

# NETWORK FABRICS FOR THE MODERN DATA CENTER

New Data Centers Require a New Network

## Table of Contents

Executive Summary .....	3
Introduction .....	3
The Modern Data Center .....	3
Old Architecture Can't Support the New Requirements .....	4
New Traffic Patterns Drive New Network Designs .....	5
Three Classes of Traffic in the Data Center .....	6
Mathematics of Complexity .....	6
The New Network Requirements .....	8
Simplification .....	8
Lowest Possible Latency, Highest Possible Bandwidth, Lowest Possible Congestion .....	8
The New Network Requires a Scalable Network Fabric .....	9
Juniper Delivers Fabric for Networks .....	9
Extending Fabrics Across Data Centers .....	9
Single Operating System, Single Orchestration Platform .....	12
Conclusion .....	12
About Juniper Networks .....	12

## Table of Figures

Figure 1: Legacy tree structures impose varying levels of delay when delivering packets across the network.....	5
Figure 2: The number of potential interactions grows disproportionately with each new managed device.....	7
Figure 3: Traditional multitier data center network environment.....	7
Figure 4: Juniper solutions deployed as a fabric dramatically simplify the data center network, reducing the number of devices, improving performance, and reducing costs.....	10
Figure 5: Virtual Chassis technology reduces complexity associated with interactions between disparate switches in the data center network.....	11
Figure 6: Juniper Networks' vision for the ultimate simplification of the data center is Project Stratus, which will deliver a single fabric that unites Ethernet, Fibre Channel, and Infiniband networks.....	11

## Executive Summary

It has been almost 50 years since computers combined with networks to create the first online transaction processing system. This development ushered in the modern data center, where local processing and storage came together with network-connected users. In a very real sense, the network created today's data center; now, in an ironic twist, modern data centers require a new network, one that can help drive the data center into a new era.

This paper is intended for technical decision makers who are interested in understanding Juniper's vision of how to evolve the data center. It examines the changes occurring in the data center and how these changes mandate a different design approach. This paper explains Juniper's approach to building data centers using fabrics and how this will evolve to meet the increasing challenges of scale and complexity. The reader will learn about Juniper's current fabric solution for data centers and get a view of our vision for future evolution.

## Introduction

*Draw the network of the past and you'll see a complicated hierarchy of switches—essentially a tree. Now draw the network of the future and you'll see a fabric—a simple, flat tapestry that unites all network elements as a single, logical entity.*

As business requirements drive information technology in new directions, the network is defining the data center of the future. The early mission of providing user access to applications remains, but modern networks are also connecting systems and even data centers to create cloud computing infrastructures that share a common set of global processing, storage, and application resources. Storage networking is exploding as enterprises look to consolidate their storage environments with their data center networks. Interprocess communications (IPC), the glue that binds service-oriented architectures (SOA), is also growing as a network application, to the point where it rivals user connectivity.

Networks and data centers are entering into a new, tighter, and more exciting partnership.

*Traffic, once bidirectional, is now omnidirectional.*

Enterprise sensitivity to delay and packet loss, once confined to a few vertical markets, is spreading horizontally to all sectors. Virtualization, cloud computing, SOA, and the ever expanding operational demands for improved productivity through information technology are all driving the data center into a new era. Applications, information, storage, processing, and networking are now equally important elements in the data center, and each must evolve to optimize and support the others. New requirements create new demands, and new network capabilities open avenues for improvement in IT process costs and efficiencies.

Juniper Networks® was founded on the principle that the Internet changed everything. Networks must be built in a fundamentally different way to meet the evolving needs of the Internet. Toward that end, Juniper Networks has introduced a new "3-2-1" data center network architecture based on innovative fabric technology that enables multiple interconnected networking devices to behave as a single logical device. This fabric technology dramatically reduces complexity and consolidates the data center network from 3 layers to 2, moving eventually to a single layer.

## The Modern Data Center

The modern data center has undergone a series of changes that have significantly impacted business operations.

First, the advent of service-oriented architectures—a method of structuring applications as a series of connected components that can be distributed and run on multiple servers, or assembled in different ways to efficiently share workload—forever changed traffic patterns within the data center. While SOA turned what were once fairly monolithic applications into a set of highly agile, flexible, and scalable services, it also dramatically increased the number of servers. This overprovisioning of servers led to a dramatic underutilization of these critical—and expensive—resources; in fact, according to some CIO estimates, 60 to 80 percent of an enterprise's total computing power is wasted due to underutilized server capacity.

To address this underutilization, the concepts of virtualization (which defines the hardware and software needed to host an application as a logical machine image that can then be allocated to one of many compatible physical servers in a data center) and cloud computing (an extension of virtualization principles that weaves data centers into large pools of dynamically shared resources) were introduced.

With virtualization and cloud computing, which introduced new “server-to-server” or “server-to-storage” (referred to as “east-west”) traffic flows into the data center, existing servers could be logically grouped into a single, large pool of available resources. Aggregating the capacity of these devices into a single pool of available storage enabled much more efficient utilization—so much so that the total number of servers could be reduced dramatically without impacting performance—resulting in a related reduction in both capital and operational expenses.


### Old Architecture Can't Support the New Requirements

Unfortunately, networking technology hasn't kept pace with the other innovations in the data center. LAN technology, and Ethernet in particular, have traditionally provided the connectivity required to support the growing number of applications used in business operations. However, Ethernet switches don't have an infinite number of ports or infinite backplane capacity; therefore, to accommodate universal connectivity and growth, enterprises were forced to build hierarchical “tree structure” data center networks with switches arranged in layers to provide the necessary port densities, creating a “north-south” traffic pattern. Traffic between ports in the same rack could be carried by a single switch, but as the number of racks grew, more switching layers were involved in every transaction.

This network model has a distinct deficiency: the need for multiple *switch transits* to achieve universal connectivity. Every time a data packet for an application passes through a switch, it must be read in, decoded and routed, then written out. Each such passage imposes additional latency and subjects the data packet to random collisions with other traffic that can cause even more delays or even packet loss. Network architects have sought to limit these problems by ensuring that the output connections for switches—the “trunk”—were ten times the speed of the port or input connections. That favorable output ratio was maintained by closely monitoring utilization and adding additional capacity when traffic levels crossed the 30% threshold. Higher utilization increased the risk of collisions, which in turn increased latency, delay variability or jitter, and dropped packets. The risk of performance problems multiply with each switch transit and with each layer of the switch hierarchy.

Multiple layers also increase network and traffic management requirements. Since network performance varies considerably depending on where traffic originates and terminates, it becomes more important—and more difficult—to group ports that frequently communicate with one another on the same or on nearby switches. Traffic management also becomes more difficult; as the number of layers grows, it becomes increasingly difficult to maintain the 10:1 uplink vs. downlink speed ratio, which in turn makes it difficult to engineer the network so that it performs consistently under varying loads.

Figure 1 illustrates this problem. In a typical three-layer switch hierarchy used in some large data centers, traffic between any two switch ports can pass through a varying number of switches, depending on which two ports are involved. If two switches have the same “parent” in the hierarchy, then only three switches are involved in the connection. As the two ports become more separated, however, the number of switches grows. Each increase in the number of transits adds to the total and potential delay, not to mention adding the risk of congestion and packet loss. In other words, in a tree hierarchy, packets must first travel north-south before they can move east-west—a highly inefficient model.



**Figure 1: Legacy tree structures impose varying levels of delay when delivering packets across the network.**

The evolution of the data center network has largely been driven by the data center’s mission, which is connecting users to applications. While this mission is essentially responsible for the current hierarchy-of-switches architecture, it has also helped hide that architecture’s deficiencies. That’s because user-to-application traffic is normally low speed and tolerant of both delay and packet loss; modern protocols recover well from dropped packets, and user “think time” often camouflages network delays. In other words, as long as traffic tolerates “best-effort” service, the old hierarchical switch architecture works.

However, some vertical markets such as finance are now suffering from the limitations of the multilayer switch hierarchies. Trading systems are the most demanding of all modern IT applications; they—and other SOA-based applications—will not tolerate delay or packet loss because a dropped packet or tardy transaction could cost thousands of dollars, for themselves and for their customers.

### **New Traffic Patterns Drive New Network Designs**

SOAs, virtualization, and cloud computing have had such an impact on data center traffic patterns that today, approximately 75% of it can be characterized as east-west. In order to optimize the performance of data center traffic, the traditional tree structure must be abandoned in favor of an any-to-any storage/server mesh that can eliminate the need for traffic to travel north-south before it can flow east-west.

While most enterprises still use specialized storage switches like Fiber Channel to link servers with storage devices, there is growing interest in consolidating storage networks with the rest of the data center network to realize efficiencies of scale and reduce the number of siloed networks that need to be maintained and managed. The current trends toward virtualization and cloud computing are likely to accelerate this shift, and this means that storage traffic will be added to user/application traffic in the data center.

Application components and pool resource management processes also need to be connected in order to coordinate their interactions. This type of traffic is “invisible” to the user in that it’s not something client systems and their users are directly involved with. However, the way this traffic is handled may be a major factor in the user’s experience with key applications, which means it can impact productivity.

### Three Classes of Traffic in the Data Center

Three types of traffic classes—groups of applications with similar service quality and availability requirements—are typically found in today’s data center. The first is the original *user/application traffic*, and it is still growing. The second type, *storage traffic*, is exploding as networked storage use grows. The third type, *interprocess communications (IPC) traffic*, is driven by SOA componentization and resource pool management for virtualization and cloud computing.

Each of these three traffic classes has varying degrees of tolerance for delay, delay variability (jitter), and packet loss. Businesses are increasingly employing all three types of traffic in applications like business intelligence and data mining to help them make intelligent, informed decisions. These applications often require analysis of large amounts of data, and if this data handling is impacted by packet loss and delay, the transactions can take so long as to make them useless to decision makers.

### Mathematics of Complexity

The variability of potential traffic paths, combined with the multiplicity of traffic classes, poses significant challenges for network operations with respect to planning and managing hierarchical data center networks. Generally speaking, network complexity grows disproportionately as the number of devices in the network increases, so deep hierarchies create bewildering complexity levels very quickly

The complexity in managing the network is not just a function of the number of devices to be managed, but also of the number of device interactions between those devices. There is an accelerating increase in the number of interactions that occur between switches as a network grows, making it necessary to not only manage the switches in the network, but also to manage the number of interactions between those switches.

Networking protocols like Spanning Tree Protocol (STP), link aggregation groups (LAGs), routing, and security rely on effective interactions between switches, and these grow more numerous with each added device. This can be expressed in the following formula, where  $i$  is the number of potential user configured interactions, and  $n$  is the number of devices.

$$i = n \times \frac{(n - 1)}{2}$$

For example, 10 switches can generate 45 interactions. With 100 switches, however, the number of potential interactions grows to nearly 4,950, while 1,000 switches can produce a possible upper limit of 499,500 interactions.

The best way to address complexity, therefore, is to reduce or even eliminate the number of interactions altogether by getting the data center network to behave as a single device. In other words, the goal is to reduce  $n$  to 1, with no interactions between devices.

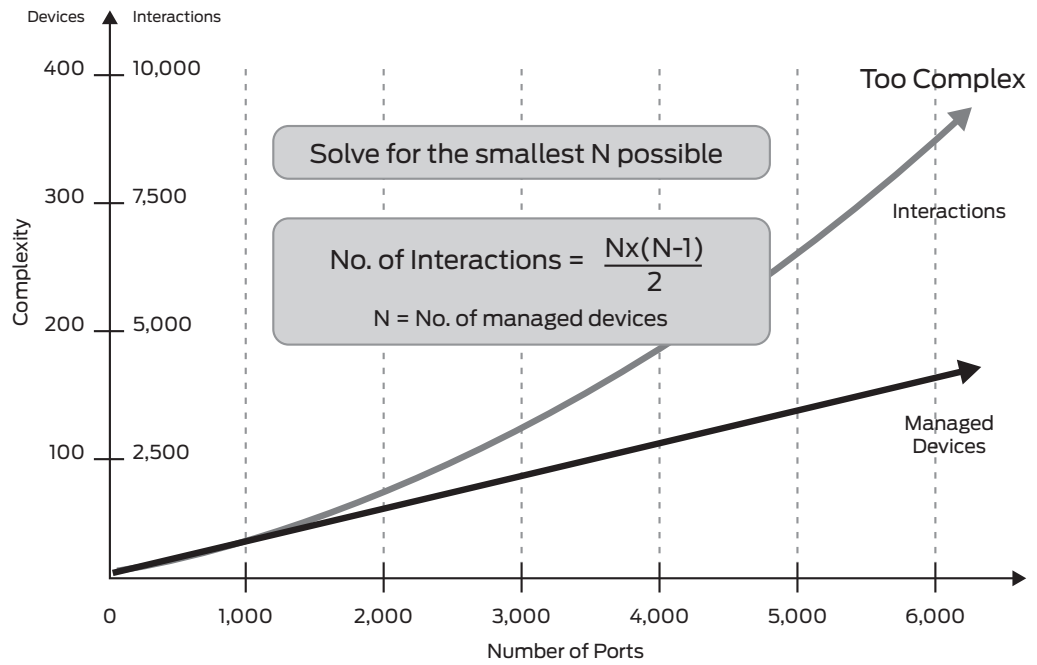


Figure 2: The number of potential interactions grows disproportionately with each new managed device.

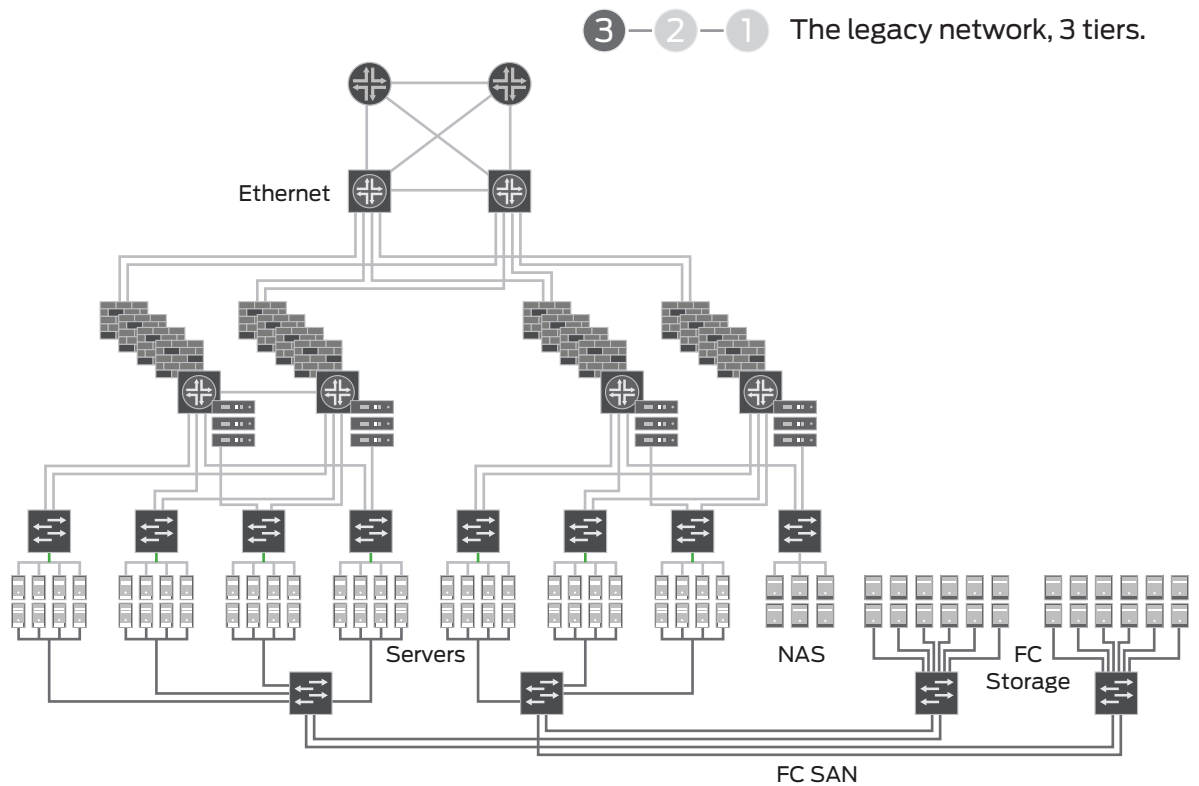


Figure 3: Traditional multitier data center network environment

## The New Network Requirements

### Simplification

Enterprises report that errors and problems related to application setup, device changes, scheduled maintenance, and failure response account for as much as one-third of application downtime. Manual provisioning of data center networks is growing impossibly complex, but the tools used to automate the provisioning process are themselves subject to errors—caused by the very network complexity these tools are designed to support. It's virtually impossible to foresee all of the possible failure modes of a three- or four-tier network structure, and to assess how traffic flows in each switch might impact packet delay and loss. Complexity threatens to make the management of application performance a guessing game.

Virtualization and cloud computing are placing even greater demands on today's data center network. To be truly effective, these high-performance technologies demand the lowest possible network latency and number of packets. Therefore, the data center network must evolve to meet these new requirements or risk becoming part of the problem rather than part of the solution.

### Lowest Possible Latency, Highest Possible Bandwidth, Lowest Possible Congestion

No matter how good or how fast an individual data center switch may be, a multitiered hierarchical collection of those switches will impose complexity and performance problems.

*The logical solution to the data center switching problem is to achieve the concept of “universal connectivity” more directly, without introducing a deep switch hierarchy.*

This approach, called “flattening” the network, reduces the number of layers and weaves the remaining components into a common “fabric” that provides any port with dependable, high capacity connectivity to any other port, from any location.

Such a fabric technology enables multiple networking devices such as switches to be operated and managed as a single, logical device. By fundamentally reducing the number of networked devices to manage, fabric technologies dramatically reduce the cost and complexity associated with large data center networks, while improving performance and efficiency.

By reducing switching layers, a fabric will also reduce the number of switches that traffic has to pass through, eliminating the primary sources of delay and packet loss present in today's hierarchical networks. This is the single most important step an enterprise can take to improve network performance.

Since a fabric is also managed and provisioned as a single device, it vastly reduces the operational complexity of the data center network, accelerates application and user provisioning, and reduces common errors associated with switch configuration management that can impact application availability.



## The New Network Requires a Scalable Network Fabric

### What is a Network Fabric?

Creating a common network fabric requires interconnecting the individual, independent devices that currently populate the data center so that they “collectively emulate” the behavior of a single, logical device that reflects the following characteristics:

- a) **Any-to-any connectivity:** A true fabric must provide a direct connection between any two ports.
- b) **Flat:** A fabric must support single step/lookup-based processing.
- c) **Simplified management:** A fabric must present itself as a single entity with a single point of management.
- d) **Single, consistent state:** Regardless of its various components, a fabric must appear on the outside as a single, logical device with a single, consistent state.

The fabric concept is nothing new; switches have always had a fabric in the form of a backplane that connects all of the ports on the switch. The key is to extend that per-switch fabric to achieve full connectivity, deterministic performance, and unified operations throughout the entire data center. To accomplish this, it is essential that the devices that comprise a fabric behave more like a single logical device than a set of switches connected in the traditional hierarchical model.


### Juniper Delivers Fabric for Networks

Juniper Networks has developed a comprehensive data center fabric strategy that achieves this goal—a blueprint for success that has helped some of the world’s most demanding data centers reduce complexity by eliminating switch layers, lowering power consumption, and minimizing up-front investments. Juniper’s fabric technology fundamentally changes the dynamics of networking in the data center, leading the evolution to a virtualized cloud computing-dominated framework for the enterprise.

In fact, Juniper has been delivering products based on network fabric technology for more than five years. The Juniper Networks TX Matrix Plus family, for instance, supports network fabrics composed of up to 16 interconnected Juniper Networks T Series Core Routers. Two years ago, Juniper extended the same fabric concept to its Juniper Networks EX Series Ethernet Switches with Virtual Chassis technology, allowing up to 10 Juniper Networks EX4200 Ethernet Switches to be interconnected over a 128 Gbps backplane to create a single, logical device. Both the TX Matrix Plus and Virtual Chassis technology provide inter-device connectivity through an ultrafast backplane extender with orders of magnitude greater capacity than is available with traditional Ethernet switch-to-switch connectivity. These Juniper technologies implement fabrics today that effectively reduce the number of layers by allowing more interconnectivity per layer. Latency and packet loss are reduced, and the high inter-switch capacity makes it much easier to manage traffic by reducing inter-switch trunk utilization.

### Extending Fabrics Across Data Centers

Juniper solutions deployed as a fabric can flatten a data center network while improving performance and lowering costs. In this example (Figure 4), a typically complex multitier switching hierarchy for a data center has been replaced by a flattened and simplified fabric based on a combination of Juniper Networks MX Series 3D Universal Edge Routers and EX Series switches, dramatically reducing the number of devices and simplifying overall network maintenance.

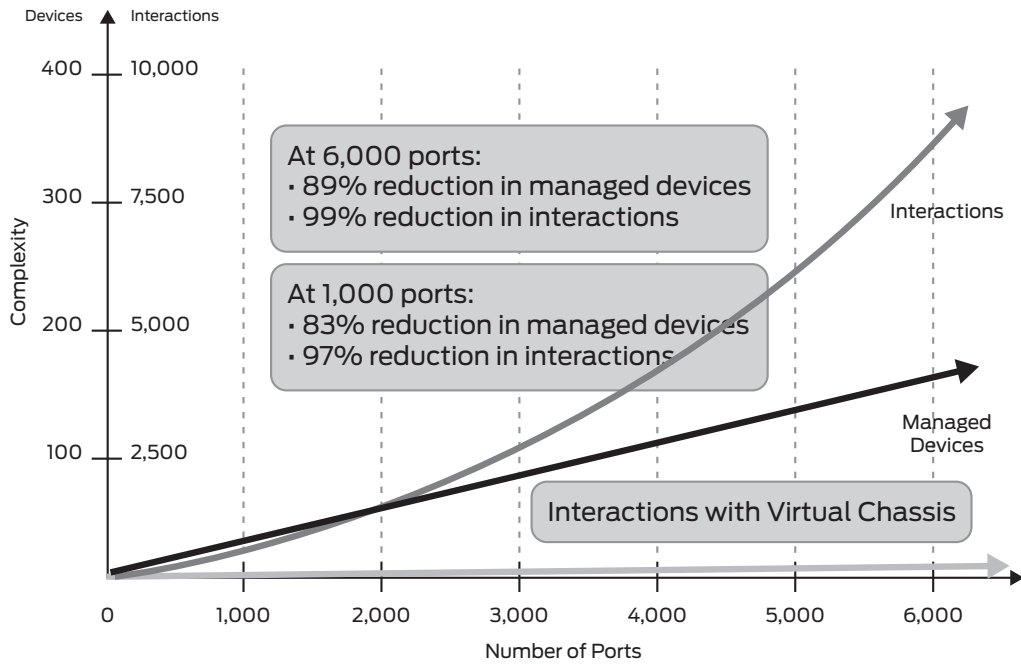


**Figure 4: Juniper solutions deployed as a fabric dramatically simplify the data center network, reducing the number of devices, improving performance, and reducing costs.**

Inside these Virtual Chassis configurations, the connectivity map is maintained without the need for STP, and all three classes of traffic introduced earlier—*user/application traffic*, *storage traffic*, and *interprocess communications (IPC) traffic*—can be carried simultaneously with performance guarantees based on their respective sensitivity to loss and delay. In addition, backup failover routes can be established to ensure that all traffic classes have connectivity during maintenance, upgrades, device diagnosis and redeployments, and network failures.


In the example shown here, a change from a three-tier to a two-tier hierarchy reduces the worst case path from five switch transits to three. If four switch tiers were collapsed onto two, the worst case path would be reduced from seven transits to three—less than half the original number.

Figure 5 illustrates how a fabric based on Juniper’s Virtual Chassis technology can fundamentally reduce complexity in the data center.



**Figure 5: Virtual Chassis technology reduces complexity associated with interactions between disparate switches in the data center network.**

In the near future, Juniper will deliver enhanced fabric technologies that reduce two-layer networks to just one by delivering a single, logical switch that will provide what evolving data centers need: the lowest possible latency, greatest possible bandwidth, and lowest possible congestion (see Figure 6).



**Figure 6: Juniper Networks’ vision for the ultimate simplification of the data center is Project Stratus, which will deliver a single fabric that unites Ethernet, Fibre Channel, and Infiniband networks.**

## Single Operating System, Single Orchestration Platform

Juniper also provides operations support with its 3-2-1 data center network architecture through the integration of its fabric technologies and its single Juniper Networks Junos® operating system. This support also includes Juniper's orchestration model with a new network automation platform, Juniper Networks Junos Space.

Junos Space is open, extensible network automation software that makes it easier to manage and administer the data center by simplifying repetitive and complex tasks, defining and implementing policies within the network, and orchestrating implementation across multiple systems using network-based software. This greatly reduces operational expenses by reducing configuration errors, measurably improves reliability, and frees up labor resources to innovate rather than administer.

The Junos Space network application platform provides end-to-end visibility and control of the network to enable network resources to be orchestrated in response to business needs. Operators can significantly simplify the network life cycle, including configuration, provisioning and troubleshooting with an open, automation platform.

Junos Space includes a core set of collaborative applications from Juniper and third parties that help managers improve operational efficiencies, rapidly scale their infrastructure, and increase the reliability and agility of their network.

## Conclusion

The IT world is changing, and that alone will drive mandatory changes in the data center. Juniper has invested in fabric technology innovations and automation tools to address this emerging data center/network dynamic and can meet the demanding requirements today and in the future with simplicity, performance, and control.

Juniper's fabric approach to the data center not only addresses the structural limitations of hierarchical switching as it relates to application performance, but it also dovetails with the emerging practices in virtualization and cloud computing—all of which visualize the data center network as a virtual device.

Simplicity delivered by fabric and a 3-2-1 data center network architecture is the focus of Juniper's strategy in the data center network and beyond—a model that begins in the data center and enables the cloud.

## About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. Additional information can be found at [www.juniper.net](http://www.juniper.net).

### Corporate and Sales Headquarters

Juniper Networks, Inc.  
1194 North Mathilda Avenue  
Sunnyvale, CA 94089 USA  
Phone: 888.JUNIPER (888.586.4737)  
or 408.745.2000  
Fax: 408.745.2100  
[www.juniper.net](http://www.juniper.net)

### APAC Headquarters

Juniper Networks (Hong Kong)  
26/F, Cityplaza One  
1111 King's Road  
Taikoo Shing, Hong Kong  
Phone: 852.2332.3636  
Fax: 852.2574.7803

### EMEA Headquarters

Juniper Networks Ireland  
Airside Business Park  
Swords, County Dublin, Ireland  
Phone: 35.31.8903.600  
EMEA Sales: 00800.4586.4737  
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

Copyright 2010 Juniper Networks, Inc. All rights reserved. Juniper Networks, the Juniper Networks logo, Junos, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. Junos is a trademark of Juniper Networks, Inc. All other trademarks, service marks, registered marks, 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.