



White Paper

Remove IT Constraints and Speed Response to Business Changes with Clustered Data ONTAP

Mike McNamara, NetApp
June 2013 | WP-7186

Executive Summary

IT teams are working harder than ever in the face of massive data growth, complex operations, and time and budget constraints. NetApp created the clustered Data ONTAP[®] operating system to address the limitations of traditional approaches to IT, removing IT constraints and allowing you to respond to business demands more readily. Clustered Data ONTAP decouples data access from physical storage devices to create a software-defined storage (SDS) infrastructure with virtualized storage services, application self-service, and the ability to incorporate multivendor hardware. NetApp continues to enhance these capabilities, helping you stay ahead of increasingly rapid changes in the IT world. This white paper explains how clustered Data ONTAP enables nondisruptive operations, increases operating efficiency, and delivers seamless scalability. It also includes highlights of the Data ONTAP 8.2 release.

TABLE OF CONTENTS

1	Does Your IT Infrastructure Respond Quickly Enough?	3
2	Clustered Data ONTAP Increases IT Agility	3
3	Nondisruptive Operations	4
3.1	A Better Approach to Maintenance and Lifecycle Operations	5
3.2	Eliminate the Pain of Technology Refresh	7
3.3	Simplify Data Protection	7
4	Proven Efficiency	7
4.1	Native Multi-tenancy and Quality of Service (QoS)	8
4.2	Storage Efficiency and Flash	8
5	Seamless Scalability	9
5.1	A More Flexible Approach to Scale-Out	10
6	Conclusion	10

LIST OF TABLES

Table 1)	Examples of nondisruptive operations	6
----------	--------------------------------------	---

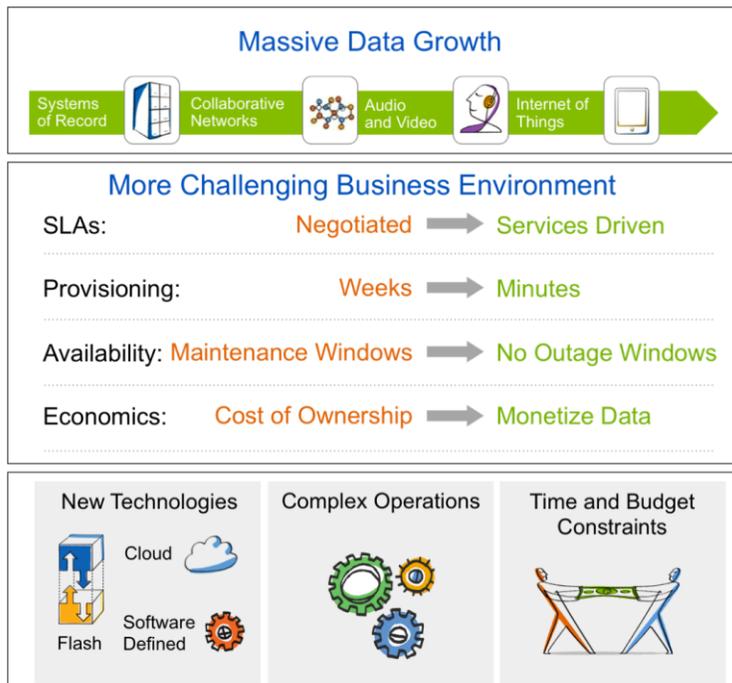
LIST OF FIGURES

Figure 1)	IT storage challenges	3
Figure 2)	Clustered Data ONTAP architecture and features	4
Figure 3)	Nondisruptive operations protect your infrastructure against planned and unplanned downtime	5
Figure 4)	Storage efficiency technologies change the data growth curve and drive down costs	9
Figure 5)	Clustered Data ONTAP scales in several dimensions	9

1 Does Your IT Infrastructure Respond Quickly Enough?

Your IT team is working harder than ever in the face of massive data growth, complex operations, time and budget constraints, and the need to accommodate new technologies such as flash and cloud. At the same time, the business functions you support demand higher service levels, faster provisioning, and increased availability.

Figure 1) IT storage challenges.



Traditional storage designs simply can't keep up. Silos of storage are expensive, become increasingly complex with scale, and can't deliver the performance or availability that modern operations require. NetApp designed clustered Data ONTAP to address these limitations to allow you to eliminate IT constraints and speed response to business changes.

2 Clustered Data ONTAP Increases IT Agility

NetApp clustered Data ONTAP is a revolutionary approach to storage software that is uniquely well suited to a broad range of storage needs, including virtualized environments, scale-out NAS, and enterprise applications including mission-critical business processes.

With innovative storage virtual machines (SVMs), clustered Data ONTAP decouples data access from physical storage devices to create a software-defined storage (SDS) infrastructure with virtualized storage services, application self-service, and the ability to incorporate multivendor hardware.

With clustered Data ONTAP you can:

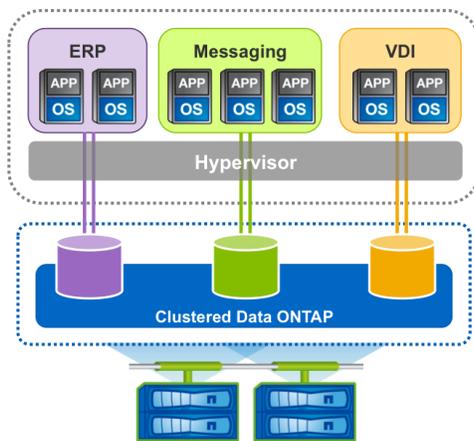
- Eliminate planned downtime
- Scale and adapt your storage environment to growing and changing business needs
- Consolidate and share the same infrastructure for numerous applications, departments, or tenants
- Add capacity on the fly without reconfiguring applications
- Simplify complex tasks
- Increase IT agility across your entire infrastructure

The unified cluster architecture of clustered Data ONTAP adapts to your changing needs, reducing risk and cost and allowing you to service your infrastructure at any time—even during normal business hours—without disrupting access to user or application data.

Proven operational efficiency helps you simplify your overall storage environment and manage storage infrastructure at scale by automating important processes and increasing administrator productivity. You can add capacity as you grow across both SAN and NAS environments, and you can start small and grow big without the disruptive hardware upgrades required by other storage vendors.

Figure 2) Clustered Data ONTAP architecture and features.

- Built-in multi-tenancy
- Integrated data protection
- Continuous data access
- Common software
- Common systems
- Unified management



NetApp continues to add value to clustered Data ONTAP, helping you to stay ahead of increasingly rapid changes in the IT world and today's more dynamic business environments. Our latest release—clustered Data ONTAP 8.2—provides a host of new features, including:

- Improved nondisruptive operations
- Quality of service (QoS)
- Next-generation backup
- Ease of deployment
- Higher scaling
- Enhanced virtualization

These enhancements make it even easier for IT teams to take advantage of clustered Data ONTAP's nondisruptive operations, proven efficiency, and seamless scalability.

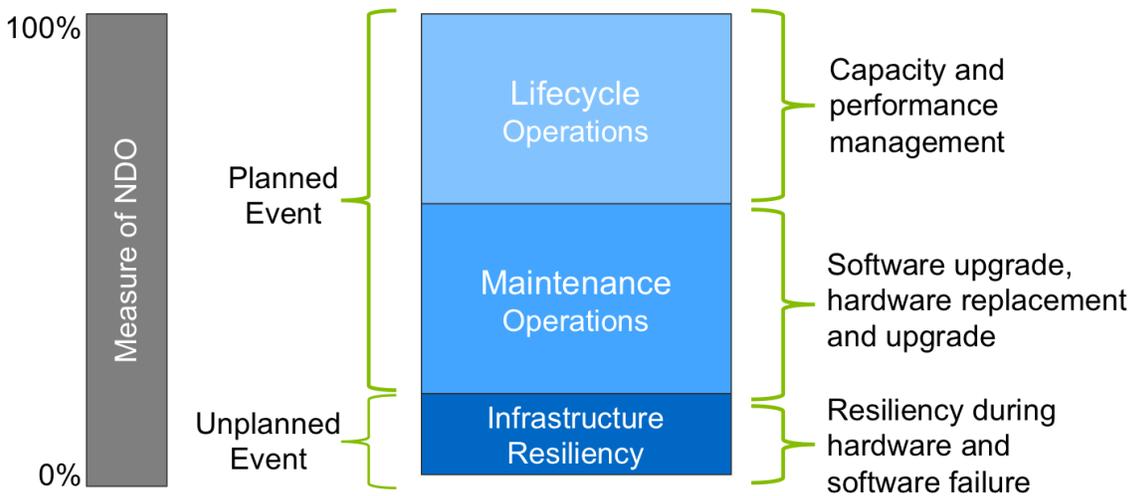
3 Nondisruptive Operations

Growing competitive pressure and 24-hour business cycles require that your mission-critical business processes and data be accessible around the clock. As much as 90% of downtime results from planned or predictable events such as software and firmware upgrades, replacing components, maintenance, adding capacity, and so on. With IT now an integral part of your business operations, the impact of downtime goes beyond dollars or productivity lost. Your company's brand might be at stake. Clustered

Data ONTAP eliminates sources of downtime and protects your critical data against disaster so important business processes stay online.

All clustered Data ONTAP configurations (with the exception of single-node clusters) are designed to be fully resilient to component and software failures. By making it possible to perform maintenance as it's needed—rather than deferring it to rare planned downtime windows—clustered Data ONTAP further reduces the risk of unplanned downtime. By using NetApp AutoSupport™ tools to identify potential issues and performing maintenance as needed, you can reduce the number of priority-1 support cases by 80%.

Figure 3) Nondisruptive operations protect your infrastructure against planned and unplanned downtime.



3.1 A Better Approach to Maintenance and Lifecycle Operations

In the traditional storage world, workloads are assigned to inflexible physical assets. When those physical assets need maintenance, IT has to schedule downtime not only for that asset but for all the workloads, applications, and users that rely on it.

Clustered Data ONTAP offers a unique capability to support nondisruptive maintenance operations and lifecycle operations such as load balancing, capacity expansion, and technology refresh so you can perform critical tasks without interrupting your business. The ability to dynamically assign, promote, and retire storage resources lets you improve service levels over the lifecycle of an application.

Because clustered Data ONTAP can transparently migrate active data and network connections anywhere within a cluster, data can be redistributed any time and for any reason. These activities are transparent and nondisruptive to NAS and SAN hosts, allowing the storage infrastructure to continue serving data during operations such as rebalancing capacity usage, optimizing changing performance requirements, or isolating one or more controllers or storage components for upgrade. On average, NetApp customers perform approximately 500 volume moves a year per cluster.

If you consider the typical consumption pattern for storage, it becomes immediately clear why this capability is so important. In most data centers, you buy a storage system and, when that nears its limits, you add another one. Because migrating data between storage systems has traditionally been such a difficult and disruptive task, old data tends to stay where it is while new data gets written to the new storage system. Because new data is usually more active, the storage workload shifts to the new storage system, which ends up busy—perhaps even overloaded—while the old storage system becomes underutilized from a performance standpoint.

With clustered Data ONTAP, you can easily redistribute storage workloads across the storage nodes in your cluster so these types of problems go away. When you add new nodes, you simply move new and

existing storage workloads to those nodes to balance the load. A cluster can contain several types of media—high-capacity HDDs, high-performance HDDs, and SSDs—so you can provide multiple tiers of storage in a single cluster and move datasets to the storage tier that best meets your needs—even when your needs change on a daily basis. Because you can move data dynamically to the most appropriate tier of storage, your storage costs are lower than under the old paradigm, in which high-performance storage tends to be overprovisioned and old and/or cold data tends to stay there indefinitely.

With the latest software release, clustered Data ONTAP 8.2, you can move data volumes even while replication or other activities are taking place, and you can perform controller upgrades without moving data volumes off the controller. That makes these operations even faster and easier to accomplish. Support for SMB 3.0 provides greater availability for Microsoft® Windows® Hyper-V™ environments as well. Earlier versions of the Microsoft SMB protocol were resilient to volume moves but failover events required clients to reconnect. A number of additional features further enhance Windows support, including:

- **ODX**, to offload data transfers from hosts to NetApp storage
- **Auto Location**, a capability unique to NetApp that enables clients to have the most direct path to storage
- **BranchCache**, to allow Windows clients to cache data locally to increase performance, particularly over WAN connections
- **FPolicy**, for quota management and file auditing
- **File access auditing**, to support monitoring, evidence, compliance, and recovery needs

Table 1) Examples of nondisruptive operations.

Lifecycle Operation	Benefit
<ul style="list-style-type: none"> • Rebalance controller performance, capacity, and/or disk performance and utilization 	<ul style="list-style-type: none"> • Lower capex (upfront outlay) • Better lifecycle management • Eliminate hot spots (better performance, less risk of unplanned outage)
<ul style="list-style-type: none"> • Add storage controllers or disk shelves • Add hardware to controllers 	<ul style="list-style-type: none"> • Improve performance and density • Improve resiliency
<ul style="list-style-type: none"> • Upgrade storage controllers, disk shelves, cluster switches 	<ul style="list-style-type: none"> • Avoid disruptive technology refreshes
Maintenance Operations	
<ul style="list-style-type: none"> • Upgrade storage software 	<ul style="list-style-type: none"> • Leverage new features sooner
<ul style="list-style-type: none"> • Upgrade system, disk, switch, firmware 	<ul style="list-style-type: none"> • Eliminate potential risks
<ul style="list-style-type: none"> • Replace failed controller or component within a controller such as a NIC, HBA, or failed storage components such as cables, drives, I/O modules 	<ul style="list-style-type: none"> • Reduce administrative overhead

3.2 Eliminate the Pain of Technology Refresh

When it comes time to replace your storage hardware, the process has always been disruptive, time consuming, and expensive. In fact, [a recent study¹](#) suggests that migrating data from an old storage array to a new one takes about five months to accomplish on average and adds almost 50% to the cost of owning an array.

Clustered Data ONTAP lets you accomplish full hardware refreshes quickly and easily without taking data offline, thereby avoiding these hidden costs. Because clustered storage systems aren't required to be the same generation or model, you can replace one FAS platform with another—or change your entire storage infrastructure—without interrupting running applications or busy users. No other storage vendor can do this.

3.3 Simplify Data Protection

To further protect your operations, clustered Data ONTAP provides integrated data protection (IDP) technologies with near-instant backup and recovery, replication for disaster recovery, and best-in-class integration with enterprise backup vendors and leading applications. IDP extends to include integrated and unified disk-to-disk backup and DR in a single process for VMware® and Microsoft virtualization.

Clustered Data ONTAP 8.2 includes support for SnapVault® disk-to-disk backup at the volume level, allowing you to replicate data to secondary storage for online backup and recovery. Because SnapVault only replicates changed blocks and preserves compression and deduplication savings it uses network bandwidth and secondary storage very efficiently.

4 Proven Efficiency

Clustered Data ONTAP is a highly virtualized storage operating system that abstracts physical storage into a set of storage virtual machines in order to deliver policy-based storage services for SAN and NAS with native multi-tenancy. Quality of service and data protection can be managed at the SVM level, so you can be more responsive, provision based on the required service level, and realize operational benefits from preset policy-based security and delegation.

Clustered Data ONTAP makes both your storage and your IT staff more productive and efficient, so you can scale your storage infrastructure without scaling your IT organization. A common set of features and procedures simplifies complex tasks so your IT staff can focus on solving higher-level problems and your operational efficiency grows as scale increases.

With clustered Data ONTAP you can consolidate and share the same infrastructure for workloads or tenants with different performance, capacity, and security requirements; achieve storage cost reductions with comprehensive storage efficiency; and take advantage of flash to optimize both capacity and performance.

Management processes scale so that twice as much storage no longer means twice as much work. All storage services running on clustered Data ONTAP—whether deployed for SAN, NAS, or both—use a common set of management tools. NetApp OnCommand® software allows you to automate, virtualize, and manage service delivery and SLAs. You can rapidly deploy resources and redeploy as your business and IT needs change, while shrinking backup windows and reducing infrastructure requirements.

¹ http://wikibon.org/wiki/v/Software-led_Infrastructure:_Moving_to_Virtualized_Clustered_Storage

4.1 Native Multi-tenancy and Quality of Service (QoS)

A storage cluster can be subdivided into distinct SVMs, each governed by its own rights and permissions. SVMs are used to securely isolate individual tenants—for instance, in a service provider environment—or individual applications, workgroups, business units, and so on. Because an SVM isn't tied to particular physical resources, you can move an SVM or adjust its resources without disruption.

Each application or tenant typically has its own SVM, and that SVM can be managed by the application owner or tenant. Application-driven storage services, available through our OnCommand plug-ins and APIs, allow application owners to automatically provision, protect, and manage data through the application management tools they are already familiar with.

QoS workload management allows you to control the resources that can be consumed by storage objects such as volumes, LUNs, VMDKs, or SVMs to manage performance spikes and improve customer satisfaction. You can create limits in terms of bandwidth (MB/sec) or throughput (I/O operations per second) to achieve fine-grained control. When a limit is set on an SVM, the limit is shared for all objects within that SVM. This allows you to set a performance limit on a particular tenant, but it leaves the tenant free to manage the assigned resources however they choose within that limit.

QoS gives you the ability to consolidate many workloads or tenants on a cluster without fear that the most important workloads will suffer or that activity in one tenant partition will affect another.

4.2 Storage Efficiency and Flash

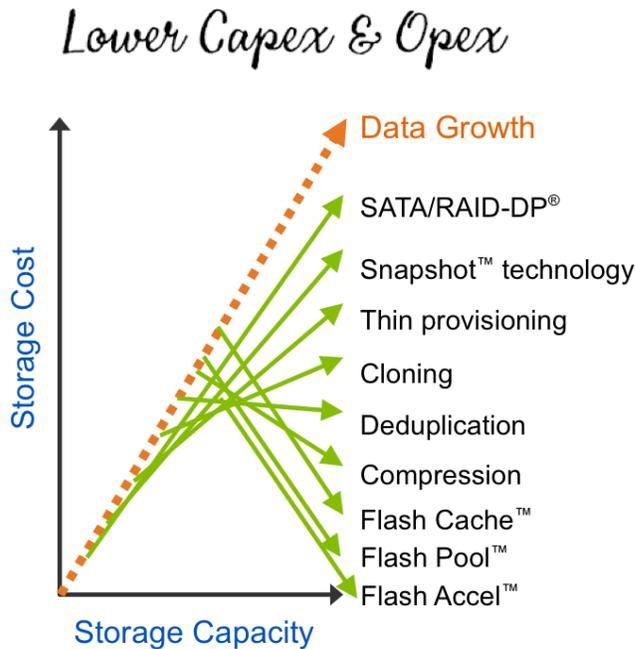
Pervasive cost reductions in clustered Data ONTAP result from the most comprehensive storage efficiency offering in the industry, including innovative Snapshot™ copies, replication and cloning technologies, thin provisioning, compression, and deduplication. NetApp is the only storage provider to deliver proven efficiencies for both SAN and NAS on entry-level, midtier, enterprise, software-based, and virtualized third-party arrays.

NetApp uses a multipronged approach to drive up your efficiency and drive down costs. We combine dual-parity RAID (RAID-DP® technology) for protection from disk failures with proven data reduction techniques and increase overall storage utilization. Our cloning technology reduces the need for full data copies, dramatically reducing your storage consumption during development and test, data mining, and other activities while making it easier and faster for people to access the datasets they need.

Clustered Data ONTAP supports a full range of flash technologies. The NetApp Virtual Storage Tier results in hybrid storage that combines the performance of flash with the capacity of hard disk drives. Hot data is automatically cached in flash, accelerating applications. You achieve optimal performance with little or no tuning, no moving data across media, and no ongoing management. Virtual Storage Tier technologies work in concert with other NetApp storage efficiency technologies such as deduplication to deliver optimal results and maximum performance from your flash investment.

When used together, these efficiency technologies significantly affect the data growth curve, allowing you to do more with less storage and driving down both your capital expenses (capex) and operational expenses (opex).

Figure 4) Storage efficiency technologies change the data growth curve and drive down costs.

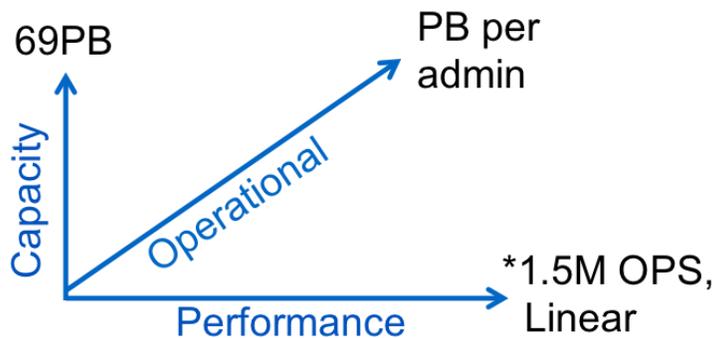


5 Seamless Scalability

A clustered storage environment running Data ONTAP scales in three dimensions: capacity, performance, and operations. You can scale SAN and NAS capacity from terabytes to tens of petabytes transparently and without reconfiguring running applications. Performance scales linearly as cluster nodes are added; a single administrator can manage petabytes of storage. You can start with a single cluster node and seamlessly expand your cluster up to 24 nodes as your business needs grow.

Figure 5) Clustered Data ONTAP scales in several dimensions.

Scale in 3 dimensions



*For more information, visit <http://www.spec.org/sfs2008/results/>.

The nondisruptive-operations capabilities of clustered Data ONTAP make your storage environment much more flexible. You can easily rebalance capacity and storage workloads as needed. You can improve service levels by dynamically redeploying workloads and avoid hot spots by moving volumes to less active disks or spreading workloads across more controllers. Each dataset gets the right technology to meet your performance and cost targets.

In clustered Data ONTAP 8.2 we've also significantly increased many of the limits to make the platform even more scalable. This includes support for:

- Up to 100,000 NFS clients
- Aggregates up to 400TB in size
- 12,000 volumes in a 24-node NAS cluster
- 49,152 LUNs in an 8-node SAN cluster
- Multiple Infinite Volumes, each up to 20PB in size
- Microsoft BranchCache v2, which allows Windows servers to cache data from a NetApp cluster locally to improve performance

5.1 A More Flexible Approach to Scale-Out

Most scale-out storage solutions provide a single large repository with limited control. Essentially, this means a single class of service for every workload sharing the storage.

The virtualized storage services and unique capabilities of clustered Data ONTAP give you much more flexibility and greater control. With NetApp Infinite Volume, you can have a single repository or several large repositories to match your requirements, each capable of scaling to billions of files and petabytes of capacity. Any SVM can include an Infinite Volume, and you can control and adjust the resources available to the SVM—and the Infinite Volume—on the fly.

Clustered Data ONTAP gives you the ability to isolate particular workloads or tenants and offer different levels of service. You can mix certain controller types, offer different storage tiers, and define QoS policies to address a wide variety of storage needs, all from the same unified infrastructure. You can also virtualize third-party arrays with NetApp V-Series, and when it's time to retire a storage system, you can simply upgrade the controllers—keeping data in place.

6 Conclusion

Clustered Data ONTAP is a revolutionary approach to storage software that can help you remove IT constraints and speed response to business changes. With clustered Data ONTAP you can eliminate planned downtime and quickly adjust your storage infrastructure to meet the needs of your business. The result is an infrastructure that works without interruption, keeps costs under control, facilitates new projects, and delivers capacity and services instantly.

NetApp provides no representations or warranties regarding the accuracy, reliability, or serviceability of any information or recommendations provided in this publication, or with respect to any results that may be obtained by the use of the information or observance of any recommendations provided herein. The information in this document is distributed AS IS, and the use of this information or the implementation of any recommendations or techniques herein is a customer's responsibility and depends on the customer's ability to evaluate and integrate them into the customer's operational environment. This document and the information contained herein may be used solely in connection with the NetApp products discussed in this document.

Go further, faster®



www.netapp.com

© 2013 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, AutoSupport, Data ONTAP, Flash Accel, Flash Cache, Flash Pool, OnCommand, RAID-DP, Snapshot, and SnapVault are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Microsoft and Windows are registered trademarks and Hyper-V is a trademark of Microsoft Corporation. VMware is a registered trademark of VMware, Inc. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. WP-7186-0613