot, the scenario will have to include the 'reboot_node' command as appropriate.

The NCCM selection options include 'previous', 'last', 'poll', 'marked', <jobid>, and <timestamp>. The 'marked' option refers to the NCCM GUI 'config diffs' tool where an operator can manually select the NCCM config to restore and mark it as such.

Hardware storage device

When transferring files between network devices and the NetYCE server, the location of a file often needs the storage device (eg: flash:, disk0: bootdisk:) to be included in the command. This storage device name was defined per vendor and could not be changed by the user.

Some vendors have products using several storage devices or changing names per model. To accomodate user-definable storage locations, the Hardware form now offers an option to change the storage device used for each hardware model.

Initially only the Cisco IOS vendor module support this setting, but others will follow shortly.

File transfer NAT-address

When NetYCE executes jobs involving file transfers (and all configuration backups are), it is the device that initiates the transfer by connecting the NetYCE server. However, in situations that the NetYCE server must be connected on a different IP-address than the server address due to NAT translations, these transfers will fail.

The workaround so far has been the modification of the NetYCE server configuration file, but that had to be repeated every software update. As of now the NAT-address can be explicitly configured by running the `net_setup.pl` script which will prompt for it.

Any subsequent software update will during the execution of the `yce_setup.pl` script create the NetYCE server configuration file with the transfer ip-address set to the NAT-address. If no NAT-address is configured, the transfers will continue to use the NetYCE server IP-address.

New relations form

The relations form has now been changed to the new style, including syntax highlighting. Another new addition is that you can move your cursor in the editor to a table name, allowing you to select its columns.



Change

Service-type logging

Although the Task-logs give very detailed information on the execution of service-types started using the NetYCE API, no logging was preserved when executing service-types using the Web-GUI.

The logging of service-types actions have now been extended to log both API and manual service-type execution in the regular 'Logs' (`action_logs`). Since the Task-logs still maintain a detailed log of the service-types, the amount of data per log item has been reduced to 16KB.

Service-type custom-variables

Service-types use variables to specify values for locating and creating NetYCE database objects. When executing service-types using the API, these variables can be extended to refer to custom variables defined in the API request.

When service-type variables cannot be resolved using custom variables or aliases, the default behaviour was that the variable value would be left unchanged. This resulted in assigning the NetYCE database object attribute values referring to an API custom name like “(Core_node)”.

This default behaviour has now been changed to set unresolved service-type variables to blank ("") values instead. To revert the the original behaviour the Lookup tweak "Keep_unresolved_st_variables" can be modified.

Relations check

When executing a NetYCE software upgrade one of the finalizing actions is the verification of the existing Relations. Relations are SQL statements that refer to the NetYCE database tables and columns. Because software updates frequently include database alterations, customer-created relations could have become unusable.

To catch the Relations that use tables or columns that cause SQL errors, all Relations are tested after each software update or database restore (which also trigger database updates). In some environments, the execution of this Relation verification can take several minutes. Since there is no benefit of having the operator to wait for the verification results, the verification process has now become a background action that will report in a Task-log entry when it is done.

The operator will have to take down the Task-id listed in the update/migration log and verify the results in the Task-log manually.

HFC Combine

A new wizard for splitting cmts node has been added to the Service config menu. This can be used to move CMTS hosts for a selection of CPEs. You can search per CMRS hostname and select multiple at a same time, and at the end a job will be kicked off that handles all changes. This job takes the scenario called CMTS_combine

Vendor modules rename

The vendor modules have been reorganized resulting in some vendor name changes. The 'CI6' and 'CI8' modules are now named 'Ciena_CI6' and 'Ciena_CI8' respectively. Likewise, the 'F5' module is now named 'F5_BigIP' and the 'Avaya_SW' now better reflects the supported device family using the 'Avaya_ERS' name.

Lookup and Logs forms

The 'Lookup' form and the 'Logs' forms have been replaced with new forms with the same functionality but with the newer look-and-feel of the Angular-based forms.

These changes are part of an effort to update the entire NetYCE web-interface using the Angular-Mojolicious technologies.

Cisco Nexus management ports

The Cisco Nexus now has support for the 'management' ethernet port.

Existing Nexus devices in a Fabric-path environment that previously used the GigabitEthernet port 0 in a blank (") slot, are migrated to this new management port-type.

Deleted Avaya 8K vendor support

Avaya 8K vendor support is deleted

IPv4 and IPv6 DHCP forms

IPv4 and IPv6 subnets now have the option of being extended with a set of DHCP parameters. By default no DHCP set is created, but in their respective forms you can now add these when if you set the Dhcp_type to something other than "none". Deleting a lpv4 or lpv6 subnet Dhcp set is done by setting the Dhcp type back to "none".

The DHCP parameters can be assigned manually using the GUI or automatically using Service-types.

IPv6 table and column_names

Our IPv6 tables and columns were all prefixed with the string "IPv6_". This was quite confusing, because it was based on a ruleset different from the norm. With this fix all IPv6 fields are prefixed with "lpv6_" instead, leading to more consistency in our names.

As a censequence, any existing Relations using the lpv6 tables and parameters will fail to locate these references. Patch number 18042302 will automatically correct these for the existing relations. When creating new relations, this change should be kept in mind.

This change has no impact on Templates or Scenarios since these are designed to be case insensitive.

Auth_permissions names

Our permissions were saved in the auth_permissions table. This was confusingly the only table in the database with only lowercase names. The names are now fixed to Auth_permissions to be more in line with our other naming policy.

This change has no visible impact on front-end or any of the database objects.

lpv6 service-types

Several of the lpv6 service-types have been reworked to behave similar to their lpv4 counterparts.

New lpv6 service-type commands have been added to support lpv6 custom subnets manage the

Ipv6 subnet attributes which now include the same range of Ipv6 point-to-point, four-corners, gateway and loopback addresses as do the Ipv4 subnets, albeit with different names.

These Ipv6 attributes, similar to Ipv4, can be modeled using ipv6-subnet-plans and modified using service-types. The service-type values for these Ipv6 addresses support ipv6-offsets and ipv6-addresses, including negative offsets to assign addresses from the end of the Ipv6 range.

Fix

Juniper peek timeout

Jobs on Juniper devices occasionally time-out on cli actions. The cause was identified as parsing the response too soon after the command, resulting in faulty behaviour and timeouts.

Tasker next-statement

The next-statement in scenarios has been fixed.

Custom data par vals

Editing custom attributes (Par_vals) in the custom data forms now works again, where it was previously impossible to select one that had multiple keys (for example with nodes).

Tasker Heredoc Variable Fix

Fixed a bug where in a scenario, a variable for a function containing a heredoc argument wasn't substituted.

IPv4 plan delete segments fix

Fixed the bug where the "delete" button for segments didn't work in the IPv4 plans form.

Tasker relation fix

Fixed a bug where in the scenario, the -n option for the 'relation'-command wouldn't accept runtime variables.

IPsec GRE login redirect bug fixed

There was a bug in the IPsec GRE form that redirected you to the login form when trying to save an entry. This bug has now been fixed.

Fixed node bootimage reset button bug

There was a bug in the Nodes form, where clicking on "Reset", instead of resetting the boot image it just removed the value. This is fixed now.

DHCP jobs reported as 'Aborted'

DHCP jobs would, when they take longer than a few minutes, be reported as 'Aborted' - although the job was still running properly. This was caused by the process monitoring that failed to locate the process name in the servers process list. By setting the process name to the job-id the process monitor can now detect these jobs and manage their status correspondingly.

Job logs operator filter

When filtering job logs by operator, the Infoblox jobs would not be included in the displayed results. The issue is caused by Infoblox jobs logging their actions using the operator full name instead of the userid. By correcting the Infoblox logging operator name and by extending the job logs filter to use both userid and full name this issue was resolved.

IPv4 subnet standard create

The standard ipv4 subnet create form now refreshes after adding a subnet, updating a current count of assigned subnets.

Custom data footprints

The Footprint_map tables are now editable in the custom data forms

Scenario operator permissions

Operators can now create and edit scenarios

Patch management

Release updates from NetYCE consists of two parts, software distribution files and patches. The patches make modifications to the database and the server configuration. Because of multi-server

architectures we maintain patchlevels for the database and each of the servers. Patches are executed in sequence for every version, allowing for seamless upgrades between any NetYCE release.

These patches are checked at update installation time and when restoring a database. A problem was identified when restoring a database from an older revision when the server patchlevel exceeded the restored database patchlevel. Depending on the version number, some patches would be mistakenly skipped.

The patchlevel administration has been modified to prevent this situation from occurring.

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

Permanent link:
https://wiki.netyce.com/doku.php?id=maintenance:releases:7.0.4_20180528

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

