



A NETAPP IT EBOOK FOR STORAGE ADMINISTRATORS

7 TIPS & TRICKS TO MAKE STORAGE ADMINISTRATION FASTER & EASIER



7 Tips & Tricks to Make Storage Administration Faster & Easier

BLOGS

- 1: [Introduction](#)—Jeff Boni 3
- 2: [Automating cDOT Configuration: How to Make a 4-Hour Process Only 5 Minutes](#)—Ezra Tingler 4
- 3: [The Importance of IO Density in Delivering Storage as a Service \(Part 1\)](#)—Eduardo Rivera 6
- 4: [The Role of QoS in Delivering Storage as a Service \(Part 2\)](#)—Eduardo Rivera 8
- 5: [The Demystification of NetApp Private Storage](#) —Mike Frycz 10
- 6: [Using AltaVault and StorageGRID to Replace Tape Backups](#) —Pridhvi Appineni 12
- 7: [How We Balanced Power & Capacity in Our Data Centers with Flash & Other Innovations](#)—Randy Egger 14
- 8: [How NetApp Products Help NetApp IT Overcome Data Protection Challenges](#) —Dina Ayyalusamy 16



Introduction

JEFF BONI, VP FOUNDATIONAL SERVICES, NETAPP IT

The storage admin role is changing rapidly thanks to virtualization and automation.

The role of a storage administrator in IT is changing rapidly as virtualization and automation improve efficiency and productivity. These trends give face to a series of new challenges: storage service levels that adjust to changing application performance/capacity

requirements and integration with other layers of the stack (compute, networking and the cloud). The new storage environment requires a new skill set that relies on proactive data analysis and looks to a future of hyper-converged infrastructures and hybrid cloud.

This ebook examines the many issues present in today's storage environment. NetApp IT practitioners share their experiences, with an emphasis on the best practices IT has adopted to improve its service delivery to the business:

- Automating ONTAP configurations ([ONTAP®](#))
- Designing storage service levels ([OnCommand® Insight/ONTAP](#))
- Adopting the hybrid cloud (Data Fabric)

- Demystifying NetApp Private Storage (NPS) for the hybrid cloud ([NPS](#), Data Fabric)
- Replacing tape backup with cloud and object storage ([AltaVault®](#) and [StorageGRID®](#))
- Overcoming data protection challenges ([Snap](#) products, [FlexClone®](#))

We invite you to take the next step and ask NetApp IT experts to share their real experiences in using NetApp products and services, including All Flash FAS and OnCommand Insight, in the NetApp production environment. Ask your NetApp sales team to arrange an interactive discussion with us soon.



*Jeff Boni, VP Foundational Services,
NetApp IT*



Automating ONTAP Configurations: How to Make a 4-Hour Process Only 5 Minutes

EZRA TINGLER, SENIOR STORAGE ENGINEER, NETAPP IT

I wrote a script to do a 4-hour storage cluster configuration in 5 minutes.

As a senior storage engineer in our Customer-1 organization, I am responsible for storage lifecycle management including the installation, decommission, and capacity management of our ONTAP® and 7-Mode storage controllers. Our group is in the midst of moving all data hosted on 7-Mode storage controllers to ONTAP storage clusters.

Business Challenge

As part of our migration, we are installing additional ONTAP clusters and nodes. The configuration of each high availability (HA) pair took about four hours, spread out over 2 to 3 days. The four hours did not include the time needed to configure the cluster inter-connect switches or initialize disks; this takes 2 to 12 hours depending on disk type. Plus typical office interruptions added more time as I had to figure out where I had left off. This sporadic schedule seemed to result in some configuration inconsistencies.

The Solution

I challenged myself to see if I could automate the process to save time and reduce errors. Although I'm not a developer, I found it easy to write the script using the [NetApp Software Development Kit](#) (SDK). I run the

script after the disks are initialized, cluster setup is complete, and the cluster inter-connect switches are properly configured. The script reads configuration information from a file, then applies the configuration to the cluster nodes. It does so by accessing the nodes via ZAPI calls, which is why it is fast.

The results have been amazing. The four-hour process now takes about five minutes to complete 99% of the configuration. It is now possible to install 24 nodes in two hours rather than 96 hours, a time savings of 94 hours or 2½ work weeks. Errors caused by interruptions have been eliminated. Automating this process has freed up my time to work on other projects.

If you are a storage admin, you can easily do this yourself with the SDK. I used an SDK tool called [Z-Explorer](#) that contains a complete list of all ZAPI calls for the cluster. With Z-Explorer most of the development work is done for you. It took me just three weeks to automate all the builds. This KnowledgeBase article is a good place to start.

It was a fun project because I could write the script without feeling like I had to be a developer. I wrote the scripts in PERL, but the SDK works with any language you are familiar with. I also used the [SDK online forum](#) to get advice from others. People were quick to answer my questions.

The Future

I'm now using the SDK to automate and streamline other storage tasks to save time and reduce errors. My next project is a quality assurance (QA) script that will login to a cluster and verify if nodes are properly configured per NetApp IT Standards and NetApp best practice guidelines. I plan to automate the cluster interconnect switch configuration in the same way as well as e-Series configuration.

Read the [story](#) in the Tech ONTAP newsletter.

Check out the [NetApp Software Development Kit](#).

Find the script in the NetApp Automation Store. Search: Day-0 c-mode cluster build setup.



*Ezra Tingler, Senior Storage Engineer,
NetApp IT*



The Importance of IO Density in Delivering Storage as a Service (Part 1)

EDUARDO RIVERA, SENIOR STORAGE ENGINEER, NETAPP IT

Can NetApp IT deliver storage as a service?

NetApp IT posed this question to itself more than a year ago. Our goal was to find a new way to offer our business customers a method by which they could consume storage that not only met their capacity requirements, but also their performance requirements. At the same time, we wanted this storage consumption model to be presented as a predictive and easily consumable service. After consulting with enterprise architects for NetApp's cloud provider services, we developed a storage service catalog leveraging two main items: IO Density and NetApp ONTAP's QoS (quality of service).

In this first part of this two-part blog, we will discuss how NetApp OnCommand Insight's IO Density metric played a key role in the design of our storage service catalog. (You can also hear this as a [podcast](#).)

The Role of IO Density


IO Density is a simple, yet powerful idea. The concept itself is not new, but it is essential to building a sound storage consumption model. By definition, IO Density is the measurement of IO generated over a given amount of stored capacity and expressed as IOPS/TB. In other words, IO Density measures how much performance can be delivered by a given amount of storage capacity.

The next step in storage management is storage service design.

[Here's an example of how IO Density works. Suppose we have a single 7.2K RPM drive. By rule of thumb, a single drive of this type can deliver around 50 IOPS @ 20ms response time. Consider, however, that 7.2K

RPM drives today can range anywhere from 1TB to 8TB in size. The ability of the drive to deliver 50 IOPS does not change with its size. Therefore, as the size of the drive increases, the IOPS/TB ratio worsens (i.e. you get 50 IOPS/TB with 1TB drive and 6.25 IOPS/TB with an 8TB drive).

Applying the same logic, we can divide the amount of capacity that we provision to a given application by the amount of IO that an application demands from its storage. The difference is that at the array level, there are many other technologies and variables at play that can determine the IO throughput for a given storage volume. Elements like disk type, controller type, amount of cache, etc., affect how many IOPS a storage array can deliver. Nonetheless, the general capabilities of a known storage array configuration can be estimated with a good degree of accuracy given a set of reasonable assumptions.]



Using OnCommand Insight we were able to gather, analyze, and visualize the IO Density of all the applications that run on our storage infrastructure. Initially, what we found was surprising. Some applications that anecdotally were marked as high performance were demonstrating very low IO Density rates, and thus were essentially wasting high-performance storage capacity. We also saw the reverse, where applications were pounding the heck out of lower performance arrays because their actual IO requirements were incorrectly estimated at the time of deployment. Therefore, we started to use NetApp OnCommand Insight's aggregated IO Density report to profile application performance across the entire infrastructure and establish a fact-based architecture.

Ultimately, OnCommand Insight's IO Density report helped us to identify the range of service levels (defined as IOPS/TB) that the apps actually needed. With this information, we created a storage catalog based on three standard service levels:

1. Value: Services workloads requiring between 0 and 512 IOPS/TB.
2. Performance: Services workloads requiring between 512 and 2048 IOPS/TB.
3. Extreme: Services workloads requiring between 2048 and 8192 IOPS/TB.

Based on our own understanding of our application requirements (as depicted by our IO Density reports), the above three tiers would address

99 percent of our installed base. Those workloads requiring something other than these pre-defined workloads are easily dealt with on a case-by-case basis since there are so few of them.

A New Perspective on of Application Performance

IO Density gave us a new perspective on how to profile and deploy our applications across our storage infrastructure. By recognizing that performance and storage capacity go hand in hand, we were able to create a storage catalog with tiers that reflected the actual requirements of our installed base.

Our next step was placing IO limits on volumes to prevent applications from stepping on the performance resources of other applications within the same storage array. Stay tuned for part two of this blog where I will discuss how we used ONTAP's adaptive QoS feature to address this issue.

Tune into the Tech ONTAP [podcast](#) for more details.



*Eduardo Rivera, Senior Storage Engineer,
NetApp IT*



The Role of QoS in Delivering Storage as a Service (Part 2)

EDUARDO RIVERA, SENIOR STORAGE ENGINEER, NETAPP IT

Using ONTAP QoS policies we can more efficiently manage our storage service levels.

NetApp IT is on a journey to offer its customers storage as a service. In part one of this blog, I discussed how we embraced IO Density to help us better profile and deploy our applications across our storage infrastructure. We developed a three-tier service catalog that offers storage as a predictive and easily

consumable service to our customers. The second step in this journey was tapping into the power of clustered Data ONTAP®'s adaptive Quality of Service (QoS) feature to assure performance stability.

QoS—Corralling the Herd

The adoption of ONTAP's QoS feature is a key component of our storage-as-a-service model. In a nutshell, QoS enables us to place IO limits on volumes (it can also work at the storage virtual machine (SVM) or file level) in order to keep the applications using those volumes within their IOPS "swim lane." This prevents one application from starving other applications of performance resources within the same storage array. QoS can be implemented dynamically and without interruption to application data access.

In our storage catalog model, we assign a QoS policy per volume for all the volumes that exist within a given cluster. The QoS policies themselves

enforce a particular IOPS/TB objective. Hence, if we have a volume that is consuming 1TB of capacity and the service level objective (SLO) is to provide 2048 IOPS/TB, the QoS policy for that volume would set an IOPS limit of 2048. If that same volume in the future grows to 2TB of consumed space, then the QoS policy would be adjusted to 4096 IOPS/TB to maintain an effective 2048 IOPS/TB. In a live environment with hundreds, or even thousands, of individual volumes and where storage consumption continuously varies (as the application writes/deletes data), manually managing all the QoS policies would be close to impossible. This is where Adaptive QoS comes in.

Adaptive QoS is a tool developed by NetApp. Its sole purpose is to monitor consumption per volume and dynamically adjust each volume's QoS policy so that it matches the desired IOPS/TB SLO. With this tool, we are able to provision volumes at will and not worry about having to manage all the necessary QoS policies.

With QoS and Adaptive QoS, we are able to easily provide predictive storage performance tiers upon which we can build the actual storage service catalog.

Building the Storage Service Catalog

With the pre-defined service levels and the ability to manage IO demand with Adaptive QoS, we were able to build a storage infrastructure that not only delivers capacity but also predicts performance. Leveraging ONTAP's ability to cluster together controllers and disks that offer various combinations of capacity and performance, we built clusters using different FAS and AFF building blocks to deliver varying tiers of performance. Then Adaptive QoS was used to enforce the performance SLO per volume depending on where that volume resides.

Moving a volume between service levels is also quite simple using ONTAP's vol-move feature. Adaptive QoS is smart enough to adjust the policy based on where the volume sits. By defining a service level per aggregate, we are also defining a multitude of service levels within a particular cluster through which we can move our data around. Addressing changes in performance requirements is easy; we move the volume to a higher performing high availability (HA) pair using vol-move.

Data-Driven Design

Together, IO Density and QoS have revolutionized how we view our storage. It has made us much more agile. The IO Density metric forces us to think about storage in a holistic manner because we operate according to a data-driven—not experience-based—storage model. We don't need to look at whether we have enough capacity or performance, but can check to see if we have enough of both. If we nail it, they run out at the same time.

The same is true with the QoS service level approach. Our storage infrastructure is much simpler to manage. ONTAP gives us granular control of resources at the volume level; our QoS policies now act as the controller. Best of all, this new storage approach should enable us to deliver a storage service model that is far more cost efficient than in the past while supporting application performance requirements.

Tune into the Tech ONTAP [podcast](#) for more details.



*Eduardo Rivera, Senior Storage Engineer,
NetApp IT*



The Demystification of NetApp Private Storage (NPS) for Cloud

MIKE FRYCZ, BUSINESS SYSTEMS ANALYST, IT SUPPORT & OPERATIONS, NETAPP IT

If you follow NetApp IT, you'll know that we talk a lot about our hybrid cloud strategy and roadmap. But one question comes up repeatedly: How do we operationalize the cloud? How does our cloud strategy translate into operational excellence, especially when it comes to taking advantage of hyperscaler cloud resources?

Our current cloud operations are actually fairly simple. They rely on three primary elements:

- NetApp Private Storage (NPS) for Cloud enables us to access cloud resources while maintaining complete control over our data.
- Cloud-connected colocation facilities, such as Equinix, allow the data to remain private just outside the cloud.
- Hyperscalers, such as Amazon Web Services (AWS), Microsoft Azure, and others offer flexible compute resources.

NPS for Cloud Architecture

To better understand what this all means, let's look at the physical architecture of NetApp IT and [NPS for Cloud](#), as shown in the graphic. NetApp's FAS system connects to the AWS and Azure compute via a dedicated network connection within an Equinix data center. We will connect to other hyperscalers in the future.


Our FAS system is physically deployed in racks located inside a NetApp cage, similar to that shown. The minimum is two nodes for high availability. The FAS system is managed by an off-site storage team.



The FAS system connects to a layer 3 network switch, patched to an Equinix patch panel through a cross connect.



The Equinix cross-connect uses single-mode fiber (SMF) optic cables that run through a large, yellow overhead tray and down the aisles of the Equinix facility to the cloud peering switch in the AWS and Azure cages.





The cable directly connects to AWS and Azure inside their respective cages. Given the close physical proximity of the storage and data to the hyperscaler, we now can access a high bandwidth (10GB) Ethernet connection from our data center by way of NPS to the cloud.



Our data resides in NetApp storage, but our compute is accessed in AWS or Azure. We currently operate our legal, human resources, branding, customer service support, and various other portals using NPS for Cloud.

Keeping Control of Our Data

The single most important benefit to NetApp IT of using NPS for Cloud is that we keep control of our data. We use the [SnapMirror®](#) feature of ONTAP® to replicate the data from our on-premises data centers to NPS then to AWS or Azure. The NetApp Data Fabric enables us to connect to and switch cloud providers at any time. We avoid vendor lock-in and costly data migrations.

Is NPS for Cloud really that simple? Yes. And its benefits are numerous:

- Ability to keep control of data at all times
- High-throughput, direct connections to the cloud
- Ability to rapidly scale our compute or secure run-time resources for peak workloads
- Centralized storage intelligence using OnCommand® Insight and data management through NetApp ONTAP® software
- Compliance with the security and privacy requirements of companies and governments
- Migration flexibility so applications can be easily moved between different clouds

Our next phase is to work with Business Apps to build cloud-aware apps that take advantage of the many benefits of the cloud, such as platform-as-a-service (PaaS) and DevOps. The cloud is definitely a key part of our strategy to excel at IT service delivery inside NetApp.

Download the [infographic](#) from NetAppIT.com.



Mike Frycz, Business Systems Analyst, IT Support & Operations, NetApp IT



Using AltaVault and StorageGRID to Replace Tape Backups

PRIDHVI APPINENI, SENIOR MANAGER, IT STORAGE SERVICES, NETAPP IT

One of the business challenges NetApp IT faces is archiving our legal, finance, and Sarbanes-Oxley (SOX) compliant data. Backing up this data is important for legal, HR, and tax reasons. In some cases, the data must be secured for seven years for tax purposes. Like most companies, we have relied on tape backups to secure this data. Tapes are reliable, inexpensive, and present very little risk to our operations.

My IT storage team was intrigued by the use case that NetApp® [AltaVault](#)® and NetApp [StorageGRID](#)® offered. AltaVault cloud-integrated storage functions as a cloud gateway for backup and archive applications. StorageGRID provides an enterprise-grade object storage solution that supports widely adopted cloud protocols such as Amazon S3. The combination of AltaVault and StorageGRID would enable us to efficiently back-up application data while optimizing data protection and simplifying disaster recovery.

NetApp IT's core tenet is to adopt technologies only when they satisfy a business use case. We evaluated these products and came to the conclusion that AltaVault and StorageGRID would be a powerful combination to modernize our backup procedures, reduce our costs, and, most importantly, improve the speed with which we can restore data for our customers.

The business case for using AltaVault and StorageGRID to replace tape backups is compelling.

Powerful Combination

Because they take advantage of the cost and scale benefits of cloud storage, AltaVault and StorageGRID are designed for an enterprise like NetApp with locations worldwide.

AltaVault delivered benefits such as 30:1 inline deduplication,

compression and encryption technology, which makes archived data easier to transport and store, and faster to retrieve. It offers complete security for the data. We can categorize that data into buckets to make it more easily retrievable. We are currently seeing 22 times deduplication savings from the data stored in AltaVault. As we push more data through AltaVault, we will benefit from even greater savings.

StorageGRID enables us to store and manage these massive datasets in a repository in the hybrid cloud. It enables us to abstract storage resources across multiple logical and/or physical data centers. We also create data policies to manage and protect our data according to our requirements.

Changing Our Archiving Architecture

Previously, critical data from all our locations was backed up in our Sunnyvale, California data center. We use backup software to manage the flow of the data from backup storage into a tape library in Sunnyvale. We defined when and how the archiving took place. The tapes were regularly transported to an off-site location for safekeeping. When data needed to be restored, we had to order the tapes from the vendor and physically transport them back to our site, which took at least 24 hours.


Under the new target architecture, the process remains virtually the same. First, AltaVault provides a local optimized disk cache for the application to store backup data, resulting in faster restores for end users and saving bandwidth. One copy is stored in the StorageGRID nodes in Sunnyvale. Then the data is copied to StorageGRID nodes in our Raleigh, North Carolina data center, which serves as a repository for the offsite data copy. The complete cutover process took about five months.

The benefits have been numerous. We eliminated the cost of transportation, storage, tape library support, and physical tape procurement. Bringing backup in-house has enabled us to automate many of the day-to-day operational processes, resulting in a more agile service. We can also retrieve the archived data in one to six hours, depending on the data set size, or 3 times faster than before. This translates to a much faster turnaround for our customers. We anticipate improved backup speeds and significant cost savings in the future

We also gained flexibility. We can more easily modify archive policies and add new archives on an ad-hoc basis. AltaVault allows us to scale our archive/SOX environment much faster than we could with a tape library. For example, we can spin up a virtual machine with AltaVault using existing disk shelves to gain capacity as opposed to purchasing more tape drives and a new frame for the tape library. Long term, the combined software will make it much easier to transition our backup procedures to new data centers as our data center footprint evolves.

Faster, More Reliable Data Archiving

One of the most satisfying parts of this project has been seeing firsthand the impact NetApp products can have on our operations. Not only will we improve efficiency and reduce costs, but we also will improve the data archiving services we provide to our business customers. As a team that constantly strives to use the best technologies to improve service delivery, that is the best result of all.




Pridhvi Appineni, Senior Manager, IT Storage Services, NetApp IT



How We Balanced Power & Capacity in Our Data Centers with Flash & Other Innovations

RANDY EGGER, DATA CENTER LEAD, NETAPP IT



One of the biggest trade-offs in any data center is power and capacity, the two biggest expenses of any data center. The golden rule is that these two costs increase together—the more racks of hardware you have, the more power you need to run it. This means when you need more capacity, you need more power, which could result in a cooling issue. If you have enough cooling and power, you could run out of rack capacity.

NetApp IT was able to address the power and cooling costs in a multitude of ways. We started by making changes to the facility itself. We installed non-traditional raised floors. We introduced overhead cooling, economization, and cold aisle containment over six years ago. These changes have helped control our power and cooling costs.

Changing Relationship between Power and Capacity

A NetApp IT data center operation analysis compiled over the past decade shows that the relationship between power and capacity is evolving due to other factors as well. We are seeing that while our compute and storage capabilities are increasing, our power costs have been actually dropping. This shift is due to several reasons: the availability of the cloud, smaller form factors offering more horsepower, and virtualization, among others.

The chart illustrates this point. Our power requirements peaked in mid-2011 when we opened a new NetApp production data center, the Hillsboro Data Center (HDC). As we moved operations into HDC and closed other data centers, power consumption dropped while storage and compute increased significantly. Since then we've seen this trend continuing.

The following factors are contributing to this change:

- **Virtualization.** In the past, each app had its set of hardware and its own power supply, which translated to thousands of servers, an expensive model to maintain. Because of virtualization, the same applications can be hosted on 10 to 20 physical machines in a few racks using around 20 kilowatts (kW). NetApp IT's compute is 75% virtualized now.
- **All Flash FAS adoption.** Our solid-state disks (SSD) take very little power (1-2kW as compared to 5-6kW for traditional disks per rack); our Flash hardware even less. As result, full storage racks aren't even close to reaching their power limits. Using [Flash](#) for all non-archival storage going forward means even lower power consumption.
- **High-density storage rack design.** HDC has high-density, taller-than-normal racks, 52U as opposed to traditional 42U or 47U racks with more power (10kW/rack). Hardware that used to take four racks

now takes half of a rack, thanks to higher density disks and higher IO capability clusters/filers. This unique design has reduced the number of infrastructure/connection points, shrinking the data center footprint and enabling a build-as-you-grow approach.

- **FlexPod® datacenter.** We have eight [FlexPod](#) systems hosting hundreds of applications in a rack connected to the Cisco fabric for networking and compute. The applications are hosted on thousands of machines, but thanks to virtualization and cloud, that doesn't mean thousands of physical servers. Most of the applications are hosted on virtual machines. These footprints will continue to shrink as compute core processor power increases, hardware size shrinks, and power consumption requirements fall due to technology advancements.
- **Smart power design.** The Starline busway system supports 'anywhere' power and connector types, and with our smart layout we can utilize power borrowing that enables us to share power across multiple racks.

We can draw power from a parallel busway if a rack needs more than 9kW. We have effectively removed power as a consideration in our hardware installations.

Our analysis shows that the relationship between storage/compute capacity to deliver applications and power will

continue, even as we begin to take advantage of the hybrid cloud. Instead of building arrays to meet peak workloads--which translates to idle capacity--we will be able to take advantage of the cloud's elasticity. This, in turn, will reduce operational, licensing, and management costs.

Adopting a hardware lifecycle management strategy is a key factor in reducing power consumption and improve capacity management. In our HDC migration, we were able to decommission 96 of 163 systems and 40 filers (of 2 PB of storage); more than 1,000 servers were either migrated or decommissioned. The configuration management database (CMDB), NetApp IT's single source of truth for everything in IT operations, also plays a major role in helping us track, manage, and analyze power and capacity over time.

Each company faces its own challenges in controlling its power consumption costs while maximizing its storage and compute. However, as we have seen, adopting a hardware lifecycle management strategy and leveraging innovations in technology and power design can make a significant difference.

While our capacity is increasing, our power costs are dropping.



Randy Egger, Data Center Lead, NetApp IT



How NetApp Products Help NetApp IT Overcome Data Protection Challenges

DINA AAYYALUSAMY, LEAD DATABASE ADMINISTRATOR, NETAPP IT


As a database administrator (DBA), I face very different challenges than I did five years ago. The adoption of the hybrid cloud, rising security and backup concerns, and the expectation of non-disruptive service are rapidly changing the IT environment. Like many IT shops, NetApp IT is doing more with less, including supporting a large number of databases (400 plus) on a smaller budget.

While performance management remains a top priority, just as critical are issues such as disaster recovery, auditing, compliance, and maintenance.

I rely on a variety of tools, including

NetApp's data protection products, to manage our databases. In this blog, I'll share how our team uses NetApp products in these four areas:

- Business continuity/disaster recovery
- Performance management
- Capacity management
- Auditing & Sarbanes Oxley (SOX) compliance



How NetApp Products Help Us Overcome These Challenges

Our team uses a variety of NetApp's data protection products, including [SnapCenter](#)®, [SnapMirror](#)®, [FlexClone](#)®, and [SnapVault](#)® in our everyday routine. ONTAP® is the underlying software that helps automate our enterprise application lifecycle tasks. I will start with a discussion of

SnapCenter, our management tool.

SnapCenter, NetApp's data protection and clone management tool combines management and backup of our many databases into one central platform. We use SnapCenter to simplify our storage planning, backup, and restore operations. For example, in the past we manually specified daily file backups, which was both time consuming and prone to errors. This process is now

completely automated. We also use SnapCenter to:

- Provide automatic scheduling of backups at the volume level, not file level. This ensures regular and quality backups and makes it easier to scale our operations.

How NetApp Products Help NetApp IT Overcome Data Protection Challenges

*NetAppIT uses
NetApp data
protection products
in its IT operations
every day.*

- Perform daily database refreshes. Using an end-to-end workflow from production through backup eliminates the many manual tasks associated with tracking and backing changes.
- Ensure automatic backups of Sarbanes Oxley (SOX) and other compliance-related data. With SnapCenter we send the backup data to the cloud using AltaVault®. (See [blog](#).)
- Grant users the ability to manage their application-specific backup/restore/clone jobs without our intervention.

For the past eight years we have used [SnapManager](#)® for SQL Server® to run all our SQL database backups in one location. Currently, we run the SQL feature as a separate product, but we will be moving to a new SQL plug-in for SnapCenter, which means one less tool to manage and more efficient SQL server management.


SnapMirror, a data transfer feature of ONTAP, is a critical tool in our database management arsenal because of its ability to compress and replicate data between Point A and Point B. We use it to ensure block-level changes are replicated in the database. It is an invaluable tool for generating multiple copies for application development/test, reporting, cloning, and disaster recovery replication. We also use SnapMirror to:

- Set up automated scheduling for the recurring refreshes of critical databases, such as those that are compliance-related (SOX, audit), making the process both consistent and predictable.
- Support high availability (HA) requirements; we can recover a database in minutes instead of hours thanks to SnapMirror's replication and compression features.

- Copy files during a data center migration; SnapMirror can copy files in a fraction of the time, reducing application downtime.
- Provide lifecycle management for database clones, accelerating application development/test.

FlexClone is a fast, efficient functionality we rely on for automation and capacity management. The thin provisioning capability delivers multiple instant, point-in-time, space-efficient copies of replicated volumes and logical unit numbers (LUNs), saving terabytes of disk space. SnapMirror and FlexClone work hand in hand to enable stored data to be used for dev/test or data mining. While SnapMirror tracks changes over time, FlexClone generates point-in-time copies that are essential for rapid application development/test. We use SnapMirror to replicate data on a different filer, then spin these off using FlexClone. And because a clone uses only a small amount of space for metadata, it only consumes additional space as data is changed. We can use it with both Oracle and SQL databases. We use FlexClone to maintain performance and automate many of our services including:

- Spin off a copy when we have a critical issue in a large database. The FlexClone version is ideal for troubleshooting while the production version keeps running.
- Generate copies of disaster recovery volumes for use during application testing so we don't need to break the SnapMirror relationship with the database, eliminating the loss of data backups.
- Create a database copy from which management reports can be generated, enabling application development/test to use the untouched production database.



- Migrate very large databases without business interruption for pre-cutover testing.
- Provide a quick return to service if a server or storage goes down in our Oracle 11G environment; FlexClone's schema-as-a-service solution enables point-in-time recovery of individual schemas.
- SnapVault, the HA backup and recovery feature of ONTAP, is used to protect data with unique requirements, such as compliance data. In the past, we had to manually move a database to storage, then move it to a vault. In the latest release we can transfer from production directly to the vault, which is much more efficient and requires no manual intervention. With SnapVault we can store data on a filer and then capture user changes over time.

SnapVault is also used for keeping multiple copies of production databases for code releases. If developers want to retrieve a database from three releases ago, they can take multiple snapshots of a database, store it in a vault, then restore it to a point-in-time as needed.

ONTAP as the Foundation

Our use of NetApp products relies on the underlying ONTAP software. ONTAP supports features such as self-provisioning to automatically handle growth in filer volumes without human intervention. It also enables transparent migration between nodes in the event of any storage changes or failures without disruption to the databases. Its non-disruptive feature is essential to ensuring continuous access to databases during updates, migrations, and other volume-related changes.

These NetApp products have been instrumental in helping our database team work more efficiently and provide fast, efficient data replication and disaster recovery. We can meet recovery-point objectives ranging from minutes to hours. We keep both the active mirror and prior backup copies to enable selective failover points in the disaster recovery copy. These products—along with rigorous work processes—help us protect our data while maximizing our database performance in a wide variety of business and IT environments.



Dina Ayyalusamy, Lead Database Administrator, NetApp IT



7 Tips & Tricks to Make Storage Administration Faster & Easier

*7 ways to make your storage administration faster & easier.
#NetAppIT*

The NetApp on NetApp Program shares its real-world IT experiences in using NetApp products and services in a global enterprise IT environment. Our subject matter experts speak with representatives from other IT organizations about common IT challenges and best practices and the business cases driving product adoption. To learn about NetApp IT or to speak with one of our subject matter experts, talk to your NetApp sales representative or visit www.NetAppIT.com.

Read our [blogs](#).

Read our other ebooks:

- [Building a Foundation for Business Apps Agility](#)
- [7 Perspectives on the Future of IT: The Drive Toward Business Agility](#)

Refer to the Interoperability Matrix Tool (IMT) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 1994–2016 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

Trademark Information

NetApp, the NetApp logo, AltaVault, ASUP, AutoSupport, Campaign Express, Cloud ONTAP, Clustered Data ONTAP, Customer Fitness, Data ONTAP, DataMotion, Fitness, Flash Accel, Flash Cache, Flash Pool, FlashRay, FlexArray, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexVol, FPolicy, GetSuccessful, LockVault, Manage ONTAP, Mars, MetroCluster, MultiStore, NetApp Insight, OnCommand, ONTAP, ONTAPI, RAID DP, RAID-TEC, SANtricity, SecureShare, Simplicity, Simulate ONTAP, SnapCenter, Snap Creator, SnapCopy, SnapDrive, SnapIntegrator, SnapLock, SnapManager, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapValidator, SnapVault, StorageGRID, Tech OnTap, Unbound Cloud, WAFL and other names are trademarks or registered trademarks of NetApp Inc., in the United States and/or other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. A current list of NetApp trademarks is available on the Web at <http://www.netapp.com/us/legal/netapptmlist.aspx>.

NA-20161003