Multiconfig compliance

The Multiconfig compliance is intended for a group of two to four nodes, whose configuration needs to be equal at all times.

The Multiconfig compliance uses it own Rule types as they will compare the node's new config to each of the current configs in its node group. A number of restrictions apply:

- The node group cannot contain more than four nodes. Any more will be ignored.
- The node group must contain nodes of the same vendor-type
- You can provide multiple node groups for a policy, provided that the total number of nodes does not exceed four. All nodes will have their configs checked against each other
- It is not advisable for the rule's policy to be run at config change, because this would mean a non-compliance every time one of the nodes changes its config. Instead, we recommend Multiconfig compliance to be run periodically, for example daily after office hours. You can use policy schedules for this.

The Rule type has to be set to "multi-config", and its Vendor type to the vendor required. Rules of this type do not have any conditions.

Like with regular rules, you can also select part of the configs to be compared to each other. You can use Rule start and Rule end to find either config blocks that start and end with the lines respectively (regex supported), or search the config for the lines in between and including the Rule start and Rule end lines that you supplied, depending on which selection option you checked in your rule (Search based on lines or Search based on config blocks).

If no Rule start is defined, the daemon will take the whole config to compare.

If a rule start is specified, the daemon looks for blocks to match these conditions. Any it finds will be compared to the other configs in the policy's node group. If it can't find any, it's automatically compliant.

F5 Bigip

The multiconfig compliance functionality was created for the F5 BIGIP vendor module and has been extended to function for other vendor modules as well. The F5 is still a special case however. For these BIGIP configurations some extensive pre-parsing takes place that will remove from the configuration tree all unreferenced segments before comparisons are made. This greatly improves the compliance results given that these configurations can be over 400,000 lines.

In the case of a F5 BIGIP node, all orphans will be logged in the rule's report, to serve as a guide for the operator to help clean them up.

NOTE that a number of blocks are meant to be orphaned. For now these are hard-coded, however if desired we can make this user-generated.

There are a number of special lines that are different in each config, for example timestamps and cyphers. These values will be cleaned out for comparison, and substituted with the text "CLEANED". There are a number of special cases to note here:

- hostname
- management-ip
- base-mac

These define the values pertaining the node itself, and therefore cannot be equal. However these values also can appear in other blocks, and these will also be cleaned. To test these lines for compliance, it is best to create a regular configuration rule that tests for these values specifically.

Do note that F5 codes contain code segments, including comments. For now, we do not filter away these comments, and the indentation of these code segments. This could change in the future if there is desire for it.

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

Permanent link:
https://wiki.netyce.com/doku.php?id=guides:reference:compliance:cmpl_multiconfig

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