

NETWORKING BENEFITS

- » Scale-out the network with the ease and pace of converged infrastructure and storage
- » Create and enforce application- level SLAs thanks to storage, data and application workload awareness
- » Identify, classify and intelligently distribute five new types of traffic introduced by converged infrastructure on the IP network
- » Deliver storage, data and business critical application traffic without compromise across a single fabric
- » Utilize all of your network capacity, not just some, with a software-definable fabric that creates an elastic pool of network resources

PLEXXI HCN™ FOR CONVERGED INFRASTRUCTURE AND STORAGE

WHITE PAPER

Traditional Networks Fall Short

As scale-out converged infrastructure and storage deployments are growing, the data center network is experiencing unprecedented stress and strain from the newly introduced storage traffic and scaled- out applications. Traditional leaf and spine data center architectures, which have remained unchanged for 25+ years, are ill equipped to handle the large volume of highly bursty and unpredictable east/west and north/ south traffic patterns that scale-out applications and storage introduces.

These legacy networks are static in nature, defined by their cabling, fragile and unaware of the newly introduced traffic types. They struggle to support growing pools of storage capacity that span across racks, rows or multiple data centers.

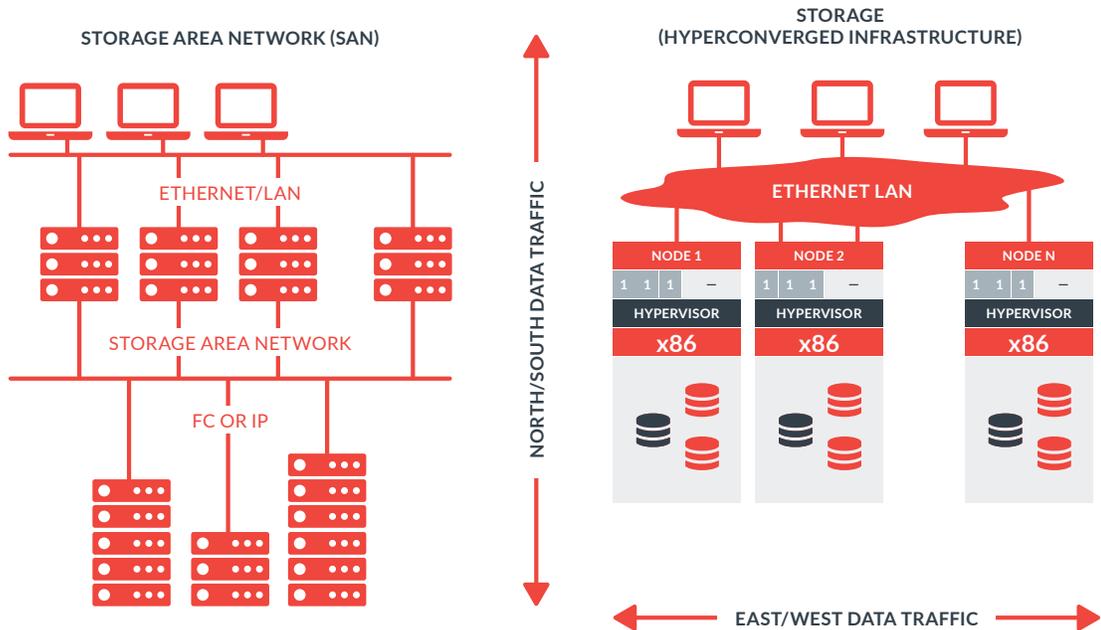
Converged Infrastructure and Storage Change Rules of the Game

The very nature of storage traffic is distinctly different from traditional client/ server network traffic. Existing network architectures were built in an era where network traffic was predominantly north/south or client-to-server. As compute began to scale-out (multi-core), applications and storage also began to scale-out in order to unlock the potential of flash as well as to simplify the deployment of massive PetaByte (PB) sized file systems, object store or arrays. With this scale-out, the bottleneck has shifted to the network.

The requirements placed on the network by modern flash systems, software defined storage systems (vSAN, ScaleIO, etc.) or scaled-out file systems like HDFS/Hadoop, NDFS, pNFS, are fundamentally different. Not only is storage traffic moving east/west and north/south, but also the requirement for consistent, low latency (low jitter) is exceptionally high.

Storage Networks of the future must offer elasticity and low latency regardless of physical location (any rack, any row, any data center), but most importantly must be under software control with an awareness of the ever-changing needs of storage. Storage networks of the future must to be capable of avoiding the “Microbursts” that are seen on traditional networks that were designed twenty years ago to solve a different set of problems.

Fig. 1 Platform 2 was about scale-up storage area networks which created north/south client to server traffic. Platform 3 moves storage onto the IP network and scales-out, generating significant east/west as well as north/south traffic.



Storage on IP networks introduces several new types of traffic, some creating very large volumes of unpredictable traffic into the data center.

- » Management traffic (include all traffic related to management including logs)
- » Control traffic (cluster node communication for node failover etc.)
- » Client-to-Server Storage traffic (VM to storage nodes—for accessing storage)
- » Server-to-Server Storage traffic (storage node to storage node—rebalancing etc.)
- » Meta data—the data about the data that file systems, object stores and distributed block systems (like vSAN) use to organize the data.

More Bandwidth and Bigger Buffers Are Not the Answer

Throwing more bandwidth and bigger buffers at the problem is costly and doesn't solve the fundamental issue. The natural IO size for modern storage transactions has rapidly grown from what were once 512 bytes to modern file system transactions of 4,816KB, and even larger size transactions for video and other types of rich media. Traditional data center networks are built, configured and managed as separate technology domains from storage and compute. Cloud-based environments need the ability to unify workflows across storage, compute AND network domains through tools, automation and orchestration. This requires a next generation approach to data center networking to address these problems.

Network Considerations Per Storage Technology

TECHNOLOGY	ADVATANGES	CHALLENGES	DATA (IP) NEWTWORK CONSIDERATIONS
Network-Attached Storage (NAS)	Better resource efficiency than DAS, relatively simple to operate	File-level (versus block) access has performance issues for applications	Performance can be impacted if network is congested or undersized
Block SANs (e.g. Fibre Channel)	Performant at scale, resource efficient	Complex and costly to scale and operate, requires separate operational domain	Typically isolate from the network as a distinct storage area network (SAN)
IP SANS (e.g. iSCSI)	Less complex and easier to operate than FC while still providing block-level access by applications for performance	Still requires separate operational domain	Opportunity for Data/Storage network convergence—but requires a networks that can provide storage network policies and isolate traffic from rest of network
Server SAN/ Hyperconverged	No separate storage arrays/ devices (part of compute nodes), incrementally scales when adding new computer nodes	Converged by design—high volumes of bursty node to node traffic can be disruptive to the rest of the network	Difficult to scale unless the network has a way of understanding data intensive tasks and ensuring adequate bandwidth and minimal latency is available

It Takes a New Approach

Plexxi offers a software-definable, programmable fabric capable of dynamically allocating bandwidth. Through storage, data and application awareness the fabric adjusts, in real time, to the needs of specific traffic types enabling the creation and enforcement of application-level SLAs.

The intent-based and software-definable fabric creates an elastic pool of network capacity. Data, application AND storage workload traffic is intelligently, efficiently and securely distributed across the entire fabric allowing you to use all of the network capacity you purchased, not just a small fraction of it, like legacy leaf and spine approaches. With Plexxi, the scale-out of storage and converged infrastructure resources is seamless, delivering a solution that lowers overall CAPEX/OPEX and boosts utilization.

The information contained herein is subject to change without notice. Plexxi, Plexxi HCN and the Plexxi logo are trademarks of Plexxi Inc. Other company, product or service names may be trademarks or service marks of their respective owners.