Learn About Data Center Fabric Fundamentals

This Learn About will introduce you to data center fabrics, describe what they are, and explain why they are becoming increasingly necessary to the success of businesses competing in a global, rapidly changing economy.

What Is Data Center Fabric?

A data center fabric is a system of switches and servers and the interconnections between them that can be represented as a fabric. Because of the tightly woven connections between nodes (all devices in a fabric are referred to as nodes), data center fabrics are often perceived as complex, but actually it is the very tightness of the weave that makes the technology inherently elegant.

A data center fabric allows for a flattened architecture in which any server node can connect to any other server node, and any switch node can connect to any server node (server refers to both compute and storage).

The flattened architecture of fabrics is key to their agility. Data center fabric architectures typically use only one or two tiers of switches as opposed to data centers that implement multi-tier data center network architectures

In almost all data center fabric architectures traffic can be transmitted between server nodes by traversing a set number of switches, which results in extreme efficiency and low latency (a few microseconds of latency at each hop can become seconds of latency per transaction).

Software automation and management can determine the path that traffic takes from its originating server node to its destination server node based on availability. The high bandwidth and availability of multiple direct paths in fabric architecture eliminates the need to block traffic while avoiding the transmission slowdown caused by network bottlenecks.

In some architectures, the fabric is represented by a spine and leaf design in which the fabric mesh incorporates devices on the edge (the leaves) of the fabric and switches on the spine. Fibre channel designs often allow for redundancy and failover by using two fabrics. In this case, the fabrics share the edge nodes; otherwise they are not connected. If a link or one of the fabrics is not functioning, the other fabric takes over traffic transmission. The spine and leaf model allows for facile expansion limited only by the number of supported devices and their ports.

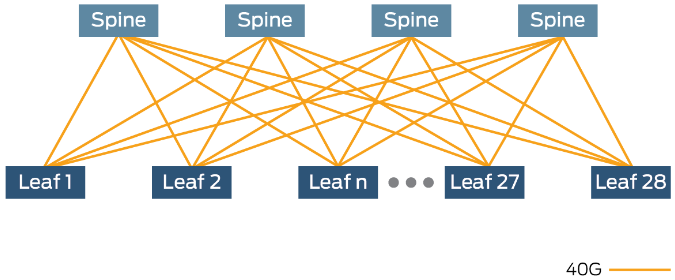

Figure 1 *Spine and Leaf Data Center Fabric*

The Rise in East-west Traffic

Cloud computing, virtualization, big data, and the increasing use of dynamic applications have exerted pressure on the older multi-tier network data center architecture, which originally evolved to serve client-server and web applications. Data centers that implement the multi-tier network model cannot efficiently manage the volume and scale of data that these new technologies process, generate, analyze, transmit, and store.

East-west traffic, also referred to as horizontal traffic, is traffic that occurs within the data center. East-west traffic dominates in nearly all cases of cloud computing, virtualization, and big data. It is one of the driving forces behind the development of data center fabrics.

The rise in east-west traffic was spurred by increased use of applications using services that resided on different servers; as the single application that handled everything fell into disfavor and distributed, more collaborative applications took hold.

For cloud and big data applications, clusters of servers with different functionalities are required to collaborate in handling a single request, generating large amounts of east-west traffic in the process. For example, big data applications typically handle complex problems that require them to use multiple compute and storage servers. They search through large sets of data and move data into and out of server memory and on to other devices for analysis.

Multi-tier network data center architectures were optimized for north-south traffic, traffic from the data center to the end user, not east-west traffic.

The data center fabric architecture model unites holistically all data center resources from processor cores to memory—servers, storage, the network, and peripherals. Its architecture is designed to handle the increase in east-west traffic while maintaining north-south traffic connectivity to end users.

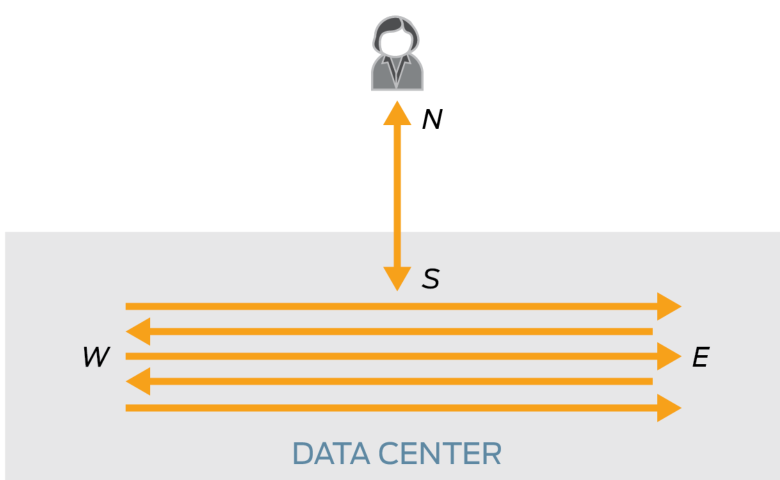

Figure 2 *East-west and North-south Traffic*

Data Center Fabric Characteristics

To better understand this dramatic change in the data center and the increase of east-west traffic, let's summarize the main characteristics of data center fabric architecture:

- It is exemplified by a flattened infrastructure with a meshed hyper-interconnectivity of its resources—compute, storage, and network. This infrastructure gives storage and compute server nodes direct access to one another through any switch to facilitate optimal communication and processing. In this architecture all nodes are only one hop away from each other, in many cases traversing only a single switch.
- It relies on high bandwidth links, such as 10-Gigabit Ethernet (10Gb) and InfiniBand, to loosely couple its processors, memory, and peripherals. Its high-speed switches interconnect these storage servers and the network in a grid-like fashion making the fabric appear as a single logical whole. It supports distributed applications and applications that generate outsized amounts of data requiring many compute and storage devices to process their tasks.
- It unites switch fabric and storage area network (SAN), and, for the first time, it brings to them high-performance computational capabilities and traffic-forwarding efficiency.

Gartner
"Market and industry trends are changing the way enterprises approach their data center strategies. Several factors are driving enterprises to look beyond traditional technology infrastructure silos and transform the way they view their data center environment and business processes. These include aging data center infrastructure that are at risk for not meeting future business requirements, an ongoing cost-consciousness, and the need to be more energy efficient."

The unique characteristics of data center fabric architecture provide the following benefits:

- Network performance is significantly improved. The use of high bandwidth switches (such as Juniper's QFX10000 series) to interconnect the devices allows for even faster I/O speeds that support converged data and storage traffic which were previously isolated.
- The open ecosystem of most data center fabric architectures lends itself to flexible reconfiguration that makes room for future technologies and scaling. For example, Juniper Networks inclusive MetaFabric technology guards against locking in data center resources to a single provider.
- Fabric's fundamental support for virtualization, shared resources, and scaling can adjust to the dynamic requirements of the applications that utilize them.

An Indifferent Approach to Locality

It is well known that multi-tier data center networking, which started out with three tiers – access, aggregation, and core – grew to include at least two additional tiers. Virtualization brought to the network hierarchy the *virtual switch tier*, and blade servers and blade switches added yet another tier. Data center fabric architectures flattened out the multi-tier network architecture into no more than two tiers by using any-to-any switching fabric to maximize performance and optimize paths for the change in traffic flow from north-south to east-west.

Data center fabric deployments, physical and virtual, typically span across multiple premises, and they can efficiently adapt to the execution of distributed applications. Their flat structure and hyper-interconnectivity allow any application to be run from anywhere in a distributed data center, which means an application can be run from wherever best aligns with its execution requirements. This low-latency, indifferent approach to locality is central to the success of existing, implemented technologies

Accenture Technology • Paul Daugherty, Prith Banerjee, Walim Negm, and Alton Alter • 2014:

"In the future, successful companies will use the Industrial Internet of Things to capture new growth through three approaches: boost revenues by increasing production and creating new hybrid business models, exploit intelligent skills to fuel innovation, and transform their workforce."

but also to emerging ones, such as the Internet of Things (IoT), which is taking hold in the data center and which continues to encompass increasingly more autonomous "intelligent" machines.

IoT embraces machine-to-machine communications that do not involve a person at either endpoint, autonomous devices such as devices with embedded sensors, and devices devoted to analytics—any "thing" that has an IP address extending beyond what is considered standard systems.

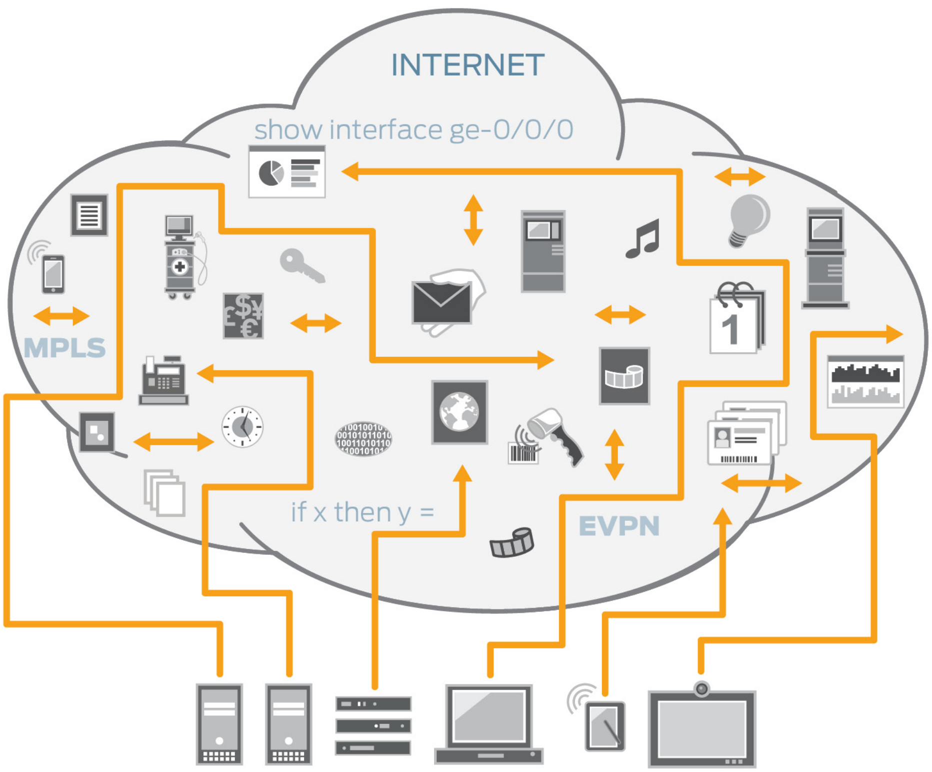

Figure 2 The Internet of Things

Gartner • February 27, 2014 • The Impact of the Internet of Things on Data Centers:

"Processing large quantities of IoT data in real time will increase as a proportion of workloads of data centers, leaving providers facing new security, capacity and analytics challenges.

IoT requires a magnitude of network connections, as well as the infrastructure to handle massive amounts of data generated from globally distributed sources. Transferring that data to a single location for processing is neither feasible nor economically sound. Rather, distributed data centers with efficient, distributed management that supports application execution at any point are required.

Virtualization: Doing More with Less, More or Less

For a long time it was well understood that to manage costs, data centers had to do more with less. The expansion of Internet services and activity, the widespread use of complex interactive applications, and the incessant advancement of technology overall brought with them the need for faster data rates, more compute resources, more storage to handle the volume of data being generated, and automatic "smart" data center networks that could adapt to accommodate the dynamic processing requirements of distributed applications. At the same time, data center rack space, power, and cooling were becoming increasingly costly.

Gartner • David Cappuccio:

"The fact that the data center was very costly in order to support high availability was seen as the cost of doing business. While this may have been a nice side effect of design principals in prior years, when looked at prudently the obvious question emerges; is there a better way to design data centers?"

The solution seemed to come with the introduction of virtualization and cloud computing. As virtualization and cloud computing technologies took off and matured, IT departments began to take advantage of them. They virtualized their data centers and implemented private and hybrid clouds, or utilized public clouds.

To further handle their resource problems and control costs they also relied on colos, in which data center service providers leased server space, resources, and cloud access. The data centers were called colos because the resources and systems of multiple businesses were co-located at the same site. Cloud computing multi-tenancy made it possible for even the same computers to be shared.

Virtualization of enterprise data center infrastructures quickly moved from being the exception to the rule. But it was not the complete solution as it, too, had its problems.

IT departments that virtualized their data centers had to deal with virtual machine (VM) stall and sprawl, the changing definition of a private cloud, the impact of VM backup and recovery, and the load put on the compute resources by running real-time applications such as anti-virus scans on VMs. These and other consumers of processing power slowed down performance and introduced management complexity. A single server might have enough memory and CPU power to run thirty or so VMs, but its saturated network interface cards (NICs) could cause network problems that kept VMs from being able to communicate with one another.

Data center fabric had been hinted at by switch fabric and foreshadowed by grid computing. Not only did data center fabrics offer more for less, but they also responded to the problems introduced by the cloud computing and virtualization "solutions."

Figure 3 shows a data center fabric communications network model that fully supports virtualization without the slow-down glut that occurs in hierarchical multi-tier networking data centers. Its nodes – in this case, servers running multiple VMs – are directly connected to one another through high-speed switches allowing for high-performance, low-latency communication.

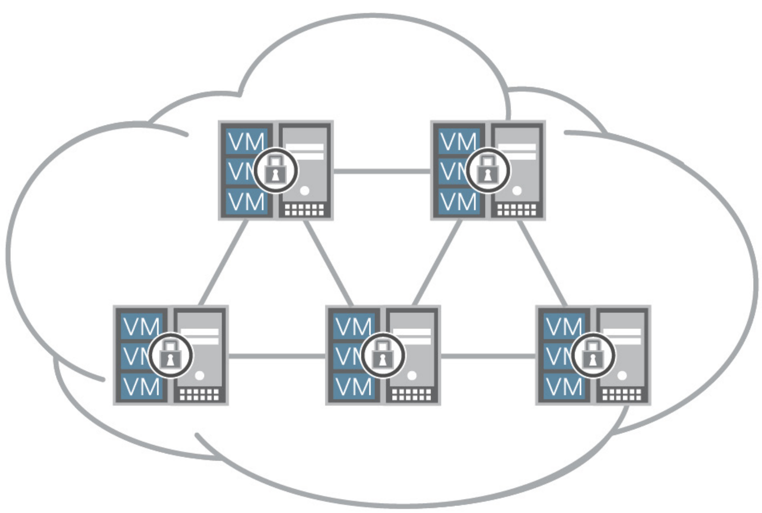

Figure 3 *Virtualization in the Data Center Fabric Network*

Juniper Networks Data Center Fabric Solutions

Today's network operators want the ability to keep up with and effectively handle changes wrought by the going-global phenomenon, cloud computing, and virtualization. They want "comprehensive agility across their networks" and their data centers, whether those data centers are centrally located or are geographically dispersed across multiple premises. And they need to improve performance and simplify management.

To meet the requirements of these network operators, Juniper Networks developed data center fabric solutions that:

- Accommodate physical and virtualized infrastructures and provide a simplified management solution that integrates both.
- Create a perfect fit for the requirements of virtualization, cloud computing, and big data – diminishing the problems introduced by these technologies while taking full advantage of their capabilities.
- Give beta data rates and results because they are designed for faster speed.
- Allow for the rapid, automated deployment of applications and services.
- Offer improved cost control and centralized management.
- Provide enough flexibility to scale for future expansion, both in software and hardware, with low impact to business operations.

Juniper Networks Virtual Chassis

For an excellent overview of Juniper Networks' Virtual Chassis begin on this page in the Juniper TechLibrary: http://www.juniper.net/techpubs/en_US/junos14.1/topics/concept/virtual-chassis-ex4200-overview.html.

For field instructions on configuring and deploying Virtual Chassis in EX Series Ethernet switches, see Chapter 2 of Day One: Configuring EX Series Ethernet Switches, 3rd Edition: http://www.juniper.net/us/en/training/jnbooks/day-one/fabric-switching-tech-series/config-ex-series/.

Juniper Networks Virtual Chassis Fabric

The Juniper TechLibrary has myriad documents and support for Virtual Chassis Fabric. Start here: http://www.juniper.net/techpubs/en_US/junos14.1/information-products/pathway-pages/qfx-series/virtual-chassis-fabric.html.

For an in-depth review of Virtual Chassis Fabric published by O'Reilly Media: http://www.juniper.net/us/en/training/jnbooks/oreilly-juniper-library/.

Juniper Networks MetaFabric Architectures

For an in-depth curated selection of data center solutions and training, see the Juniper TechLibrary's Data Center solutions: http://www.juniper.net/documentation.

Juniper Networks Junos® Fusion

Junos Fusion is Juniper Network's newest data center architecture, so new that specifics are still being evaluated: http://www.juniper.net/us/en/products-services/switching/data-center-switching-architectures/.

http://www.juniper.net/assets/us/en/local/pdf/solutionbriefs/3510502-en.pdf.

Learn About Data Center Fabric Fundamentals

by Judy Thompson-Melanson

A data center fabric allows for a flattened architecture in which any server node can connect to any other server node, and any switch node can connect to any server node. This flattened architecture of fabrics is key to their agility. Learn more about the fundamentals of data center fabrics and their importance in our daily lives.

Judy Thompson-Melanson is a Juniper Networks staff technical writer with over twenty-five years in the industry. She has written API documentation, design guides, and networking and security documentation for many companies including Apple, Sun Microsystems, Cisco Systems, and Intuit.
The author thanks the following for their engagement in this project: Patrick Ames, Editor in Chief; Karen Joice, illustrator; Nancy Koerbel, editor; and project promoter, Mindy Isham

For more information see: www.juniper.net/documentation

© 2015 by Juniper Networks, Inc. All rights reserved.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners. Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

ISBN: 978-1-941441-12-1 Version History: First Edition, April 2015 2 3 4 5 6 7 8 9

ISBN 978-1941441121

50200

9 781941 441121