

Improving Database Performance with QLogic Gen 6 (32Gb) Fibre Channel

Introduction

Databases are among the most critical services deployed in the data center. Many important enterprise applications leverage one or more databases to store and process business information. Database application owners who are serious about performance have already invested in flash storage. The problem is that flash is faster than most storage area networks (SANs). Previous generation Fibre Channel (FC) technology is holding back high-end storage and this hinders database performance.

Current attempts at solutions include replacing SAN storage with local flash, but this sacrifices the flexibility and reliability of networked storage and is expensive when widely deployed. The other option is increasing the number of ports across the SAN. This increases the complexity and management costs of the SAN and does nothing to address latency added by older networking technologies.

Gen 6 (32Gb) Fibre Channel (FC), released March 1st of this year, doubles the bandwidth of Gen 5 (16Gb) FC to 32GFC for a single lane of traffic while reducing latency compared to previous generations. These advances should be attractive to any enterprise looking to increase the value of high-end storage. However, data center managers and application owners will certainly weigh the cost of a SAN upgrade with the return on investment (ROI) delivered through increased infrastructure usage and the business advantage realized by stronger application performance.

QLogic commissioned Demartek to evaluate the advantages that Gen 6 FC—in the server and at the switch—would bring to common database workloads. We deployed the QLogic QLE2742 Gen 6 FC HBA and a Brocade G620 Gen 6 FC switch with existing high-speed storage devices to compare performance metrics of transactional and data mining applications, plus a database maintenance activity, against older Gen 5 FC technology. In each use case, identical workloads showed a clear benefit to upgrading host HBAs and switching to Gen 6 FC.

QLogic Gen 6 FC HBAs

QLogic has been an industry leader in FC technology for years, through direct sales and OEM relationships. Many popular storage and server vendors deploy or rebrand QLogic HBAs as FC initiators and targets. The QLogic 2700 series Gen 6 FC Adapters double the bandwidth of previous Gen 5 FC Adapters while reducing response times to maximize performance of FC SANs. They also introduce several new features that enable a highly resilient network and improve IT staff productivity.

Enterprises with a business need for Gen 6 FC can deploy the QLogic QLE2742 right now. Backwards compatibility to 8GFC and Gen 5 FC means existing SANs can be upgraded one system or component at a time without needing to redesign the infrastructure. Doubling the bandwidth over Gen 5 FC will immediately reduce server FC port counts and simplify SAN management, even when connecting to legacy Gen 5 FC switches. The QLE2700 Series HBAs include the following features:



- ◆ Single, dual, and quad-port versions
- ◆ Up to 2.6 million IOPS
- ◆ Port isolation to provide enhanced reliability and scalability across all ports on a card (i.e.: QLogic HBAs are intentionally designed so that in the rare event of a hardware issue on the card, only a single FC port is at risk of being disabled, not the entire HBA)
- ◆ 36% improvement in power efficiency measured by data moved/Watt¹
- ◆ The complete set of QLogic StorFusion™ features for Gen 6 FC:²
 - Diagnostic port (D_Port), read diagnostic parameters (RDP) and link cable beacon (LCB) for easy, remote troubleshooting of SAN infrastructure
 - Fabric-assigned port worldwide names (FA-WWNs) and boot LUN discovery for simplified deployment and reduced fabric reconfiguration

¹ QLogic, "The Performance Benefits of Gen 6 Fibre Channel", http://www.qlogic.com/Resources/Documents/WhitePapers/Adapters/WP-The_Performance_Benefits_of_Gen6_FC.pdf

² A detailed summary of all QLogic StorFusion features can be found at : <http://www.qlogic.com/Products/adapters/Pages/QL2600SeriesFW.aspx>

- Forward error correction (FEC) to improve performance and link integrity
- Integration with Brocade[®] Network Advisor (BNA) and Fabric OS
- Enhanced fabric device management interface (FDMI), Fibre Channel Ping (FC ping), and Fibre Channel traceroute (FC traceroute) for checking the connectivity of SAN devices

Database Applications

Relational databases are at the core of many enterprise applications, including retail, financial services, and customer-relationship management systems. Data centers often host dozens, if not hundreds of databases. Writing and retrieving data drives a lot of I/O between servers and storage. There are many types of database workloads, but this evaluation focused on two of them as representative of common I/O drivers in most data centers. These two are transactional, also referred to as online transaction processing (OLTP), and data warehousing. We also chose a common database maintenance process for performance comparison: snapshot replication.

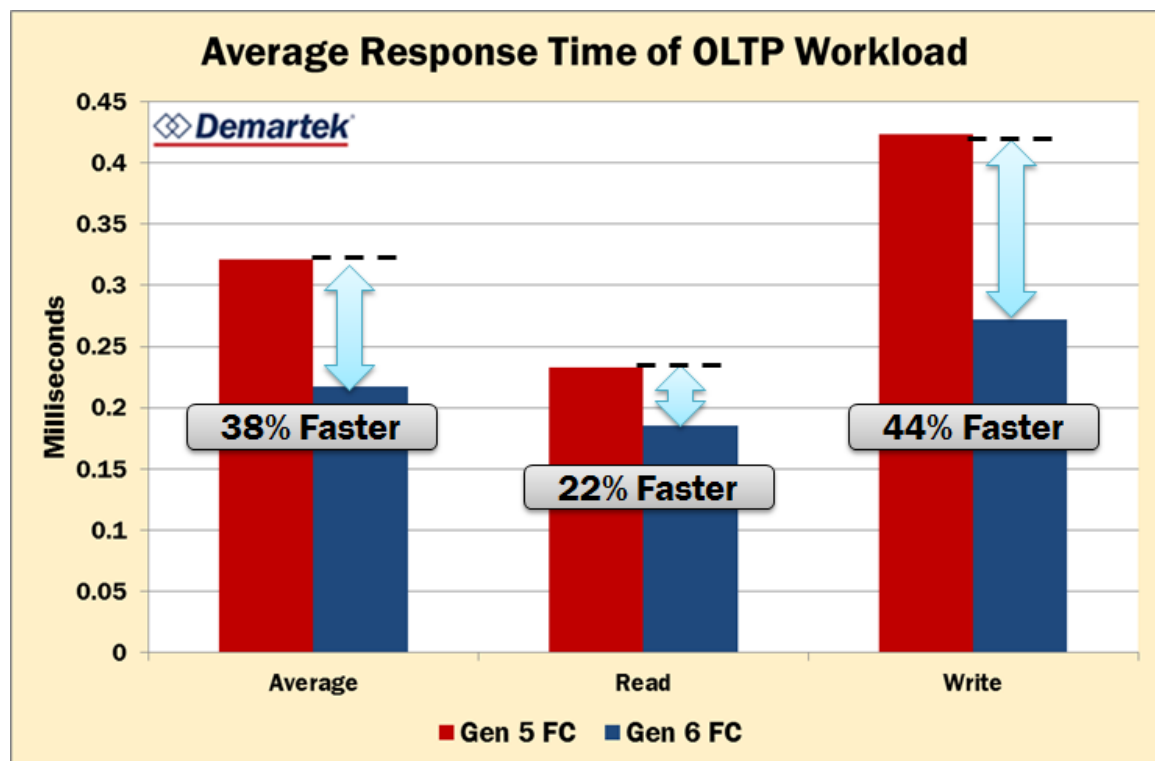
The database software chosen for this evaluation was Microsoft SQL Server running on modern Intel[®] Xeon[®] Processor E5-2600 v3 Series servers and high-speed storage with flash and dynamic random-access memory (DRAM) sufficient to support high bandwidths and IOPS. Workloads were created and executed with the HammerDB open source workload engine. HammerDB runs on a real database, creates real tables, rows, indexes and stored procedures, and executes real queries and transactions, all in the same manner as a bona fide application. The data populating those tables was randomly generated within guidelines defining structure and quantities. We deployed HammerDB OLTP and data warehouse databases for this evaluation. These databases are scalable to hundreds of millions of rows. For our testing, we deployed a 120Gb (data and logs) OLTP database for transactional testing. Data warehousing ran against a much smaller database, strictly to cut down on query times. As our testing was bandwidth and response time oriented, sizes of databases were somewhat immaterial.

We collected performance data at the server, to best represent the user or application experience, with Windows Perfmon. The exact same workload parameters were used with Gen 5 FC or Gen 6 FC in the HBA and switch. The test environment is documented in more detail in the appendix at the end of this paper.

Use Case 1 - A Transactional Database

The formal definition of a transactional, or OLTP, database is a database with the ability to roll back write operations that were not properly completed. It follows that transactional database applications perform both read and write operations to a backend database, as opposed to strictly single direction I/O. Data written is immediately available for retrieval and processing through read I/O requests. Some examples of transactional database applications would be real-time trend modeling, such as weather forecasting, accounting applications, financial fraud detection, and retail sales.

Most OLTP databases create I/O requests in the 8-16KB block size range. This can generate a respectable number of IOPS, but depending on the amount of processing performed by the application engine, may not always drive high bandwidth. IOPs and bandwidth aside, critically important to many transactional database applications is the response time, or latency, of those I/O requests.



The workload never reached Gen 5 FC limits for bandwidth or IOPS, so neither the HBA nor switch was saturated by I/O requests. The response time improvements can be attributed strictly to the upgrade from Gen 5 FC to Gen 6 FC. Latency was sub-millisecond in all cases, but the deployment of Gen 6 FC was responsible for


measurable improvement. As the chart shows, write I/O saw the most significant reduction in latency.

In a highly transactional environment, even small improvements in latency can mean the difference between meeting service-level agreements (SLAs) and driving strategic business advantage versus not. In this test case, the Gen 6 FC response times were a significant fraction of the Gen 5 FC baseline. Remembering that this is a transactional application, processing work occurs based on database I/O. Faster writes place data in the database for retrieval by read requests, and faster read requests reduce the I/O wait state of the processor cores. Freeing cores from a wait state essentially returns them to the system for additional processing, which in turn allows the server to do more work. Combined, these effects drive up ROI on the storage and the server. Aggregated across the data center, this can result in considerable cost savings.

Use Case 2 - Data Warehousing

Data warehousing, or data mining, is a read I/O-intensive series of processes that run against databases referred to as warehouses. These applications examine large amounts of data for trends, patterns, and exceptions—just about any type of analysis. There is very little write I/O to a data warehouse, except perhaps batch jobs to periodically populate tables with new data. Bulk data retrieval is common, with larger I/O sizes than OLTP applications; 32-128KB I/Os are often employed for this purpose. Reviewing sales data for customer purchasing patterns to create targeted advertising is an example of a data mining application.


Due to the large I/O size and sequential nature of the reads, latency is a lesser concern for data warehousing applications. These workloads tend to be more bandwidth intensive. We intentionally saturated the Gen 5 FC infrastructure by instructing several virtual user processes to execute a number of queries designed to generate significant read I/O. As the results show, the Gen 5 FC components hit a single channel line rate most of the time for the duration of the workload. This same workload came close to achieving a steady line rate across the Gen 6 FC HBA and switch, but did not quite saturate it. The end-result is that the requested work completed nearly twice as fast.



Let's not minimize the importance of this; there is a clear business advantage to be gained here. Implementing Gen 6 FC effectively doubles the number of queries the system can run over the test interval. The output from those queries will be imported into other applications for processing. Downstream applications now have twice the input for processing over that same interval and now those systems do more work as well. The cascading benefits of more I/O are obvious, particularly if we again extrapolate them across the data center. Large scale-out storage technologies, including Hadoop, were invented, in part, to get a lot of data to application engines quickly. We just did the same thing using legacy FC array targets by doubling server bandwidth with the QLE2742 HBA. No new server or storage investment was needed to make this happen.

Use Case 3 – Database Snapshot Replication

Database snapshot replication creates a “point-in-time” image of a database, or selected portions of one. Snapshot replication generates large block sequential I/O. The database snapshotted in this evaluation was the same database used to create OLTP traffic. IOPS ranged from 64-256KB reads of the target database to 1MB writes at the snapshot destination. The snapshot replication process is not sensitive to latency. As with the data warehouse workload, this use case focuses on bandwidth and time to complete.



The replication process began with a massive copy request of all database objects defined in the replication parameters. Large block reads fed 1MB writes to the snapshot repository. The higher bandwidth of Gen 6 FC was leveraged by the snapshot process, resulting in completion of the entire process roughly thirty seconds sooner than with Gen 5 FC. That is nearly a third faster.

Backups are unavoidable and no business would ever incur the risk by omitting them, but the processes do consume server and storage resources. There are many ways to backup or copy a database and snapshotting is just one of them – albeit one that can consume a lot of storage bandwidth. A well-designed backup strategy will not require constant, full point-in-time snapshots, but snapshots will likely occur at least occasionally (perhaps as an initial backup image or to create production database copies for test and development). While a backup process is running, other database I/O is impacted. Getting that process out of the way as quickly as possible reduces the length of time that transactional or warehouse I/O is degraded.

Summary and Conclusion

The use cases demonstrated in this paper highlight the response time and bandwidth benefits of Gen 6 FC over previous generations. OLTP experienced a 22%-44% reduction in I/O latency. By responding to I/O requests faster, processor cores are kept busy doing application work instead of remaining idle, waiting for data. System use goes up along with customer satisfaction, which is good for ROI and good for business.

Doubling bandwidth saves time, as shown by the large block, sequential data warehouse and snapshot replication workloads. Time really *is* money for competitive businesses. With Gen 6 FC, we saw data warehouse queries complete in half the time it took for the same work on Gen 5 FC. No business is mining data in a vacuum; queries against a data warehouse feed downstream applications to deliver tactical or strategic business advantages. Bringing data in twice as fast means *those* applications are now that much faster as well. The same is true for database maintenance, such as snapshot replication. By cutting that process' time by a third, the database system reclaimed CPU cycles and I/O capacity tied up for unproductive, but necessary, administrative functions, minimizing the time application I/O and processing would be impacted.

Application owners who demand high performance invest in flash storage. They also know that flash's potential is wasted without a SAN fast enough to deliver the bandwidth and IOPS flash is built for, at latencies low enough to support real-time processing. A SAN that doesn't throttle performance maximizes ROI in high-end storage systems. Gen 6 FC makes that network a reality.

Gen 6 FC, supported with QLogic QLE2700 Series HBAs and a Gen 6 FC switch, forms a platform tailor-made to support all-flash systems, unleashing the blazingly fast performance of existing all-flash devices, and future-proofs a SAN for new technologies. We expect to see storage vendors start rolling out Gen 6 FC storage targets soon. As technologies like Non-Volatile Memory Express (NVMe) make their way into storage products, SAN performance will become a bottleneck. Deploying Gen 6 FC HBAs in servers now as a future-proofing strategy will save effort and expense throughout the next few years as the industry delivers newer and faster products.


Benefits continue to cascade across the data center as I/O response time reduction drives up server usage by reducing or eliminating the time processor cores wait for I/O. This returns processing power to servers for more work, through additional transactions or new workloads. In virtualized environments, this can mean increased virtual machine

(VM) density. Less idle processing time on a server is a very real ROI enhancement. When extrapolated across the data center, the value of these benefits adds up quickly.

It would be myopic to consider server and storage performance improvements as the only benefits of upgrading host HBAs to Gen 6 FC. From an administration perspective, QLogic StorFusion features reduce the effort and time to deploy and manage an FC fabric. The diagnostic and analytic capabilities of QLogic D_Port (Brocade ClearLink™) and RDP minimize or eliminate outages caused by formerly common infrastructure issues. Add in the potential to simplify SAN architecture through the consolidation of ports and cables, while still achieving similar or substantially better I/O, and we can see significant cost savings over time.

Enterprises deploying flash in the SAN need to be thinking about Gen 6 FC. Older generations of FC are holding back high-performance storage, and that's limiting the business advantage that storage was purchased for. If storage performance is critical, we recommend contacting QLogic now to begin upgrading legacy SANs to Gen 6 FC.

Appendix



A private SAN was implemented for this evaluation, with Gen 6 FC technology only deployed at the host and Brocade G620 Gen 6 FC switch. High-speed storage targets hosted database and replication volumes with a combined 14 lanes of 8 Gbps FC³.

An Intel Xeon Processor E5-2600 v3 Series (codenamed Grantley) server supported the PCI Express (PCIe) requirements (PCIe 3.0 x8 minimum to support full dual-port line rate) and delivered enough processing power to drive the I/O needed to effectively demonstrate Gen 6 FC capabilities. We executed each test scenario twice, with the

³ In addition to being the default Fibre Channel technology of the targets used in this evaluation, older generation Fibre Channel in the storage tier also proves the benefit of upgrading to Gen 6 in the server and switch alone, even if the storage remains at a legacy Fibre Channel version.

QLogic QLE2742 Gen 6 HBA or a Gen 5 HBA serving as the storage initiator. A single port was active on either HBA.

QLogic is a registered trademark of QLogic Corporation in the United States and/or other countries.

Microsoft and Microsoft SQL Server are registered trademarks of Microsoft Inc. in the United States and/or other countries

Demartek is a registered trademark of Demartek, LLC.

All other trademarks are the property of their respective owners.