



Contribution ID: 244

Type: **not specified**

## Towards Real-Time NVMe Monitoring (nvme-top)

*Friday 12 December 2025 10:00 (45 minutes)*

Monitoring NVMe devices and paths in production is currently limited to static snapshots via `nvme-cli`. While sufficient for basic inspection, this model falls short in NVMe-oF (fabrics) deployments, where path conditions can change dynamically due to fluctuating network latency, congestion, or link failures. Administrators troubleshooting fabric multipath environments often need continuous visibility into path state, ANA status, queue depth, link speed, and error counters, but today they are forced to repeatedly run commands or rely on ad-hoc tooling.

This motivates the idea of `nvme-top`, a tool providing real-time monitoring of NVMe fabrics paths and devices, similar in spirit to `iotop` or `top`. The goal is to give administrators a continuously updating view of device and path health, enabling faster detection of link degradation, imbalances in multipath I/O, or transient failures.

Today, `nvme-cli` builds a static in-memory tree from `sysfs`. This works for one-off queries but does not update dynamically. To enable real-time monitoring, we need mechanisms to refresh the topology at regular interval. For an initial proof-of-concept, a simple polling-based refresh (e.g., once per second) may be sufficient to demonstrate the value of a real-time monitor. Longer term, community input will be needed on whether kernel-assisted notification (e.g. `inotify/fanotify` or `uevents`) mechanisms are desirable and maintainable.

This BoF aims to:

1. Identify the most useful attributes for monitoring in real time (e.g., ANA status, path state, queue depth, error statistics).
2. Explore trade-offs between polling and kernel notification mechanisms.
3. Discuss expectations around responsiveness, overhead, and presentation (e.g., text vs. rich TUI like `btop`).

During the session, we also plan to show pre-built mockups of a potential `nvme-top` dashboard to make the discussion concrete and gather input on usability and presentation. By focusing on NVMe-oF multipath monitoring, this discussion seeks to shape a tool that improves troubleshooting and operational visibility in fabric-based NVMe deployments.

Link: <https://github.com/linux-nvme/nvme-cli/issues/2904>

**Primary author:** SHROFF (IBM), Nilay

**Co-author:** WAGNER (SUSE), Daniel

**Presenters:** WAGNER (SUSE), Daniel; SHROFF (IBM), Nilay

**Session Classification:** Birds of a Feather (BoF)

**Track Classification:** Birds of a Feather (BoF)