



White Paper

THE BIG DATA REVOLUTION

Is Your Network Ready?

May 2012

SPIRENT

1325 Borregas Avenue
Sunnyvale, CA 94089 USA

Email: sales@spirent.com
Web: www.spirent.com

AMERICAS 1-800-SPIRENT • +1-818-676-2683 • sales@spirent.com

EUROPE AND THE MIDDLE EAST +44 (0) 1293 767979 • emeainfo@spirent.com

ASIA AND THE PACIFIC +86-10-8518-2539 • salesasia@spirent.com

© 2012 Spirent. All Rights Reserved.

All of the company names and/or brand names and/or product names referred to in this document, in particular, the name “Spirent” and its logo device, are either registered trademarks or trademarks of Spirent plc and its subsidiaries, pending registration in accordance with relevant national laws. All other registered trademarks or trademarks are the property of their respective owners.

The information contained in this document is subject to change without notice and does not represent a commitment on the part of Spirent. The information in this document is believed to be accurate and reliable; however, Spirent assumes no responsibility or liability for any errors or inaccuracies that may appear in the document.

The Big Data Revolution

Is Your Network Ready?

CONTENTS

- Executive Summary 1
- Background 2
- Big Data is Proliferating 4
- Understanding Network Challenges with Big Data 6
- Business Challenges 8
- Network Testing for Big Data 9
- Summary 11

EXECUTIVE SUMMARY

Big data has emerged as one of the most visible topics in technology today. As with other trends such as mobile, social and cloud computing, big data represents both business opportunities and technical challenges. Organizations of all types are extracting increased value from large data sets that come from real-time data streams, video, mobile Internet, and many other sources. However, handling this data is not easy.

Big data must not only be stored. It must be transmitted, processed and reused by a variety of applications. This means big data impacts everything in the entire IT stack, including servers, storage and networking as well as operating systems, middleware and applications. In order to meet the requirements of big data, a number of different approaches are used: Architects are replacing hard drives with solid state drives; operations teams are scaling infrastructure up and out; and computer scientists are developing new algorithms to process data in parallel. However, even with all these approaches, the network often fails to get the attention it needs.

To keep up with the demands of big data, the network must be properly designed, often using a number of advanced techniques. Network engineers may work around the bandwidth limitations of legacy networks by using new 2-tier fabric designs. They may optimize packet forwarding using software defined networking technologies like OpenFlow. They may also enhance security by introducing next-generation, application-aware firewall and threat prevention solutions. Yet, even with a solid design, the complexity of these networks requires careful testing and validation.

Networks used for big data demand testing at scale with realism. Test teams must select tools and processes that enable testing at the magnitudes encountered with big data: millions of concurrent users, billions of packets-per-second, tens of thousands of connections-per-second. They must also select methodologies such as PASS testing which helps ensure proper testing across the dimensions of performance, availability, security and scale. Products and services from Spirent are used by many of the world's most successful companies to enable realistic testing at scale of big data applications.

BACKGROUND

Big data is a relative term. Rather than representing a specific amount of data, big data refers to data sets that are so large that they introduce problems not typically encountered with smaller data sets. Depending on the type of data, as well as storage, processing and access requirements, challenging problems can arise with just a few terabytes or many petabytes of data. Table 1 is a summary of the most common data quantities along with rough guidelines on which amounts of data are likely to introduce big data challenges.

The 2011 IDC Digital Universe Study estimated that 1.8 zettabytes of data would be created and replicated during the year. With a single zettabyte representing a trillion gigabytes of data, the industry is already handling staggering amounts of information. By 2015, IDC estimates this number could climb to about 8 zettabytes. With exponentially increasing quantities of data throughout the world, the demand for handling big data is also growing exponentially.

These are just some of the many scenarios in which big data is regularly encountered today:

Table 1. Common Data Quantities and Challenges

Name	Symbol	Value	Challenge
kilobyte	KB	10 ³	None
megabyte	MB	10 ⁶	None
gigabyte	GB	10 ⁹	Minimal
terabyte	TB	10 ¹²	Moderate
petabyte	PB	10 ¹⁵	Acute
exabyte	EB	10 ¹⁸	Extreme
zettabyte	ZB	10 ²¹	Extreme

- **Real-time audits** of credit card data and customer transactions for fraud
- **Data mining** of customer databases and social media interactions
- **Business intelligence** derived from customer behaviors and spending patterns
- **Video** storage, streaming and archival
- **Mobile Internet** support for billions of mobile users
- **Scientific research** through modeling and number crunching
- **Predictive analytics** for spotting emerging risks and opportunities

While Internet companies such as Amazon, Facebook, Google and Yahoo pioneered many of the approaches to big data used today, organizations of all types must also now apply them. Businesses, government agencies, universities, science labs and others are encountering big-data-related requests on a daily basis. However, keep in mind that big data is not simply defined by data volume.

IT research and advisory company Gartner has defined three important dimensions of big data. They are:

1. **Volume:** Measuring not only bytes of data but also quantities of tables, records, files and transactions
2. **Variety:** Introducing data attributes such as numeric, text, graphic and video as well as structured, unstructured and semi-structured
3. **Velocity:** Suggesting variation for processing time requirements and including categories such as batch, real time and stream processing

When dealing with large magnitudes, each of these dimensions on its own represents potential difficulties. Making matters worse, high measures of volume, variety and velocity can happen concurrently. When massive amounts of data in countless formats must be rapidly processed, big data challenges are at their most extreme. Amazon provides an outstanding example of this scenario.

Amazon has been pushing the limits of big data for several years with its Amazon Web Services (AWS) offerings. The AWS object storage service called Simple Storage Service (S3) already holds a volume of 905 billion objects and is currently growing by a billion objects a day. In terms of variety, objects in S3 may contain any type of data and object sizes can range from 1 byte to 5 terabytes. Data velocity is also enormous at 650,000 requests per second for the objects, a figure that Amazon says increases substantially during peak times.

BIG DATA IS PROLIFERATING

By the end of this decade, IDC expects the number of virtual and physical servers worldwide to grow by 10 times. Yet, during that same time period, the amount of information managed by enterprise datacenters should grow by 50 times and the number of files handled by datacenters by 75 times. These growth rates along with increasing demands for rapid processing of data are leading to previously unheard of challenges for all kinds of organizations, not just Internet companies.

Addressing the Basic Technical Challenges

When data sets are massive, existing equipment, tools and processes often fall short. In order to meet the demands of big data, the entire IT stack—from applications to infrastructure—must be properly designed and implemented.

Applications, databases and middleware must all support the requirements of big data. Since relational databases have limitations on the number of rows, columns and tables, the data must be spread across multiple databases using a process called sharding. Applications and middleware also have scale and performance limitations, typically requiring a distributed architecture to support big data. Even operating systems have inherent limitations. Getting beyond 4 gigabytes of addressable RAM requires a 64-bit operating system.

Datacenter infrastructure, including servers, storage and network devices, presents another set of obstacles. Infrastructure challenges are typically addressed by using one of two approaches—scale up or scale out.

Scale Up: Infrastructure is scaled up by using equipment with higher performance and capacity. For servers, this means more CPUs, cores and memory. For storage it could mean higher density platters that spin faster, or perhaps solid state drives (SSD). For networks, it means fatter pipes such as 40Gbps Ethernet rather than 10Gbps. Unfortunately, the ability to scale up for all these types of infrastructure is always limited by the infrastructure components with the highest performance and capacity.

Scale Out: Infrastructure is scaled out by using more equipment in parallel, with little regard to the performance and scale attributes of the individual devices. When one server can't handle all the processing, more servers are added. When one storage array can't handle storing all the data, more arrays are added. Networks can also be scaled out to some degree. However, this requires more discussion in part because networks are better viewed as an interconnected fabric rather than a set of independent components.

While the scale-up approach works for some applications, it introduces physical limitations and becomes cost prohibitive on a processing-per-dollar basis. When it comes to big data, scale-out servers and storage combined with distributed software architectures forms the dominant design pattern. Frameworks based on this pattern as well as additional technologies are becoming widely available and helping organizations of all kinds adapt to the zettabyte world.

Hadoop

The Apache Software Foundation has a framework called Hadoop which is designed for running applications on a large cluster built of commodity hardware. The Hadoop framework transparently provides both reliability and data motion to applications. It implements a computational paradigm called MapReduce, where an application is divided into many small fragments of work, each of which may be executed or re-executed on any node in the cluster.

The Hadoop framework also provides a distributed file system called Hadoop Distributed File System (HDFS) that stores data on the compute nodes, providing very high aggregate bandwidth across the cluster. The Lustre™ file system is also commonly used with Hadoop and scales to tens of thousands of nodes and petabytes of storage. MapReduce, HDFS and Lustre are designed so that node failures are automatically handled by the framework.

Companies from Adobe to Yahoo use Hadoop to process big data in a variety of ways. Adobe uses Hadoop in several areas including social media services and structured data storage and processing. At Yahoo, the Hadoop production environment spans more than 42,000 nodes and is used behind the scenes to support Web search and advertising systems. The popularity of Hadoop makes it useful for considering some of the deeper challenges encountered when processing big data, particularly related to the network.

UNDERSTANDING NETWORK CHALLENGES WITH BIG DATA

Scale-out infrastructure, distributed software architectures, and big data frameworks like Hadoop go a long way toward overcoming the new challenges created by big data. Yet—in order to achieve the greatest scale and performance—compute, network and storage need to function together seamlessly. Unfortunately, the network does not always receive sufficient attention and can become the weakest link, making big data applications significantly under-perform.

The Network is Different

For optimal processing, big data applications rely on server and storage hardware that is completely dedicated to the task. Within a big data processing environment, the network is generally dedicated as well. However, networks extend well beyond these boundaries. As raw data is moved into the processing environment it can consume the bandwidth of the extended network, impacting other business-critical applications and services.

The Extended Network

Moving data into a big data processing environment is not always a one-time activity. Rather than relying on a fixed repository, some big data applications consume continuous streams of data in real time. Mobile Internet, social media, and credit card transactions are all examples of never-ending data streams. Similarly, there may be multiple big data applications pulling huge amounts of raw data from a shared repository. Transferring data between multiple processing environments and distributing intermediate or final results to other applications also place a heavy burden on the extended network.

The Dedicated Network

Even when the network within the processing environment is dedicated to a single big data application, it can be pushed to its limits. For example, a Hadoop application is spread across an array of servers, so getting raw data to the right servers, transferring intermediate results, and storing final data is no easy task. Traditional datacenter network architectures may work well for more common three-tier applications. However, the any-to-any data transfer requirements of Hadoop and other big data applications are not easily met.

Part of the problem with any-to-any communication is that the Layer 2 data plane is susceptible to frame proliferation. This means switches must use an algorithm to avoid creating data loops, an occurrence which would impact the entire bridged domain or processing network. The Spanning Tree Protocol is typically used to prevent loops. However, this loop-free restriction prevents Layer 2 from taking full advantage of the available bandwidth in the network. It may also create suboptimal paths between hosts over the network.

Network Options

On the surface it may appear that a scale-out approach for the network may provide a solution. Unfortunately it isn't that simple. Adding more connections—including faster connections—is not only expensive, but it still might not work. While the Spanning Tree Protocol prevents loops, it also tends to increase latency between some network nodes by forcing multi-hop paths.

What Hadoop and other big data applications really need in order to meet demand is a well-designed network with full any-to-any capacity, low latency, and high bandwidth. Further, since Hadoop only scales in proportion to the number of networked compute resources, the network must sustain the inclusion of additional servers in an incremental fashion. A number of new networking approaches are becoming available to meet the emerging requirements of big data. Here are a few:

- **2-tier Leaf/Spine.** One approach to designing a scalable datacenter fabric that meets these requirements is called a 2-tier Leaf/Spine fabric. This design uses two kinds of switches—one that connects servers and another that connect switches. Leaf switches are used to connect servers and Spine switches are used to connect the Leaf switches.
- **TRILL and SPB.** Transparent Interconnect of Lots of Links (TRILL) and Shortest Path Bridging (SPB) create a more robust Layer 2 topology by eliminating Spanning Tree while supporting both multipath forwarding and localized failure resolution.
- **OpenFlow.** This is a communications protocol that gives access to the forwarding plane of a network switch allowing paths for packets to be determined by centrally running software. OpenFlow is one approach to software defined networking (SDN).

These technologies can be used in various combinations along with high density 10G, 40G and 100G Ethernet. Care must also be taken to account for virtual switches on servers where CPU cycles are used to forward traffic without new hardware offload technologies as edge virtual bridging (EVB) virtual port aggregator (VEPA).

BUSINESS CHALLENGES

As with other technologies, big data creates a number of business and economic challenges. Businesses, government agencies and other organizations are not working with big data just for fun. They are working to produce valuable information from raw data which raises the issue of return on investment (ROI). When it comes to big data, as well as other IT challenges, the goal is rarely to maximize performance, scale or other technical attributes. More often than not, the goal is to optimize a number of variables, including cost and time. This is where big data gets even more difficult.

The cost of producing information or intelligence from raw data must not exceed its value. In fact the cost must be significantly less than the value in order to produce a proper return on investment. Time is also a constraining factor. Insights must be produced in time for them to matter. For example, credit card companies must determine fraudulent activities before accepting a credit card. Otherwise it may simply be too late.

Computer scientists develop algorithms that tend to have three competing characteristics: cost, time and quality. This implies the need for tradeoffs. For example, the highest performing solution may also be expensive and/or take a long time to develop. Due to the need for tradeoffs, business leaders are responsible for picking the two most important characteristics and accepting what that implies to the third.

A specific example arises when designing a network for big data. CLOS and Fat Tree are network designs with different oversubscription ratios. A 1:1 ratio of full-rate bisectional bandwidth between servers requires the most networking hardware and performance. If the business cares most about high-quality results and speed of delivery, then the highest performing CLOS architecture found in some supercomputers can be chosen. The implication is that the business is willing to pay for quality and speed. If instead the business is willing to accept slower processing times, then the less-expensive Fat Tree ratio architecture could be used.

Ultimately, the network must be designed to meet performance requirements without under-provisioning or over-provisioning. However, exact performance demands are not always known in advance. They can also quickly change as data volume, variety, and velocity change. In order to build a network that delivers results based on the right business tradeoffs, network engineers must first design them properly. This requires testing individual components—from virtual switches running on servers to physical switches and routers. It also requires testing the resulting network with realism based on known traffic patterns along with various levels of extreme traffic conditions which simulate the addition of big data to the network.

Network security is another business challenge that should not be overlooked. As an example, the Sony PlayStation network was hacked in 2011. The attack brought down the service for several weeks and exposed personal information from about 77 million user accounts, including the names, addresses, birthdates and e-mail addresses for its users. As more data is aggregated and the number of access points for that data grows exponentially, there is a higher risk for sophisticated attacks. To avoid such a negative business impact, high volumes of traffic through firewalls and IDS/IPS unified thread management (UTM) devices must be handled while simultaneously maintaining big data performance demands.

NETWORK TESTING FOR BIG DATA

Network test teams should use a proven approach to perform network testing for big data. The recommended methodology is PASS testing which is used to validate the four core pillars of performance, availability, security and scalability. After all, if the network has weaknesses in any of these dimensions, then big data applications—and ultimately the business—may be impacted. Network test teams should also keep in mind that PASS testing can and should be applied to network components as well as the end-to-end network. Data produced in testing network components can be fed to network engineers who use it as design input.

The following real-life test scenarios demonstrate each element of the PASS testing methodology at the magnitudes required for big data.

Performance

Independent test lab Network Test recently used Spirent's datacenter testing solution to demonstrate that Juniper Networks' QFabric™ can deliver unprecedented performance at scale. While legacy switch fabrics scale to hundreds of ports, the QFabric solution scales to thousands. A fully loaded QFabric system supports 6,144 ports configurable as a single device using any-to-any port topology with low latency and the appearance of unlimited bandwidth. The test bed for this industry-first test comprised 1,536 x 10 Gbps Ethernet ports connected to the same number of Spirent TestCenter test ports and 128 redundant 40 Gbps fabric uplinks. Because of the full-mesh topology, Spirent TestCenter generated and analyzed throughput and latency over 2.3 million streams in real time and over 15 Tbps of traffic at line rate.

Availability

In an exclusive Clear Choice test using Spirent test equipment, Network World tested Arista Networks' DCS-7508 data center core switch for availability and other attributes of PASS. The 7508 has a highly redundant design including six fabric cards. The time needed to recover from the loss of a fabric card was measured while using Spirent equipment to drive 64-byte unicast frames to all 384 ports and physically removing one of the fabric cards. The Spirent TestCenter calculated frame loss in order to determine that the system covered in 31.84 microseconds.

Security

Another Clear Choice test performed by Network World and using test equipment from Spirent focused on the Palo Alto Networks PA-5060 firewall. Spirent Avalanche was used to measure the performance tradeoffs encountered when using different security capabilities. When the PA-5060 was configured as a firewall, it moved data at around 17Gbps. However, Avalanche determined throughput to be about 5Gbps when also running anti-spyware and antivirus. Impact on performance of adding SSL decryption was also assessed.

The Spirent Mu Dynamics Security and Application Awareness test solution was used to test the advanced capabilities of the PA-5060. Specifically the Studio Security application tested the PA-5060's ability to handle Spirent's extensive published vulnerabilities. Studio Security generated 1,954 vulnerabilities and determined that the PA-5060 blocked 90% of the attacks in the client-to-server direction and 93% of the attacks in the server-to-client direction. Further functional tests examined the firewall's application awareness features, verifying its ability to block applications based on identification rather than port numbers. The PA-5060 successfully blocked 83% of the applications.

Scale

Crossbeam Systems, Inc. offers a proven approach to deploying network security that meets extreme performance, scalability and reliability demands. The company used Spirent Avalanche to effectively test the firewall capability of the Crossbeam X-Series under extraordinarily demanding real-world conditions. It achieved over 106 Gbps of throughput while simultaneously maintaining 270,000 open connections per second and 685,000 transactions per second with tens of milliseconds of impact on Web browser page render time without packet loss. In order to support the growing requirements of big data, security must be maintained while achieving these performance levels.

PASS represents a powerful methodology for testing networks used for big data. At the same time, effective PASS testing must be backed by test equipment that can achieve the extreme magnitudes required by big data applications.

SUMMARY

When it comes to big data, the network often fails to get the attention it needs. To keep up with the demands of big data, the network must be designed properly. Network engineers may work around the bandwidth limitations introduced by the Spanning Tree Protocol by using fabric topologies such as Leaf/Spine. They may optimize packet forwarding using software defined networking technologies like OpenFlow. They may also enhance security by introducing next-generation, application-aware firewall and threat prevention solutions. Yet, even with a solid design, the complexity of networks requires them to be carefully tested and validated.

Networks used for big data call for testing at scale with realism. Test teams must select tools and processes that enable testing at the magnitudes encountered with big data: millions of concurrent users, billions of packets-per-second, tens of thousands of connections-per-second. They must also select methodologies such as PASS testing which helps ensure proper testing across the dimensions of performance, availability, security and scale. Products and services from Spirent are used by many of the world's most successful companies to enable realistic testing at scale of big data applications.

