

NMIS8 and fping or just ping

NMIS has multiple poll cycles, every 5 minutes an SNMP poll cycle is run, this is also when services are polled. If fpingd.pl is being used it will poll nodes every ~1 minute by default. It is recommended to use fpingd.pl when using NMIS, which will use ICMP to poll nodes with higher frequency than the regular NMIS poll cycle. During the fping poll cycle, the escalation process is also run, this helps to ensure that notifications are sent out with minimal delay, and enables highly granular escalation policies.

NMIS will try to make sure it's own fpingd.pl is running every time a collect or update is run if configured to use it (detailed below). If fpingd is not enabled or the loss value for the node to ping cannot be found in the fping status file nmis-fping.nmis or nmis-fping.json (on newer systems) then NMIS will run a ping of it's own. This code can be found in nmis.pl in sub runPing.

**Do Not Edit nmis-fping.json**

The fping status file nmis-fping.nmis, or nmis-fping.json (on newer systems) is automatically created and updated by fpingd. This file should NEVER be hand edited. At the very least hand editing will be overwritten by fpingd and at the worst a bad edit could result in disabling fpingd completely.

fpingd

fpingd is enabled by default, the config setting (in Config.nmis for NMIS8) to control this is:

```
'daemon_fping_active' => 'true',
```

In NMIS9 the configuration option is nmisd_fping_worker

```
'nmisd_fping_worker' => 'true',
```

**Not Using DNS?**

If you are NOT using DNS make sure to set daemon_fping_dns_cache in /usr/local/nmis8/conf/Config.nmis to false, the default value is true.

fpingd runs a constant cycle of pinging all nodes updating the NMIS event table, sleeping and then doing it all over again. Every 10 poll cycles the Nodes.nmis file is re-read to look for changes.

Nodes are split into groups of size fastping_node_poll and then fping is run on these sets, one after another. After all the pings are complete fpingd takes the results and updates the NMIS event table, when that is complete NMIS escalations are run and then the daemon will sleep for fastping_sleep + a random value between 0 and 10.

fastping_timeout is an important configuration variable when investigating ping loss, as any ping value higher than this number will be considered loss, 300 ms may be too low if you have a high latency network and are experiencing high ping loss on nodes that appear pingable.

Here is a table of the config variable and what it controls, these are all set in Config.nmis where you can find the default values:

Configuration Variable	Default (if not set at all)	Controls
fastping_timeout	300	Controls the -t setting for fping which is: -tn Initial target timeout in milliseconds (default 500). In the default mode, this is the amount of time that fping waits for a response to its first request. Successive timeouts are multiplied by the backoff factor.

fastping_packet	56	Controls the -b setting for fping which is: -bn Number of bytes of ping data to send. The minimum size (normally 12) allows room for the data that fping needs to do its work (sequence number, timestamp). The reported received data size includes the IP header (normally 20 bytes) and ICMP header (8 bytes), so the minimum total size is 40 bytes. Default is 56, as in ping. Maximum is the theoretical maximum IP datagram size (64K), though most systems limit this to a smaller, system-dependent number. If this is set to < 24 it will be set to 24.
fastping_retries	3	Controls the -r setting for fping which is: -rn Retry limit (default 3). This is the number of times an attempt at pinging a target will be made, not including the first try.
fastping_count	3	Controls the -C setting for fping which is: -C Similar to -c, but the per-target statistics are displayed in a format designed for automated response-time statistics gathering. For example: <pre>% fping -C 5 -q somehost</pre> <pre>somehost : 91.7 37.0 29.2 - 36.8</pre> shows the response time in milliseconds for each of the five requests, with the "-" indicating that no response was received to the fourth request.
fastping_sleep	60	The number of seconds fpingd will sleep before running another ping cycle + random(10)
fastping_node_poll	200	The number of nodes to ping in one invocation of fping. We strongly recommend that this value is set to no more than 250; otherwise you risk exceeding the linux command line length limits which will cause incorrect operation of fpingd.pl.
fastping_interval	1	The interval between sending ping packets. Some firewalls like to have larger gaps between packets, setting fastping_interval to 80 and fastping_target_interval to 100 could be used to pacify these devices.
fastping_target_interval	1	The interval between ping packets to one target. See fastping_interval.