
Intergenerational Energy Efficiency of Dell EMC PowerEdge Servers

This white paper compares the energy efficiency of the refreshed 14th generation PowerEdge 1U rack server, based on the Intel Xeon Scalable processor product family, to that of its direct predecessors.



Bruce Wagner

Dell EMC | PG Enterprise Solutions Performance Analysis Team

This document is for informational purposes only and may contain typographical errors and technical inaccuracies. The content is provided as is, without express or implied warranties of any kind.

© 2018 Dell Inc. All rights reserved. Dell EMC and its affiliates cannot be responsible for errors or omissions in typography or photography. Dell, the Dell logo, and PowerEdge are trademarks of Dell Inc. Intel and Xeon are registered trademarks of Intel Corporation in the U.S. and other countries. Microsoft, Windows, and Windows Server are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. 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. Dell disclaims proprietary interest in the marks and names of others.

SPEC® and the benchmark name SPECpower_ssj® are registered trademarks of the Standard Performance Evaluation Corporation. For more information about SPECpower, see spec.org/power_ssj2008/

April 2018 | Version 1.1

Contents

Executive summary	4
Introduction	4
Key findings	4
Performance/watt	4
Performance	4
Power	4
Methodology	4
Typical configuration.....	5
Results	6
Summary.....	9
Outlook	9
Appendix A—Test methodology.....	11
SPECpower_ssj2008 standard.....	11
BIOS settings.....	12
OS tuning.....	12
SPECpower_ssj2008 configuration.....	13
Power and temperature configuration.....	13
Appendix B—SPECpower_ssj2008 results.....	14

Tables

Table 1. Detailed configuration for energy efficiency comparison.....	5
Table 2. BIOS settings	12

Figures

Figure 1. Normalized SPECpower_ssj2008 results for PowerEdge R640, R630+ and R630	6
Figure 2. Performance per watt ratios across workload intervals.....	7
Figure 3. R640 subsystem power consumption across SPECpower workload intervals.....	8
Figure 4. PowerEdge energy efficiency progress.....	9
Figure 5. PowerEdge energy cost and performance improvement trend.....	10
Figure 6. SPECpower_ssj2008 results for Dell PowerEdge R640	14
Figure 7. SPECpower_ssj2008 results for the Dell PowerEdge R630+	15

Executive summary

Introduction

With power and cooling costs accounting for an increasingly large portion of IT budgets, IT departments looking to minimize total cost of ownership (TCO) are finding it advisable to make energy efficiency a priority when choosing server hardware. In this white paper, we examine the intergenerational energy efficiency improvements in the latest Dell PowerEdge server family focusing on the popular two-processor, 1U rack form factor platform, configured just as it is typically specified by large data center customers.

The Dell Solutions Performance Analysis (SPA) team compared the Dell PowerEdge R640 configured with the latest Xeon SP microarchitecture product family CPU versus one configured with the preceding Broadwell-EP microarchitecture Xeon E5-2600 v4 ones. Using the industry-standard SPECpower_ssj2008[®] benchmark, the two servers were tested for outright performance, performance/watt and input power consumption.

The results showed **the Dell PowerEdge R640 delivered substantially better performance than its one-year old predecessor with the same overall energy efficiency.**

Key findings

Performance/watt

The PowerEdge R640 achieved **the same, excellent overall performance-to-power ratio** as its R630+ predecessor in the configuration commonly ordered by data center customers.

Performance

The **PowerEdge R640 provided 13% better raw performance** than the R630+ at the 100% workloading level.

Power

The PowerEdge R640 consumed 2.5% less power when idle potentially **saving 14KWh of electricity annually.**

Test methodology and detailed result reports are documented in this paper.

Methodology

SPECpower_ssj2008 is an industry-standard benchmark created by Standard Performance Evaluation Corporation (SPEC) to measure a server's power and performance across multiple utilization levels. Appendix A—Test details the test methodology used by Dell and Appendix B SPECpower_ssj2008 provides the detailed report data that supports the results in this paper.

Typical configuration

The two systems were configured according to enterprise data center customer requirement feedback and Dell's new product marketing projections. The differences between the two configurations being due to the natural advancement in technology, commodity component price/availability and industry performance agency accepted benchmarking environment standards that occur over time.

The configurations used are summarized in Table 1.

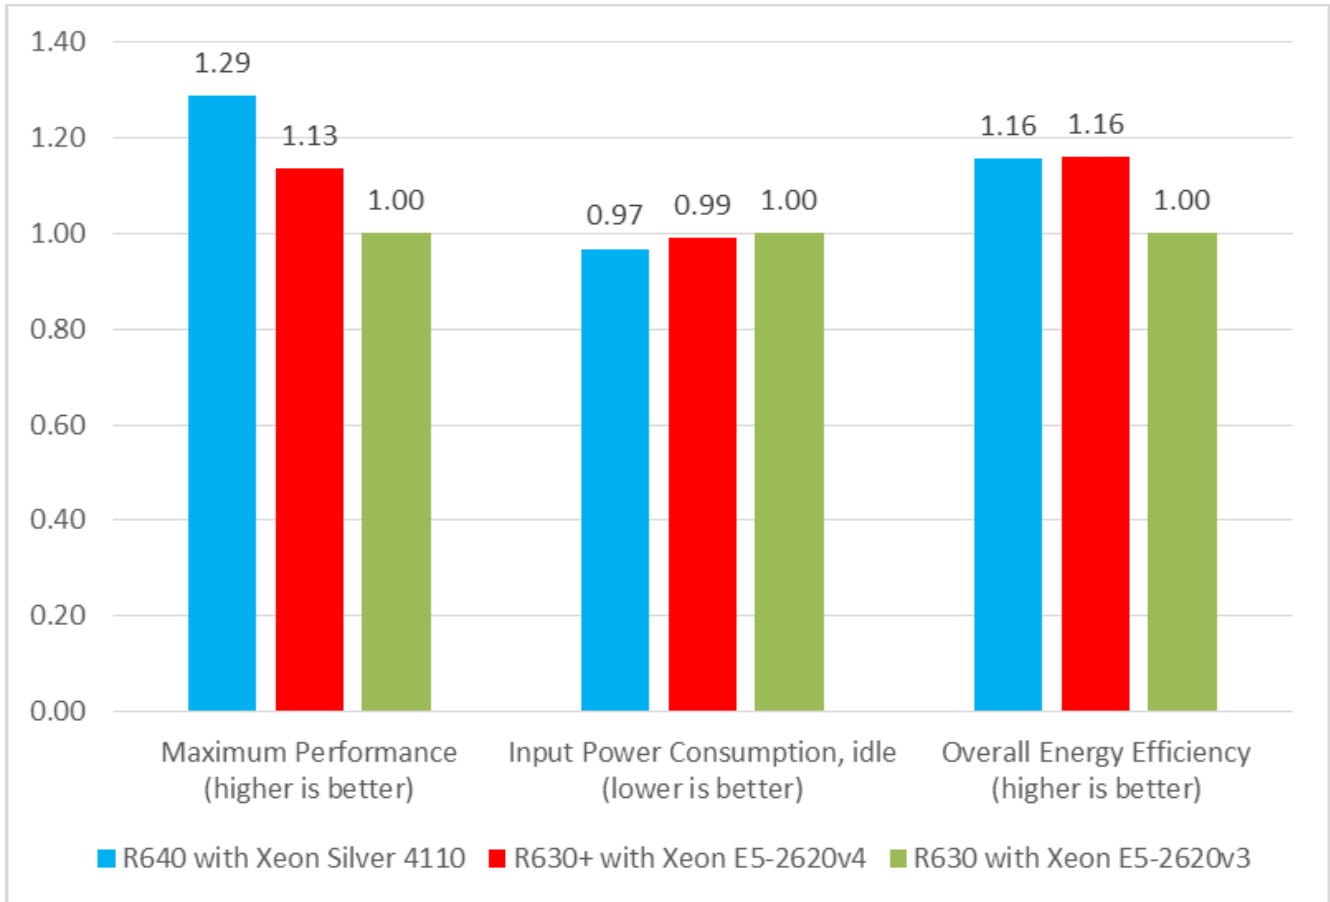
Table 1. Detailed configuration for energy efficiency comparison

Configuration	PowerEdge R630+	PowerEdge R640
Sockets/form factor	2S/1U	2S/1U
Processors	2 x Intel® Xeon® E5-2620 v4, 8 physical/16 logical cores, 2.10GHz	2 x Intel Xeon Silver 4110, 8 physical/16 logical cores, 2.10 GHz (Dell P/N 7KW7T)
Memory	64GB, 8 x 8GB dual-ranked PC4-2133P RDIMMs (Dell P/N H8PGN)	192GB, 12 x 16GB dual-ranked PC4-2666V RDIMMs (Dell P/N VM51C)
Hard drives	2 x 300GB 10K RPM 6Gb SAS RAID 1 (Dell P/N PGHJG)	2 x 400GB 12Gbps SAS SSD RAID 1 (Dell P/N YT53C)
Storage controller	Dell PERC H730 1GB cache (Dell P/N KMCCD)	Dell PERC H730P 2GB cache (Dell P/N 7H4CN)
Power supply quantity/rating	2 x 495W (Dell P/N 2FR04)	2 x 750W (Dell P/N 5RHVV)
Network adapter	1 x quad-port Broadcom® 5720 1GBase-T (Dell P/N FM487)	1 x quad-port Broadcom® 5720 1GBase-T (Dell P/N FM487)
Operating system	Microsoft® Windows Server® 2012 R2 Datacenter, Version 6.3.9600.17196	Microsoft Windows Server 2016 RS1 Datacenter, Version 10.0.14393.1378
System BIOS FW	2.1.7	1.3.7
Board management FW	2.30.30.30	3.00.00.00

Results

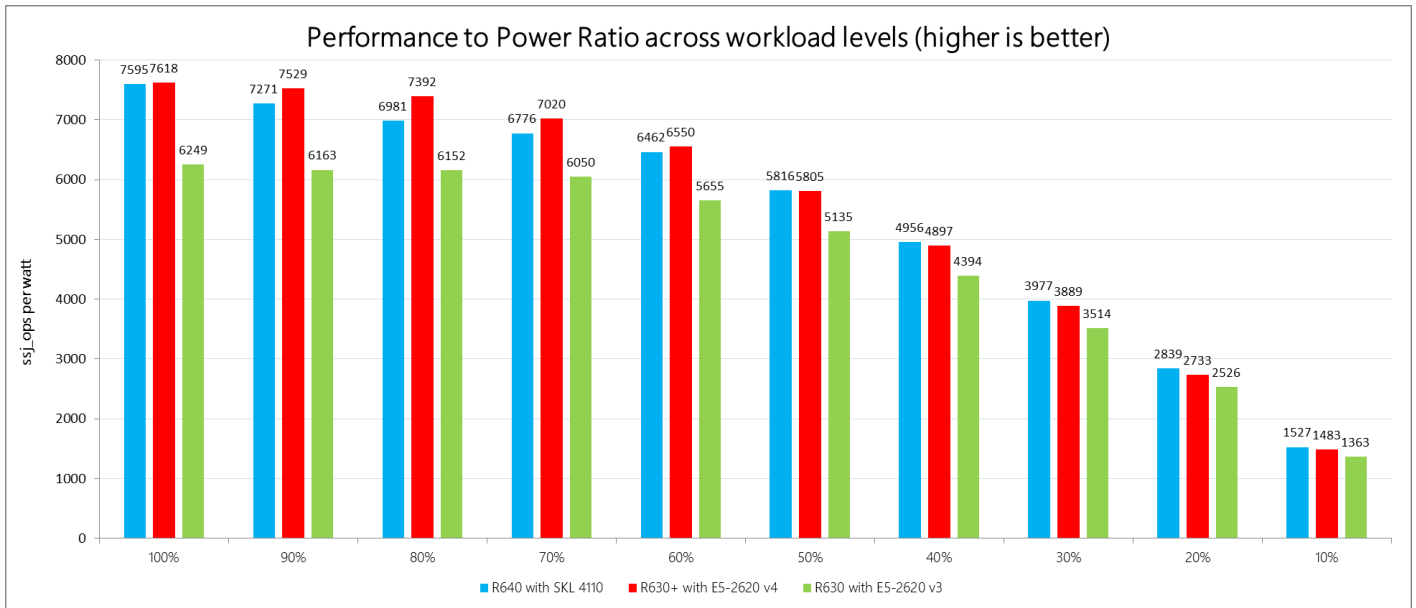
In the like-for-like configurations detailed in table 1, the PowerEdge R640 in comparison to its immediate predecessor: R630+ delivered 13% more work and lowered power consumption 2.5% while sustaining the same excellent overall energy efficiency.

Figure 1. Normalized SPECpower_ssj2008 results for PowerEdge R640, R630+ and R630



SPECpower_ssj2008 reports the server’s performance to watt ratio at workload levels from 10% to 100% CPU utilization. These are calculated by dividing total computational output for the given workload level by the measured average input power consumption (ssj_ops/watts). Figure 2 shows the PowerEdge R640 has a higher performance to power ratio than the R630+ up thru 50% CPU utilization levels and substantially higher performance to power ratio across all possible CPU utilization levels than the R630.

Figure 2. Performance per watt ratios across workload intervals¹




¹ Required SPEC score disclosure information: R640: (1,625,307 ssj_ops and 214W) at 100% target loading and 5583 vs. R630+: (1,432,168 ssj_ops and 188W) at 100% target load and 5597 overall ssj_ops/watt vs. R630: (1,262,314 ssj_ops and 202W) at 100% and 4825 overall ssj_ops/watt. Comparison based on Dell lab results from Mar’18, Aug’16 and Aug’14. For more information about SPECpower, see spec.org/power_ssj2008/.

The following figure shows the R640's power consumption breakout by major subsystem over a cross section of workload levels. You can see how the PowerEdge Energy Smart design efficiently manages CPU and memory consumption in direct proportion to the workload required as well as minimizes cooling power. The iDRAC9 systems management user interfaces provide real-time, averaged and long-term sensor logging of the major subsystem and overall input power consumptions.

Figure 3. R640 subsystem power consumption across SPECpower workload intervals

R640 Subsystem Power Consumption across SPECpower Loading Intervals



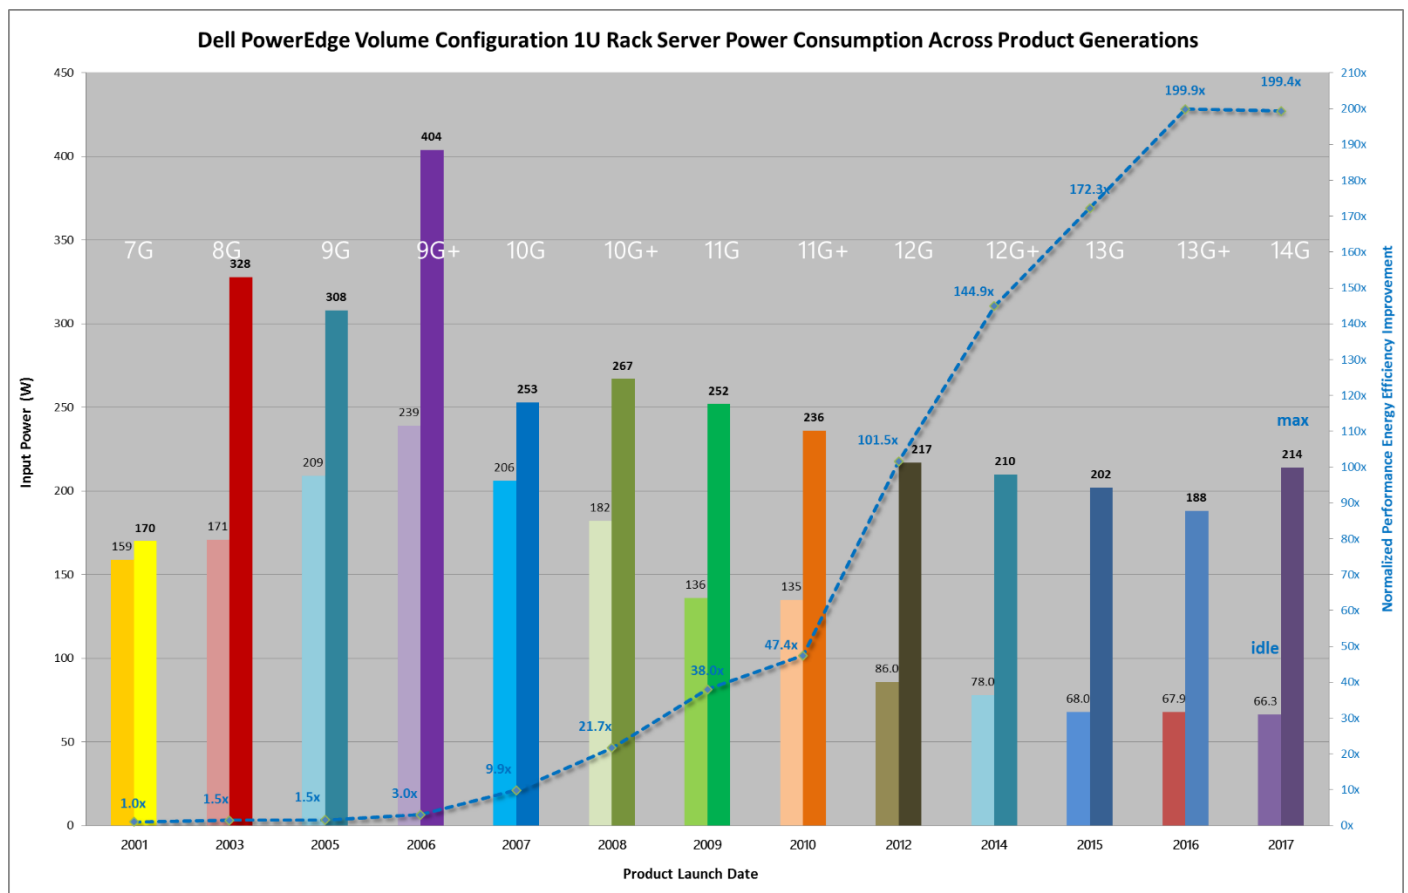
Summary

The PowerEdge R640 1U rack server orderable now with Xeon Silver 4110 CPUs produces 13% more work with equivalent overall energy efficiency to the solution it replaces with Xeon E5-2620 v4 CPUs.

Outlook

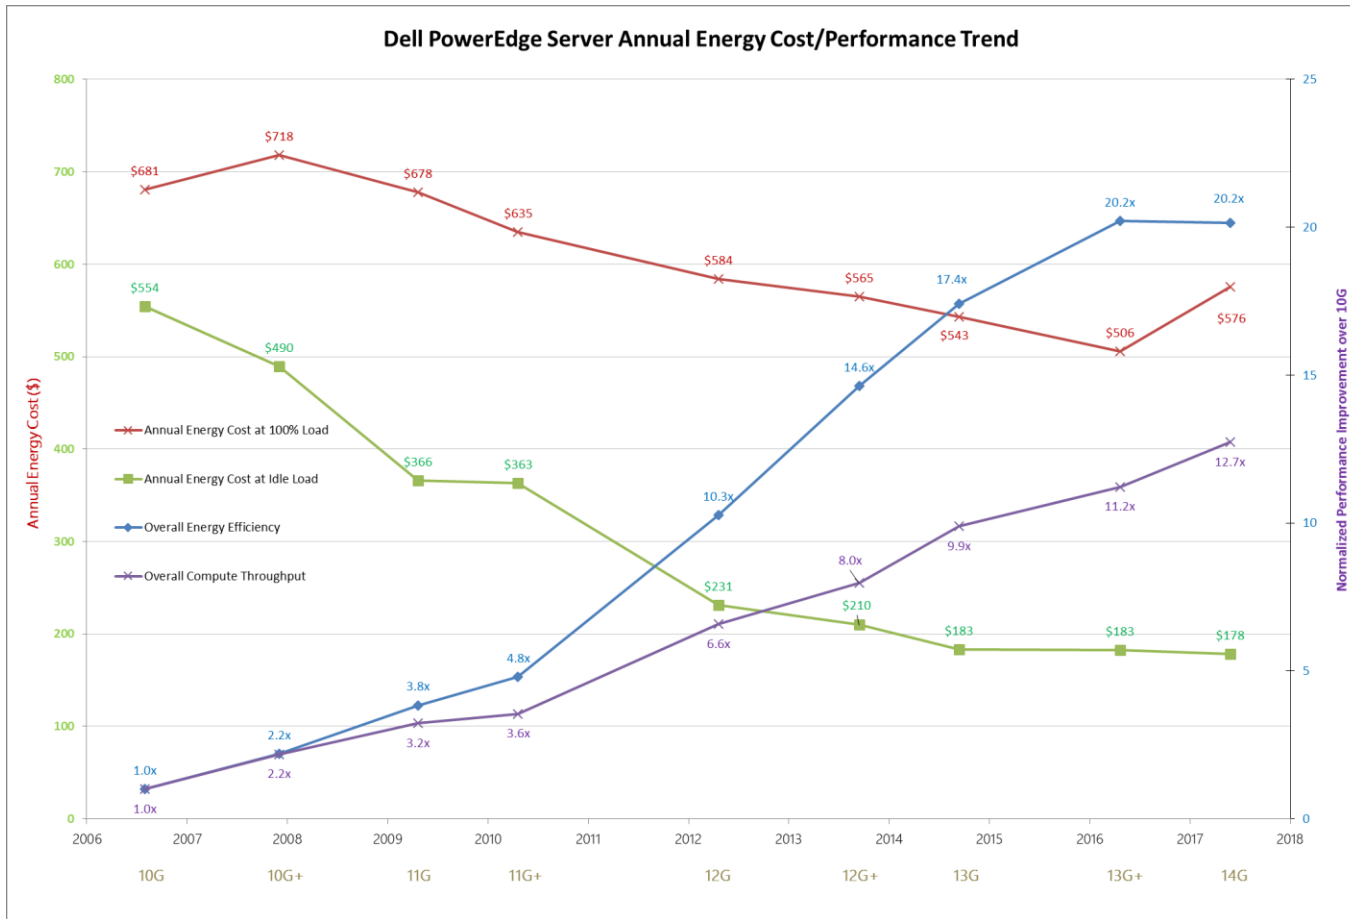
As charted in Figures 3 and 4, the energy efficiency of Dell's typically-configured PowerEdge servers has improved a whopping 200-fold over the past twelve years. Given IT customers' demand for servers that can perform more work while also reducing the data center's physical, electrical, cooling and TCO requirement; Dell continues to supply those solutions.

Figure 4. PowerEdge energy efficiency progress²



² Based upon Dell internal lab recurring evaluation of typically-configured product performance.

Figure 5. PowerEdge energy cost and performance improvement trend³



³ Based upon:

- 1) Enterprise server industry prediction of the fully-burdened cost of a watt (annualized) of **\$2.69**
<http://perspectives.mvdirona.com/2008/12/06/AnnualFullyBurdenedCostOfPower.aspx>
- 2) US EIA 2014, 2016, 2017 average retail price of electricity for commercial sector in (\$ per KWh) of **\$0.11**
http://www.eia.gov/electricity/monthly/epm_table_grapher.cfm?t=epmt_5_6_a

Appendix A—Test methodology

SPECpower_ssj2008 standard

SPECpower_ssj2008 consists of a server-side Java™ (SSJ) workload along with data collection and control services. SPECpower_ssj2008 results portray the server’s performance in ssj_ops (server-side Java operations per second) divided by the power used in watts (ssj_ops/watt). SPEC created SPECpower_ssj2008 to accurately measure the power consumption of servers in relation to the performance that the server is capable of achieving with ssj2008 workload.

SPECpower_ssj2008 consists of three main software components:

- **Server-Side Java (SSJ) Workload** — Java database that stresses the processors, caches and memory of the system, as well as software elements such as OS elements, and the Java implementation chosen to run the benchmark.
- **Power and Temperature Daemon (PTDaemon)** — Program that controls and reports the power analyzer and temperature sensor data.
- **Control and Collect System (CCS)** — Java program that coordinates the collection of all the data.

For more information on how SPECpower_ssj008 works, see spec.org/power_ssj2008/

All results discussed in this white paper are from “compliant runs” in SPEC terminology, which means that although they have not been submitted to SPEC for review, Dell is allowed to disclose them for the purpose of this study. All configuration details required to reproduce these results are given in table 1, 2. The run report summaries are included in the Appendix B SPECpower_ssj2008 results section and the full reports are attached to the pdf version of this white paper.

Both servers ran fresh installs of Microsoft Windows Server Datacenter operating system editions with all Windows update service patches, supplements and hotfixes that were available at the time of running. In the R640 case, Server 2016 RS1 was used. In the R630+ case, Server 2012 R2 (Service Pack 1) was used. Operating systems were installed on a two-drive RAID 1 configuration, choosing the “full installation” option for each.

BIOS settings

The BIOS menu settings were those that the SPA team identified as being the best efficiency practices for the select PowerEdge server model.

Table 2. BIOS settings

PowerEdge R630+	PowerEdge R640
Memory Snoop mode set to Home Snoop	Logical Processor Idling disabled
QPI speed set to 6.4GT/s Data Rate	UPI speed set to 9.6GT/s Data Rate
Adjacent Cache Line Prefetch disabled	Adjacent Cache Line Prefetch disabled
Hardware Prefetcher disabled	Hardware Prefetcher disabled
DCU Streamer Prefetcher disabled	DCU Streamer Prefetcher disabled
DCU IP Prefetcher enabled	DCU IP Prefetcher enabled
CPU Power Management set to System DBPM (DAPC/PowerSaver)	CPU Power Management set to System DBPM (DAPC)
Turbo Boost enabled	Turbo Boost enabled
Energy Efficiency Turbo enabled	Energy Efficiency Turbo enabled
Memory DDR Freq Limit set to Maximum Performance	CPU Interconnect Bus Link Power Management enabled
C1E enabled	C1E enabled
C States set to Autonomous	C states set to Autonomous
Memory Patrol Scrub disabled	Memory Patrol Scrub disabled
Collaborative CPU Performance Control disabled	Collaborative CPU Performance Control enabled
Uncore Frequency set to Dynamic	Uncore Frequency set to Dynamic
Energy Efficiency Policy set to Energy Efficient	Energy Efficiency Policy set to Energy Efficient

OS tuning

To improve Java performance, large pages were enabled by entering **Control Panel > Administrative Tools > Local Security Policy > Local Policies > User Rights Assignment > Lock Pages in Memory**. An option was changed to add Administrator.

The Operating System Power Plan was set to Power Saver for the R630+ configuration testing and Balanced for the R640 configuration testing.

Both servers were configured with a separate IP address on the same subnet as the SPECpower_ssj2008 controller system where the Director, CCS, and PTDaemon components were located. Both servers were connected directly to the controller system through an available network interface port.

SPECpower_ssj2008 configuration

For the R630+ E5-2600 v4 configuration, the Oracle® HotSpot 64-bit server Java Virtual Machine (JVM)⁴ was used with the following run options:

```
-server -Xmx13g -Xms13g -Xmn11g -XX:SurvivorRatio=1 -XX:TargetSurvivorRatio=99  
-XX:ParallelGCThreads=16 -XX:AllocatePrefetchDistance=256 -XX:AllocatePrefetchLines=4  
-XX:LoopUnrollLimit=45 -XX:InitialTenuringThreshold=12 -XX:MaxTenuringThreshold=15  
-XX:InlineSmallCode=9000 -XX:MaxInlineSize=270 -XX:FreqInlineSize=6000 -XX:+UseLargePages  
-XX:+UseParallelOldGC -XX:+AggressiveOpts -XX:+OptimizeStringConcat -XX:+UseStringCache
```

And these logical processor to JVM application thread bindings: **start /NODE[0,1] /AFFINITY [0xFFFF]**

For the R640 4110 configuration, the Oracle® HotSpot 64-bit server Java Virtual Machine (JVM)⁵ was used with the following run options:

```
-server -Xmx21g -Xms21g -Xmn19g -XX:SurvivorRatio=1 -XX:TargetSurvivorRatio=99 -XX:ParallelGCThreads=16 -  
XX:AllocatePrefetchDistance=256 -XX:AllocatePrefetchLines=4 -XX:LoopUnrollLimit=45 -  
XX:InitialTenuringThreshold=15 -XX:MaxTenuringThreshold=15 -XX:InlineSmallCode=3900 -XX:MaxInlineSize=270 -  
XX:FreqInlineSize=2500 -XX:+AggressiveOpts -XX:+UseLargePages -XX:+UseParallelOldGC
```

And these logical processor to JVM application thread bindings: **start /NODE[0,1] /AFFINITY [0xFFFF]**

Power and temperature configuration

A Yokogawa WT210 Digital Power Meter was used for the actual power measurement of the servers, as this was the most commonly used analyzer for SPECpower_ssj2008 publications at the time that this study was conducted. The WT210 unit used was within its one-year calibration certification period to ensure accurate power consumption measurements. Input line voltage supplying both systems varied by less than 1V.

To ensure a fair comparison, the systems were run within a temperature controlled chamber with inlet temperatures measured at the front chassis bezel of both systems using a Digi® International Watchport®/H temperature probe. As the attached reports show, temperatures over each run were held constant to within 0.5 °C.

⁴ Build 24.80-b11, mixed mode, version 1.7.0_80

⁵ Build 24.80-b11, mixed mode, version 1.7.0_80

Appendix B—SPECpower_ssj2008 results

Figure 6. SPECpower_ssj2008 results for Dell PowerEdge R640

SPECpower_ssj2008

Copyright © 2007-2018 Standard Performance Evaluation Corporation

Dell Inc. PowerEdge R640 (Intel Xeon Silver 4110 2.10 GHz)			SPECpower_ssj2008 = 5,583 overall ssj_ops/watt		
Test Sponsor:	Dell Inc.	SPEC License #:	55	Test Method:	Single Node
Tested By:	Dell Inc.	Test Location:	Round Rock, TX, USA	Test Date:	Mar 7, 2018
Hardware Availability:	Dec-2017	Software Availability:	Nov-2017	Publication:	Unpublished
System Source:	Single Supplier	System Designation:	Server	Power Provisioning:	Line-powered

Benchmark Results Summary

Performance			Power	Performance to Power Ratio
Target Load	Actual Load	ssj_ops	Average Active Power (W)	
100%	99.5%	1,625,307	214	7,585
90%	90.0%	1,468,668	202	7,284
80%	80.0%	1,305,396	187	6,990
70%	70.1%	1,145,167	169	6,790
60%	60.2%	982,214	152	6,459
50%	50.2%	820,039	141	5,799
40%	40.1%	654,243	132	4,960
30%	30.0%	489,136	123	3,972
20%	20.0%	326,512	115	2,827
10%	10.0%	163,340	107	1,530
Active Idle		0	66.3	0
$\Sigma \text{ssj_ops} / \Sigma \text{power} =$				5,583




Figure 7. SPECpower_ssj2008 results for the Dell PowerEdge R630+

SPECpower_ssj2008

Copyright © 2007-2016 Standard Performance Evaluation Corporation

Dell Inc. PowerEdge R630 (Intel Xeon E5-2620 v4 2.10 GHz)		SPECpower_ssj2008 = 5,597 overall ssj_ops/watt			
Test Sponsor:	Dell Inc.	SPEC License #:	55	Test Method:	Single Node
Tested By:	Dell Inc.	Test Location:	Round Rock, TX, USA	Test Date:	Jul 12, 2016
Hardware Availability:	Mar-2016	Software Availability:	Apr-2015	Publication:	Unpublished
System Source:	Single Supplier	System Designation:	Server	Power Provisioning:	Line-powered

Benchmark Results Summary

Performance			Power	Performance to Power Ratio
Target Load	Actual Load	ssj_ops	Average Active Power (W)	
100%	100.1%	1,432,168	188	7,604
90%	91.6%	1,310,033	174	7,539
80%	80.1%	1,145,739	155	7,397
70%	69.7%	996,778	142	7,012
60%	60.0%	858,096	131	6,536
50%	49.9%	713,972	123	5,802
40%	40.1%	572,927	117	4,903
30%	29.9%	427,761	110	3,878
20%	19.9%	284,274	104	2,741
10%	10.0%	142,919	96.4	1,482
Active Idle		0	67.9	0
$\sum \text{ssj_ops} / \sum \text{power} =$				5,597

