

Data Center Evolution Is Dependent on a Network Fabric



by Zeus Kerravala | May 2010

I. Evolution of the Data Center


The data center has undergone several major computing transitions over the past few decades. It has evolved from a single compute and storage resource, to distributed client/server computing, to Internet computing. Now it finds itself on the precipice of yet another transition—this time to the virtualized data center. The transition to a virtualized data center will be the single biggest IT transformation of this generation and will reshape not only the compute environment but also the network environment. The biggest impact and challenges to the network are due to the following drivers:

- **Virtual machine (VM) mobility:** As organizations look to use virtualization more strategically, the VM will evolve from a static resource that resides on a fixed computing platform to a mobile resource that can be moved from server to server in real time. Most current data center networks are designed in multiple tiers that add too much overhead and latency to allow VM mobility to be used broadly.
- **Resource pooling:** Most of the focus of virtualization has been on servers. Over time, however, all of the resources in the data center will become fully virtualized and “pooled” (see Exhibit I). 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 it needs, governed by business policies. For cost/benefit reasons and optimum utilization of existing resources, the ideal scenario is to create large, dynamically shared pools. Unfortunately, a network large enough to handle this using current best practices would be far too complex to support the on-demand vision.
- **Cloud computing:** Moving and using compute resources on the fly is a great concept. Cloud computing lets enterprises pay only for the resources they need, shifts investment from capex to opex and offloads much of the complexity of managing compute resources. While all of this is true, networks built using traditional best practices can’t support the real-time interaction needed between the cloud and the private data center.

All of the current trends being driven by the transition to a fully virtualized data center are highly dependent on a seamless, real-time network. However, significant challenges exist with current network infrastructure and architectures.

Exhibit I: Resource Pooling in the Future Data Center

Source: Yankee Group, 2010



This custom publication has been sponsored by Juniper Networks.

Limitations of Existing Data Center Network Architectures

Current data center networks are architected to have three distinct tiers: a first access tier that connects all of the server/storage devices in a rack; a second tier that aggregates access switches to the core of the network; and a third core network tier that connects the routers and allows communications to the outside world. This hierarchical design is optimized for efficient North-South traffic flow, in that network traffic moving from one server to another passes through each tier of the architecture, through the core, and then back out through each tier until it reaches the destination server (rather than East-West traffic flow, in which traffic travels from one server directly to the other with no hops in between).

Traditional best practices also call for the use of Spanning Tree Protocol (STP) to prevent loops (traffic that takes a circular path and never reaches its destination) and ensure broadcast radiation. STP is a protocol that operates at Layer 2 of the OSI stack and creates a “spanning tree” within a mesh network of Ethernet switches. STP disables the links that are not part of the spanning tree, creating a single, active path between any two nodes in the network. The disabled links become active only when one of the active links fail.

While this architecture was optimal for legacy data centers, it does not meet the challenges of the virtual data center for the following reasons:

- **Predominant East-West traffic flows:** With the advent of new service-oriented architecture (SOA) applications and virtualized environments, the predominant traffic flow is now East-West. Since the system is optimized for North-South, this leads to poor performance. For example, when traffic moves from one device in a rack to another, the traffic must first traverse the entire path, pass through the core and then come back. Each of these hops adds latency to the transaction, since processing is repeated at each hop. Ideally, traffic would flow so that no port is more than one hop away from another.
- **Inefficient use of network infrastructure:** Because STP disables up to half the ports in the network, the enterprise cannot fully utilize all its network resources, and it ends up needing to purchase far more network equipment than is necessary. It would be better if all of the ports could be active at the same time without the fear of routing loops.
- **Complexity and inability to scale:** The larger a network becomes, the more tiers need to be built, the more interconnections are required and the more management overhead is needed. Furthermore, since the devices don't share state and instead communicate with one another using protocols like STP, the number of device interactions increases exponentially as devices are added. Exhibit 2 shows how the number of devices, device interactions and packet processing stages grow as companies add more servers to the data center.

Exhibit 2: Data Center Network Complexity Grows Exponentially With Server Growth

Source: Yankee Group, 2010

Number of servers	Number of ports (4 per server)	Number of top-of-rack network devices	Total number of network devices	Device interactions	Packet processing stages
25	100	3	4	6	4
125	500	11	13	78	4
250	1,000	21	23	253	4
500	2,000	42	48	1,128	6
750	3,000	63	69	2,346	6
1,000	4,000	84	90	4,005	6
1,250	5,000	105	111	6,105	6
1,500	6,000	125	131	8,515	6
1,750	7,000	146	152	11,476	6

Organizations need to radically simplify the way the data center network is designed and operates if the virtual data center is going to become a reality. But if current architectures do not scale to meet the challenges, what is the solution? The answer is that the network needs a fabric architecture, rather than the traditional structure.

II. Defining a Network Fabric

The term “fabric” has been thrown around quite liberally in the network industry for the past few years, but the fact remains that Ethernet networks are not fabrics today (although they are with storage networking).

Fabrics actually exist today in the Ethernet world, but they are limited to the physical confines of a single device. A switch or router fabric enables efficient device performance and management through the following characteristics:

- Any port in the network is directly connected to any other port. This significantly reduces transmission latency.
- Packet processing is done once to reach the final destination.
- The fabric has a single, shared state table for all ports.
- The network is flat; all ports in the device are connected in a single tier.
- All ports in the fabric can be managed from a single point.

A physical switch is made up of a number of connectivity ports and a fabric that connects the ports together. However, when multiple devices are connected together to form a traditional network, all the problems listed in the previous section come into play. This is because the network fabric is currently limited to the physical confines of a switch. Each device might have its own fabric, but the fabrics are islands that are connected over cables to form the network.

A true network fabric built using a switch fabric architecture would be like extending the fabric out of the switch, over the cables to the other switches. In essence, the entire data center network would behave and be managed like a single device and take on the characteristics of a fabric.

Additionally, since every port would be connected to each other, there would be no use for STP, meaning all of the ports could be active at the same time without concerns about routing loops.

III. Benefits of a Network Fabric

There are numerous benefits to implementing a data center network that has a fabric architecture versus a traditional hierarchical architecture. The fact that the overall complexity of the network has been greatly reduced has many technical benefits as well as business benefits. The primary benefits are as follows.

Business Benefits

- **Better ROI for virtualization technology:** Many organizations have invested millions of dollars in virtualization technology and have experienced a solid ROI on that investment. However, many of the advanced features of virtualization, such as workload mobility, will not function properly because of the latency in the network due to traffic moving North-South and having to pass through each tier. Since the fabric architecture enables East-West movement of data, traffic would never be more than one hop away from its destination, reducing the latency by an order of magnitude. This would allow organizations to use virtualization for more than just tactical initiatives like server consolidation. Virtualization could be used for strategic initiatives such as data center automation and disaster recovery, further increasing the virtualization ROI. A modest investment in a network fabric can result in a significant savings in server/storage infrastructure.
- **Lower capital expenditure:** Since all the ports are active, fewer network devices are required. This reduces the capital outlay required to build or refresh the data center network.
- **Lower operational costs:** Fewer managed devices and much fewer device interactions that are exposed to the customer result in simplified management. Thus, network managers can spend more time on strategic initiatives and less time managing the day-to-day operations of the network.
- **Better application performance:** Since no port is ever more than one hop away from any other, application traffic can traverse the network much faster than with a traditional network architecture. This significantly enhances the performance of latency-sensitive applications.

- **More scalable networks:** The biggest inhibitor for scaling networks is the associated complexity that primarily comes from managing multiple devices. With fabric technology, multiple devices can be operated as a single logical device. As a result, there is a significant decrease in complexity, allowing networks to scale much more.

Technical Benefits

- **Better network performance:** Fewer packet processing steps, plus the elimination of STP, mean all network traffic can traverse the network much faster; in large networks, this turns out to be orders of magnitude faster.
- **Simplified network design:** Today, fabric at the access and core layers allows network managers to remove the aggregation tier from the architecture, creating a flatter network. In the future, the network will look like a single network device, dramatically simplifying the design by reducing the network to a single flat tier.
- **A network that scales faster:** Since all the devices in the network contribute to the “fabric,” scaling the network is as simple as installing new devices. Each new device simply adds to the fabric and becomes part of the network’s virtual “single-device” architecture.

Almost all the benefits of a fabric architecture are derived from the fact that the network has been simplified. Exhibit 3 compares the number of devices, device interactions and packet processing stages required for various numbers of servers in both traditional and fabric architectures.

Using a fabric architecture in a network with 1,750 servers (not a huge network in terms of today’s data centers), reduces the number of devices that need to be managed by almost 90 percent versus a traditional architecture. Additionally, the device interactions are reduced by 99 percent, and the number of packet processing stages is typically half of what it is for a traditional network. This is a significant reduction in the complexity of the network, and it will allow the network to grow and scale as fast as the organizations requires.

IV. What to Look for in a Network Provider

Making a purchase decision in times of technology transitions can be difficult. In many cases, it’s not the incumbent vendor that leads the industry through the transition. Consider the shift from mainframes to client/server computing. The dominant computing vendors at that time were companies like IBM, Honeywell, Univac and other giants. None of these companies led the industry into the age of desktop computing. Instead, small companies like Microsoft and Intel led the charge and allowed organizations to scale their compute environments faster than anyone thought possible.

Exhibit 3: Traditional vs. Fabric Architectures

Source: Yankee Group, 2010

Number of servers	Number of ports (4 per server)	Number of top-of-rack network devices	Total number of managed devices		Device interactions		Packet processing stages	
			Traditional	Fabric	Traditional	Fabric	Traditional	Fabric
25	100	3	4	1	6	0	4	1
125	500	11	13	3	78	3	4	3
250	1,000	21	23	4	253	6	4	3
500	2,000	42	48	6	1,128	15	6	3
750	3,000	63	69	8	2,346	28	6	3
1,000	4,000	84	90	10	4,005	45	6	3
1,250	5,000	105	111	12	6,105	66	6	3
1,500	6,000	125	131	14	8,515	91	6	3
1,750	7,000	146	152	16	11,476	120	6	3

With that in mind, network managers must carefully consider their choice of network solutions provider—and not automatically use the incumbent vendor. Network managers must evaluate their network infrastructure providers on the following criteria:

- **Ability to simplify:** Look for a vendor that fundamentally simplifies the architecture. Complexity is the biggest killer for achieving fast, scalable and economically efficient data centers. Solutions need to have simplification built in, rather than forcing customers to manage the complexity on their own.
- **History of technology innovation:** Technology transitions require innovation, not simply finding a cheaper way of solving a problem someone else solved years ago. Choose a vendor with a history of innovation in other markets.
- **Completely open environment:** No solution provider can do it alone, and as the compute and network layers in a data center become codependent, it's critical to use a solution provider that has partnerships with providers of all the related technologies. This will help ensure a standards-based solution.

V. Conclusions and Recommendations

The dawning of a new computing era—the virtual data center—brings with it many benefits that will allow companies to reach new levels of productivity at dramatically lowered cost. However, it's not just the computing architecture that needs to change; the network needs to be rearchitected as well to optimize the effectiveness of virtualization technology.

A fabric architecture is required to build a network that is scalable, simplified and able to meet all of the challenges of the virtual data center. To have a successful deployment, Yankee Group makes the following recommendations:

- **Evaluate vendors on criteria specific to a fabric.** Historical evaluation criteria were specific to an individual network device, such as port density and backplane capacity. Fabric evaluation criteria need to be based on factors such as total number of active ports and transaction latency.
- **Consider virtualization a strategic technology, not just a tactical one.** The majority of virtualization deployments are used to consolidate servers. While this does provide some value, server virtualization is primarily tactical in nature. Using it to mobilize workloads for disaster recovery and other factors, however, can raise the value of virtualization.
- **Treat the network with the same importance as the compute layer.** The network is often referred to as the plumbing or the pipes of a company, and many network organizations buy network equipment from the incumbent because it's the path of least resistance. However, the data center network is now the backplane of the virtual data center and a key enabler for data center virtualization. It can help IT obtain optimum usage while lowering the cost of server and storage resources. It needs to be evaluated with that in mind. Buying network infrastructure that's "good enough" or vendor-proprietary will not allow an organization to meet future business challenges.

Yankee Group—the global connectivity experts

The people of Yankee Group are the global connectivity experts—the leading source of insight and counsel trusted by builders, operators and users of connectivity solutions for 40 years.

We are uniquely focused on the evolution of Anywhere, and chart the pace of technology change and its effect on networks, consumers and enterprises.

For more information, visit <http://www.yankeegroup.com>



Research

Leverage qualitative research to make informed business decisions today and plan for the future.

Data

Gain quantitative insight into current markets and new opportunities via monitors, surveys and forecasts.

Interaction

Connect with analysts to gain deeper insight into research and trends.

Consulting

Get in-depth analysis and actionable recommendations tailored to your needs.

Events

Access world-class events live and online with industry leaders and Yankee Group experts.

Yankee Group's Link products and services provide clients the insight, analysis and tools to navigate the global connectivity revolution.

Zeus Kerravala, Distinguished Research Fellow



Zeus Kerravala, distinguished research fellow, leads the Research Council and is chartered with the responsibility of providing thought leadership to the research organization. Comprising senior research leaders, the Research Council provides outreach to clients and the broader Yankee Group community, as well as ensures that the company's research agenda addresses the needs of business leaders. Zeus drives the strategic thinking of the research organization and helps shape the research direction. Much of his expertise involves working with customers to solve their business issues through the deployment of infrastructure technology.

Headquarters