White Paper
High-Performance Computing | Network Performance Monitoring
July 2022

Find and Remediate Network Bottlenecks in Your HPC Cluster

Identify network performance issues and tune workloads for optimum performance in HPC clusters, using Intel® Deep Insight Network Analytics Software

Authors

Babu Peddu
Product Marketing Manager,
Switch & Fabric Group

Mateusz Podlasinski
Cloud Software Development Engineer,
Switch & Fabric Group

Thomas Fletcher
Cloud Software Development Engineer,
Switch & Fabric Group

Table of Contents

Executive Summary

As more organizations in almost any industry you care to name turn to artificial intelligence (AI) and advanced data analytics to make better business decisions, high-performance computing (HPC) is also on the rise. A recent Technavio survey revealed that the HPC market is expected to grow by USD 27.15 billion from 2021 to 2026, progressing at a compound annual growth rate (CAGR) of 11.31%.[1] But as HPC clusters increase in size, network performance becomes a crucial factor in overall cluster health. Legacy network monitoring tools are not designed to manage networks of this size and speed.

Intel's silicon portfolio enables an end-to-end programmable network, which in turn enables real-time telemetry and analytics with Intel® Deep Insight Network Analytics Software. This approach helps users to detect and resolve issues as they happen, which is not possible with traditional network troubleshooting and telemetry approaches that rely on historical data and manual investigation.

This white paper provides an overview of the benefits of using Intel Deep Insight Network Analytics Software for network monitoring and management. It also illustrates some real-world examples of how this unique solution can help you quickly and precisely pinpoint the cause of performance issues like latency or dropped packets. Once you know what the problem is, you can then take remedial action, such as rerouting network traffic or adding parallel switches. Using Intel Deep Insight to troubleshoot the HPC cluster network can help you get the most out of your HPC investment.

Business Challenge

It is tempting to consider only computing resources when looking at performance bottlenecks in a high-performance computing (HPC) cluster. After all, HPC clusters must process TBs of data, often in near real time. But the reality is that all this data must be delivered to the cluster and back to clients, which makes network performance nearly as critical as compute performance. To complicate matters, today's HPC clusters often run in a multi-cloud environment, which introduces a new level of complexity that raises network visibility challenges.

Let's consider an example: Your HPC cluster has several hundred nodes, and the network is dropping packets somewhere, but you do not know where. All you know is that dropped packet counters keep increasing, indicating that there's a problem. But you can't determine which flow, which application or which server is causing the problem or why. You could try sending ping probes every second, but if the problematic event is short-lived—say 20 microseconds—the every-second ping probes are unlikely to capture the event.

Today's growing applications of HPC to new use cases and evolving multi-cloud adoption—combined with an explosion of IoT devices and data, increasingly distributed workloads and continued digital transformation—create a perfect storm that mandates a new approach to managing HPC network performance. Network architects need more than rudimentary tools like ping probes and traceroute techniques. Today's HPC clusters need to unlock the power of advanced analytics and intelligent network programming to provide deep insights into the network infrastructure. Only by doing so can HPC reach its full potential to drive innovation in every industry.

Solution Overview

To tune an HPC cluster network for optimal performance, you need to know four ground truths—that is, information acquired by direct observation rather than by inference.

- How did the packet get here? (The sequence of devices it visited on its path.)
- Why is the packet here? (Identify the set of rules that it satisfied at every device along the way.)
- How long was the packet delayed? (Determine the time it spent buffered in every switch—to the nanosecond—that is 100,000 times more precise than a ping probe.)
- Why was the packet delayed? (Visualize the packets, flows and applications it shared each queue with.)

However, traditional network monitoring tools cannot provide this level of granularity.

Intel has developed a solution that uses programmable components to provide a new level of network monitoring and management capabilities. This solution is Intel® Deep Insight Network Analytics Software, which uses in-band network telemetry (INT) to enable end-to-end network visibility that can remediate issues quickly and optimize packet flow across a wide variety of uses cases, including HPC.

As shown in Figure 1, Intel Deep Insight takes data from a variety of INT sources to create actionable insights.

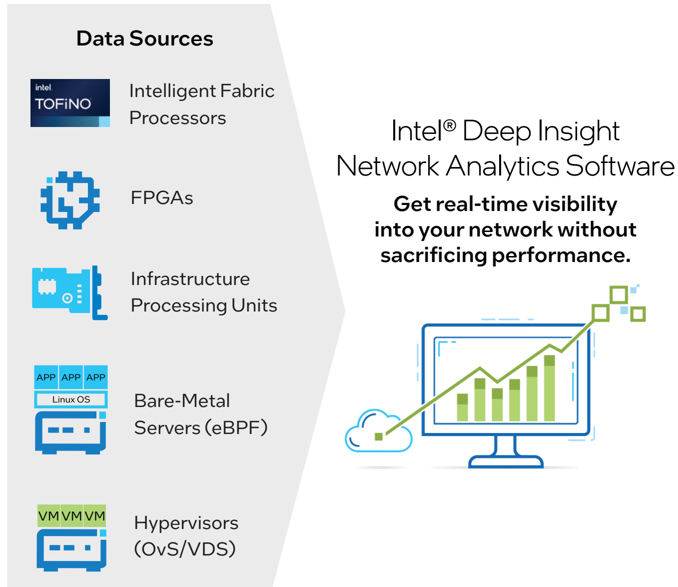

Figure 1. Real-time visibility into network performance is based on a number of diverse sources.

Network architects can use Intel Deep Insight to perform the following tasks to optimize network traffic in an HPC cluster:

- **Analyze.** Observe every packet from every switch and router in the network and server, in-band at line rate.
- **Detect and report.** Enable an intelligent, flexible triggering mechanism to detect and report events in real time, with nanosecond accuracy.
- **Interpret and pinpoint.** Run diagnostics on industry-standard servers and pinpoint—in real time—the many conditions that impeded packet flow.
- **Track and monitor.** Identify every packet drop to know when, why, where and to whom it happened.
- **Baseline and track.** See the path and latency for every packet in the network and server.

Putting Intel Deep Insight to Work: Two Test Cases

While the above discussion describes generally what is possible with Intel Deep Insight in HPC environments, the power of the solution is realized by seeing it in action. Therefore, Intel performed two tests that used Intel Deep Insight to pinpoint network issues that would have been time-consuming and tedious to find using traditional ping probes and rudimentary dashboards—time that HPC workloads don't have to waste.

Test #1: Diagnosing Slow Performance

To simulate standard HPC network behavior, we ran an All-to-All traffic test across 111 nodes.[2] As shown in the top half of Figure 2 (bottom of the Events & Anomalies window), this test revealed more than 2,000 packet drops. When we looked deeper into the test behavior, Intel Deep Insight's network traffic visualization capabilities showed that the queue on the switch was building up and releasing quite slowly (bottom half of Figure 2). Such performance issues can hamstring an HPC cluster.

To improve performance, we ran a second iteration of the test, where we set PSM3_FLOW_CREDITS=8 (the default value is 64). This modified setting enabled us to throttle the hosts and alleviate network pressure. The second iteration was 2.5X times faster than the first iteration, with no dropped packets and a much faster queue drain (see Figure 3).

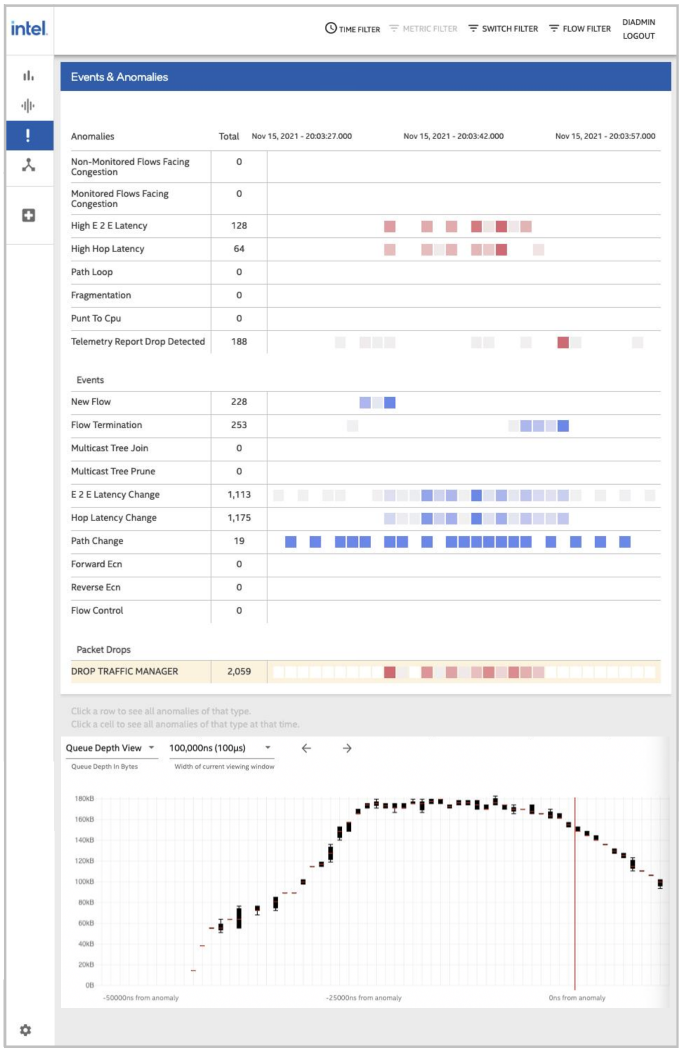

Figure 2. Intel® Deep Insight shows events and anomalies, such as inconsistent network reliability, as measured by dropped packets and queue behavior.

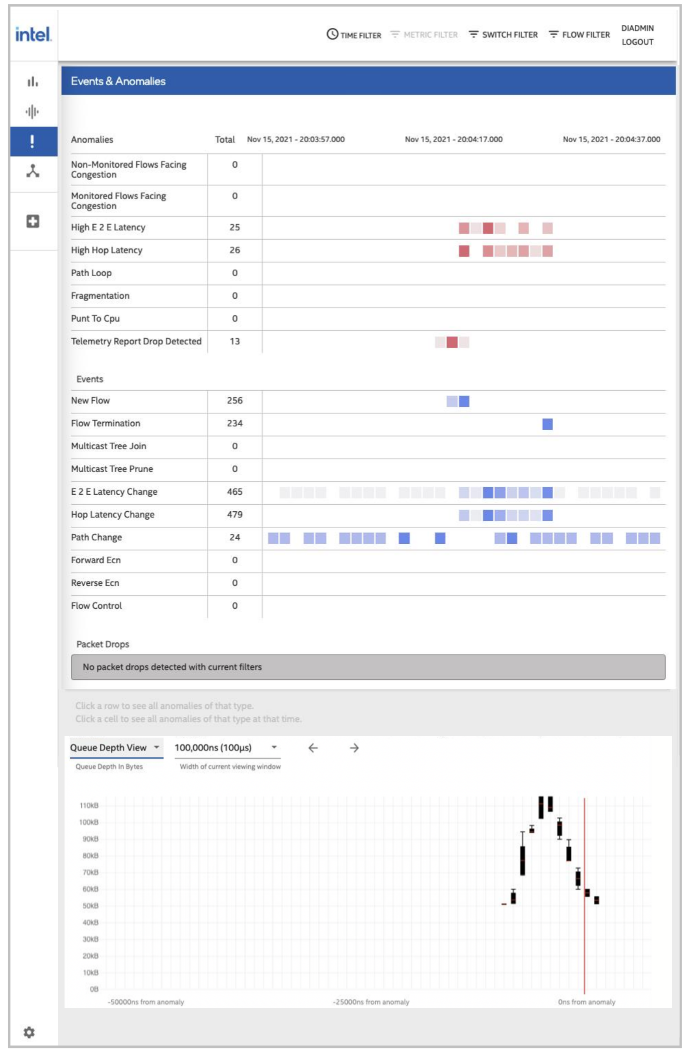

Figure 3. Modifying a network traffic parameter enabled us to eliminate dropped packets and increase network performance.

Test #2: Diagnosing Performance Degradation

For the next test, we ran pairwise traffic, from N nodes to N nodes. We increased N until we saw a performance drop-off in the benchmarking output. In this test, we observed that performance continued to improve as we increased the number of nodes, up to N=4. However, with N=8 pairs, the performance decreased by half. As illustrated in Figure 4, Intel Deep Insight shows that the problem is associated with Spine 1001. When we moused over the problematic spine, we could see a maximum latency of 67us, while the others showed a latency around 700ns (a difference of about 100X). Although not directly related to this particular test, we can also see this spine was connected to all leaf nodes, but traffic is only flowing to/from two of them.

We used Intel Deep Insight to take a closer look at that underperforming spine (see Figure 5). The highlighted row reveals that the congestion is on a single port, 43/1. (Note that the latency for this port is 67.22µs, compared to other ports around 0.5µs.)

With Intel Deep Insight's rich visualization capabilities, we were easily able to pinpoint that two pairwise flows are using the same inter-switch-link (in this case, the link from spine 1001 to leaf 8222). This is a classic example of how once we identify the problem, the solution is straightforward. Changing the routing to make one of these two flows go through a different spine, or using another parallel link between these two switches, would improve the performance for this exact set of nodes.

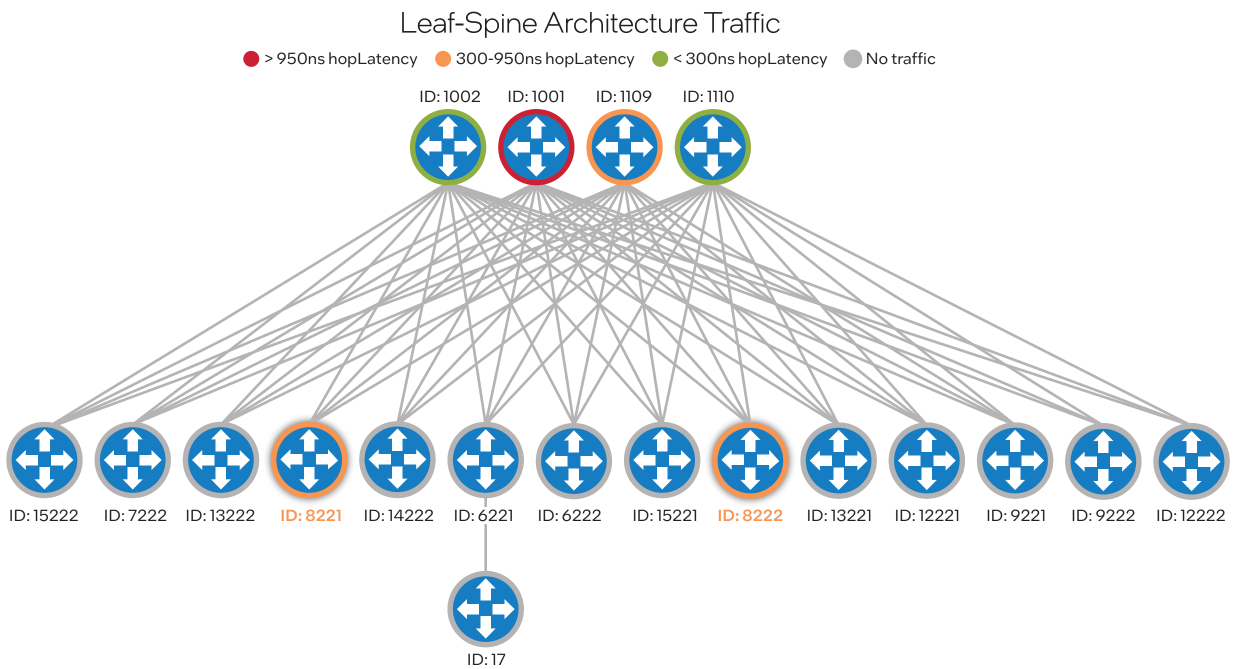

Figure 4. An analysis of the leaf-spine architecture reveals that all traffic is being routed through only two leaf nodes.

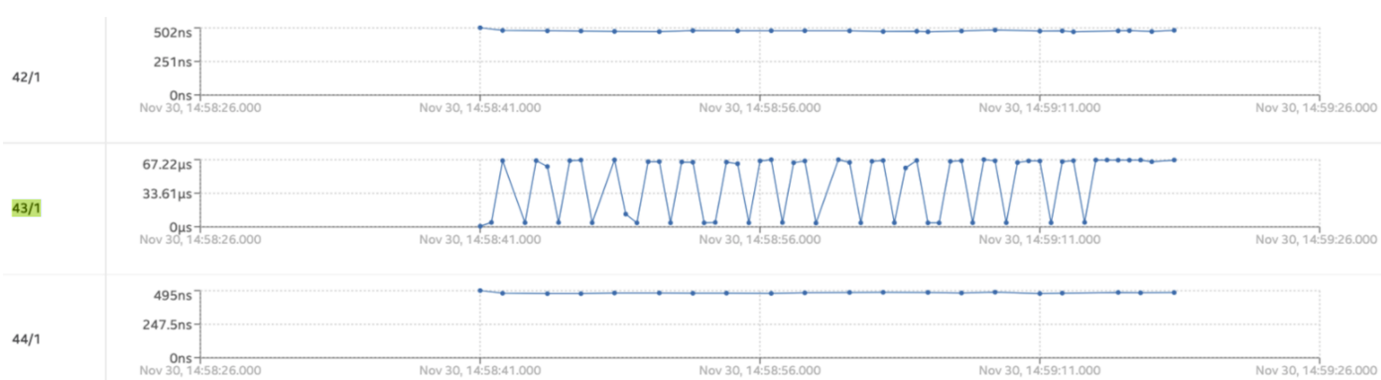

Figure 5. Using Intel® Deep Insight, you can drill down to the port level to better understand traffic congestion.

We can look at the events with high end-to-end (E2E) latency, which are shown in Figure 6. Doing so helps us determine which flows are being delayed and pinpoint the node in the network that is contributing to the latency. As shown in Figure 6, we can see two flows on the congested port (43/1 in Figure 5) are being delayed:

- from 10.48.9.1 to 10.68.4.1
- from 10.48.7.1 to 10.68.1.1

Specific occurrences

	Source IP	Destination IP	Ingress Timestamp	Egress Timestamp	Hop Latency	Queue Occupancy (bytes)
⌄	10.48.9.1	10.68.4.1	Nov 30, 2021, 14:58:34.103070188	Nov 30, 2021, 14:58:34.103074388	4200	51120
⌄	10.48.9.1	10.68.4.1	Nov 30, 2021, 14:58:34.104548111	Nov 30, 2021, 14:58:34.104551403	3292	50880
⌄	10.48.7.1	10.68.1.1	Nov 30, 2021, 14:58:34.107500256	Nov 30, 2021, 14:58:34.107504340	4084	55120
⌄	10.48.9.1	10.68.4.1	Nov 30, 2021, 14:58:34.108221313	Nov 30, 2021, 14:58:34.108225656	4343	51120

Figure 6. Investigating events associated with high end-to-end latency can help further diagnose network issues.

Intel Deep Insight also allowed us to see what was in the queue on the problematic switch to identify flows that are noisy or exhibit other issues. As shown in Figure 7, the “green” and “purple” flows have large (4,120-byte) packets, while the other flows all have packets under 350 bytes.

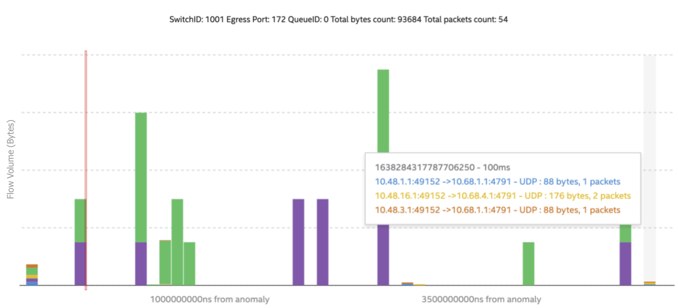

Figure 7. Using data from Intel® Deep Insight, over-sized packets can be identified, providing an explanation of unusually high latency.

Conclusion

As HPC clusters grow—both in size and in popularity—it becomes increasingly important to maximize network performance. INT from Intel Deep Insight Network Analytics Software can provide the intelligence, performance, visibility and control necessary to tune HPC cluster networks to provide predictable low latency. And when issues do arise, INT has the means to identify and mitigate the problem before it gets worse. From traffic congestion to oversized packets to dropped packets, the granular, intelligent data provided by Intel Deep Insight can help organizations meet their HPC goals and drive innovation.

Learn More

You may find the following resources helpful:

- In-band Network Telemetry Detects Network Performance Issues white paper
- Intel® Deep Insight Network Analytics Software
- Intel Deep Insight Network Analytics Software video
- Intel® Tofino™ Intelligent Fabric Processors
- In-band Network Telemetry 0.5 specification
- P4.org
- Open Networking Foundation
- High Precision Congestion Control white paper

Find the telemetry solution that is right for your organization. Contact your Intel representative or visit intel.com/deep-insight.

[1] “High-Performance Computing (HPC) Market – 49% of Growth to Originate from APAC,” https://www.prnewswire.com/news-releases/high-performance-computing-hpc-market---49-of-growth-to-originate-from-apac--driven-by-increasing-utilization-of-big-data-analytics--technavio-301535574.html

[2] Testing by Intel as of November 2021. **Testing methodology.** Software: Intel® Deep Insight Network Analytics Software v4.2, OS = Linux 4.18.0-147.8.1.el8_1.x86_64, MPI v3.1.
Run a 111-node “All-to-All” benchmark, targeting the SONiC plane
Command used: mpirun -np 111 -ppn 1 -hostfile myhosts -genv FI_PROVIDER=psm3 -genv PSM3_NIC=irdma-cvl1 -genv PSM3_ALLOW_ROUTERS=1 $I_MPI_ROOT/bin64/IMB-MPI1 Alltoall -npmin 999 -msglen 4mbX10.