

# EMC Celerra Multi-Path File System (MPFS) with Aspera

Joint solution for ultra high-speed bulk data transfer

## The Big Picture

- ▶ **Distance-independent performance:** MPFS with Aspera *fast*™ technology breaks free of the packet loss and latency bottlenecks of standard TCP transport solutions like NFS, CIFS, FTP and HTTP over the wide area network (WAN), transferring large files and file sets up to 1000X faster. Multi-gigabit throughput performance allows for incredible data restore times or highly concurrent bulk data access at global distances, over commodity and dedicated IP networks.
- ▶ **Optimized environment:** combined, MPFS and Aspera *fast* allow for fast data sourcing and resource control, with dynamic, policy-based WAN bandwidth utilization.
- ▶ **Improved efficiency:** MPFS with Aspera's *fast* protocol allows full utilization of public and private WAN bandwidth regardless of network distance, at high performance rates. This reduces wait times for data increasing resource utilization, service and process efficiency.

## Overview - Celerra Unified Storage with MPFS and Aspera *fast* High-speed Data Transfer

Celerra MPFS allows clients to communicate over an IP network to share metadata and control information using the NFS or CIFS protocols. When the MPFS client receives a data request from the application, it uses Celerra-provided metadata to access the storage array directly using low-latency and high-performance SAN.

TCP-based applications (like FTP, HTTP, CIFS and NFS), ensure the reliable delivery of data, but over distance (WAN networks) throughput performance can be severely constrained by latency and packet loss. This creates significant inefficiencies in moving data, resulting in underutilized resources and lower productivity. Aspera's patented *fast* protocol is an innovative bulk-data transport technology, providing a superior alternative to traditional TCP-based transport technologies. *fast* works over any public and/or private IP network; it is implemented at the application layer, allowing easy deployment within the enterprise by re-using existing client, server, operating system and networking technologies. *fast* is designed to deliver 100% bandwidth efficient transport of bulk data – independent of network delay, packet loss, and data size.

Celerra MPFS combined with Aspera's *fast* technology provide the optimal solution for fast data retrieval within the SAN, and efficient and scalable bandwidth utilization over the WAN. Lab testing with 10Gb Ethernet WANs has resulted in single file and multiple file transfers reaching transfers of 3.Gbps. This was achieved with Aspera software running on commodity computers, over networks with worst-case global Internet conditions of 5% packet loss and 300ms of latency. Such performance would allow delivery of a 1TB of aggregate file data in 50 minutes, over any geographical distance.

## High-Performance NAS Test Environment

Performance testing of the combined solution was conducted to demonstrate WAN transfer rates with varying packet loss and latency combinations. Additionally, to stress the performance of both MPFS and Aspera a specific a range of very large file sizes was chosen. Data was transferred between source and destination file systems using MPFS and Aspera's *fast* protocol.

The test environment consisted of the following hardware and software (Figure 1):

- ▶ EMC Celerra NS-480 unified storage platform
- ▶ 2 x Dell 1950's
  - ➔ 1U blade servers with Intel Xeon 2GHz Quad-Core, 16 GB of RAM
- ▶ Linux Red Hat Enterprise
  - ➔ 64-bit, kernel version 2.6.18-92.el5 operating system
- ▶ 1 x Emulex 4Gb fibre channel (FC)
- ▶ 1 x 1GbE and 1 x 10GbE Intel NIC
- ▶ FC connectivity to the drives within the NS480 (SAN delivery)
- ▶ File system consists of 45 disks
  - ➔ 4Gb FC, 450 GB, 15K RPM
- ▶ Apposite Linktropy 10 Gb WAN emulator

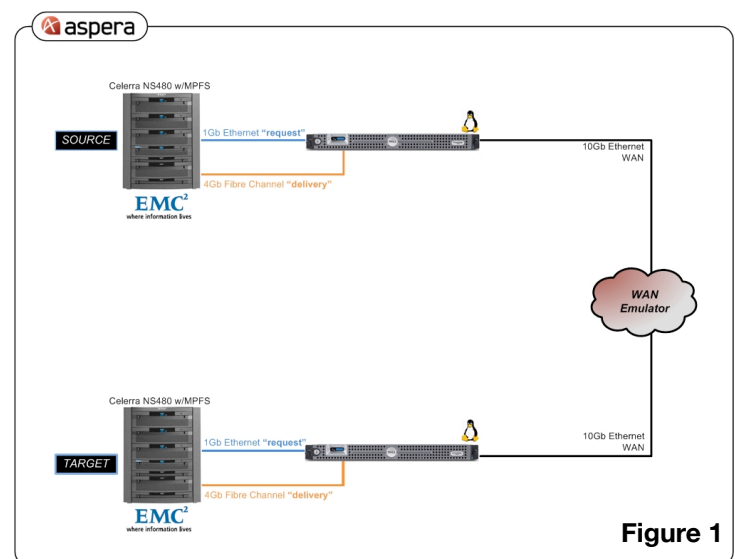



Figure 1

### Results

The performance test bed simulated bulk-data transfers over a 1Gbps and 10Gbps Ethernet WAN infrastructure, differing packet loss and latency conditions were introduced and the test metrics delivered were: data throughput, efficiency and transfer times. The expectation from this testing was two-fold: 1) Because WAN latency and packet loss greatly diminish data throughput at high speeds, the presence of both (at varying degrees) would quantitatively demonstrate Aspera's uniform transport performance and 2) The MPFS file system would be able to enhance storage throughput at rates many times faster than accessing the same data using an NFS-only file system or other NAS protocols.




**Figure 2**

**Figure 2:** The effective file transfer throughput achieved across all 10GbE WAN packet-loss, and latency conditions, was consistently over 2.9Gbps; even for worst-case global Internet conditions (300 ms / 5% packet loss). At this speed, 1TB of data could be moved over the WAN in 50 minutes, which creates innovative ways of architecting storage systems, data management applications and DR scenarios.

System throughput performance was measured using both small-file large directory structures totaling 25 GB and single large files of 25 to 100 GB. Files were read from, and written to the NS480 storage systems over 10GbE WAN. To maximize storage throughput, four concurrent *fasp* sessions were initiated concurrently. For reference, a single *fasp* session transferred data at 1.9 Gbps.

Based on this throughput performance of MPFS with *fasp*, the transfer times for 1TB of data, over the best and worst case WAN latency and packet loss combinations (300ms/5%), is ~50 minutes. Compare this to a TCP transfer times of 286 days, with the same WAN conditions, or 1 day in a favorable WAN environment.

**Figure 3:** Shows the throughput performance of FTP, SCP & RSYNC, and standalone TCP. The throughput performance of these applications when using high-bandwidth latent & loss-prone WANs can be extremely poor, if not unusable. In comparison, the MPFS with *fasp* solution shows a 1250X throughput gain over the best performing equivalent. In worst-case network scenarios (300ms latency/5% loss), MPFS with *fasp* performance is shown to be robust and exceptionally resilient, delivering near-linear throughput performance with little degradation.



**Figure 3**

### Conclusion

As enterprises grow to support distributed employees, partner setups, collaborative projects and global development, data availability and accessibility become a cornerstone to high productivity and efficiency. Business processes that necessitate the movement of/or access to bulk data require delivery times that meet project or revenue deadlines. With legacy technologies unable to meet demand, future network, application and storage architectures will need to carefully balance cost, redundancy and global application accessibility. Celerra Unified Storage with MPFS technology provide all the traditional benefits of NAS, in addition to the raw performance, and scalability benefits of block storage access. The Celerra MPFS and Aspera solution shows how highly scalable, high-speed network, storage and transport technologies can be deployed in combination to meet the needs of today's digital data-driven business.

**About Aspera** — Based in Emeryville, CA, privately held Aspera has been delivering bulk-data movement solutions for 6 years. Aspera's server and client software products are powered using *fasp*<sup>TM</sup>, a patent-pending transport protocol. *fasp* is specifically designed for high-speed bulk-data transfers, allowing companies to move data at maximum throughput, regardless of file size, transfer distance and network conditions.

**About EMC** — EMC Corporation (NYSE: EMC) is the world's leading developer and provider of information infrastructure technology and solutions that enable organizations of all sizes to transform the way they compete and create value from their information. Information about EMC's products and services can be found at [www.emc.com](http://www.emc.com)