

10 Gigabit Ethernet: Enabling Storage Networking for Big Data

Hadoop and other big data frameworks introduce a data storage model which is fundamentally different from the NAS or SAN environments many enterprise applications rely on. With any new architecture, new requirements emerge. With years of experience dealing with Hadoop workloads and the network challenges of HDFS and similar file systems, Arista offers the industry's best networking solutions for big data, and IP storage architectures. Only Arista switches have the performance, extensibility, and deep buffering necessary to not only ensure optimal performance under these demanding workloads, but also to provide the visibility necessary to operate and optimize big data, and storage-enabled networks.

Big data requires big networks.

Explosive data growth is a reality and the trajectory is continuing to be strong. In order to accommodate and support this level of intensification, more robust and powerful networks are becoming more important than ever before. Data generation and the diversification of data use drive the adoption of more role-based storage solutions within the data center. These factors, coupled with the transition to highly virtualized data center environments, affects how organizations buy and manage server, storage, and network assets and are key drivers in what is propelling big data into an everyday reality. The outlook is big data in the Cloud.

Data is everywhere, whether it is from users, applications, or machines and it's growing exponentially with no vertical or industry being spared. Due to this reality, IT organizations everywhere are forced to come to grips with storing, managing and extracting value from every piece of it -- as inexpensively as possible. This begins the real race to cloud computing where the framework needs the ability to process data increasingly in real-time and in greater orders of magnitude -- and at a fraction of what it would typically cost.

Application Drivers of Storage Networking

There are many key application drivers related to storage networking. The ones that rise to the top of the list are in the Cloud Computing, Hadoop Cluster, Storage Virtualization, High Performance Computing, and Rich Content/Video arenas. Furthermore, the rise of new web-based application architectures in the data center, the increasing use of virtualization tools to consolidate servers, and the utilization of HPC in core mission-critical applications all drive the need for high performance, low latency, and highly available networks for big data.

The key players driving big data are predominantly Internet service providers, application service providers, storage service providers, and now large enterprises. These players are among those who need 40/100 Gigabit Ethernet (GbE) bandwidth at the core of the network to support applications as well as for remote replication and disk-to-disk backup. Also guaranteed to benefit from 10/40/100GbE are applications such as video editing and rich content applications that consume or generate huge amounts of data in a short time.

Hadoop Clusters

Data analytics has become a key element of the business decision process over the last decade. Classic reporting on a dataset stored in a database was sufficient until recently, but yesterday's data gathering and mining techniques are no longer a match for the amount of unstructured data and the time demands required to make it useful. The common limitations for such analysis are compute and storage resources required to obtain the results in a timely manner. A network that is designed for Hadoop applications, rather than standard enterprise applications, can make a big difference in the performance of the Hadoop cluster. Testing has repeatedly shown that Hadoop clusters deployed in a leaf/spine topology gain performance benefits from having deep buffers at both the leaf and spine layers. Arista Networks switches satisfy Hadoop cluster network requirements and have been successfully deployed in many Hadoop environments. As costs fall and companies think of new ways to correlate and analyze data, big data analytics has become more common. Businesses will especially benefit given their low-cost ability to manage and analyze big data.

Hadoop is a very powerful distributed computational framework that can process a wide range of datasets, from a few gigabytes to petabytes of structured and unstructured data. Use of Hadoop has quickly gained momentum, and Hadoop is now the preferred platform for various scientific computational and business analytics. While availability of commodity Linux based servers makes it feasible to build very large clusters, the network can be a bottleneck, resulting in congestion, dropped traffic, and less efficient use of the cluster.

Hadoop allows enterprises to easily explore complex data using custom analyses tailored to their information and queries. The big data problem is not just all about size of the data; it is also about performance and how fast the data can be processed.

Arista offers high performance 10/40/100GbE non-blocking, ultra-low latency solutions that can scale from a few racks to some of the largest Hadoop deployments of thousands of devices. In addition, Multi-Chassis LAG (MLAG) offers true active/active uplink connectivity from each rack, allowing the full bi-sectional bandwidth of the network to be utilized in a flat layer 2 network for smaller deployments. Arista's Extensible Operating System can easily be integrated with third party and customer developed management tools. Lastly, Arista's networking solutions offer a true flat-line growth when it comes to price for server-interconnect bandwidth. These factors make Arista's networking solutions ideal for any Hadoop deployment.

Cloud Computing

Cloud computing helps organizations store, manage, share, and analyze their big data in a reasonable and simple-to-use way. Today's cloud Infrastructure-as-a-Service (IaaS) providers, supported by on-demand analytics solution vendors, make big data analytics very affordable. As location-independent computing entails shared servers providing resources, software, and data to systems and devices on demand, cloud computing is a very strong use driver for high performance storage networking. Widespread adoption of virtualization and utility computing has caused this natural evolution and has resulted in customers who no longer need expertise or control over the technology infrastructure that supports them. However, in order for those customers to achieve the

best performance and service possible, robust and high performance networking solutions must be in place.

Storage Virtualization

Virtualization is a requirement for cloud computing as virtual nodes only need a subset of the performance of modern CPUs. Storage is now following a similar trajectory. As appealing as this convergence to virtual compute has been, it has also come with some stumbling blocks. For one, it has equated to an over-provisioning of storage and network capacity, specifically via the creation of an excess of Fibre Channel (FC) storage. Software Defined Storage (SDS) and iSCSI over Ethernet are ideal for virtualization due to the fact that they allow administrators to move virtual servers among physical machines without reconfiguring the zoning and logical unit number (LUN) masking in the classical FC storage network. However, the continued use of a dedicated FC SAN reduces the benefits inherent in virtual server deployment by continuing to use expensive dedicated hardware to transport storage traffic. Arista's 10 Gigabit Ethernet solutions ensure the full benefits of storage virtualization are met and realized by providing the high performance network required for successful storage deployment while leveraging a common Ethernet infrastructure.

High Performance Computing (HPC)

High performance computing involves the use of supercomputers or clusters of powerful processors to solve computationally intensive problems. In order for a large number of processors to work together, clusters require interconnects that support high-bandwidth and low latency communication. Arista Networks has the highest performance switches in the market today which are used for HPC cluster interconnects to process large amounts of data in such industries as meteorology, genomics, finance, oil and gas, biological research or any other type of business that needs to crunch large amounts of data. As with big data, in HPC environments, there is often a need for deep buffers in the switches to accommodate the bursts of traffic seen on the network in these extreme environments.

Rich Content / Video

Arista Networks is a key player in delivering a fully redundant, cloud based content origination service that hosts all video and high resolution photo content. Arista Networks provides the Ethernet network to interconnect the streaming servers, which require high-performance 10GbE switching with very large packet buffering. This is perfectly suited to support high-bandwidth video streaming and storage interconnection while maintaining the highest levels of system availability with both device and system level resiliency designed into the rich content / video infrastructure.

Storage Networking Options for Big Data

10 Gigabit Ethernet deployments have been rapidly growing as price and performance targets are met and as new optics enable broader deployments. Additionally, the aggregate growth of new applications continues to increase bandwidth requirements. Although 10 Gigabit Ethernet is the optimal network interface, its role is really within a much larger overall picture of a switching solution. Successful 10 Gigabit Ethernet deployments must incorporate leading intelligent switching services such as high performance, low latency, high availability, and enhanced manageability to provide the necessary support for new applications.

A wide range of storage solutions exist in the market today, utilizing various approaches and a wide range of technologies. 10 Gigabit Ethernet prevails as the mainstream technology for Cloud Storage with iSCSI based block storage and network attached storage (NAS). With non-blocking throughput, record density, low latency, massive buffers, and leading total cost of ownership, Arista Networks switches are ideal for cloud storage applications.

Depending on their access method, storage systems are categorized as Storage Area Networks (SAN) and Network Attached Storage solutions (NAS). In SAN environments, storage devices, although remote, appear as locally attached to the client, and access to storage is block-based. In contrast, in NAS environments, clients access files remotely using a network-based file system.

Storage Area Networks

A Storage Area Network (SAN) is an architecture whereby servers access remote disk blocks across a dedicated Interconnect. Most SANs use the SCSI protocol to communicate between the servers and the disks. Various interconnect technologies can be used, each of them requiring a specific SCSI mapping protocol as shown in table 1.

Table 1: SCSI mapping protocols for SANs

Interconnect Technology Protocol	SCSI Mapping
Fibre Channel (FC) Protocol (FCP)	Fibre Channel
TCP/IP over Ethernet	iSCSI
Ethernet	FCoE
Infiniband	iSER

A SAN is a specialized network that enables fast, reliable access among servers and external or independent storage resources and is the answer to the increasing amount of data that needs to be stored in an enterprise network environment. By implementing a SAN, users can offload storage traffic from daily network operations while establishing a direct connection between storage elements and servers. SAN interconnects tie storage interfaces together into many network configurations and across large distances. Interconnects also link SAN interfaces to SAN Fabrics.

Switched SCSI, FCS, and Switched SSA form the most common legacy SAN fabrics. With gateways, SANs can be extended across WAN networks as well. Switches allow many advantages in building centralized, centrally managed, consolidated storage repositories shared across a number of applications. Building a SAN requires network technologies with high scalability, performance, and reliability in order to marry the robustness and speed of a traditional storage environment with the connectivity of a network. As Ethernet technology has matured, it can now achieve all of these metrics at or better than a traditional Fibre Channel SAN at greatly reduced cost.

Fibre Channel Protocol (FCP)

Today, many SANs still use FCP to map SCSI over a dedicated Fibre Channel (FC). Enterprises deploying Fibre Channel deploy multiple networks including the LAN network, which typically uses Ethernet technology (Ethernet is a basic component of 85% of all networks worldwide, and is one of the most ubiquitous network protocols in existence), as well as the dedicated FC network.

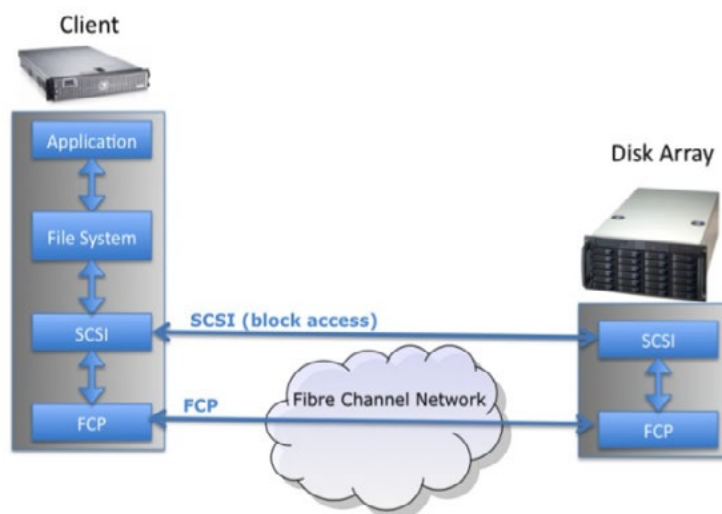


Figure 1: Fibre Channel (FC) Storage Area Network (SAN)

iSCSI

iSCSI has gained traction and attention as data centers look to lower costs for robust storage. iSCSI rides on 1/10/40/100Gigabit Ethernet transport, alleviating the complexity of a separate traditional Fiber Channel SAN. Historically billed as an ideal solution for many small and medium enterprise organizations, iSCSI relies on TCP/IP protocols, making it a natural communication for private and public cloud communications. With the proliferation of 10G server connectivity, and 40/100 Gig connectivity between switches, iSCSI is no longer considered a small to medium enterprise solution. It is now deployed in the most demanding of enterprise environments.

The performance advantages of iSCSI are compelling. Storage arrays must keep up with new multicore processors and stack software that are now capable of generating millions of iSCSI IOPS. New iSCSI arrays with 10 Gbps Ethernet controllers such as Dell Equal Logic's 6100 arrays, EMC VNX storage line, and Netapp's FAS 6000 appliances combined with Arista Networks 7000 switches can offer non-blocking storage access with performance characteristics, which meet and often exceed those of classic Fibre Channel networks.

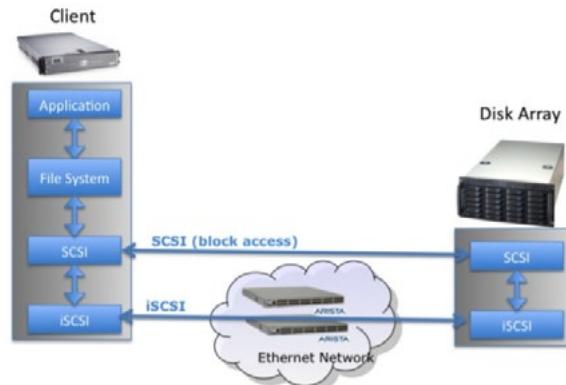


Figure 2: iSCSI Storage Area Network (SAN)

Network Attached Storage (NAS)

An increasingly popular method for consolidating storage resources is Network Attached Storage (NAS). A NAS appliance is a server, which has the purpose of supplying file-based data storage services to other devices on the network. NAS is a remote file system I/O where the file request is redirected over a network.

NAS is recognized for three principal benefits, which in combination lower overall TCO:

- Storage consolidation
- Deployment simplicity
- Ease of management

NAS systems have evolved to support, via a standard Ethernet network, the storage tiering, high performance and high availability that had previously only been available in SANs. This combined with its TCO advantages has made NAS an increasingly adopted solution in the enterprise.

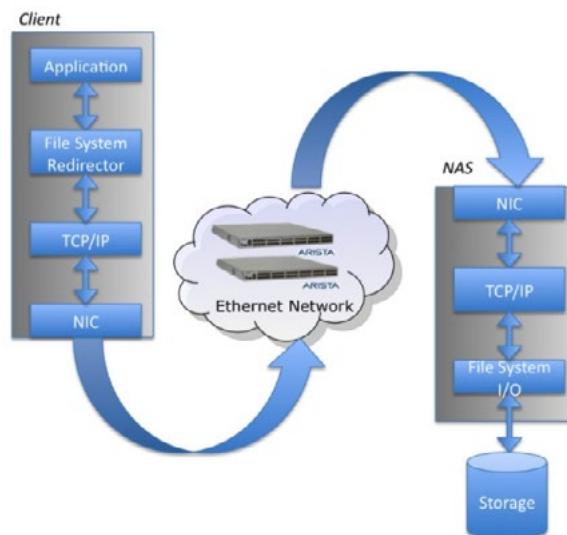


Figure 3: Network Attached Storage (NAS)

Fibre Channel over Ethernet (FCoE)

The FCoE protocol is essentially an encapsulation of FCP over Ethernet. FCoE enables enterprise customers accustomed to Fibre Channel to run the Fibre Channel Protocol directly over their LAN Ethernet network, hence allowing them to consolidate their LAN and Storage network over the same network infrastructure. FCoE is aimed at organizations interested in keeping an FC SAN, yet are interested in LAN and SAN convergence. FCoE has not seen the widespread deployment expected and suffers from the need to be constrained to an L2 network and support for DCB. Although Arista switches support DCB, demand for FCoE has been weak as the same results can be achieved via iSCSI without sacrificing the benefits of an L3 network.

At Arista, we have viewed storage as a natural driver for 10GbE cloud networking. The 7000 Family is optimized for datacenter solutions, providing key storage characteristics such as resilience, high throughput, large buffers for handling loss-less 1/10/40GbE traffic, PFC/ DCBX support and predictable low latency. Arista has designed a compelling 10 Gigabit Ethernet architecture for iSCSI and NAS storage.

Table 2: Cloud Storage Networking vs FCoE

Ethernet Storage Networking	FCoE
Supported by all Ethernet vendors	Brocade and Cisco only
Interoperable across multiple vendor solutions	Brocade and Cisco implementations do not interoperate
Leading cost / performance	Fibre Channel (FC) traditionally expensive. FCoE has similar cost / performance to FC.
Benefit for cost savings of convergence	No cost benefit due to convergence

In summary, iSCSI and NAS are the optimal options for “cloud storage”, as opposed to FC and FCoE, due to the industry standard, multi-vendor support, and cost-effectiveness. Tables 2 and 3 show the strengths of iSCSI and NAS and illustrate why they are the best choice.

Arista Networks Advantages & Differentiating Features

Arista Networks is the leader in building scalable high-performance and low latency networks for today’s data center and cloud computing environments. Purpose-built hardware and the Arista Extensible Operating System provide a single binary system image across all platforms, maximum system uptime, stateful fault repair, in- service upgrades, and a fully accessible Linux shell. Arista switches are the perfect network solution for the most demanding workloads. With support for VMware Virtualization and hundreds of Linux applications integrated into hardware platforms specifically designed to meet the stringent power and cooling requirements of today’s most demanding data centers, Arista delivers the most energy efficient and best performing 10/40/100GbE platforms.

Storage at its very core is simply data. The more data, that can be moved from host to target, the more efficiently the host will process applications. Arista switches have proven to be the highest performance, and top rated 10GbE switches available. The lower the latency between a storage write and the acknowledgement from the storage target, the more efficiently the host can process data and applications. Arista 10GbE switches are among the lowest latency available, ranging from 350ns for 24-ports to 1.2usec in the 1RU 48-port 10GbE switch and 4.5usec at 1152-port density with the Arista 7500E Series.

If the storage and host are connected at different speeds, when storage is consolidated onto a subset of switches, and when deploying hyper-converged architectures, buffering becomes a very important consideration. Arista has a line of switches with extremely large buffers, over 100x what other vendors deliver. The buffers absorb bursty reads/writes and do not drop frames, reducing retransmissions, improving efficiency and application performance.

Adding to the simplicity of Arista and its solutions comes the ability to provision your networks without the manual intervention typically associated with doing so. Arista gives customers the ability to have Zero Touch Provisioning (ZTP) - Network Automation for Cloud Data Centers. Using standards based protocols (e.g. DHCP, T/FTP, HTTP) the network can be rapidly provisioned. Advanced

scripting capabilities allow the administrator to tailor boot configurations based on a variety of parameters, meeting the needs of even the most complex data center deployments. Combined with other Arista features such as VMTracer's adaptive VLAN configuration, data center managers can fully automate the turn-up of network elements and virtual servers. Arista is unique with its 'hands-off' provisioning to enable and automate the Cloud Data Center.

Table 3: Comparison of Storage Networking Technologies

Capability	Ethernet	FCoE
Port Speed	1 Gbps / 10 Gbps 40 Gbps / 100 Gbps (future)	10 Gbps
Layer 3	Yes	No
Switch Latency	~ 1 µsec	3.5 – 10's µsecs
Interoperability with Ethernet- based apps	Seamless	Not operable today
Manageability	Same tools as traditional LAN	Same tools as traditional LAN
IEEE 802.1Q bb/az/au	Capable	Capable
PFC, Enhanced Transmission Selection, QCN	Yes (L2 & L3)	Yes (L2 only)

Summary

Coinciding with the developments of multi-protocol SAN and NAS, Arista Networks has taken concrete steps in extending the value of 10GbE storage network assets in fast and evolving storage data center environments. Arista Networks has made great strides and considerable investments in developing strong storage infrastructure business partners around the globe. These partners focus on large enterprises with advanced data center and storage requirements. With this collection of network based storage services offerings that bring advanced capabilities to new environments, Arista Networks is the best-in-breed choice for 10 GbE deployments.

Table 4: Facts About Cloud Storage for Big Data

Cloud Storage for Big Data

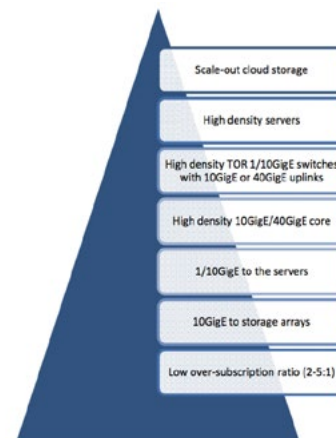
iSCSI and NAS are predominate choices as opposed to FC or FCoE

iSCSI and NAS is where most of the opportunity for 10GbE is today

Cloud storage driving 10GbE in the data center today

Arista has strong partnerships with many cloud storage vendors such as BlueArc, Coraid, Dell EqualLogic, EMC, Isilon (now EMC), NetApp, and Panasas

10/40/100GbE empower companies to expand application capabilities, reduce time to solve complex financial, Arista has strong partnerships with many cloud storage vendors such as BlueArc, Coraid, Dell EqualLogic, EMC, Isilon (now EMC), NetApp, & Panasas cloud, and clustered applications problems, and quickly respond to changing customer needs and market conditions. High bandwidth, low-latency, and energy efficient are the key differentiators of Arista Networks switch offerings, which address all of these drivers. Arista's 7000 family of Ethernet switches is uniquely suited to address the needs of cloud storage, featuring non-blocking performance on every port with sub-microsecond latency and the highest port density per RU on the market for fixed and modular form factors. Arista's switches offer highly resilient Extensible Operating System (EOS) software, featuring self-healing and live-patching capabilities, and are designed for massive and dynamic buffering to deal with speed mismatch or avoid dropped packet or loss.



In summary, Arista's product family is designed for storage data centers - implementing high-performance compute clusters; and enterprise networks - consolidating disparate SAN solutions into a common, low-cost and high-speed IP storage network while revolutionizing the way big data is networked. The time is now, and with Arista you will always have the highest performance and best of breed solution for your big data needs.

Santa Clara—Corporate Headquarters

5453 Great America Parkway,
Santa Clara, CA 95054

Phone: +1-408-547-5500

Fax: +1-408-538-8920

Email: info@arista.com

Ireland—International Headquarters

3130 Atlantic Avenue
Westpark Business Campus
Shannon, Co. Clare
Ireland

Vancouver—R&D Office

9200 Glenlyon Pkwy, Unit 300
Burnaby, British Columbia
Canada V5J 5J8

San Francisco—R&D and Sales Office

1390 Market Street, Suite 800
San Francisco, CA 94102

India—R&D Office

Global Tech Park, Tower A & B, 11th Floor
Marathahalli Outer Ring Road
Devarabeesanahalli Village, Varthur Hobli
Bangalore, India 560103

Singapore—APAC Administrative Office

9 Temasek Boulevard
#29-01, Suntec Tower Two
Singapore 038989

Nashua—R&D Office

10 Tara Boulevard
Nashua, NH 03062



Copyright © 2016 Arista Networks, Inc. All rights reserved. CloudVision, and EOS are registered trademarks and Arista Networks is a trademark of Arista Networks, Inc. All other company names are trademarks of their respective holders. Information in this document is subject to change without notice. Certain features may not yet be available. Arista Networks, Inc. assumes no responsibility for any errors that may appear in this document. 12/14