



NetApp®



## *Technical Case Study*

# Drilling Down: Speeding Time to Discovery in the Oil and Gas Industry with Fast Access to Data

How Apache Corporation met intensive performance requirements for geologic modeling and seismic interpretation with NetApp

One of the world's largest independent oil and gas exploration and production companies, Apache Corporation (Apache) is growing quickly, and its data is growing even more quickly. With operations in Argentina, Australia, Canada, Egypt, the United Kingdom, and the United States, the company is continually seeking new drilling opportunities and strives to extract oil and gas resources in the most environmentally friendly way possible.

Limited natural resources make oil and gas a fiercely competitive industry, and Apache's success depends on quick access to large volumes of reliable seismic and interpretation data. Both speed and accuracy are vital, and with drilling rights costing millions of dollars to secure, mistakes could be devastating.

To quickly evaluate and explore potential oil and gas discovery sites, Apache's more than 200 geoscientists feed huge amounts of seismic, well, and other data into numerous applications, including the Schlumberger Petrel seismic-to-simulation software application, for analysis. Improving time to discovery—a key metric in Apache's business—requires giving staff near-instantaneous access to highly accurate information about ideal drilling locations 24/7. Keeping up with growing data and supporting advanced technologies for 3D seismic interpretation and reservoir simulation are an ongoing challenge that Apache must meet without adding headcount to its four-person IT team.

### **Supporting a Shared Seismic Working Environment**

For Apache, time is money, particularly as it applies to geoscientists making the most informed scientific evaluations. Each new version of Schlumberger Petrel adds valuable functionality and gives Apache a better understanding of the subsurface under the earth, but also increases storage and performance demands.



Apache's seismic exploration data has grown by 700% in recent years to more than 18PB. To accommodate this information explosion and make the best business decisions about where and how to drill, Apache requires:

- Fast and easy scalability
- Robust performance for computing, file sharing, and network systems
- High availability for data and applications

The high-performance Petrel software places huge demands on the infrastructure, requiring a 64-bit Microsoft® Windows® 7 desktop environment for increased memory and processing power and the Server Message Block 2.0 and 3.0 (SMB2 and SMB3) file-sharing protocols for efficiency and performance. Providing optimal performance for its compute environments and the 10GbE connections among desktops, servers, and storage systems and meeting aggressive service-level agreements (SLAs) for data availability require a unified, clustered storage environment that allows engineers to perform maintenance on its infrastructure without disrupting data access.

“We have a huge data flow coming into a very complex environment,” says Bradley Lauritsen, director of Exploration Applications at Apache Corporation. “We have drilling operations going on every day, 24 hours a day.”

### How Apache Accelerated Time to Discovery with NetApp Solutions

Apache met its goals by deploying NetApp® FAS6280, FAS6290, and FAS6080 storage systems at its primary (Houston headquarters) and secondary (CyrusOne) data centers for its production environment, with additional NetApp storage systems supporting other operational areas (Figure 1). The NetApp storage is connected to Cisco Unified Computing System™ (Cisco UCS®) blade servers to support Apache's VMware vSphere® virtualized server environment, as well as HP ProLiant servers running Red Hat Linux®.

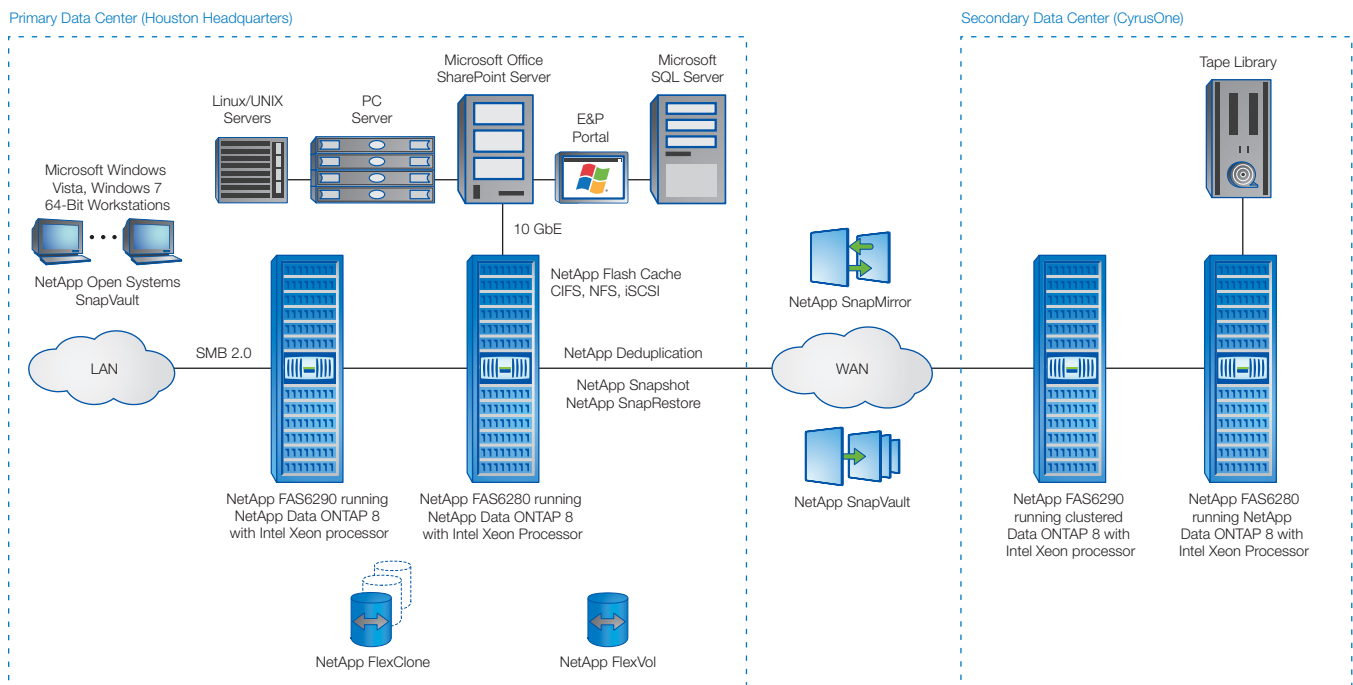


Figure 1) Apache Corporation storage infrastructure. Live data and backups are efficiently replicated to a secondary data center for disaster recovery.

The Apache IT team selected NetApp for several reasons, including:

- The ability to achieve true nondisruptive operations using NetApp clustered Data ONTAP®
- Multiprotocol support, including the SMB2 and SMB3 file-sharing protocols, which deliver as much as 100% performance gains in loading data into systems across the 10GbE network
- Storage controllers powered by the Intel® Xeon® processor, giving Apache more processing power on a per-server basis
- NetApp's relationships with numerous leading independent software vendors (ISVs) serving the oil and gas industry, including Schlumberger, giving Apache the flexibility to optimize storage for its applications (see sidebar, "Deep Technology Integration")

### **Storage Innovations Behind Seismic Interpretation and Workflow**

NetApp technologies play a major role in helping Apache optimize performance and data availability to support the best business decisions.


#### **Optimizing performance and costs with intelligent caching**

One of the technologies behind faster data access is NetApp Flash Cache™, which accelerates data access by caching recently read user data or NetApp metadata. Flash Cache is highly effective at improving I/O performance for random read-intensive workloads and allows Apache to take full advantage of powerful exploration and production applications unhampered by storage bottlenecks.

Two terabytes of Flash Cache cards in each controller improve storage efficiency and, by using Flash Cache with cost-effective SATA drives, reduce Apache's cost per terabyte by 52%. Intelligent caching and NetApp storage efficiency technologies enable the virtual storage tier, which promotes hot data to performance storage in real time without moving data. The company is realizing a 66% power savings in its storage environment by using fewer, larger hard drives for less frequently accessed data, and the storage footprint is 59% smaller with high-capacity SATA drives instead of high-RPM SAS drives.

#### **Improving storage and management efficiency**

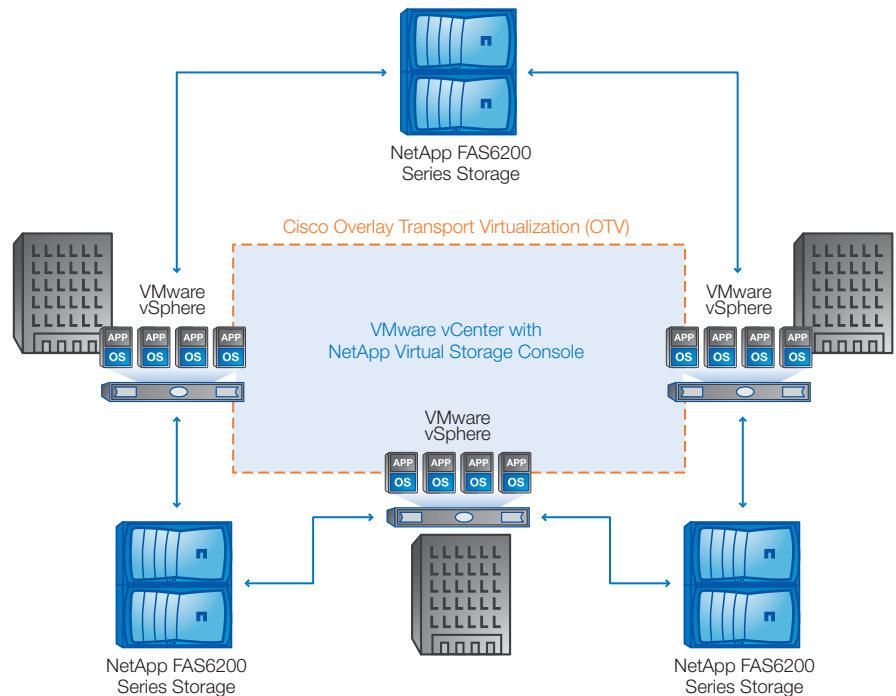
Apache uses NetApp Snapshot™ and SnapRestore® technologies to create hourly Snapshot copies of storage volumes and quickly restore them when needed, enhancing local data protection. NetApp FlexClone® efficiently clones and replicates data volumes and reduces storage requirements in Apache's development and testing environment, helping both operational IT staff and developers be more efficient.



To increase storage efficiency in its production environment by as much as 53%, Apache uses NetApp deduplication to eliminate redundant data blocks within a volume.

To reduce administrative overhead amid rapidly expanding data volumes, Apache uses NetApp OnCommand® Operations Manager. The centralized management console delivers comprehensive monitoring and management and allows the IT staff to easily track performance metrics and utilization statistics to verify that they are meeting SLAs and plan for future needs.

NetApp Virtual Storage Console for VMware vSphere allows Apache IT to access and execute storage management capabilities from VMware® vCenter™, including discovery, health monitoring, capacity management, provisioning, cloning, backup, restore, and disaster recovery (see Figure 2).



**Figure 2) NetApp Virtual Storage Console for VMware vCenter. Apache IT can manage its NetApp storage from within VMware vCenter.**

Apache leverages Data ONTAP 8 and its 64-bit aggregates to store entire databases—such as those that support Apache exploration groups—on a single aggregate while simplifying database management. The unified NetApp storage platform enables the IT team to manage petabytes of storage as a single scalable entity; as a result, the team can manage more storage in less time.

### Enhancing disaster recovery and data protection

Apache uses NetApp SnapMirror® to replicate compressed data from remote sites in Argentina, Egypt, the UK, and Canada as well as U.S. sites in Tulsa, Oklahoma, and Midland, Texas, to its Houston data center. Data is then replicated from the primary data center to a managed hosting facility, where Apache maintains 90 days' worth of backup copies. NetApp Open Systems SnapVault® (OSSV) transfers incremental data changes, reducing bandwidth requirements and increasing storage efficiency.

This approach to data protection has reduced recovery time for large seismic datasets from days to hours, improving both IT and user productivity. Tapes are now only required for quarterly backups, reducing costs and improving reliability. "NetApp SnapVault and SnapMirror have been very efficient at moving data across low-speed or high-latency connections from our global sites," says Lauritsen.

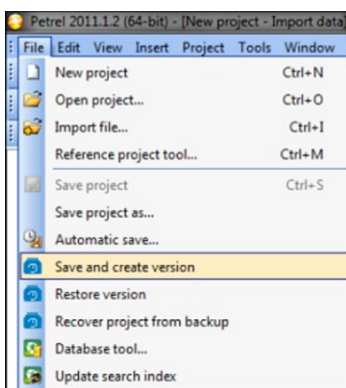
The commitment from NetApp for an open systems API allowed IT consultancy Blueback Reservoir to create a plug-in using NetApp SnapVault and NetApp OSSV agents on Petrel workstations. This enables users to realize considerable application performance benefits while working locally on projects, while enabling IT to regain control over Petrel data (see sidebar, "Deep Technology Integration").

"One of the greatest benefits of NetApp Open Systems SnapVault technology is integrating with the current environment," says Ivan Ganev, senior assistant staff analyst of Apache's Geocomputing Group. "The users see the Snapshot and SnapVault copies from within the Petrel application, and restoring is as easy as right-clicking the data and requesting a previous version."

### Providing reliability and performance for growing exploration data

Apache is in the process of migrating to a NetApp environment built around clustered Data ONTAP, which will give the company a single logical pool of virtualized, clustered storage resources that can scale to tens of petabytes and thousands of volumes. The unified NetApp cluster architecture meets Apache's needs for both scalability and availability while improving performance by rebalancing volumes across controllers/spindles.

"NetApp clustered Data ONTAP enables us to increase performance by utilizing the performance of all of our storage systems throughout the organization, as well as gives us better reliability, because we can now expand the cluster across multiple data centers," says Lauritsen.



### Deep Technology Integration

Blueback Reservoir, a NetApp technology partner and provider of consulting services and software solutions to the oil and gas industry, worked with NetApp and Schlumberger to develop a plug-in for the Petrel application that leverages the capabilities of NetApp OSSV to allow users to back up and restore datasets themselves, without contacting IT for assistance.

Data protection options such as "save and create version," "restore version," and "recover project from backup" are integrated directly

into the Schlumberger Petrel user interface. Snapshot copies of local workstations are created using the OSSV agents and transferred to NetApp storage as a SnapVault archive.

The creation of the custom plug-in has helped eliminate problems with users locally storing and working on data. Apache IT can now provide reliable backup of content previously stored on individual workstations, and because users no longer need to continually save data to new locations for project versioning, data growth can be better controlled.

## IT Benefits and Business Impact

Performance gains from NetApp storage are helping Apache reservoir engineers, geologists, geophysicists, production engineers, and others effectively and efficiently perform their day-to-day jobs. Reliable, unprecedented access to vital data is supporting rapidly growing operations worldwide with measurable performance benefits.

Apache is achieving a 70% cache hit rate and has reduced latency by a factor of 10 or more with NetApp Flash Cache, delivering data in seconds rather than minutes. This means that 70% of the time, data is already stored in memory, eliminating the need for the system to retrieve data from the SATA drives. Combined with the performance advantages of the Intel Xeon processor, this reduction in latency enables Apache to handle large amounts of simultaneous requests on Schlumberger Petrel and other exploration and production systems and almost instantaneously transfer large amounts of data that employees request from the systems.

The SMB2 and SMB3 protocols supported by NetApp have delivered performance increases on Windows 7 systems. The NetApp storage and 64-bit Windows 7 environment have decreased the amount of time it takes geoscientists to run tests during huge survey projects from 10 to 2 minutes. “We’ve been able to load seismic data into cache memory on the workstations at two to three times greater performance than we could without SMB2,” says Lauritsen.

The effects of the performance enhancements on staff activities are significant. One office runs geological models that used to take 20 minutes to open and load into the Schlumberger Petrel application. It now takes less than 5 minutes. “NetApp has given us the ability to grow our storage as we need while providing our geoscientists with reliable, fast access to data,” says Lauritsen. “That makes it possible for them to make highly accurate determinations about where we should drill for oil and gas resources.”

In an increasingly competitive industry where mistakes can cost billions of dollars, Apache’s Exploration Applications team is giving the business fast, uninterrupted access to the data it needs to turn drilling opportunities into assets.

## What’s Next

Apache is continuing to partner with NetApp to improve performance, provide faster access to data, and support other business needs. Some of Apache’s plans include:

- **NetApp Flash Accel™**. Using NetApp intelligent caching at the server level to eliminate the need to store cold data on server flash media, thus reducing server latency and further improving application performance
- **NetApp Flash Pool™**. Mixing SSD and HDD at the aggregate level and automating storage tiering between the two
- **Remote visualization to improve decision making in the field**. Providing data access to remote users on various clients and mobile devices by leveraging the technology integration NetApp has established with solutions from NVIDIA and Citrix for remote visualization of Petrel data
- **Eliminating tape backups**. Using OSSV to efficiently back up data to the cloud
- **Exploring further systems integration**. Uncovering more areas where application-aware integration could improve efficiency for IT and geoscientists

“Storage challenges in the oil and gas industry can be daunting, and NetApp provides excellent solutions to enable us to overcome those challenges and keep moving forward,” says Lauritsen. “Before we had NetApp, we had ongoing issues with downtime and lost data. NetApp provides reliable systems that we count on to help make sure critical data is available for our geoscientists to make accurate decisions.”

### About NetApp

NetApp creates innovative storage and data management solutions that deliver outstanding cost efficiency and accelerate business breakthroughs. Discover our passion for helping companies around the world go further, faster at [www.netapp.com](http://www.netapp.com).

Go further, faster®

#### NetApp Products

- NetApp FAS6290, FAS6280, FAS6040, FAS3270, FAS3240, FAS3170, FAS3140, FAS3050, FAS2240, FAS2040, and FAS2020 storage systems
- NetApp Flash Cache
- NetApp clustered Data ONTAP 8.0 operating system
- NetApp Data ONTAP operating in 7-Mode
- NetApp OnCommand management software
- NetApp deduplication
- NetApp FlexClone
- NetApp FlexVol®

- NetApp SnapMirror
- NetApp Snapshot and SnapRestore technologies
- NetApp Open Systems SnapVault
- NetApp Virtual Storage Console for VMware vSphere

#### Third-Party Products

- Cisco UCS blade servers
- HP ProLiant servers
- Cisco® switches
- Cisco Overlay Transport Virtualization (OTV)
- VMware vSphere
- Red Hat Linux



[www.netapp.com](http://www.netapp.com)

© 2013 NetApp, Inc. All rights reserved. No portions of this document may be reproduced without prior written consent of NetApp, Inc. Specifications are subject to change without notice. NetApp, the NetApp logo, Go further, faster, Data ONTAP, Flash Accel, Flash Cache, Flash Pool, FlexClone, FlexVol, OnCommand, SnapMirror, SnapRestore, Snapshot, and SnapVault are trademarks or registered trademarks of NetApp, Inc. in the United States and/or other countries. Microsoft, SharePoint, SQL Server, Windows, and Windows Vista are registered trademarks of Microsoft Corporation. Linux is a registered trademark of Linus Torvalds. Intel and Xeon are registered trademarks of Intel Corporation. VMware and VMware vSphere are registered trademarks

and vCenter is a trademark of VMware, Inc. Cisco and Cisco UCS are registered trademarks and Cisco Unified Computing System is a trademark of Cisco Systems, Inc. UNIX is a registered trademark of The Open Group. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such. NA-179-0913

Follow us on:     