HP ProLiant Gen8 vs Gen9 Server Blades on Data Warehouse Workloads

Gen9 Servers give more performance per dollar for your investment.

Executive Summary

Information Technology (IT) organizations face increasing demands on their infrastructure in order to support business growth. In September of 2014, HP introduced its new line of HP ProLiant BladeSystem Gen9 Server Blades with the goal of improving the price/performance ratio of data center computing to empower the data-driven enterprise. These new servers take advantage of the new Intel Xeon E5-2600 v3 processors, DDR4 Smart Memory, and improved Virtual Connect FlexFabric 20Gb. This allows HP to deliver an infrastructure that meets IT needs for a more service driven platform. The HP ProLiant BladeSystem Gen9 Server Blades are purpose-built for enterprise workloads such as virtualization and cloud computing, delivering lower cost, faster time, and higher value of service delivery. These next generation blades are specifically focused on ensuring that there is a close match between the application being run, the environment in which it is running, and the business outputs being sought.

When IT Managers consider upgrading older generation HP hardware or purchasing new hardware to help manage the increased demands of their datacenters, they want to know that the next generation offers solid, measurable performance to drive business improvements.

With this in mind, HP commissioned Demartek to compare Gen8 and Gen9 performance of their HP ProLiant BL460c Server Blades using a Data Warehousing workload. The workload was chosen to provide comparative information as a decision support vehicle. It consists of a suite of business oriented ad-hoc queries and concurrent data modifications. The queries and the data populating the database have been chosen to have broad industry-wide relevance. This workload illustrates decision support systems that examine large volumes of data, execute queries with a high degree of complexity, and give answers to critical business questions. Two variations of data warehousing workloads were run: a single sequence of queries and then a set of streamed multiple query sequences concurrently.

Key Findings

- The BL460c Gen9 Server Blades outperformed the BL460c Gen8 Server Blades in all tests, despite the BL460c Gen9 Server Blade processor having a slower clock speed than the BL460c Gen8 (2.7 GHz vs 2.3 GHz).
- The largest performance improvement was shown when multiple streams of database workload queries were performed. The Gen9 blades performed the streamed database workload queries an average of 31% faster and the single sequence of database workload queries an average of 10% faster.
- While executing the database workload queries, the Gen9 Server Blade had more processor headroom available to handle other processes or other workload applications.
- When comparing like configurations of the blades used for these tests, the cost for Gen9 Server Blade hardware was 10% lower than the cost for the Gen8 Server Blade (using HP Internet List Price). We did not include the price for the network adapters, as they were the same 20Gb CNA model and quantity for both servers.
- An important consideration when evaluating servers for a data center environment is the trade-off involving acquisition costs vs. lifecycle Total Cost of Ownership (TCO). In this database application example, we run headlong into the world of server hardware costs plus software application licensing vs. increased performance, lower operational costs, and business needs over the lifecycle of the solution.
    - Performance considerations – You will find that a higher performance system that involves more processor cores, larger on-chip cache sizes and the ability to effectively use those cores can produce significant business benefits in faster execution times, earlier access to information, and the potential of increased business agility and service delivery. You will see those performance numbers in this report. That does come with a significant per processor core license cost when using Microsoft SQL Server 2014 with Intel E5-2600 v3 processors that have up to 18 cores per processor. The list price of Microsoft SQL Server 2014 options are available here: http://www.microsoft.com/en-us/server-cloud/products/sql-server/purchasing.aspx
    - Acquisition cost considerations – The other consideration for the data center decision-makers is balancing the overall acquisition costs that

include application software licensing with the performance needed by the business for immediate requirements and to support future growth, competitive issues, new services that may be deployed, and overall internal and customer satisfaction. Each company will have to evaluate the trade-offs between performance and the overall system costs necessary to meet business goals. In this report we concentrated primarily on attaining and evaluating performance of the Gen8 and Gen9 Server Blades. Realize that the entire system of servers, storage, and networking contribute to the performance of a system. Just as the processor speed and cores come into play, so too does the SSD drives we used in the iSCSI storage array and the VC FlexFabric components that provided 20 Gb FlexibleLOMs between the servers and the VC FlexFabric modules, along with 40 GbE uplinks to the HP 5900 Series Switch which connected to an iSCSI SSD array.

New Features and Total Cost of Ownership

The improved Gen9 Server Blade hardware innovations from Gen8 Server Blades include the following:

- DDR4 memory with increased memory speeds, increased bandwidth, decreased latency, and improved error correction.
- The Intel Xeon Processor E5-2600 v3 product family with increased cores, increased cache, and Quickpath Interconnect (QPI).

The following server and server option prices were quoted by HP as the Hardware Internet List Price. We did not include the Microsoft SQL Server 2014 license cost as we were testing for performance to see if more processor cores helped or if the processor frequency was the determining factor when running this load test. For the overall system, we wanted a fast storage subsystem hence the use of Virtual Connect (VC) FlexFabric interconnects that could provide the 10GbE uplinks to the Top of Rack HP 5900 Series Switch and then connect via 40GbE to the iSCSI SSD Storage Array. This provided the high bandwidth pipes that would be needed to provide data to the servers during these tests.

Prices for our configured test servers:

	Gen8		Gen9	
Blade	BL460c	$1,961	BL460c	$2,011
CPU (2 sockets)	Intel® Xeon® E5-2680 2.70 GHz, 8-core	$5,198	Intel® Xeon® E5-2670 v3 2.30GHz, 12-core	$4,458
Processor cache size	L1: (8 x 32 KB) x 2 L2: 8 x 256 KB L3: 20 MB		L1: (12 x 32 KB) x 2 L2: 12 x 256 KB L3: 30 MB	
DIMM	512 GB, DDR3 1333MHz	$15,984	512 GB, DDR4 2133MHz	$14,384
Total Cost (without adapters)		$23,143		$20,853

Test Setup

For the tests, each HP ProLiant BL460c Server Blade was equipped with a FlexFabric 20Gb 2-port 630FLB adapter and two FlexFabric 20Gb 2-port 630M adapters.

The blade enclosure included a pair of HP Virtual Connect (VC) FlexFabric-20/40 F8 Modules in interconnect bays 1 and 2, and two pairs of HP VC FlexFabric 10Gb/24-Port Modules in bays 3 and 4, 5 and 6.

HP OneView was used to provision 14x10Gbps uplinks from the blade chassis to the HP 5900 Series Switch JG838A. The switch had 2X40Gbps uplinks to iSCSI flash storage.

A Data Warehousing database and TempDB database were stored on the flash storage targets for use in the Microsoft SQL Server Database Warehousing workload.

Two diagrams are shown below. The physical diagram provides the major hardware components and connections. The logical diagram shows the individual logical uplink connections and information for the database storage.

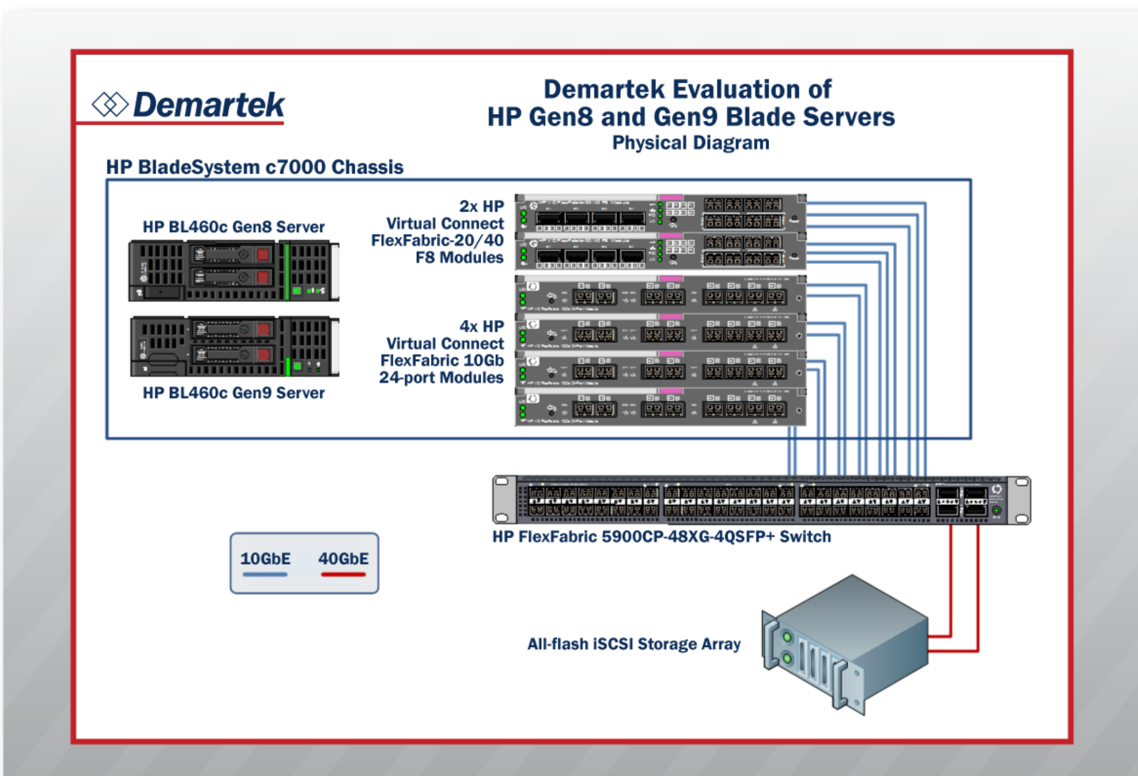

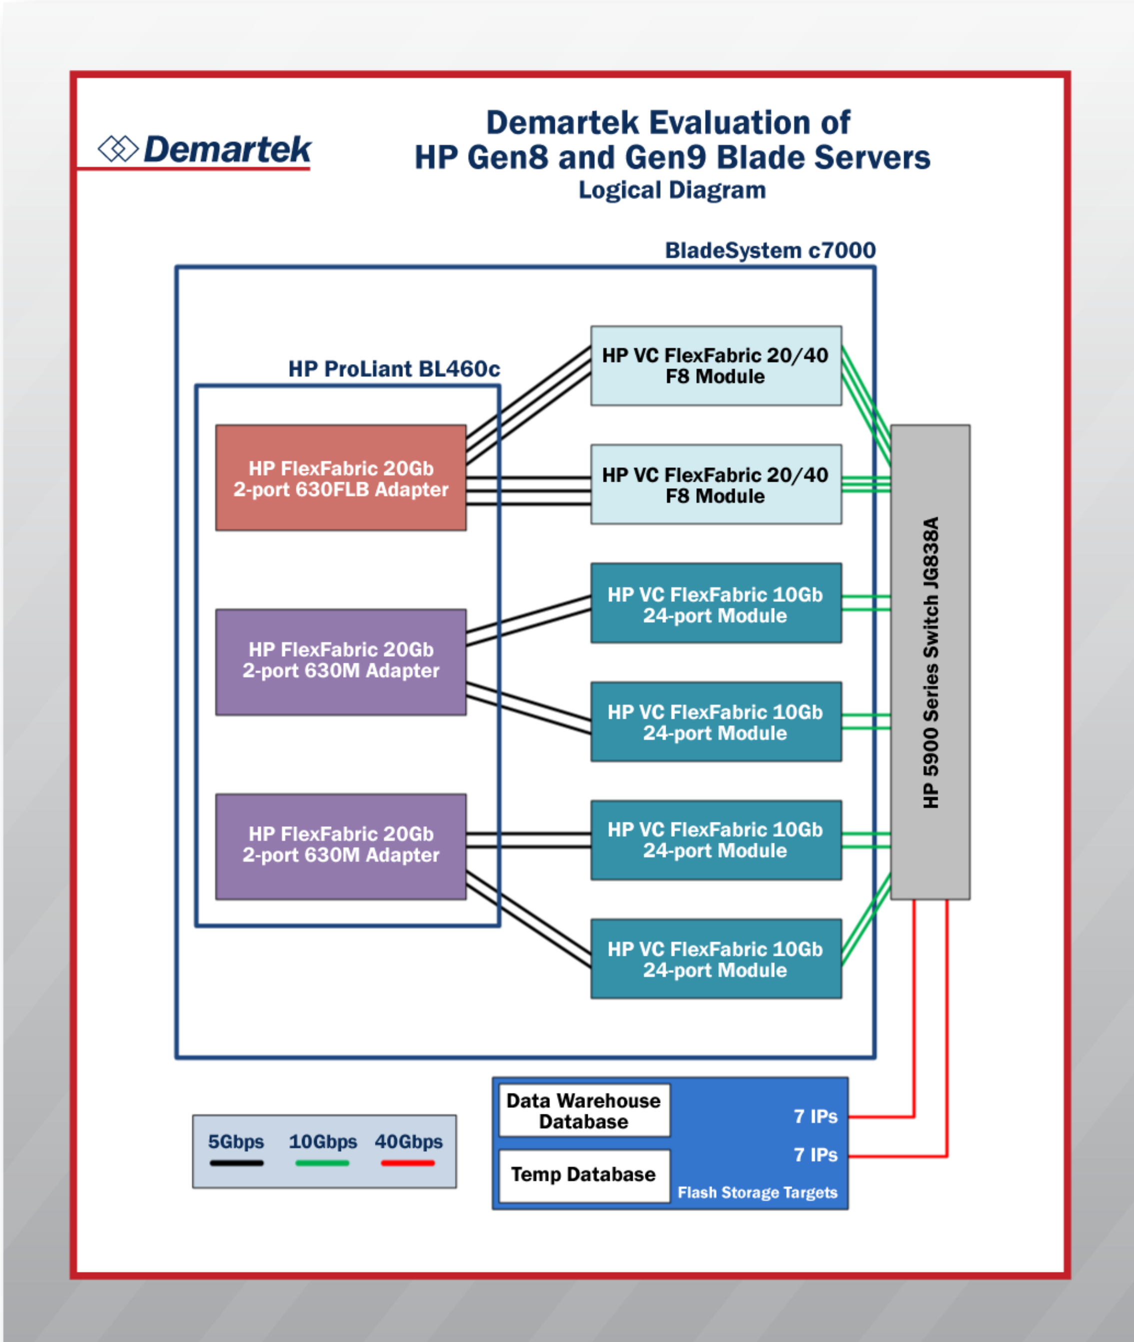

Two different types of VC modules were used in our setup. This impacted our available bandwidth. Each FlexFabric 20Gb adapter has 2 ports of 20Gb. The top adapter, colored red, connects those ports to the 20Gb ports on the HP VC FlexFabric 20/40 F8 module, so it has 20Gb of throughput available on each port. However, the bottom two adapters, colored purple, connect their 20Gb ports to 10Gb ports on the HP VC FlexFabric 10Gb 24 port Modules. While the server blade has 20Gb available on each port, the VC Modules have only 10Gb, so only 10Gb of links were used on each of these ports.

In order to ensure that the uplinks and networks were not adversely impacting performance, two measures were taken:

- To ensure there was enough uplink bandwidth, more uplinks were allocated than what was necessary to support the server. The server blade has a total of 2x20Gb + 4x10Gb = 80Gb of bandwidth available. 14x10Gb uplinks made 140Gb of throughput available to the server blade.
- In an effort to ensure packets used all available uplinks to reach the top of rack switch, each uplink was assigned a different HP OneView network. As there were more uplinks than server blade ports, the ports were split into 5Gb logical links in HP OneView. Each link was on a separate network with a separate uplink. With this setup, each 5Gb logical server blade port had a different, dedicated 10Gb uplink available to it. Configuring Multi-Path IO (MPIO) on the server blade ensured that all server blade ports and all their corresponding uplinks were used.

Gen8 and Gen9 with Database Queries in Single-File

Three database workload tests were run where a single sequence of 22 queries is completed. Each query is focused on stressing a different part of the database infrastructure, from IOPS to throughput to processor. Therefore a metric won't be maxed out for the entire duration of the test, but only for specific queries that are designed to stress that particular metric. The server resource that is maxed out for the most queries is the server processor bottleneck for database workloads.

Both the Gen8 and the Gen9 Server Blades have a significant portion of their runs where all cores approach 100% utilization. The corresponding query is noted on the graphs. For the Gen8, queries 18, 1, and 5 max out the processing capability for an extended period of time.

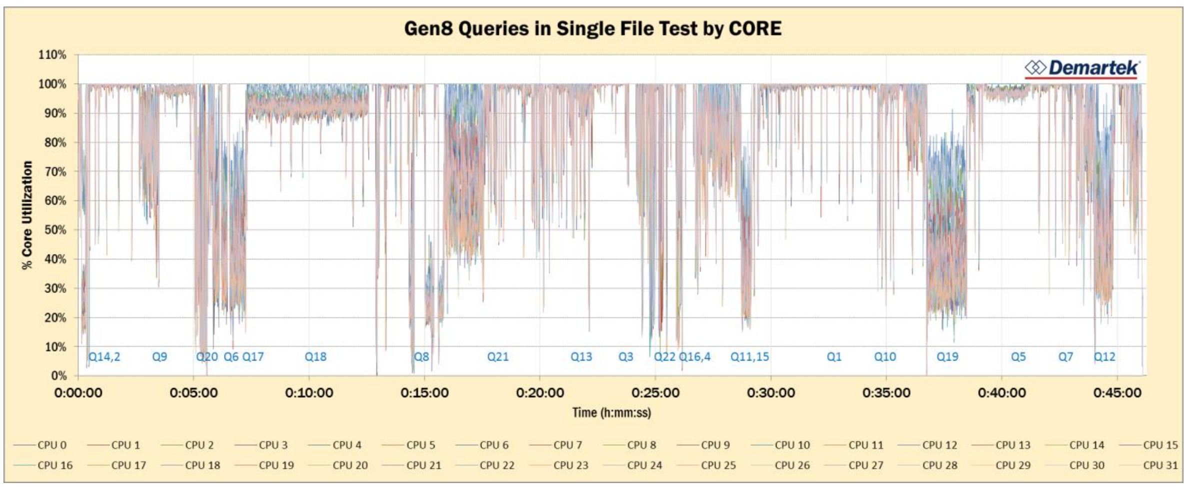

The Gen9 server spends less time at full CPU utilization. Queries 18, 1, and 5 push the CPU to near 100% utilization just like in the Gen8, however with multiple cores working in parallel the duration to complete these processor intensive queries is reduced.

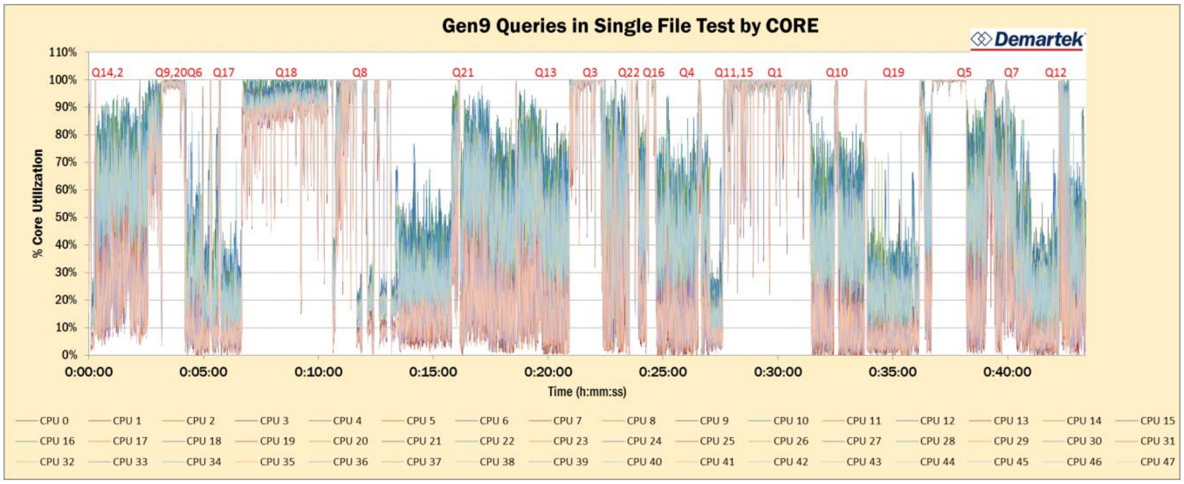

Completion Times (mm:ss)	Gen8	Gen9
Query 18	5:40	3:58
Query 1	5:10	3:35
Query 5	4:29	3:38

The tests were run 3 times on each system and the Gen9 was consistently faster, by an average of 4:55 minutes, or approximately 10%.

Completion Times	Gen8	Gen9
Test 1	0:52:19	0:43:03
Test 2	0:46:15	0:43:34
Test 3	0:46:12	0:43:24

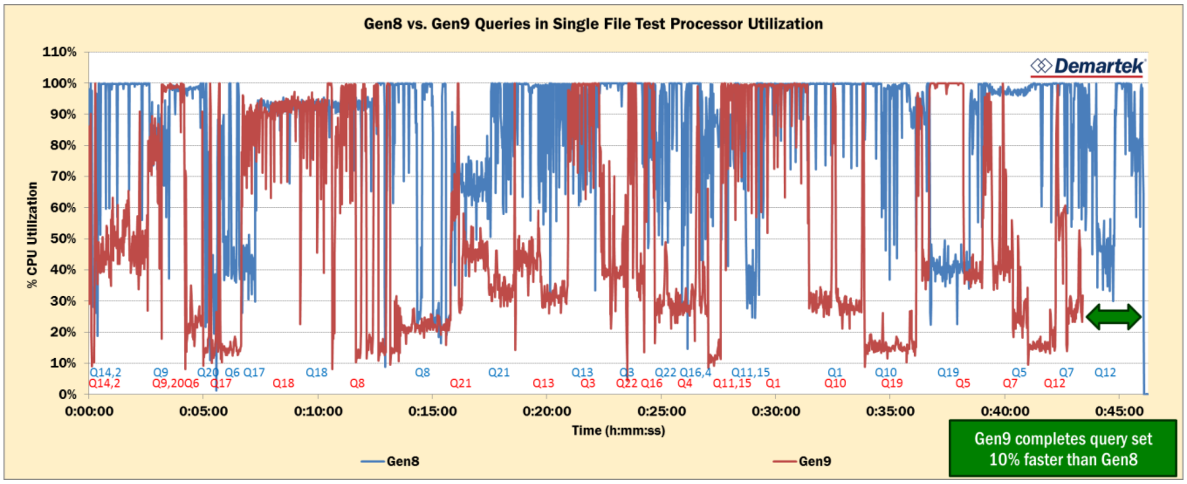

The average processor use over the entire test run was 85% for the Gen8 server and 53.5% for the Gen9 server. The Gen9 server still had plenty of processor capability to spare for other workloads.

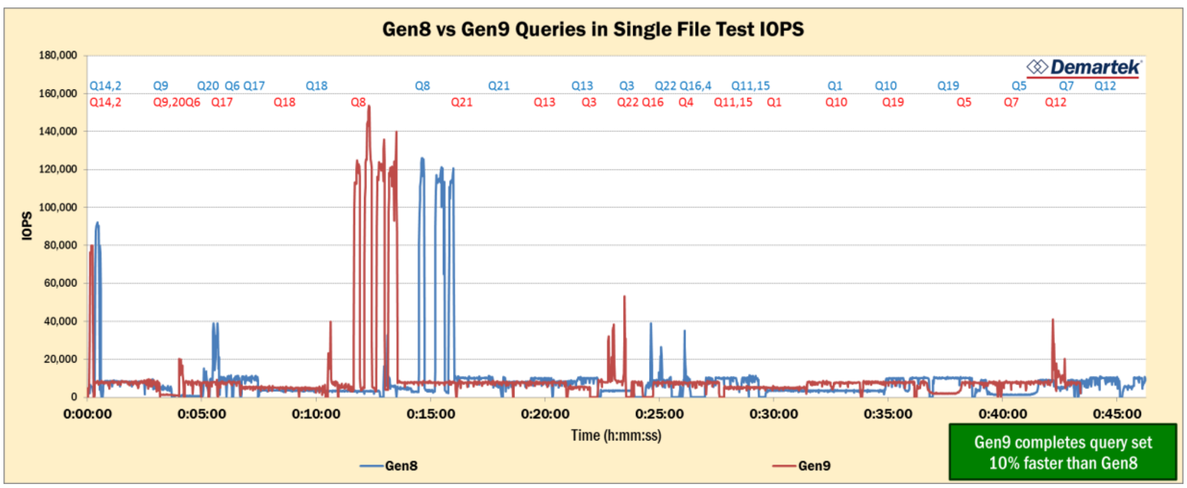

Query 8 stresses the IOPS. Note that the IOPS peaks are higher for the Gen9s, most likely due to the enhanced capability to process the data received.

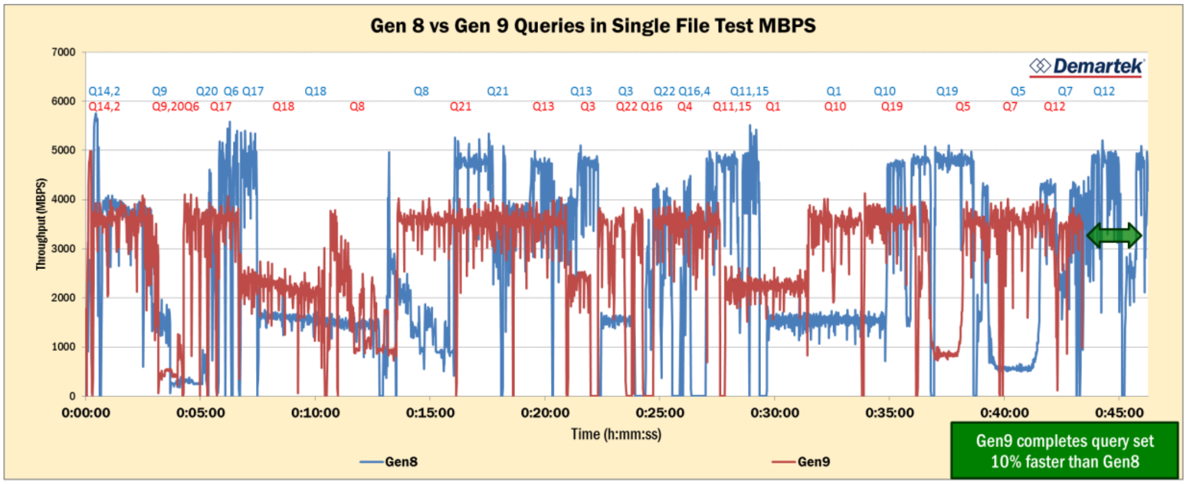

The results show that the Gen9 server is able to produce a more steady and predictable load to the storage subsystem, while the Gen8 server has many short, high peaks in throughput.

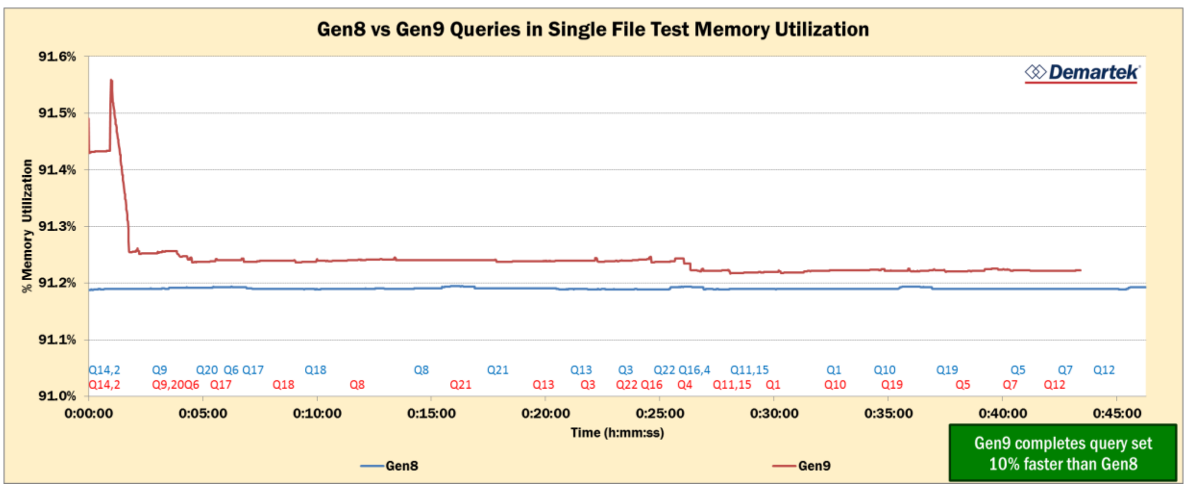

Both systems had 16x32GB DIMMs, giving a total of 512 GB of memory available, and MSSQL Server was capped at 461GB (approximately 90% of available memory on each system). If we look at the scale of the graph, we can see that the memory utilization is mostly due to the 90% allocated to MSSQL server.

Gen8 and Gen9 Server Blades with Multiple Concurrent Queries

Three database workload tests were run where 7 concurrent streams of 22 queries were run. A total of 154 queries are completed in the test, with seven queries running at any one time.

All cores were utilized to a greater extent on both servers. Again the Gen8 server maxed out the processor compute power more frequently than the Gen9 server.

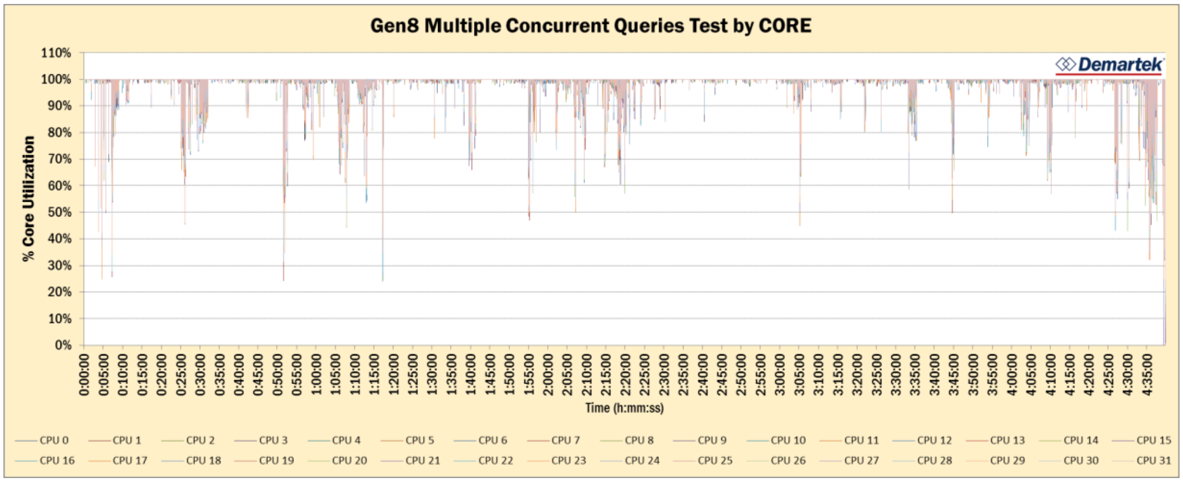

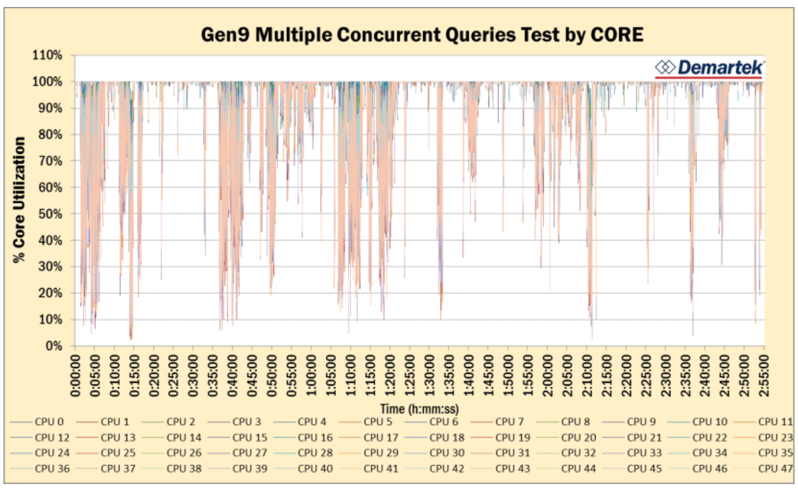

Running 7 queries concurrently ensures that there should be several processor-intensive queries available to take advantage of CPU resources at any point during the test.

The tests were run 3 times on each system and the Gen9 server was consistently faster, by an average of 1 hour and 25 minutes, or approximately 31%. This is most likely due to the increased number of cores available on the Gen9 system that enable the server to complete more tasks in parallel.

Completion Times	Gen8	Gen9
Test 1	4:40:41	3:17:40
Test 2	4:39:55	3:12:32
Test 3	4:41:17	3:15:47

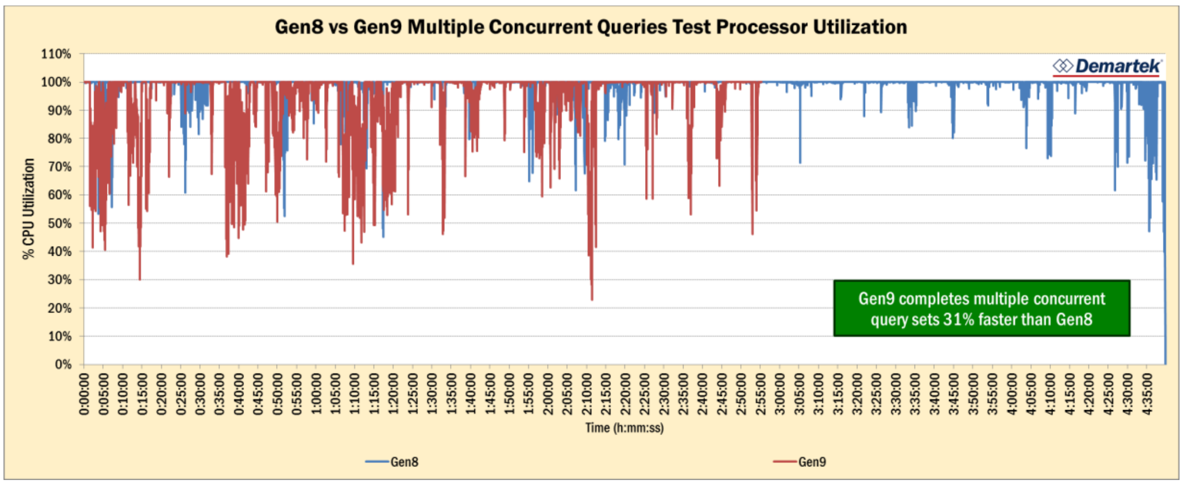

The average processor use over the entire test run was 99.5% for the Gen8 server and 94.6% for the Gen9 server. The Gen9 server still had more processor capability that could be used for other workloads.

We see the same patterns for memory, throughput and IOPS as we did in the single query set tests.

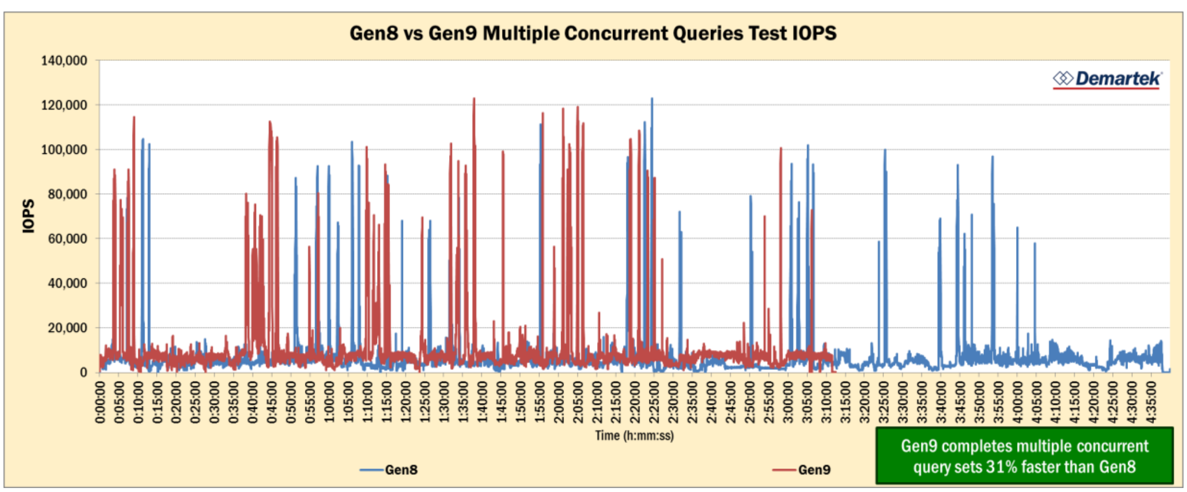

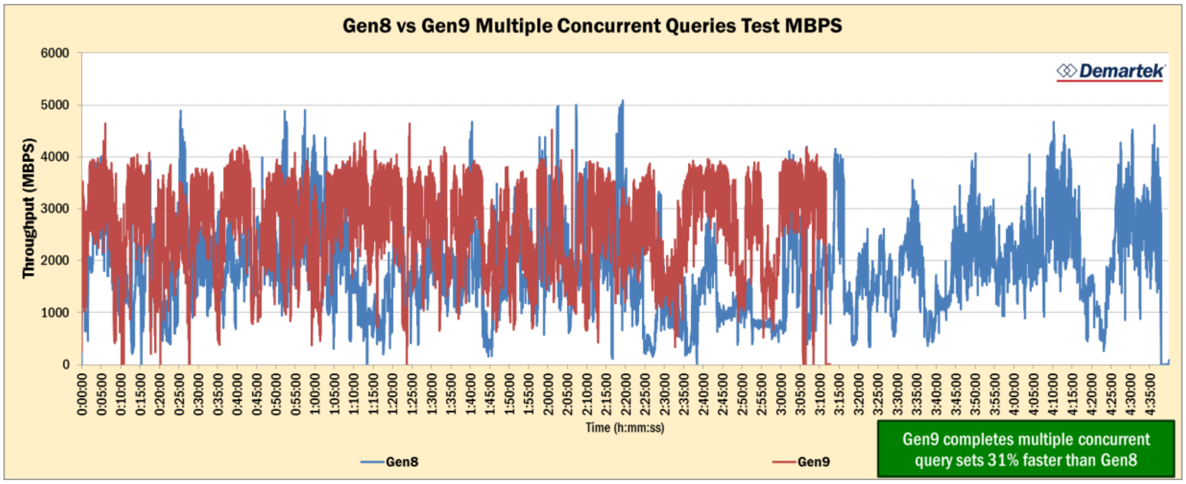
Gen8 vs Gen9 Multiple Concurrent Queries Test MBPS
Throughput (MBPS)
Time (h:mm:ss)
Gen8
Gen9
Gen9 completes multiple concurrent query sets 31% faster than Gen8

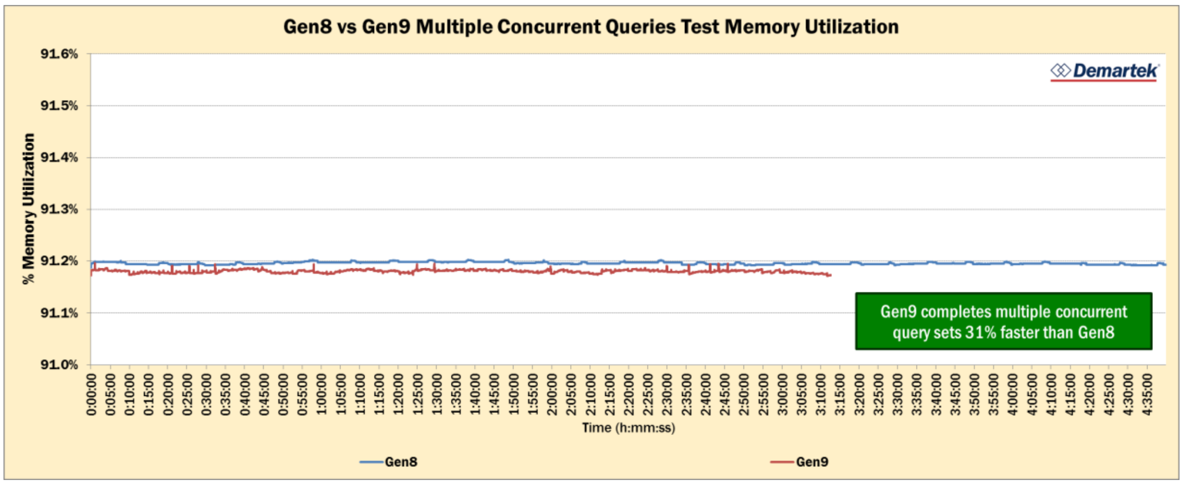
Gen8 vs Gen9 Multiple Concurrent Queries Test Memory Utilization
% Memory Utilization
Time (h:mm:ss)
Gen8
Gen9
Gen9 completes multiple concurrent query sets 31% faster than Gen8

Summary and Conclusion

The HP ProLiant BL460c Gen9 Server Blades outperformed the HP ProLiant BL460c Gen8 Server Blades, especially in multi-threaded database throughput tests. Should an IT manager need to purchase more servers to handle an increased database workload, the Gen9 servers would be a smarter choice. The Gen9 Server Blade comparative configuration hardware costs 10% less and offers a 31% increase in performance.

Performance improvements drive more savings. If we quadrupled our workload to include 28 streams and 616 total queries, it would take four HP ProLiant BL460c Gen8 Server Blades to complete the workload while it would take 25% fewer (3) HP ProLiant BL460c Gen9 Server Blades to complete the same task. In short, three Gen9 Server Blades could do the same job as four Gen8 Server Blades while saving $30,003 (33% reduction) in pure capital acquisition costs. We would also anticipate data center savings in power and cooling, space, and maintenance operational costs.

Putting some of this into quantifiable terms:

- Saving of time per database run = 85 minutes
- Doing two database runs per day = 170 minutes of time savings per day.

Enterprise data centers are generally available 365 days per year and the lifecycle of a server is typically 3 years or 1095 days. A data center would save in operational time over those 3 years:

- 170 minutes/day X 1095 days = 186,150 minutes or 3102.5 hours /database server

That is over a man-year's worth of OPEX savings. It appears that a typical database engineer earns approximately $104K per year including benefits. If we use $50/hour for our calculations, then a business could save up to $50,000 per database engineer per year in time now available for use doing other revenue generating tasks.

It is a potential win-win for your business and your customers.

The most current version of this report is available at http://www.demartek.com/Demartek_HP_BladeSystem_Gen8_Gen9_Evaluation_2015-07.html on the Demartek website.