Solution Brief

virtana

Benefits of Performance Hardware Probe Monitoring vs. Software Switch Monitoring for Data Fabric Observability

Introduction

The network switch integration is an agentless software integration that utilizes storage and network data collected from SMI-S and SNMP Management Information Base (MIB) to gather switch performance and link error statistics in a non-intrusive manner. These critical switch statistics are correlated with other system-wide metrics and presented within the intuitive Infrastructure Monitoring (VirtualWisdom) interface as entities.

Our hardware-based SAN Performance Probe is the most advanced, high capacity full line-rate data inspection and analysis device available. It inspects, processes, and analyzes every frame or packet header in real-time. It captures the true, unaltered I/O profile of of the actual application traffic, detecting application performance slowdowns and transmission errors by measuring every I/O transaction from start to finish.

Summary

There are significant differences between the hardware-based Performance Probe using Traffic Access Points (TAPs) and the software-based switch integration for data fabric observability. The quick summary is that hardware-based monitoring helps prevent more problems, finds problems faster, sees problems that software-based monitoring cannot, and helps optimize costs through right sizing and balancing SAN resources.

There are three key differences in hardware-based performance monitoring:

- **Sees things at a lower level so it finds more problems and finds them faster.** Metrics like frame errors, code violation errors, excessive aborts, SCSI reservation conflicts, and excessive cancelled transactions all indicate problems in the SAN infrastructure, and are useful in predicting problems that have not yet impacted the application.
- **Sees Exchange Completion Time (ECT).** ECT is the measurement of how long it takes to complete a command like a read or write operation. These operations can be negatively impacted by incorrectly set queue depths, slow draining devices, buffer credit starvation, unbalance workloads, transaction size, or device impatibility. This is a critical metric that enables faster problem detection, workload balancing, and optimized CAPEX.
- **Sees things with greater granularity.** Since software-based monitoring uses averages, it sometimes misses problems that are detected using the wire-data provided by hardware-based monitoring. This means that your application users may experience latency issues hidden by averages, and your storage team may be blind to the problem. Since hardware-based monitoring sees every I/O, it sees things missed by software monitoring.

If finding most problems is good enough, or performance is never an issue then software monitoring is good enough. If your applications are business-critical and you are measured on availability and performance, then you need the wire data provided by hardware monitoring.

To quote Gartner: “Machine data extracted from the nodes of an IT system will continue to provide information about the IT system state and behavior. However, wire data gleaned from the movement across the packets across network links ... provide the core input to any availability and performance analysis.”

Reports

Software switch monitoring provides port-level health and utilization statistics for the SAN so you can observe how much traffic has gone through a port during the 5-minute polling interval, and by applying logic based on the direction of the traffic you can guess if it's read or write traffic.

Adding hardware-based Performance Probe monitoring gives you additional details along multiple dimensions:

- From a workload characterization perspective, it definitively tells you whether each IO is a read or a write. It also gives you the count of read and write I/O per second (as opposed to just MB/s) and their sizes.
- Provides latency data, telling you how long each read or write operation took to complete and whether any delay was imposed by the array. It records the response time and size of every read and write operation for display in a histogram chart, and provides second-by-second summaries of the IOPS, IO size, and response time metrics.
- Enables insight into the data going across the wire so you can report on the number of concurrent operations at the port level, or the granularity of an individual Initiator-Target-LUN relationship. It tracks the low level Fibre Channel primitives including Buffer Credit utilization in both directions, at sub-microsecond granularity.

Analytics

When combined with hardware monitoring, the following analytics provide additional visibility into performance and workload over software monitoring:

Balance Finder

Works with software monitoring but adding the Performance Probe allows for the clasiification of traffic in NPIV environments where there are multiple physical devices using a single link.

Event Advisor

Works with software monitoring for the metrics that are collected by the software integration. Adding the Performance Probe increases the number of Entity/Metric combinations Event Advisor sees. For slow drain or flow control use cases, adding the Performance Probe gives you the ability to identify flow control issues in both directions and shows you the duration of these events.

Trend Matcher

Works with software monitoring but adding the Performance Probe gives you a much broader set of use cases and matches, including IO type, size, and rate, as well as concurrency.

Queue Solver

Does not work with software monitoring as it requires wire data to analyze response times and queue depths. These metrics are only available with the hardware Preformance Probe.

Seasonal Trend

Works with software monitoring for the metrics that are collected by the software integration, but the real power lies in identifying seasonal behavior of response times and workloads, which requires the Performance Probe.

Discovery and Dependency Mapping

Software monitoring discovery and dependency mapping is limited to those entities that are discovered by the software integration. Adding the Performance Probe provides visibility into additional entities like ITL conversations and storage ports and lets you build application entities based on ITLs. A common use case is isolating DB instances or anything that requires LUN-level visbility.

Event Correlation

The following alerts rely on hardware probe data and cannot be used:

- Bad SCSI status
- Exchange Performance
- Link Buffer-to-Buffer Credits
- Link Transmission Errors
- Performance Probe Bandwidth
- Queue Depth
- Seasonal Trend Deviation - no response times, IOPS, IO sizes, etc.

Value of Software Monitoring

The key benefit of software switch monitoring is the continuous SAN data fabric monitoring and alerting it provides to help you avoid upstream problems caused by physical layer issues, congestion and high utilization, and to plan future capacity requirements.

Key Metrics/Use Cases

- Multi-path verification: Find issues like failed or misconfigured or redundant links.
- Top/bottom talkers: For servers, storage ports, and ISLs (assuming the environment is healthy, which can be confirmed via the hardware probe).
- Class 3 Discards: The switch receives a frame and is unable to pass it along to the next point on the path to its destination.
- CRC Errors: A frame containing a bad CRC checksum is received by a switch indicating that a portion of the frame or its data has been corrupted and must be resent.
- Loss of Synch: Two devices lose the synchronization signal between them for a period.
- Loss of Signal: Light on a receiving port drops to a point too loow for the data to be considered valid.
- Link Failure: Loss of Sync or Loss of Signal persists beyond the Receiver/Transmitter Timeout value.
- Link Reset: Port needs to force a reset of the link, forcing the receiving device to also reset.

Value of Hardware Monitoring

Hardware monitoring provides the best performance analysis possible. Hardware monitoring lets you prove conclusively whether the SAN is the cause of performance problems. Hardware monitoring provides real-time latency tracking and alerting to avoid issues before they impact the appllication. The queue depth optimization it provides solves many common performance issues, and link optimization helps you optimize and control CAPEX.

Key Metrics/Use Cases

- Frame Errors: A frame with an embedded CV or other bit level error. This will result in a CRC Error.
- Excessive Aborts: Generation of Abort Sequence FC Frame (ABTS) by a host to terminate an exchange.
- SCSI Reservation Conflicts: When the SCSI protocol uses the reservation to allow exclusive access to a LUN.
- Excessive Cancelled Transactions: Exchanges cancelled due to SCSI and FC link events.
- Queue Depth/Pending Exchanges: On storage arrays: the max number of SCSI exchanges that can be open on a storage port at a given time. On HBA: specifies how many exchanges can be sent to a target or LUN at one time.
- Exchange Completion Time: Measurement of how long it takes to fully complete specific types of commands from the initiator to the target to the LUN.
- Storage port and LUN demand: To determine proper fan-in ratio; reducing the number of ports to reduce costs at the same time; improving performance through better load balancing.
- Target Busy/Queue Full: The target cannot complete a command at that time, i.e., there is no room for a new ITL conversation.