

Building the Virtual Data Center

The advent of fine-grained modular computing and switch-fabric communication architectures is enabling the next evolution in enterprise server computing—the virtual data center. This article describes how virtual data centers differ from traditional data centers and focuses on the many benefits of this emerging model.

BY J. CRAIG LOWERY, PH.D.

The virtual data center is a truly evolutionary computing model that hides all hardware and software components behind an abstraction layer, presenting the computing environment as a collection of managed services rather than as a collection of machines. Previously, creating such an abstraction layer was prohibitively expensive if not impossible to implement; partial solutions built on proprietary technologies have only recently appeared. However, the ongoing cycle of refining both technical components and business models has recently yielded standard components with which to build the virtual data center:

- **Low-cost commodity hardware:** Server manufacturers—most notably Dell through its direct business model—have driven down the cost of server hardware by supporting industry standards, avoiding the built-in costs of proprietary systems, and selling products in mass quantities directly to consumers.
- **High-density components:** Mobile computing platforms and the need to reduce space and power footprints in the data center have resulted in smaller, more efficient devices.
- **Remote storage:** Storage area networks (SANs) and network attached storage (NAS) allow for the dissociation and dynamic reassociation of processors and persistent (disk) storage.
- **Switch-fabric interconnects:** High-throughput, low-latency, and low-contention communication architectures allow hundreds to thousands of networked components to communicate at speeds comparable to those obtained through the traditional motherboard bus.
- **Communication protocol standards:** The IP suite is now driving standardization in data center interconnects, just as it has in wide area networking. Initiatives such as Internet SCSI (iSCSI) and remote direct memory access (RDMA) are examples of how the IP suite is contributing to the development of virtual data centers.

The motivation for using these components to build the next-generation data center lies in the need to solve the same problems that data center managers have faced for decades:

- **Cost:** Reducing total cost of ownership (TCO) of computing resources has become a dominating goal, targeting all life-cycle phases and sources of inefficiency—hardware, software, administration, management, resource utilization, and reliability.
- **Capacity:** Providing sufficient computing resources at exactly the time they are needed is a difficult task, especially when it must be balanced with cost containment.

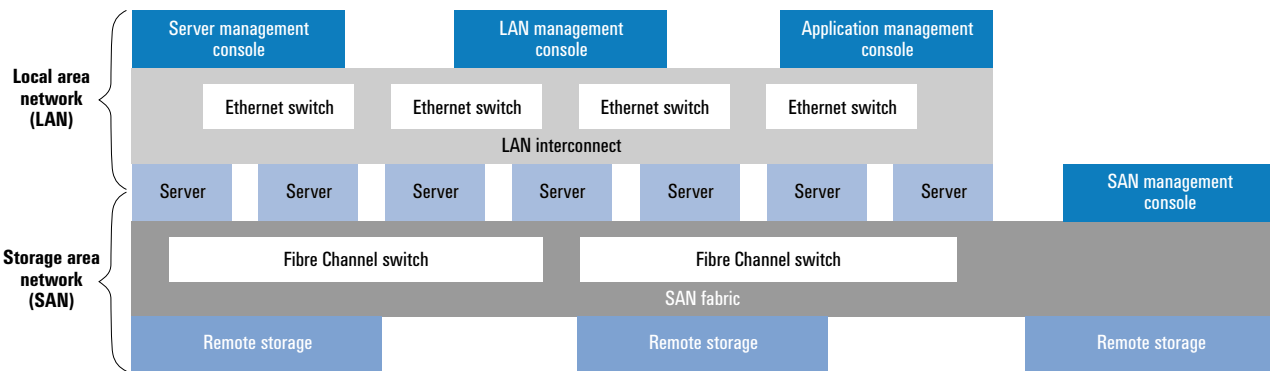


Figure 1. A contemporary data center architecture

- **Security:** Designing systems that achieve high availability, ensure data integrity, and maintain privacy is essential to guaranteeing business continuance.

Examining the contemporary data center

Figure 1 shows the architecture most prevalent in contemporary data centers. It comprises two independent networks: the LAN and the SAN. The LAN is a general-purpose network for server-to-server and server-to-Internet communication, typically employing the IP suite atop high-speed Ethernet (Fast Ethernet or Gigabit Ethernet¹). The SAN is a special-purpose network for server-to-storage communication, typically implemented using Fibre Channel technology because of its high-performance characteristics.

Servers, which are connected to both networks, use the SAN to access remote storage and the LAN for all other communication. This architecture offers considerable benefits, derived chiefly from the dissociation of storage from the physical server. If applications store state information in virtual (remote) storage, then a spare server can replace a failed server simply by connecting to the SAN. Similarly, applications can access spare processing capacity if idle servers are added, and the virtual storage is used to coordinate the work of the new servers. This improved availability and manageability makes scaling out the data center possible, and scalability is necessary to support growing client demands.

Although the LAN/SAN architecture provides increased flexibility and functionality, this approach also has drawbacks. The fact that it consists of two networks—and is thus a disjointed communication infrastructure—implies considerable support overhead. Because the LAN and SAN differ in their basic technologies and usage, most data centers have necessarily evolved two separate technical cultures for their respective support. This schism also manifests in the software abstractions required to manage the

data center. Separate administrative consoles exist for server and application functions and for storage services. Although the LAN/SAN segregation does not cause this fragmentation of management interfaces, it does not mitigate the fragmentation either. Finally, maturity of SAN standards significantly lags behind that of the more general-purpose LAN, diminishing return on investment (ROI).

The LAN/SAN design is state of the art and a vast improvement over previous architectures, but it is an intermediate step in data center evolution. The virtual data center is the next step—it will extend virtualization to all resources, not just storage. The emergence of new interconnect technologies will provide a single, standards-based infrastructure for both general-purpose and high-performance networking requirements.

Introducing the virtual data center

The virtual data center, as shown in Figure 2, comprises two kinds of components: processor components (*compute nodes*) and storage components (*disk nodes*). Processor components contain

The ongoing cycle of refining both technical components and business models has recently yielded standard components with which to build the virtual data center.

general-purpose processors and RAM with optional local disk storage; storage components contain disks or other forms of persistent auxiliary storage with enough local processing power to enable the components' communication and management.

A high-performance switch fabric then interconnects all the components. The fabric is bound to the components (by using a software layer in the form of agents on each node) and to the intervening switches. This binding

¹Gigabit Ethernet indicates compliance with IEEE® 802.3ab and does not connote speeds of 1 Gbps.

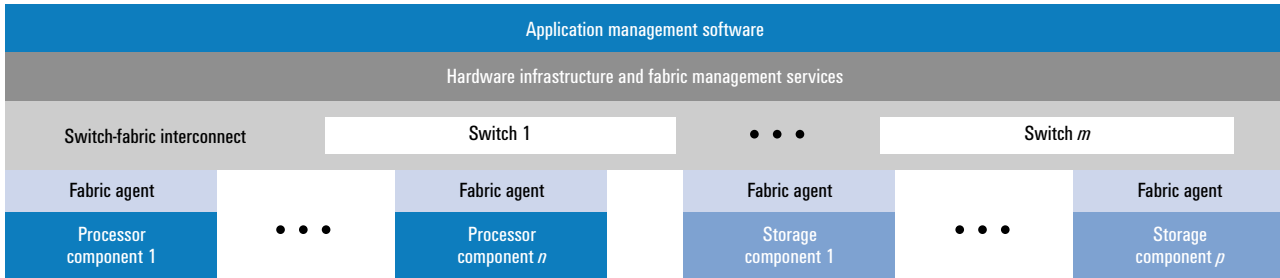


Figure 2. A virtual data center architecture

allows individual components to become members in one or more communicating, cooperating logical groups.

Another software layer monitors the health of the hardware components and manages the fabric to create these dynamic associations in real time and on demand. These services are made available to top-level software that manages applications and—more indirectly through automation—manages the underlying switch fabric and hardware components. This software can be staged on a central entity similar to a master node of a cluster, or more sophisticatedly, it can be staged as distributed agents working in concert.

This architecture offers several capabilities that transcend or fully realize those of previous generations:

- **Dynamic, software-based reconfiguration:** The switch fabric and fabric management layer cooperate to enable soft reconfiguration of the hardware components' communication affinities. At any time, administrators can use software commands to place any subset of processor components and storage components in a communicating group, rather than reconfiguring physical cables to common switches.
- **Interchangeable components:** Although components in a particular class (processor or storage) are not strictly homogeneous, they are similar enough to be interchangeable in most application environments. Through dynamic reconfiguration, a virtual data center can scale out applications across additional components. Similarly, administrators can quickly replace failed components through reconfiguration, ensuring that service level agreements (SLAs) are met.
- **High-performance communication:** Fabric interconnects, such as the one described in this article, differ from traditional LANs and wide area networks (WANs) in that they

The virtual data center architecture offers several capabilities that transcend or fully realize those of previous generations.

are designed to deliver high-speed virtual point-to-point connections between any two components with little or no switch hardware contention. Not only can such an interconnect support traditional networking protocols, but it also makes possible previously impractical or unwieldy mechanisms such as RDMA.

- **Application-centric view:** From the management perspective, the virtual data center is a collection of services rather than a collection of hardware components (see Figure 3). The focus is on the software stack—generically referred to as an *image*—which resides on a storage component. Administrators can replicate images across multiple processor components by several methods, ranging from simple bit copies in purely homogeneous environments to fully scripted remote installations. The traditional approach is to replicate an image from shared storage to the local storage of a processor component. A more flexible method, however, is to replicate instances of the image on remote storage, which processor components then access directly across the fabric.

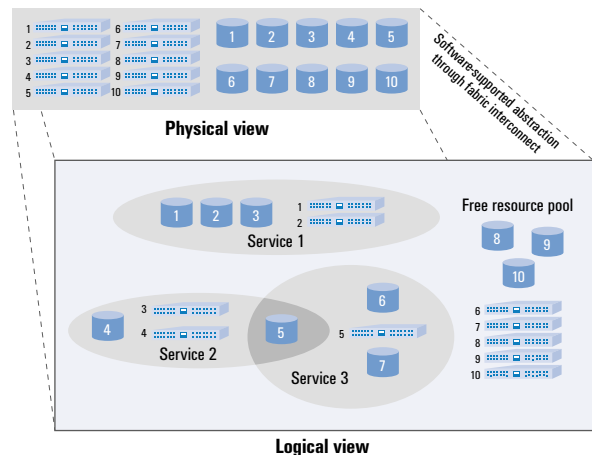


Figure 3. Service-layer abstraction of fabric-based hardware

With the virtual data center in place, enterprises can choose from several computing models that address each of the three classic IT problems: cost, capacity, and security.

Optimizing cost, capacity, and security

Although the prices of commodity components have dropped, managing these components is still costly. Current-generation data centers require a large workforce of expensive system administrators to perform the same tasks repeatedly for multiple servers. The key to reducing this cost is achieving a few-to-many ratio of administrators to components for all aspects of data center management.

Much of the work performed by system administrators is installation and configuration of servers for a particular application. Administrators accomplish this work by repeating a proven sequence of steps at each server. In the virtual data center, a single expert creates a reference software stack that implements the application; therefore, the installation and configuration steps are performed only once. An image of the stack is placed in common storage for replication across computing components. Subsequent upgrades require updates to the reference image.

The top-level management software of a virtual data center allows administrators to execute the same operation—such as image deployment and capture, health monitoring, or mass command invocations—across components. This layer also provides both prepackaged (best practices) and customized (site-specific) logic to automatically respond to asynchronous events, such as Simple Network Management Protocol (SNMP) traps received from components.

Another benefit of the virtual data center's top layer is the presentation of a unified data center. Current data centers are fragmented in this regard, exporting management and user interfaces at each layer of abstraction. The unified model of the virtual data center presents a single, rich interface at the top level. Model view controllers can be designed to present different user-customized views of the unified model. For example, the database administrator, a clerk in the business office, and a self-service Internet access subscriber would each have a different view of the same system.

The dynamic reconfiguration of components afforded by the fabric interconnect enables the top-level management software to allocate resources to application groups on demand. For example, a pool of idle components can be held in reserve until increased demand or the failure of an active component triggers its automatic deployment into an application group. If processor

With the virtual data center in place, enterprises can choose from several computing models that address each of the three classic IT problems: cost, capacity, and security.

components execute an image directly from shared storage, the deployment is immediate. This capability minimizes underutilization of resources while delivering high availability and capacity on demand. Finally, the virtual data center retains many of the security benefits gained through segregating applications on separate physical hardware.

Virtual data centers: Becoming a reality

The virtual data center would appear to solve many, if not most, of the problems faced by enterprise data centers. It reduces costs by dynamically allocating several small commodity components in real time and by simplifying the management of these components through technology consolidation, all without introducing new risks. The virtual data center focuses on application management—a higher level of abstraction of the data center than the current generation's hardware-centric view. It automates several labor-intensive tasks and allows administrators to replicate software, issue commands, and monitor health across thousands of devices from a single interface.

A substantial number of software vendors have entered the market to deliver products that, among them, offer all these features. High-density server systems, such as the Dell™ PowerEdge™ 1655MC server blade system, are the first in a generation of products that may eventually provide the standard, low-cost hardware upon which to build the virtual data center.

Although not yet fully realized, the virtual data center is becoming a reality. Even though it is an unconventional approach, the virtual data center will likely succeed because it solves problems at the level of abstraction that has the most relevance, impact, and appeal, while delivering secure, sufficient computing services with the lowest possible cost. ☞

J. Craig Lowery, Ph.D. (craig_lowery@dell.com) is a software architect and strategist in the Software Development Group of the Dell Product Group. Craig has an M.S. and a Ph.D. in Computer Science from Vanderbilt University and a B.S. in Computing Science and Mathematics from Mississippi College. His primary areas of interest include computer networking, security, and performance modeling.