

Scheduled jobs

Overview

The 'Scheduled jobs' tool is found in the **Operate - Jobs** menu. Its purpose is to review the currently scheduled jobs and manipulate the scheduling of these jobs.

Since all NetYCE (font-end) systems have their own scheduler, the jobs are collected from all servers in succession, a process that is listed at the top of the page. Servers that fail to connect will be disregarded when manipulating the jobs until the tool is re-opened from the menu.



Using this tool, the user can perform three kinds of actions on the submitted jobs: manipulation, filtering, and approvals. Whether these options are available for him and for which of the jobs depends largely on the user's permissions and the job-permissions setup. Generally, options not selectable are not available, although in some cases the non-concurrence is shown as an error message.



Job manipulation

The generic Job manipulation options (Cancel, Now, Suspend, and Resume) operate on the set checkboxes in the 'Select' column. If the checkbox is missing, the user has no auditor permissions on this job. After selecting the appropriate jobs the user can perform one of those four actions on all these jobs.

See the article on Job-configuration for details on [Job configuration](#) and the role of auditors.

However, not all actions are available given the current job state. If a job is in the 'suspend' state (waiting approval), it is not possible to have it executed Now or changing it to Suspend state. Trying to do so results in an error message:



The various states the NetYCE scheduler maintains has its own set of possible entry and exit actions associated. These are summarized in a 'state-machine':



The colours for each state in the diagram correspond to the colours used in the Job list.

See the article [Queue Operation](#) for the behaviour and configuration of the queues.

Job Filtering

The user can limit the jobs listed by applying filters. There are filters for User name, Change-id, Job state, Job-id and Server name. The default is not to filter any jobs.

These five filters can be used in combination without restriction. After selecting the appropriate filters, press **Find** to apply them. When no jobs comply, the list will be blank and not show the "No jobs" message that indicates no jobs are present.

The JobID filter accepts a regular expression in its text box. Normally it is sufficient to type part of the JobID and click **Find**. An entry like `19_001` locates jobID's like `0119_0017` and `0119_0018`, but also `1219_0010`.

Do not use wildcards like `*` and `?` unless you are familiar with the regular expressions (where `.`, `*` and `.` are similar).

All filters can be defaulted simply by using the **Reset** button. The full list is then shown.

Each time the job list is modified, the actual list is updated and shown. It is reported directly by the scheduler by each of the actions. An update of the list can explicitly be requested using the **Refresh** button.

Job Approvals

NetYCE supports approval workflows where operational jobs require a second pair of eyes to confirm correct submission of network change. This second pair of eyes can be configured to be a peer-level operator or a higher-level engineer.

The job owner cannot approve his own jobs, but he is allowed to cancel them when in 'pending', 'scheduled' or 'suspended' states.

Submitted jobs requiring approval will be put in 'pending' state. Once approved, the job is in the 'scheduled' state. Pending jobs not approved before their scheduled time will be placed in the 'suspended' state. Scheduled and pending jobs can be cancelled by the job owner.

Per job-creating tool (like `command-jobs`) the approval requirements are configurable. For each tool the user-levels that are allowed to do the approvals (e.g. designers and managers) can be set and the number of jobs that can be submitted by a user before an approval is required for each level can be defined independently (e.g. operators may schedule up to 5 jobs unapproved and engineers 20). See the article on [Job configuration](#) for details.

The job approval configuration includes the operator levels (a list of levels) of the users that can approve the job. These are dependent on the operator level of the job-owner. In each case, the level refers to the effective level of the user, i.e. the operator level for the node's client-type of the user-group.

Email notifications of the following events will be sent:

- to the job owner when a job is rejected (*note 1*)
- to the job owner when a job expires and is suspended

- to the job owner when a job is canceled
- to the owner's peer group when a job is submitted for approval (*note 2*)

To prevent mailbox flooding, email notifications will be sent at an appropriate interval (2 min) combining all jobs for the same user, grouped by change-id.

note 1) When a job is rejected or cancelled, a reason for the rejection is prompted for.

note 2) The list of mail addresses that should be informed on pending jobs is the list of emails in the user-group where the operator is a member of.

Notifications can be enabled or disabled per job-type (or in the inheritable default settings) for three states: pending (for approval), canceled (removed by auditor) or suspended (when expired waiting for a queue slot).

See the article on Job-configuration for details on [Job configuration](#)

Job locks

Jobs that are submitted for the same node at the same time could cause configuration dependencies or conflicts when committing these configurations. To prevent this, NetYCE uses a global locking mechanism on jobs.

Before launching a job the scheduler(s) will check if the node is locked by another job. If so, the job is put in a 'waiting' state. It will not start while the lock is in place. Jobs occasionally can take a long time and to prevent waiting jobs to be delayed overly long, a maximum lock time of **10 minutes** is imposed. That means that a job locked-out by another job may start anyway after 10 minutes. It in turn will take the lock and may keep it for 10 minutes. Of course, when a job finishes the lock is removed.

Should many have jobs been scheduled for the same node, only one will be allowed to run at the time and all will be executed in their original scheduled order. However, to prevent a job at the back of the queue being started hours after the original scheduling time, a limit is imposed on the time a job can be kept 'waiting'. After 1 hour (3600 seconds) in a the 'waiting' state, the job is suspended. This setting is controlled by the scheduler queue definition as described in the [Queue operation](#) article.

The locks are maintained globally and are shared between all schedulers ensuring that a job on one server will honour the lock set on that node by another server.

Note that the lock is set for the 'anchor' node for which the jobs was scheduled. If a job scenario connects one or more other nodes these connections will not be subject to locking. The locking permits the scenario to start (or not) based on the selected node when scheduling the job.

Finally, if the job-locking by node is undesired, it can be disabled - globally! - by setting the Lookup Tweak "**Sched_ignore_locks**". The job activation mechanism will then function as in previous releases.

See [Lookup Tweaks](#) for details on this tweak.

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

Permanent link:
https://wiki.netyce.com/doku.php?id=menu:operate:job_status:scheduled_jobs

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

