

Ethernet Everywhere: Accelerating Data-Center Migration to the Cloud

[Electronic Design](#)

[Martin Nuss](#)

Thu, 2014-12-04 15:49

By simplifying storage networking and reducing total cost of ownership, the object-oriented, Ethernet-based architecture should streamline the data-center network and accelerate the shift of applications to the cloud.

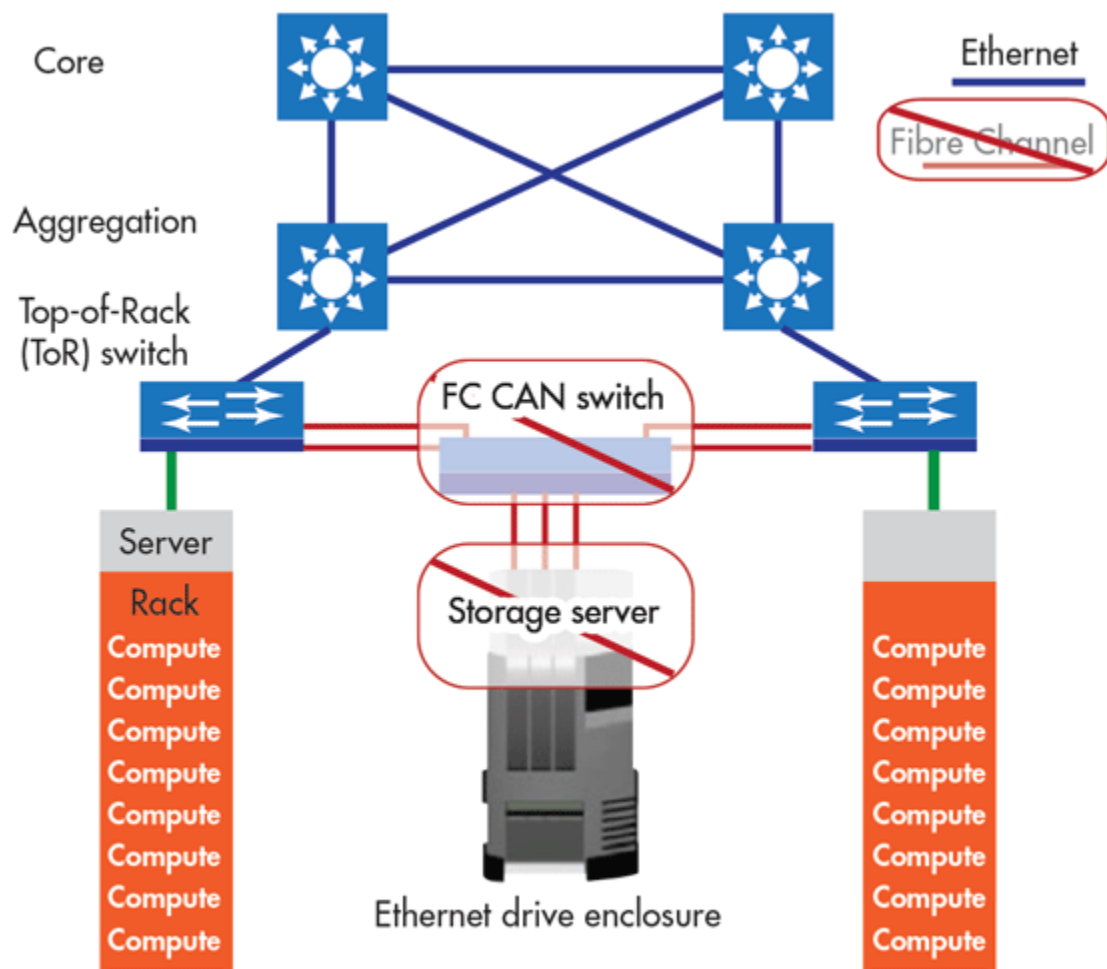
Since its introduction as a local-area-network (LAN) technology over 40 years ago, Ethernet has evolved into the communications backbone for all kinds of networks. This includes enterprise LANs transitioning from semi-proprietary protocols long ago, to the more recent shifts among telecommunications service providers, to packet-based carrier Ethernet for wired and wireless infrastructure. More recently, industrial Internet of Things (IoT) networks are migrating to Ethernet as well.

Historically, storage networks have relied on traditional protocols like SAS/SATA and Fibre Channel (FC), which are ill-suited for networking outside a small deployment. While Fibre Channel over Ethernet (FCoE) and iSCSI permit storage networking over Ethernet, they still require specialized networks and adapters to encapsulate storage before finally mapping to Ethernet.

However, a new Ethernet-based approach to storage networking has emerged, driven primarily by a paradigm shift in the data center toward object-oriented storage. Not only will it reduce costs and simplify operations, but can also potentially accelerate the migration of IT resources and applications to the cloud.

Object-Based Architecture: A New Storage Paradigm

Storage systems traditionally used a file/directory approach, in which data is stored in files, and files in directories, much like a personal computer. Today's data centers, though, now largely rely on a key/value, object-oriented approach. This model uses handles to reference data objects, rather than file/directory location. For example, on a social media website, a user's name would be the value in the "name" key field, age in the "age" key field, etc. With the key/value approach, data objects are written once and read multiple times, but are rarely deleted or even modified.

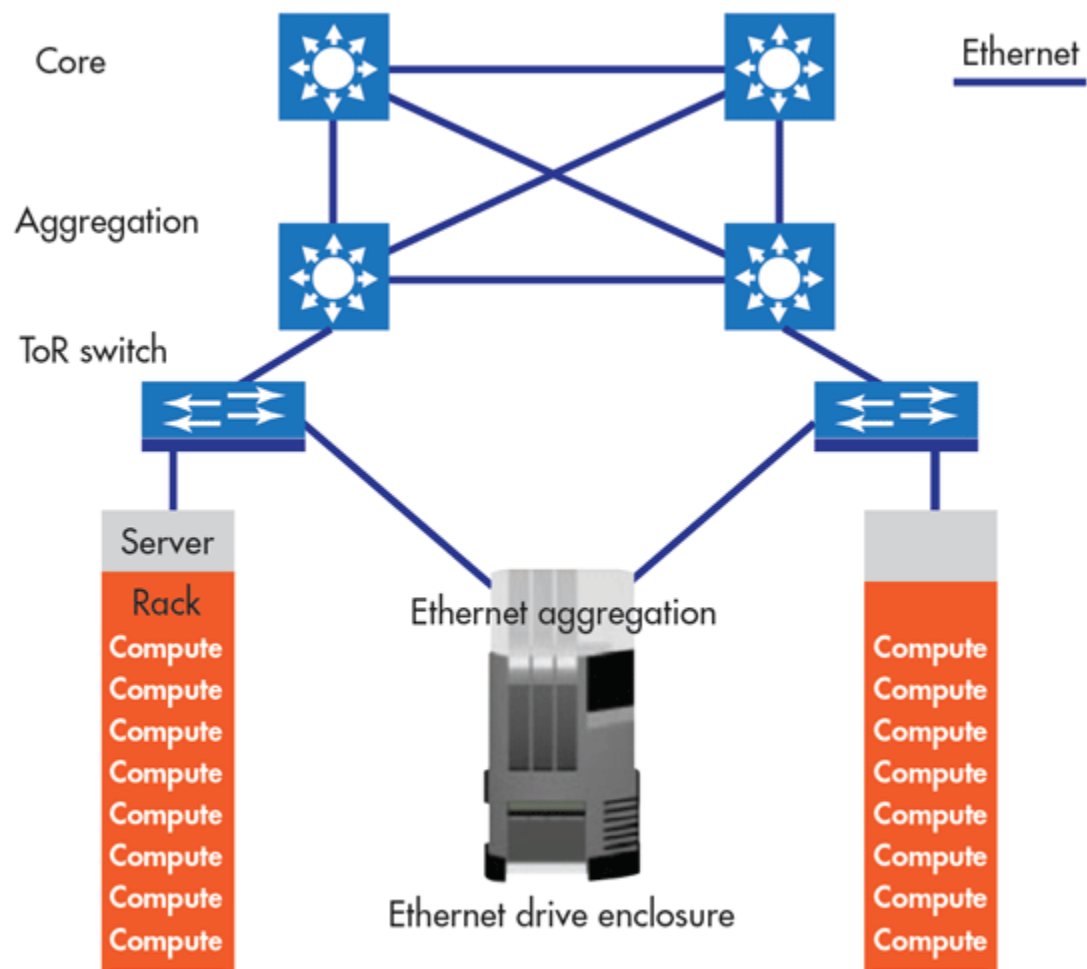


Ethernet-based storage takes into account the object-oriented model used in today's data centers and dramatically simplifies network architecture (*Fig. 1*). How? It obviates the need for file semantics and a file system for space-management handling on a storage device. Applications can communicate directly with storage devices using Ethernet. This eliminates multiple layers of software and hardware, including the FC storage area network (SAN) switch and storage server. Ethernet interfaces replace SAS on storage drives, and similarly, Ethernet switches supplant SAS/SATA expanders within the storage enclosures.

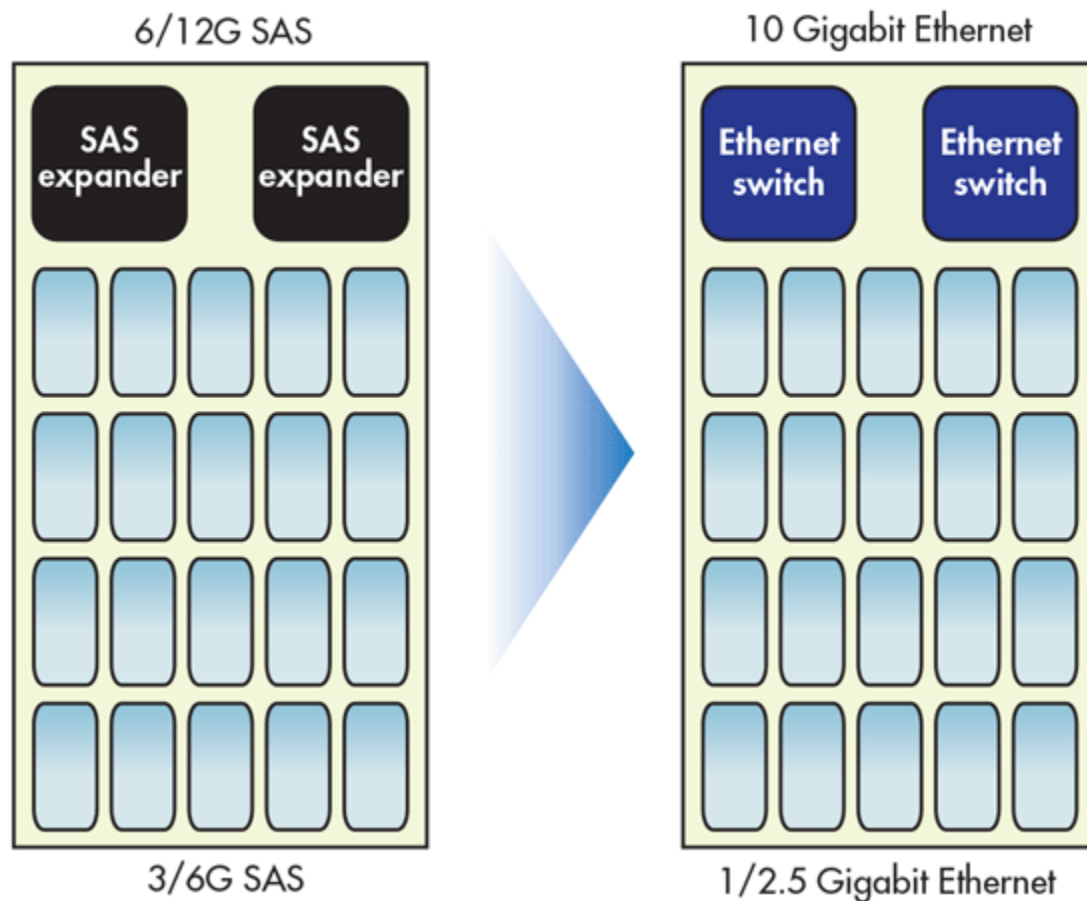
With all communication via Ethernet, neither FC nor FCoE are needed. This model is not simply input/output consolidation, like FCoE, but true network integration. More importantly, this Ethernet-based approach enables data-center design with far greater operational efficiencies.

The Ethernet-Based Storage Rack

A traditional storage rack has a storage server that communicates directly with the Top of Rack (ToR) switch via 10 Gigabit Ethernet and with SAS expanders via a 6G or 12G SAS interface. In turn, these SAS expanders interface with individual storage drives via 3G or 6G SAS.



The new architecture eliminates the storage server altogether and replaces SAS expanders with Ethernet switches supporting this architectural approach (*Fig. 2*). Ethernet switches communicate directly with the ToR switch via 10 Gigabit Ethernet. Storage drives can remain the same as those used traditionally, with the same connectors. However, instead of 3G or 6G SAS for communications, the connectors interface via 1 or 2.5 Gigabit Ethernet (*Fig. 3*).



This significantly reduces total cost of ownership by:

- Dispensing with the storage server.
- Eliminating multiple protocol and processing layers, improving efficiencies.
- Using Ethernet versus SAS for communications. Unlike Ethernet, which is widely deployed and enjoys broad ecosystem support, SAS is a fairly specialized protocol supported by a limited number of manufacturers.

Another benefit is unprecedented flexibility. With Ethernet as the basis, compute and storage devices can be mixed as needed within the same rack. This is easily configurable to support applications that, for example, require significant compute resources alongside small amounts of storage, or vice versa.

Open Environment Storage Virtualization Becomes Reality

Designed correctly, replacing the SAS expander with the Ethernet switch enables real-time reassignment of virtual storage to different storage devices without disrupting communications. Equally important, it can do so in a way that easily interoperates with devices from multiple vendors. While commonplace for compute storage virtualization, such virtualization for storage is much less universal.

This kind of open environment virtualization would give storage the type of flexibility seen with compute resources. It would also give enterprises impetus to migrate more applications to the data-center cloud.

Related

[Storage And Computation Capacity Continues To Grow](#)

[Direct Oriented Disk Drive Goes Direct To Ethernet](#)

[Ethernet Evolves Again To Meet The Internet Of Things](#)

However, to fully realize the on-demand, cloud-based model for IT resources and applications requires one additional piece of the puzzle: Connectivity and networking capacity between virtualized compute and storage resources that's configurable on-demand. Fortunately, software defined networking (SDN) for the data-center network, which can now be completely Ethernet-based, should make that possible. By centralizing the control plane separate from individual network devices, SDN enables two critical differences:

- Creation of a flat, interconnect Ethernet “fabric” from aggregation/leaf and core/spine switches.
- Parsing of that fabric into multiple virtual networks, allowing the correct virtual storage and compute resources to interconnect as required.

This makes SDN ideal for supporting on-demand changes to quality-of-service (QoS) parameters, bandwidth, etc. Furthermore, it opens up real possibilities for data-center operators to truly implement on-demand IT, where network, storage, and computing resources are fully virtualized and connected end-to-end with standard Ethernet communications.

Storage-System Design Considerations

Making the on-demand IT vision a reality requires careful planning and design. As system designers look to implement an Ethernet-based architecture for storage, several factors must be considered:

- *Power*: Especially in large data centers, power is always a critical consideration. Air flow within storage enclosures can be problematic, and space constraints may not allow for enough cooling fins. Designers should look for Ethernet switch ICs with minimal power requirements.
- *SGMII connectivity*: The ideal Ethernet switch should have enough SGMII interfaces to support the number of drives in a storage enclosure, typically 60-90 or more. Double-speed SGMII support is ideal, in order to support future capacity upgrades.
- *Flexibility*: In cases where there's a limited number of traces across the mid-plane between the storage drive and Ethernet aggregation cards, designers should look for Ethernet switch ICs that can multiplex the interface to higher speeds (e.g., 2.5G, 5G, or 10G) and de-multiplex again on the other side before connecting to the storage drive.
- *Scaling ability*: Storage enclosures vary in capacity and size, ranging from about 24 drives for smaller enclosures up to 96 drives in larger systems. Ethernet switch capacity should scale according to your needs, without requiring a system redesign with different ICs and software. A good way to address this is to leverage the stacking capability of Ethernet switches, allowing the port count, power, and cost to scale linearly with the number of supported drives.
- *Embedded CPU*: External processors add power and cost. The ideal Ethernet switch should have an embedded CPU to run the entire operating system and switch protocols, including stacking software, directly on the chip.
- *Switch manageability*: An approach to consider for advanced data architectures is to view the storage enclosure aggregation switch as a port aggregator into the large ToR switch. Treated as an extension of the ToR switch, the aggregation switch also needs to be manageable. Open APIs like JSON-RPC manage virtual hypervisor “vSwitches” from the ToR switch in much the same way.

Flexible VLAN management: Flexible VLAN management in the switch is crucial, enabling VLANs and IP addresses to move from one storage drive to another without impacting communications.

- *Advanced QoS, packet-classification capabilities:* Looking forward, future Ethernet-based storage applications will likely involve both East-West and North-South traffic. The Ethernet switch will need to support advanced QoS and packet classification to enable precise policy control.

Summary

The new Ethernet-based, object-oriented architecture dramatically simplifies storage networks, effectively outmoding overly complex and costly communications protocols and architectures. No longer are the traditional FC SAN switch and storage server requisite in the data center. Ethernet can supplant all—storage-drive SAS interfaces; Ethernet switches instead of SAS/SATA expanders; and multiple protocols such as SAS/SATA, iSCSI, FC, and FCoE. Designed correctly, this optimizes operational efficiencies in the data center, generating substantial total cost-of-ownership savings. With data-center compute *and* storage based on Ethernet, we're that much closer to making the data-center cloud a reality.

Martin Nuss is the CTO and vice president, technology and strategy, for [Vitesse Semiconductor](#). He also serves on the board of directors for the Alliance of Telecommunications Industry Solutions (ATIS), and is a fellow of the Optical Society of America and a member of IEEE. He holds a doctorate in applied physics from the Technical University in Munich, Germany. He can be reached at nuss@vitesse.com.

Source URL: <http://electronicdesign.com/communications/ethernet-everywhere-accelerating-data-center-migration-cloud>