A NETAPP IT EBOOK

NETAPP ON NETAPP: SHARING OUR IT EXPERIENCES AS CUSTOMER-1

NetApp

NetApp on NetApp: Sharing our IT Experiences

BLOGS

www.NetAppIT.com **Twitter: @NetAppIT**

Introduction to NetApp on NetApp

As the corporate IT shop for a global enterprise, NetApp IT faces many of the same challenges as NetApp's customers. Under the NetApp on NetApp program, IT subject matter experts share their knowledge in both operationalizing NetApp products and adopting IT best practices. We have experience in using NetApp products, including:

- Flash, AltaVault, E-Series, FAS, StorageGrid
- Data Fabric, Hybrid Cloud
- Storage management (ONTAP, OnCommand Insight)
- IT best practices: Office 365, CMDB/ServiceNow, knowledge management

This eBook offers the perspectives of some of our most experienced NetApp IT practitioners on subjects ranging from the cloud to Flash to automating storage capacity management. We hope you can learn from their experiences as Customer-1 and apply them to your own IT environment.

To read our other eBooks or subscribe to our blog, visit www.NetAppIT.com.

Sincerely,

Bill Miller, CIO, NetApp

Three Steps to Defining a Cloud-First Strategy

MATT BROWN, CUSTOMER ENGAGEMENT AND NETAPP ON NETAPP PROGRAM, NETAPP IT

As NetApp's business goals evolve, so must IT. We constantly look at ways to streamline IT processes to become more efficient in terms of cost and agility. When we looked to the cloud as a possible solution, we discovered the real conversation has evolved from 'to cloud or not to cloud?' but 'how to cloud and why.'

The cloud decision framework is an effective tool to standardize support of core vs. non-core systems and embrace a dynamic hybrid cloud strategy.

Inside NetApp IT we previously looked at the cloud as a static end state where apps were placed in the cloud on an individual basis. We then realized that we could leverage the cloud in a far more dynamic way. A dynamic strategy enables us to move apps in, out, and within clouds; adjust our course as workloads and risk factors change; and continually evaluate what apps are in the cloud and why. Our strategy is based on a blend of private and public clouds (which includes SaaS providers) and aims to eliminate the legacy data center concept.

One of the biggest mistakes an IT shop can make is to manage the cloud as if it were a project with an end date. IT organizations must acknowledge that they need the cloud as part of their new operating model. Adopting this new model requires thoughtful planning in terms of agility, scalability, and supportability. Below we offer an approach to starting your own cloud journey.

Next Stop: Cloud

Many CIOs I talk with ask me why NetApp IT decided to move to the cloud. My response is that all enterprises need to go to the cloud, whether it is private, public/SaaS or, most likely, a blend of all three. NetApp IT sees the cloud as enabling the agility to maximize technology investments, deliver business capabilities more rapidly, ensure greater supportability, and enhance the ability to leverage future technologies and IT services.

With the goal that all applications will end up in a cloud, we needed a framework to guide those decisions. A cloud decision framework is how we incorporated the cloud into the evolution of our daily operations. NetApp IT seeks to maintain flexibility in the cloud without vendor lock-in to meet the every-changing needs of the business and maintain a competitive advantage. The framework addresses cloud placement for new apps and the evolution of legacy apps to the cloud.

In part one of this blog series, we will discuss the development of our cloud decision framework to transition IT from a traditional operating model to one that is cloud-centered. In future blogs, we will explore the other phases of implementing a cloud-first strategy, including policies, governance, and execution. We start with our first step:

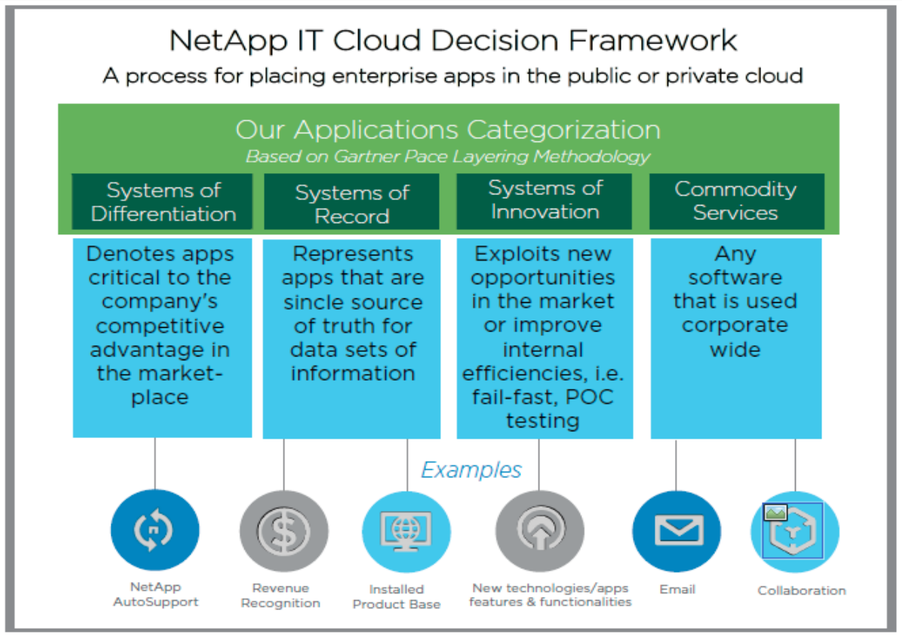

Recognize that you're already in the cloud and learn from it.

Many IT organizations see the cloud as new technology when the reality is that they are already leveraging cloud-based SaaS business services, such as payroll, HR benefits, lead generation, and the help desk, which have been common cloud services for decades. It is also likely your business is using the cloud without IT involvement. Examples can be found in marketing, finance, and sales.

The goal here is to learn why your organization chose to use SaaS, validate that criteria, and then apply it to the rest of your application environment. More than likely you've chosen SaaS because the apps/services were not core to your company's competencies. Core competency becomes the foundation of your cloud decision framework.

Develop your cloud decision framework.

You will need a cloud decision framework to standardize and bring consistency to your delivery and operational process for the cloud. We recognized that core competency was an output of the first assessment. Now let's table that for a moment to ensure we can apply that criteria in the most meaningful way possible to your entire application portfolio.

In our environment, we were already using the Gartner Pace Layering methodology for IT governance to prioritize technology investments. We recognized it could also help us define our cloud decision framework and prioritize the order of cloud adoption. This methodology offers a structured approach based on roles that the applications play in the business, based on three categories. We added a fourth category called Commodity that captures common-use desktop- or service-based applications throughout the corporation. These categories are:

- Systems of Record — These apps are the single source of truth for data sets of information. Examples include payroll, revenue recognition, installed base, etc.
- Systems of Differentiation — These apps are critical to the company's competitive advantage in the marketplace. One example is NetApp AutoSupport™ (ASUP™), our proprietary customer support system.
- Systems of Innovation — These apps are developed to exploit new opportunities in the market or improve internal efficiencies and require a fail-fast, proof-of-concept approach for testing new capabilities. The objective is to decide about investing in them as fast as possible.
- Commodity Systems — Any software that is used corporatewide. This includes email, help desk, collaboration, office tools, and web conferencing.

Place your apps into categories.

Now that you understand the classifications, the next step is to place your IT-supported apps into one of the four categories. This process should be simple and take only a couple of hours. This is a not a detailed exercise. The goal here is to identify the first categories of focus in the simplest way possible.

For NetApp IT, the first logical focus was on our commodity systems and systems of innovation. Commodity systems are not a core competency and should already be in the cloud. NetApp IT is taking a very aggressive approach to making this a reality.

As you move commodity and innovation systems to the cloud, you will recognize that systems of record and differentiation are more complex to analyze. This is because these systems are legacy systems for most companies. These systems are usually highly customized and highly integrated, with very specific proprietary features. They should be part of your private cloud strategy.

Many of the applications categorized as systems of record are not a core competency and should be evaluated against all your working business requirements to make a final determination for placement. As you go through this exercise, you will discover many of these systems will be candidates for SaaS-based solutions.

Now that you have these categories, you can start developing your policies, such as requiring that all new apps be cloud-ready, with no proprietary constraints or customizations. In our next blog, we will explore developing policies for both new and legacy applications, including the importance of maintaining control of our data regardless of which cloud an application resides in.

Supporting Business Agility

In this blog, we described how NetApp IT developed a cloud decision framework, including developing criteria based on our current cloud presence and applying a structured approach to categorizing apps based on their cloud fitness and business requirements. The framework is an effective tool for IT to standardize its approach for supporting core vs. non-core systems and utilizing the cloud as part of a dynamic hybrid cloud strategy. Just as importantly, the framework is a solid foundation for NetApp IT to evolve from a technology provider to a broker of services.

Matt Brown, Customer Engagement and NetApp on NetApp Program, NetApp IT

Why the Data Fabric is Critical to Managing Data in the Hybrid Cloud

KAMAL VYAS, SENIOR IT SERVICE ARCHITECT, NETAPP IT

Data management is critical to any successful hybrid cloud strategy, especially when using multiple clouds. In a true hybrid environment, enterprises should have the flexibility to move applications across different public and private environments based on their business requirements and cost considerations.

Business applications are relatively stateless and can be easily brought up and down in various environments. However, data must be managed differently because it has its own unique characteristics:

- ***Life*** – Data needs to be maintained, synchronized, audited, archived, etc., throughout its lifecycle.
- ***Value*** – Data is a corporate asset that must be protected.
- ***Mass*** – Data requires time to move in and out of environments.

Cloud companies (such as Amazon Web Services (AWS) and Microsoft Azure) are aware of the demands of data management. They entice enterprises with a variety of perks (including free data uploads) to gain control of their data. The more data that cloud providers can control, the more they lock in their customers for the long term and the more revenue they generate. Our primary goal is to use the Data Fabric —NetApp's vision for the future of data management—to control our data and avoid the vendor lock-in, among other things.

Data Management Considerations

Where data is hosted is critical to an enterprise data management strategy. It defines what type of choices and options an enterprise such as NetApp has as its cloud footprint grows. To provide IT with the maximum benefit of cloud services, a data management strategy should address the following five key areas:

- Secure control and governance of data regardless of its location, and guaranteed data privacy, as mandated by government policies and in-country laws such as Privacy Shield (previously Safe Harbor) laws.
- Access to data where and when applications need it to satisfy business use cases for disaster recovery, business continuity, and archiving.
- Flexibility to migrate data and applications between different cloud providers, locations, etc., to avoid vendor lock-in.
- Data compliance with company requirements (e.g. Sarbanes Oxley (SOX), HIPPA, etc.) and the ability to satisfy audit and other governance processes.

- Lower total cost of ownership, including storage costs, personnel costs, storage efficiencies, lower data transfer costs, etc.

NetApp IT Data Fabric Overview

NetApp IT's Customer-1 program is the first adopter of NetApp products and services into our IT production environment. Customer-1's goal is to provide feedback to Product Engineering on a product's performance so that a more stable product can be delivered to customers.

Customer-1 implemented the Data Fabric as the underlying architecture of our hybrid cloud strategy. The Data Fabric enables IT to manage data across multiple environments using standard tools, processes, and governance methodologies, independent of cloud providers or locations.

The Data Fabric delivers three major benefits to NetApp IT:

Data storage. NetApp® ONTAP® is the foundation of our Data Fabric, including management of all our public and private cloud data. Our private cloud leverages ONTAP-enabled FAS systems. As mentioned in my previous blog, we leverage NetApp Private Storage (NPS for Cloud) and ONTAP Cloud (a software-only version of ONTAP) for our public cloud workloads. NPS takes advantage of high-speed direct connections from a nearby co-location provider to leading cloud providers such as AWS and Azure. It also provides private storage options to augment elastic compute capabilities from these cloud providers. ONTAP Cloud delivers the same enterprise class data management as on-premises storage. This enables NetApp IT to retain full control of our enterprise data at all times, irrespective of the data location or cloud provider

Data replication. Within the Data Fabric, we use the SnapMirror® replication and SnapVault® backup feature of ONTAP to move data between end-points. This seamless transfer of bulk data provides the underlying transport for our data across various public and private clouds to deliver a variety of business use cases such as application migration, data replication, disaster recovery, etc.

Data archiving. We use NetApp's AltaVault® and StorageGRID® storage solutions for hybrid cloud data protection and archiving. AltaVault integrates with our NPS and FAS storage systems to back up data to a multi-site StorageGRID object data store. This combination provides us a truly scalable and tapeless backup solution to meet our data archiving and compliancy requirements.

Data Fabric in Action

The Data Fabric enables the Customer-1 program to use the cloud as a flexible component in its integrated IT environment. We can choose the cloud that offers the right service level at the right price for that business customer. This framework opens up many benefits in how we manage our production environment:

- We have complete control of data at all times, irrespective of the application location/cloud. The same on-premises data governance, security, privacy, and compliance methodologies are applied to cloud workloads as well, enforcing consistency across our IT environment and minimizing risk.
- We can map the right workloads to the right clouds. When requirements change, due to performance or cost, we can easily move workloads in and out of a cloud without worrying about data migrations. This helps us avoid vendor lock-in, cloud data transfer delays, and extra charges.

The Data Fabric drives NetApp IT's hybrid cloud data management strategy.

- We support a variety of cost-saving use cases. For example, the Data Fabric has enabled us to migrate our disaster recovery (DR) applications to the cloud and remove rarely used and costly compute from our on-premises data centers. We synchronize the data between our on-premises data center and our public cloud through NPS. Compute from public cloud providers is only used during a DR or testing event.
- The Data Fabric allows us to use consistent storage standards and policies across various cloud stacks. By providing a homogeneous storage layer we can easily expand to meet data management capabilities across various technologies and cloud stacks, including AWS, Azure, and OpenStack.
- Finally, the NetApp Data Fabric/NPS provide a rich set of enterprise features and capabilities that are not available in public storage. This allows us to standardize data management across all the platforms and eliminate the need for application re-design, new storage skills, new process development, etc., as technology changes.

The landscape of enterprise IT is changing rapidly with the rise of the cloud. Data management is a critical factor to consider in this journey. As Customer-1 for NetApp's hybrid cloud strategy, we are using the Data Fabric to gain greater visibility and control over our enterprise data, regardless of where it physically sits. The Data Fabric enables us to combine on-premises capabilities with cloud provider resources to take advantage of a whole new level of compute power and automation for our business customers. More importantly, it supports the evolution of our data management strategy to meet the demands of the future.

This is the third blog in a series on NetApp IT and the hybrid cloud. To read previous blogs, click on these links:

Blog 1: How to Define Your Hybrid Cloud Strategy with These Six Goals

Blog 2: Improving the User Experience When Adopting the Hybrid Cloud

Kamal Vyas, Senior IT Service Architect, NetApp IT

7 Reasons Why Moving Apps from the Public to the Private Cloud Shouldn't be Hard

MIKE FRYCZ, CLOUD OPERATIONS AND SERVICES LEAD, NETAPP IT

Start by automating the provisioning and decommissioning of your cloud servers.

As the cloud becomes an integral part of our IT operations, I have read a lot of articles about leveraging the public cloud as part of a data center strategy. What's just as interesting, but less discussed, is the move from the public cloud into a private cloud in a data center.

In our case, moving apps between clouds is an integral part of our strategy to retain control of our data and make it available where it makes the most sense, both from a cost/efficiency and usage perspective. It's a component in our broader Data Fabric journey that gives NetApp IT the flexibility it needs to excel at service delivery.

7 Lessons Learned in Cloud Migrations

As a cloud business analyst, I am responsible for moving applications from the public cloud, managed by a service provider, to the NetApp private cloud using our NetApp Private Storage for Cloud (NPS) solution. In the past 18 months, I have moved or decommissioned more than 100 applications, both small and large. We are migrating the apps from service-provider leased space to our private data center to both reduce overhead costs and utilize available data center space.

Here are a few things I learned during this process that are worth sharing with anyone facing any type of cloud migration:

1. ***Automated provisioning.*** Cloud adoption is really a story of automation. Most cloud benefits lie in the ability to eliminate manual tasks. We automated three major tasks so that provisioning went from days/weeks to minutes and required much fewer resources.
   - Added a self-service portal that was integrated to our service management software;
   - Automated decommissioning of cloud servers; and
   - Created a DMZ for basic approved security firewall rules.
2. ***Start small.*** Any IT project manager knows that every application has its own unique migration path; some apps are more complex to migrate than others. We started with small apps to test the process and determine what automation was needed. We learned a few lessons along the way: develop a template to follow (and don't deviate from it), clearly outline requirements prior to deploying the infrastructure, and open communications with end users as early as possible. Our first few migrations taught us what we needed to know to avoid future issues.

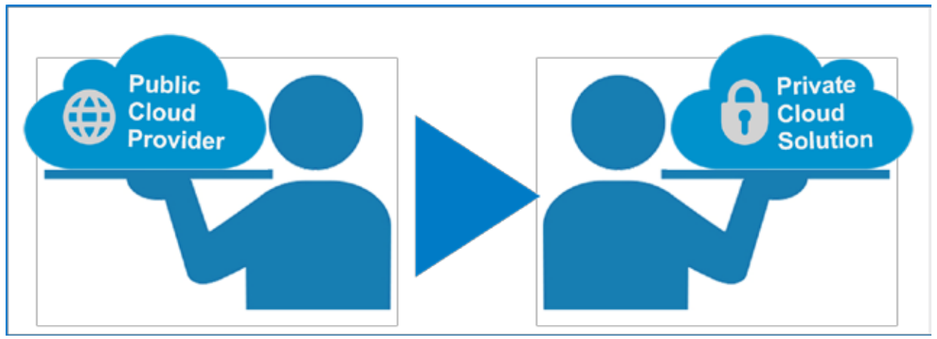

3. ***Avoid customization.*** You will be tempted to adopt manual exceptions in the interest of time. In our case, it is the highly customized apps that cause issues with firewall security. In addition, one-time customizations cause long-term problems. In this case, our technical support team didn't have a process that support these exceptions. Non-automated (i.e. manual) rules require more time and resources to mitigate problems.

4. ***Internal support.*** Moving applications into a private cloud also means moving from cloud vendor support to internal support. This has both benefits and drawbacks. IT gains a more knowledgeable support staff and more control over its operations. Issues can be resolved more quickly. On the minus side, internal tools are not as mature as cloud providers. For example, users can't scale or resize their compute/storage as quickly. However, these shortcomings should not outweigh the long-term benefits.

5. ***Hidden cost savings.*** Moving to a private cloud reduces many expenses. Excess capacity is being utilized and public cloud vendor costs are eliminated. IT can piggyback on the existing IT infrastructure—tools and staff in the data center and command center/support. Expenses may start high, but the more apps that are moved, the more expenses will decline. In addition, any investments can be managed more closely.

6. ***Get the right backers.*** Why change things if they are working? Human behavior resists change especially if the current situation isn't broken. That's why it's important for executive leadership to endorse a cloud migration so the business understands that long-term business benefits outweigh the short-term hassles.

7. ***Automate decommissioning.*** Via our self-service portal, we can automatically decommission the compute, storage, and monitoring in one push of a button. This results in time savings, accurate record-keeping, and the elimination of manual intervention.

The Data Fabric in Action

My project to move apps from the public to private cloud migration is part of a journey that is not possible without a Data Fabric that ensures we keep control of our data. We leverage NetApp Private Storage for Cloud to retain data control, avoid vendor lock-in, and dodge costly data migrations. NPS enables us to connect to and switch cloud providers at any time, whether it's to move from a public cloud to a private cloud or vice versa. NPS also provides the high-throughput, direct connections to the cloud, the ability to rapidly scale our computer or run-time resources for peak workloads, and a common set of capabilities for managing data.

Operationalizing the cloud is a major challenge for any enterprise, including NetApp. But the secret to success lies in patience, flexibility, and processes that enable IT to use the cloud on its terms.

Mike Frycz, Cloud Operations and Services Lead, NetApp IT

The Business Case for Flash Is Compelling When Looking at Data Center Efficiency

RANDY EGGER, DATA CENTER LEAD, NETAPP IT, AND EDUARDO RIVERA, SENIOR STORAGE ARCHITECT, NETAPP IT

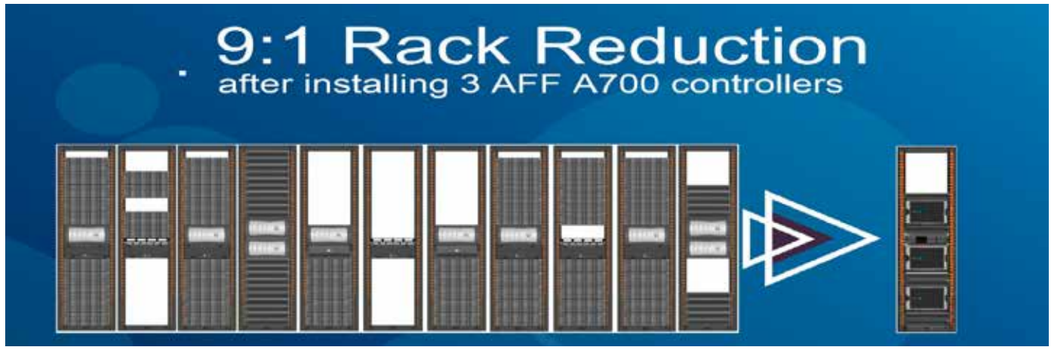

Technology is evolving faster than anyone anticipated and having a significant impact on our ability to reduce our storage footprint as well as operational costs. We are seeing this evolution play out in our storage lifecycle management strategy, which is undergoing a dramatic upgrade thanks to the benefits of All Flash FAS (AFF).

As part of our lifecycle management strategy, we are decommissioning some older FAS nodes running on ONTAP 9 and replacing them with AFF. We are excited to migrate to AFF for a variety for reasons, including improved application performance, higher density, and economies of scale that will help us transition to the future.

The biggest change we see will be in the provisioning strategy. In the past, we had to over-provision to meet performance requirements. AFF takes that issue off the table. One AFF solid state drive (SSD) equals the I/O performance of multiple shelves of hard disk drive (HDD) or spindle technology. The sizing of the array is much simpler because we can provision based on capacity, not on future performance requirements. The AF700's powerful controllers and incredibly dense drives will dramatically improve the application performance while minimizing the footprint.

Another benefit is the non-disruptive migration made possibly by ONTAP. We anticipate no interruption of service for our end users, something we could not have achieved with older 7-mode technology. To non-disruptively move the applications requires three steps. First, we will add the new A700 nodes to the existing cluster, then we will use ONTAP's Volume-Move feature to non-disruptively migrate all the data in from the older nodes to the new AFF nodes. Finally, we will remove the older nodes from the cluster. We will repeat this process to refresh all our existing storage infrastructure without an outage.

A smaller storage footprint reduces data center costs while significantly improving efficiency.

- Power requirements of 36.5kW are being reduced to 6.5Kw, an 82% drop. Based on a $0.08 cent/kWh electricity rate and PUE of 1.5, costs will drop from almost $32,000 to just $6,800. Our savings would be double or triple if we were in a traditional leased or outsourced facility. We also eliminate the additional costs and latency issues that come with adding a second cage to house the legacy and/or new hardware during the migration.

- Space requirements are dropping 91%, as we move from 380 rack units (RU) of hardware to 36RU to house the three A700s.
- On the greenhouse gas side, our CO2 savings are projected to be 185 metric tons, equal to taking 36 cars off the road per year.

Next Steps

We will be upgrading to AFF in phases. Initially, we will be replacing approximately nine racks of FAS equipment with one rack containing three A700 15TB SSDs. Along with the migration, we will introduce 40Gb fiber into our clusters, which should dramatically improve data throughput. The combination of new, faster solid-state drives, faster storage controllers, and faster connections will have an enormous impact on capacity, performance, and efficiency---and the service we delivery to our business customers.

We expect to see an 82% drop in our power costs and a 9:1 reduction in rack space.

The adoption of Flash brings us one step closer to the vision of a data center that fits into a living room, not a football field. AFF plays a major role in achieving this goal by improving performance, simplifying operations, increasing data center efficiency, and streamlining the integration of any future infrastructure changes. A small footprint offers exciting possibilities to reduce IT overhead costs, especially space and power, without sacrificing performance.

Randy Egger, Data Center Lead, NetApp IT, and
Eduardo Rivera, Senior Storage Architect, NetApp IT

Adding OnCommand Unified Manager to Gain Single Pane Event Monitoring in IT

ED WANG, SENIOR MANAGER, AUTOMATION AND MONITORING TOOLS, NETAPP IT, AND TIM BURR, SR. MANAGER, INFRASTRUCTURE OPERATIONS, NETAPP IT

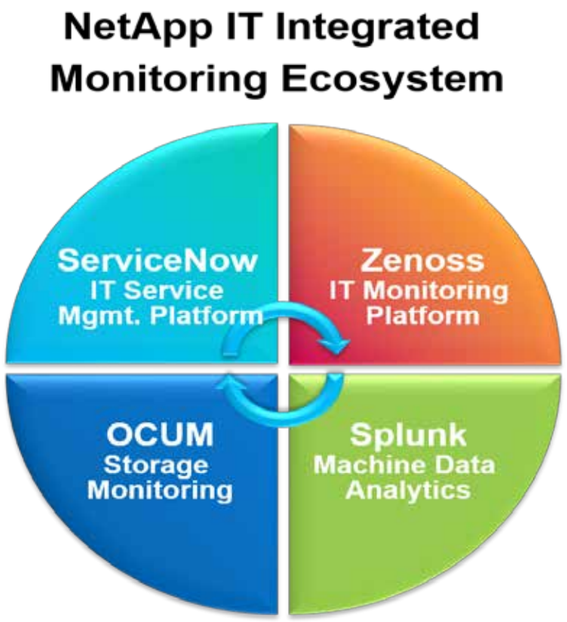

Necessity is the mother of all invention. At least when you work in IT and support global resources that include five data centers, 5,300 servers, and 52 PBs of data center storage. As a result, NetApp's IT environment generates a constant flow of alerts. Our eternal and ongoing challenge is to quickly identify the root cause of the issues and prevent them from happening again.

Our event monitoring strategy plays an important role in addressing this challenge. We want to ensure critical alerts quickly rise to the top for immediate attention, while informational alerts can be analyzed separately for later action. To support this strategy, we needed to consolidate our alerts into a single ecosystem made up of individual, best-in-class components. This would feed alerts into our incident management software for auto-ticketing.

This 'single pane of glass' strategy enables the NetApp resources on the infrastructure support team, called the Command Center, to quickly resolve critical issues 24x7 across the globe and not be sidetracked by non-urgent alerts. This approach improves IT's responsiveness and focus, ultimately resulting in increased operational stability.

Our first step was developing an alerting process. Like most IT shops, we have a two-tier alerting system, but we classified our alerts in a slightly different way:

- Reactive: This alert is the only type of alert to automatically be forwarded to the Command Center for immediate action. It is defined as "actionable" and requires attention by the team.
- Proactive: These alerts are typically performance related, but less urgent and are not immediately forwarded to the Command Center for action. Dashboards are used to manage thresholds for the alerts at a broader level. The Command Center monitors the dashboards to proactively address issues, such as storage capacity or CPU utilization, with partner application support teams. These types of alerts remain a key volume driver for the Command Center, but teams continue to focus on streamlining and automating these responses over time.

Over the course of about nine months, process and support teams focused on understanding what existing alerts, thresholds, and events were most important and "actionable." The result of this work was to position NetApp IT to implement a single, integrated service management and alerting ecosystem, with significantly less noise for those accountable for responding to the alerts.

Building an Ecosystem

Our plan was to create an event monitoring ecosystem that fed alerts into central incident management software. A single ecosystem would

enable the sorting, tracking, and accurate routing of alerts from our IT systems into our incident management software through auto-ticketing. For storage events specifically, this required we integrate multiple tools--Zenoss, Splunk, and NetApp OnCommand® Unified Manager (OCUM)--into our ServiceNow incident management platform.

A major hurdle occurred once we began integrating our storage environment into the ecosystem. The storage alerts were a necessary part of the end-to-end alerting strategy. And yet, configuring each individual storage controller to connect with Zenoss and ServiceNow presented an administrative and management challenge. OCUM, however, offered the ability to connect a single storage management tool into our existing monitoring ecosystem. OCUM gave us the advantage of managing thresholds and administering alerts in a single tool, and Zenoss provided the ability to analyze and dedupe the critical alerts before they were auto-ticketed.

By creating a strategy that enables fast action on high-priority issues, we've improved the efficiency ad effectiveness of our Command Center.

We created a Zenoss ZenPack (to be published soon in the Zenoss community), a plug-in module that outlines the business rules for OCUM to pass its monitoring events to Zenoss. Zenoss screens the alerts, dedupes them, then identifies the critical alerts for auto-ticketing. This integration brings storage alerts into the ecosystem, along with similar alerting configurations for server virtualization, network, and security components in the data center. It also enables NetApp IT to achieve another critical step toward consolidated event and incident management.

Improving IT Operational Stability

The new alerting strategy offers many benefits. The Command Center has greatly reduced its dependency on email for event notifications. Team members don't need to sort through alerts to find the critical ones, dedupe alerts about the same issue from multiple sources, or risk assuming that someone has already addressed the issue. The team only receives alerts that specifically require action.

When a device goes offline, a storage volume becomes unavailable, or a storage system experiences a hardware failure, the team is positioned to respond appropriately. Therefore, urgent infrastructure issues are identified and fixed more rapidly, before they cause havoc in our IT environment, reducing the overall number and impact of P1 incidents.

Regardless of the incident management or event monitoring software being used, any IT organization can benefit from rationalizing the number of actionable alerts and adopting an integrated event monitoring ecosystem. By creating a strategy that enables fast action on high-priority issues, we've improved the efficiency and effectiveness of our Command Center. Ultimately, this approach has a direct impact on the operational stability of IT operations for our customers, partners, and employees.

Tim Burr, Sr. Manager, Infrastructure Operations, NetApp IT

Ed Wang, Sr. Manager, Automation and Monitoring Tools, NetApp IT

Why Storage Plays Such a Critical Role in the Success of the NetApp Multimedia Team

ED CHO, PRINCIPAL MULTIMEDIA ARCHITECT, NETAPP IT, AND BART SAUNDERS, MANAGER ONSITE MULTIMEDIA SERVICES, NETAPP IT

Last year was a busy year for NetApp IT AV & Multimedia Services. Within this group, the Multimedia team is responsible for architecting the video production infrastructure as well as capturing, editing, and delivering video content for the entire enterprise. Our charter covers a wide array of activities:

- Produce an average of 370 live events and 400 studio recordings per year on a wide range of topics, from executive presentations to product launches, corporate alliances, celebrations, community support, and corporate all-hands meetings.
- Manage and edit live video streaming and on-demand recordings from a variety of campus venues, including a 1,200-seat auditorium, a 400-seat theater, a green-screen sound stage, and training rooms, all connected by a fiber-optic network.
- Consume approximately 30TB of storage per year, which is steadily climbing with the advent of 4K video. An average one-hour studio recording uses four recording feeds—three cameras and one switched program cut which translates to approximately 60 GB of video content.
- Support six editing workstations.
- Provide on-demand viewing of more than 800 videos per year. Most of this content is then edited into additional segments.

The E-Series improved our team workflow and process efficiency.

With just five people involved in video production, we need to be smart and highly efficient in the way we work. Reliable data acquisition and storage is critical to our daily workflow. Even the slightest performance issue can severely impact our ability to capture, edit, and deliver the content of our clients. Any interruption to this high-volume data workflow and our daily operations grind to a halt.

Storage Capabilities

Professional video production and delivery require very different storage capabilities. One emphasizes high performance; the other reliability and scalability. Below, we describe how we use two NetApp storage solutions in our multimedia workflow.

Production: E-Series

Our production activities require high-capacity, high-performance storage that can accommodate bandwidth-intensive content and is simple to deploy and manage. We require real-time, multi-stream video access from workstations across multiple campus locations. These workstations must be able to simultaneously access the same content in the post-production environment without retrieval lags, which slow down editing. To cope

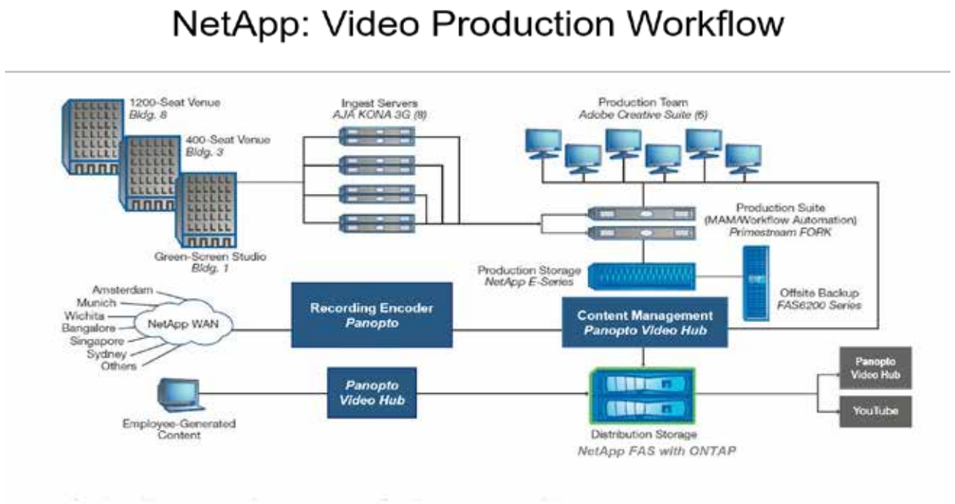

with such demanding and potentially adverse conditions, we utilize the NetApp E-Series.

Installing the E-Series was a huge step forward for our team. We previously used four Sony XDCAM-formatted optical disks to facilitate a one-hour recording. Once the recording was finished, we had to manually transfer the content from each disk into an edit station, which took three to four hours. When the transfer was complete, the content could only be edited at that station. All the content had to be ingested before any production work could start. If the workstation suffered a failure, all content was lost. The process was laborious, unreliable, and not secure.

We updated our workflow with the NetApp E-Series hybrid storage system with 180TB of capacity and roughly 2 GBs of video throughput. The E-Series supports high-performance mixed workloads, including bandwidth-intensive video streaming, with ease. It also enables us to capture up to eight broadcast-quality live streams of content and place it directly into E-Series storage. Because of the E-Series' scalability, we are able to easily expand our ingest capabilities without sacrificing quality.

The content is now available to all six edit bays simultaneously within seconds and with no storage delays. Our ingest process is fully automated and integrated, thanks to the Quantum StorNext file system and Primestream FORK Production Suite media asset manager. The E-Series is provisioned using the NetApp SANtricity® operating system, which safeguards the ingest and delivery of data.

In addition to improving our team's workflow and process efficiency, the E-Series delivers other benefits, including:

- High performance. The E-Series handles the live ingest of videos in a broadcast-quality video format, without any dropped frames, directly into storage. This throughput is critical to the editing process, especially during live feeds.
- Shared storage workflow. We can edit live-streamed content roughly 30 seconds after a recording starts.
- Collaborative post-production. We can work simultaneously at different edit stations, dramatically increasing the pace at which the final output can be delivered. We can edit in RAW 4K format without any editing lags.
- Increased productivity. After adopting E-Series, we eliminated hours of manual tasks in the production workflow. This gave us more time to focus on other value-added activities.
- Enterprise-grade resiliency. When a hard drive failed under the old system, the content was lost. Now, if we incur a hard drive failure, we can swap out the failed drive and the content remains available via the E-Series.
- Rock solid operation. While other parts of our video ecosystem have failed, we have not had any issues with the E-series since it was installed four years ago.

FAS enables us to mange our com-plex workflow and maximize storage efficiency.

- Easy storage expansion. Storage capacity can be added without requiring any type of special configurations or major design changes.

Delivery: FAS & ONTAP

To address video delivery and the video-on-demand (VOD) needs of end-users--including employees, partners, customers, and investors—we have a very different set of storage requirements. In this case, storage must be scalable, flexible, cost-efficient, and reliable. For this, we rely on NetApp FAS hybrid storage and ONTAP® storage management software to quickly and easily distribute content to a growing audience using a wide array of devices, including laptops, cell phones, tablets, and video monitors. Maintaining an agile encoding solution is critical to meeting the demands of these diverse distribution channels.

In addition to our own video content, we support the internal groups that want to produce their own video content using the Panopto enterprise video content management solution. Users can record, upload, encode, and post their own videos. To conserve storage space, we automatically compress files so they can be easily streamed to web servers then delivered to desktop computers and mobile devices.

We use the NetApp FAS6000 hybrid storage series to store our finished media assets and archive our source files. NetApp FAS enables us to manage our complex workflow, maximize storage efficiency, and automatically adapt to changing workloads. In addition, we leverage NetApp FlexPod®, a converged infrastructure that combines storage, networking, and compute into one architecture. FlexPod serves as a DMZ area for secure access to assets outside our VPN. The combination of FAS, ONTAP, and FlexPod delivers numerous benefits, including:

- ***Non-disruptive management.*** ONTAP enables non-disruptive operations and eliminates planned downtime; we plan to expand this capability to the hybrid cloud in the future.
- ***Simple backups.*** We use ONTAP to automate backup and recovery with Snapshot™ software with no performance impact.
- ***Flexible storage.*** Content can be made available securely to both internal and external audiences.
- ***Lower costs.*** Power, cooling, and floor space as well as operational costs are much lower than that of traditional, large-scale storage.

The Future Is Shared Storage

As a company-wide service organization, the Multimedia team has one goal: excellent service for its customers. We rely on a robust and reliable NetApp storage infrastructure to operate our video production and delivery workflows at top efficiency, otherwise we could not keep up with the programming demands of our global enterprise.

With the adoption of NetApp products into our workflow, we can devote more energy and resources to creating high-quality content. A diversified foundation of NetApp storage solutions is one way for us to continually adapt to the growing needs of our many stakeholders while managing and protecting NetApp's critical video assets.

Ed Cho, Principal Multimedia Architect, NetApp IT, and Bart Saunders, Manager Onsite Multimedia Services, NetApp IT

www.NetAppIT.com **Twitter: @NetAppIT**

Automating Capacity Management: How to Get Rid of Those Pesky Storage Capacity Alerts

EZRA TINGLER, SR. STORAGE ENGINEER, NETAPP IT

NetApp IT's global enterprise environment has a large storage footprint. One of our biggest challenges is in capacity management. As a senior storage engineer in our Customer-1 organization, I am responsible for the storage capacity and performance management of our ONTAP® systems.

Managing storage capacity is a daily grind. It involves analyzing utilization trends, balancing existing storage pool utilization, and forecasting future storage needs of applications. While all those tasks are certainly important and interesting aspects of my job, I also spend a lot of time addressing everyday issues that hamper my ability to effectively manage our ever-growing environment.

This new process has significantly improved storage team productivity.

I consistently see alerts for full aggregates, which NetApp IT has defined as any aggregate which exceeds the 70% capacity threshold. This is a low threshold by most IT organization standards. Even so, whenever an aggregate exceeds this threshold a support ticket is opened. These events cause a lot of extra work for our operations team as they work to resolve the incident.

In addition to the workload generated by aggregates running out of space, full aggregates are a real problem. Full aggregates can result in performance and data accessibility issues. All volumes within an aggregate (FlexVols) are thin-provisioned, which allows us to provision the full amount of storage at time of creation, but allocate it when needed. In many cases, the aggregate full conditions are the result of volumes that do not conform to our standards. These are easily resolvable by simply applying our standard thin provisioning settings.

If the aggregate full condition is not resolved by a simple adjustment of the FlexVol's properties, then we need to look for a way to lower the consumed space within the aggregate. In these cases, we look for a good destination aggregate for the volume and plan a volume move. We manually perform a series of checks prior to and after the move. This is a time-consuming and repetitive task, with lots of back and forth between the engineering and operations teams.

The Solution

My goal was to automate capacity management to reduce the number of capacity-based alerts that require manual intervention. To do this, I needed to automate the process of finding potential capacity issues and resolve them by moving volumes to aggregates with more available free space. I'm not a software developer, but once again I found myself writing PERL scripts for a few basic steps:

- Pull information from the OnCommand® Unified Manager (OCUM) reporting database via a wrapper script;

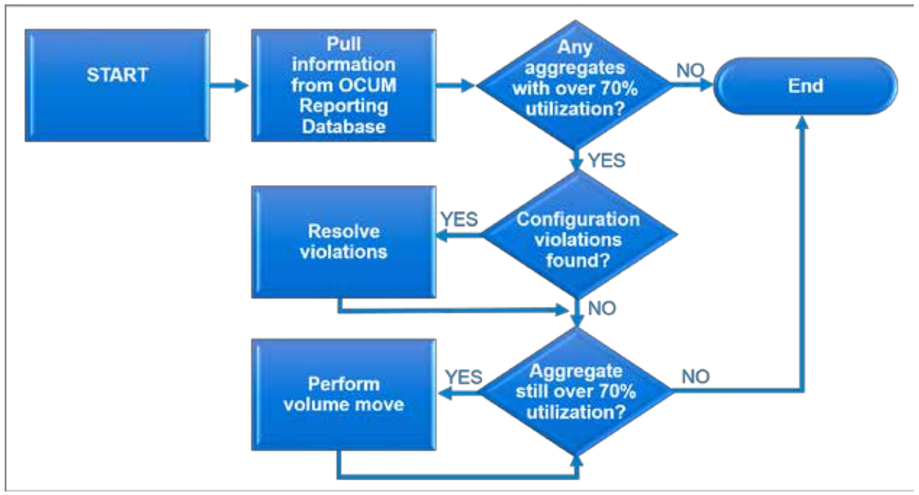

- Check the settings for the volumes and logical unit number storage (LUNs) on all clusters and resolve violations via a separate configuration enforcement script;
- Check the aggregate capacity, perform volume move pre-checks, execute volume move, and perform the post-checks via a capacity redistribution script.

I used a similar methodology to write a script for automating cluster ONTAP configuration, using the NetApp Manageability Software Development Kit (NM SDK). The scripts check for violations by performing several steps to identify and resolve issues and evaluate possible targets. The redistribution script moves the affected volume to another like aggregate of the same service level. NetApp IT uses three levels for delivering storage services. Using the OCUM database enabled me to minimize the number of connections to the clusters.

The process is kicked off by the wrapper script which runs as an hourly cronjob on our administrative hosts. It pulls a list of all clusters in our environment from the OCUM reporting database and calls the configuration enforcement script to check and resolve cluster configuration issues. The wrapper script also calls the capacity redistribution script to evaluate the target cluster and determine a source volume and destination aggregate. The wrapper script is also used for other storage automation tasks.

Specific users and roles as well as SSL certificate authentication provide secure access. The capacity redistribution script writes to the standard syslog facility on the administrative host and generates a service incident for tracking and reporting purposes.

Results

Following the implementation of this automation, the NetApp IT storage environment has greatly improved. We found that just under 30% of the volumes in our development environment and about 20% in production had configuration issues or conflicted with our current volume setting standard, a common issue in large storage environments. The scripts automatically resolved these standard conflicts.

This new process also significantly improves storage team productivity. Each time an aggregate exceeded 70% capacity it generated a storage service incident. We had an average of 15 aggregate capacity incidents per month. These incidents took an average of 6.25 days to resolve. These scripts enable us to resolve these incidents in minutes or hours, not days.

Another benefit is that the scripts automatically implemented thin-provisioning on all volumes, which instantly resulted in space savings of more than 5PB. We have been able to keep our capacity growth almost flat and eliminate the immediate need to invest in additional storage capacity.

The Future

Now that version two of the scripts are deployed, I am working on some other automation tasks. The next step will be importing the scripts into OnCommand® Workflow Automation (WFA) as an automation script. I should have more time to work on this project now that I don't need to worry about as many capacity issues.

Ezra Tingler, Sr. Storage Engineer, NetApp IT

www.NetAppIT.com **Twitter: @NetAppIT**

NetApp on NetApp: Sharing our IT Experiences

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 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 our web site: www.NetAppIT.com.

Read our regular blogs with insights into our everyday operations and strategies.

Read our other ebooks:

- Building a Foundation for Business Apps Agility
- 7 Tips & Tricks to Make Storage Administration Faster & Easier
- 7 Perspectives on The Future of IT: The Drive Toward Business Agility

www.NetAppIT.com **Twitter: @NetAppIT**

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.