


COMPUTING PODS: LARGE-SCALE BUILDING BLOCKS FOR INTELLIGENT, AUTOMATED DATA CENTER DEPLOYMENTS



By Timothy Sherbak
Chris Auger

Using advanced virtualization and software tools together with latest-generation server, storage, and network infrastructure, IT organizations can build large-scale computing pods designed to aggregate server, storage, and network nodes into a simple, cooperating set of computing resources. This approach enhances enterprise efficiency by streamlining operations while facilitating fast, flexible response to increasingly complex service-level requirements.

IT infrastructures tend to become ever more complex as administrators support growth with incremental hardware expansion, which in turn calls for increased real estate, power, cooling, and maintenance. Faced with shrinking budgets, many organizations end up dedicating a majority of their IT spending to simply maintaining data center resources—placing serious constraints on their ability to advance new initiatives.

To reverse the tide of complexity caused by piecemeal resource expansion, IT executives are looking for ways to make strategic capital investments geared to reduce overall operating expenses. In particular, transitioning to a virtualized environment can advance the efficiency of the overall IT infrastructure by enabling self-managing data center resources to scale quickly and flexibly to meet specific organizational needs.

CREATING BUILDING BLOCKS FOR AN INTELLIGENT, AUTOMATED DATA CENTER

Many of the elements required to build a highly scalable, automated data center are now available. Strategic integration and coordination of key hardware elements such as standards-based servers, storage, switches, and processors are making these data center

advances possible. The required software elements are also available today, including mature operating systems, hypervisors, and management solutions.

Another key requirement of an automated data center is policy-based automation utilizing real-time intelligence. Feedback loops are needed to help recognize problems and optimize the core features of each layer of the infrastructure. Multipathing built into some of today's server and storage infrastructures is one example: if server and storage communication fails on one network path, the infrastructure has the intelligence to try different paths. Other examples include automatic load balancing of workloads across multiple physical servers and storage devices, and migration of workloads to help save energy by minimizing the number of actively running servers.

Many platforms are now optimized for virtualization. To advance data center flexibility, intelligence, and automation, IT managers must be able to integrate virtualization across multiple layers of the infrastructure to create virtualized server, storage, and network resource pools. A key requirement is that specific workloads not be bound to a specific physical resource.

Keeping workloads independent of specific hardware devices allows the flexibility and adaptability needed to migrate workloads, reallocate

Related Categories:

Blade servers

Dell EqualLogic storage

Dell PowerEdge blade servers

Storage

Virtualization

Visit DELL.COM/PowerSolutions for the complete category index.

resources, and expand resources on the fly. These capabilities help simplify proactive management scenarios to support real-world business processes. For example, IT managers may dynamically reallocate data center resources when end-of-month and financial reports are due, or quickly make resources available for a new project that has to be up and running in 30 days.

Such flexibility and adaptability allows organizations to build as they grow rather than over-provision infrastructure up front based on anticipated needs. An intelligent, self-managing data center that integrates virtualization across multiple layers of the infrastructure is also designed to avoid prolonged service outages because of infrastructure failures, or disruptive system upgrades and manual performance tuning in response to unforeseen performance and growth needs.

**THINKING ABOUT IT
INFRASTRUCTURE IN NEW WAYS**

With many elements of the intelligent, automated data center now available, IT executives are starting to think about the IT infrastructure in innovative ways. They are using existing data center components to build architectures around the concept of a *computing pod*. The pod can be defined as a self-contained building block having the following attributes:

- An optimized power and cooling footprint to facilitate scalability
- Integrated compute, storage, and network nodes of significant size and processing power
- Simplified data center deployment and management to help reduce operational costs

IT managers traditionally think of servers, storage area network (SAN) arrays, and switches as core building blocks. Now they can consolidate these elements into a single large-scale building block that becomes the basic unit of data center deployment (see Figure 1).

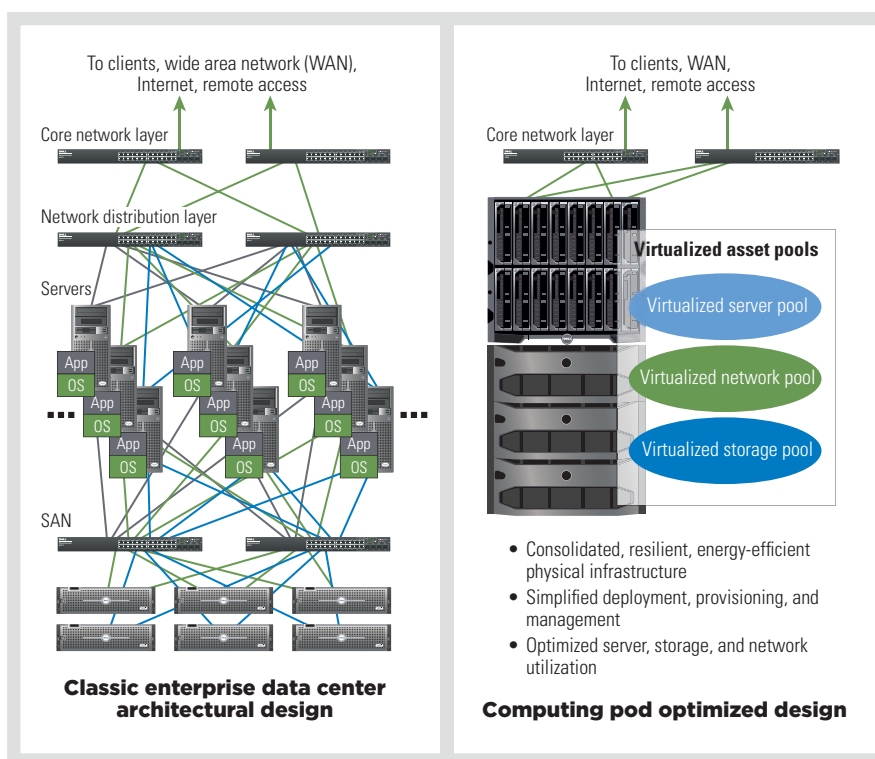


Figure 1. Comparing a traditional data center architecture with a computing pod that consolidates server, storage, and network resources

DESIGNING A COMPUTING POD ARCHITECTURE

To design a computing pod architecture, IT organizations can take advantage of advanced virtualization and software tools together with the latest generation of server, storage, and network infrastructure. Dell recommends that each pod include the Dell™ PowerEdge™ M1000e modular blade enclosure with PowerEdge blade servers, SAN arrays, hypervisor, management module, and switches with stacking capability.

Dell PowerEdge enclosures and blade servers

The Dell PowerEdge M1000e enclosure enables administrators to slide in blade servers and switches to expand or upgrade infrastructure as needed. Expansion is simple and internal to the chassis, allowing nondisruptive growth. An important element of the PowerEdge M1000e enclosure is the Dell FlexAddress feature, which effectively virtualizes the physical

addresses of ports within the infrastructure. Newly swapped-in components can automatically inherit the attributes of the prior components. This feature enables administrators to modify or change physical network infrastructure without an impact on the production environment.

Organizations using PowerEdge M1000e enclosures can also optimize pod design using the Dell M1000e Rack and Cable Advisor and Dell Virtualization Advisor online tools, available at DELL.COM/RackAdvisor and DELL.COM/Virtualization. The M1000e Rack Advisor guides administrators through the configuration process, including the use of power distribution units and enclosure stacking, while the Virtualization Advisor provides information to help size blade and storage infrastructure according to the aggregate workloads to be hosted. PowerEdge M610 and PowerEdge M710 blade servers are designed to reduce space and power requirements, which in turn can contribute

to reduced total cost of ownership and environmental impact.¹

Virtualized server and storage platforms

Virtualized server platforms are essential to the computing pod model, enabling administrators to create a flexible pool of server resources that can be allocated dynamically. If one physical server fails, virtual machines can automatically fail over to a different physical system to help ensure continuous availability. IT administrators can also quickly and easily start up virtual machines at a disaster recovery site. In addition, virtualization allows the consolidation of multiple servers onto a reduced number of physical platforms, enabling the computing pod to provide significant processing power without sprawling hardware.

Virtualizing storage extends the benefits of high availability, scalability, and optimization of resources. Internet SCSI (iSCSI)-based virtualized storage arrays such as Dell EqualLogic™ PS6000 series virtualized iSCSI SANs are recommended because they use the same Ethernet technology as other pod elements.² The EqualLogic SAN also contains built-in intelligence—such as automatic load balancing, simplified storage provisioning, thin provisioning, replication, and multipathing—that is designed to integrate with pod environments utilizing hypervisor technologies.

IT organizations have a variety of choices for mature virtualization software, including the VMware® vSphere™, Microsoft® Hyper-V™, and Citrix® XenServer™ platforms.

Network stacking and virtual blade switching

Enhanced scalability and simplified management are key objectives of computing pod design. Server stacking and virtual blade switching help deliver these capabilities. Stacking allows IT managers to incrementally grow network infrastructure

by adding bandwidth and I/O ports in a seamless way. Cisco switches that include virtual blade switch (VBS) technology, such as the Cisco Catalyst Blade Switch 3100 series switches, enable IT administrators to interconnect multiple switches so that they interact and behave as a single logical switch, managed together as a virtualized pool of network resources.

Stacking and virtual blade switching also enable the computing pod model by creating a self-contained network distribution layer.³ With its single virtual switch, the blade enclosure becomes the basic computing pod infrastructure unit. The network distribution layer resides entirely within the enclosure, helping avoid the need for traffic among the server and

DELL BUSINESS-READY CONFIGURATIONS: BUILDING BLOCKS FOR GROWTH

As data centers continue to grow in complexity, many traditional approaches to incremental expansion may not be able to keep up with rapid changes in business conditions. IT organizations can enhance business response by taking advantage of advanced software tools and infrastructure such as the following:

- Dell PowerEdge M1000e modular blade enclosure
- Dell PowerEdge M610 and PowerEdge M710 blade servers with the Intel® Xeon® processor 5500 series
- Dell EqualLogic PS6000 series virtualized Internet SCSI (iSCSI) storage area network (SAN)
- Cisco Catalyst Blade Switch 3100 series
- VMware vSphere, Microsoft Hyper-V, and Citrix XenServer virtualization platforms
- Dell Management Console

By capitalizing on large-scale building blocks that help simplify and automate IT operations, the computing pod model helps enterprises create innovative, self-contained data centers that can scale quickly, flexibly, and cost-effectively.

Dell Business-Ready Configurations provide a detailed reference architecture for deploying and implementing a computing pod architecture using Dell blade server and iSCSI storage environments. Based on extensive engineering work in architectural design and certification, the architecture is designed for rapid deployment into production environments, helping to eliminate much of the costly and time-consuming trial and error often encountered during complex infrastructure design and implementation. The solution is optimally configured to run virtualized workloads and is designed to provide redundancy without single points of failure, as well as scalability and manageability. It includes the network architectures, storage configurations, and best practices necessary for deploying and configuring the solution.*

*For more information, see "Business-Ready Configurations for Dell PowerEdge Blade Servers, Dell EqualLogic Storage, and VMware Infrastructure," by Dell Virtualization Solutions Engineering, May 2009; and "Networking Best Practices for VMware Infrastructure 3 on Dell PowerEdge Blade Servers," by Dell Virtualization Solutions Engineering, April 2009, both available at DELL.COM/Virtualization.

¹ For more information on PowerEdge M610 and PowerEdge M710 blade servers, see "Data Center Workhorses: New Dell PowerEdge Rack and Blade Servers," by Edward Yee, Indrani Paul, Robert Tung, Truc Nguyen, and Chad Fenner, in *Dell Power Solutions*, June 2009, DELL.COM/Downloads/Global/Power/ps2q09-20090246-Nguyen.pdf.

² For more information on Dell EqualLogic PS6000 series arrays, see "Boosting SAN Performance with Dell EqualLogic PS6000S Solid-State Drive Arrays," by Dylan Locsin, Ujjwal Rajbhandari, and Wendy Chen, in *Dell Power Solutions*, June 2009, DELL.COM/Downloads/Global/Power/ps2q09-20090236-Locsin.pdf.

³ Large enterprise data centers, which can potentially house thousands of physical servers, may use a three-tier network architecture. The servers may be connected together through an "access layer" network of switches, which connect to the distribution layer network (which, in turn, connects to the core network layer). In this case, the same conceptual argument for a pod architecture applies. In the pod architecture, server communication and data transfers between servers across the access layer network are consolidated on the pod's backplane-based network infrastructure, isolating it from other parts of the network.

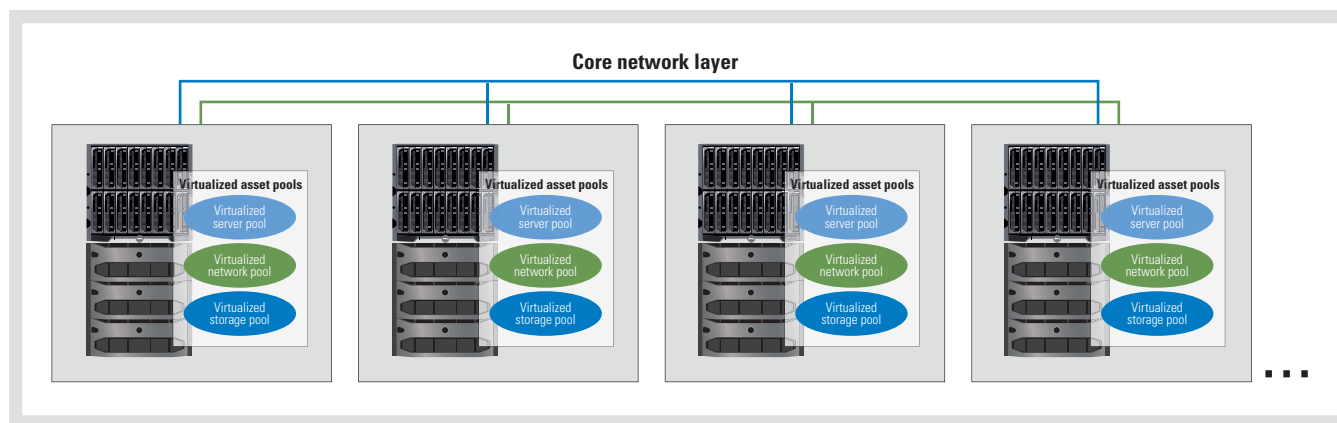


Figure 2. Scaling enterprise data center architectures with multiple computing pods

storage nodes to go out onto the data center core network layer. Servers, switches, and the management module are contained within the blade server enclosure environment (see the “Dell Business-Ready Configurations: Building blocks for growth” sidebar in this article).

ENHANCING DATA CENTER CAPABILITIES WITH COMPUTING PODS

Computing pods are designed to offer IT organizations significantly enhanced data center capabilities, including efficient aggregation and networking, increased space savings, unified maintenance and training, and simplified infrastructure design and deployment.

Redundant, high-bandwidth connectivity internal to blade enclosures within the pod helps speed communication and reduce congestion on the core data center network layer. For example, suppose an administrator needs to place a blade server into standby mode and move the virtualized servers to a different blade server in the same stack. None of that traffic actually leaves the backplane or affects the core data center network layer. Even large workloads can move among blade servers without disrupting productivity.

Likewise, access to storage is isolated from the data center’s core network infrastructure. The iSCSI SAN networking

infrastructure can be integrated directly into the blade enclosure’s Ethernet switch infrastructure, either as a virtual LAN or as an independent network. With Dell EqualLogic PS Series arrays, storage nodes are “stacked” as well—that is, multiple cooperating arrays are integrated as a logical pool of storage to help optimize I/O and storage utilization. In this way, administrators can use each pod to aggregate dozens or even hundreds of server and storage workloads in a minimum amount of space.

Data protection procedures may also be consolidated and centralized within the computing pod. Administrators begin by designating one or two blade servers per pod as media servers for backup and recovery with direct connectivity to tape and/or disk resources dedicated to backups. Data backup traffic is limited to the pod’s computing, storage, and network infrastructure without affecting external data center resources.

Another blade server within each pod can be designated as a configuration server for pod management. Using the Dell Management Console,⁴ hypervisor management tools, and other management applications, this designation enables administrators to manage the production servers in the pod as a unit, from configuration and reimaging to patch management and antivirus updates. As a result, IT staff

members need not spend time managing individual servers, which helps reduce IT support costs substantially.

Organizations can maintain a single logical entity—the pod—instead of performing time-consuming, one-to-one server management, storage management, and network management. Managing the complete computing pod environment as a unit not only helps optimize servers and switches, but also helps simplify backup and patch management. IT staff can be efficiently trained to manage the standard pod and its components, using logically consolidated tools and procedures rather than multiple disparate systems.

Computing pods also help simplify infrastructure design and deployment. Designing a data center using the computing pod as a basic building block helps simplify the design target. Organizations using pods are able to reduce the time and cost requirements of data center planning, deployment, and implementation.

DEPLOYING COMPUTING PODS FOR FLEXIBLE, SCALABLE DATA CENTERS

The building block approach to computing pods enables outstanding flexibility and scalability. Blade server, storage, and network nodes within the pod are designed to be easily scalable, and the

⁴For more information on the Dell Management Console, see “Systems Management Simplified,” by Pascal Nicolas, Jeanne Feldkamp, and Tom Kolnowski, in *Dell Power Solutions*, March 2009, DELL.COM/Downloads/Global/Power/ps1q09-20090191-CoverStory.pdf.

pod itself becomes a building block for scaling out the data center.

Organizations of different sizes have distinct opportunities to benefit from the computing pod model. For example, midsize organizations can start with a single pod as a convenient, all-inclusive computing infrastructure in a single rack enclosure and add nodes to the pod or additional pods as requirements grow. In large enterprise environments, IT managers may deploy numerous pods to help meet comprehensive data center infrastructure requirements, assigning different pods for different usage scenarios. For example, pods may be designated for production, development, and Web presence—all sharing a common environment (see Figure 2).


The computing pod model also facilitates virtual desktop infrastructures by allowing IT administrators to easily add a pod to virtualize and aggregate desktops. End users can remotely access the computing pod server infrastructure and run their personal computer environments there.

In addition, a growing number of organizations are utilizing computing clouds to deliver flexible access to compute and storage resources. The cloud computing approach enables organizations to leverage the scalability and manageability of pods to build large data center infrastructures with a pool of flexible resources. The pod design helps minimize acquisition and operating costs,

maximize energy efficiency, and enable rapid scalability.

INTEGRATING INFRASTRUCTURE ELEMENTS INTO EFFICIENT, AUTOMATED DATA CENTERS

The computing pod is an exciting model for IT infrastructure that helps significantly advance enterprise efficiency. By elevating data center intelligence through advanced virtualization and software tools in combination with latest-generation server, storage, and network infrastructure, organizations can create highly scalable, flexible, and self-managing data centers. By consolidating data center functionality into a basic unit that also serves as a building block for large-scale data center deployments, computing pods help reduce IT infrastructure complexity, simplify management, and increase cost-efficiency.

Moreover, computing pods are designed to provide a comprehensive, integrated management infrastructure that can be customized to specific organizational requirements. They help reduce and consolidate network traffic by isolating it to the specific resources requiring that traffic. And the large-scale building blocks provide for a flexible computing infrastructure that can quickly and easily adapt to evolving needs and requirements. 

Timothy Sherbak is senior manager of virtualization solutions marketing in the

Dell Large Enterprise Business Unit, and has over 20 years of experience in computer systems and storage networking technologies. He has a B.S. in Systems and Information Science from Syracuse University and an M.S. from the Massachusetts Institute of Technology's Sloan School of Management.

Chris Auger is manager of enterprise technology in the Dell Public Business Unit, leading a team of enterprise technologists helping government, education, and health care organizations architect world-class data center solutions. He has a B.S. in Mathematics from the University of Texas at Austin.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell Efficient Enterprise Portfolio:
DELL.COM/Switch

Dell M1000e Rack and Cable Advisor:
DELL.COM/RackAdvisor

Dell Virtualization Advisor:
DELL.COM/Virtualization

Dell storage:
DELL.COM/Storage