A Data Center Fabric Is Critical to a Next-Generation Data Center

January 2014

Prepared by:

Zeus Kerravala

A Data Center Fabric Is Critical to a Next-Generation Data Center

by Zeus Kerravala

January 2014

ZK Research
A Division of Kerravala Consulting

Introduction: The Changing Data Center

The data center has gone through many major evolutionary changes over the past several decades, and each change has been defined by major shifts in architectures. The industry moved from the mainframe era to client/server computing and then to Internet computing. In 2011, another major shift began: the shift to a virtual data center. This has been the primary driver in enabling customers to transition to the cloud and ultimately IT as a service. The shift to a virtual data center will be the single biggest transition in the history of computing. It will reshape all the major data center tiers: applications, storage, servers and the network. The drivers of this transition are as follows:

- **Evolving applications:** As the workforce has become increasingly more distributed, the nature of applications has changed. Legacy client/server applications were designed for workers to use on a corporate device that was attached to the private company network. Today's workers need the ability to access any application from any location, meaning the applications must evolve to be a more dynamic resource compared to static legacy applications.
- **IT resource pooling:** To date, most of the focus of virtualization in the data center has been on servers. Over time, however, all the resources in the data center will become fully virtualized and "pooled" (see Exhibit 1). This pooling will bring to fruition the vision of on-demand computing, in which any application or service can access as much of any resource as it needs, governed by business policies. For cost/benefit reasons and to promote optimum utilization of existing resources, the ideal scenario is to create large, dynamically shared pools. These pools are driving the demand for a flexible data center architecture.
- **Changing IT service models:** Historically, IT delivered resources in tightly defined silos to optimize application performance and to simplify management. However, deploying infrastructure in discrete silos is highly inefficient and one of the reasons server utilization historically has been under 20% and storage utilization has been approximately 25%. As IT budgets were slashed and CIOs were tasked with doing more with less, new service models arose. Today, companies are using cloud computing to deliver IT more efficiently and provision resources faster than ever before. Over the next several years, more and more companies will build private or hybrid clouds to change the way IT is delivered to the business.

Every trend toward a fully virtualized data center is highly dependent on the data center moving away from a number of vertically integrated silos. Next-generation data centers will see a greater coupling of hardware and software to bring together both physical and virtual infrastructure. The result will be a fluid, dynamic "fabric" capable of moving IT resources where the business needs them.

zeus@zkresearch.com

Cell: 301-775-7447
Office: 978-252-5314

Influence and insight through social media

Exhibit 1: The Role of the Data Center Fabric

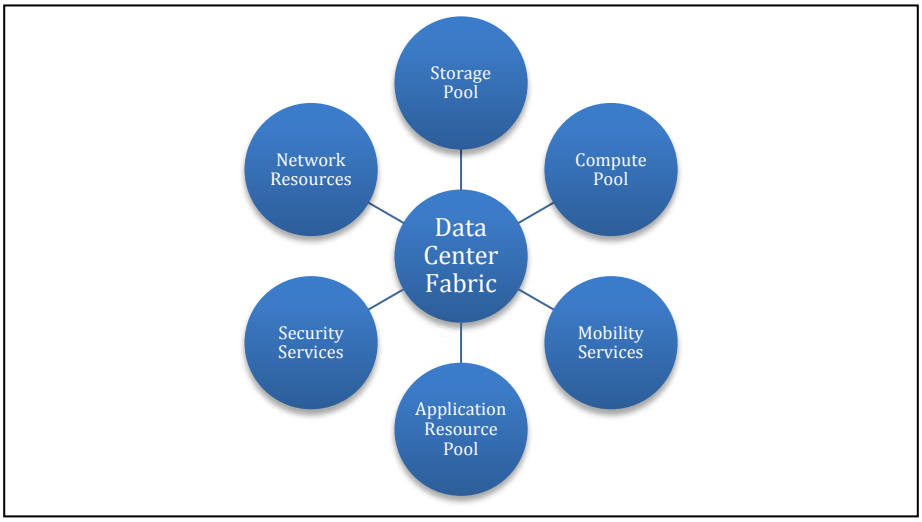

Source: ZK Research, 2014

Section II: The Challenges of Legacy Data Centers

Legacy data centers were optimized for performance, not cost. Each application typically has its own dedicated network, storage and compute infrastructure. This means that each application has finite resources; it cannot use idle resources from any other application to improve performance. Over time, the siloed nature of the legacy data center has led to the following inefficiencies:

- **Lack of agility to meet business demands:** Poor resource utilization combined with IT silos results in long lead times in deploying the infrastructure required to deliver new applications and services. Typically, application owners must contact each area of operations independently for the needed resources, resulting in typical lead times of six to eight weeks for resourcing each part of a new application. In a worst-case scenario, an application may not be ready until about six months after the business asks for it.
- **Poor resource utilization:** Deploying applications in silos means IT managers must provision all resources for peak utilization and then add a little more for overhead on the peak days, resulting in highly underutilized data center infrastructure. Based on ongoing research, ZK Research estimates server utilization in a legacy data center environment is about 20%, while storage utilization is approximately 25% and network utilization is about 30%. Considering data center infrastructure is among the most expensive IT resource, this kind of waste is very costly.
- **Organizational structure not optimized for delivering IT as a service:** IT infrastructure's vertically integrated nature has driven IT operations to be highly siloed as well. Each area of operations tends to work independently from the others. The resulting lack of cross-silo awareness can create situations where one area of operations makes a change, and the other areas must scramble to adapt to it. This is one reason organizations spend an estimated 83% of their IT budget on simply maintaining the status quo.
- **High total cost of ownership (TCO):** Poor resource utilization combined with a lack of agility and dedicated operational teams results in a lack of efficiency that drives the cost of running a data center to incredibly high levels. Almost every CIO's current mandate is to find a more cost-effective way of running the data center.
- **Difficulty meeting current data center demands:** Traditional networks struggle to meet the demands of varying applications (bare metal or virtual) and increasing traffic requirements in terms of performance, scale and capacity.

Clearly, this expensive and inflexible legacy way of running a data center will not enable an organization to leverage the trends of virtualization and cloud

computing. Instead, organizations should look to migrate to a data center.

Section III: Characteristics of a Data Center Fabric

"Data center fabric" is a widely used industry term referring to a network-centric system that links storage and compute functions through high-bandwidth interconnects to function as a single logical infrastructure unit. Data center fabrics can be easily built, dissolved and rebuilt as necessary to provide data center resources where and when they are needed. They are the result of several major trends coming together to deliver a unified data center, offering organizations the ability to manage and orchestrate all the elements for physical, virtual or cloud-based deployments seamlessly. Data center fabrics help dissolve both IT and organizational silos to increase efficiency, simplicity and business agility.

The data center fabric does more than just cobble together infrastructure components. The fabric actually operates as a single system to deliver data center resources to any application or service as business policy dictates. The characteristics of a data center fabric are as follows:

- **Open:** No solution provider can deliver an entire data center. Being open is important to enable third-party integration across solution providers.
- **Integrated and simplified:** Traditional IT silos need to be integrated into a single fabric to ensure that the storage, compute and network resources can be optimized and maintained as the environment changes.
- **Flexible and dynamic:** The application and virtualization environments supported by the data center are not static, and the fabric can't be static either. A data center fabric has holistic control over data center resources and can more easily orchestrate the configuration changes required as the application environment evolves, in order to support automation and service orchestration.
- **Scalable:** It's difficult to predict the growth of compute, storage and network requirements. IT managers must be able to scale up the data center fabric in real time to ensure there is enough capacity to deliver a high-quality experience to both workers and customers.
- **Resilient:** Today's business environments require non-stop operations. Downtime means lost customers, lost opportunity and lost revenue. The data center fabric needs to be built with resiliency and self-healing in mind to ensure maximum uptime with minimum disruptions.
- **Secure:** The data center delivers a wide range of information to applications and services, including data that is highly confidential and critical to the organization. Security is an area where data center architects simply cannot compromise. The data center fabric provides the highest level of integrated security to not only the data center but also end-user devices including mobile endpoints.
- **Application centric:** The data center fabric provides control APIs for network and compute resources, enabling applications to directly interface with the infrastructure that supports it. Through the APIs, the infrastructure can be provisioned or reconfigured on the fly according to business policy.

Evolving the data center to a fabric might seem like a nearly impossible task for an organization, but it's necessary in order to leverage the trends of virtualization and cloud computing, and truly fulfill the vision of next-generation computing.

Section IV: Benefits of a Data Center Fabric

A network-centric data center fabric offers organizations many benefits. The most obvious is the consolidation and simplification of network, compute, storage and application resources. However, the value proposition is much broader. Other benefits of a data center fabric include:

- **Faster time to deliver services:** Legacy infrastructure often has long deployment times, meaning it can take weeks or even months to deploy a new service. The dynamic nature of a data center fabric combined with centralized management can shorten the deployment time of a new service from weeks to days or even minutes in some cases.
- **Better utilization of data center resources:** A data center fabric virtualizes physical resources into logical pools. Instead of each application having its own set of IT resources, applications fetch needed resources from the logical pools. If more resources are needed, more are added to the pool. This optimizes IT resource utilization.
- **The elimination of IT silos:** The limitations of the current data center architecture drove IT into its current silos. A data center fabric allows IT organizations to combine all the different operational silos into one common management

area. This eliminates many redundant tasks and streamlines data center operations.

- **More productive IT staff and corporate workers:** A data center fabric improves efficiency and creates more manageable environments. This means less downtime for IT and less ad-hoc troubleshooting, leading to a more productive IT staff that can focus on strategic initiatives. Corporate workers are also more productive because applications are more available and perform better.
- **Lower overall TCO:** Improved utilization means less hardware to buy, maintain, power and cool. Organizations that deploy a data center fabric will see a reduction in both operational expenses and capital costs, significantly reducing TCO.

Section V: Evaluating a Data Center Fabric

Moving to a data center fabric will be critical for almost every organization over the next five years. However, organizations might not understand how to evaluate solution providers because the term "fabric" is relatively new in data center operations. When considering a move to a virtualized data center, key decision criteria for IT leaders include:

- **A holistic data center fabric:** Choose a fabric solution that encompasses application, storage, servers and network infrastructure. Many vendors offer solutions in only one of these areas, but a data center fabric must address all of the changing needs of a virtualized data center. Ideally, the fabric should provide a converged storage, server and network solution.
- **Scalability:** Data centers are not static entities; fabrics can't be static either. Organizations need the ability to start small, migrate at a comfortable pace and then expand the fabric as the data center grows. Additionally, it's important that the solution scale between data centers to enable data center fabrics in separate data centers to be managed as a single entity.
- **Location-independent virtual resources:** Location-independent virtual resources provide the ability to deliver a consistent fabric experience, independent of whether the workload is in the same blade, within a data center or across data centers. Location independence improves visibility, manageability and ease of use.
- **Integrates physical and virtual infrastructure:** Much of data center innovation has focused on virtualizing infrastructure. However, physical infrastructure still plays a key role in how applications perform. Unfortunately, many data center solutions are confined to just virtual or physical infrastructure. The fabric needs to integrate both physical and virtual infrastructure to ensure optimal resource utilization and application performance.
- **Cloud ready:** Within five years, more than two-thirds of all workloads will reside in the cloud and will account for well over 50% of data center traffic. An agile, efficient network with features such as multi-tenancy and high availability is required to prepare the enterprise for the era of cloud.

Section VI: Data Center Fabric Case Studies

The Organization: Apollo Group

Based in Phoenix and founded in 1973, Apollo Group is a provider of higher education programs for adults currently in the workplace. The University of Phoenix, Apollo's marquee university, has over 400,000 students and 24,000 faculty members across its more than 200 campus locations and learning centers.

The Problem

Apollo Group's applications reside on 7,500 physical servers located in three interconnected data centers. Its data centers were designed in very typical legacy fashion: Each application had its own dedicated server, storage and network resources. Additionally, its network supported 1 Gbps speeds, which is common for legacy data centers. Although this speed was sufficient five years ago, it's not workable today. The current architecture led to the following problems:

- **Poor user experience:** Apollo Group is seeing a significant increase in network traffic from applications such as videoconferencing, Web-sharing tools and social media. Students tend to expect any application to work anywhere on campus, whenever needed—but this was not the case. During peak traffic times, the network faltered, creating a sub-par experience for students. The organization became concerned that if students' experience did not improve, they would choose another online educator.
- **Long provisioning time for new applications:** Because each application had its own dedicated hardware, adding new applications required procuring new storage, servers and network resources. This created provisioning times that

were less than ideal. For example, if a business unit requested infrastructure for a new application, it would take two to four weeks to purchase and build the new servers alone.

- **Inconsistency in OSs and patches across servers:** Apollo Group's infrastructure experienced a significant amount of "drift"—when applications and OSs are not consistent from server to server. This happens when a server or OS is patched or modified to react to an urgent condition, such as a security warning or business change. Too much drift can lead to intermittent outages and long troubleshooting times.
- **Lack of scalability:** Because each application or service was deployed in a dedicated silo, scaling the infrastructure meant having to add staff. Typically, doubling the IT infrastructure would require a 20% to 30% increase in IT staff.

The Solution

Apollo built its next-generation data center network on a foundation of Cisco Nexus switches. Apollo used a modular design to build out computing "blocks" that could be partitioned and provisioned as the business required. Each block contains a pair of Cisco Nexus 7010 switches at the core. The core switches are connected to Cisco Nexus 5010 or 5020 switches at the end of each row. The end-of-row switches are connected to Cisco Nexus 2248 Fabric Extenders deployed at the top of each rack. Each of the interswitch links uses dual 10 Gig-E links, creating the high-speed, redundant network foundation needed to support the organization.

Apollo uses the Cisco data center fabric to create virtual data center "pods." A data center pod is a self-contained unit made up of physical/virtual storage, server and network resources that can be provisioned on an on-demand basis from a centralized management console. When a business unit requests infrastructure for a new application, a "pod" is created and assigned to the business unit. The allocation and provisioning of the storage, server and network resources is fully automated, removing the risk of human error that existed in the old, manual process. The virtualized compute resources deployed on the high-speed network foundation create a data center fabric. This fabric addresses all the technical challenges that could potentially keep the IT infrastructure from meeting the demands of the universities.

The Benefits

Apollo Group realized many benefits from its move to a data center fabric:

- **Greater agility to support business growth:** The data center fabric's modular design allows the IT department to respond to business units much faster. Prior to using the fabric, the IT department could only provision one server at a time and it would take two to three weeks to complete the task. Now the infrastructure can be provisioned on the fly, giving IT the ability to create multiple servers simultaneously in minutes.
- **Lower infrastructure costs:** The Cisco Nexus 7000 Series switches have a port density that is 10 times greater than other comparable options. Combined with a flatter fabric design, this ultimately means fewer switches to buy and less infrastructure to procure, manage, power and cool.
- **Reduced operational costs:** Because the Nexus 2248 switches are extensions of Cisco's 5000 Series switches, the IT department does not have to manage each 2248 switch individually. All management can be done on the 5000 Series switch, providing a single point of management. Apollo can now manage twice the number of devices with the same size IT staff. Eventually, Apollo will actively monitor and automate more tasks with Cisco Data Center Manager software to drive operational costs down further.
- **Improved user experience and application availability:** Legacy networks utilize Spanning Tree Protocol (STP) to prevent routing loops and broadcast storms. STP works by having only one unique path for traffic to take at any given time, with the rest of the paths being disabled. When the primary path becomes unavailable, STP recalculates the best path to make active. In a large network, such as Apollo's, this process would create periodic downtime. The data center fabric removed the dependency on STP, increasing the overall uptime of the network. Additionally, the high-speed 10 Gig-E backbone ensures peak performance even during peak usage periods.

"The data center fabric built on a Cisco Nexus foundation allows our IT department to respond to the business units at Apollo Group in near real time," said Dennis Crowe, director of network engineering at Apollo. "The data center fabric has given us better uptime with an easier to manage infrastructure."

The Organization: NTT DATA Corporation

NTT DATA Corporation is a global systems integrator with a successful track record in developing administrative systems as well as

building and operating shared data centers. NTT DATA owns more than 100 systems itself and is aggressive with its own IT strategies. The company has bases in 145 cities and 34 countries, and has a vision to be a "Global IT Innovator."

The Problem

NTT DATA's internal IT systems last had a major upgrade and redesign in 2006. The systems were aging and the company faced the problems of expiring hardware and software maintenance contracts. Additionally, the age of the systems had caused the cost of maintenance and management to skyrocket.

Historically, the company had optimized and standardized its platforms in an ad-hoc manner, making it very difficult to manage the systems. In an effort to lower the overall total cost of running IT and make more effective investments, NTT DATA had a goal of optimizing the infrastructure by means of integrated virtualization. The ultimate goal would be to have a standardized, shared, company-wide platform on which corporate applications could be run (Exhibit 2). NTT DATA would have its own private cloud where IT resources could be shared.

Exhibit 2: NTT DATA's Vision of Private Cloud

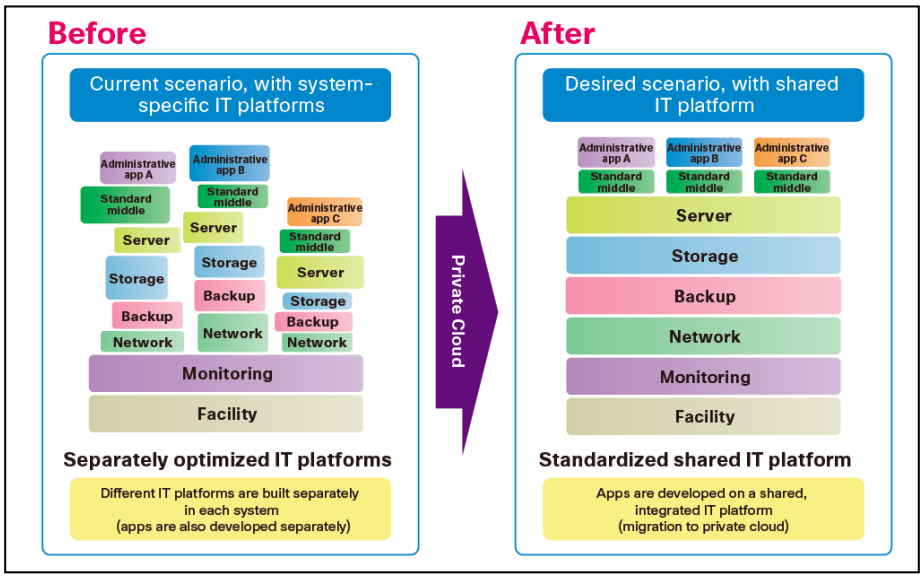

Source: ZK Research, 2014

The Solution

To fulfill on the private cloud vision, NTT DATA implemented a data center fabric using Cisco's Unified Computing System (UCS), Cisco Nexus 7000 core data center switches, Cisco Catalyst 3560 and 2960, Cisco MDS 9100 Fiber Channel Switches and Cisco ASA 5500 Adaptive Security Appliances.

NTT DATA chose the Cisco solution for its data center fabric because Cisco offered functions such as service profile and fabric integration that allowed the company to adopt a new approach in the development and operating phases of data center operations. In addition, the Cisco data center fabric had the following benefits:

- **Simple system structure:** The Cisco UCS has been optimized for architectural simplicity, and the integrated system has far fewer elements than a traditional blade server. In addition, cabling is minimized because the network is delivered via the fabric interconnect. The storage network can also be integrated with the fabric, so fewer SAN switches and cables are required. Also, the entire system is managed by UCS Manager, making data center operations less complex and increasing the cost savings.
- **Network affinity:** To build a private cloud using integrated virtualization, reducing network complexity is a must. A Cisco data center fabric minimizes network complexity by utilizing fabric extender technology. This reduces the number of

management points in the network, facilitating network interoperability.

- **Virtualization affinity:** The Cisco UCS hardware was designed to maximize the benefits of virtualization technology. For example, it has virtual interface cards for controlling the quality of service for each virtual NIC. Additionally, by using the service profile function, the operational burden of installing and swapping out servers can be reduced.

The Benefits

The NTT DATA private cloud built on a Cisco data center fabric resulted in the following benefits for the organization:

- **Reduced TCO:** The total cost of running the environment has been cut in half compared to the pre-installation level. This was achieved through integrated virtualization of servers and network integration using the fabric interconnects, which significantly reduced the number of hardware elements required. Compared to the previous system, the initial investment was reduced by 58%, maintenance costs were lowered by 17% and rack costs were cut by 58%.
- **Shortened server provisioning time:** Central management of the servers combined with the use of the profiler function has increased operational efficiency. The provisioning lead times for servers were reduced by more than 50%.
- **Green data center:** The reduction of hardware components combined with other power-saving products has reduced CO_2 emissions. The reduction currently stands at 79%. NTT DATA estimates that over a five-year period, the cumulative reduction in CO_2 will amount to 3,540 tons.
- **Reference architecture for customer-facing services:** NTT DATA is now using the lessons learned and the experience and expertise gained from this implementation to process its own total cloud service, BizXaaS. This will enable the company to suggest customer-facing solutions with NTT DATA serving as innovation partner to BizXaaS.

Section VII: Data Center Fabric Best Practices

Building a data center fabric involves more than just infrastructure. Organizations should follow best practices to minimize the risk and maximize the ROI of the technology deployment. ZK Research recommends the following:

- **Build an agile infrastructure.** Traditional data center architectures are highly inflexible, are difficult to scale, and do not adapt easily to business changes. Building an agile infrastructure allows IT organizations to respond immediately to changes in the business climate, freeing IT from being the roadblock to competitive differentiation.
- **Converge all data center operations under one management leader.** Although much of data center evolution has focused on the rigid nature of the infrastructure, the organizational groups that support the operations are equally rigid. Converging all data center operations under a single management domain ensures the storage, application, network and server teams are working toward a common goal, and all changes are well orchestrated.
- **Deploy a fabric that addresses storage, servers and network infrastructure.** The evolution to a virtual data center needs to span across servers, storage and the network. A next-generation fabric architecture needs to extend across all three domains. Migrating just a single area will significantly limit the benefits and ROI.
- **Automate as much as possible.** Human error accounts for the largest component of downtime in data centers today. A major cause is the fact that administrators need to perform repetitive tasks under tremendous pressure from the business, which leads to errors. The ability to automate and coordinate many of the tasks required to operate a data center can significantly increase availability, reduce the amount of downtime due to human error, and allow IT to respond faster and more accurately.
- **Strive for 100% virtualization.** Virtualization has been used extensively to reduce the number of physical servers in a data center. However, ZK Research estimates that only about 45% of workloads are virtualized today, with the majority being Tier 2 and Tier 3 applications. Virtualization technology has evolved rapidly over the past few years and can now support even the most mission-critical applications. Additionally, having the workload virtualized provides a higher level of resiliency because the VM can be moved between servers or data centers in real time. The goal for every organization should be to virtualize all workloads over time.

Section VIII: Conclusion

The industry is on the verge of another major computing revolution—this time to a fully virtualized unified data center. However, the path to the unified data center requires more than just using virtualization technology to consolidate servers. Applications, servers, storage and network infrastructure must work together to create an agile, flexible environment that can adapt to the short- and long-term needs of the business. A data center fabric is critical to supporting the journey to a fully unified data center, and should be one of the most important initiatives of all organizations looking to gain a competitive advantage.