HP ProLiant BladeSystem Gen9 vs Gen8 and G7 Server Blades on Data Warehouse Workloads

Gen9 server blades 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 server blades 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 deployment models such as virtualization and cloud computing, delivering lower cost, faster time, and higher value of service delivery. These next generation server 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 outcomes 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 increases to drive business improvements.

With this in mind, HP commissioned Demartek to compare Gen9 vs Gen8 and G7 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 blade outperformed the BL460c Gen8 server blade and BL460c G7 server blade in all tests, despite the BL460c Gen9 server blade processor having a slower clock speed than the BL460c Gen8 and BL460c G7 (2.7 GHz and 2.93GHz vs 2.3 GHz). We believe this can be attributed to the new features of the Intel® Xeon® Processor E5-2600 v3 product family as well as other new Gen9 server blade features.
- The largest performance improvement was shown when multiple streams of database workload queries were performed. The Gen9 server blade completed the same multiple stream workload 30.4% faster than the Gen8 server blade and 56.2% faster than the G7 server blade. In other words, the Gen9 server blade can perform more than twice the work of the G7 server blade, or almost 1½ times the work of the Gen8 server blade in the same amount of time.

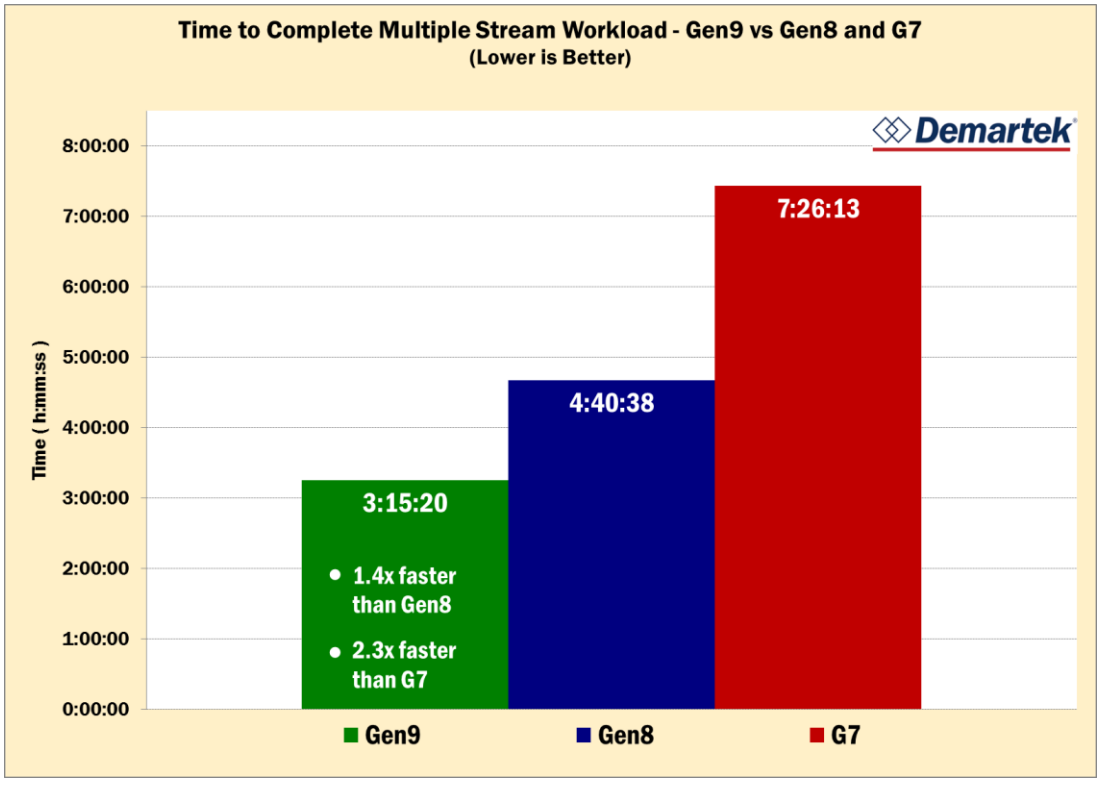

- The single query set of database workloads also showed large differences in time to complete. The Gen9 server blade completed the same single stream workload 10.2% faster than the Gen8 server blade and 43.6% faster than the G7 server blade.
- When comparing like configurations of the server blades used for these tests, the cost for Gen9 server blade hardware was 9% lower than the cost for the Gen8

server blade and 15% lower than the cost for the G7 server blade hardware (using HP internet list price, see footnote on the bottom of page 5).

- 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® Xeon® 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 G7, Gen8 and Gen9 server blades. Realize that the entire system of server blades, storage, and networking contribute to the performance of a system. Just as the processor speed and cores come into play, so too do the SSD drives we used in the iSCSI storage array and the VC FlexFabric components that provided 20 Gb FlexibleLOMs between the server blades 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 includes 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, increased processor system bus, increased memory channels, increased memory bandwidth, and Quickpath Interconnect (QPI).
- 20 Gb FlexibleLOMs between the server blades and the VC FlexFabric modules

The following server blades and server blade 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 server blades during these tests.

Prices and comparison for our configured test server blades:

	Gen9		Gen8		G7	
Server Blade	BL460c + 1 FlexFabric 20Gb 2-port 630FLB Adapter	$2,011	BL460c + 1 FlexFabric 20Gb 2-port 630FLB Adapter	$1,961	BL460c + Embedded 10Gb 2-port adapter + 1 Processor	$4,249
CPU (2 sockets)	Intel® Xeon® E5-2670 v3 2.30GHz, 12-core	$4,458	Intel® Xeon® E5-2680 2.70 GHz, 8-core	$5,198	Intel® Xeon® X5670 2.93GHz, 6-core, add for second socket only*	$2,186
Processor cache size	L1: (12 x 32 KB) x 2 L2: 12 x 256 KB L3: 30 MB		L1: (8 x 32 KB) x 2 L2: 8 x 256 KB L3: 20 MB		L1: (192KB) L2: 1536KB L3: 12288KB	
Processor System Bus	9.6 GT/s		8GT/s		6.4GT/s	
Instruction Set	AVX 2.0		AVX 2.0		SSE4.2	
Max # of Memory Channels	4		4		3	
Max Memory Bandwidth	68 GB/s		51.2 GB/s		32 GB/s	
DIMM	512 GB, DDR4 2133MHz	$14,384	512 GB, DDR3 1333MHz	$15,984	384 GB, DDR3 1066MHz	$18,504
NIC	2 x FlexFabric 20Gb 2-port 630M Adapter	$1,997	2 x FlexFabric 20Gb 2-port 630M Adapter	$1,997	2 x NC553m 10Gb 2-port FlexFabric CNA	$1,798
Total Cost (without adapters)		$22,850		$25,140		$26,737

*The G7 Hardware is past obsolescence, which made the HP internet list price for an additional X5670 CPU for the BL460c G7 unavailable. Instead, the HP internet list price for a similar processor, the X5660, was quoted. The X5660 has the same cache and other features, but differs in clock speed. The X5660 is 2.8GHz and the X5670 is 2.93GHz. Had the correct processor been quoted, the BL460c price would most likely have been higher resulting in a larger difference in price between the G7 and Gen9 server blades.

Test Setup

For the tests, the Gen8 and Gen9 HP ProLiant BL460c Server Blades were equipped with a FlexFabric 20Gb 2-port 630FLB adapter and two FlexFabric 20Gb 2-port 630M adapters. The G7 HP ProLiant BL460c Server Blade was equipped with a FlexFabric Embedded Ethernet 10Gb 2-port and 2 NC553M 10Gb 2-port FlexFabric Converged Network Adapters.

The BladeSystem 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 BladeSystem 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 sets of diagrams are shown below. The physical diagram provides the major hardware components and connections. The logical diagrams show the individual logical uplink connections and information for the database storage.

Physical Diagram

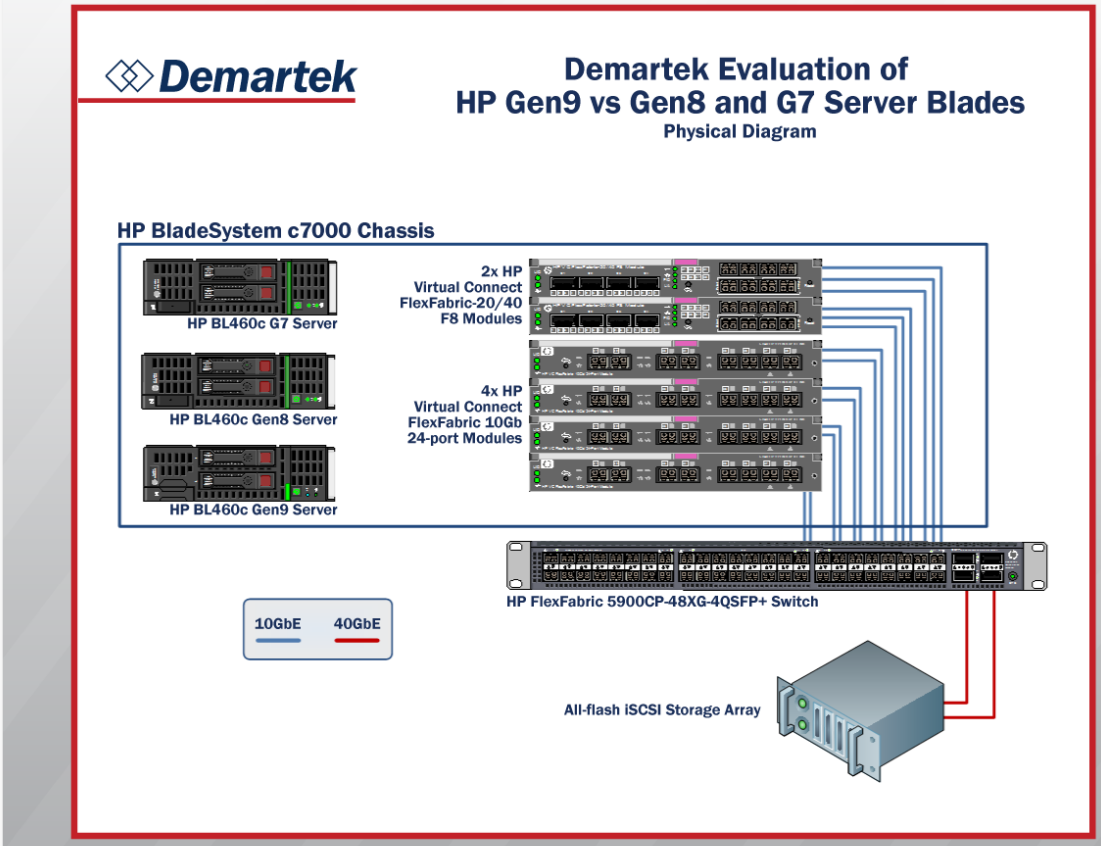

Logical Diagrams

G7 Server Blade

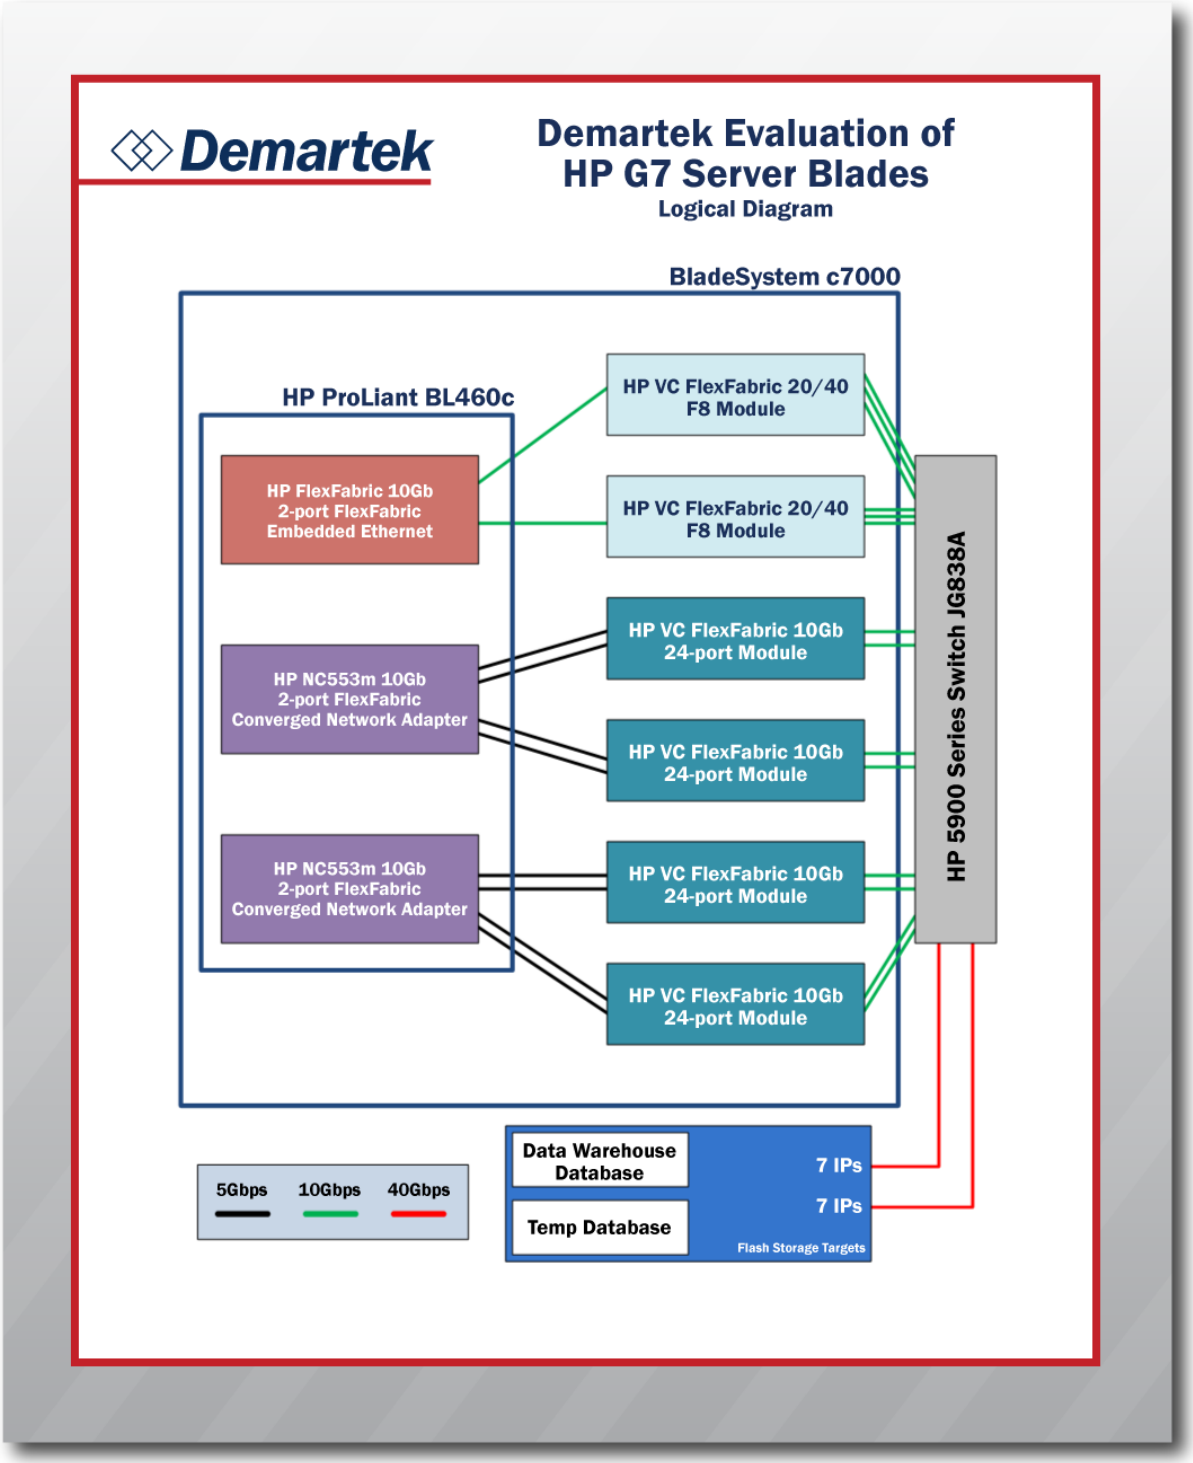

Two different types of VC modules were used in our setup, but all 6 ports on the G7 are 10Gb ports. The top two ports on the embedded NIC, shown above in red, connect to a 20Gb port on the HP VC FlexFabric 20/40 F8 module. However, since the embedded NIC is only capable of 10Gb, it does not take full advantage of the interconnect 20Gb bandwidth capability. The bottom two adapters, shown in purple, are 10Gb and connect to 10Gb on the switches, so the adapters take advantage of the full bandwidth of these 10Gb VC Modules.

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

- 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, four 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 and each 10Gb logical link also had a dedicated 10Gb uplink available to it. As we had eight 5Gb logical connections and two 10Gb logical connections, there was a total of 10x10Gb = 100Gb of uplinks available to the G7 server blade. 4 uplinks were not used for the G7 and were reserved for use on the Gen8 and Gen9 systems. Configuring Multi-Path IO (MPIO) on the server blade ensured that all server blade ports and all their corresponding uplinks were used.

- To ensure there was enough uplink bandwidth, more uplinks were allocated than what was necessary to support the server blade. As discussed before, the G7 server blade had 100Gb worth of uplinks available, but the G7 server blade only possessed a total of 6ports x 10Gb = 60Gb bandwidth capability.

Gen8 and Gen9 Server Blades

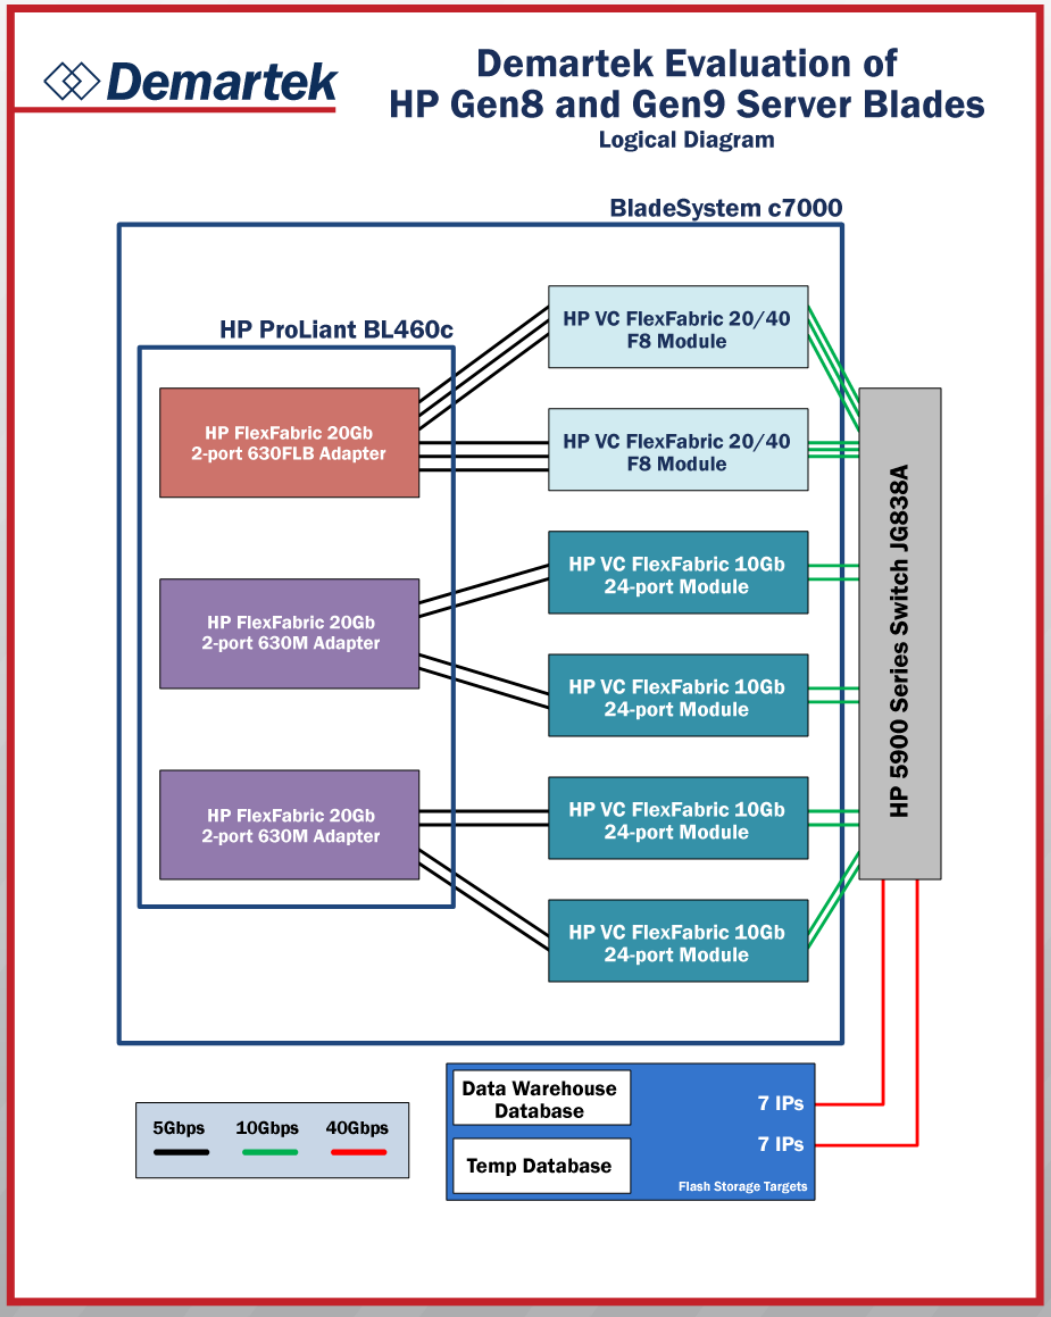

Two different types of VC modules used in our setup impacted our available bandwidth for the Gen8 and Gen9 server blades. Each FlexFabric 20Gb adapter has 2 ports of 20Gb. The top adapter, colored red, connects those ports to the 20Gb ports on the HP VC FlexFabric 20/40 F8 module, so it has 20Gb of throughput 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 blade. The Gen8 and Gen9 server blades each have a total of 2 x 20Gb + 4 x 10Gb = 80Gb of bandwidth capability. 14 x 10Gb uplinks made 140Gb of uplinks available to these server blades.
- 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, 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.

Gen9 vs Gen8 and G7 with Database Queries in Single-File

Three database workload tests were run on each server blade 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 resource that is maxed out for many queries is the processor; processor can be a bottleneck for many database workloads.

All server blades have a significant portion of their runs where all cores approach 100% utilization. The corresponding queries are noted on the graphs. Queries 18, 1, and 5 max out the processing capability for an extended period of time.

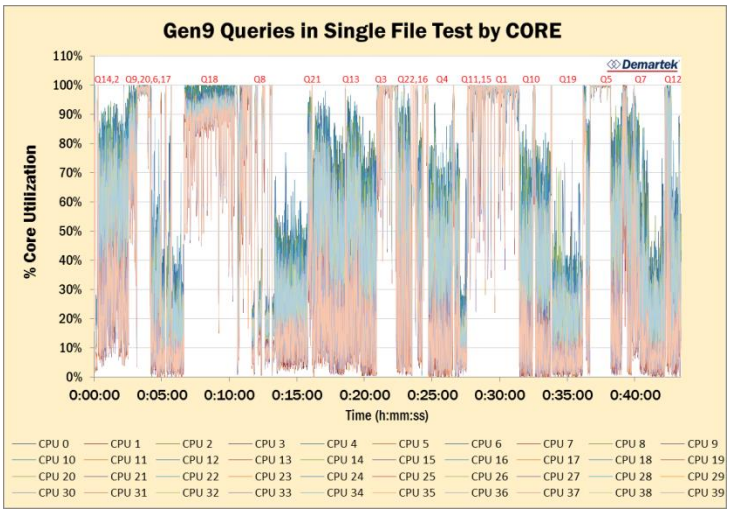
Gen9 Queries in Single File Test by CORE
% Core Utilization
Time (h:mm:ss)
Demartek
Q14,2 Q9,20,6,17 Q18 Q8 Q21 Q13 Q3 Q22,16 Q4 Q11,15 Q1 Q10 Q19 Q5 Q7 Q12
CPU 0 CPU 1 CPU 2 CPU 3 CPU 4 CPU 5 CPU 6 CPU 7 CPU 8 CPU 9
CPU 10 CPU 11 CPU 12 CPU 13 CPU 14 CPU 15 CPU 16 CPU 17 CPU 18 CPU 19
CPU 20 CPU 21 CPU 22 CPU 23 CPU 24 CPU 25 CPU 26 CPU 27 CPU 28 CPU 29
CPU 30 CPU 31 CPU 32 CPU 33 CPU 34 CPU 35 CPU 36 CPU 37 CPU 38 CPU 39

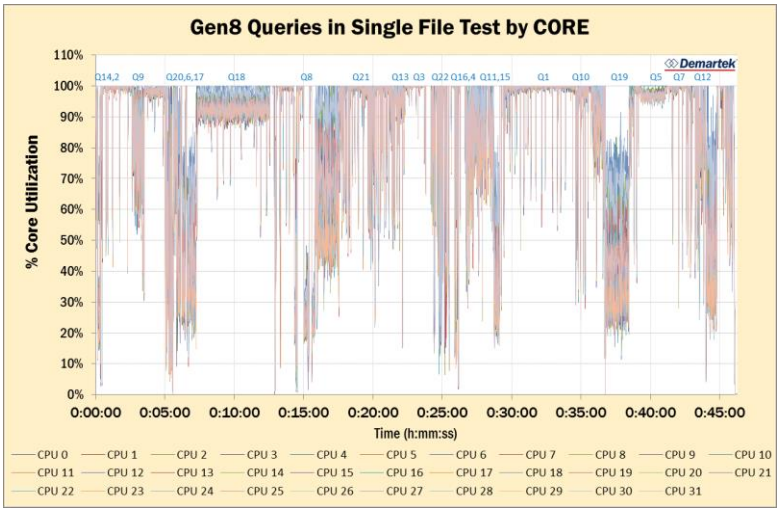
Gen8 Queries in Single File Test by CORE
% Core Utilization
Time (h:mm:ss)
Demartek
Q14,2 Q9 Q20,6,17 Q18 Q8 Q21 Q13 Q3 Q22 Q16,4 Q11,15 Q1 Q10 Q19 Q5 Q7 Q12
CPU 0 CPU 1 CPU 2 CPU 3 CPU 4 CPU 5 CPU 6 CPU 7 CPU 8 CPU 9 CPU 10
CPU 11 CPU 12 CPU 13 CPU 14 CPU 15 CPU 16 CPU 17 CPU 18 CPU 19 CPU 20 CPU 21
CPU 22 CPU 23 CPU 24 CPU 25 CPU 26 CPU 27 CPU 28 CPU 29 CPU 30 CPU 31

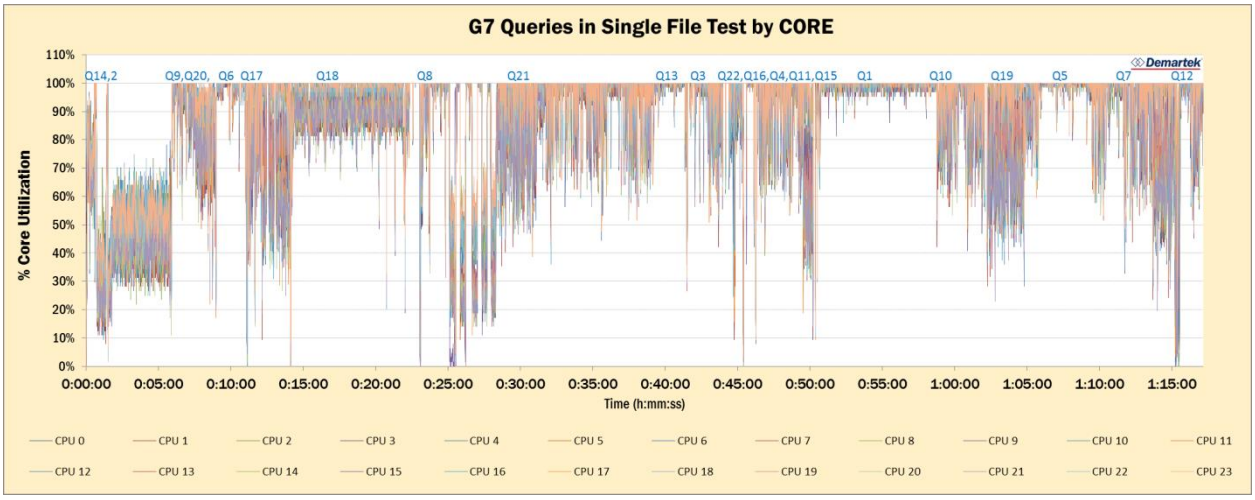
G7 Queries in Single File Test by CORE
% Core Utilization
Time (h:mm:ss)
Demartek
Q14,2 Q9,Q20, Q6 Q17 Q18 Q8 Q21 Q13 Q3 Q22,Q16,Q4,Q11,Q15 Q1 Q10 Q19 Q5 Q7 Q12
CPU 0 CPU 1 CPU 2 CPU 3 CPU 4 CPU 5 CPU 6 CPU 7 CPU 8 CPU 9 CPU 10 CPU 11
CPU 12 CPU 13 CPU 14 CPU 15 CPU 16 CPU 17 CPU 18 CPU 19 CPU 20 CPU 21 CPU 22 CPU 23

The Gen9 server blade spends less time at full CPU utilization. Queries 18, 1, and 5 push the Gen9 CPU to near 100% utilization just like in the Gen8 and G7, however with multiple cores working in parallel, the completion time for these processor intensive queries is reduced.

Completion Times for Select Processor Intensive Queries			
Completion Times (mm:ss)	**Gen9**	**Gen8**	**G7**
Query 18	3:58	5:40	8:56
Query 1	3:35	5:10	8:05
Query 5	3:38	4:29	6:59

This type of performance difference adds up over the course of a 22 query test:

Single-File Query Test Run Completion Times			
Completion Times (h:mm:ss)	**Gen9**	**Gen8**	**G7**
Test 1	0:43:03	0:52:19	1:17:10
Test 2	0:43:34	0:46:15	1:16:49
Test 3	0:43:24	0:46:12	1:16:27
Average	0:43:20	0:48:15	1:16:49

The tests were run three times on each system and the Gen9 server blade was consistently faster than the Gen8 server blade, by an average of 4 minutes and 55 seconds. In other words, the Gen9 was on average 10.2% faster than the Gen8. Even better, the Gen9 server blade took an average of 33 minutes and 28 seconds less, or was 43.6% faster, than the G7.

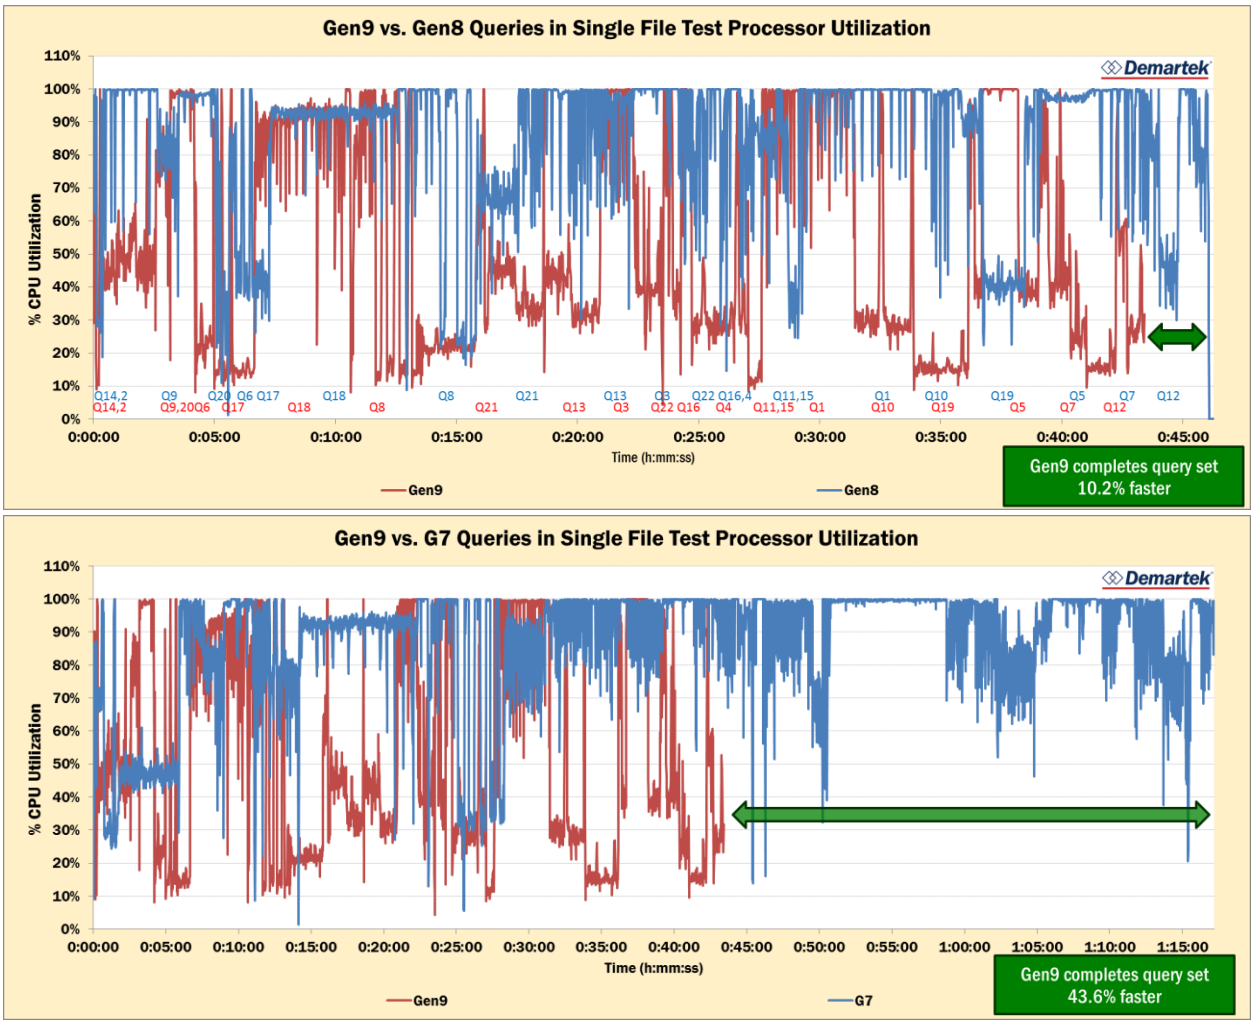

The average processor use over the entire test run was 87% for the G7 server blade, 85% for the Gen8 server blade and 53.5% for the Gen9 server blade. The Gen9 server blade still had plenty of processor capability to spare for other workloads.

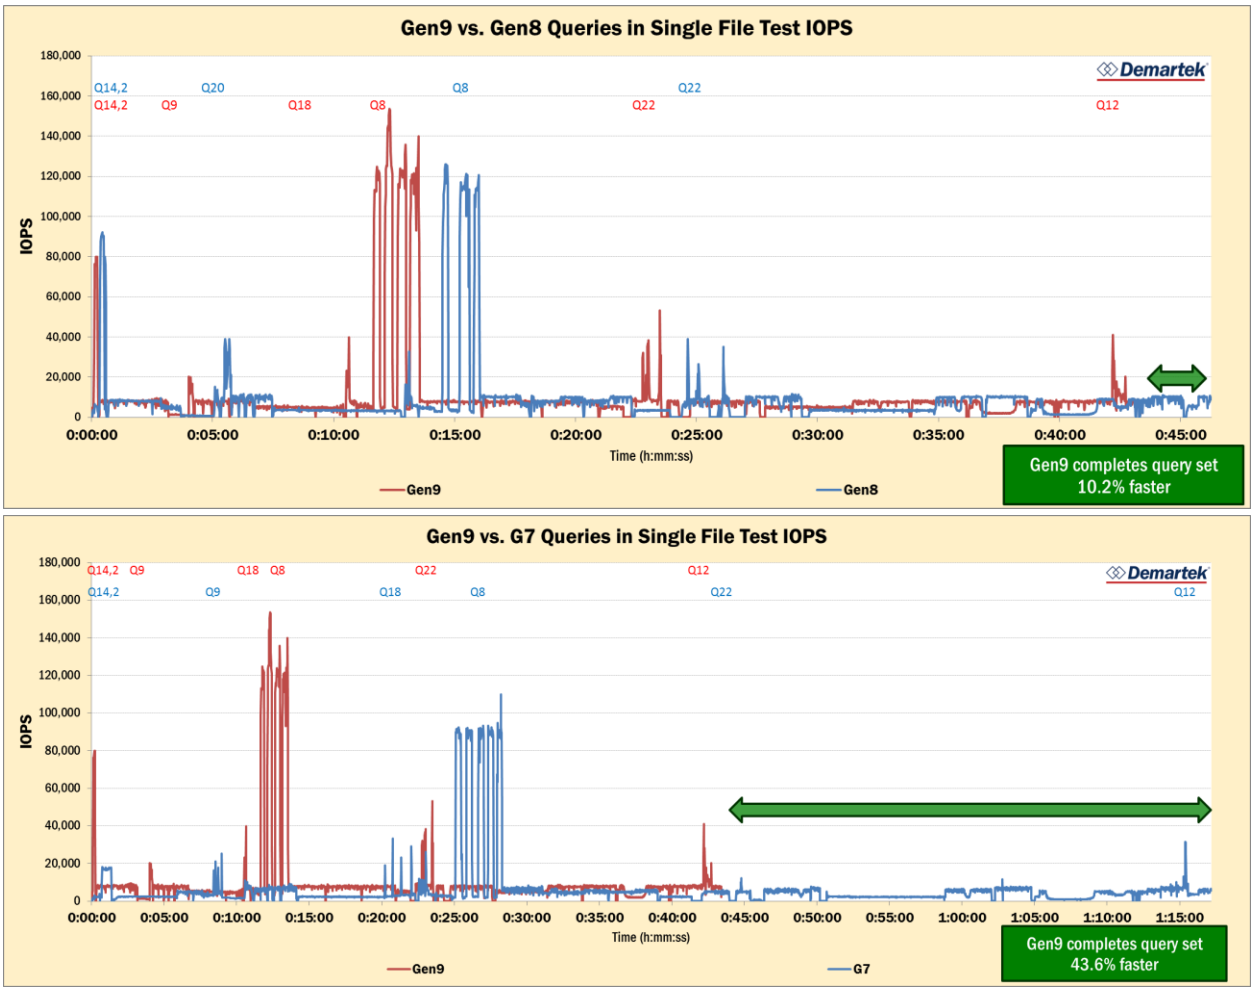

Query 8 stresses the IOPS. Note that the IOPS peaks are highest for the Gen9s and lowest for the G7s, most likely due to each generation having enhanced capability to process the data received as compared to the last.

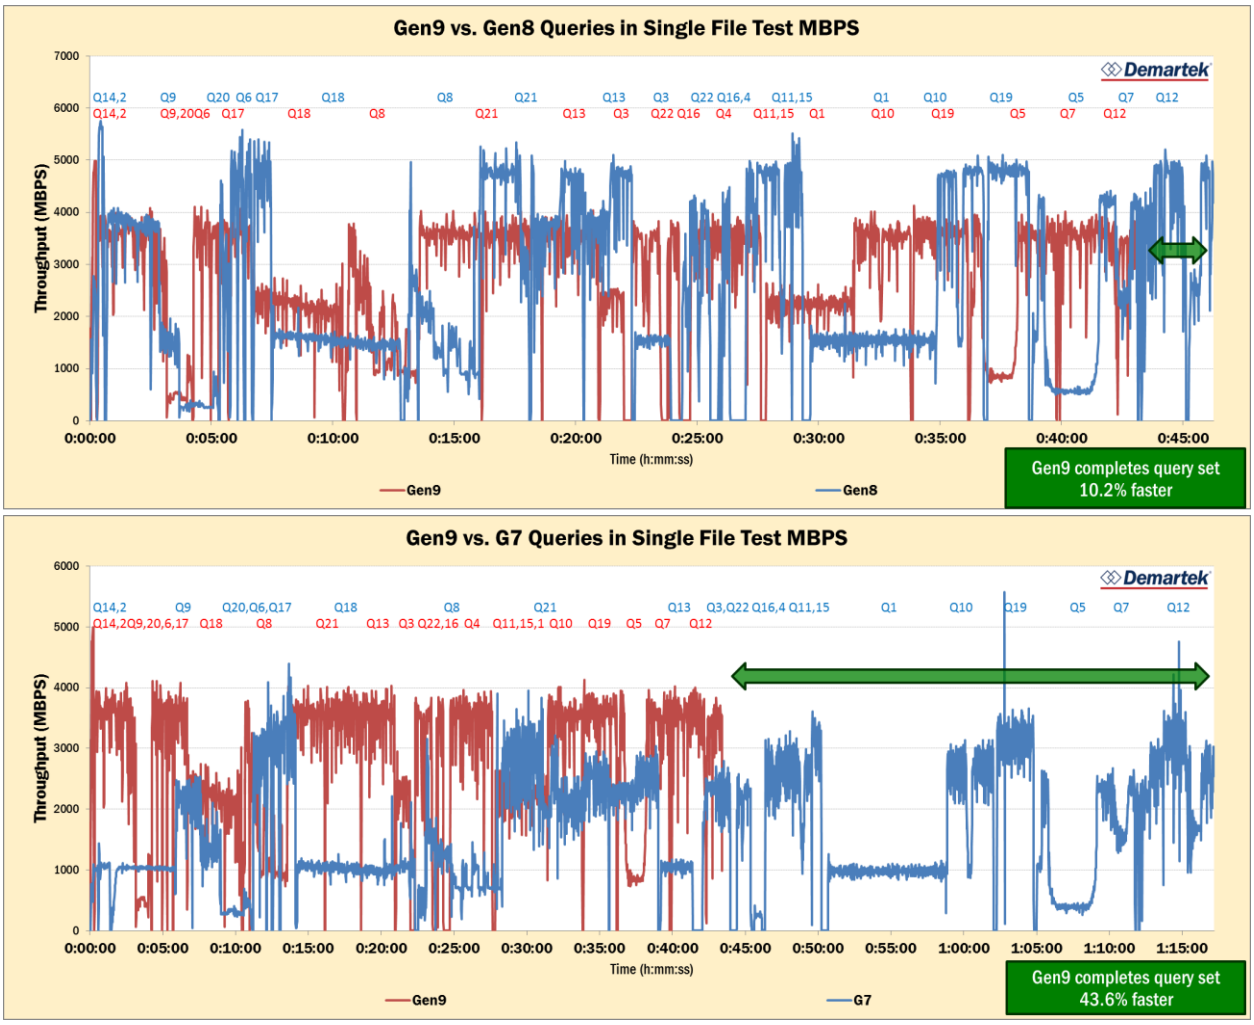

The results show that the Gen9 server blade is able to produce a more steady and predictable load to the storage subsystem, while the Gen8 server blade has many short, high peaks in throughput.

There are some isolated points where the G7 had throughput peaks similar to the Gen8, but for the most part the G7 had a hard time creating the same peaks seen in the Gen8, most likely due to reduced processor capabilities and reduced memory capacity making the G7 unable to drive the throughput for most queries.

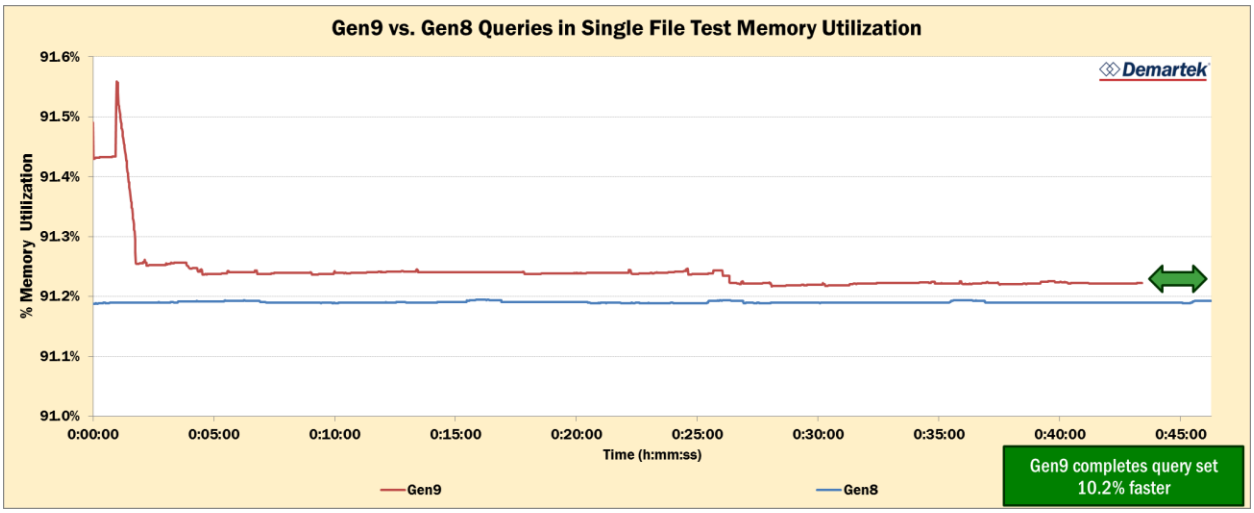

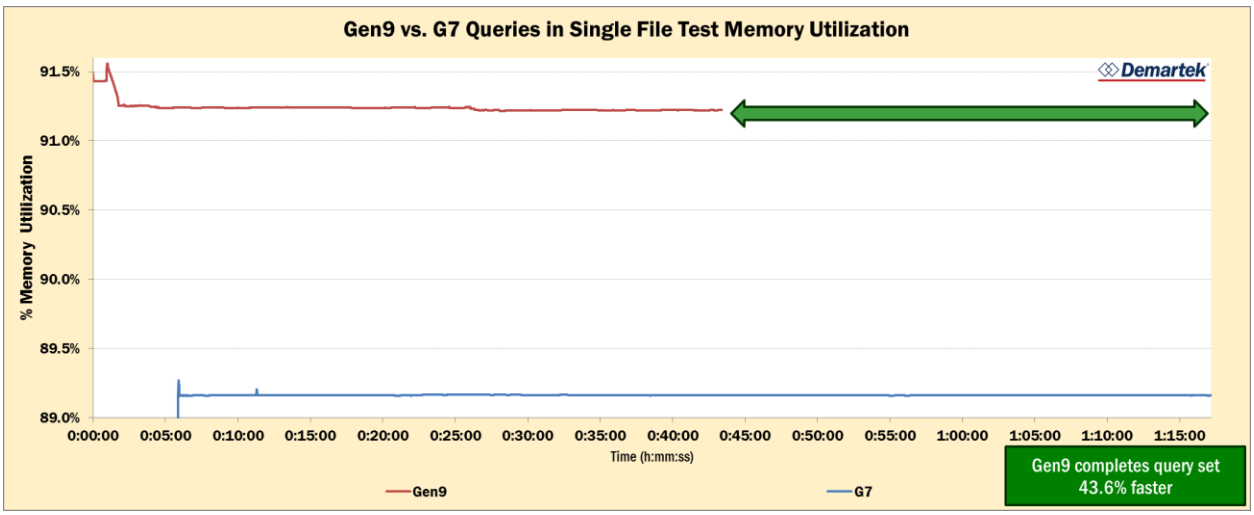

Both Gen9 and Gen8 systems had 16x32GB DIMMs, giving a total of 512 GB of memory available, and MSSQL Server was capped at 461GB (approximately 90% of available memory on each system). If we look at the scale of the graph, we can see that the memory utilization is mostly due to the 90% allocated to MSSQL Server. On the other hand, the G7 server blade had 12x32GB DIMMS, giving a total of 384 GB of memory available. MSSQL Server was capped at 372GB, a bit higher at 96%, but SQL was not able to use all of the memory allocated to it. This is most likely due to other memory use not recorded on the committed memory metric limiting the memory allocated to SQL.

Gen9 vs Gen8 and G7 Server Blades with Multiple Concurrent Queries

Three database workload tests were run on each server blade where 7 concurrent streams of 22 queries were run. A total of 154 queries are completed in each test, with seven queries running at any one time. All cores were utilized to a greater extent on both server blades. Again the G7 and Gen8 server blades maxed out the processor compute power more frequently than the Gen9 server blade.

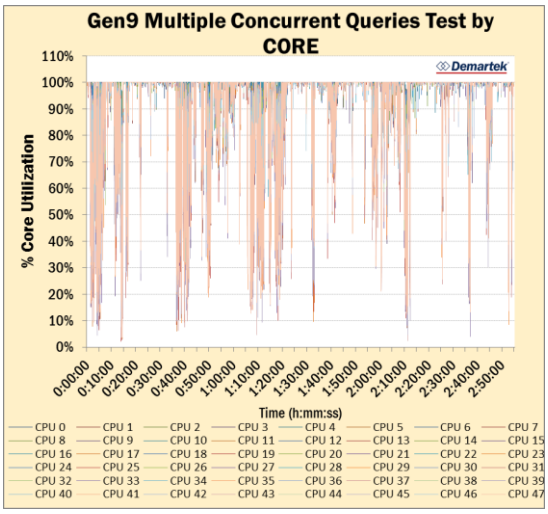

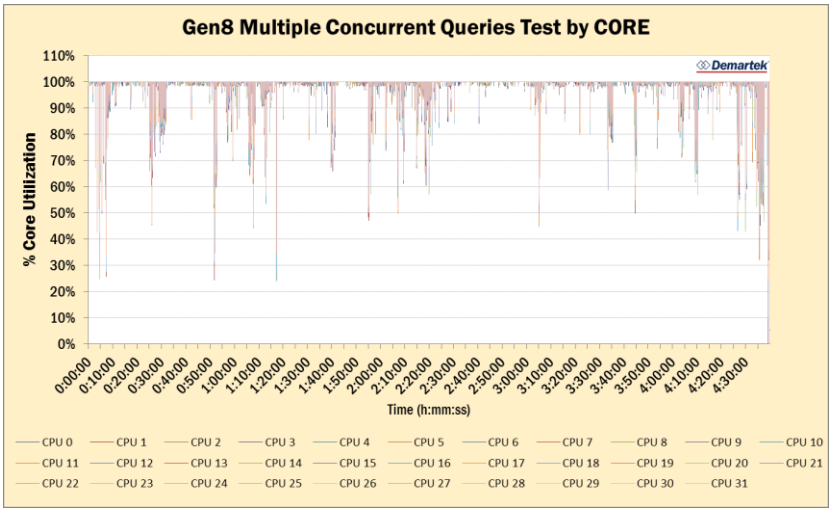

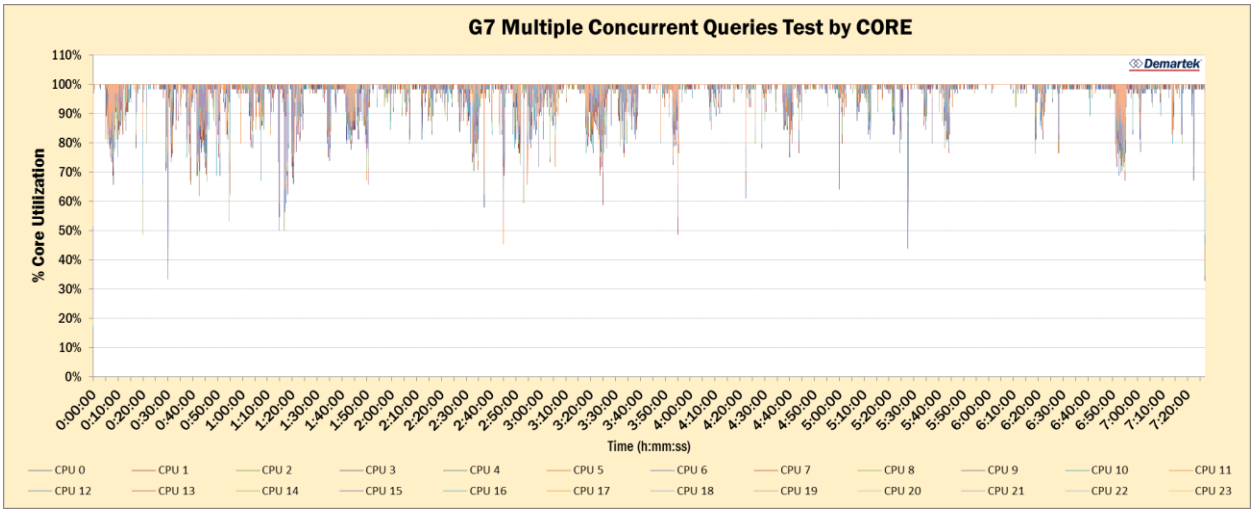

Running 7 queries concurrently ensures that there should be several processor-intensive queries available to take advantage of CPU resources at any point during the test.

The tests were run three times on each system and the Gen9 server blade was consistently faster than the Gen8 server blade, by an average of 1 hour and 25 minutes. This is most likely due to the increased number of cores available on the Gen9 system that enable the server blade to complete more tasks in parallel. Similarly the Gen9 server blade was consistently faster than the G7 server blade, by an average of 4 hours and 10 minutes. Again this is most likely due to the increased number of cores. In summary, the Gen9 completed same set of multiple concurrent queries 30.4% faster than the Gen8 and 56.2% faster than the G7. As stated before, the Gen9 reduced the G7 query completion time to less than half.

Multiple Concurrent Query Test Run Completion Times			
Completion Times	Gen9	Gen8	G7
Test 1	3:17:40	4:40:41	7:27:01
Test 2	3:12:32	4:39:55	7:29:43
Test 3	3:15:47	4:41:17	7:21:55
Average	3:15:20	4:40:38	7:26:13

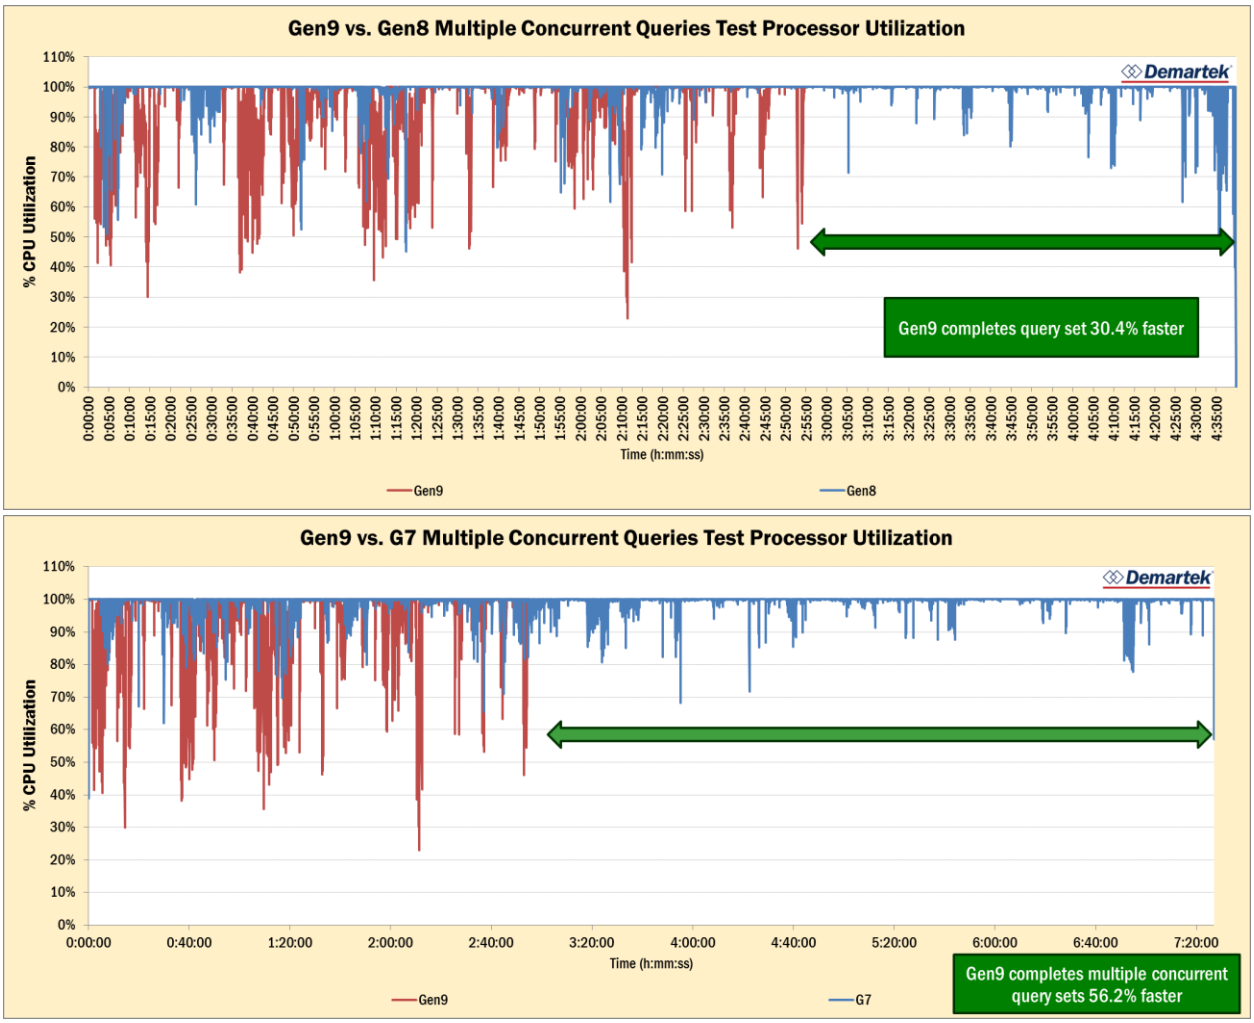

The average processor use over the entire test run was 99.7% for the G7 server Blade, 99.5% for the Gen8 server Blade and 94.6% for the Gen9 server Blade. The Gen9 server blade still had more processor capability that could be used for other workloads.

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

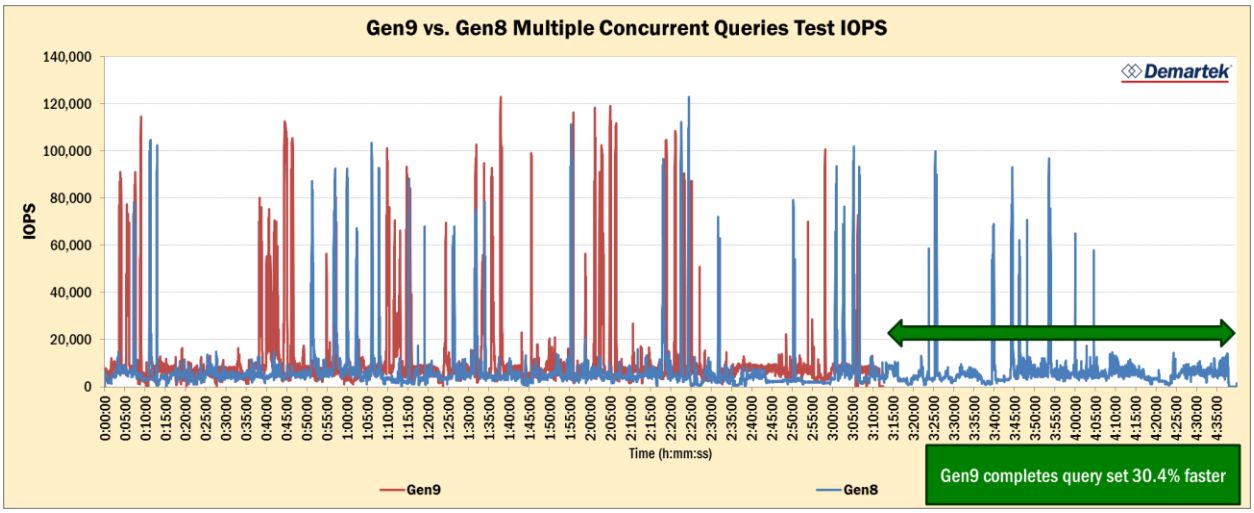
Gen9 vs. Gen8 Multiple Concurrent Queries Test IOPS
IOPS
Time (h:mm:ss)
Gen9
Gen8
Gen9 completes query set 30.4% faster

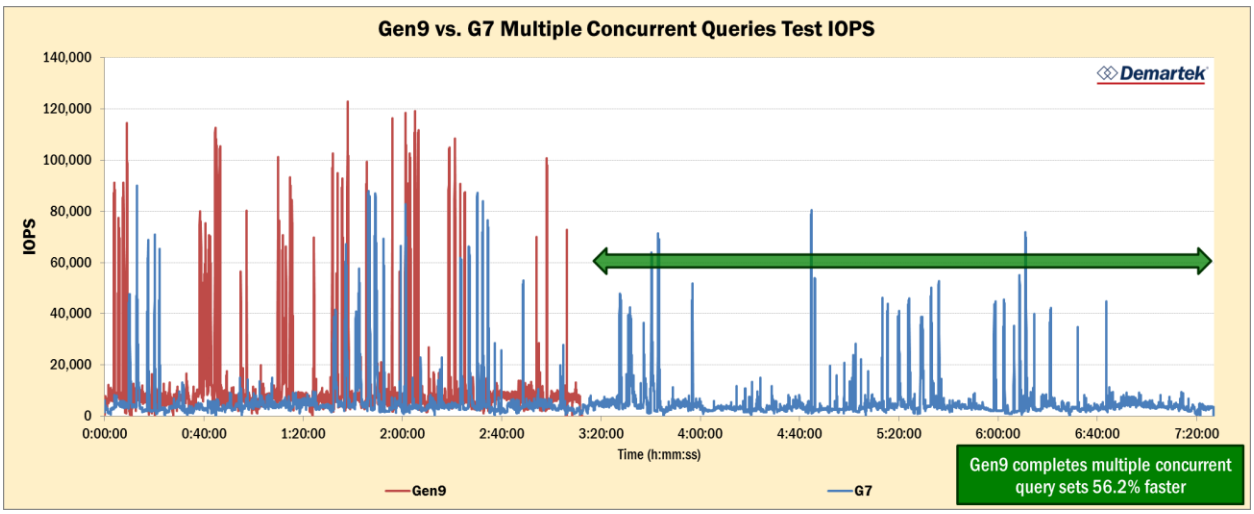
Gen9 vs. G7 Multiple Concurrent Queries Test IOPS
IOPS
Time (h:mm:ss)
Gen9
G7
Gen9 completes multiple concurrent query sets 56.2% faster

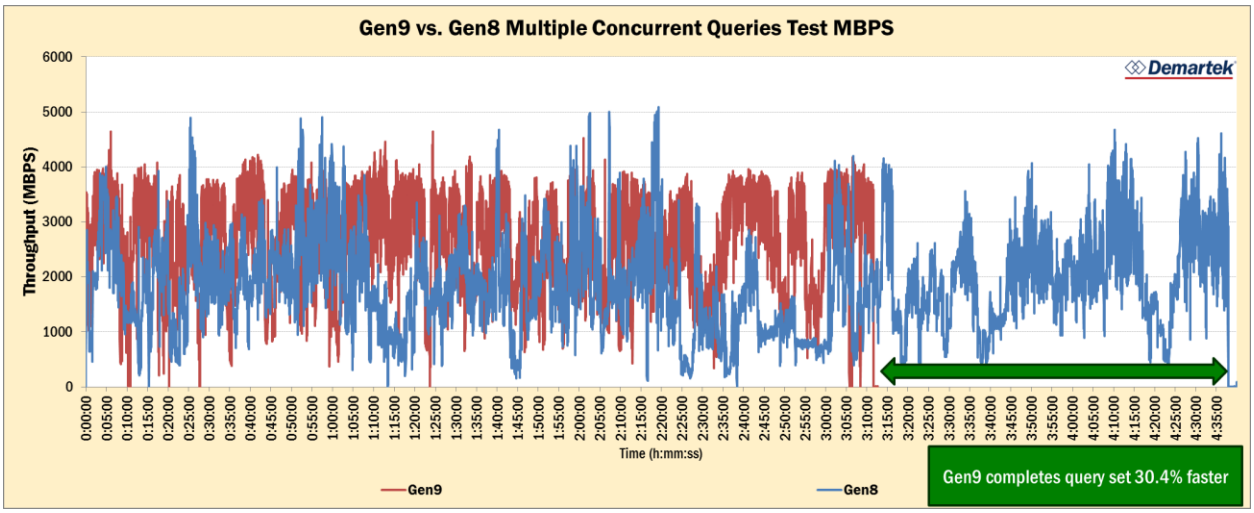
Gen9 vs. Gen8 Multiple Concurrent Queries Test MBPS
Demartek
Throughput (MBPS)
6000
5000
4000
3000
2000
1000
0
Time (h:mm:ss)
Gen9
Gen8
Gen9 completes query set 30.4% faster

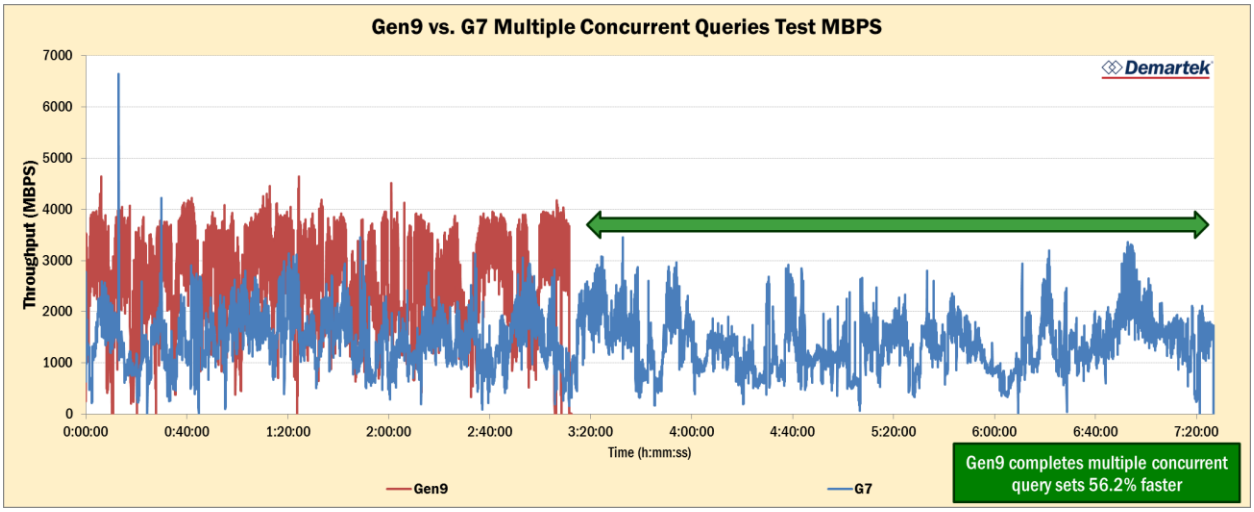
Gen9 vs. G7 Multiple Concurrent Queries Test MBPS
Demartek
Throughput (MBPS)
7000
6000
5000
4000
3000
2000
1000
0
0:00:00
0:40:00
1:20:00
2:00:00
2:40:00
3:20:00
4:00:00
4:40:00
5:20:00
6:00:00
6:40:00
7:20:00
Time (h:mm:ss)
Gen9
G7
Gen9 completes multiple concurrent query sets 56.2% faster

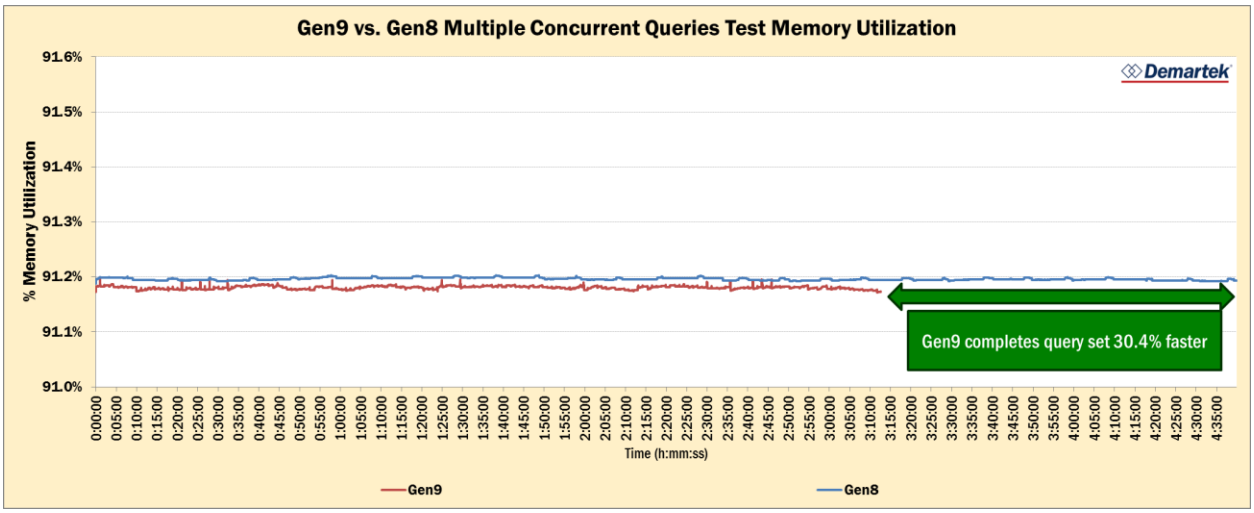
Gen9 vs. Gen8 Multiple Concurrent Queries Test Memory Utilization
Demartek
% Memory Utilization
91.6%
91.5%
91.4%
91.3%
91.2%
91.1%
91.0%
Time (h:mm:ss)
Gen9 completes query set 30.4% faster
Gen9
Gen8

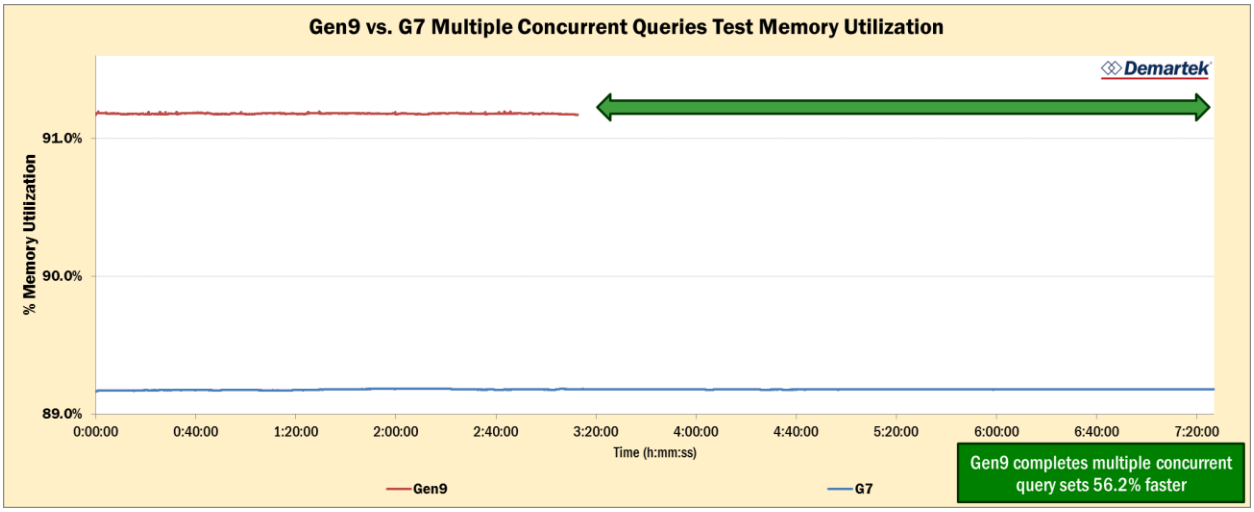
Gen9 vs. G7 Multiple Concurrent Queries Test Memory Utilization
Demartek
% Memory Utilization
91.0%
90.0%
89.0%
0:00:00
0:40:00
1:20:00
2:00:00
2:40:00
3:20:00
4:00:00
4:40:00
5:20:00
6:00:00
6:40:00
7:20:00
Time (h:mm:ss)
Gen9 completes multiple concurrent query sets 56.2% faster
Gen9
G7

Summary and Conclusion

The HP ProLiant BL460c Gen9 Server Blades outperformed the HP ProLiant BL460c Gen8 Server Blades and HP ProLiant BL460c G7 Server Blades, especially in multi-threaded database throughput tests. Should an IT manager need to purchase more server blades to handle an increased database workload, the Gen9 server blades would be a smarter choice. The Gen9 server blade comparative configuration hardware costs 9% less than the Gen8 and offers a 30% increase in performance. Even better, the Gen9 server blade comparative configuration hardware costs 15% less than the G7 server blade and offers a 56% increase in performance.

Performance improvements drive more savings. If we quadrupled our workload to include 28 streams and 616 total queries and split those queries between 3 HP ProLiant BL460c Gen9 Server Blades, it would take approximately 4 hours and 20 minutes to complete. It would take 4 HP ProLiant BL460c Gen8 Server Blades to complete the same workload in under 5 hours (4:40). It would take 6 HP ProLiant BL460c G7 Server Blades to complete the same workload in under 5 hours (4:57).

Putting some of this into quantifiable terms:
Switching from Gen8 to Gen9:

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

Switching from G7 to Gen9:

- Saving of time per database run = 251 minutes
- Doing two database runs per day = 502 minutes of time savings per day (over 8 hours).

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 on a switch from Gen8 to Gen9 server blades. That is over a man-year's worth of operational expense (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.
- 502 minutes/day X 1095 days = 549,690 minutes or 9161.5 hours/database server on a switch from G7 to Gen9 server blades. Using the same statistics as above, that is

over 4.5 years of processing time per database server, and 3 times the OPEX savings to over $150,000+ per database engineer that would be available for use doing other revenue generating tasks. Also contributing to the case for replacing G7 server blades with Gen9 server blades, is the fact that the G7 server blades are past their obsolescence date and are at the end of their server lifecycle.

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_Gen9_vs_Gen8_and_G7_Evaluation_2015-09.html on the Demartek website.