

INTRODUCTION TO BIG DATA: INFRASTRUCTURE AND NETWORKING CONSIDERATIONS

Leveraging Hadoop-Based Big Data Architectures for a
Scalable, High-Performance Analytics Platform

Table of Contents

Executive Summary	3
Introduction	3
What Is Big Data	4
Big Data Use Cases	4
Big Data Technologies (Hadoop)	4
What Does Big Data Mean to IT?	5
Understanding Data Traffic Flows	5
A Scalable Data Infrastructure	6
IBM and Big Data	8
A Scalable Switching Fabric Network Infrastructure	8
Considerations When Implementing Big Data Solutions	10
Conclusion	10
Additional Resources	10
About Juniper Networks	11

List of Figures

Figure 1: High level structure of Hadoop deployment	5
Figure 2: Granular look at Hadoop node architecture	7
Figure 3: Juniper Networks switching solution	9

Executive Summary

Big data is certainly one of the biggest buzz phrases in IT today. Combined with virtualization and cloud computing, big data is a technological capability that will force data centers to significantly transform and evolve within the next five years. Similar to virtualization, big data infrastructure is unique and can create an architectural upheaval in the way systems, storage, and software infrastructure are connected and managed. Unlike previous business analytics solutions, the real-time capability of new big data solutions can provide mission critical business intelligence that can change the shape and speed of enterprise decision making forever. Hence, the way in which IT infrastructure is connected and distributed warrants a fresh and critical analysis.

This paper provides a general overview of traditional Hadoop architectures designed to deliver high-performance and scalable big data analytics. It is intended to provide a basis of understanding for interested data center architects and as a starting point for a deeper implementation engagement. This document assumes little to no background in big data or horizontally scaled query infrastructure, but rather it represents a starting point for the big data journey. Links to additional resources are provided at the end of this paper that help provide logical steps in this journey. Ultimately, the goal is to empower IT to enable an infrastructure that can provide immediate and deep business intelligence for decision making and agility. Every journey begins with a single step, and this document is the first step in recognizing the value of the big data process.

Introduction

Numerous technological innovations are driving the dramatic increase in data and data gathering. This is why big data has become a recent area of strategic investment for IT organizations. For example, the rise of mobile users has increased enterprise aggregation of user statistics—geographic, sensor, capability, data—that can, if properly synthesized and analyzed, provide extremely powerful business intelligence. In addition, the increased use of sensors for everything from traffic patterns, purchasing behaviors, and real-time inventory management is a primary example of the massive increase in data. Much of this data is gathered in real time and provides a unique and powerful opportunity if it can be analyzed and acted upon quickly. Machine-to-machine interchange is another often unrecognized source of big data. The rise of security information management (SIM) and the Security Information and Event Management (SIEM) industry is at the heart of gathering, analyzing, and proactively responding to event data from active machine log files. At the heart of this trend is the ability to capture, analyze, and respond to data and data trends in real time.

Although it may be clear that new technologies and new forms of personal communication are driving the big data trend, consider that the global Internet population grew by 6.5% from 2010 to 2011 and now represents over two billion people. This may seem large, but it suggests that the vast majority of the world's population has yet to connect. While it may be that we never reach 100% of the world's population online (due to resource constraints, cost of goods, and limits to material flexibility), increasingly those that are online are more connected than ever. Just a few years ago, it was realistic to think that many had a desktop (perhaps at work) and maybe a laptop at their disposal. However, today we also may have a connected smartphone and even a tablet computing device. So, of today's two billion connected people, many are connected for the vast majority of their waking hours, every second generating data:

- In 2011 alone, mankind created over 1.2 trillion GB of data.
- Data volumes are expected to grow 50 times by 2020.
- Google receives over 2,000,000 search queries every minute.
- 72 hours of video are added to YouTube every minute.
- There are 217 new mobile Internet users every minute.
- Twitter users send over 100,000 tweets every minute (that's over 140 million per day).
- Companies, brands, and organizations receive 34,000 "likes" on social networks every minute.

International Data Corporation (IDC) predicts that the market for big data technology and services will reach \$16.9 billion by 2015 with 40% growth over the prediction horizon. Not only will this technology and services spend directly impact big data technology providers for related SQL database technologies, Hadoop or MapReduce file systems, and related software and analytics software solutions, but it also will impact new server, storage, and networking infrastructure that is specifically designed to leverage and optimize the new analytical solutions.

What Is Big Data

Big data refers to the collection and subsequent analysis of any significantly large collection of data that may contain hidden insights or intelligence (user data, sensor data, machine data). When analyzed properly, big data can deliver new business insights, open new markets, and create competitive advantages. Compared to the structured data in business applications, big data (according to IBM) consists of the following three major attributes:

- **Variety**—Extends beyond structured data and includes semi-structure or unstructured data of all varieties, such as text, audio, video, click streams, log files, and more.
- **Volume**—Comes in one size: large. Organizations are awash with data, easily amassing hundreds of terabytes and petabytes of information.
- **Velocity**—Sometimes must be analyzed in real time as it is streamed to an organization to maximize the data's business value.

Big Data Use Cases

There are many examples of big data use cases in virtually every industry imaginable. Some businesses have been more receptive of the technologies and faster to integrate big data analytics into their everyday business than others. It is evident that organizations embracing this technology not only will see significant first-mover advantages but will be considerably more agile and cutting edge in the solutions and adaptability of their offerings.

Use case examples of big data solutions include:

- Financial services providers are adopting big data analytics infrastructure to improve their analysis of customers to help determine eligibility for equity capital, insurance, mortgage, or credit.
- Airlines and trucking companies are using big data to track fuel consumption and traffic patterns across their fleets in real time to improve efficiencies and save costs.
- Healthcare providers are managing and sharing patient electronic health records from multiple sources—imagery, treatments, and demographics—and across multiple practitioners. In addition, pharmaceutical companies and regulatory agencies are creating big data solutions to track drug efficacy and provide more efficient and shorter drug development processes.
- Telecommunications and utilities are using big data solutions to analyze user behaviors and demand patterns for a better and more efficient power grid. They are also storing and analyzing environmental sensor data to provide insight into infrastructure weaknesses and provide better risk management intelligence.
- Media and entertainment companies are utilizing big data infrastructure to assist with decision making around customer lifecycle retention and predictive analysis of their user base, and to provide more focused marketing and customer analytics.

There are productized use cases and concrete examples of big data for every industry and company size. Therefore, whether or not your business currently is using a big data solution, your competitors likely are. The real question is how can you better optimize your environment to create a faster, more efficient solution that gives you a competitive edge? Why is this so pressing? According to research by McKinsey Global Institute (MGI) and reported by McKinsey's Business Technology, analyzing large data sets will become—and has already become for a large number of businesses—a key planning tool. With the caveat that the correct policies and enablers must be considered and implemented, big data will become a critical tool in developing plans for:

- Competitive planning and research
- Future productivity and product growth
- Product and services innovation
- Customer satisfaction (or as delineated in the study, "Consumer Surplus")

Big Data Technologies (Hadoop)

The driving force behind an implementation of big data is the software—both infrastructure and analytics. Primary in the infrastructure is Hadoop. Hadoop is the big data management software infrastructure used to distribute, catalog, manage, and query data across multiple, horizontally scaled server nodes. Yahoo! created it based on an open source implementation of the data query infrastructure (originated at Google) called MapReduce. It has a number of commercially supported distributions from companies such as MapR Technologies and Cloudera. Hadoop is a framework for processing, storing, and analyzing massive amounts of distributed unstructured data. As a distributed file storage subsystem, Hadoop Distributed File System (HDFS) was designed to handle petabytes and exabytes of data distributed over multiple nodes in parallel. Figure 1 illustrates an overview of Hadoop's deployment in a big data analytics environment.

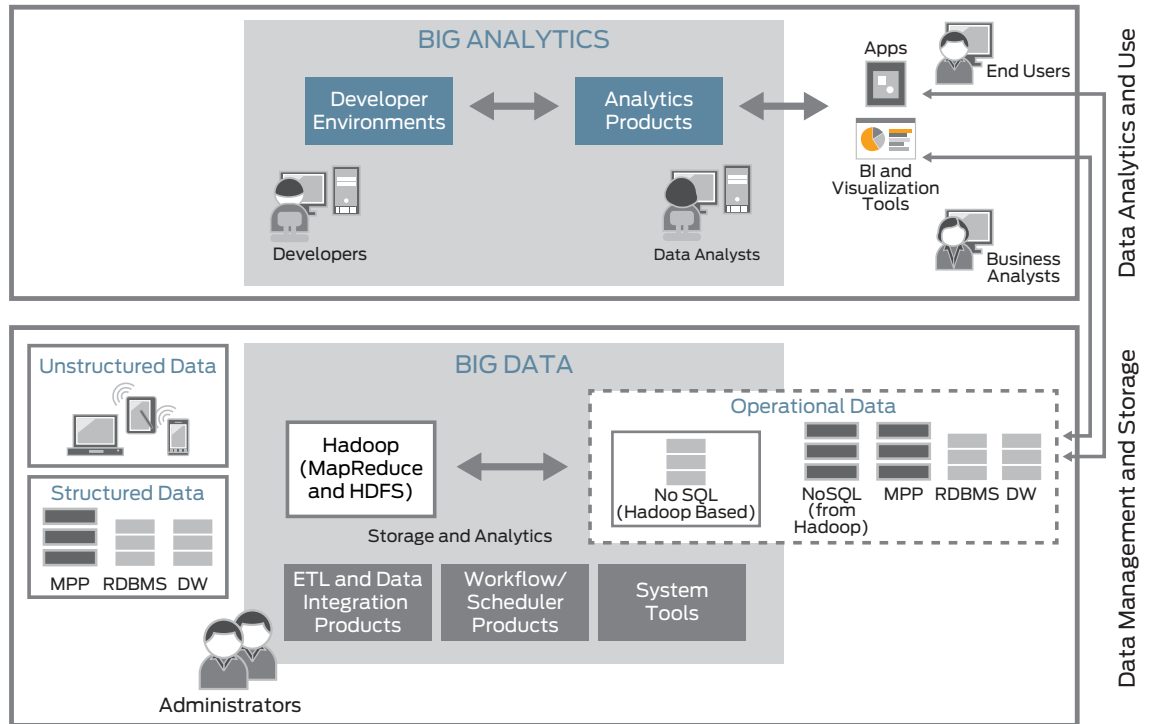


Figure 1: High level structure of Hadoop deployment

What Does Big Data Mean to IT?

Driven by a combination of technology innovations, maturing open source software, commodity hardware, ubiquitous social networking, and pervasive mobile devices, the rise of big data has created an inflection point making real-time data collection and analysis mission critical for businesses today. However, given that the data and its structures are fundamentally different, it is increasingly evident that the infrastructure, tools, and architectures to support real-time analysis and insight from this data also must be different.

As an IT solution, big data mirrors the growth in both content and data source, as well as the pervasiveness of technology in our every day lives. As more and more of what we do is both connected to and often empowered by a network—and the devices that we connect to are themselves powered by an array of sensors—we should expect that the ongoing stream of data will grow. Within data centers, every node (servers, storage, and applications) generates a tremendous number of log files and isolated data streams that also can be collected, collated, and analyzed. With storage costs dropping, the cost associated with saving and leveraging even the most mundane data becomes a nonissue.

Understanding Data Traffic Flows

Within the last 20 years, data center infrastructure has been designed in a manner that closely aligns data, applications, and end users to provide secure, high-performance access. These silos have become rather commonplace, and network administrators can safely assume that the biggest consumer of an application and its corresponding data is an intelligent endpoint that can provide dedicated resources for compilation, execution, and display. This infrastructure has often been referred to as a three-tier architecture. The computing, storage, and networking to support this tiered architecture is largely optimized to deliver data and corresponding network traffic up and down the integrated stack to an end user and back to the database or storage (often referred to as north-south traffic).

During the last few years, this predominant traffic pattern has changed dramatically. Big data represents just the latest application environment to drive this architectural shift. As data becomes more horizontally scaled and distributed throughout network nodes, traffic between server and storage nodes has become significantly greater than between servers and end users. Interestingly enough, the data itself can be generated by servers, applications, or storage environments as opposed to an external source (system log files, for example). This machine-to-machine network traffic and data sharing is often referred to as east-west traffic. Building a data center optimized to provide high-speed connections optimized for east-west traffic is critical in developing scalable, high-performance big data implementations.

A Scalable Data Infrastructure

Another unique characteristic of big data is that, unlike large data sets that have historically been stored and analyzed, often through data warehousing, big data is made up of discretely small, incremental data elements with real-time additions or modifications. It does not work well in traditional, online transaction processing (OLTP) data stores or with traditional SQL analysis tools. Big data requires a flat, horizontally scalable database, often with unique query tools that work in real time with actual data (as opposed to time delineated snapshots). Table 1 compares traditional data with big data.

Table 1: Big Data Versus Traditional Data Types

Components	Traditional Data	Big Data
Architecture	Centralized	Distributed
Data volume	Terabytes	Petabytes to exabytes
Data type	Structured or transactional	Unstructured or semi-structured
Data relationships	Known relationship	Complex/unknown relationships
Data model	Fixed schema	Schema-less

An entire cottage industry is driving big data implementations. There are new forms of supporting databases, some utilizing traditional SQL for queries (often called NewSQL), and some that have largely abandoned SQL for new queries libraries (often called NoSQL). Other businesses looking to leverage their large SQL infrastructure have sharded their existing database infrastructure to create more flexible, horizontally scaled environments to leverage big data tools and capabilities. This has added complexity and created some distinct decision points for IT organizations in planning for their big data implementations. In this same vein, the impact on high-performance, real-time data analytics has a profound impact on many aspects of data center technologies and architectures. We address some of these technologies and architectures later in this document.

The Hadoop Cluster

Hadoop, which includes a distributed file system known as Hadoop Distributed File System (HDFS) and MapReduce, is a critical big data technology that provides a scalable file system infrastructure and allows for the horizontal scale of data for quick query, access, and data management.

At its most basic level, a Hadoop implementation creates four unique node types for cataloging, tracking, and managing data throughout the infrastructure: data node, client node, name node, and job tracker. For more information on Hadoop and related technologies along with implementation best practices, please see the Conclusion for additional resources. The capabilities of these four types are generally as follows:

- **Data node**—The data nodes are the repositories for the data, and consist of multiple smaller database infrastructures that are horizontally scaled across compute and storage resources through the infrastructure. Larger big data repositories will have numerous data nodes. The critical architectural concern is that unlike traditional database infrastructure, these data nodes have no necessary requirement for locality to clients, analytics, or other business intelligence.
- **Client**—The client represents the user interface to the big data implementation and query engine. The client could be a server or PC with a traditional user interface.
- **Name node**—The name node is the equivalent of the address router for the big data implementation. This node maintains the index and location of every data node.
- **Job tracker**—The job tracker represents the software job tracking mechanism to distribute and aggregate search queries across multiple nodes for ultimate client analysis.

Within each data node, there may exist several tens of server or data storage elements and its own switching tier connecting each storage element with the overall Hadoop cluster. Big data infrastructure purposely breaks the data into horizontally scaled nodes, which naturally adds latencies across nodes. This is important as locality and network hops represent potential latencies in the architecture.

For example, Figure 2 shows a granular representation of a more sophisticated Hadoop big data implementation that illustrates a typical architecture for the individual data nodes in a cluster. Not only is it particularly more resource and potentially management intensive than a simple diagram may indicate, but from a pure performance perspective, the additional hops between client, job tracker, and individual data nodes are more significant. The job tracker is required to keep track of the five hops associated with top-of-rack switch 3 in order to access data with top-of-rack switch 8. These individual hops also represent latencies and potential performance bottlenecks.

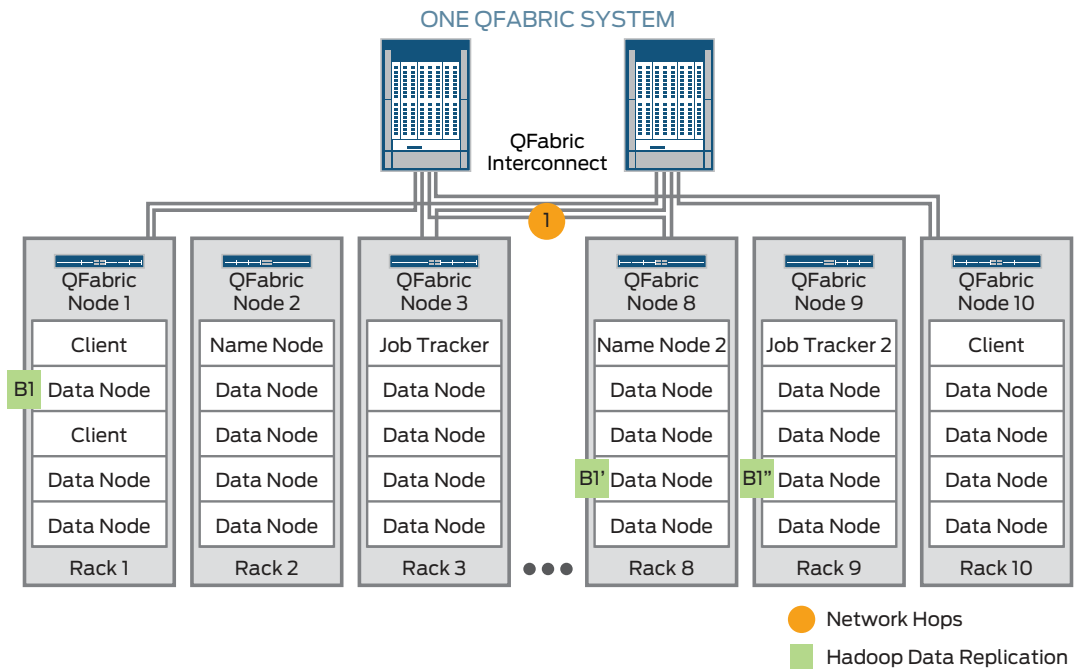


Figure 2: Granular look at Hadoop node architecture

In addition, Figure 2 highlights the directional flow of network traffic in the big data infrastructure. Quite prevalent in modern cloud-based environments, big data generates far more east-west (server-to-server or server-to-storage) network traffic than north-south (server-to-client or server-to-outside) network traffic.

Why is this important? Distributed data of today's big data and cloud architectures places a tremendous burden on connecting nodes rather than connecting clients. For every one-client interaction, there may be hundreds or thousands of server and data node interactions. This is counter to the original client/server network architectures built over the last 20 years. Those architectures assumed that the client, rather than the backend infrastructure, supplied much of the computational overhead. With that in mind, network administrators should consider distributed systems such as Hadoop.

Hadoop and most cloud computing infrastructures today run as a cluster and handle huge data sets, which are distributed over multiple nodes. Hadoop clusters run tasks in parallel and scale-out fashion. Although Hadoop is agnostic to network infrastructure, it places the following requirements on the network:

Data locality—The data shuffle-and-sort operation between the distributed Hadoop nodes running parallel jobs causes east-west network traffic that can be adversely affected by suboptimal network connectivity. The network has to provide high bandwidth, low latency, and any-to-any connectivity between the nodes for optimal Hadoop performance.

Scale-out—Deployments might start with a small cluster and then scale out over time as the customer realizes initial success and then needs change. The underlying network architecture also should scale seamlessly with Hadoop clusters and provide predictable performance.

Increased east-west traffic—As previously described, traffic patterns range from 1-to-1, 1-to-many, many-to-1, and many-to-many. These flows are demanded by a combination of unicast/multicast flows between multiple Hadoop nodes that run in parallel. This requires a high bandwidth, low latency network infrastructure for efficient communication between Hadoop nodes.

As a result, it is critical in the network architecture to prioritize locality, high-performance horizontal scalability, and direct server node to server node connectivity. In addition, with the latencies prevalent in tiered node hops, it is evident that a new network architecture is required for high-performance scale.

IBM and Big Data

IBM is promoting another big data solution that represents an alternative to the Hadoop architecture. Realizing IBM's unique position and legacy, it is important to consider its solution carefully.

To support big data processing, IBM brings enterprise capabilities such as high availability, performance, and ease-of-use to the Hadoop ecosystem through IBM InfoSphere BigInsights. BigInsights extends Hadoop by eliminating potential single points of failure events in the Hadoop name node. Many customers consider this extension a mission critical feature. IBM's InfoSphere Streams solution is a comprehensive analytics computing platform for big data. From a packaging perspective, this appeals to many customers looking for a one-stop shop. IBM provides the integration, services, and end-to-end components to get the infrastructure online with relative ease.

A Scalable Switching Fabric Network Infrastructure

One of the biggest evolutions in the networking industry in the last few years is the introduction of point-to-point switching fabric network solutions. The best way to describe a network switching fabric and its benefits is to refer to Wikipedia's definition:

“**Switched fabric, switching fabric, or just fabric**, is a network topology where network nodes connect with each other via one or more network switches (particularly via crossbar switches, hence the name). The term is popular in telecommunication, Fibre Channel storage area networks, and other high-speed networks. The term is in contrast to a broadcast medium such as early forms of Ethernet. Switched fabrics can offer better total throughput than broadcast networks because traffic is spread across multiple physical links.

“A switched fabric should be able to function as a simple point-to-point interconnect and also scale to handle thousands of nodes. The term fabric derives its name from its topological representation. As the data paths between the nodes of a fabric are drawn out, the lines cross so densely that the topology map is analogous to a cloth.

“The advantage of the switched fabric is typically one of overall system bandwidth and performance versus connectivity between individual devices.”¹

Understanding the benefits of a data center fabric is critical in creating the highest performance/lowest latency connections between big data nodes. The benefit of switch fabric infrastructure compared to more traditional tree architecture is obvious and startling. First, since fabric architectures create a point-to-point connection between nodes with a single hop in the switching infrastructure, inherently this architecture will significantly reduce latencies between nodes. Second, for those looking to provide a seamless high-performance interconnect with the easiest switching management, virtualizing the switching fabric is also an essential feature, as it allows multiple networking components to behave as a single component.

Flat network scalability also is a major benefit of implementing a data center fabric. The inherent virtual domains and point-to-point connections allow a company to seamlessly merge new data sources into the cluster without rewiring the entire data center. This benefit significantly eases the expansion of the system and allows for rapid enhancement of the analytic environment.

Whether you choose the cutting edge of Juniper Networks® QFabric™ family of products or the more traditional Juniper Networks EX Series Ethernet Switches with Virtual Chassis technology, Juniper's innovative data center networking solutions provide ideal network infrastructure for big data deployments. These solutions deliver, among other things, location independent simplicity, scalability and high performance.

The QFabric family provides a unique, cutting edge data center fabric technology ideal for big data deployments. The QFX Series allows for location independence, scalability, performance, convergence, and large resource pools:

- **Location independence**—Allows big data clusters and applications to be placed anywhere in the data center and still achieve optimum performance.
- **Scalability**—Offers unprecedented scalability, from tens to 6,144 10GbE ports, all managed as a single switch.
- **Performance**—Provides any-to-any connectivity within a data center, increases bandwidth, and deterministic low latency provides substantial improvement in performance (40 Tbps capacity) compared with traditional network architectures. (In the **Executive Summary of STAC-M2™ Test Report**, the QFabric System, combined with the IBM server and middleware with Mellanox ConnectX®-2 EN Ethernet network interface cards, deliver deterministic performance similar to single switches enabling transaction times of 10 µsec under high load, and display very linear and deterministic performance despite spanning three network devices.)

¹ Switched Fabric (Wikipedia), http://en.wikipedia.org/wiki/Switched_fabric, accessed August 27, 2012.

- **Convergence**—Allows Hadoop cluster, storage area network (SAN), and network access server (NAS) to easily communicate across the fabric with converged access support for Fibre Channel/Fibre Channel over Ethernet (FC/FCoE).
- **Operational simplicity**—Helps to manage a group of resources with one switch abstraction dynamically and easily to apply policies across the entire data center.

For more information on the dramatic benefits that switch fabrics and, more specifically, Juniper Networks QFabric System can provide to data center deployments, refer to the final section of this document for links to a number of additional white papers that provide both the top down benefits and specifics of the solution.

While network fabrics solve all of the problems associated with east-west traffic, data locality, and network latencies, along with providing overall ease of management, some customers may simply be more inclined to stick with tiered networking. To accommodate these types of customers, Juniper offers the EX Series platform that provides a more traditional networking topology along with Virtual Chassis capabilities to provide significantly reduced management overhead and overall implementation simplicity.

Juniper's unique EX Series with Virtual Chassis technology provides flexible 1GbE and 10GbE mixed infrastructure for big data deployments:

- **Location independence**—Keeps inter-rack traffic local within the access layer (up to 10 racks) and delivers east-west traffic that outperforms traditional stacking solutions.
- **Scalability**—Offers scalability from tens to thousands of 1GbE/10GbE ports, and can scale across multiple data center locations.
- **Performance**—Provides unprecedented core switch capacity with up to eight core switches in a single Virtual Chassis configuration (10.24Tbps capacity). Virtual Chassis technology in the core also eliminates Spanning Tree Protocol (STP) and increases bandwidth.
- **Flexibility**—Offers seamless, nondisruptive, plug-and-play migration from 1GbE to 10GbE data center networks.
- **Operational simplicity**—Reduces significantly the number of managed devices.

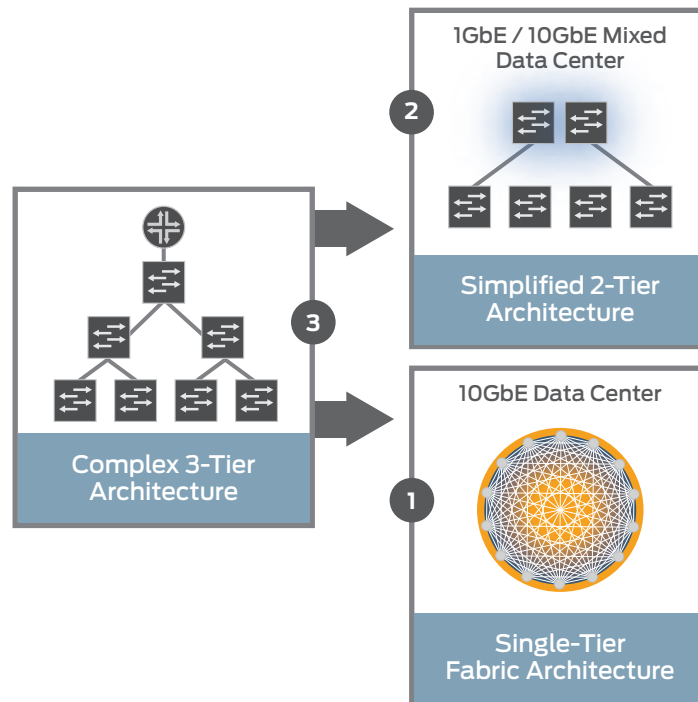


Figure 3: Juniper Networks switching solution

Considerations When Implementing Big Data Solutions

Big data represents a tremendous opportunity for businesses of all sizes to capture and analyze huge amounts of data in a manner never available to them before. As IT organizations begin to test and evolve these solutions, it is critical for network administrators to consider the impact of these technologies on their server, storage, networking, and operations infrastructure.

The outstanding question for business and IT organizations is how to best develop this new infrastructure to leverage and analyze the increasing flow of big data? How can data centers optimize for and build true next-generation infrastructures to support the increasing data demands? What are the critical technologies and what impact will they have on existing infrastructures, and more importantly, on data center networks?

In order to help you through this transition, consider the following questions carefully to provide the priorities and guidance in developing the proper infrastructure and networking topology to optimize investments and outcomes:

- What business insights are you trying to achieve?
- Are the data sources being used differently or in the same way as existing production data sources?
- If the pilot is successful, how big will the cluster need to be?
- How important is it to easily and quickly add more capacity?
- Will big data applications require access to and integration with the other applications in your data center?

The answers to these questions will help you frame the discussion with internal constituents as well as outside hardware and software vendors. As suggested, these questions also will likely indicate how a (or your) big data infrastructure will influence traditional data center architectures and interconnect requirements both between nodes and between racks. Similar to new storage discussions, big data is pushing the envelope of networking requirements, and it is forcing many new strategies to provide the kind of real-time business analytics and higher levels of infrastructure agility needed to react to new technologies, as well as new business insights.

Conclusion

With careful planning and predetermined expectations, creating an optimized big data deployment is relatively straightforward. Keep in mind only three or four years ago, broad commercial appeal for big data implementations was not a key requirement in data center design. However, you should be developing this infrastructure with an eye toward the horizon. Adopting technologies and underlying infrastructure, including networking that can provide the scale, performance, and headroom for tomorrow's technologies is critical in offering the highest levels of investment protection, business agility, and time to market.

Additional Resources

For a better understanding of how a packet-switched networking topology differs from traditional switching architectures and the fundamental benefits of switch fabric technologies, refer to *Defining Characteristics of QFabric* by Pradeep Sindhu at www.juniper.net/us/en/local/pdf/whitepapers/2000384-en.pdf.

For further information on big data implementations and networking considerations, refer to *Understanding Big Data and the QFabric System: QFabric System Enables a High-Performance, Scalable Big Data Infrastructure with Simplicity* at www.juniper.net/us/en/local/pdf/whitepapers/2000483-en.pdf.

For a better understanding of the impact of switching fabric and overcoming the inherent limitations of traditional networking architectures, including security limitations, be sure to review the CIO Playbook: *The Role of the Network Has Changed. Are You Ready?* at www.juniper.net/datacenter.

About Juniper Networks

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at 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

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 2012 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. 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.

2000488-001-EN Sept 2012

 Printed on recycled paper