



# Maximize the All-Flash Data Center

Lenovo Gen6 Fibre Channel

---

June 2017



[Lenovo.com/systems](http://Lenovo.com/systems)

---

## Introduction

Growing numbers of storage architects are recognizing the benefits of an investment in flash-based data storage. Just a few years ago, flash usage was reserved for the most demanding, performance-sensitive applications, while today it is being deployed for general use with primary data throughout the data center. In fact, the adoption rate of solid state arrays has accelerated to the point that Gartner predicts the market will eclipse hybrid- and HDD-based storage by 2020, or possibly sooner.<sup>1</sup>

It is well known that flash offers considerable cost savings over disk-based storage arrays by reducing rack space requirements and energy consumption. However, flash also dramatically reduces Input/Output (IO) latency, improving host server CPU efficiency. This reduces the overall number of servers needed and provides potential savings on expensive software application licenses based on CPU or core count. With performance gains driving direct economic benefit, it is clear that enabling flash storage to perform at its highest level of efficiency will allow organizations to reap the greatest economic benefit and return on their investment.

To maximize the performance value of their flash storage investment, administrators must consider not only the storage, but also the connectivity to the host servers. Lenovo ThinkSystem Gen 6 Fibre Channel maximizes the performance of flash-enhanced storage, even when connecting at speeds below the maximum 32 Gbps data rate. In addition to faster throughput speeds and increased Inputs/Outputs per Second (IOPS), Gen 6 Fibre Channel provides enhanced monitoring and diagnostic capabilities that enable visibility into network latency levels and IOPS, from the storage level all the way to the Virtual Machine (VM) level. These functions ensure the fast and consistent performance of Fibre Channel, as well as greater visibility into network performance to prevent issues—from the host servers down to the storage—from affecting operations.

## Why Lenovo Gen 6 Fibre Channel for Flash?

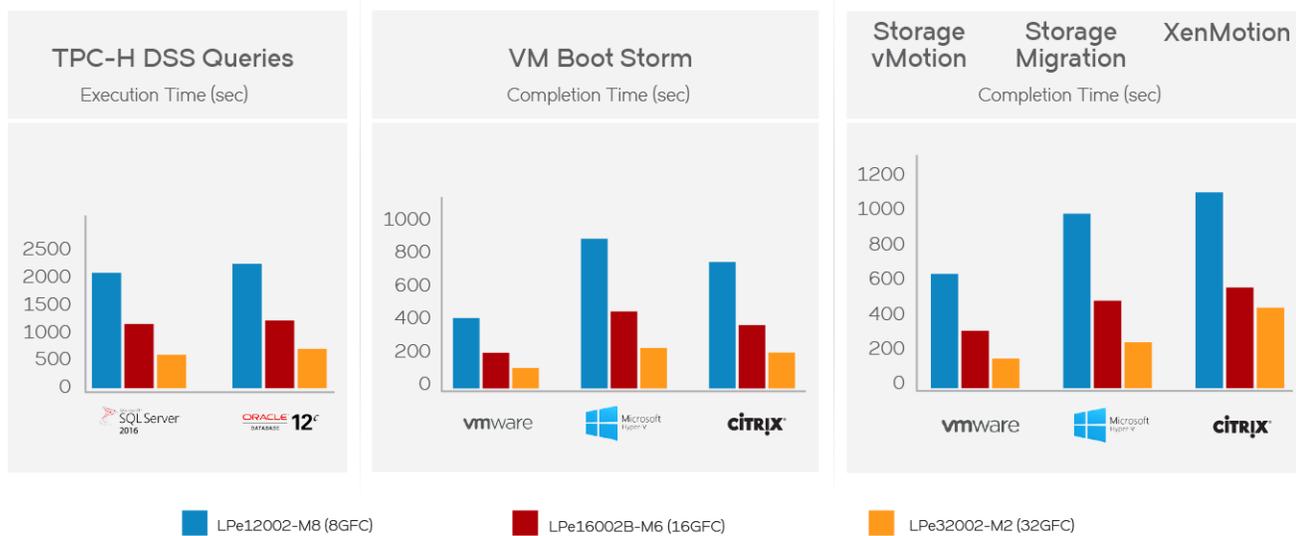
To get the maximum value out of the low-latency, high-throughput designs of all-flash and flash-enhanced hybrid storage, administrators need both fast connectivity and consistent performance through highly deterministic data delivery. Benchmark testing with simulated workloads shows that faster network speeds improve the overall ability of flash storage to maximize IOPS, even if the storage connectivity speed is less

---

<sup>1</sup> Gartner, *Forecast Analysis: External Controller-Based Storage, Worldwide, 3Q16 Update*, December 12, 2016.

than the maximum network throughput. Integrating a flash array with 8 Gbps Fibre Channel connectivity into a gen 6 storage fabric with 32 Gbps connectivity to servers, for example, enables up to four times more queries to be completed in the same amount of time.

Furthermore, Gen 6 connectivity provides on average a 50 percent reduction in application response times compared to 16 Gbps storage networks and Host Bus Adapters (HBAs), and a 71 percent reduction over 8 Gbps networks and HBAs. The incremental performance gain achieved by flash storage ultimately enhances the application response time, the efficiency of the IT infrastructure, and the overall return on storage investment (see Figure 1).



**Figure 1:** Emulex/Broadcom TPC-H benchmark testing: Accelerate 8 Gbps flash storage with 32 Gbps networking.

For workloads such as Online Transaction Processing (OLTP), this increase in performance translates into an increase in the number of transactions executed and a gain in revenue generation for the organization. From this perspective, a small incremental investment in an optimized server and network infrastructure enables maximum utilization of the array performance, delivering a real increase in ROI from capital expenditures. As the need to scale the storage environment grows, and more flash storage is added into the environment with faster connectivity speeds, Gen 6 networks can support storage growth without becoming a bottleneck to application performance.

In addition to offering fast throughput speeds, Fibre Channel fabrics ensure a highly reliable network with deterministic data delivery, a critical aspect for ensuring consistency of data transfer with predictable

---

performance under normal use. As network speeds increase, avoiding Fibre Channel frame retransmission retries due to bit errors becomes critical in order to meet the Fibre Channel standard for bit error rates and to maintain optimal performance levels.

Brocade pioneered the adoption of Forward Error Correction (FEC) in Gen 5 Fibre Channel for enhanced link reliability. FEC is integral to the ThinkSystem Gen 6 standard of sustaining 32 Gbps throughput with OM4 cable. This ability to correct bit errors without discarding the frame is vital at higher data transfer speeds for sustaining or achieving the reliability of transmission, particularly with older cabling or longer links, and almost eliminates the need for performance-disrupting IO retransmissions. This consistency of response time allows application administrators and storage vendors to more accurately optimize their tuning of response times for specific workloads. This results in faster performance levels and greater value from the flash-enhanced array.

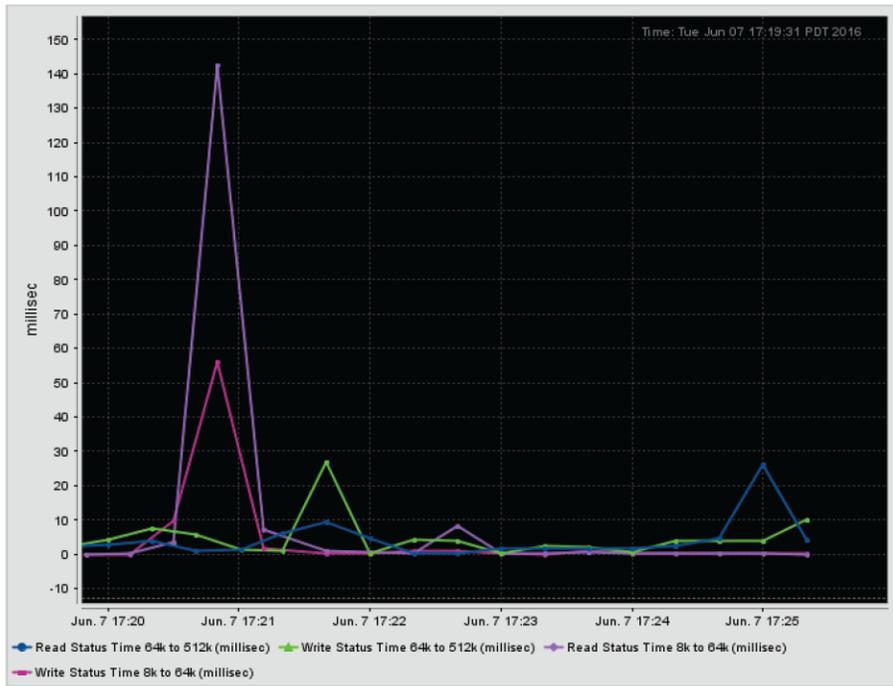
## Maximizing Network Efficiency with Brocade Fabric Vision Technology

Flash storage significantly improves IO response times, with claims of response times being up to 10 times faster. If the connectivity from the host to the storage is not functioning optimally, however, the impact is amplified when high congestion impacts data flows. The resulting performance issues are often extremely difficult to diagnose. The widespread use of host virtualization has yielded mixed workloads resulting from multiple applications (whose own individual workload profiles are difficult to characterize), which can cause unanticipated congestion and affect performance. In such situations, it is unclear who is to blame: the application owner or the storage administrator? Also, it might take days or weeks to identify the root cause of the problem.

Flash storage is often used to mitigate the congestive effects of the so-called “IO blender” effect, but the use of flash can be further optimized by ensuring that no limitations prevent line-speed operations during peak operations. Brocade Fabric Vision® technology, an integral part of Lenovo Gen6 Fibre Channel products, is a suite of tools that allows administrators to pre-validate, monitor, and troubleshoot storage environments to prevent issues from occurring and to mitigate their impact when they do occur. Lenovo Gen 6 Fibre Channel extends these capabilities with the introduction of Brocade IO Insight and Brocade VM Insight, the industry’s only integrated network sensors that provide deeper visibility into the IO performance of storage infrastructure, down to the individual VM. This level of granularity enables quick identification of degraded application performance at the host VM and storage tiers, reducing time to resolution.

## IO Insight

IO Insight proactively monitors IO performance and behavior to provide unparalleled insight into problems and to ensure service levels. With IO Insight, storage administrators can baseline the performance profiles for IOs and storage latency. In addition, they can use the monitoring capability and activity dashboards provided by the Brocade Monitoring and Alerts Policy Suite (MAPS) to quickly identify threshold spikes in IOPS activity. For activity that does not comply with expected behavior, preventative actions can be defined—ranging from admin notifications to port fencing—to avoid greater negative impact (see Figure 2).



**Figure 2:** IO Insight metrics displayed in a Brocade Network Advisor real-time performance graph.

IO monitoring allows greater insight into storage performance and provides demonstrable evidence of achieving critical Service Level Agreements (SLAs). IO Insight enables proactive IO monitoring of:

- Total IOs at a flow level to monitor workload profiles over time
- First response times (maximum and average) for an IO request
- IO latency for Exchange Completion Time (ECT), maximum and average
- Outstanding IOs in the queue, maximum and average

## VM Insight

VM Insight further extends IO Insight by enabling visibility of IO statistics at the level of the individual VM. By applying the interpretation of Gen 6 Fibre Channel standards-based VM identification to IO Insight analysis, VM Insight allows the storage administrator to pinpoint issues not just on a physical server, but down to the specific VM workload. VM Insight further enables IO profiles for individual VMs to be baselined and monitored over time with Fabric Vision MAPS for continual performance monitoring and optimization. In addition to a Gen 6 fabric, VM Insight requires the use of ESX 5.5 or later and driver support from host and target adapters. Customers should check with their HBA or storage vendor to verify support.

---

## Brocade Analytics Monitoring Platform

For additional insight into end-to-end application performance across the entire fabric, storage administrators can leverage the Brocade Analytics Monitoring Platform. The Analytics Monitoring Platform can monitor up to 20,000 data flows and millions of IOPS, with benchmark metrics that include latency, IOPS, and pending IO. By monitoring these metrics across the fabric, administrators can identify behavioral anomalies and help ensure predictable performance, while avoiding downtime to the network and to applications.

## Increasing Performance with Seamless Integration of NVMe over Fibre Channel

Looking forward, the next wave of all-flash arrays will provide connectivity of flash media to Non-Volatile Memory Express (NVMe) over Fibre Channel, offering the opportunity for performance gains and further reductions in latency. Using flash media connected to NVMe enables direct communication from the host server to the storage subsystem via Peripheral Component Interconnect express (PCIe) commands. This avoids the Small Computer Systems Interface (SCSI) translation layer that impacts storage response latency. While this translation time is not significant in a traditional spinning disk-based storage device relative to the overall disk response time, the impact is relatively higher in a flash-based array where seek time and rotational delays are dramatically reduced, creating an opportunity to improve response times by as much as 10 times.

Carrying NVMe traffic over ThinkSystem Gen 6 Fibre Channel enables storage administrators to gain the performance advantage of NVMe-connected storage media while using their existing Fibre Channel network. Fibre Channel networks are ready for use with NVMe, as they offer the same design characteristics of fast throughput, resilient operations, and high reliability that make Fibre Channel the storage connectivity option of choice for enterprise data centers. As a result, organizations can seamlessly integrate ThinkSystem Gen 6 Fibre Channel networks with next-generation NVMe without a disruptive rip and replace operation. By integrating the efficiency of NVMe with the high performance and low latency of ThinkSystem Gen 6 Fibre Channel, organizations can scale IOPS to deliver the performance, application response time, and scalability needed for next-generation data centers.

ThinkSystem Gen 6 Fibre Channel builds on these qualities. As storage response latencies are reduced, server efficiencies and IOPS increase further, shifting more demand onto the network. The 32 Gbps throughput allows support for more VMs and higher IOPS, while introducing minimal latency, ranging from sub-microsecond to low single-digit microseconds.

As more data centers adopt NVMe integrated with Gen 6 Fibre Channel, the technology can realize additional gains by leveraging 128 Gbps Parallel Fibre Channel, also part of the Gen 6 standard. Parallel Fibre Channel stripes four 32 Gbps channels into a single high-speed link that uses a Quad Small Form-Factor Pluggable (QSFP) connector. Support for enabling Parallel Fibre Channel into a future Brocade Fabric OS® release is already designed into ThinkSystem Gen 6 Fibre Channel switches, and will enable the fastest connectivity option available to a high-speed storage device.

---

## Summary

Data centers are recognizing the value proposition of all-flash and flash-enhanced storage, not only to support their demanding, high-performance workloads, but also for mixed workloads that include Tier 2 and Tier 3 applications. ThinkSystem Gen 6 Fibre Channel with IO Insight and VM Insight complements and enhances the value of flash storage by ensuring maximum throughput and the lowest latency. In addition, Gen 6 feature enhancements, such as IO Insight and VM Insight's deeper visibility into the IO performance of the storage infrastructure, ensure the consistently highest level of performance and system utilization. A ThinkSystem Gen 6 network infrastructure enables consolidation of servers and storage infrastructure while maximizing performance and utility, all of which will pay significant dividends on an investment in flash-enhanced storage.

For more information about Lenovo ThinkSystem Fibre Channel products, visit [www.lenovo.com](http://www.lenovo.com).

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. This information could include technical inaccuracies or typographical errors. Changes may be made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any performance data contained herein was determined in a controlled environment; therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems, and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk.

Lenovo, the Lenovo logo, System x, and ThinkServer, are trademarks of Lenovo in the United States, other countries, or both.