Table of Contents

Compliance scenarios

This is an overview of the four main scenarios for netYCE compliance: through polling groups, a config change, through the xch API and multiconfig compliancy.

The purpose of this page is to give an overview from start to finish of the whole compliance process.

Schedule via polling groups

1. The user creates policies, with rules and conditions (or imports them) and links the correct node groups to them.
2. The user creates a polling group.
3. The user adds nodes or node groups to the polling group, and sets the Next poll interval and the Max retries.
4. The nccmd daemon wakes up every five minutes. It updates the schedule times for all polling groups. This schedules the nodes in the newly created polling group for NCCM.
5. The nccmd daemon looks at the list of nodes that need to be polled for nccm for its server.
6. The nccmd daemon takes this list of nodes, and performs an nccm check for each of them. If it can't do all nodes with in 5 minutes, the remaining ones will be picked up by its next cycle (or by a daemon on another server).
7. Each node whose configuration has been changed and who has been flagged for a compliance upon config change (a Task_mask value of 1), will be scheduled for a compliance check, its schedule time will be set to the current time.
8. The nccmd daemon wakes up every 5 minutes, and takes a look a the list of nodes that still need to be checked for compliance for its server.
9. The nccmd daemon takes this list, and performs a compliance check on each of them. If it can't do all nodes with in 5 minutes, the remaining ones will be picked up by its next cycle (or by a daemon on another server).
10. Depending on the the policy's signalling settings, on basis of the compliance check's result a signal is sent. This can be either a syslog, trap, email or API call (multiple are also possible). These can be used to take action.

Schedule via node config change

1. The user creates policies, with rules and conditions (or imports them) and links the correct node groups to them.
2. A node changes its configuration and sends out a syslog message
3. yce_events retrieves this message, filters it out of all other irrelevant syslog messages and schedules the node for an NCCM poll for the upcoming daemon cycle.
4. The nccmd daemon wakes up every 5 minutes and looks at the list of nodes to be polled for nccm on its server.
5. The nccmd daemon takes this list of nodes and polls each of them for compliance. Any nodes that aren't polled within the time limit of five minutes are released back and can be picked up in the next cycle, or another daemon.
6. Each node whose configuration has been changed and who has been flagged for a compliance upon config change (a Task_mask value of 1), will be scheduled for a compliance check, its schedule time will be set to the current time.
7. The nccmd daemon wakes up every 5 minutes, and takes a look a the list of nodes that still

need to be checked for compliance for its server.
8. The nccmd daemon takes this list, and performs a compliance check on each of them. If it can't do all nodes with in 5 minutes, the remaining ones will be picked up by its next cycle (or by a daemon on another server).
9. Depending on the the policy's signalling settings, on basis of the compliance check's result a signal is sent. This can be either a syslog, trap, email or API call (multiple are also possible). These can be used to take action.

Schedule via xch

1. The user creates policies, with rules and conditions (or imports them) and links the correct node groups to them.
2. The user sends a message to the yce_xch to denote whether he wants the node to be scheduled for an nccm poll (nccm_run) or a compliance check (cmpl_run).
3. yce_xch receives this request and schedules the node either for nccm or compliance, depending on the instructions.
4. The nccmd daemon wakes up every 5 minutes and looks at the list of nodes to be polled for nccm on its server.
5. The nccmd daemon takes this list of nodes and polls each of them for compliance. Any nodes that aren't polled within the time limit of five minutes are released back and can be picked up in the next cycle, or another daemon.
6. Each node whose configuration has been changed and who has been flagged for a compliance upon config change (a Task_mask value of 1), will be scheduled for a compliance check, its schedule time will be set to the current time.
7. The nccmd daemon wakes up every 5 minutes, and takes a look a the list of nodes that still need to be checked for compliance for its server.
8. The nccmd daemon takes this list, and performs a compliance check on each of them. If it can't do all nodes with in 5 minutes, the remaining ones will be picked up by its next cycle (or by a daemon on another server).
9. Depending on the the policy's signalling settings, on basis of the compliance check's result a signal is sent. This can be either a syslog, trap, email or API call (multiple are also possible). These can be used to take action.

F5 multiconfig compliancy

1. The user creates policies, with rules and conditions (or imports them) and links the correct node groups to them.
2. The user creates multiconfig rules. These are special types of rules that do not have any conditions and only make sense for node groups with four or less nodes. This is not mandated by the GUI. If you do have a node group of more than four nodes, only the first four nodes are used.
3. The user uses one of the above ways to schedule a node for nccm and compliance. Often this means that all of the nodes in its node group are also scheduled for compliance.
4. The nccmd daemon wakes up every 5 minutes and looks at the list of nodes to be polled for nccm on its server.
5. The nccmd daemon takes this list of nodes and polls each of them for compliance. Any nodes that aren't polled within the time limit of five minutes are released back and can be picked up in

the next cycle, or another daemon.

6. With the correct settings (see the above three scenarios) the nccmd daemon then schedules the nodes for a compliancy check for the upcoming daemon cycle.
7. The nccmd daemon wakes up every 5 minutes, and takes a look a the list of nodes that still need to be checked for compliance for its server.
8. The nccmd daemon takes this list of nodes and performs a compliance check on each of them. If any multiconfig rules are present, the configs of the other nodes in the node groups are retrieved from the NCCM. Important: no nccm polls are ran on these nodes. Also note that the same process happens for each node in the node group.
9. Depending on the the policy's signalling settings, on basis of the compliance check's result a signal is sent. This can be either a syslog, trap, email or API call (multiple are also possible). These can be used to take action.

- When an nccmd poll is unsuccessfull, the counter for failed polls is raised with 1. Once this is equal to the polling group's maximum retries, it is automatically disabled. It can be enabled again in the NCCM Selection form in the GUI. Setting the max retries to '0' will always keep the node enabled.