Application Fluency in the Data Center

Table of contents

Introduction

The need for data center transformation has taken on a new urgency. From the now mainstream and widespread adoption of server virtualization to the more forward-looking opportunities of cloud computing and desktop virtualization, companies are taking the initial steps necessary to deploy a next-generation data center switching network, one that is more agile and adaptable to the changing needs of the enterprise, and at reduced costs.

In order for the data center to truly evolve, the network must now be modernized. This modern data center network needs to respond to the new requirements of server and desktop virtualization. It also needs to deliver a high-quality user experience, especially for real-time applications such as video, collaboration and video surveillance that executives now consider essential to their organizations' responsiveness, creativity and security. As the number of mobile devices begins to overtake deskbound and office-bound equipment, the network must also support the plethora of smartphones, tablets and other mobile devices infiltrating the organization that are often no longer under control of the IT team and thus cannot be tuned for application delivery.

Virtualization, new applications and new devices require moving away from the old multi-tier network architecture in the data center to a true switching fabric that provides low-latency, any-to-any connectivity. An innovative data center switching fabric can form a "mesh" network that enables a range of innovative data center deployment models — from dedicated virtual data centers to multi-site private clouds to a hybrid cloud environment — while providing a high quality user experience, agility and reduced costs.

In this white paper, you'll learn why the traditional data center network infrastructure is under extreme pressure, what changes need to happen to modernize the data center network infrastructure, and how an application-fluent approach can help ensure a successful step-by-step transition toward the next-generation enterprise data center switching fabric.

> *"Virtualization and distributed applications are transforming every part of the data center. To maximize the potential of virtualization, the network must also transform."*
>
> **- Yankee Group**

A Perfect Storm of Network Pressures

While everyone was focused on new applications, server virtualization, sustainability, mobility, security and other key initiatives, the pressures on the very heart of the data center grew insidiously and significantly. Of course, we are referring to the data center network, where the types of traffic and density of data loads are continually changing. At one point in time, most network traffic flowed between a server in the data center and the desktop (north to south traffic). Now server-to-server traffic (east to west traffic) is expected to surpass server-to-desktop traffic in the enterprise. Current data center networks were never architected to be efficient for server-to-server traffic. This is creating tremendous stress on the existing network infrastructure, threatening the ability of the data center to continue meeting the increasing expectations of the business.

Virtualization

Virtualization, while bringing proven benefits to the organization, is a major culprit in putting additional stress on the network. According to Gartner, between 2009 and 2010, the number of virtual machines in corporate data centers increased more than 50 percent.[1] Yet the traditional data center network is not optimally designed for server and desktop virtualization. For instance, there was a time when a Spanning Tree Protocol (STP) made sense in terms of balancing asset utilization. But in today's highly virtualized environment where low-latency, server-to-server connectivity is needed, a spanning tree is no longer viable. Of course, the notion that an application may automatically and dynamically change location in the data center is completely new and something that data center networks were never designed to accommodate.

[1] "Gartner's Numbers for Virtualization Say it All," by Alex Williams, December 20, 2010, ReadWriteCloud.

Real-Time Applications

At the same time, new applications — many of them requiring real-time communications — are also pushing the network to its limits in terms of bandwidth. Social media sites such as Facebook, YouTube and Twitter have redefined the way we interact in our personal lives, but they've also made a significant impact on business. In a recent Frost & Sullivan survey of 200 C-level executives in North America, nearly half of the respondents reported that their organization uses social media.[2] The rise in social media use has led to increased employee and executive expectations for the availability of capabilities such as on-demand media, high-bandwidth connectivity and flawless audio/video playback (in many cases, high definition).

Mobile Devices

In a 2010 CIO magazine survey, more than two-thirds of IT leaders believed that mobile technology facilitates business innovation in their organizations.[3] Whether smartphones, tablet PCs or other mobile devices, these technologies are not only changing the way we work, but they are also changing the requirements for application delivery. With the rise of end user-owned devices, IT can no longer rely on the ability to tune application delivery at the endpoint. The network needs to become application fluent, that is, taking over the function of understanding the needs of the application and tuning for application delivery performance.

Video

With the rise of Web and video conferencing, video surveillance and streaming media, video is now commonplace in the work environment, which creates substantial new demand for quality-controlled bandwidth. In the past, IT departments increased capacity as needed by simply adding more raw bandwidth capability to the network, but that is not economically viable in today's media-rich environment. A better approach is required where the network can dynamically allocate available bandwidth based upon business priority.

Modernizing the Network: An Application-Fluent Approach

With bandwidth-intensive video applications, virtualization, and new devices being introduced into the enterprise at a rapid pace, it's critical that the network, including the data center, understands how to accommodate and dynamically adapt to these increasingly demanding workloads. Building on the model that Gartner calls "application fluency," PVA and Alcatel-Lucent have adopted an application-fluent approach toward architecting the next-generation data center network.

The Application Fluent Network is based upon a resilient architecture with automated controls in which the network dynamically adapts to the application, user and current device to provide a high-quality user experience, as well as simplified operations. This is achieved through a design that is built on three core pillars and is applied to the data center as follows:

- Resilient architecture: Simplifies the network through a data center fabric that provides low-latency connectivity, has a small footprint, and is ready for the convergence of storage traffic. Virtualization of the network allows for any-to-any connectivity, supports the definition of virtual data centers, enables coexistence with the cloud and supports Ethernet Virtual Bridging. It also ensures resiliency due to localization of individual component failures, and offers built-in security.
- Automatic control: Includes the ability for applications to be managed as services, where the network understands each application, adapts to virtual machine movement, and dynamically allocates bandwidth.
- Streamlined operations: Offers automated, low-touch provisioning of top-of-rack switches and applications. The design provides vendor-agnostic integration between the application virtualization platform and network management platforms, with the lowest power consumption possible.

According to Gartner, an application fluent and scalable network can help enterprises maintain business continuity and meet user SLAs[4] by addressing both the internal and external forces that can impact application delivery. At the same time, an application fluent network can serve as a foundation for data center transformation, empowering the enterprise to further evolve towards a more flexible, powerful and simplified computing environment.

A Roadmap for Application Fluency in the Data Center

Figure 1 depicts the path for moving from a client-server computing model to application fluency in the data center. The path includes two important milestones for corporations: virtualizing the network itself with deployment of a switching fabric, followed by convergence of data and storage networks. The diagram also highlights the network requirements imperative for success at each stage.

Many organizations have already made the critical first step toward application fluency with server virtualization. Unfortunately, many have not been able to reap all of the benefits of server virtualization because virtual machine (VM) movement requires manual intervention to modify network provisioning. That is one reason why network virtualization is the next step on the path.

Network virtualization enables the data center switching network to route traffic based on the optimal path in the network without being constrained by the underlying physical connectivity. Virtualization of the network delivers a switching fabric with low-latency, any-to-any connectivity. Equally important, a true data center fabric must automatically adapt to virtual machine movement to relieve IT of the burden of manually provisioning the network.

The choice of virtualization technology for the data center fabric is key to efficiently enabling a multi-site data center and connectivity to public cloud services. Ideally, the virtualization technology would allow the data center network to appear as a single, logical fabric capable of being physically spread across several corporate sites, as is the case with Shortest Path Bridging (SPB). This ensures that IT has a more efficient, unified framework to manage operations across multiple data center sites. Having all of the enterprise data centers connected in this way, in effect, transforms the corporate data center into a private multi-site cloud.

Convergence of storage and data onto Ethernet is expected to become more prevalent as 40GigE connectivity becomes more readily available. This will be due to the combined data and storage bandwidth requirements of many virtual machines placed on a single server. Key technologies, which include lossless Ethernet (also referred to as Data Center Bridging [DCB] and Fibre Channel over Ethernet [FCoE]), are still being finalized in the standards committees and are a key requirement to support the reliable delivery of storage over Ethernet. Any data center fabric deployed today should be ready to support these protocols.

4 "User Survey Analysis: Network Challenges and Opportunities in Data Centers Through 2011," by Naresh Singh, November 22, 2010, Gartner.

Ultimately, new embedded security technology and 100GigE connectivity will enable a true hybrid cloud model for corporations — one that allows organizations to mix and match resources as they see fit, while keeping mission-critical data on premise.

The Fabric for Change: Bringing Application Fluency to Data Center Switching

What does it take to move from a multi-tier switching hierarchy to a true fabric in the data center? It requires an innovative blueprint for application fluent data center switching that offers the low latency, high density and sustainable design that enterprises need as they evolve their data center network. It is essential that the blueprint offer an incremental ability to deploy the new fabric in the data center, while also incorporating the essential capabilities for optimizing the user experience, improving network manageability, increasing agility and reducing costs. To deliver on these capabilities, the following design concepts provide a unique blueprint for application fluent data center switching.

The Pod

Virtualization in the data center requires enterprises to optimize server-to-server traffic while striving to reduce costs. The Pod concept employs a unique direct-connect architecture for top-of-rack switches as shown in Figure 2. The Pod is a highly dense structure that allows server-to-server traffic to be delivered without the need for a core switch. The example shown in Figure 2 interconnects six top-of-rack switches delivering 240 server-facing ports while keeping latency between servers in the same pod at less than two microseconds. The network virtualization technology used is Shortest Path Bridging (SPB) and Virtual Chassis, to allow all top-of-rack switches to be managed as a single device with one IP address.
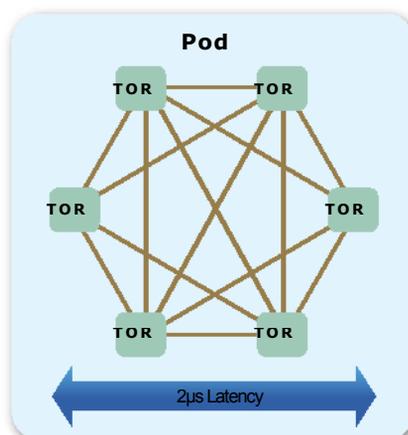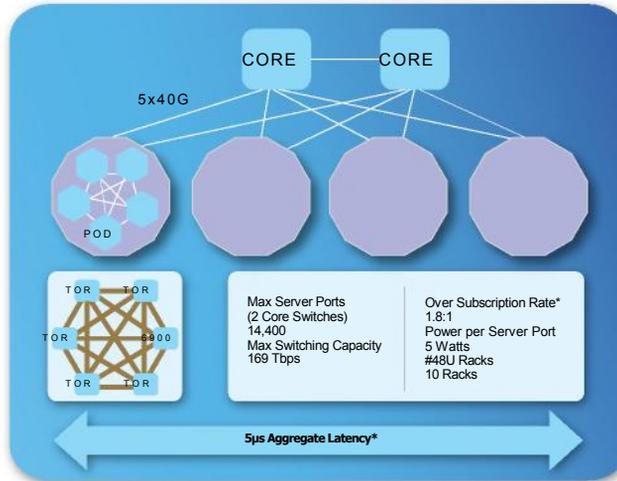

The Mesh

The Mesh delivers a complete switching fabric that can bring together more than 14,000 server-facing ports with only two core switches, delivering aggregate end-to-end latency of five microseconds and unmatched resiliency [5]. The Mesh is constructed by interconnecting Pods and core switches as shown in Figure 3, again using SPB for network virtualization. SPB provides full interoperability with data center interconnect technologies for multi-site private and public cloud deployments. The Mesh enables enterprises to create virtual data centers supporting defined workgroups or departments.

Also, the Mesh is ready for storage to be converged onto the same fabric with lossless Ethernet, FCoE or native Fibre Channel interfaces, and all switching is done in the fabric with Ethernet Virtual Bridging (EVB).

[5] Assuming 70 percent of server-to-server traffic within a Pod, 20 percent between Pods and 10 percent via core

Figure 3. PVA Global/Alcatel-Lucent Mesh

5x40G

POD

TOR TOR
TOR 6900
TOR TOR

Max Server Ports
(2 Core Switches)
14,400
Max Switching Capacity
169 Tbps

Over Subscription Rate*
1.8:1
Power per Server Port
5 Watts
#48U Racks
10 Racks

5µs Aggregate Latency*

*Assuming server-to-server traffic 70% within a Pod, 20% between Pods and 10% Via Core

What is a network fabric?

Simply put, it's a flat network where every port is virtually connected to every other, providing high-speed, low latency interconnectivity. It's operationally simple, allowing the network to be managed as a single entity rather than as individual components.

Virtual Network Profile (vNP)

The Virtual Network Profile (vNP) is shown in Figure 4. Using vNP, the virtual machine applications are managed as services where the network understands each application and can automatically adapt to optimize application performance, including automating the movement of virtual machines within the fabric, agnostic to the choice of server virtualization platform. The vNP contains the critical information the network needs to understand each application, including provisioning requirements, security profiles such as access control rights and VLAN assignment, expected quality-of-service levels, the priority of the application with respect to other applications, and specific latency and jitter requirements. With this knowledge, the vNP can manage applications as services, enabling the network to automatically discover the location of each virtual machine, bind a specific vNP to that virtual machine, and provision the network for the applications, including modifying the network configuration to follow virtual machine moves.

Figure 4. Virtual Network Profile (vNP)

Virtual Network Profile (vNP)

Network Provisioning	Security Profile
	Application	
Quality of Service Requirements	Priority

vNP

Top of Rack Switch
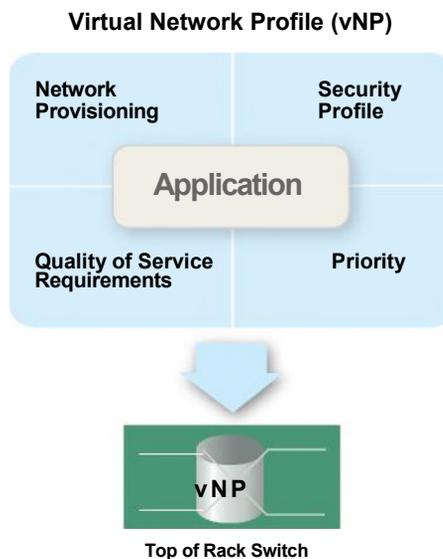
Enabling Multi-Site Data Centers and Private and Hybrid Cloud Models

As enterprises move further along the application-fluency roadmap, their data center fabric will enable them to begin taking greater advantage of the benefits of cloud computing: increased agility, faster provisioning and rollout of new services, simplified management and reduced costs.

Enterprises can begin preparing to leverage public cloud services by migrating their data center to a private cloud architecture, where IT resources are dynamically allocated across the business in an infrastructure-as-a-service (IaaS) model. This brings the benefits of fast deployment and scalability expected of the cloud while still delivering the control and security that organizations desire. This blueprint lets enterprises transform the corporate data center into a multi-site, private cloud by enabling the data center switching fabric to become one logical structure that can be physically spread across several corporate sites.

With the fabric acting as one entity, it helps simplify management and improve security. The fabric can be partitioned to support virtual data centers for specific departments. An application or service can exist in more than one virtual data center and security can be applied differently within each virtual data center and at virtual data center boundaries. The virtual data center boundary is adjusted automatically to account for virtual machine movement.

It is critical for data center decision makers to make informed technology choices, such as SPB, to enable the corporate data center to become a multi-site private cloud. It is believed that with the drive to the SPB standard for network virtualization, the optimal strategy to enable a hybrid cloud model, depicted in Figure 5, is achieved.

In the hybrid cloud model, public cloud services are seamlessly delivered onto the data center fabric where they can be combined with local services to provide composite applications for users. The hybrid cloud model lets enterprises broaden the flexibility, availability and scalability of the IT environment without sacrificing the security and control available with a private cloud. To enable the flexibility needed for the hybrid cloud, the choice of a service provider-compatible virtualization technology, SPB, and the ability to manage applications as a service with vNP are essential.

Figure 5. Hybrid Cloud Model
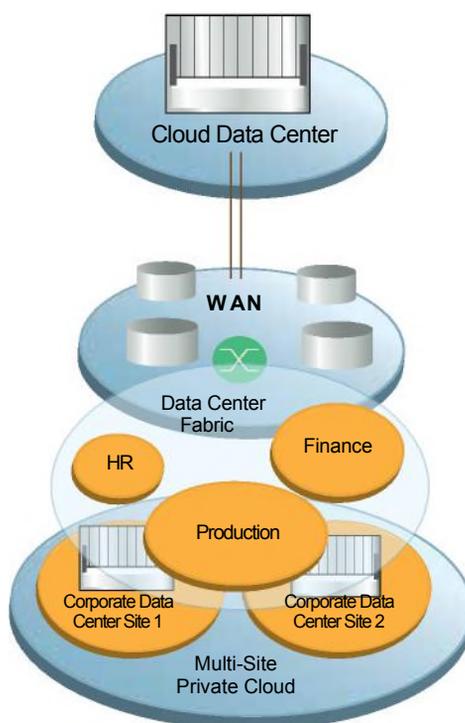
The Benefits of the Application Fluent Data Center Network

With the Application Fluent Network, enterprises get the flexibility they need to move beyond costly client-server computing in a step-by-step, incremental manner starting with the deployment of a single Pod. Organizations can start to benefit from an application fluent data approach and manage applications as a service across a range of data center deployment models, including multi-site private clouds, dedicated virtual data centers and a hybrid cloud that integrates service provider offerings.

For enterprises, application fluency delivers benefits across the organization, from the end user to IT to the corporate bottom line. With an Application Fluent Network, organizations can:

Deliver a high-quality user experience:
- Meet the needs of virtualized applications with market-leading low latency
- Automatically optimize application performance through management-as-a-service capability
- Minimize downtime by localizing the effect of network failures with resilient direct-connect architecture

Increase agility:
- Optimize server utilization more rapidly and with fewer errors via automated virtual machine movement
- Simplify application deployment and disaster recovery with a multi-site private cloud
- Accelerate the rollout of new services through a hybrid cloud model

Reduce Costs:
- Reduce capital and operational costs with a high-density, low power-consuming network fabric
- Gain application performance visibility to reduce troubleshooting effort
- Streamline IT operations through integration with standard hypervisors

Conclusion

The data center infrastructure is undergoing a rapid transformation to drive down costs and improve the end user experience in the face of rapidly evolving technology trends. This is causing fundamental changes in how data centers are designed and how data center networks need to evolve. Not only do data center networks have to adapt to the higher bandwidth requirements of media-rich applications like video, they also need to support a wider variety of devices and connectivity methods. At the same time, the network must be optimized for server and desktop virtualization, while enabling network virtualization as the next step in modernizing the network.

Application Fluency provides a new approach to networking in the data center. Following this approach, Enterprises have a new blueprint for a complete data center switching fabric which extends the boundaries of the data center through virtualization of the network with an innovative direct-connect architecture. Leveraging market-leading scalability, low latency and low power consumption, enterprises can move beyond costly client-server computing by managing applications as a service across a range of data center deployment models, including multi-site private clouds, dedicated virtual data centers and a hybrid cloud that integrates service provider offerings.

The Application Fluent Network helps enterprises ensure a high-quality end user experience and more simplified operations. With the right approach and a trusted partner to work toward application fluency, enterprises can feel confident moving forward in their data center transformation.