SOLUTION BRIEF

Datera Low Latency Layer 3 Routed Storage with Mellanox

Enabling storage to take advantage of network performance improvements by integrating at Layer 3

OVERVIEW

Ethernet has emerged as the interconnect protocol for cloud services. The ubiquity and prevalence of Ethernet has made it easy to use and understand, and thus the technology of choice.

Advances including faster storage media and Ethernet speeds have exposed the limitations of aging protocols and frameworks (Fibre Channel, and complicated Layer 2 topologies). Legacy networked storage systems are stymied by the inability to fully utilize the performance of new media types including NVMe and 3D XPoint. Fibre Channel with its protocol limitations and link speeds simply can't keep up. Low latency media within a storage server is not translated effectively to the network or across the network. All this is addressed by the joint Datera-Mellanox solution.

DATERA DATA SERVICES PLATFORM (DSP)

Datera DSP is a software-defined data infrastructure for virtualized environments, databases, cloud stacks, DevOps, and micro-service deployments. Like other cloud solutions, Datera DSP is built on the foundation that resources are aggregated into shared, scalable pools. For Datera scale-out cloud storage, the datacenter network effectively becomes the "backplane."

In modern flat networks, proliferating "floating" network endpoints necessitate a matching flat network namespace. This effectively requires pulling native Layer 3 network support into the network leaf nodes. For instance, in a container environment, it is common for each node to have its own /31 network with 10-100s of containers that can dynamically be moved around (each node is a small data center in its own respect). Datera nodes can act as native Layer 3 network endpoints, and thereby seamlessly participate in all route changes. This architecture provides the ability to deliver the low latency storage performance (including NVMe and 3D XPoint) directly to the applications while using standard x86 storage nodes.

Datera system spans and block placement methods can use fault domains (such as rack awareness) by placing replicas/spans in different fault domains. This allows tremendous flexibility by providing data center scaling rather than array or rack level scaling from traditional storage architectures. In a standard monolithic array, the storage is defined and limited by the array boundary. In old-style scale out storage, clusters are limited to a Layer 2 subnet (or rack). Datera's DSP can be distributed in nodes across racks with provision made for rack and switch failure handling. Datera is aware of the rack boundaries but is not limited by them.

KEY BENEFITS

- High performance with standard hardware
- Makes use of Layer 3 for storage nodes within DSP
- Ultra-low latency and high throughput for all application use cases
- Optimized server utilization with up to 100GbE networking and RDMA offload
- APIs and automation enable software-defined data centers
- Application, storage and front-end network in one network fabric

"Working with Mellanox to build a complete low latency solution enables Datera to fully accomplish the objectives of our architecture."

– Flavio Santoni, CRO, Datera

Benefits of Layer 3 networking for cloud storage are:

- **Optimal link capacity utilization**: Layer 3 protocols such as OSPF or BGP utilize all available links
- **Fast convergence**: On a link failure or node failure, Layer 3 allows for fast convergence and strict adherence to the IO timeout values
- **Dynamic network balancing**: With equal cost multi-pathing, and dynamic Layer 3 routing, the network adapts to topology and bandwidth changes delivering optimized performance
- **Security**: With the ability to control Layer 3 Access Control Lists, and allow for white lists communication, a Layer 3 topology is far more secure.
- **Quality of Service** (QoS): This architecture allows for granular QoS controls across the network thereby allowing fine grain controls for traffic engineering
- **High scalability**: As cloud deployments are constantly growing, Layer 3 topologies allow for growth of nodes or bandwidth or services without disruption to the existing cluster.

DevOps teams create one or more templates that determine policies that are applied to the various nodes but are not restricted to given nodes. Changes in the templates propagate throughout the cluster automatically resulting in real time data migration and enabling an always on storage service. The DSP was developed from the kernel up to take advantage of new media types and new network topologies with serving the intent of each application as the highest value. As a result, Datera can make use of the advances provided through the Mellanox Ethernet solution.

MELLANOX END-TO-END ETHERNET SOLUTION

The Mellanox end-to-end Ethernet solution provides a simple, efficient and high-performance network fabric for modern data centers and clouds. Built on the best-in-class switching ASIC, Mellanox Spectrum™ switches are modern switches with a rich set of Layer 2/3 features and deliver line-rate throughput and ultra-low latency at any speed with zero packet loss. With dynamically shared buffer and RDMA over Converged Ethernet (RoCE) optimization, Spectrum switches support the fastest storage traffic with high and resilient performance. (See Figure 1.)

In particular, the Mellanox SN2010/SN2100 switches are the ideal TOR switches for storage and hyperconverged deployments. In a half-width, 1U form factor, SN2100 provides 16 100GbE ports whereas SN2010 provides 18 10/25GbE ports plus 4 100GbE ports. Two of these switches can be installed in a 1U rack space for high-availability. And using break-out cables, the SN2100 supports up to 64 10/25GbE ports and enables the highest-density rack design.

Equipped with intelligent ASICs, Mellanox Ethernet adapter cards provide offloading mechanisms such as RDMA over Converged Ethernet (RoCE), Erasure Coding, T10/DIF, TCP and UDP offloads, and overlay offloads. By bypassing the CPU with hardware offloads, Mellanox ConnectX® NICs free up the CPU's resources, especially in a CPU-bound environment, for the necessary storage and compute tasks, allowing for higher scalability and efficiency within the data center.

The Mellanox Ethernet solution also includes automated network provisioning, monitoring and troubleshooting through Ansible integration and Mellanox's network orchestration and management software, NEO™. Network automation and zero-touch provisioning allow rapid deployment and intelligent network management, and also reduce network downtime and improve operation efficiency.

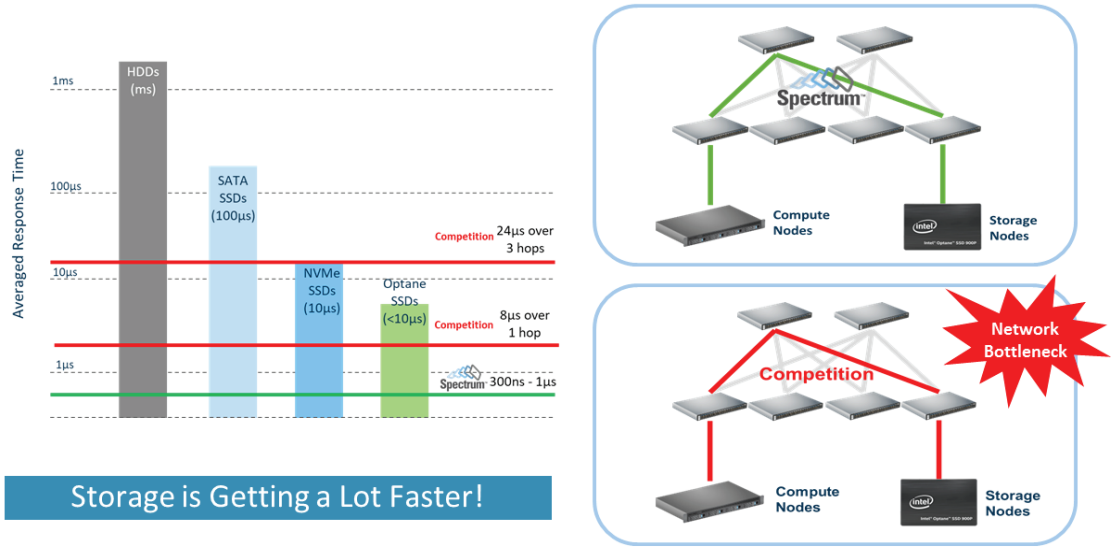

Figure 1. Mellanox Ethernet Storage Fabric unlocks the maximum storage performance

JOINT SOLUTION

Datera nodes built from standard x86 servers, with low latency storage media such as NVMe and 3D XPoint combined with a Mellanox low latency network backbone can deliver unprecedented IO performance and scale to the applications. This solution can deliver IO at sub-80 µs latency at 100,000 IOPS per node to the application layer. In addition, Datera's system intelligence ensures that data is served from the fastest available node within the Mellanox fabric ensuring consistently high performance across a large and highly utilized cluster.

Datera's DSP is a subscription-based software product that is available on a terabyte per month or per year basis. The cluster and subscription fees can be scaled up or down based upon usage. All software and support are included in a single subscription SKU. For customers who prefer an appliance approach, Datera works with integration partners who can deliver a turnkey appliance using the server vendor of the customer's choice containing Mellanox NICs and top of rack switches.

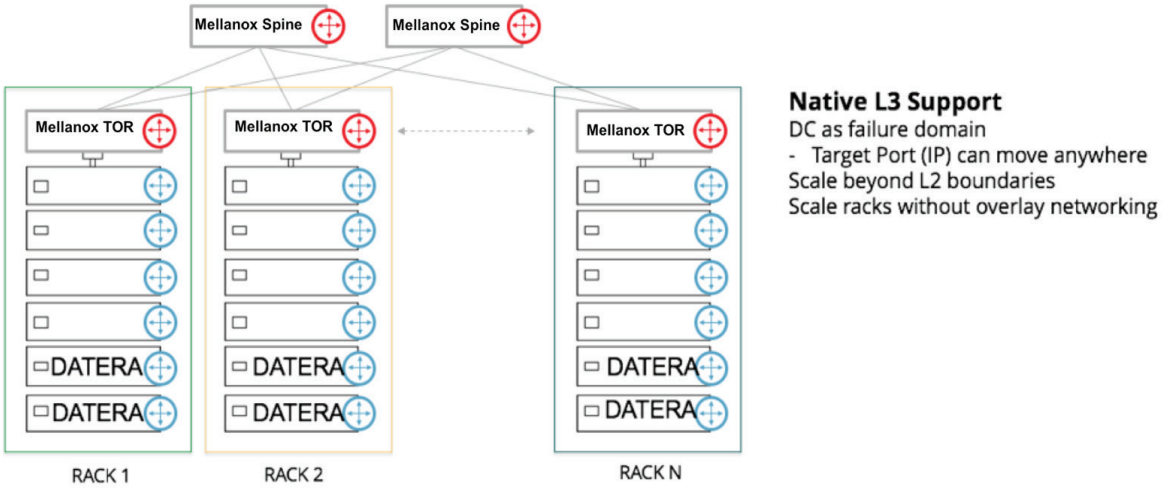

Figure 2. Joint Mellanox-Datera solution includes native Layer-3 support

About Datera

Datera is the only 100% software-based elastic data services platform that powers high performance, data-intensive applications @scale with breakthrough orchestration and automation in partnership with the leaders in server-based infrastructure, all at a 70% lower total cost of ownership and operation.

Data Automation – Datera enables IT teams to easily "program" the level of service needed for individual applications and tenants, including performance, efficiency, availability, security and protection, and leaves the on-going monitoring and dynamic service optimization to us, freeing IT teams to focus on accelerating application development.

Self-Driving Data Orchestration – Datera orchestrates the placement of data within an overall data environment down to the location, device and storage media – 3D Xpoint, NVMe, Flash, and SATA – to achieve, optimize and exceed service levels and extend these higher performance benefits to the S3 object protocol, increasingly important to the next generation of applications.

Maximum Performance – Enterprises are driven by the "need for speed" for their application environments. Designed with data-intensive, database-driven applications in mind, Datera's predictive intelligence engine automatically maximizes environment-wide performance to run at 180K+ IOPS per node and less than 200 microseconds of latency through efficient request routing (to minimize latency) and continuous refinement of data placement down to the storage medium to account for "hot" data.

Scale with Operational Control – Datera enables continuously available data environments of up to 6 petabytes in capacity distributed across an active-active, multi-site environment built with multi-tenancy and protection zoning in mind.

Freedom – Datera delivers you the freedom to mix and match the full complement of application serving choices and infrastructure appliances to enable your optimal data center environment. Datera supports choice in the application environment through certification with VMware, Docker, Rancher, and Kubernetes and choice of server infrastructure through our partnerships with the leaders in infrastructure-based technology – Intel's Optane Group, Samsung's Flash Memory business group and the server divisions of HPE, Cisco, Dell and SuperMicro – to turn it into reality.

About Mellanox

Mellanox Technologies (NASDAQ: MLNX) is a leading supplier of end-to-end Ethernet and InfiniBand intelligent interconnect solutions and services for servers, storage, and hyper-converged infrastructure. Mellanox intelligent interconnect solutions increase data center efficiency by providing the highest throughput and lowest latency, delivering data faster to applications and unlocking system performance. Mellanox offers a choice of high performance solutions: network and multi-core processors, network adapters, switches, cables, software and silicon, that accelerate application runtime and maximize business results for a wide range of markets including high performance computing, enterprise data centers, Web 2.0, cloud, storage, network security, telecom and financial services. More information is available at www.mellanox.com.

350 Oakmead Parkway, Suite 100, Sunnyvale, CA 94085
Tel: 408-970-3400 • Fax: 408-970-3403
www.mellanox.com

060219SB
Rev1.0