

NASCIO 2018 Enterprise IT Management Initiatives

1. **Title: Utah's DTS Storage, Data Backup/Restore Infrastructure Initiative**

2. **Category: Enterprise IT Management Initiatives**

3. **Project Initiation and completion date:**

Apr 2016 - Nov 2017

Project Website: <https://dts.utah.gov/product/backup-archive-and-restore-services>

4. **Organization and primary point of contact**

Organization(s): Utah Department of Technology Services (DTS)

Other contacts: Darrus McBride, darrus@utah.gov



The screenshot shows the top navigation bar of the DTS website. On the left, there is the Utah.gov logo and the text 'Utah.gov'. On the right, it says 'Department of Technology Services'. Below this is a dark blue navigation menu with the DTS logo (a hexagon with 'DTS' inside) and links for 'Get Help', 'Products', 'My Links', 'News', 'Security', 'MDM', and 'Policies'. The main content area has a large orange heading 'Backup, Archive and Restore Services'. Below the heading, it states: 'DTS' data backup and restore environment is available for State agencies and government entities for the following applications:'. This is followed by a bulleted list: '• Servers hosted in the Salt Lake City or Richfield data centers.' and '• Servers with State of Utah WAN connectivity hosted at agency sites outside the Salt Lake City and Richfield data centers.'. Below the list, there is a paragraph: 'Server backups are sent to the backup appliances and our in-house cloud (object) storage located in the Salt Lake City and Richfield data centers. Agent software is needed on physical servers but not on virtual servers and backup and restores are managed through a web console. An initial full backup or snapshot is performed and then followed by daily incremental snapshot backups determined by the SLA policy. Archives will be sent to object storage for long term retention. System Administrators may restore data through the management console as needed. DTS personnel are available 24x7x365 to provide assistance.'

Executive Summary

When the Utah State Legislature formed the Department of Technology Services (DTS) in 2005, the Legislature ruled in code that one of DTS's responsibilities would be to:

"Lead state executive branch agency efforts to establish and reengineer the state's information technology architecture with the goal of coordinating central and individual agency information technology in a manner that:

- (a) ensures compliance with the executive branch agency strategic plan; and
- (b) ensures that cost-effective, efficient information and communication systems and resources are being used by agencies to:
 - (i) reduce data, hardware, and software redundancy;
 - (ii) improve system interoperability and data accessibility between agencies; and
 - (iii) meet the agency's and user's business and service needs;"

With this legislative mandate, DTS has made it a top priority to ensure that cost-effective, efficient information and communications systems and resources are being used by agencies. Part of this core responsibility, DTS continually reviews and evaluates provided systems and services. One of those foundational infrastructure services provided that requires constant review is storage. Therefore, in February of 2015, the Department of Technology Services (DTS) began to review the State's storage consumption and notice a dramatic increase in the use of utilized server storage and found it necessary to conduct an internal study that would identify the trends of data storage over the previous 5+ years (12/09 - 02/15). The study revealed that the storage pool utilization growth was 4,783 GB per week. The rate growth served as a warning that DTS's current technology infrastructure would not be able to adequately meet the growing demand for reliable and always available access to data.

DTS found the need to begin addressing the growing storage concern, fast. Despite the speed at which DTS addressed the impending change, it was imperative that the efforts were precisely orchestrated. With an investment to be made, it was crucial that the acquired solution remained open enough to meet the expanding need, not only for the present, but for the future as well, at a cost that the State of Utah could sustain.

Concept


The Department of Technology Services has the responsibility to continually review and evaluate provided systems and services. A major part of that evaluation is to ensure that cost-effective, efficient information and communication systems and resources are being used by agencies to reduce data, hardware, and software redundancy.

Being that storage is such a crucial and foundational part of the State's infrastructure that spreads across all state agencies and continues to grow daily, it is an obvious place to continually monitor for available areas of consolidation, cost savings, capacity and performance enhancement.

The State of Utah has been on a journey for multiple years, working toward achieving a Software Defined Data Center (SDDC). With such, in 2015, Department of Technology Services (DTS) found it necessary to conduct an internal study that would identify the trends of data storage over the previous 5+ years (12/09 - 02/15) in an effort to see what . It was determined that a review of the State's storage consumption was necessary and it was noticed that there was a dramatic increase in the utilization of server storage and it was revealed that the storage pool utilization growth was 4,783 GB per week.

The study identified areas of growth and focus, which lead DTS to propose a plan for:

- The upgrading and modernization of the State's two data centers
- Decrease the Data Center footprint through the use of newer software defined/hardware minimalistic components
- Implement new solutions that are more energy efficient
- Implement future upgrades in a more modular fashion with little or no downtime.



This screenshot depicts the rate at which storage needs are growing

Storage Component:

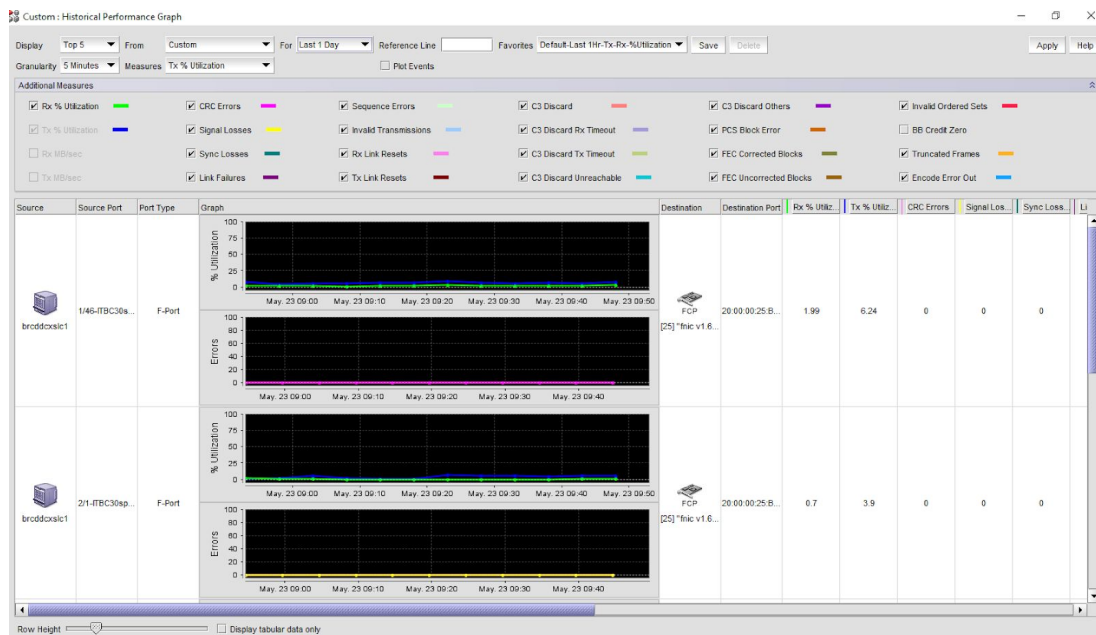
The first building block that needed to be addressed was storage modernization. DTS's existing spinning disk storage system provided by a large datacenter storage company, was no longer capable of providing DTS's projected future storage demands within reasonable budgetary projections related to both infrastructure upgrade costs and customer agency service chargeback costs.

By the end of March, 2016 DTS entered into a proof-of-concept meetings with a narrowed down list of six storage solution vendor's encompassing a variety of technology solutions (spinning, flash, hybrid, adaptive flash hybrid disk). After deeper vetting and analysis of various "proof of concept" trial results, it became very clear that the newer and more cost effective adaptive flash hybrid array storage technologies were the best fit for the future needs of DTS and the State agencies it serviced. It was determined by our technical architecture and engineering teams that an adaptive flash hybrid array would be the most cost effective, performance enhanced and labor efficient solution DTS could obtain as part of the core infrastructure and meet the needs of the customer.

On April 5, 2016 the installation of the new adaptive flash hybrid arrays storage began and by mid-May, 2016 DTS started to realize the benefits of selecting the adaptive flash hybrid arrays and begin migrating off of its existing spindle disk storage and onto the new storage solution.

The new adaptive flash hybrid storage solution offers DTS a number of benefits:

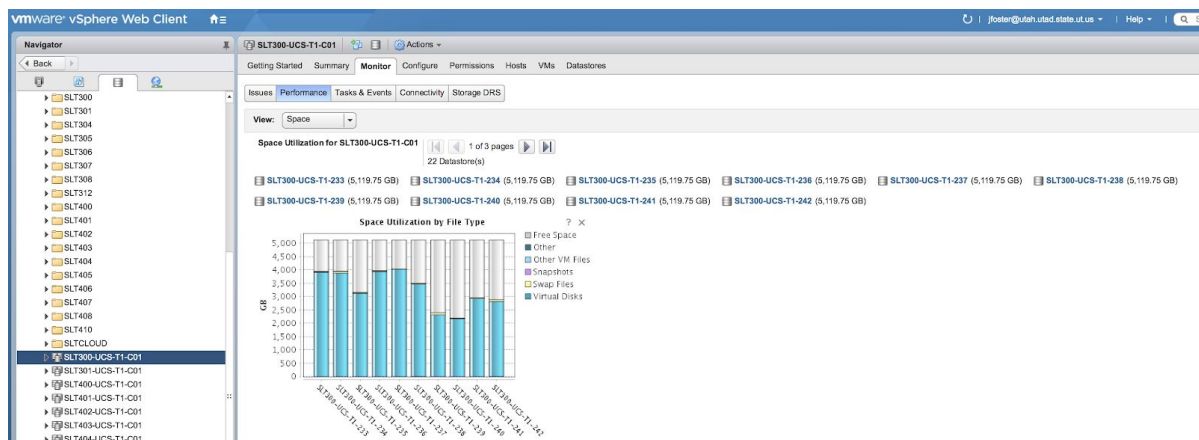
1. The chosen vendor utilizes predictive analytics that helps the system anticipate hardware failures, demand spikes and IO patterns.
2. The adaptive flash hybrid solution highly abstracts the storage components, removing the tedious burden of traditional storage system administration.
3. Vendor support is a very fast and responsive, the majority of issues are resolved through email or webforms.



Minimal utilization during production hours, which demonstrates optimal sizing for peek capacity

The newly implemented storage solution provides the State with tremendous flexibility, as it relates to increasing storage capacity and demanding workloads. To handle the large throughput requirements State workloads have, DTS combines the Brocade Fibre Channel switches and Cisco UCS Fibre Interconnects to create a storage network. One of the keys to having a maintainable storage network is to make everything congruent, modular and adaptable, therefore, if there is a need to expand capacity, simply add more servers to UCS and deploy the storage templates already created. If additional compute capacity is needed, compute resources can be added to the storage system in 60 minutes after being racked and placed into production, without interruption or downtime.

It is imperative to the state and those that manage it to find and create systems that manage themselves. Therefore, Software Defined Storage became a solution that needed to be implemented. The team leveraged VMware's Datastore cluster technology to minimize management and monitor tasks. Datastore clusters aggregate luns into a larger pool of storage. With Datastore enabled, as soon as a VM has been moved to a storage cluster, our Software Defined Storage solutions closely monitor usage patterns and behaviors, to act intelligently. If a lun starts to get completely used, the datastore cluster technology automatically moves VMs around to optimize space. Vsphere also closely watches IO patterns that the VMs create. If two VMs produce a large amount of storage traffic, then the system will separate them to different luns and hypervisors to minimize the impact to other VMs.



VMware vSphere capacity view of a production datastore cluster

The Department of Technology Services makes a concerted effort to virtualize every workload the State has. However, due to political and or technical considerations, this is not always possible. For physical workloads, there needed to be a simple way to convert them from spindle to adaptive flash hybrid arrays without interruption. The team used Sansymphony by Datacore to minimize downtime that can be required when replacing large storage systems. Sansymphony adds a storage abstraction layer to physical server. This enables the storage and sysadmin teams the ability to clone data over to the new system while maintaining application availability and integrity. Upon successful data transfer, Sansyphony's mappings and abstraction are removed and the physical servers are repurposed and simply mapped to their new storage partitions on the adaptive flash hybrid system.

Backup and Archives Component:

As part of the initial study mentioned earlier, the DTS identified that there was also room to improve service delivery, as it relates to backups and archival of data. DTS had been using a backup and archive system created by a Large well established software/hardware company, in one iteration or another, for backup and archive since 2000. It became evident to DTS that the current solution's pricing model and lack of compatibility with newer infrastructure technologies (both implemented and planned) was becoming an impediment to the successful introduction of newer innovated technologies, in turn, making it difficult to fit with DTS architecture. As a result of this challenge, DTS knew at it was important to find a replacement Backup, Archive and DR technology solution that would fit into the new DTS technology paradigm and be adaptive enough to change and grow as needs changed. It was decided that a new solution was needed to solve the State's backup and archive needs.

In September, 2015 DTS began focusing more seriously into a backup, archives, DR solution from a relatively new company (January, 2014) Rubrik. A POC with Rubrik began on December 9, 2015. After successful POC results it was decided to purchase Rubrik. The first Rubrik devices (Briks) arrived on-site and were installed in late March, 2016.

In April, 2016 DTS began the migration of 1.92 Petabytes of data off of its aging backup and recovery solution, to its new backup & recovery solution replacement "Rubrik." The migration to Rubrik permitted DTS to provide its customers with a newer, more simplified, robust, backup and recovery offering that best meets the needs of DTS's conversion to a more modern, efficient and integrated software defined data center. Due to Rubrik's advanced deduplication and compression algorithms we are seeing an average of a 58% data reduction.

Rubrik works fundamentally different than traditional backup software. Rubrik's key to success is understanding that new software and hardware technologies could enable them to combine storage and backup functions. Storage management in rubrik is as simple as click and apply when adding new Briks for additional capacity. The actual backup process is completely different than legacy software. In the past backup software relied on agents installed on servers, Rubrik understood that is burdensome, inefficient and prone to inconsistent backups. The new approach to backups is to leverage API's built in to vSphere and operating systems to pull backups.

During 2017, the storage and backup teams began pushing backup archives to object storage. With the addition of object storage, DTS has been able to infinitely increase the storage capacity available to archival backups. Due to the extremely low price of object storage, DTS has been able to dramatically decrease the cost of cold storage.

Significance:

Storage Component - Physical presence and power consumption of the new storage system is significantly less than the old system. Currently the adaptive flash hybrid array is using the equivalent of two full racks (spread across four racks so that we can easily add more capacity when needed.) In comparison the spindle system occupied 9 full racks, just in our primary datacenter. We have estimated that the power consumption of the adaptive flash hybrid system

is approximately 40 to 60 percent less than its spindle counterpart. The operational costs due to the reduced size and power consumption has been significant for the State of Utah's datacenter.

Backup and Archives Component - While the environmental and space savings are significant with our primary storage system, the savings with our data protection system have been even more significant. With all the components of old backup system it consumed more than 15-20 racks of datacenter space. Rubrik simply lives in two partially filled racks with room to grow. We have estimated that Rubrik uses 90 percent less energy consumption than the old backup system it replaced. Again, this is only for our primary datacenter in Salt Lake City, including our alternate datacenter makes the savings even more significant.

Through the purchase and implementation of these new technologies, DTS was able to complete the infrastructure upgrade of its two data centers in an accelerated time frame with little or no service interruptions due to the nature of the redundant architecture, all the while saving the state and its agencies money. This same architecture design will benefit future upgrades and technology enhancements.

Impact:

Storage Component -

By undertaking this significant change, DTS has been able to reduce core storage infrastructure and CAPEX by migrating from nine (9) fully loaded spindle storage server racks to two (2) full adaptive flash hybrid server racks. As a result, there have been impressive cost reductions for state agencies. In FY2018, DTS was able to realize a 50% rate reduction in primary storage. The savings were so dramatic, DTS issued a significant rebate for the previous fiscal year to state agencies.

Backup and Archives Component -

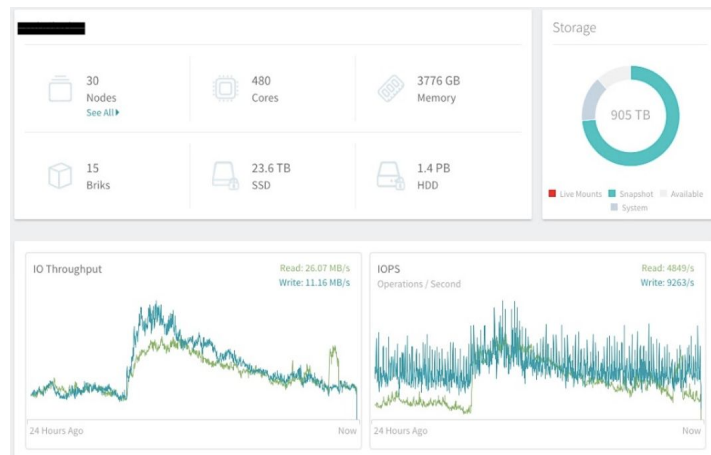
The physical CAPEX saving of the backup solution
Additionally, DTS was able to retire:

- 13 racked backup servers (10 from its Salt Lake data center and 3 from its Richfield data center)
- 2 Oracle ACSLS servers (1 from each of its data centers)
- 2 SL8500 tape libraries
- 2 DD990 Data Domains

Current Rubrik infrastructure is now composed of:

- 15 Rubrik Briks in its Salt Lake data center.
- 7 Rubrik Briks in its Richfield data center.

OPEX in Rubrik has been reduced by an estimate of 50% or more. There are two major drivers for this savings. First, Overall system administration is much more simple than the old system. Second, the time it takes to backup and recover files and servers is quick and done through a simple portal.



Rubrik backup console: Capacity and Performance dashboard