



Data ONTAP GX – NetApp’s New Horse Team and Wagon for High-Performance Computing

Analyst: Michael Fisch

Management Summary

NetApp recently announced *Data ONTAP GX*, a new high-performance operating system for its networked storage platforms. It clusters multiple platforms, so they look and function like a single system. Data ONTAP GX scales to a very large six petabytes of capacity and over 1,000,000 IOPS of performance. It adds nodes dynamically and incrementally, allowing its capabilities to grow over time with an enterprise’s storage requirements.

NetApp offers Data ONTAP GX in addition to its existing *Data ONTAP 7G* operating system. They are priced the same, and both run on the same storage platforms (*FAS3050* and *FAS6070* only for Data ONTAP GX). Data ONTAP GX is a parallel offering.

So, what is the difference between them? Consider, as an analogy: Data ONTAP 7G is like a solo horse and Data ONTAP GX is like a horse team and wagon. A solo horse can carry so much weight, including rider and pack, but there are limits. If the rider wants to carry more, he would have to switch to a larger horse, such as a Clydesdale. This may be inconvenient, and still there are weight limits. On the other hand, the horse team and wagon can carry more cargo than any solo horse. It also can increase capacity by hitching additional horses to the team – one wagon, one driver, one team of a variable number of horses. The horse team and wagon is significantly more powerful, flexible, and scalable than the solo horse.

Similarly, Data ONTAP GX is the more scalable and flexible operating system. Like the horse team, it can “hitch” from two to 24 storage nodes into a single system. It presents one file system and namespace across all nodes. It can even stripe a file across any or all nodes, so the file may leverage the aggregate system performance – drives, processors, and host ports. It supports NFS, a basic CIFS capability, Gigabit Ethernet, and Fibre Channel and SATA drives.

Data ONTAP GX is a great fit for high-performance computing (HPC) applications, which perform massive calculations against large data sets. These applications include seismic processing for oil and gas, chip design and simulation, video rendering for film and entertainment, and various lab applications. **In fact, NetApp specifically recommends Data ONTAP GX as a storage solution for HPC Linux clusters.**

At the same time, its peer Data ONTAP 7G is still NetApp’s preferred offering for other environments because of its broad connectivity and complete set of data management and protection capabilities. Read on for details.

IN THIS ISSUE

➤ Scale-Out vs. Scale-Up.....	2
➤ Data ONTAP GX	2
➤ Benefits.....	3
➤ Data ONTAP GX Applications.....	3
➤ Conclusion	3

Scale-Out vs. Scale-Up

There is a debate within the computer industry about whether a scale-up or scale-out architecture is best for increasing the scale and workload capacity of a system. Scale-up represents the traditional approach of *building a bigger box*. If an enterprise needs more compute power or storage capacity, scale-up means installing a larger, self-contained system. Examples include mainframe and n-way SMP servers, as well as monolithic, high-end storage platforms. Scale-out represents the approach of *clustering many smaller, independent systems*, so they work together as a single entity. If more compute power (capacity) is needed, scale-out means attaching more nodes to the cluster. The nodes are aware of each other and capable of sharing resources and functioning in a logically (or virtually) unified manner. Grid computing is an example.

Both scale-up and scale-out offer particular advantages. Scale-out is dynamically expandable and avoids large upfront investments and disruptive upgrades. Scale-up is self-contained and can offer more built-in functionality, especially at the hardware layer.

So, which architecture is best? That is for pundits to ponder and the market ultimately to decide. In the meantime, it is safe to say that both scale-out and scale-up will be around and that scale-out, as the newer of the two, is set to play an increasing role.

NetApp's new *Data ONTAP GX* offers a scale-out storage systems architecture for high-performance, scale-out server environments, like high-performance computing on a large-scale grid of servers.

Data ONTAP GX

Data ONTAP GX is long-anticipated fruit of NetApp's acquisition of Spinnaker Networks, which developed a distributed file system for scale-out NAS systems. NetApp has been working on a multi-year roadmap to converge its tried-and-true Data ONTAP operating system with the one from Spinnaker. The goal is to create a super operating system with the best of both worlds – the flexibility and functionality of Data ONTAP and the scale and performance of Spinnaker. Data ONTAP GX represents a milestone toward that end. It

offers a number of significant capabilities.

Clustered File System

Data ONTAP GX creates a single file system and namespace across multiple platforms (nodes). It clusters from two-to-24¹ nodes, forming a single logical system that scales to very high levels of performance and capacity – up to 6,000 TB, over 1,000,000 IOPS, and multi-GB/s throughput. This reaches well into the high end of the storage market.

All nodes have access to all data. If a node fails, applications and users can continue accessing data through another one in the cluster. Data ONTAP GX can also stripe directories and files across all nodes, delivering massive throughput for even a single file by leveraging the aggregate performance of drives, processors, and host ports.

Non-disruptive Operations

Data ONTAP GX scales performance and capacity discretely and linearly by adding nodes to the cluster. These additions are non-disruptive to applications. The system also can migrate data non-disruptively, for purposes of load balancing or lifecycle management. No software needs to run on host servers.

Standard NetApp Hardware

Data ONTAP GX supports the *FAS3050* and *FAS6070*, two of NetApp's standard hardware platforms. They can be mixed within a cluster. The system supports both Fibre Channel and SATA drives, the former for performance and the latter for large capacity and low cost. It uses Gigabit Ethernet connections for both host connectivity as well as communication between the nodes.

Select Data ONTAP 7G Capabilities

Data ONTAP GX also offers a select set of capabilities from NetApp's existing Data ONTAP 7G.

- **WAFS – Write Anywhere File Layout** is the specialized file system that NetApp created for optimizing performance.

¹ 24 nodes is a qualification limit, not a technical one. NetApp will likely offer larger clusters in the future, as market requirements grow.

- **FlexVol HPO (High Performance Option)** – Improves utilization and by virtualizing system capacity and allocating it as data is written.
- **RAID-DP** – Provides two parity disks in a RAID group to improve data protection. This is especially helpful for mitigating the possibility of data loss during SATA drive rebuilds.
- **Snapshot** – Creates differential data copies for fast restores, non-disruptive backup, and other data re-purposing.
- **NFS** – Data ONTAP GX is currently focused primarily on NFS (limited CIFS support and no block protocols).

Data ONTAP GX and Data ONTAP 7G are priced the same. Customers simply choose which operating system best suits their requirements.

Benefits

Several business and technical benefits accrue from the scale-out architecture of Data ONTAP GX.

- **Very large scale** – The level of performance and capacity it attains could serve high-performance computing applications or even multiple enterprise applications.
- **Pay as you grow** – Since the system is dynamically and incrementally expandable, an enterprise can start with a small system and grow as requirements dictate, without paying upfront for a monolithic system. This model fits better into IT budgets and avoids underutilized equipment.
- **Non-disruption to productivity** – Continuous data access translates into user productivity in a busy enterprise. The ability to add nodes, move data, and failover transparently contributes to this productivity. No need for disruptive forklift upgrades.
- **High availability** – Multiple nodes with transparent failover provide built-in high availability.

- **Centralized management** – With a single system image across multiple platforms, it offers efficiencies of management at scale and lowers storage TCO.

Data ONTAP GX Applications

With these benefits, why wouldn't a customer choose Data ONTAP GX? The fact is that Data ONTAP GX is still a somewhat specialized operating system because it does not (yet) offer the fully array of NetApp's data management and protection capabilities found in Data ONTAP 7G. For instance, it does not support full CIFS functionality for Windows environments or block protocols (i.e., Fibre Channel, iSCSI), nor does it offer application integration.

That said, Data ONTAP GX is a superb fit for high-performance computing, which is why NetApp is specifically targeting Linux clusters for applications such as:

- Seismic processing in the energy & gas industry
- Integrated circuit simulation and verification in the semiconductor industry
- Video creation and rendering in the entertainment industry
- Lab applications in the intelligence, defense, and biotech sectors
- Risk analysis in the financial services industry

For all other environments and applications, NetApp offers its staple Data ONTAP 7G.

Conclusion

If you have a Linux computer cluster running a serious computational application, then it likely needs storage with very large-scale capacity and performance to “feed the beast”. Data ONTAP GX is a great solution for this environment – and it comes from an established vendor with solid support services, to boot.



About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

Michael Fisch is Director of Storage and Networking for The Clipper Group. He brings over ten years of experience in the computer industry working in sales, market analysis and positioning, and engineering. Mr. Fisch worked at EMC Corporation as a marketing program manager focused on service providers and as a competitive market analyst. Before that, he worked in international channel development, manufacturing, and technical support at Extended Systems, Inc. Mr. Fisch earned an MBA from Babson College and a Bachelor's degree in electrical engineering from the University of Idaho.

- ***Reach Michael Fisch via e-mail at mike.fisch@clipper.com or at 781-235-0085 Ext. 211. (Please dial "211" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group Captain's Log, The Clipper Group Voyager, and "***clipper.com***" are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "***Navigating Information Technology Horizons***", and "***teraproductivity***" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.