

Dell Hybrid High Performance Computing (HPC)

Meeting the Demands of HPC and Big Data Applications by Leveraging Hybrid CPU/GPU Computing

Tolly Report #215158
Commissioned by
Dell Inc.

November 2015





Executive Summary

"Rack 'em and stack 'em."— a winning approach for a long time but not without its limitations. A generalized server solution works best when the applications running on those servers have generalized needs.

Enter "Big Data." Today's application and workload environments can be required to process massive amounts of granular data and, thus, often consist of applications that place high demands on different server hardware elements. Some applications are very compute intensive and place a high demand on the server's CPU where others in the same environment are tasked with unique processing requirements performed on specialized graphical processing units (GPUs).

Whether it is customer, demographic, seismic data — or a whole host of other uses — the number crunching and processing required across the suite of applications can result in processing demands that are radically different from demands of prior years.

Hybrid High Performance Computing (HPC)

Enter Hybrid High Performance Computing. These systems are built to serve two masters: CPU-intensive applications and GPU-intensive applications delivering a hybrid environment where workloads can be optimized and run-times reduced through ideal resource utilization.

Deploying Dell PowerEdge C6320¹ and C4130² servers in a hybrid configuration optimizes your compute density and efficiency for demanding general purpose workloads and can be combined to run high value GPU intensive workloads in the same environment. This produces the optimal blend of compute, density, and high value for key workloads at the best value in the most efficient footprint. These latest servers from Dell offer impressive generational improvements in many areas including cores, memory, density and power efficiency.

Dell HPC building blocks can be customized to client needs and can be easily deployed in a heterogeneous sub-cluster with simple scaling and using normal data center power and cooling. Dell HPC solutions are available globally and thus are a viable option for even the most far-flung organizations requiring the support and stability of a true worldwide partner.

This new generation of Dell server systems brings demonstrable improvements in key areas when compared both to prior generation Dell offerings as well as competing solutions. This paper will explore usage scenarios as well as highlight some of the significant improvements in areas such as: capacity, performance, power efficiency and remote system management and support.

THE BOTTOM LINE

The Dell Hybrid HPC solution delivers a hybrid CPU and GPU compute environment with the PowerEdge C6320 and C4130 to:

- 1 Optimize workloads across CPU/GPU servers
- 2 Deliver the highest-density, highest-performance in a small footprint
- 3 Provide significant power, cooling and resource utilization benefits
- 4 Lower cost of ownership and enhance reliability through integrated Dell Remote Access Controller (iDRAC) and Lifecycle Controller

¹ <http://www.dell.com/learn/us/en/uscorp1/press-releases/2015-06-23-dell-pushes-boundaries-of-hpc>

² <http://www.dell.com/learn/us/en/uscorp1/press-releases/2014-11-17-dell-high-performance-computing>



Dell Hybrid HPC: PowerEdge C6320 & PowerEdge C4130 Highlights

- **High Capacity**

The PowerEdge C6320 is a 2U chassis running Intel Xeon E5-2600v3 processors. A single chassis can provide up to 18 cores per socket or up to 144 cores per chassis, up to 512GB of DDR4 memory per server and up to 72TB of local storage.

The PowerEdge C4130 is a 1U chassis running Intel Xeon processors. The server is the only Intel Xeon E5-2600v3 1U server to offer up to four GPUs/accelerators³.

- **High Performance**

The PowerEdge C6320 is designed to offer four independent server nodes in a 2U chassis and can deliver 999 Gigafllops on a single server. This is a 2X LinPack spec improvement over previous generation systems⁴. The system provides up to a 45 percent improvement on the SPECint_rate benchmark over prior generation hardware.⁵

The PowerEdge C4130 can achieve over 7.2 Terafllops on a single 1U server.⁶

- **Power Efficiency**

The PowerEdge C6320 provides up to 28 percent better power efficiency on the Spec_Power benchmark.⁷

The PowerEdge C4130 has a performance/watt ratio of up to 4.17 Gigafllops per watt.

- **Sophisticated Management**

Both systems leverage the agent-free capabilities of the integrated Dell Remote Access Controller (iDRAC) with Lifecycle Controller technology.

³ Based on Dell internal analysis of number of GPUs per 1U of rack space performed against SuperMicro SYS 1028GR-TR and HP SL250s Gen8 (SL6500) in November 2014.

⁴ Based on the LinPack benchmark, the Dell PowerEdge C6320 achieved 999 Gflops using two Intel Xeon E5-2699 processors. Dell PowerEdge C6220 achieved 498 GFlops using two Intel Xeon E5-2697 processors on the same benchmark. Based on testing by Dell Labs in June 2015.

⁵ Based on testing by Dell Labs in June 2015. The Dell PowerEdge C6320 achieved a SPECint_rate of 1,3320 using two Xeon E5-2699 processors. The Dell PowerEdge C6220 achieved a SPECint_rate of 919 using two Xeon E5-2697 processors. For latest SPEC CPU2006 benchmark results, visit www.spec.org.

⁶ Based on Dell Labs testing performed November 2014 using the High Performance Linpack (HPL) benchmark test. Actual performance will vary with configuration, usage and manufacturing variability.

⁷ Based on testing by Dell Labs in June 2015. The 4 node Dell PowerEdge C6320 cluster achieved a SPECpower_ssj2008 result of 3,235,764 overall ssj_ops/Node, compared to the 4 node Dell PowerEdge C6220 cluster with a SPECpower_ssj2008 result of 1,678,974 overall ssj_ops/Node score. For the latest SPECpower_ssj2008 benchmark results, visit www.spec.org.



Hybrid Landscape & Big Data

So-called “Super-computing” is not just for research labs anymore. Processing seismic data to select the best drilling location and chromosome sequencing are just two of the higher profile examples that most of us have run across. The proliferation of “big data” across any number of commercial environments, though, brings the need for super-computing to the doorstep of many organizations that might not have considered themselves candidates for super-computing previously.

Dell commissioned Tolly to analyze how customers can leverage a hybrid solution using the PowerEdge C6320 & PowerEdge C4130 to maximize and optimize processing of large, complex data sets in conjunction with more traditional server computing.

Gartner: “Through 2015, 85% of Fortune 500 organizations will be unable to exploit big data for competitive advantage.”

According to Gartner, the vast majority of Fortune 500 companies are behind the curve when it comes to leveraging big data.⁸ If you are one of those firms, implementing the Dell hybrid solution can get you ahead of that curve in short order.

Big data has been a fact of life for industries such as oil and gas exploration with seismic research and for government and advanced military research projects for some time now and will only grow as more sophisticated data collection mechanisms gather ever more data. New to the mix, however, are more broad-based business enterprises. The massive amounts of granular data that are collected on a regular basis by common applications such as web browsing and even browsing of physical stores using WLAN and RFID tracking only have value if they can be analyzed and used to benefit strategic business goals. The 2014 IDG Enterprise Big Data study found that organizations are witnessing dramatic growth in data managed with projections of up to 76% increase by mid-2016 with nearly half the respondents already implementing big data or in the process of doing so.⁹

⁸ <http://www.gartner.com/technology/topics/big-data.jsp>

⁹ <http://www.idgenterprise.com/report/big-data-2>

Hybrid HPC Use Cases

#1 - Energy/Oil & Gas Exploration

Business Need

Locating profitable, new deepwater oil and gas fields globally. To accomplish this requires running advanced applications that place great demands on CPU and GPU resources. Applications simulate and optimize seismic and reservoir modeling, hydro geophysics, deep sea drilling, chronostratigraphy, 3D basin modeling and other engineering challenges. For example, Wide-Azimuth (WAZ) surveys and 3D imaging combined with greater numbers of advanced sensors have resulted in a dramatic increase in the data that must be processed.

Innovative Solution

Massively scalable CPU and GPU solutions with efficient rack footprint and sophisticated energy management provide cost-effective resources for handling both CPU-intensive and GPU-intensive tasks.¹⁰

Dell Benefits

- *Massive computing power and systems manageability* — Dell HPC can deliver the power and ease of mastery required for the most advanced computation tasks.
- *Scalable platform architecture* — Can allow for system expansion and growth as computational needs increase.
- *Standards-based computing* — Can streamline design and deployment, which delivers returns on your IT investment fast.



¹⁰ http://www.dell.com/downloads/global/solutions/industry_apps/oil_gas_ae_leavebehind.pdf

#2- University Supercomputer Deployment

Business Need

Raw Power: A petascale supercomputer¹¹ computer complex will support a vast range of processing requirements across a diverse set of scientific disciplines. Provide HPC power to both large and smaller scale research projects in a cost-effective, flexible and state-of-the-art fashion.

Innovative Solution

Dell HPC supercomputer comprised of PowerEdge C6320: 1,944 nodes/46,656 cores. Runs high-end applications in biochemistry, bioinformatics, large-scale data analysis, predictive analytics, quantum chemistry and other areas.

The HPC complex also runs 36 Dell high-density GPU PowerEdge C4130 nodes that provide dramatic improvements on GPU processing power per RU over previous generation technology.

Dell Benefits

Dell designed the system with the customer to enable the customer to deploy the latest processor, storage, GPU and networking technology thus maximizing the value proposition. Delivers advanced compute capacity to new research domains and users traditionally unable to leverage high-performance computing.

A partial list of customer benefits from choosing Dell include:

- Partnering on system design with Dell experts
- Leveraging “bleeding edge” technologies desired by the customer
- Compute Performance provided by PowerEdge C6320
- High-density GPU performance provided by PowerEdge C4130
- System Reliability & ongoing support
- Rapid delivery and ease of deployment



¹¹ XSEDE is an example Dell HPC deployment. See <https://portal.xsede.org/sdsc-comet>

The Rack Revisited

The strategy that worked in the past needs to be updated in anticipation of the future — and that future is now. The question: How to handle broad based, “traditional” server processing yet be able to handle the specialized processing demands of big data? The answer: Dell PowerEdge C6320 servers for the main, broad-based server processing augmented with Dell PowerEdge C4310 servers to provide the dedicated GPU/acceleration needed to handle big data in a timely fashion. This combination of systems provides benefits unique in the server market. For a systems overview, see Table 1.

Dell Hybrid HPC: PowerEdge C6320 & C4130 System Highlights

System	Role	Chassis	Processor	Max. Memory	Web Link
Dell PowerEdge C6320 Rack Server	High Density, Hyper-Converged Computing Solution	2U (Half-height sled, chassis can support up to four sleds.)	Up to 4x, 2-socket servers with Intel® Xeon® Processor E5-2600 v3 Product Family. Up to 18 cores per socket, 144 cores per 2U chassis.	512GB (16 DIMM per motherboard)	http://www.dell.com/us/business/p/poweredge-c6320/pd?oc=&model_id=poweredge-c6320
Dell PowerEdge C4130 Rack Server	High-Performance, High-Density GPU Solution	1U	Up to 2x Intel® Xeon® Processor E5-2600 v3 Product Family	256GB (16 slots)	http://www.dell.com/us/business/p/poweredge-c4130/pd?oc=&model_id=poweredge-c4130

Note: The Dell PowerEdge C4130 offers several distinct configurations of GPU to CPU connectivity to optimize the system for particular customer requirements.

Source: Dell, October 2015

Table 1

Dell PowerEdge C6320 Rack Server

Hyper-converged computing is what the PowerEdge C6320 is all about. It can run up to four independent server systems in a single 2U chassis. A single chassis system offers up to 144 CPU cores, 512GB of memory per server and 72 TB of local storage.

Dell has selected the PowerEdge C6320 to be the platform for several flagship Dell Engineered Solutions: VMware EVO: RAIL and Dell's own XC Series Web-scale converged appliance. The server has even made CRN's list of “The 10 Coolest Servers of 2015.”¹²

¹² <http://www.crn.com/slide-shows/storage/300077368/the-10-coolest-servers-of-2015-so-far.htm/pgno/0/1>

Think of the PowerEdge C6320 as the workhorse of your environment. If its predecessor PowerEdge C6220 wasn't powerful enough, the Generation 13 PowerEdge C6320 brings even more compute power into your server rack. Its hallmarks are performance, power efficiency and high density.


FEATURED IN **CRN**

CRN: "PowerEdge C6320 is one of the 10 coolest servers of 2015."

Dell PowerEdge C6320



Dell PowerEdge C6000 LINPACK Benchmark Results as reported by Dell Labs



Note: Testing conducted by Dell Labs in June 2015. PowerEdge C6320 used two Intel Xeon E5-2699 processors, PowerEdge C6220 used two Intel Xeon E5-2697 processors.

Source: Dell, June 2015


Figure 1

Performance Improvements

LINPACK Linear Algebra Calculations

In benchmarks running the LINPACK super computer benchmark, the PowerEdge C6320 delivered twice the performance of the predecessor PowerEdge C6220: 999 Gigaflops compared to 498 Gigaflops. See Figure 1.

Dell PowerEdge C6000 SPECint_rate Benchmark



Note: Testing conducted by Dell Labs in June 2015. PowerEdge C6320 used two Intel Xeon E5-2699 processors, PowerEdge C6220 used two Intel Xeon E5-2697 processors.


Source: Dell, June 2015

Figure 2

SPECint - Integer Processing

In benchmarks running the SPEC® SPECint_rate CPU integer processing benchmark, which is part of the SPEC CPU2006 test suite, the PowerEdge C6320 delivered more than 30% greater throughput than the performance of the predecessor PowerEdge C6220 at 1,330 vs 919. See Figure 2.

Dell PowerEdge C6000 SPECpower Benchmark



Note: Testing conducted by Dell Labs in June 2015. PowerEdge C6320 used two Intel Xeon E5-2699 processors, PowerEdge C6220 used two Intel Xeon E5-2697 processors. Test was server-side Java 2008 (SPECpower_ssj2008).

Source: Dell, June 2015

Figure 3

Power Efficiency

Similarly, in benchmarks of processing power as it relates to power consumption, the PowerEdge C6320 shows significant efficiency gains over the predecessor PowerEdge C6220 with respect to workload per watt of power consumed.

In the “SSJ” server-side Java test, the PowerEdge C6320 was able to achieve more than twice the number of operations per Watt than the predecessor technology. The PowerEdge C6320 completed 9,472 operations per watt of power consumed compared to less than 5,000 operations per watt for the predecessor technology PowerEdge R620. See Figure 3.



More Compute Power Rack

Because of its enhanced power efficiency, the PowerEdge C6320 can be used to build a high-density rack of up to 80 servers utilizing normal data center power and cooling. Its hot-swappable disks allow you to keep running even in the event of a disk failure.

One 42U rack = 80 Servers.

20 x 2U x 4 servers per chassis = 80 Servers

A single 42U rack can hold 20 of the 2U PowerEdge C6320 and each of those chassis can deliver four servers.

Assuming that the maximum power density of a 42U rack is 30kWh without requiring any special data center power or cooling, then a full

rack of C6320 servers (20 x 2U x 4 servers = 80 servers in a rack) should be able to fit into a 42U rack and when managed with Dell OpenManage with Dynamic Power Capping deliver optimal compute performance with normal data center power and cooling.

Dell PowerEdge C4310 Rack Server

High-density GPU computing is what the PowerEdge C4310 is all about. offers one or two CPUs and up to four GPUs(NVIDIA Tesla K80 or K40 GPU accelerators) that can, optionally, include a 96-lane PCIe switch.

Dell PowerEdge C4130

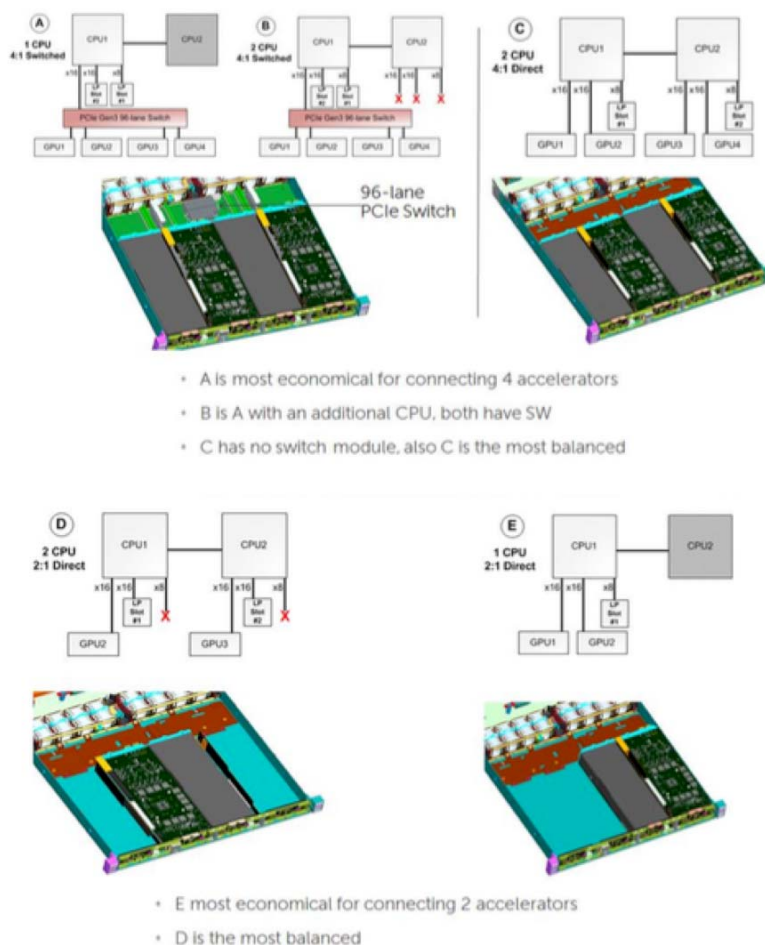


Deploying Dell PowerEdge C6320 and C4130 in a hybrid configuration optimizes compute density and efficiency for demanding general purpose workloads and can be combined with high value GPU intensive workloads running in the same environment to produce the optimal blend of compute, density, high value for key workloads at the best value in the most efficient footprint.

GPU computing is now established and widespread in the HPC community. There is an ever-increasing demand for compute power. This demand has pushed server designs towards higher hardware accelerator density. However, most such designs have a standard system configuration, which may not be optimal for maximum performance across all application classes.

The C4130 offers five carefully selected system configurations, potentially making it a better fit for a greater variety of extreme HPC applications. The motivation behind offering this configuration is to offer a system design that will best fit the application requirement. This offers users a unique opportunity to match the system design with their application characteristics. The configurations differ in the manner in which accelerators connect to the CPUs, resulting in different accelerator to CPU ratios. They vary from balanced to clustered configurations. See Figure 4. More important than studying the configuration diagram is understanding what the configuration options give you - flexibility.

Dell PowerEdge C4130: GPU/CPU Configuration Options



Source: Dell, October 2015

Figure 4

PowerEdge C4130 GPU Configuration Flexibility

- **Peer-to-Peer**

Single or dual-CPU configurations in systems outfitted with four GPU boards and the PCIe switch module are optimized for peer-to-peer operations.

- **Balanced**

Three other configurations of two or four GPU boards and single or dual-CPU without switching provide increasing levels of performance for balanced applications.

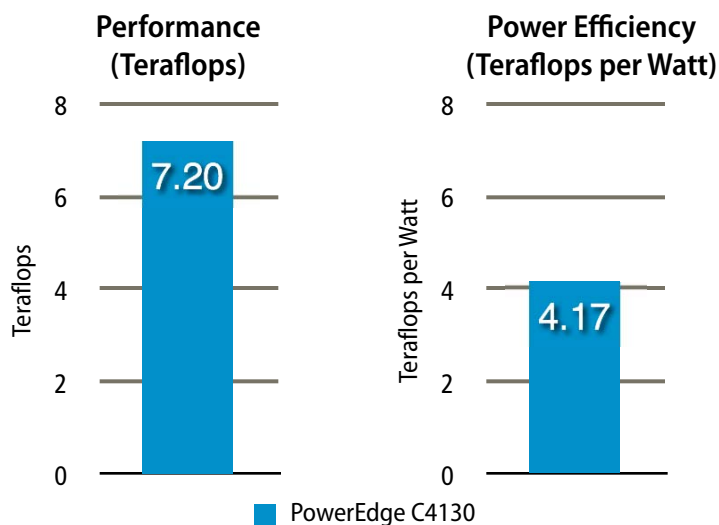
GPU Density

The PowerEdge C4130 delivers up to 33 percent better GPU/accelerator density than its closest competitors¹³ and 400 percent more PCIe GPU/accelerators per processor per rack unit than a comparable HP Gen8 system.¹⁴

Performance

The PowerEdge C4130 can achieve over 7.2 Teraflops on a single 1U server and has a performance/watt ratio of up to 4.17 Gigaflops per watt. See Figure 5.

Dell PowerEdge C4130 LINPACK Benchmark Results as reported by Dell Labs



Note: Testing conducted by Dell Labs in November 2014.

Source: Dell, June 2015

Figure 5

¹³ Based on Dell internal analysis of number of GPUs per 1U of rack space performed against SuperMicro SYS 1028GR-TR and HP SL250s Gen8 (SL6500) in November 2014.

¹⁴ Based on Dell internal analysis performed November 2014.



Dell Remote Access Controller (iDRAC) with Lifecycle Controller

Both the Dell PowerEdge C6320 and PowerEdge C4130 servers are fully supported by Dell OpenManage as well as implementing iDRAC8 - the latest version of the agent-free remote support feature.

Dell iDRAC adds value in:

- 1) Deployment, 2) Configuration Drift,
- 3) Power Monitoring/Capping

An integrated Dell Remote Access Controller (iDRAC) with Lifecycle Controller is embedded in every Dell PowerEdge server. It provides functionality that helps you deploy, update, monitor and maintain Dell PowerEdge servers with or without a systems management software agent. Because it is embedded within each server from the factory, the Dell iDRAC with Lifecycle Controller requires no operating system or

hypervisor to work. This agent-free operation means that with OpenManage technology, your new server is ready to deploy a new OS or accept new settings with no extra software to install. With powerful, easy-to-use, remote management and configuration options, iDRAC with Lifecycle Controller can alert you when an issue occurs, no matter where you are. It enables streamlined local and remote server management, and reduces or eliminates the need for administrators to physically visit the server — even if the server is not operational.

As an integral part of iDRAC, the Dell Lifecycle Controller simplifies server lifecycle management tasks like provisioning, deployment, servicing, user customization, patching and updating. It's a collection of out-of-band automation services, embedded pre-OS applications, and remote interfaces that give you deployment, update, and maintenance capabilities through managed, persistent storage. This means you can reduce or eliminate dependence on media-based systems management tools and utilities.

Lifecycle Controller reduces your time spent on management tasks, reduces potential for error, improves security, and increases overall efficiency in your IT environment.

And because it functions out of band, iDRAC with Lifecycle Controller doesn't require you to install the OpenManage Server Administrator software agent.

Deployment

Going from "out of the box" to "custom configured" can be a challenge, but iDRAC¹⁵ gives you tools that save you time and provide you with flexibility. Prior versions already provided for automatic update of server firmware. iDRAC8 brings even more benefits.

¹⁵ Not all of the features/functionality listed herein apply to the C4130/C6320; contact your Dell sales representative for details.



New with iDRAC8, it allows easy and complete server configuration with either a USB memory stick or the direct connection of a laptop to the system's USB port.

Other highlights of iDRAC8 include: 1) Easy restoration of server settings when replacing a system motherboard, 2) Agent-free system performance monitoring, for performance monitoring that doesn't impact workload performance, and 3) Out-of-band storage monitoring and management of the Dell server's PERC9 RAID controller and SAS Controller Monitoring.

Configuration Drift

Over time, servers that began life configured identically can, because of manual adjustments, become out of sync - a condition known as "configuration drift." Managing groups of servers that have drifted is challenging because the exact configuration of each can only be determined by inspecting each. Manual management of server baselines is not only prone to error by very time consuming.

Using Dell OpenManage Essentials (OME) baselines can be created against which OME can conduct regular, non-disruptive inventory to detect situations where configuration drift has occurred. According to the same, aforementioned report, using Dell OME to conduct this automated maintenance requires 91% less time than the manual approach.

Power Monitoring/Capping

When it comes to maximizing capacity and minimizing operational cost monitoring and controlling power is an essential element. When used with a 12th and 13th generation PowerEdge server with an iDRAC Enterprise license on a system with an Intel chipset, Dell OpenManage Power Center (OMPC) leverages Intel Node Manager technology for platform-level power reporting and capping of Intel chipsets. Power Center then communicates with the Dell iDRAC on each managed server to provide node, rack, row or data-center level aggregation of power-management data, as well as execution of control policy - making it easy for IT professionals to identify areas to gain efficiencies and cut wasteful costs.¹⁶

Manage Power via OMPC:

- *Measure consumption at server, rack, row or room level*
- *Set power caps, individually or in groups, to maximize up-time of critical application by reducing non-critical consumption*
- *Enable oversubscribed power supply configurations*

¹⁶ <http://en.community.dell.com/techcenter/power-cooling/w/wiki/3534.dell-openmanage-power-center>



About Tolly...

The Tolly Group companies have been delivering world-class IT services for over 25 years. Tolly is a leading global provider of third-party validation services for vendors of IT products, components and services.

You can reach the company by email at sales@tolly.com, or by telephone at +1 561.391.5610.

Visit Tolly on the Internet at:

<http://www.tolly.com>

Terms of Usage

This document is provided, free-of-charge, to help you understand whether a given product, technology or service merits additional investigation for your particular needs. Any decision to purchase a product must be based on your own assessment of suitability based on your needs. The document should never be used as a substitute for advice from a qualified IT or business professional. This evaluation was focused on illustrating specific features and/or performance of the product(s) and was conducted under controlled, laboratory conditions. Certain tests may have been tailored to reflect performance under ideal conditions; performance may vary under real-world conditions. Users should run tests based on their own real-world scenarios to validate performance for their own networks.

Reasonable efforts were made to ensure the accuracy of the data contained herein but errors and/or oversights can occur. The test/audit documented herein may also rely on various test tools the accuracy of which is beyond our control. Furthermore, the document relies on certain representations by the sponsor that are beyond our control to verify. Among these is that the software/hardware tested is production or production track and is, or will be, available in equivalent or better form to commercial customers. Accordingly, this document is provided "as is", and Tolly Enterprises, LLC (Tolly) gives no warranty, representation or undertaking, whether express or implied, and accepts no legal responsibility, whether direct or indirect, for the accuracy, completeness, usefulness or suitability of any information contained herein. By reviewing this document, you agree that your use of any information contained herein is at your own risk, and you accept all risks and responsibility for losses, damages, costs and other consequences resulting directly or indirectly from any information or material available on it. Tolly is not responsible for, and you agree to hold Tolly and its related affiliates harmless from any loss, harm, injury or damage resulting from or arising out of your use of or reliance on any of the information provided herein.

Tolly makes no claim as to whether any product or company described herein is suitable for investment. You should obtain your own independent professional advice, whether legal, accounting or otherwise, before proceeding with any investment or project related to any information, products or companies described herein. When foreign translations exist, the English document is considered authoritative. To assure accuracy, only use documents downloaded directly from Tolly.com.

No part of any document may be reproduced, in whole or in part, without the specific written permission of Tolly. All trademarks used in the document are owned by their respective owners. You agree not to use any trademark in or as the whole or part of your own trademarks in connection with any activities, products or services which are not ours, or in a manner which may be confusing, misleading or deceptive or in a manner that disparages us or our information, projects or developments.