LiquidCool SOLUTIONS

Power-to-Cool

LIQUID SUBMERSION COOLED SERVERS VS. AIR-COOLED SERVERS

LiquidCool LSS220 compared to Dell™ PowerEdge R620

A White Paper by LiquidCool Solutions, Inc
ANDY REGIMBAL | RICK TUFTY

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY. IT MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

LiquidCool is a trademark of LiquidCool Solutions, Inc. Dell and PowerEdge are trademarks of Dell Inc. Microsoft, Windows, and Windows Server are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. SPEC and the benchmark name SPECpower are registered trademarks of the Standard Performance Evaluation Corporation. For more information about SPECpower, refer to www.spec.org. Benchmark results stated in this white paper reflect results run as of December 14, 2012. For the latest SPECpower_ssj2008 benchmark results, visit http://www.spec.org/power_ssj2008/results/power_ssj2008.html

Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. LiquidCool Solutions, Inc. disclaims any proprietary interest in trademarks and trade names other than its own.

First Revision, January 2013

FOR MORE INFORMATION

Contact LiquidCool Solutions
2717 Highway 14 West, Suite D
Rochester, MN 55901
Phone: 1-866-530-0101

CONTENTS

ABOUT THE AUTHORS

Andy Regimbal is a Technology Development Engineer at LiquidCool Solutions. Andy has 25 years experience leading the development of products for the enterprise printing, server, and storage markets. He has held technical leadership positions at Dell, HP, and Tektronix and is responsible for over 30 US patents.

Rick Tufty is Vice President of Engineering at LiquidCool Solutions. Rick has 30 years of engineering and technology research leadership experience with Tier One server and storage companies including Dell, Maxtor, Western Digital, HP, and IBM. He holds 12 issued patents with several additional patent applications pending and has authored numerous technical papers.

ACKNOWLEDGEMENTS

The authors wish to thank the following individuals at LiquidCool Solutions for their important contributions to this white paper.

Daryl Lautenschlager
Lead Test Engineer, for configuring the software utilities on the test systems, running the tests, and logging the results.

David Roe
PMP, Program Manager, for developing and managing the project plan and for editing the white paper.

Gary Reed
Principal Electrical Engineer, for configuring and installing the power monitoring equipment used in the test.

Sean Archer
Development Engineer, for his contribution to developing the test plan and for designing and instrumentation of the test setup.

EXECUTIVE SUMMARY

Recent advances in high-density computing are creating new challenges for cooling servers. Air-cooled servers in particular must resort to more robust cooling systems to deal with increased thermal loads, which are also more power-hungry. As a result, the "power-to-cool," defined as the power required to cool a server's internal components, is now a significant portion of the total power consumed by a typical air-cooled server.

LiquidCool™ Solutions' patented liquid submersion technology offers a far more energy efficient cooling solution for servers than air-cooled servers can provide. This white paper documents the results of a benchmark study undertaken to compare the power-to-cool efficiency of a server employing LiquidCool's liquid submersion technology with that of a representative "best-of-breed" air-cooled server.

The results show that LiquidCool's cooling technology reduces server power-to-cool by **up to 98%** when the server is operating under heavy compute load, depending on ambient air temperature. At the same time, the CPUs in servers cooled by LiquidCool technology were typically found to operate at temperatures **20°C to 30°C** lower than the air-cooled CPUs. These lower CPU junction temperatures improve both the energy efficiency and the projected long-term reliability of LiquidCool servers compared to air-cooled servers.

Overall, the study shows that a server cooled with LiquidCool technology **uses up to 26% less total power**[1] than a similarly configured air-cooled server when operating at higher workload levels. In addition, LiquidCool technology enables a server to operate at peak computational levels at inlet-coolant temperatures of up to 56°C irrespective of ambient air temperature.

By leveraging the dramatic improvements in energy efficiency documented in this study, datacenter implementations using LiquidCool technology can virtually eliminate power-to-cool at the server level while maintaining maximum computational power without the need for chilled air or water. When deployed on a large scale, the resulting energy saving opportunities are significant.

[1] This estimate is based on reduced power-to-cool plus power savings due to reduced leakage current. It assumes the power efficiency of each server's respective electronic components to be equal under the same operating conditions.

INTRODUCTION

Almost all servers in data centers today use forced-air convection as the primary means of cooling the electronics within the server. This cooling approach has evolved to greater complexity as electronic power densities have increased with each succeeding generation of silicon. These heat loads require multiple high-speed fans within the server to force cooled ambient air from the data center through the server enclosure and remove heat from the electronics. The electrical power needed to run these fans, referred to as the “power-to-cool,” often represents a significant percentage of a server’s total power usage. In a typical data center deploying hundreds or thousands of servers, the energy cost of the aggregate power-to-cool can be enormous.

In response to these cost and energy challenges, new technologies are emerging to supplant air cooling as the primary heat removal mechanism in servers. LiquidCool Solutions’ patented liquid submersion cooling technology[2] is one such innovative approach, employing a dielectric liquid , called CoreCoolant, to cool the electronic components within the server enclosure. Because liquid submersion–cooled servers require a low level of remotely provided pump energy to move the dielectric liquid through the server enclosure, eliminating the need for fans, the power-to-cool and the associated energy costs are dramatically reduced.

This white paper documents the results of an extensive study comparing the performance and power consumption of four air-cooled Dell PowerEdge R620 servers with four comparably configured LiquidCool Solutions (LCS) LSS220 servers employing liquid submersion technology.

It will be demonstrated that:

1. The power-to-cool associated with modern server system level fans can approach 20% of the actual power consumption of each server depending on workload and data center ambient air temperature conditions, and that adoption of LiquidCool technology can reduce this energy demand by up to 98% relative to comparably configured air-cooled systems.

2. Server computing power can be reduced by up to 13% due to the reduction in leakage current (static power). This is due to the lower system temperature and the associated reduced electrical resistance provided by liquid submersion cooling.

3. LCS cooling technology allows servers to operate at full computational power, even when deployed with high ambient and inlet-coolant temperatures.

The energy saving opportunities enabled by LiquidCool technology do not stop at the server level. Because heat energy is removed from the systems in liquid form, the need for data center air conditioning and air handling provisioning is greatly reduced. This can result in phenomenal cost savings in capital investment and total cost to operate (TCO). Additionally, the waste heat removed in the fluid can be captured and used for constructive purposes at the facility level rather than simply be ejected to the external environment. Finally, it should be noted that this fundamental technology can be employed to cool other forms of high-power-density electronics, including switches, routers, high performance memory systems, and UPS units. As the technology gains acceptance and as application efforts expand, data center power requirements can be reduced even further.

[2] Refer to Appendix D: Background Information for more detail on LiquidCool submersion cooling technology.

SUMMARY OF TEST RESULTS

Tests Conducted

The benchmark study included a series of tests comparing the performance of LiquidCool LSS220 servers with similarly configured Dell PowerEdge R620 servers. The primary software utility used for the tests was SPECpower_ssj2008, an industry standard benchmarking tool. All test runs discussed in this white paper are "compliant" per SPECpower_ssj2008 definitions. At time of publication of this white paper test results have not yet been submitted to SPEC, but are valid for disclosure and discussion. A detailed description of SPECpower and the test methodology employed can be found in Appendix A. Additional software utilities were used for specific portions of the testing. These utilities provided for higher stress levels than SPECpower is capable of. Additional information about these utilities are included in Appendix A.

Unless otherwise indicated, testing was conducted in an environmental chamber to precisely control the ambient temperature. Each test was repeated at specific ambient air temperatures between 15°C and 45°C to collect data on the effect of ambient temperature on each server's performance, power usage and critical component temperatures.

Individual tests focused on the following performance elements:

Computational Performance The primary function of this test was to operate a server at defined computational loads while measuring the server's power consumption. Ambient temperature was also recorded during the tests to ensure results can be compared on a fair basis.

Power Consumption This test measured the total power consumed by each server at the various SPECpower loads and at various ambient temperatures. The purpose of this test was to determine how each server's power usage is affected by changes in computational load and ambient temperature.

Power-to-Cool This test measured the power consumed by cooling mechanisms in each server at the various power-to-compute loads and at various ambient temperatures. For the Dell servers, the power used by the cooling fans was measured. For the LiquidCool servers, the power used by the pump which circulates the coolant was measured.

Computational Performance vs. Ambient Temperature This test measured the computational performance of each server under the SPECpower loads and at various ambient temperatures. In this case, the ambient temperature of each server's cooling medium was regulated: air temperature for the Dell server and coolant temperature for the LCS server. The purpose of this test was to determine the effect of increasing ambient temperature on performance.

Reduction in Leakage Current This test measured the increase in processor core temperatures and the corresponding increase in power consumption of the LCS server at a fixed computational load as the coolant temperature was increased. The results were then used to calculate LCS power savings based on data recorded showing that LCS cooling keeps processors cooler than the air-cooled Dell servers under the same ambient conditions.

Key Findings

Following are the key findings from the tests conducted. Refer to the appendices for details on the testing methodology used, the test setup, and specific results of the tests conducted.

- ➢ LCS Liquid Submersion Technology reduced system level power-to-cool and associated energy costs by as much as 98% compared to air cooling.
- ➢ Computational performance of the LCS servers was better than or equal to the Dell servers at all ambient temperatures and all computational loads.
- ➢ The LCS servers maintained peak computational performance at ambient air temperatures between 15°C and 45°C and at delivered coolant temperatures up to 56°C. Performance of the Dell servers, however, was dramatically reduced at ambient air temperatures above 41°C.
- ➢ The LCS power-to-cool remained constant at very low power usage regardless of computational load or coolant temperature between 15°C and 45°C. The power-to-cool of the Dell servers increased significantly as computational load and ambient air temperature increased.
- ➢ Under the same ambient temperature conditions and computational load, LCS servers consistently maintained processor core temperatures 20°C to 30°C cooler than the Dell servers.
- ➢ An increase of 13% in total system power consumption associated with component leakage currents was measured as CPU core temperatures increased from 50°C to 88°C.
- ➢ Ambient air temperatures have no discernible effect on operational characteristics of LCS servers, which rely only on the temperature and flow rate of the liquid coolant for thermal management.

CONCLUSIONS

The study is conclusive in validating several clear advantages of LiquidCool submersion technology to cool servers compared to air cooling.

Computational performance is better than or equal at all ambient temperatures

Liquid submersion cooling provides more consistent computational performance than air-cooled servers when operating across a broad range of ambient temperatures. As the test results clearly demonstrate, the computational performance of the LCS LSS220 servers matched or exceeded the performance of the Dell R620 servers under all ambient test conditions supported by the air cooled systems. This is illustrated in **Figure 1** below, which compares LCS and Dell server performance at an ambient air temperature of 30°C. Note that in this case the delivered coolant temperature to the LCS server is higher than the ambient temperature and yet performance still matches that of the Dell server.

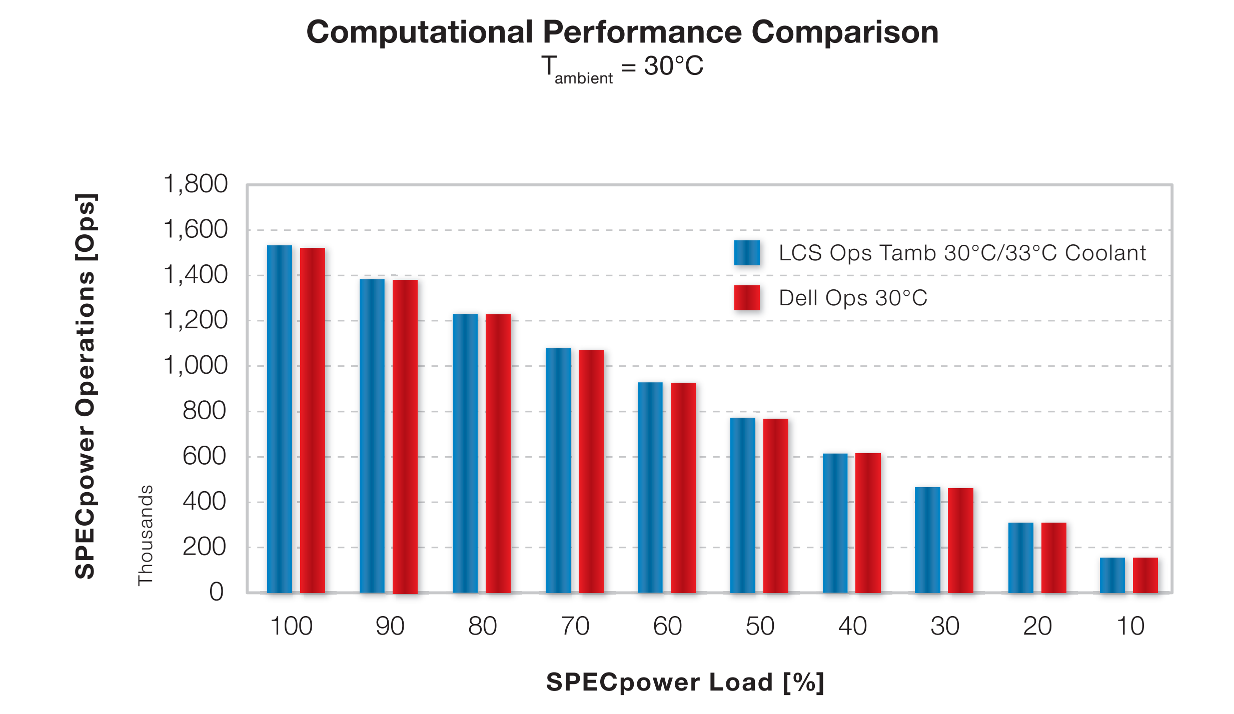

FIGURE 1 LCS vs. Dell Computational Performance Comparison at 30°C

Power-to-cool reduced by up to 98%, or as much as 18% of total server power usage

By eliminating the need for high-speed fans, the adoption of LCS submersion cooling technology can reduce server power-to-cool by up to 98%. This is illustrated in the figure below, which shows a 75W reduction in power-to-cool at computational loads of 350W and higher. This reduction represents as much as 18% of total server power.

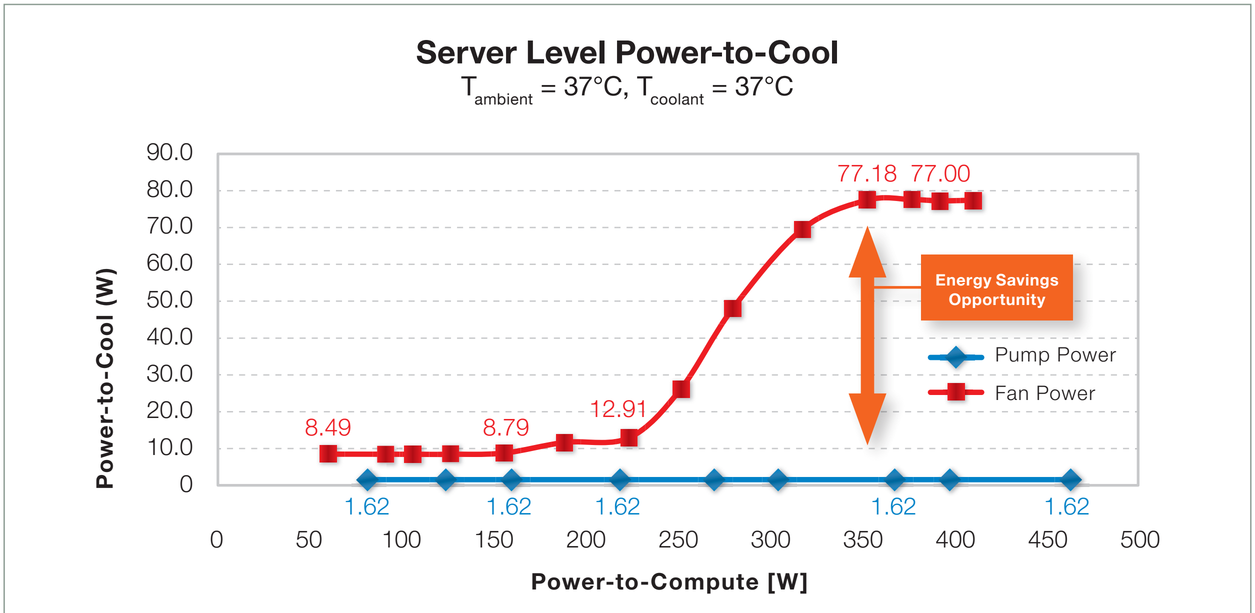

FIGURE 2 Server Power-to-Cool Comparison at 37°C

Leakage current reduction leads to additional power savings of up to 8%

As shown in the figure below, power usage due to leakage current increases by 13% at constant computational load as processor temperatures increase from 50°C to 88°C. Because processor temperatures in LCS servers are kept 20°C to 30°C cooler than in air-cooled servers at equal computational load and ambient conditions, power usage due to leakage current is reduced by up to 8%. This savings is **in addition** to savings due to reduced power-to-cool.

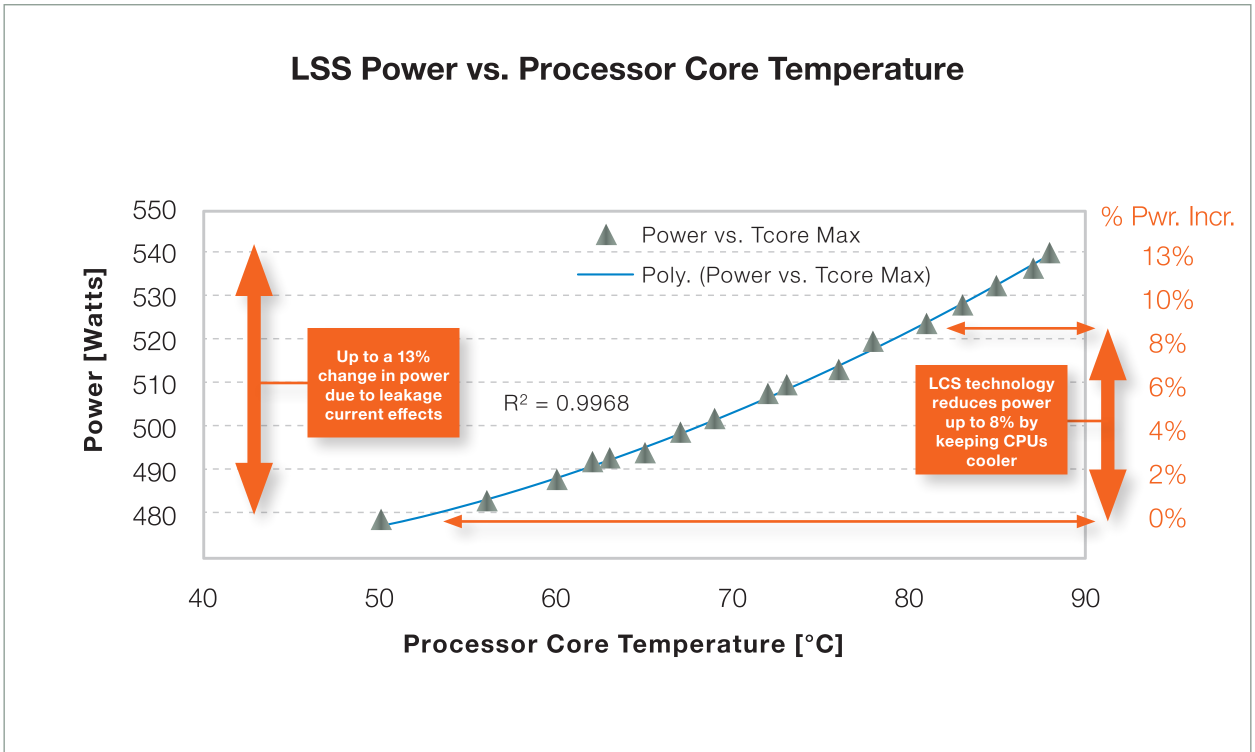

FIGURE 3 Effect of Leakage Current on Total Power Usage

By combining the energy savings due to reduced leakage current with the savings from reduced power-to-cool illustrated in Figure 2, a total reduction in server power use of up to 26% can be achieved.

Performance does not change regardless of ambient temperature or coolant temperature up to 56°C

Unlike air-cooled servers, which are subject to CPU throttling under heavy computational workload and high ambient temperature, servers cooled by LCS technology deliver consistent computational performance regardless of ambient air temperatures[3] or coolant temperatures up to 56°C.

This is illustrated in the following two figures. Figure 4 shows that the Dell R620 servers have throttled under an ambient temperature condition of 45°C. Meanwhile, the LSS220 servers continue to operate at peak performance at this same ambient temperature. In addition, as depicted in Figure 5, computational performance of the LSS220 servers remains consistent as the temperature of delivered coolant ranges from 22°C to 56°C.

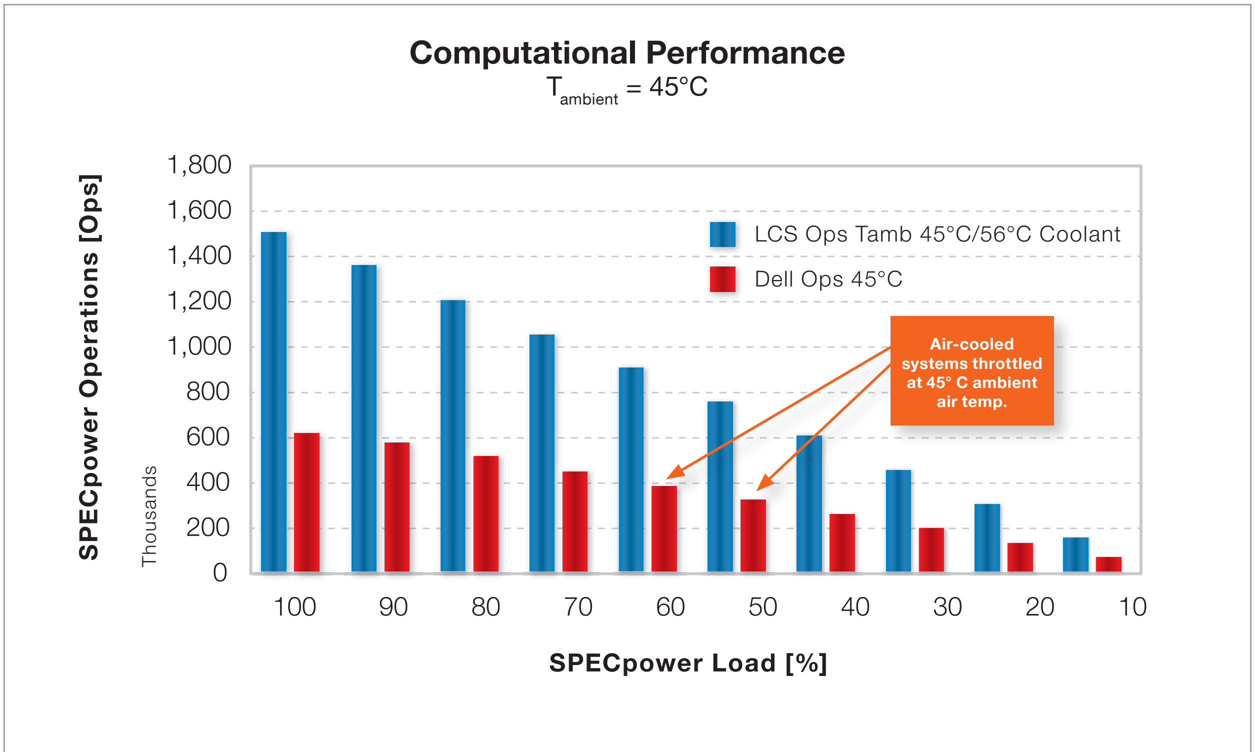

FIGURE 4 LCS vs. Dell Computational Performance Comparison at 45°C

[3] Testing at ambient temperatures above 45°C was not in the scope of this study, However, in other testing done by LCS it has been shown that ambient temperatures up to 65°C have no effect on server performance. Contact LiquidCool Solutions for details.

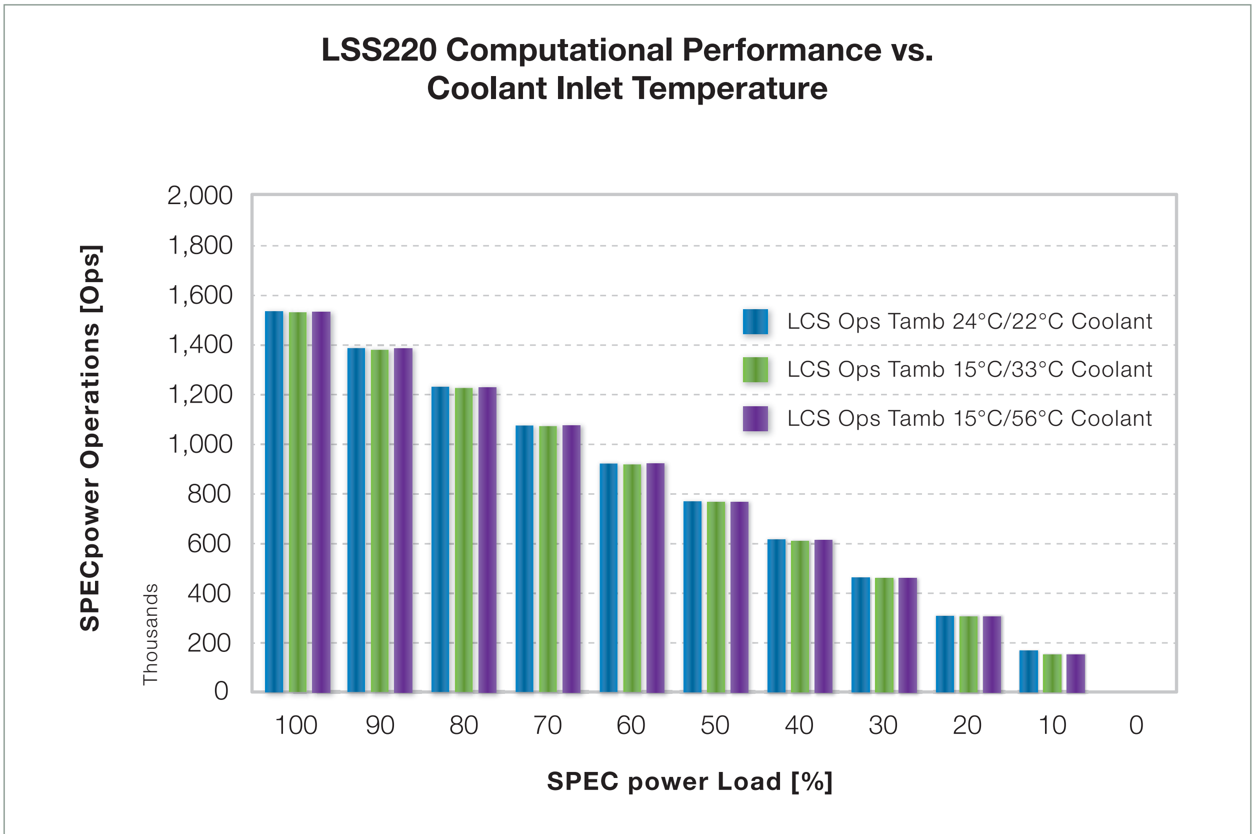

FIGURE 5 LSS220 Computational Performance vs. Coolant Temperature

Additional Benefits

In addition to the advantages of LiquidCool's submersion cooling technology at the device level that have been demonstrated in this study, LCS technology provides tremendous opportunities to further reduce energy consumption at the data center level. For example, because LCS-cooled devices are not constrained to strict ambient air temperatures, data center operators can make tradeoffs to better allocate power to the computing aspects of the operation rather than to cooling. These opportunities will be the subject of future white papers.

APPENDIX A - TEST METHODOLOGY

Test Software

Several software utilities were used to generate the results described in this white paper. These utilities are described briefly below. Links to additional information and access to the software are also provided.

Software for System Performance vs. Power Consumption Measurements

SPECpower_ssj2008 was used to develop Computational Performance vs. Power Consumption baselines for each of the Systems under test. SPECpower is an industry standard benchmarking tool. A general description of the tool and its functionality as described by SPEC states:

> *SPECpower_ssj2008 is a benchmark product developed by the Standard Performance Evaluation Corporation (SPEC), a non-profit group of computer vendors, system integrators, universities, research organizations, publishers and consultants. It is designed to provide a view of a server system's power consumption running Java server applications.*
>
> *The general approach is to compare measured performance with measured power consumption. An initial requirement, as well, was to include power measurement data of a system running at different target load levels, to reflect the fact that data center server systems run at different target loads relative to maximum throughput.*
>
> *SPEC also considers the ambient temperature during the benchmark measurement relevant to the results, and so temperature measurement is a requirement as part of a full benchmark report.*[4]

The links below provide access to the complete SPECpower_ssj2008 Design Overview document as well as the SPECpower_ssj2008 home page.

SPECpower_ssj2008 Design Overview:
http://www.spec.org/power/docs/SPECpower_ssj2008-Design_Overview.pdf

SPECpower_ssj2008 Site:
http://www.spec.org/power_ssj2008/

[4] **Standard Performance Evaluation Corporation.** Design Overview, SPECpower_ssj2008. http://www.spec.org/power/docs/SPECpower_ssj2008-Design_Overview.pdf. [Online] 2012.

System Stress Software Utilities

During testing it was determined that SPECpower does not stress the systems under test at a maximal level. Because of this, additional testing was performed using two different "stress" tools. Both of these tools stress the processors and memory to peak levels, but do not provide performance metrics.

The focus of this stress testing was to control the servers in a repeatable manner so they operated at specific power consumption levels across their full operating range.

- **Prime95** A number of tests determined that using the custom option and adjusting the size of FFTs used was the best way to maximize power consumption. The final setting was determined to be FFTs set at 4,096 and FFTs run in place.

- **BurnInTest** It was determined that configuring BurnInTest to exercise the CPU only was the best way to control the power consumption in the desired range. The power consumption was recorded at CPU utilizations varying from 0% to 100% in 5% increments.

The power consumed at each setting for each utility was measured and recorded. The power consumption values ranged from 59W to 462W per server. The same utility and settings were used to test the LCS and Dell servers across this power range to generate the data presented in the Power-to-Cool section of this white paper.

The links below provide access to the software utilities used:
Prime95 – Version26.6 build 3 64 bit
http://www.mersenne.org/freesoft/

BurnInTest – Version 6.0 Pro (1006) 64 bit
http://www.passmark.com/products/bit.htm

System Configuration and Settings for SPECpower Testing

SYSTEM CONFIGURATION

COMPONENT	LIQUIDCOOL SOLUTIONS	DELL
SYSTEM	LSS220	R620
MOTHERBOARD	ASUS ZPH-D16	Dell Proprietary
CPUs - 2 PER SYSTEM	Intel® E5-2690	Intel® E5-2690
MEMORY - 8 DIMMs PER SYSTEM	Legacy 4GB 1600 MTS VLP (LE34RV16LVH-CL1060)	Legacy 4GB 1600 MTS VLP (LE34RV16LVH-CL1060)
HDD - 1 SSD PER SYSTEM	Intel® 520 SERIES, 180 GB, 6Gb SATA	Intel® 520 SERIES, 180 GB, 6Gb SATA
OPERATING SYSTEM	Microsoft® Windows Server® 2008 Enterprise R2	Microsoft® Windows Server® 2008 Enterprise R2
JAVA VERSION	1.7.0_07	1.7.0_07

In addition to the components specified above, the LCS and Dell systems were configured with non-redundant power-supply configurations. Each system contained a single 750W PSU.

No expansion cards (PCIe Cards) were installed in either system.

The Dell R620 configuration tested supports a maximum of 4 HDDs. The significance of this is that there is no hard-drive back-plane on the right side of the system. This chassis configuration provides the lowest overall air-flow impedance of any of the R620 chassis configurations. The lower the air-flow impedance, the lower the fan power required to cool the system. In fact, this R620 configuration should require the least power-to-cool of any available.

The power-to-cool values for the LSS220 are independent of PCIe or HDD configuration. The pump power required to circulate CoreCoolant through the LSS is determined by the overall coolant flow impedance of the device. This impedance is dominated by the fluid couplings and internal flow paths in the system. The addition of a PCIe card or additional HDDs do not affect the coolant flow impedance.

BIOS Settings

SETTING	LIQUIDCOOL SOLUTIONS	DELL
HYPER-THREADING	Enabled	Enabled
HARDWARE PREFETCHER	Disabled	Disabled
ADJACENT CACHE LINE PREFETCH	Disabled	Disabled
DCU STREAMER PREFETCHER	Disabled	Disabled
TURBO MODE	Enabled	Enabled
ENERGY PERFORMANCE	Balanced Performance	Balanced Performance

Operating System Settings

SETTING	LIQUIDCOOL SOLUTIONS	DELL
LOCK PAGES IN MEMORY	On	On
POWER PLAN	Balanced	Balanced

JVM Settings for SPECpower_ssj2008

SETTING	LIQUIDCOOL SOLUTIONS AND DELL
JVM VERSION	SPEC Java VM 5.0 (build 1.2.3.4-tricore 20071111)
JVM COMMAND-LINE OPTIONS	-server -Xmx1024m -Xms1024m -Xmn853m -XX:SurvivorRatio=60 -XX:TargetSurvivorRatio=90 -XX:ParallelGCThreads=2 -XX:AllocatePrefetchDistance=256 -XX:AllocatePrefetchLines=4 -XX:LoopUnrollLimit=45 -XXInitialTenuringThreshold=12 -XX:MaxTenuringThreshold=15 -XX:InlineSmallCode=3900 -XX:MaxInlineSize=270 -XX:FreqInlineSize=2500 -XX:+UseLargePages -XX:+UseParallelOldGC -XX:+UseCompressedStrings -XX:+AggressiveOpts
JVM AFFINITY	start /affinity [3,C,30,C0,300,C00,3000,C000,30000,C0000,300000,C00000,3000000,C000000,30000000,C0000000]
JVM INSTANCES	16
JVM INITIAL HEAP (MB)	1500
JVM ADDRESS BITS	1500
BOOT FIRMWARE VERSION	64
MANAGEMENT FIRMWARE VERSION	1.2.6
BOOT FIRMWARE VERSION	1.2.6
MANAGEMENT FIRMWARE VERSION	1.2.3.4
WORKLOAD VERSION	SSJ 1.2.10
DIRECTOR LOCATION	Controller
OTHER SOFTWARE	None

Power Meter Configuration

A single Yokogawa WT210 digital power meter was used to measure the power consumed by the systems under test. This is the most common meter used for SPECpower_ssj2008 measurements. The WT210 was attached to a Voltech Universal Breakout Box which was connected to the power outlet.

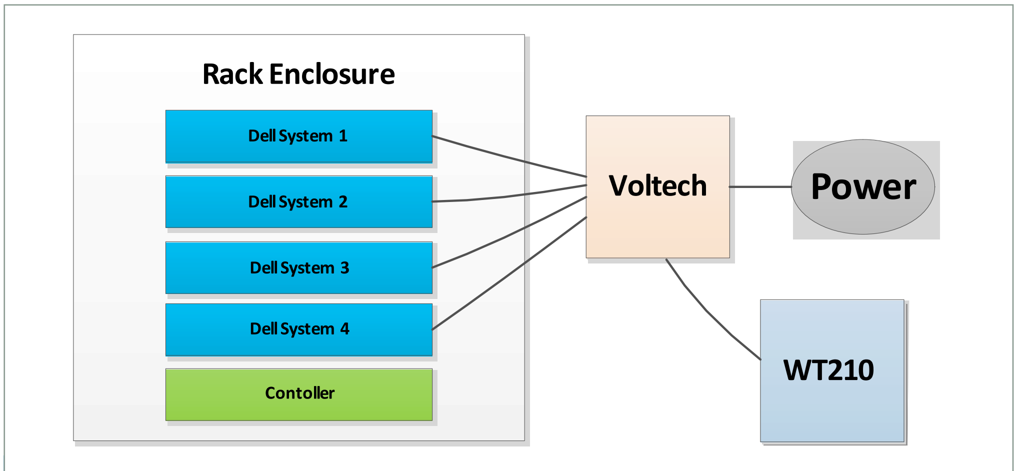

FIGURE 6 Power Meter Configuration - Dell R620 Servers

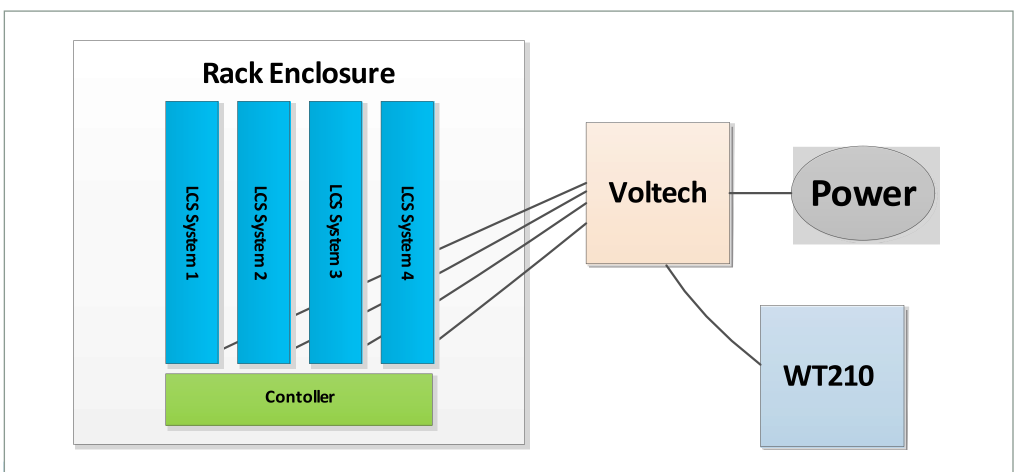

FIGURE 7 Power Meter Configuration - LCS LSS220 Servers

Leakage Current Testing

System Configuration

COMPONENT	LIQUIDCOOL SOLUTIONS
SYSTEM	LSS220
MOTHERBOARD	ASUS Z9PH-D16/QDR
CPUs – 2 PER SYSTEM	Intel® E5-2687W
MEMORY – 8 DIMMS PER SYSTEM	Legacy 8GB 1600 MTS VLP (LE38RV16LAH-ML1060)
HDD – 1 SSD PER SYSTEM	Intel® 520 SERIES, 180 GB, 6Gb SATA
OPERATING SYSTEM	Microsoft® Windows Server® 2008 Enterprise R2

APPENDIX B – SPECPOWER TEST SETUP

All systems tested were installed in a server rack as pictured below. The LSS220s were cooled by a liquid-to-liquid Cooling Distribution Unit (CDU-LL200). The water source for the CDU-LL200 was a water-to-air CDU located outside the chamber. This water-to-air CDU was capable of delivering cooling water to the CDU-LL200 at a minimum temperature of approximately 28°C. The CDU-LL200 was instrumented to provide coolant temperatures. Chamber ambient air temperature was measured using a thermal sensor located at the inlet to the servers or CDU. Testing was performed in the walk-in environmental chamber at Environ Laboratories in Minneapolis, MN.

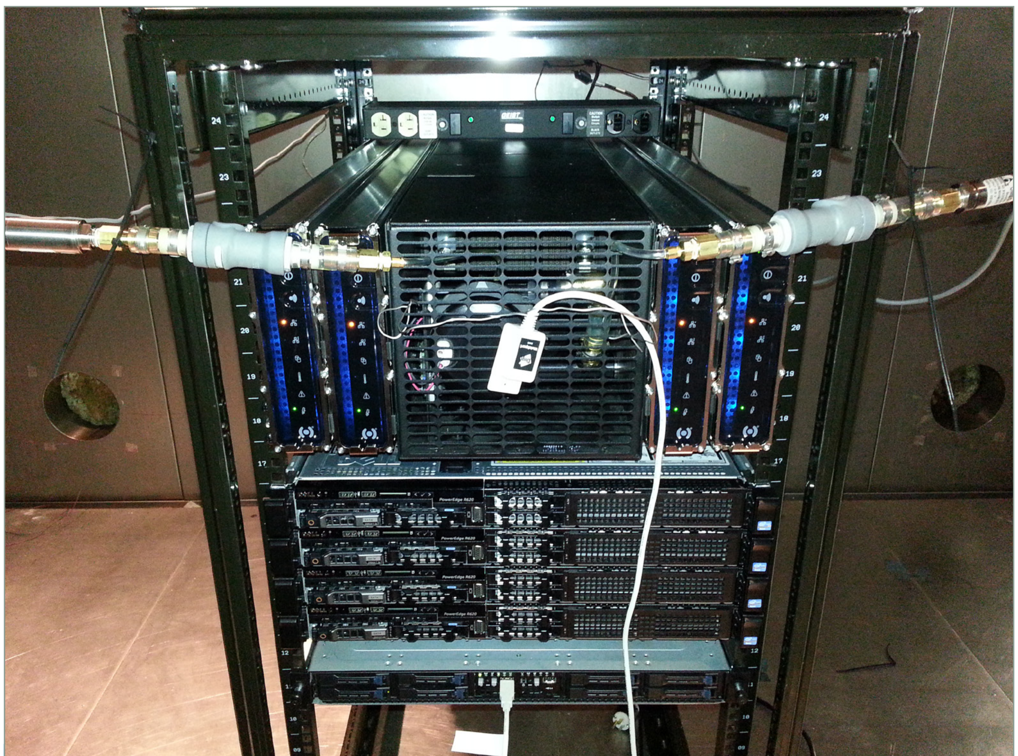

FIGURE 8 Server Rack with 4 LSS220s + CDU-LL200 and 4 Dell R620s

FIGURE 9 Walk-In Thermal Chamber

FIGURE 10 Server Rack in Environmental Chamber

The servers were tested and monitored using control computers outside the environmental chamber. The water-to-air CDU that provided the cooling water for the CDU-LL200 can be seen under the table in the figure below.

FIGURE 11 Control Computers and Water-to-Air CDU

APPENDIX C – DETAILED TEST RESULTS

Computational Performance

The first portion of the comparative testing focused on Computational Performance vs. Power Consumption. The primary function of this test is to operate a server at defined computational loads while measuring the server's power consumption. Ambient temperature is also collected during the test to ensure test results can be compared on a fair basis.

As shown in the following three figures the computational performance for the LCS and Dell servers were nearly identical for ambient temperatures between 15°C and 37.5°C.

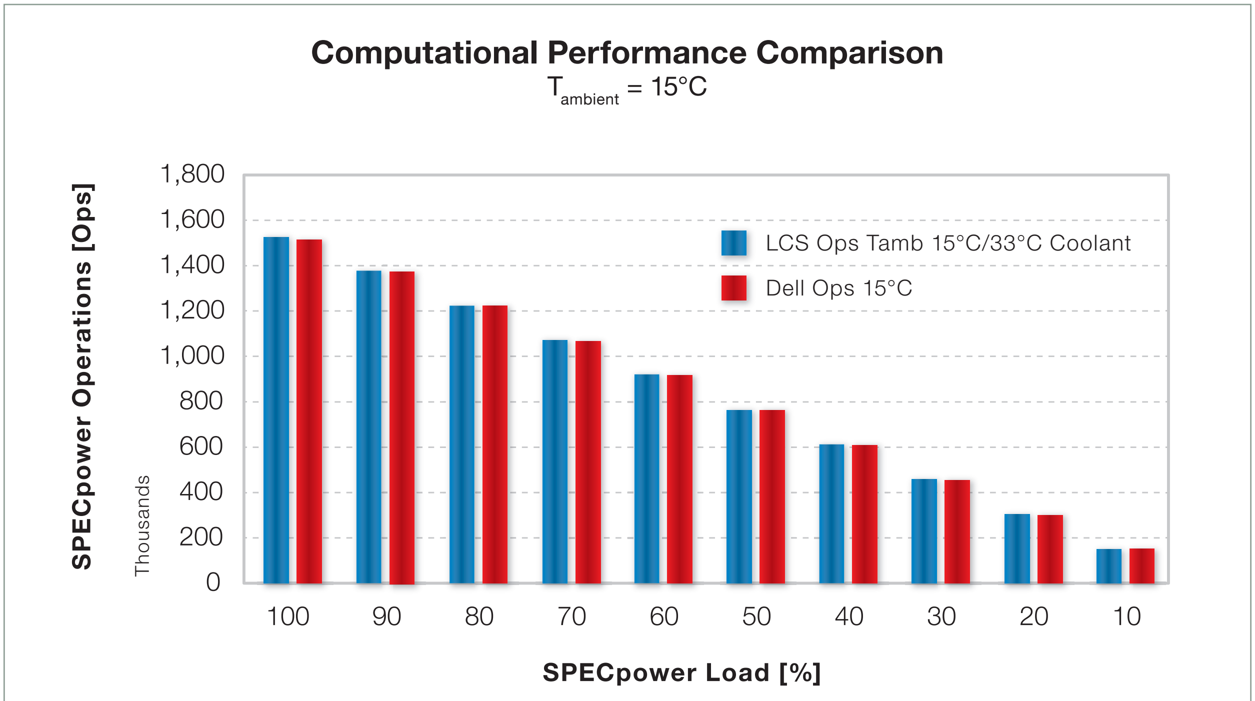

FIGURE 12 Computational Performance - Ambient 15°C

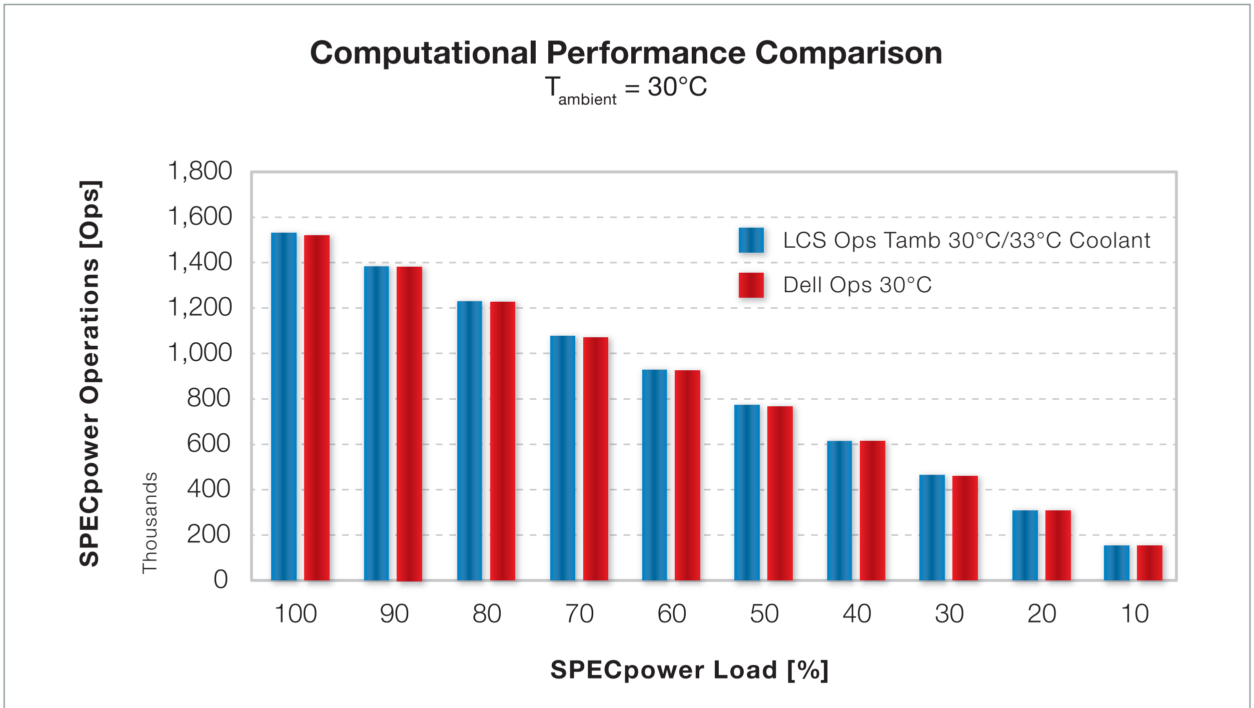

FIGURE 13 Computational Performance - Ambient 30°C

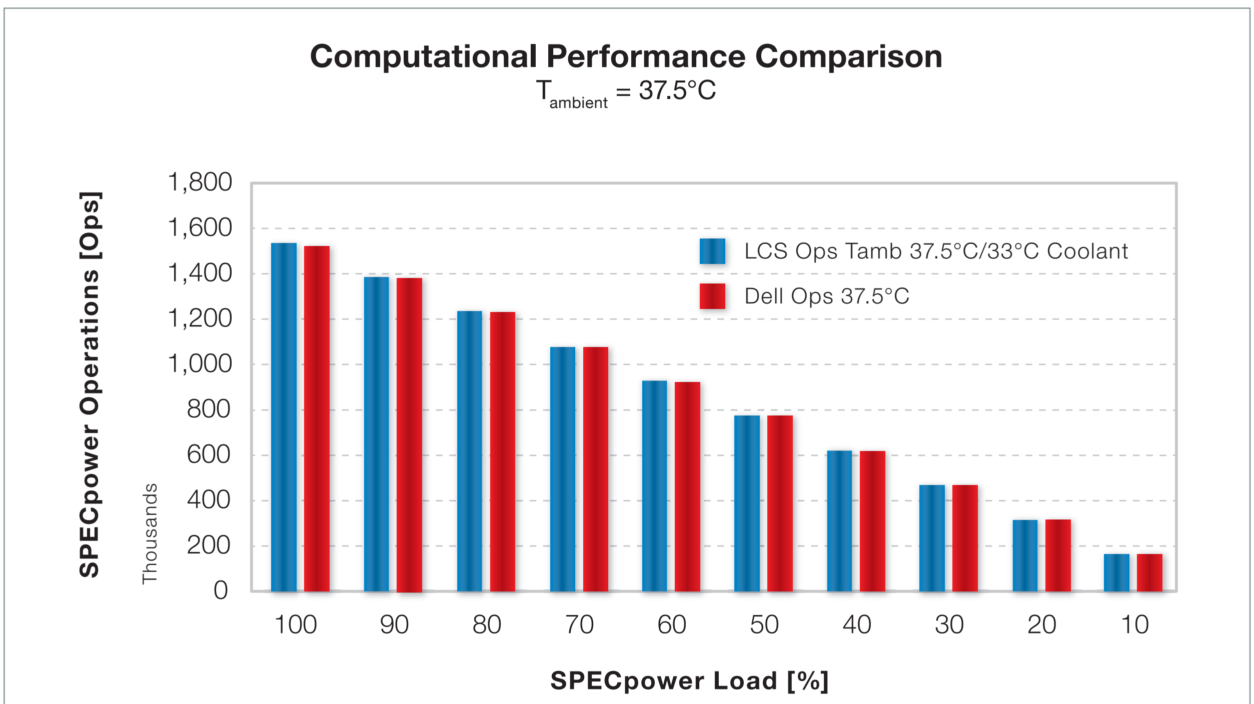

FIGURE 14 Computational Performance - Ambient 37.5°C

Figure 15 compares LCS and Dell computational performance results at 45°C. LCS computational performance experienced no noticeable degradation. The Dell server, however, had significantly reduced performance at 45°C.

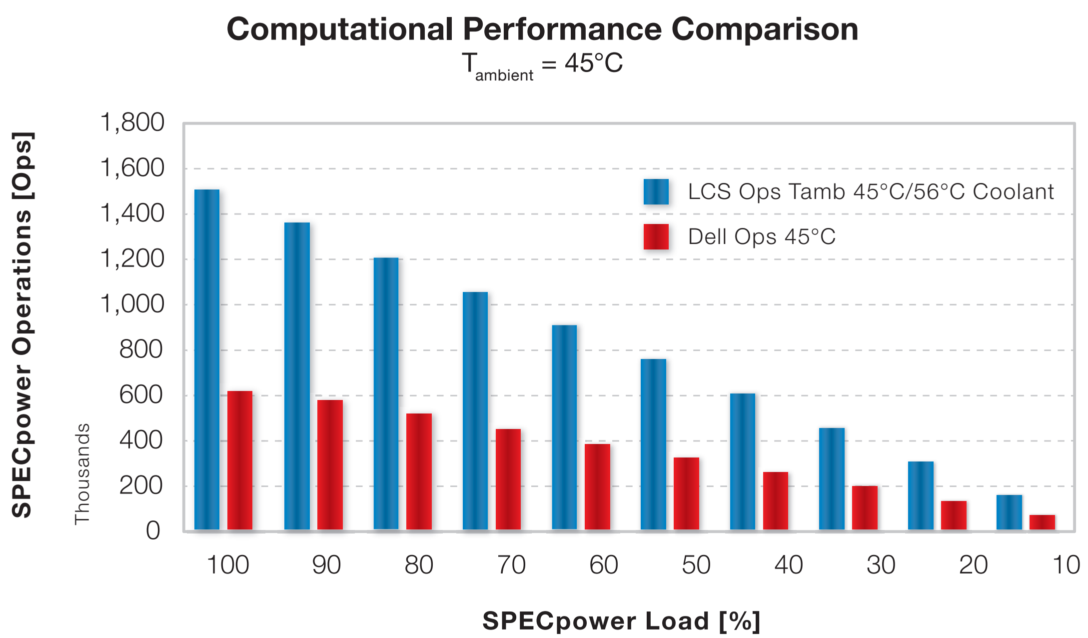

FIGURE 15 Computational Performance - Ambient 45°C

To determine the precise temperature at which Dell performance begins to degrade, computational performance was measured at one-degree intervals between 37°C. and 45°C. As shown in Figure 16 it's clear that when the ambient temperature exceeded 41°C the computational performance of the Dell servers was significantly reduced.

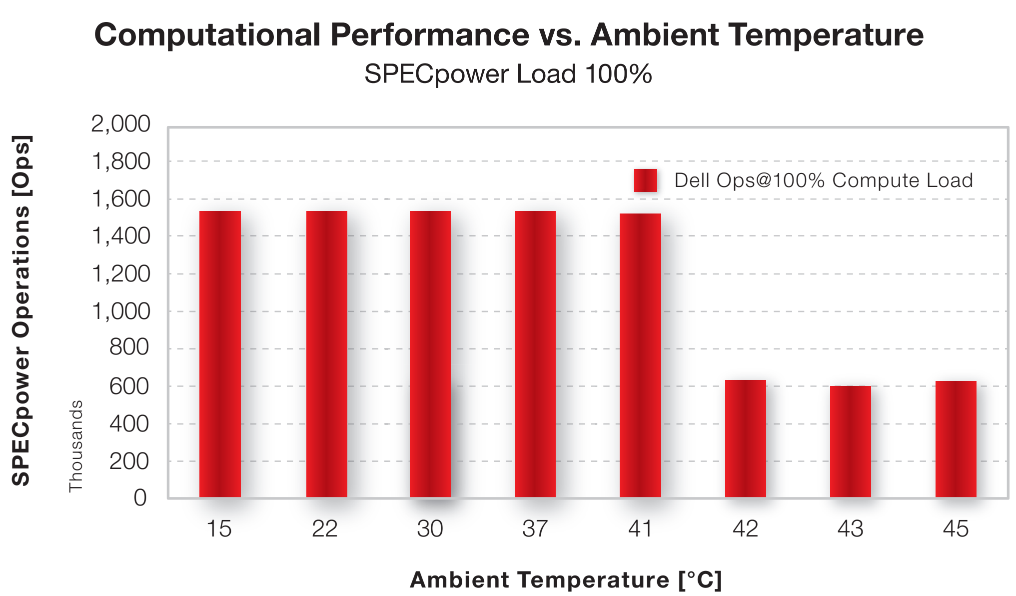

FIGURE 16 Computational Performance vs. Ambient Temperature SPECpower Load 100%

Power Consumption

When comparing power usage results from the SPECpower runs, it was found that the total power consumption of the LCS servers was higher than that of the Dell servers at lower ambient temperatures. There are multiple factors that contributed to these results. One key factor is that the fan power required is minimized due to the cold ambient air. It was also determined that there were two key differences between the LCS and Dell servers: the efficiency of the Power Supplies and Voltage Regulation (VR) circuits on the motherboards. Dell's power supply has an efficiency of greater than 90% while the efficiency of the LCS power supply used in this study is closer to 85%. The specifics regarding the VRs are more difficult to measure, but appear to be of similar magnitude.

At higher ambient temperatures and computational loads, the total power consumption of the LCS servers was lower than that of the Dell servers. This was driven mainly by the increase in power to run the fans that cool the Dell servers.

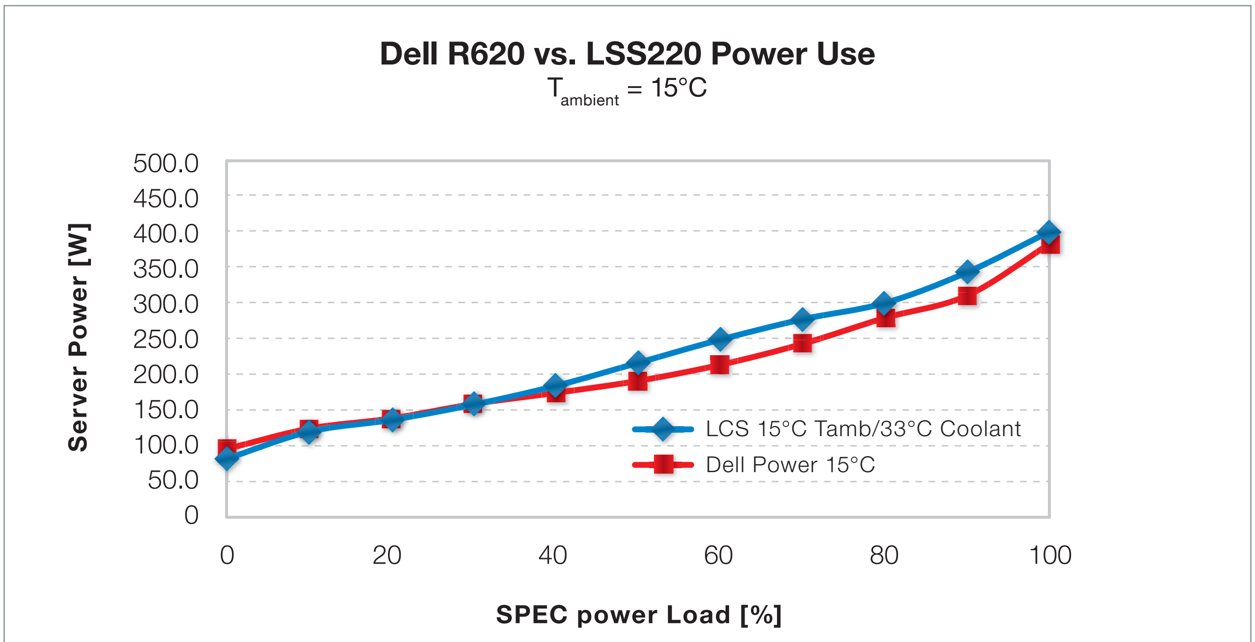

FIGURE 17 Server Power Consumption at 15°C

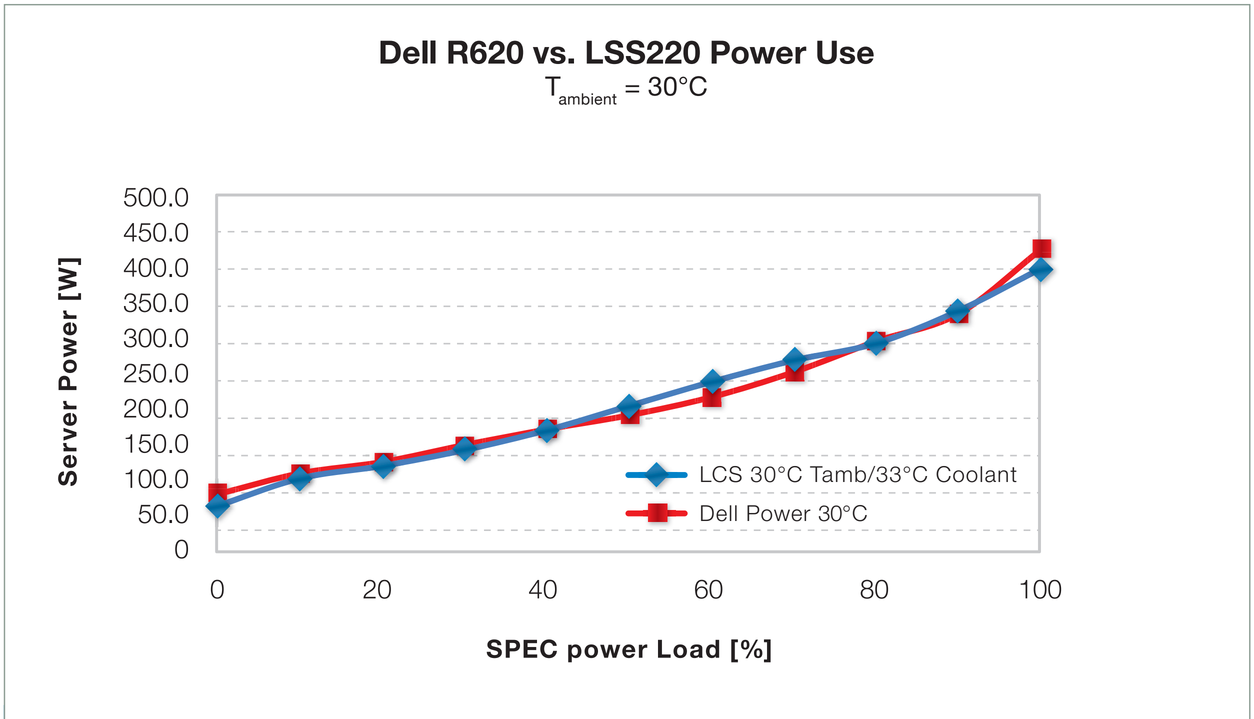

FIGURE 18 Server Power Consumption at 30°C

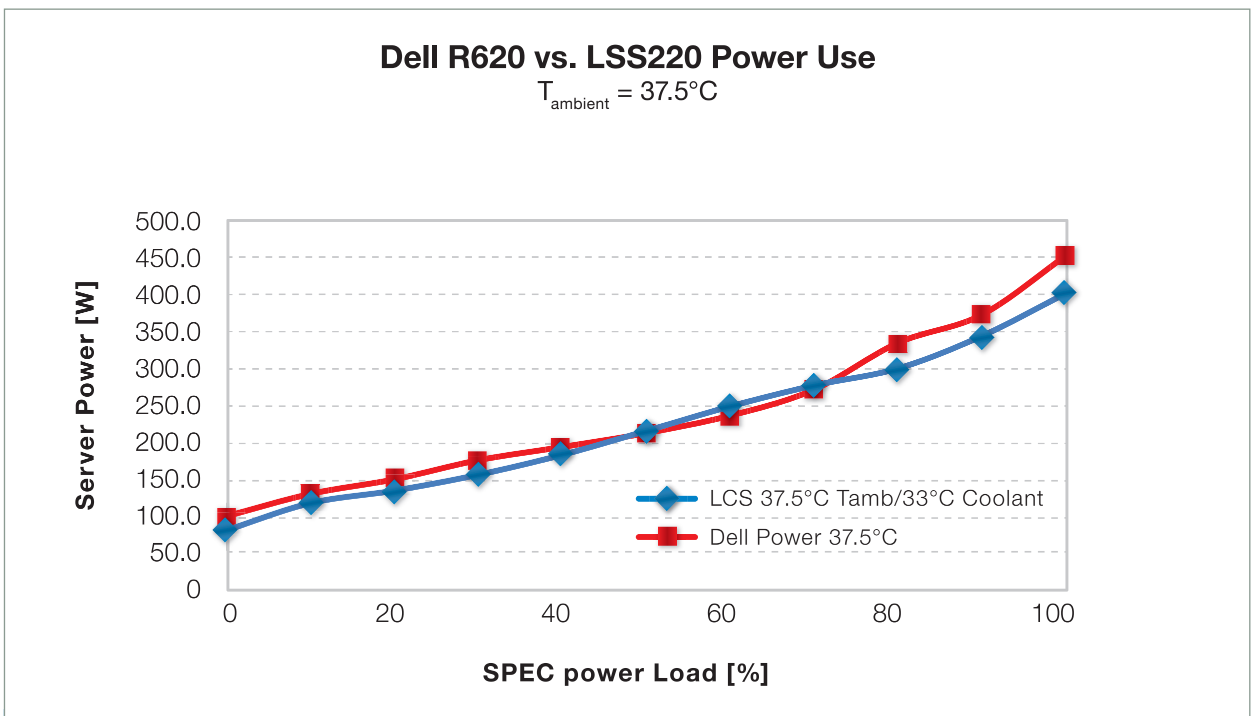

FIGURE 19 Server Power Consumption at 37.5°C

Power-to-Cool

The intention of this white paper is not to compare the total power consumed by the LCS servers to that consumed by the Dell servers. The focus is to compare the fan power used to cool a well designed air-cooled server to the pump power used to cool a liquid submerged server. The reasoning behind this is that in today's age of commodity servers, there is nothing to prevent LCS from designing a server with identical power-to-compute to that of Tier One OEMs like Dell. High efficiency power supplies and motherboards with high efficiency VRs are readily available. The real benefit to the data center comes from LCS' liquid submersion cooling technology. Air-cooled server vendors have no way to reduce their power-to-cool levels to those of liquid submersion cooling.

Since the SPECpower testing showed that the computational performance of the LCS servers is superior or equivalent to those of the Dell servers at all ambient temperatures tested, the focus of the testing was shifted to accurate measurement of the power-to-cool for each of the servers.

Power-to-Cool = Total Fan Power for Dell Server

Power-to-Cool = Coolant Pump Power for LCS Server

The SPECpower results also indicated that the servers were not being stressed to maximum levels. Peak power consumption was below 450W per server in all but one test case. To stress the servers to the highest level possible, different software utilities were used.

Two different utilities were used: BurnIn Test and Prime95. These utilities were benchmarked to determine how much computational power was required at specific program settings. A matrix was created using selected settings for each utility to provide a nearly linear scale ranging from idle to maximum computing power. Measurements were taken at the different setpoints to generate the data provided in this paper. Further details regarding the software utilities can be found in Appendix A under "System Stress Software Utilities."

Measuring the pump power required to cool the LCS servers was straight-forward. The pump that circulated the Core Coolant is DC-powered. The pump was powered by a digital DC power supply and the power consumed was measured directly. The pump power was kept constant across the varying computational levels.

Measuring the fan power consumed by the Dell Servers was slightly more complicated. The Dell R620 server uses seven 40mm dual-rotor Delta GFC0412DS-D REV00 fans. The fan power was monitored by adding a 0.1 ohm sense resistor to 12V power cable for specific fans as shown in Figure 20.

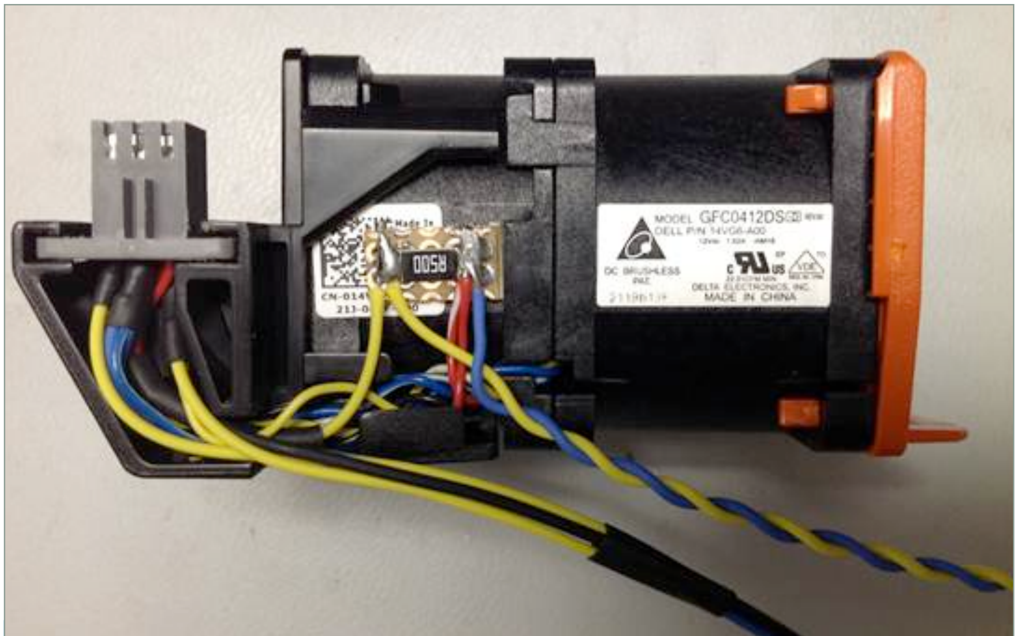

FIGURE 20 Delta Fan Modified with Current Sense Resistor

During testing, it was found that the seven fans are controlled in independent cooling zones within the Dell server. Fans 1-4 are operated at one level while fans 5-7 can be operated at a different level based on closed loop sensor input. Typically, it was observed that the fans 1-4 run at a higher speed than fans 5-7.

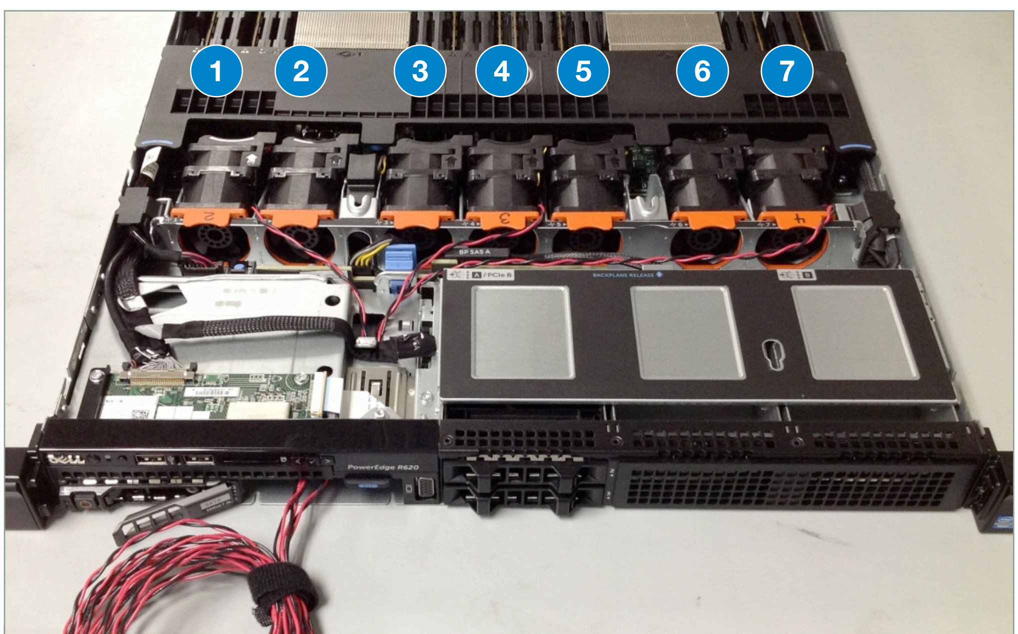

FIGURE 21 View of the Seven Dell Cooling Fans

Fans 1, 4, and 7 were instrumented with the sense resistors as described above. This provided an accurate measurement of the fan power consumed by the Dell server. To determine the total fan power consumed per server, the following formula was used:

$$P_{Total\ Fan} = 2P_{Fan1} + 2P_{Fan4} + 3P_{Fan7}$$

Once the pump and fans were properly instrumented, stress tests were performed with both the LCS and Dell Servers. Tests were performed across a range of ambient temperatures from 22°C to 45°C for the Dell servers. For the LCS servers, the temperature of the inlet coolant was set to the same level as the ambient temperature at which the Dell server was tested.

Total Server Power and Power-to-Cool were measured and recorded. From these values, the power-to-compute was calculated as:

$$\text{Power-to-Compute} = [\text{Total Server Power}] - [\text{Power-to-Cool}]$$

The figures below provide a comparison of the power-to-cool between the LCS and Dell servers based on specific levels of power-to-compute. The power-to-compute was controlled by the stress software utilities described above and in Appendix A.

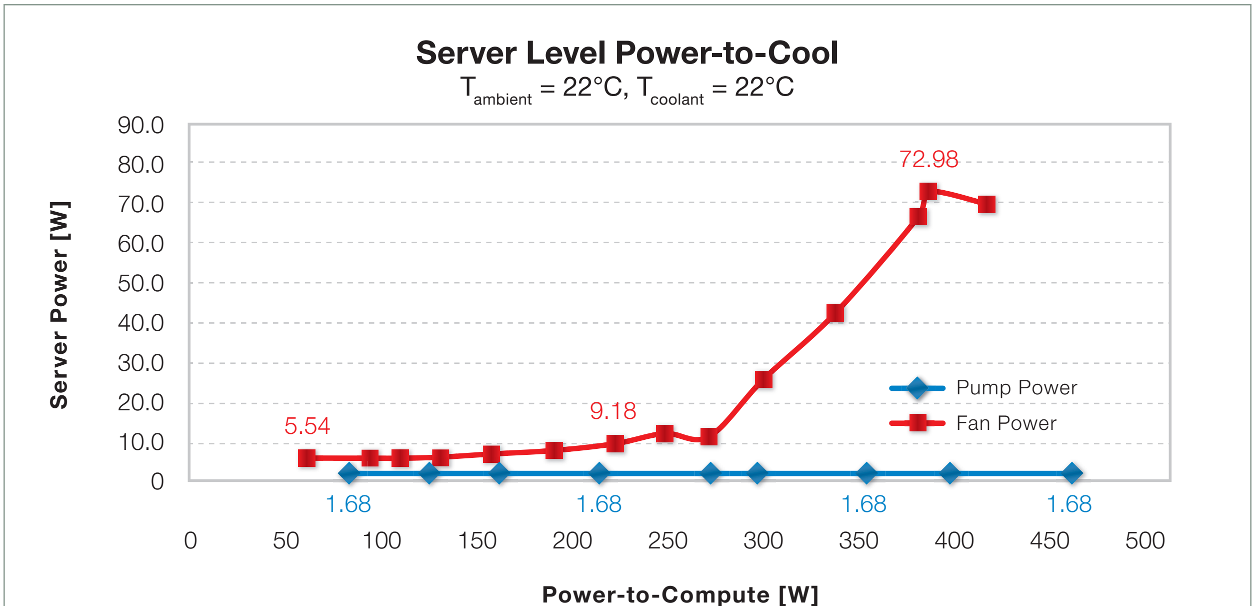

FIGURE 22 Server Level Power-to-Cool - Tamb = 22°C

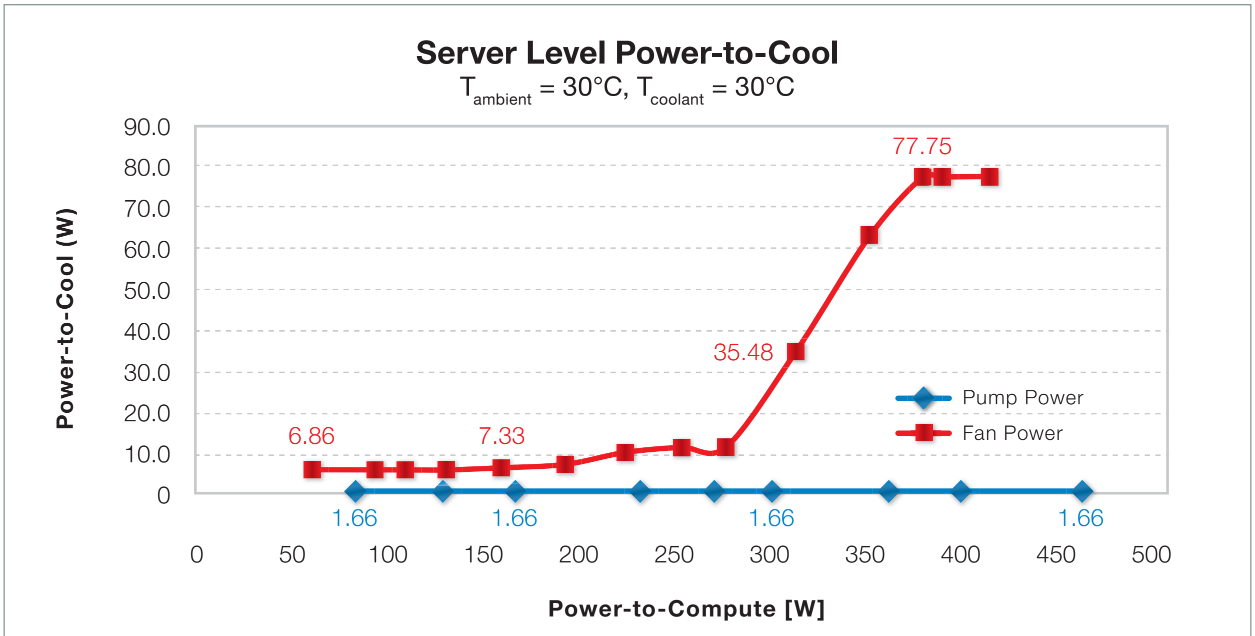

FIGURE 23 Server Level Power-to-Cool - Tamb = 30°C

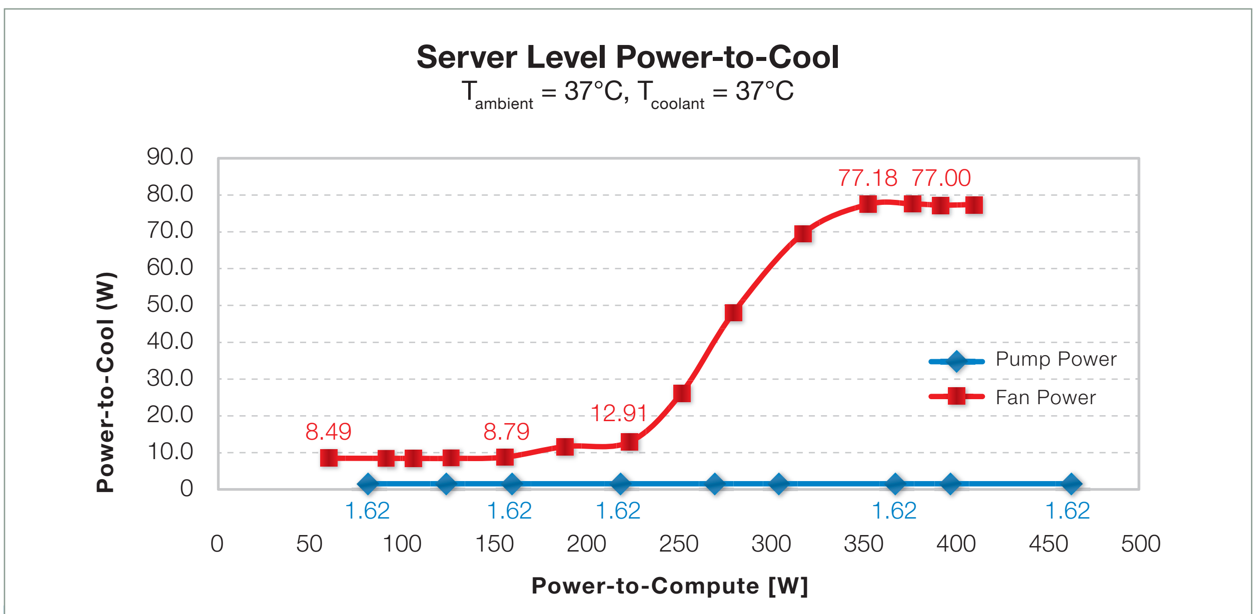

FIGURE 24 Server Level Power-to-Cool - Tamb = 37°C

Depending on the level of power to compute, liquid submersion cooling reduced the server level power-to-cool by 4W at a minimal computational load to over 75W at high computational loads. This is a reduction in power-to-cool of 70% to 98%. In terms of total server power, this translates into a savings of up to 18%.

As noted above, the Dell server throttled at 45°C so no useful data could be collected. This is reflected in Figure 25.

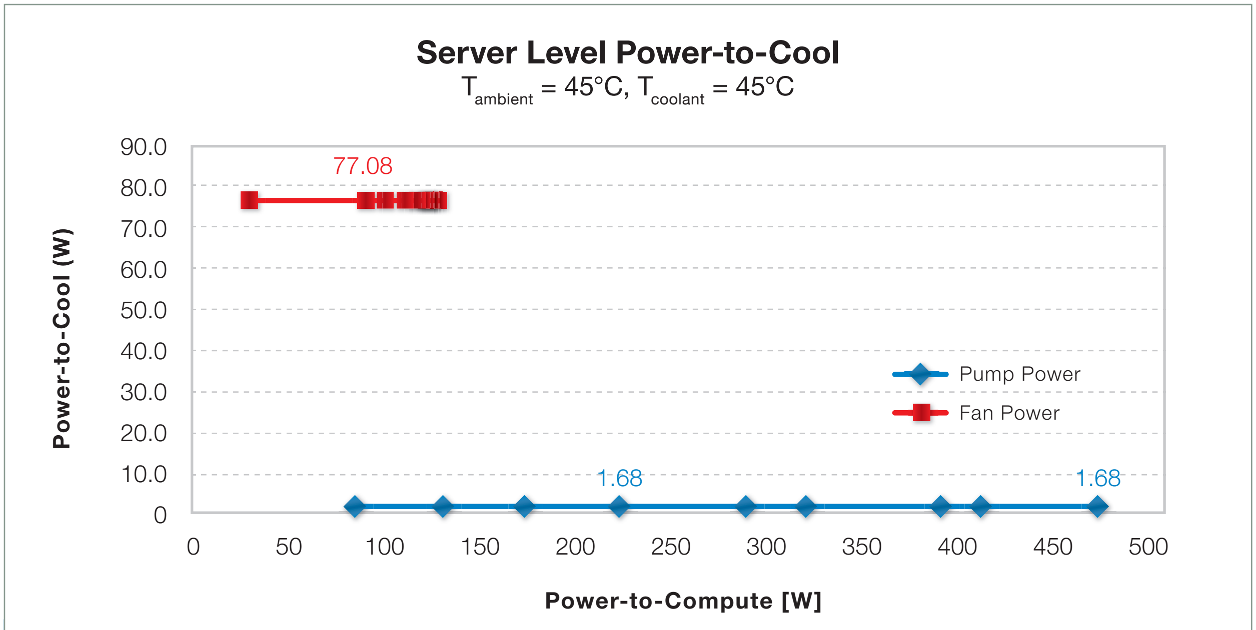

FIGURE 25 Server Level Power-to-Cool - Tamb = 45°C

Comparing CPU Core Temperatures

The CPU Core temperatures were also recorded at each power level to show the thermal margin each system had at the specific operating points.

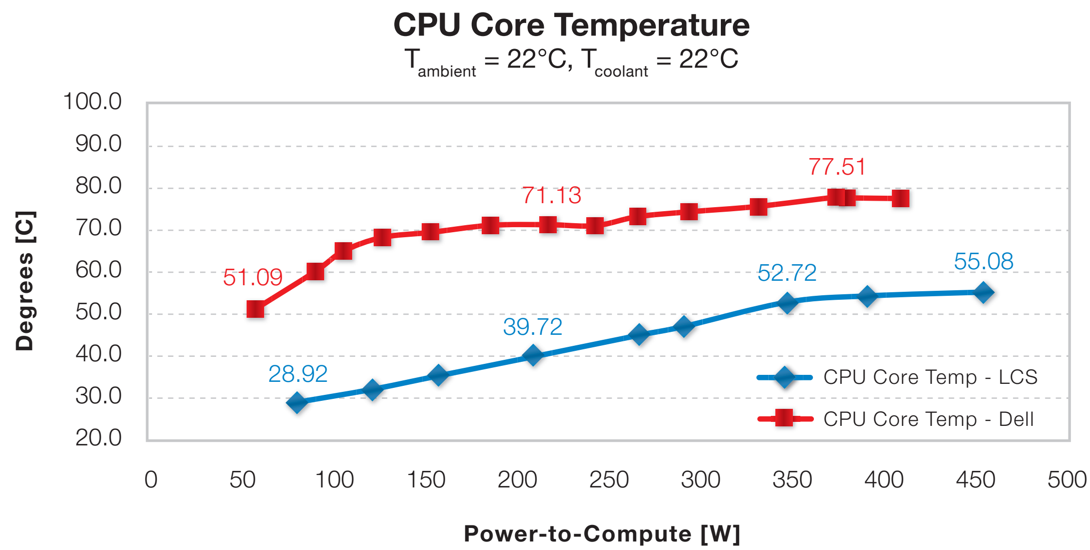

FIGURE 26 CPU Core Temperature Comparison - Tamb = 22°C, Tcoolant = 22°C

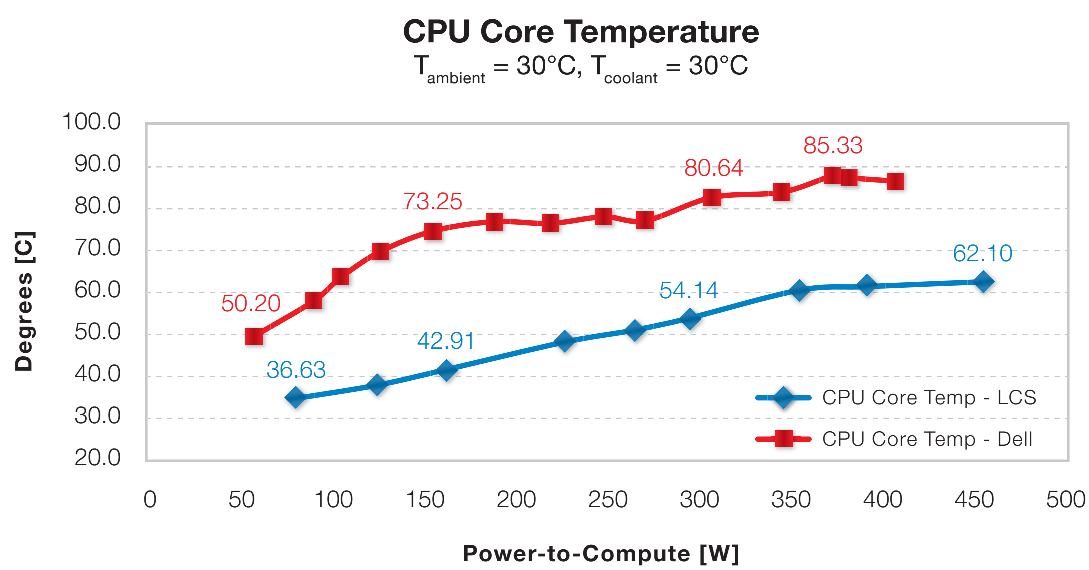

FIGURE 27 CPU Core Temperature Comparison - Tamb = 30°C, Tcoolant = 30°C

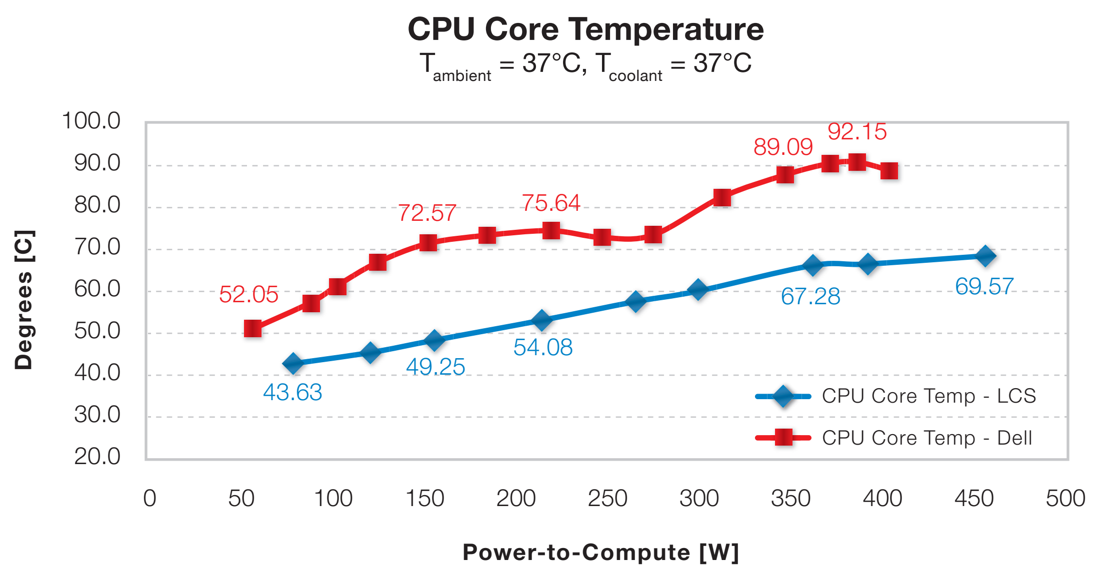

FIGURE 28 CPU Core Temperature Comparison - Tamb = 37°C, Tcoolant = 37°C

As expected, the CPU core temperatures increased with increasing levels of power-to-compute. The changes in core temperatures also increased directly with the increase in ambient/coolant temperature for temperatures from 22°C to 37°C. The core temperatures for the CPUs in the LCS servers were significantly lower than those in the Dell servers for all temperatures tested. The reduction in core temperature was from approximately 10°C to over 30°C depending on operating point.

As was seen in the SPECpower tests, the processors in the Dell servers throttled prior to reaching 45°C so no useful data could be collected on those systems under these conditions. This is reflected in Figure 29.

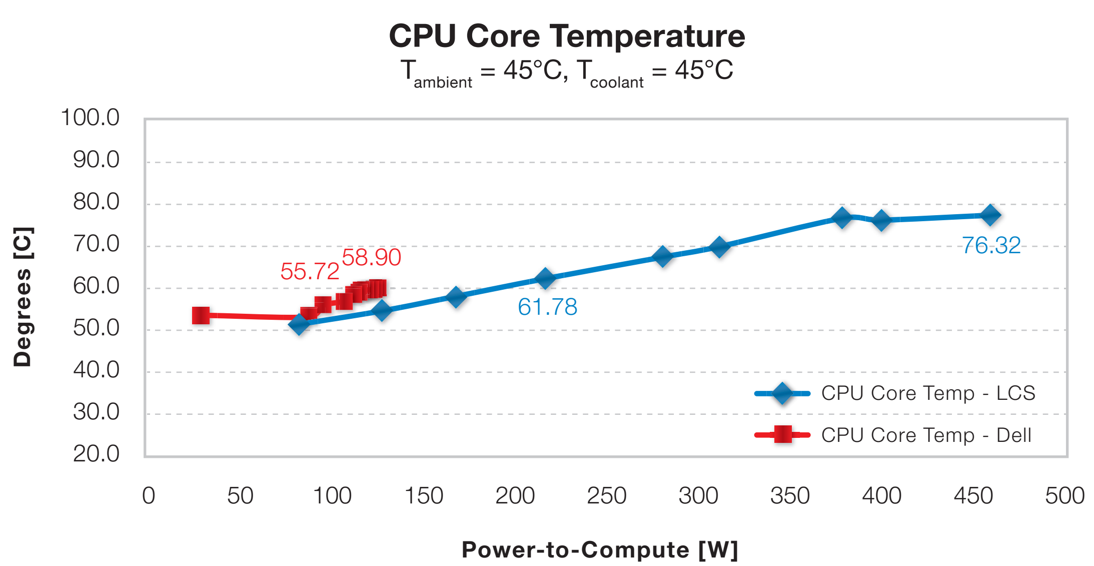

FIGURE 29 CPU Core Temperature Comparison - Tamb = 45°C, Tcoolant = 45°C

Computational Performance vs. Ambient Temperature

Liquid submersion cooling provides more consistent computational performance and server-level power consumption than air-cooled servers when operating across a broad range of ambient temperatures. SPECpower testing shows that the computational performance of the LCS server is not affected by ambient temperature.

The figure below illustrates that, with the inlet coolant temperature held constant at 33°C, the computational performance of the LSS220 was essentially constant when operated in ambient environments from 15°C to 45°C.

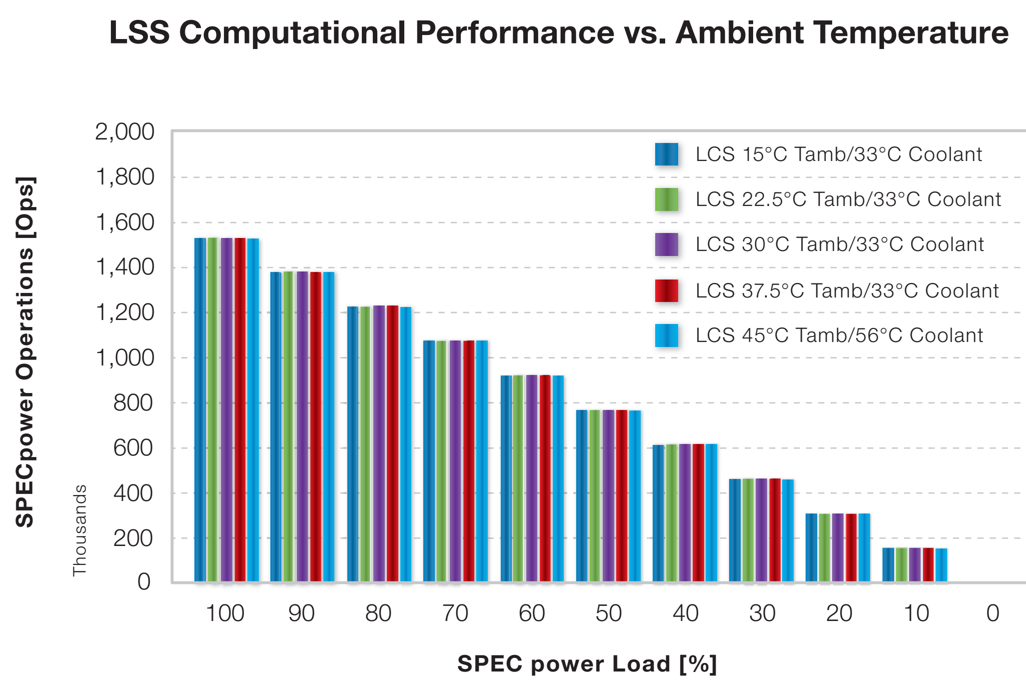

FIGURE 30 LSS220 Computational Performance vs. Ambient Temperature

The computational performance for the Dell server remained consistent until the ambient temperature exceeded approximately 41°C. At that point, the processors began to throttle and the computational performance was reduced significantly.

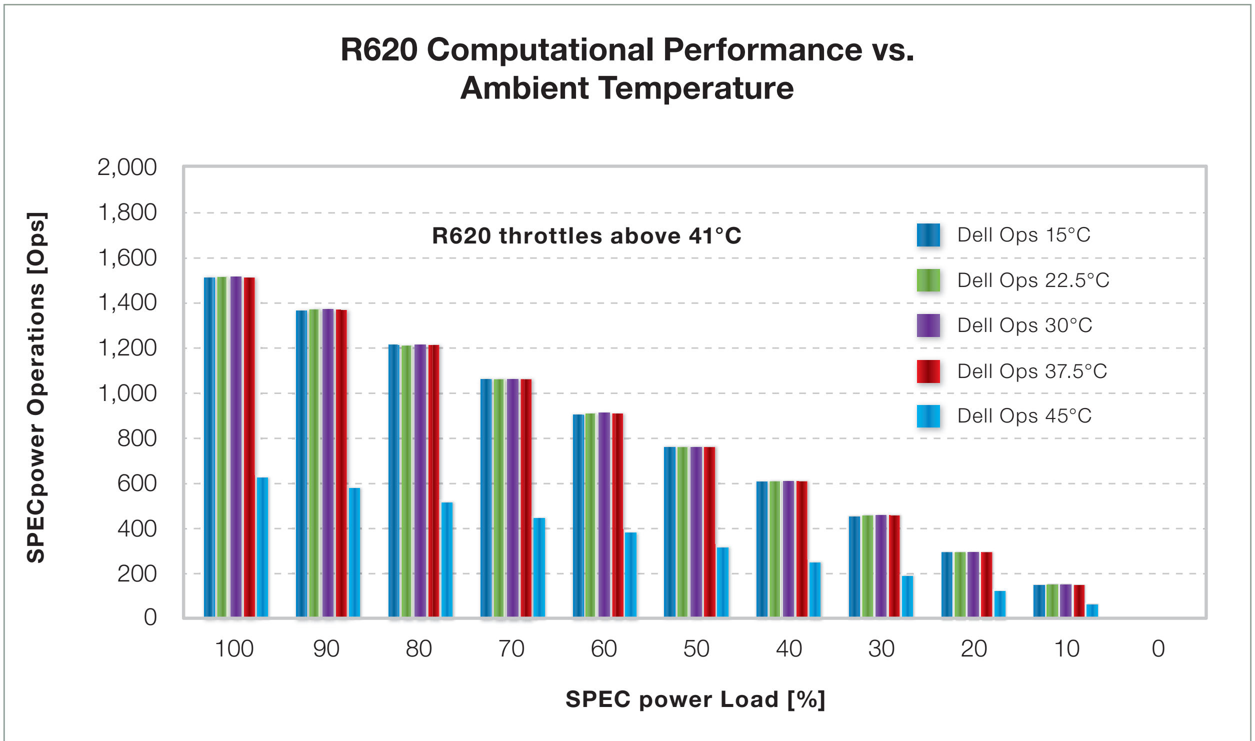

FIGURE 31 R620 Computational Performance vs. Ambient Temperature

The power consumption of the servers followed similar trends. The LCS server consumed nearly the same amount of power regardless of ambient temperature as shown in Figure 32.

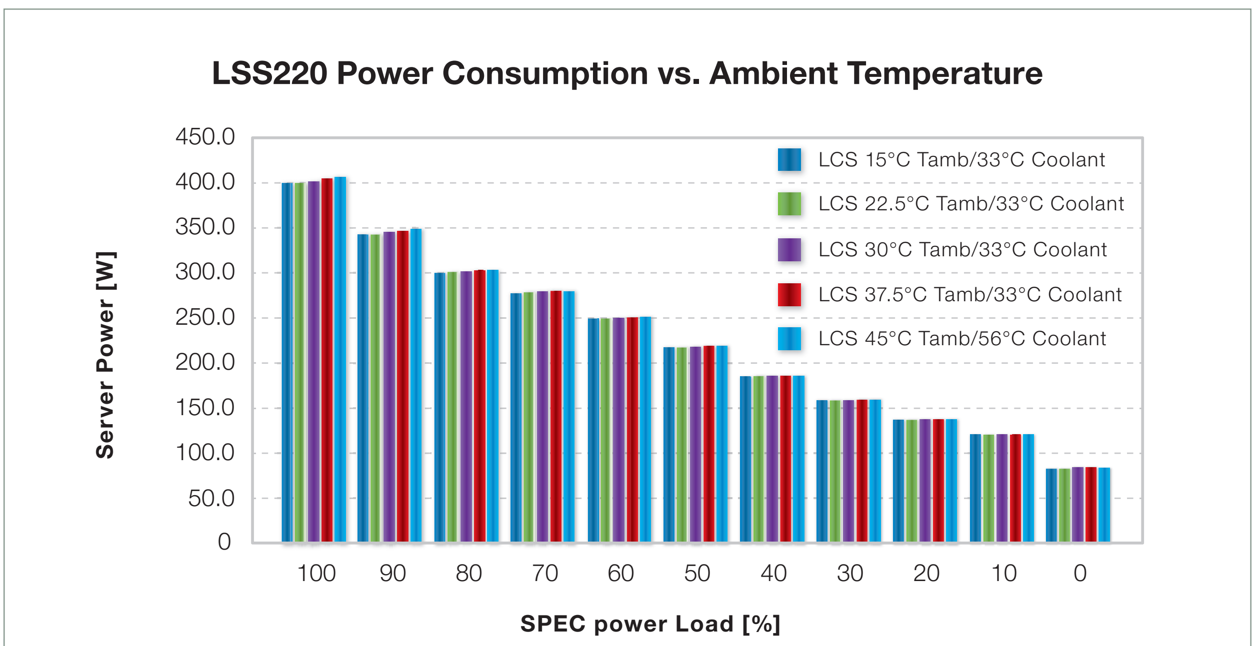

FIGURE 32 LSS220 Power Consumption vs. Ambient Temperature

As shown in Figure 33, the Dell server consumed progressively higher levels of power as the ambient temperature was increased. This trend continued until the ambient temperature exceeded approximately 41°C and the processors began to throttle. At this point, the fans continued to run at full speed, but the power consumption of the throttling processors was greatly reduced. This, in turn, reduced the overall power consumption of the Dell server.

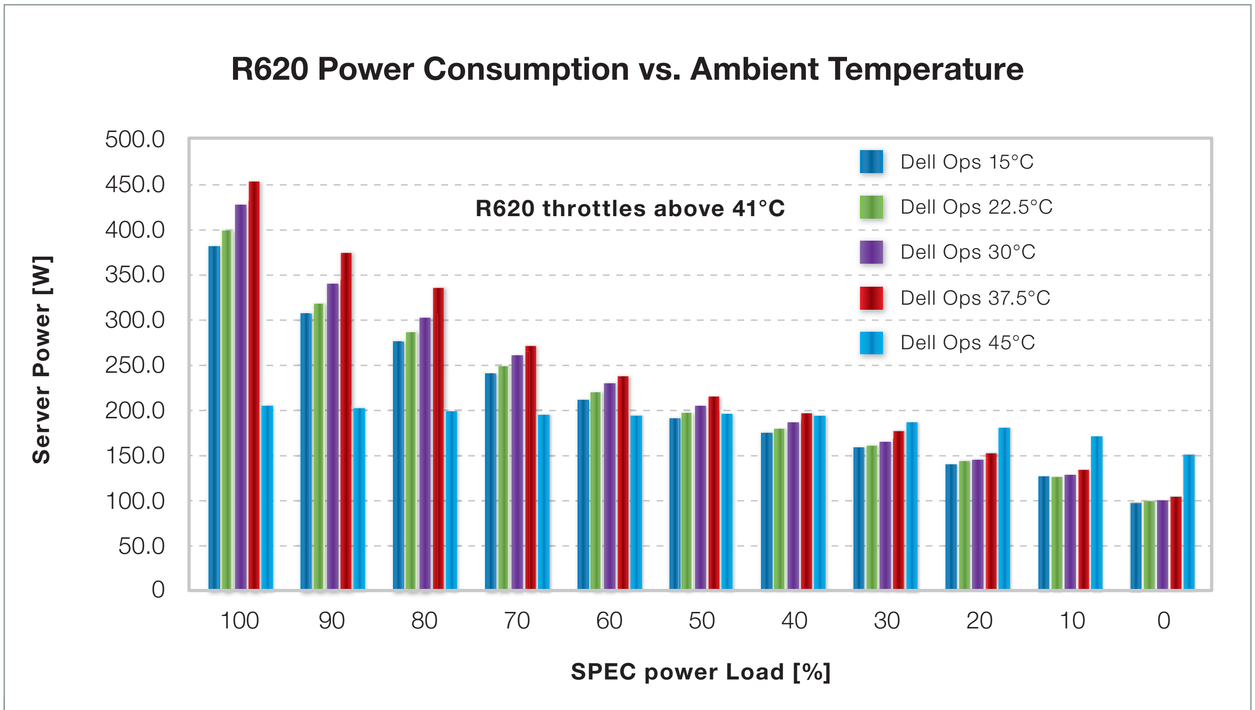

FIGURE 33 R620 Computational Performance vs. Ambient Temperature

The LCS servers were also tested across a broad range of coolant inlet temperatures. The results of this testing can be found in Figure 34. As shown in the figure, the computational performance of the servers was essentially unchanged for coolant inlet temperatures between 22°C and 56°C. This is significant as it allows the LCS Servers to be cooled with warm or even hot water. This will result in significant capital, maintenance, and energy savings for organizations that deploy servers that utilize LCS submersion cooling technology.

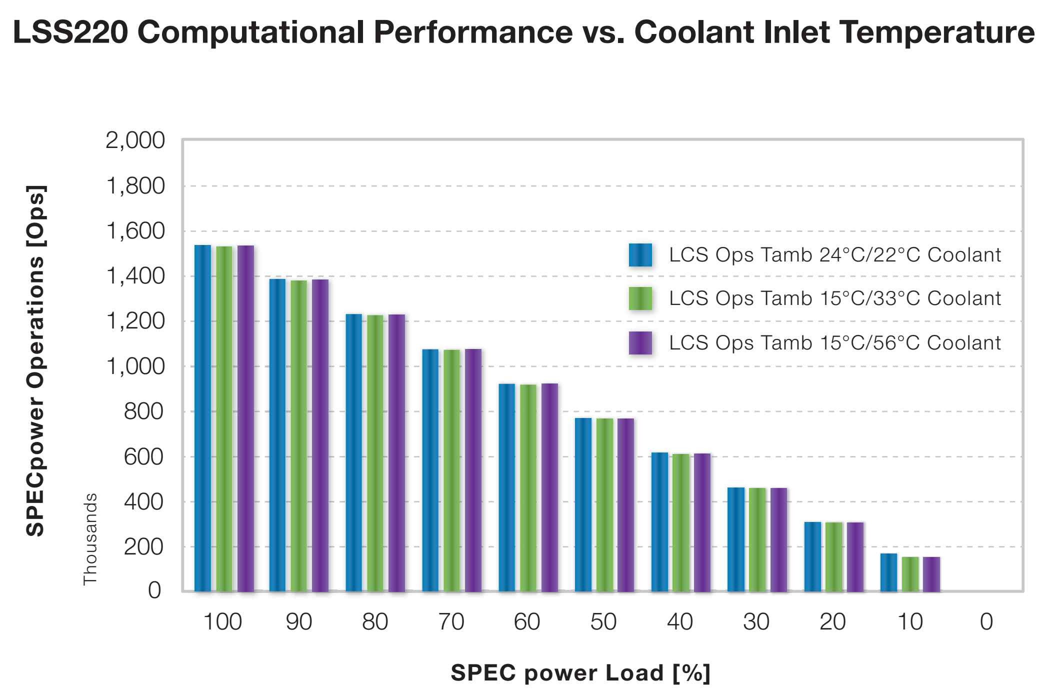

FIGURE 34 LSS220 Computational Performance vs. Coolant Inlet Temperature

Reduction in Leakage Current

An additional benefit to submersion cooling is the fact that the entire system is cooled to levels significantly lower than in air-cooled systems. Where this benefit can be seen is in the minimization of the "leakage current"[5] for the integrated circuits on the motherboard.

To demonstrate this effect, an LCS server was operated with a constant computational load. Over a two-hour period, the coolant inlet temperature was allowed to increase gradually from 19°C to 54.5°C. The coolant inlet temperature and power consumed by the server were collected at five-minute intervals. The results are displayed in the following figures.

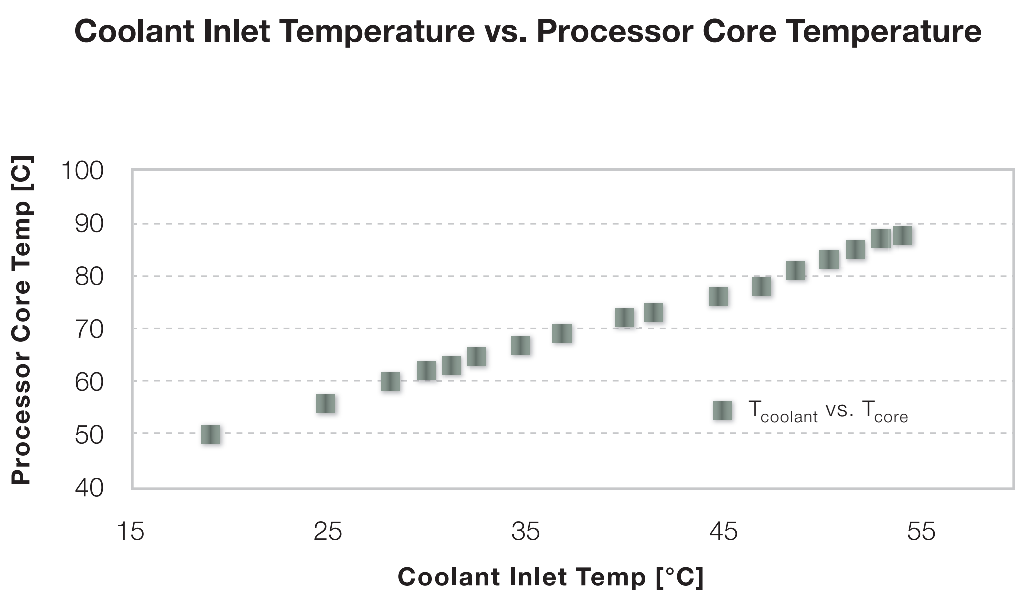

FIGURE 35 Coolant Inlet Temperature vs. Processor Core Temperature

The figure above shows that as the inlet temperature of the coolant increased, the temperature of the processor core increased as well. The core temperatures above were all well below the temperature where throttling can occur. For the processor tested, throttling typically begins above 95°C.

Figure 36 below shows the total server power usage as a function of the processor core temperature while under a constant computational load. It can be seen that the power consumed by the server varied as a second order polynomial with respect to the processor core temperature. The power consumed by the server increased by approximately 13% as the core temperatures increased from 50°C to 88°C. This corresponds to an increase in the inlet temperature of the coolant from 19°C to 54.5°C as shown in Figure 35.

[5] Background on causes and effects of leakage current in microprocessor design can be found in "Leakage Current: Moore's Law Meets Static Power" by N.S. Kim, et al., published by the IEEE Computer Society, December 2003

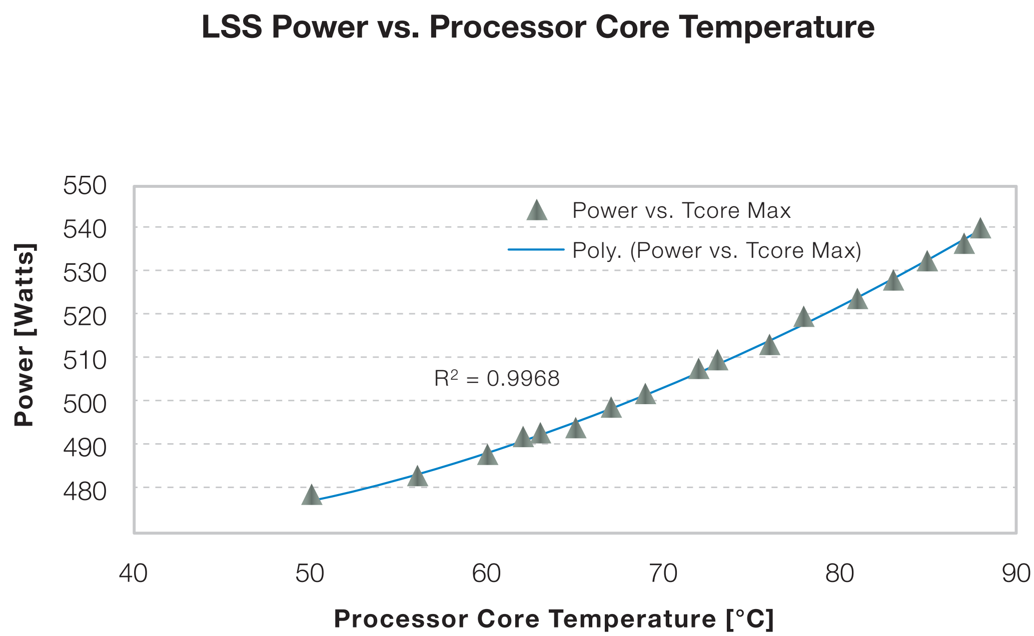

FIGURE 36 LSS Power Consumption vs. Processor Core Temperature

These findings highlight two distinct advantages of LCS cooling technology.

First, LCS technology gives data center operators the flexibility to "trade-off" cooling power for server power to minimize their total power consumption based on their specific environmental conditions. Data center operators in cooler climates can use free-cooling to deliver lower coolant inlet temperatures to help reduce the server's power-to-compute. Data center operators in hotter climates may not be able to reduce server power significantly, but will be able to use free-cooling to avoid power consumed by chillers.

The power trade-off described above is also accomplished with no impact to server computational performance. As shown previously, the LCS server computational performance is essentially constant for coolant inlet temperatures ranging from 24°C to 56°C.

Second, as this study shows, LCS servers keep processors 20°C to 30°C cooler than air-cooled servers given the same ambient conditions and computational loads. Because lower processor temperatures mean less leakage current, LCS servers use less compute power than air-cooled servers under these conditions. By combining the results of the leakage current measurements with the core temperature results presented earlier, the amount of power savings can be estimated. For example, as shown in Figure 27, at a cooling temperature of 30°C and a power-to-compute load of 350W, the difference in core temperatures between the LCS server and the Dell server is about 23°C (62°C to 85°C). Applying this result to the data shown in Figure 36, the difference in total power used is 41W (533W – 492W). On a percentage basis, the reduction is 7.6% of total server power.

APPENDIX D: BACKGROUND INFORMATION

Liquid Submersion Technology – A Better Approach to Cooling

LiquidCool Solutions' patented liquid submersion cooling technology works by immersing all of a device's electronics in a dielectric liquid. Using this approach, each device is packaged within a compact, sealed, liquid-tight enclosure that installs into a standard server rack. The dielectric coolant is distributed to the devices in the rack using a system of dripless quick-connect fittings and manifolds connected to a liquid-to-liquid cooling distribution unit (CDU) which is mounted separately, either in the server rack or a remote location convenient to the customer. A variable-speed pump in the CDU provides the right amount of coolant flow to achieve the necessary cooling while minimizing power use. Within each device, LiquidCool's directed flow technology circulates coolant to all components while providing maximum cooling to the primary heat-generating components, such as CPUs or GPUs. This is illustrated in Figure 37.

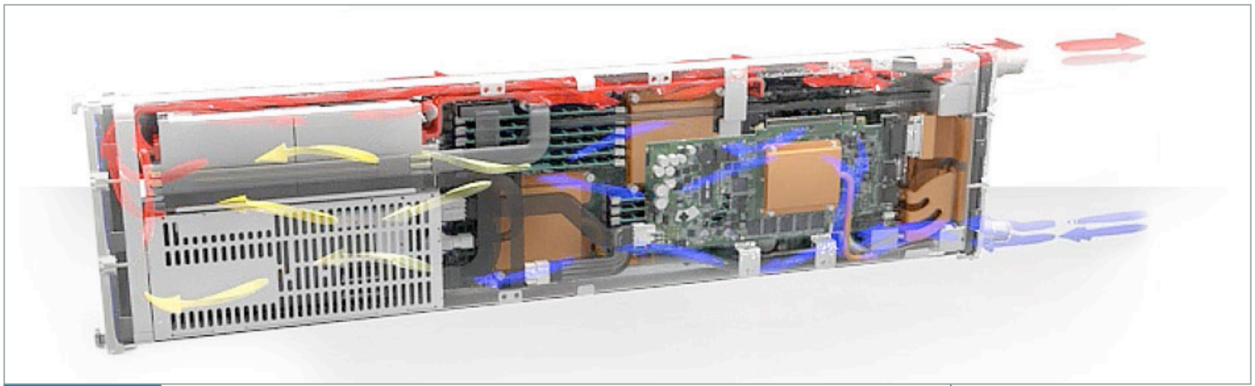

FIGURE 37 View of Directed Coolant Flow Within a LiquidCool Liquid Submerged Server

Liquid submersion cooling has a number of advantages over air-cooled approaches.

- The primary advantage is that by completely eliminating the need for fans within the device, far less energy is required for cooling.
- Because the dielectric liquid is a far more effective cooling medium than air, liquid submersion cooling keeps components operating at lower temperatures, which reduces thermal stress and improves component reliability.
- Lower operating temperatures also reduce leakage current effects, which further improves energy efficiency.
- System failures directly correlate to fluctuating temperatures. Submersion cooling reduces thermal fatigue as the temperature changes for the individual components are reduced significantly due to the higher density and heat capacity of the dielectric coolant.[6]
- Use of liquid submersion cooling opens up new opportunities to push the operating levels of the discrete electronics beyond normal thresholds for future performance gains.
- Liquid submersion cooling also eliminates failure modes associated with oxidation and corrosion, which can occur within traditional air-cooled systems and eliminates the susceptibility to ESD events.

For more information on LiquidCool technology and its benefits, visit the LiquidCool Solutions web site at: www.liquidcoolsolutions.com.

[6] **El-Sayed, N., Stefanovici, I., Amvrosiadis, G., Hwang, A.A.** Temperature Management in Data Centers: Why Some (Might) Like It Hot. *http://www.cs.toronto.edu/~nosayba/temperature_cam.pdf.* [Online] 2012.