

Installing the NVMe Host Module Patch on Compute Nodes

This section provides instructions on how to compile and install the NVMe host module patch required for each compute node. This step is independent of KumoScale software but is needed to complete a deployment.

At the appropriate time you are ready to set up your orchestration environment, you will need to apply a patch to all of your compute nodes. This patch addresses an issue where commands get stuck while Host NVMe-oF controller is in a reconnect state. This section provides a description of the patch with instructions for installing on Ubuntu and CentOS.

The NVMe controller enters into the reconnect state when it loses connectivity with the target. It tries to reconnect every 10 seconds (default) until it successfully reconnects or until it times-out. However, the host may not enforce the time out and will continue to try to reconnect for a longer period or indefinitely.

The patch provided resolves this issue and is based on the NVMe community patch. To fix this long delay, we:

- Introduced a new session parameter `fast_io_fail_tmo`. The timeout is measured in seconds from the controller reconnect, and any command beyond that timeout is rejected. The new parameter value may be passed during `CONNECT`. The default value of -1 means there is no timeout, which is similar to the current behavior.
- Added a new controller `NVME_CTRL_FAILFAST_EXPIRED` and respective delay work that updates the `NVME_CTRL_FAILFAST_EXPIRED` When the controller is entering the `CONNECTING` state, we schedule the `delayed_work` based on the failfast timeout value. If the transition is out of `CONNECTING`, we terminate the delayed work item and ensure `failfast_expired` is equal to false. If the delayed work item expires then `NVME_CTRL_FAILFAST_EXPIRED` is set to true.
- Updated the `nvme_fail_nonready_command()` and `nvme_available_path()` functions to check the `NVME_CTRL_FAILFAST_EXPIRED` controller flag.

For multipath (function `nvme_available_path()`), the path will not be considered available if the `NVME_CTRL_FAILFAST_EXPIRED` controller flag is set and the controller is in the `NVME_CTRL_CONNECTING` state. This prevents commands from getting stuck when available paths have tried to reconnect for too long.

To install the patch, follow the instructions appropriate for your environment:

[How to Install the NVMe Host Module Patch on Compute Nodes under CentOS](#)

[How to Install the NVMe Host Module Patch on Compute Nodes under Ubuntu](#)