5 Ways to Ensure **Reliability**, **Availability**, and **Serviceability** in Your Enterprise Environment

If your RISC-based architecture isn't meeting your needs, it's time to look at today's x86 platforms. Dell's future-ready PowerEdge R930 provides the performance and reliability that mission-critical applications need.

EXECUTIVE SUMMARY

Businesses dependent on some of the most demanding enterprise applications — such as in-memory databases, CRM, and ERP — have long turned to RISC-based platforms to ensure the productivity and efficiency they need. But RISC architectures are coming up short when it comes to the needs of today's computing environments. It's time to take a look at the latest generation of x86 platforms, like the Dell PowerEdge R930. This future-ready, four-socket server rivals RISC performance and delivers the reliability, availability, and serviceability (RAS) capabilities enterprises want. There are five key factors to consider when assessing an x86 platform: Does it have a foundation built on reliability, a complete systems approach, virtualization that ensures highly available data and processing, responsive management, and tightly integrated, enterprise-savvy partners? The Dell PowerEdge R930 has them all.

Dell's top-end four-socket PowerEdge server, the R930, must be on the shortlist for any enterprise that is looking at refreshing its RISC-based architecture. Enterprises whose productivity and success depend on large-scale, mission-critical applications have historically run high-end RISC platforms as the basis of their IT infrastructure. But in today's computing environment, where x86 platforms have closed any significant performance and reliability gaps, enterprises should seriously reconsider x86 options to take advantage of their well-documented efficiencies.

The R930, powered by Intel's latest E7 processors (18-core maximum), not only rivals RISC performance, but in this latest generation it also delivers the RAS capabilities required by applications that demand extremely high service levels. These sorts of applications include search-intensive data warehouses, highly concurrent online transaction processing, and high-performance computing (HPC). This paper describes five major points to help IT professionals evaluate the reliability and suitability of Dell's x86-based enterprise environment.

ONE Reliability at the Core

At the core of RAS capabilities are highly reliable system components. Dell strives to increase product design and reliability with every generation of product. We "design in" reliability with advanced engineering that occurs at all levels of the product prior to release: component, subsystem, and system. We ensure reliability with stringent processes and leading validation practices — routinely validating system components beyond expected customer usage models and environments to uncover weaknesses. We have invested $26 million in R&D, including test and measurement equipment, with more than 1,000 lab benches used for product assessment, testing, and qualification. Our suppliers are required to demonstrate reliability on an ongoing basis as well.

The Intel Xeon processor E7 line is a core building block of Dell's R930 server solutions. Not only does it provide this four-socket server with the performance necessary for traditional mission-critical applications such as ERP, CRM, accounting, sales, and HR systems, it also enables higher levels of memory capacities, making in-memory database solutions not

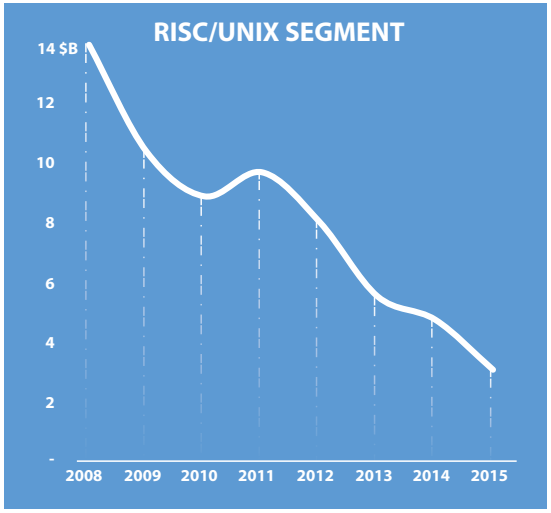

Source: IDCServer Tracker2015Q3-RISC Based System Vendor Revenues

just viable, but highly performant.

The Intel E7 processor provides a powerful engine for both in-memory and traditional workloads, with support for up to 6 TB of memory in a four-socket system. It supports both DDR4 and DDR3 memory technologies, so enterprises can balance cost versus performance more flexibly as they design and grow their data centers.

To further increase fault tolerance, E7 processors include mechanisms for memory entity sparing, mirroring, and failover. Sparing and mirroring are two RAS features that allow on-the-fly failover from a failing component to another component. Sparing allows failover to a physical spare in the same memory controller, and mirroring preserves data in the case of DRAM component failure.

TWO A Complete Ecosystem With RAS Capabilities for the Most Demanding Systems

True RAS depends on more than just reliable components. It must include a layered "ecosystem" of hardware and software along with integrated management and monitoring. Each component builds on top of others to deliver the overall RAS capability within a complete system and infrastructure.

Yes, Dell delivers mission-critical performance and reliability by designing in RAS, not only from the component level (including suppliers), but also in the platform design, at the OS and hypervisor levels, and with its industry-leading approach to systems management. Doing so results in complete solutions that help enterprise customers succeed in the most demanding environments.

Importantly, this ecosystem approach allows for the inclusion of the latest technologies; the inability to do so is one of the limitations driving the dramatic move away from RISC/Unix architecture. This approach also puts Dell at the forefront of virtualization environments.

And Dell brings unique systems management capabilities to this ecosystem, letting customers manage and monitor their entire infrastructure in a way that minimizes downtime, including planned maintenance.

Finally, this ecosystem permits us to benefit from integrating with enterprise-savvy technology partners to offer even higher levels of reliability to mission-critical operations.

THREE Better Virtualization for Higher Data and Processing Availability

In today's computing environment, virtualization allows computing resources (both processing and storage) to be accessed and shared more readily across large infrastructures. Innovations like live migration have made it easier to ensure that critical operations stay up and running. Virtualization enabled the massive scale-out computing seen now in high performance computing (HPC) installations and the ever-widening use of the big data paradigm. Initially virtualization was used strictly for server virtualization to consolidate processing resources. But over time and with the introduction of innovations, storage virtualization (software-defined storage) has come to play an important role in the design of modern enterprise data centers. And most recently, accessible memory has become large enough that it is feasible to virtualize very large, mission-critical applications.

However, scaling is not the only aspect of virtualization having an impact on IT. Virtualization, by its very nature, makes a wider set of resources available to workloads running in an enterprise. Virtualization-enabled software-defined storage has allowed higher levels of data availability, while live migration of virtual machines lets enterprises

Dell PowerEdge Four-Socket Portfolio

R930

- Four-socket server in a 4U chassis
- Intel Xeon E7 8800/4800/2800 processors
- Up to 96 DDR4 DIMMs of memory
- Up to 24 drives with hot-plug SSD or SAS hard drives
- Optional eight front-accessible PCIe Express Flash SSDs
- 10 PCIe Gen3 slots — dedicated RAID and NDC slots
- Dual PERC option
- Dual SD cards for redundant hypervisor
- New Intel Xeon E7 RAS Run Sure Technology
- Optional external PERC (H830)
- OpenManage management: iDRAC8 Enterprise with Lifecycle Controller

FC830

- Two full-width, four-socket servers in a 2U chassis
- Intel Xeon E5-4600 processors
- Up to 48 DDR4 DIMMs of memory
- Access to 32 DAS drives through FD332s
- Supports up to two Express Flash SSDs
- 8x2.5-inch HDD/SSD or six HDD/SSD plus two NVMe or 16x1.8-inch SSD
- Access to eight PCIe expansion slots
- Dual SD cards for redundant hypervisor

shift available resources to where they are needed and to pre-emptively shift processing to more stable parts of the infrastructure. Both data and processing become more available.

Dell brings a number of added-value features to the enterprise virtualization environment.

The R930 (and other PowerEdge servers) can use dual embedded SD cards to provide redundancy at the hypervisor level for VMware vSphere. If the operative copy of the hypervisor experiences corruption or is dislodged, the redundant copy can take over. This eliminates the single point of failure found in many other industry servers and ensures maximum hypervisor uptime in virtualization environments.

Dell's Fault Resilient Memory provides a protected memory zone for a hypervisor without consuming excessive memory). The protected zone keeps the hypervisor from any harm other than the most extreme memory corruption errors. Combined with using redundant SD cards to host the hypervisor, this RAS feature also allows for highly available virtualization environments.

Also, in coordination with a running hypervisor, memory faults on Dell servers are monitored. If certain regions, such as a memory page, produce recoverable errors beyond a certain threshold, the faulty page will be retired and effectively removed from the system so that it is no longer accessed. The event is logged so that memory can be replaced during scheduled service time.

FOUR Proactive, Highly Responsive Systems Management

To ensure the uptime of a data center's operations, an enterprise must be able to efficiently manage across its entire infrastructure. It must be able to rapidly detect any issues that may impact operational availability or productivity and be able to quickly act to pre-empt or contain the problem.

Dell provides its comprehensive systems management portfolio of tools, OpenManage, to help customers fully realize the efficiencies of their IT infrastructures and maintain continuous uptime and reliability. It provides wide-ranging monitoring and notifications to let administrators proactively

detect and respond to problems with hardware and software components. To that end, it provides:

- Automated and replicable processes
- IT policy-based operations
- Accelerated workload deployment and migration

Management automation is a key component of optimizing any data center environment, and the new Integrated Dell Remote Access Controller 8 (iDRAC8) with Lifecycle Controller delivers automation across every dimension of server management.

To help customers manage possible component issues and maintain uptime of critical business applications, OpenManage provides a wide range of alerts via iDRAC and other elements of the OpenManage portfolio. The iDRAC in each server monitors the status of critical subsystems and notifies system administrators if or when components reach or exceed warning and critical thresholds. With this information, IT administrators can investigate and take corrective action before a component fails, ensuring higher server uptime and greater application availability.

"We were looking for the best technology, and we found it in Dell SAP HANA. The reliability and fast time to data access were unmatched by any competitor solutions we evaluated."

- Franklin Caraballo, Manager of IT, Grupo Estrella

In many cases, Dell systems will use embedded intelligence to take action automatically. Enterprises can keep workloads running — and not lose a moment — by opening a trouble ticket or dispatching parts based on automated alerts. All of this is part of Dell's comprehensive efforts to provide intelligent automation based on embedded, agent-free management technologies.

The ability to manage (and prevent) system-wide issues extends to the domain of power. OpenManage Power Center (OMPC) provides accurate, real-time power and thermal monitoring and management for individual servers, groups of servers, racks, and IT equipment like PDUs. OMPC can even support reading power data from other server manufacturers, including HP and Cisco.

OMPC can mitigate risks from power and cooling events by allowing administrators to predefine group policies for throttling performance and set power limits at a rack, room, or row level. It lets customers extend the uptime of business-critical applications in response to events like brownouts, rolling blackouts, and the failure of room-cooling equipment. Enterprises can use rack-, row-, and room-level views to easily locate any available space and power, thus utilizing the full capacity of existing rack and power infrastructure.

Another feature that can proactively prevent unnecessary downtime is Auto Update, which performs multiple, recurring firmware updates by specifying a network repository that contains a catalog of available updates and the scheduling parameters. At the

What Dell Means by RAS

To understand how the different RAS capabilities integrate to comprise a rock-solid foundation for enterprise computing, it helps to clarify what Dell means by RAS.

RELIABILITY encompasses two components: reducing the mean time between hardware failures and ensuring data integrity. Dell has many server design and manufacturing processes that ensure the highest levels of reliability at the component and integrated system levels. Data integrity is protected through error detection and correction — or, if not correctable, error containment. Both the Intel E7 processors and PowerEdge memory design use innovative technologies to correct or contain errors — guaranteeing optimum data integrity.

AVAILABILITY refers to uninterrupted system and application operation even in the presence of uncorrectable errors. Several elements can contribute to system availability: simple hardware redundancy, data availability technologies (i.e., RAID), and service availability delivered through virtualization/migration and failover clustering. Dell's layered "ecosystem" delivers availability at all these levels.

SERVICEABILITY means a system can be maintained without disrupting operation. This capability requires both thoughtful platform design and innovative systems management. Features like the pervasive hot-swappable front-access storage in PowerEdge servers enhance serviceability, and Dell's OpenManage system management simplifies component replacement in the case of hard failures and even offers predictive failure capabilities to prevent downtime.

New Intel Run Sure Technology

Four features you'll want to know about:

eMCA GEN 1 is designed to improve system diagnostic and predictive failure analysis, providing advanced error log data to system firmware and the OS.

MCA RECOVERY EXECUTION PATH is designed to increase system uptime, extending the software-assisted recovery to include uncorrectable data errors delivered to the core execution engine.

MCA IO is designed to provide improved diagnostics and reduce service costs, providing information on uncorrected I/O errors to the OS so that it can take appropriate actions.

PCIe LIVE ERROR RECOVERY is designed to increase system uptime, extending recovery and containment from PCIe errors.

scheduled time, all applicable updates contained in the repository are applied to the server.

Using the iDRAC GUI, you can schedule firmware updates to recur on a daily, weekly, or monthly basis. Updates can either be applied automatically or can be "staged" for the next manual reboot. You can integrate the update process with the Dell Update catalogs made available by the Dell Repository Manager.

FIVE Partners That Deliver RAS Across the Enterprise

When you are considering moving to a new architecture, you want to know that the complete environment is committed to the reliability that you require. Dell is integrated tightly with enterprise-savvy partners to deliver the highest levels of availability and serviceability across enterprise infrastructures — from the chip level up through the application layer.

A case in point, the Intel Xeon E7 processors used in the R930 leverage Intel's Run Sure Technology to maximize uptime. Run Sure includes more than 40 features designed to reduce the frequency and cost of downtime, and to protect data integrity by integrating processor, firmware, and software to help diagnose fatal errors, contain faults, and automatically recover to keep the server up and running.

These technologies fall into two major categories: 1) resilient memory technologies that ensure data integrity and enable systems to keep running reliably over a longer period of time, thus reducing the need for immediate service calls, and 2) resilient systems technologies that integrate processor, firmware, and software layers to allow the system to diagnose and/or recover from previously fatal errors. The Intel Xeon E7 product line supports all previous-generation RAS features and, as part of its Run Sure Technology, adds a number of new features. (See "New Intel Run Sure Technology," at left.)

SAP HANA is an excellent example of how the application layer benefits from RAS-oriented integration. In collaboration with Intel, SAP has successfully applied MCA recovery to SAP HANA to ensure maximum resiliency of in-memory data stores.

The very large main memory in Intel Xeon E7 processors enables SAP HANA to store and process multiterabyte transactional data stores completely in memory. If an unrecoverable memory error is encountered, the processor issues an MCA recovery signal to the OS. The OS in turn determines which software applications are using the offending memory.

In the past, the OS would, at best, have to halt any application that uses that memory or, at worst, stop all processing and halt the system. But now, with Intel's MCA recovery, it is possible to use the intelligence of each layer in the hardware-software stack to handle the error and quickly return to normal processing. The OS can notify the application, which can decide what course of action to take to repair the effects of the memory error.

If an unrecoverable memory failure occurs during the SAP HANA appliance operation, and the corrupt memory space is occupied by one of the SAP in-memory tables (see 1 in diagram on p. 6), the OS signals the SAP infrastructure software (2), which in turn responds by reloading associated tables (3). SAP HANA analyzes the failure and determines whether it affects other stored or committed data, in which case it uses snapshots (built during SAP HANA operation and kept in flash memory) to recover and reconstruct the committed data in a new working memory location.

Extending the functionality of MCA recovery has made this collaboration among silicon, the OS, and the application layer possible, ensuring a

smooth return to normal operation without disruption in continuity of service or loss of information.

Get the Reliability and Innovation You Need Now

Enterprises evaluating where to move now that their RISC-based architectures aren't meeting the demands of ever-changing computing environments should give serious thought to how Dell's future-ready, four-socket PowerEdge R930 server and complete RAS ecosystem approach to enterprise computing can bring both increased reliability and new, innovative technologies to their business.

ABOUT DELL

Learn more about Dell's PowerEdge R930 server and RAS ecosystem. Also find out more about:

Dell solutions for big data and analytics

Dell's SAP HANA solution

How Grupo Estrella is using Dell's SAP HANA appliance

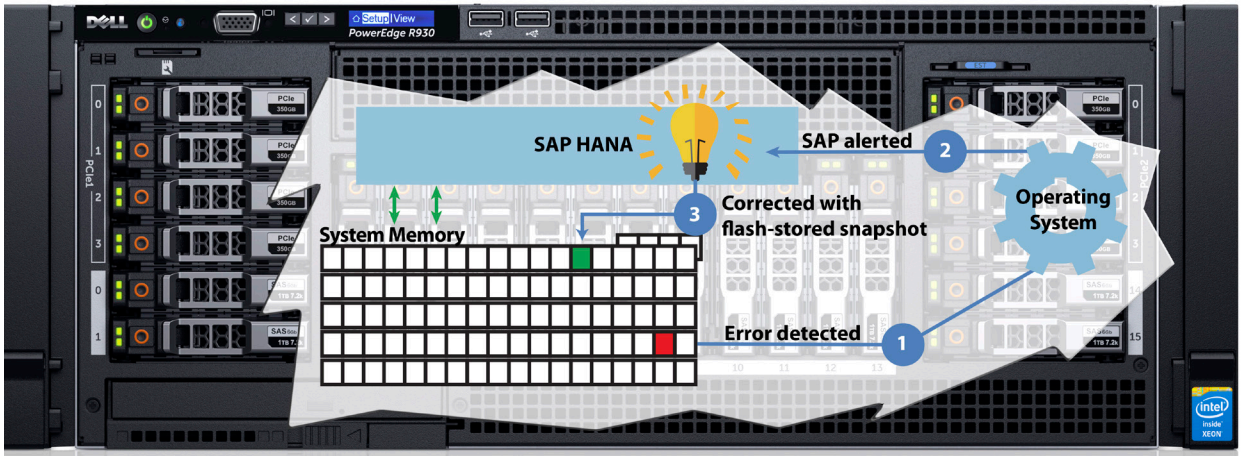

Tight integration of technology partners enhances system reliability.

If a memory failure occurs during SAP HANA operation, (1) the error is detected, (2) the OS signals the SAP infrastructure, (3) SAP responds by reloading the associated memory tables.