White paper

+1.857.444.8500
nasuni.com | sales@nasuni.com

Revolutionize Seismic Interpretation with Cloud-Based File Storage

Unleash Seismic Data in the Cloud to Reduce Time-to-Oil

NASUNI

Executive summary

Oil and gas operators slogging through the upstream workflow of seismic data interpretation have become resigned to a slow, cumbersome method of storing and sharing information collected at seismic survey sites. The seismic data they gather, sometimes in extreme environmental conditions, comprises the intellectual property that can make or break their company's fortunes. It's not just seismic files that are critical for exploration and production, but an overwhelming and fast-growing amount of other file data required to run an energy business such as engineering and CAD files, financials, legal docs, contracts, etc. Renewable energy companies are similarly burdened with the same file chaos created in site selection, design, and construction for wind, solar, and thermo investments. Eventually, all these files are stored on outdated media (i.e., tape) and managed on legacy file storage systems not designed to accommodate the ever-expanding volume of data files. Oil and gas companies can only hope that the infrastructure they depend on will stay intact long enough to preserve the invaluable intelligence these file stores contain – particularly important in an industry where merger and acquisition activity is common.

Time-to-oil is hampered in large part by a lengthy seismic interpretation lifecycle. The laborious tasks of collecting, storing, copying, sharing, analyzing and managing seismic data with legacy file-storage infrastructure often can take 12- to-18 months or more. And that's assuming these antiquated systems – engineered for an era that enterprises in other industries have long since left behind – are running perfectly. But time-to-oil also depends on the quality of interpretation of these seismic files. These older file-storage systems keep their valuable information highly siloed, preventing an enterprise from bringing its full expertise to bear on decisions representing millions of dollars of gain or loss.

Cloud-based file storage represents a wholesale shift away from the dependencies and limitations of on-premises legacy NAS, file servers, and backup systems that create silos of data

Fortunately, there's a far better way. As with enterprises in other data-intensive industries, oil and gas companies are now modernizing their approach to file storage and management through cloud-based solutions. Cloud-based file storage represents a wholesale shift away from the dependencies and limitations of siloed, on-premises legacy Network Attached Storage (NAS) and backup systems. Virtual desktop infrastructure (VDI), engineered for the high-performance compute demands of seismic interpretation, with the help of cloud-based file storage synchronized between VDI regions, makes massive seismic data files instantly accessible to an operator's data scientists wherever they are in the world.

Oil and gas companies in particular benefit from moving their data-management processes to the cloud. Cloud storage reduces inherent risks involved in data storage, transfer, backup and security that bog down the interpretation lifecycle. A cloud-based approach to data storage and on-demand compute enables oil and gas companies to navigate market volatility more easily by shifting from a CapEx to an OpEx financial model. With cloud-based storage and infrastructure, operators have the flexibility to scale up or down storage resources as needed without requiring major capital outlays or waiting for annual budget cycles.

Each enterprise will have different starting and ending points when it comes to accessing cloud storage and file-sharing. Some will use high-powered, local workstations; others will opt

for a VDI solution. Whatever the operator's specific needs, Nasuni can recommend an ideal configuration for centralizing seismic data. Each configuration will provide efficient access for a global user base while improving the scale, business continuity, and global recovery requirements such enterprises demand.

The current state: Seismic capture and interpretation is expensive and fraught with obstacles

To assess the value of cloud-based file storage and access, it's helpful to first identify the many challenges that oil and gas enterprises face in relying on legacy file storage infrastructure.

The hazards and challenges of capturing and analyzing seismic images

Operators and service providers undertaking the demanding endeavors of oil and gas exploration are hampered by information systems that are not in keeping with the sophistication of their work. Seismic data collection is difficult enough to begin with – often taking place in environmental extremes – but the real challenges begin when that data reaches IT hubs and has to be stored, shared, analyzed and presented to the user base across the enterprise.

Operators and service providers undertaking the demanding endeavors of oil and gas exploration are hampered by information systems that are not in keeping with the sophistication of their work.

Files are getting larger and larger – now file volumes of several hundred terabytes or even petabytes are required. Limited bandwidth through constricted and costly WAN lines means transmission of seismic data can be painfully slow and suffer repeated interruptions. Transferring the data from the site of collection is necessary due to either the hazards of the local environment, deficiencies in the storage system, latency issues of data transmissions or lack of expertise in proximity to the data. In some cases, oil and gas companies choose to physically move file servers storing data to regional offices where seismic experts can analyze the files.

Limited bandwidth through constricted and costly WAN lines means transmission of seismic data can be painfully slow and suffer repeated interruptions.

Relocating hardware is obviously a logistically challenging operation and it's not always an option legally. Some local, national and regional regulations prohibit data fstored outside the country of origin. In such cases, the company must send geoscientists to the country where the data was collected.

The burdensome work of sharing images within an organization

Much of the time and effort involved in seismic interpretation is spent copying these enormous files and transmitting them through constricted networks. Often, the operators send the field data to a regional office housing powerful desktop computers that run the seismic interpretation software geoscientists use in their analyses. Once there, it's a herculean effort to share the results beyond the regional office so that task is not performed, preventing an enterprise wide understanding of the geoscientists' conclusions.

The process of moving data from collection to the point at which the information can be acted upon can often take many more weeks. Much of this time can be chalked up to the inefficiencies of the process. Any improvement in the efficiency of seismic interpretation activities could significantly reduce time-to-oil. Consequently, an enterprise realizes profits faster and avoids dramatic expenditures — highly prized deliverables for the industry in these volatile times.

Seismic Interpretation via Legacy File Storage and Data Management Processes

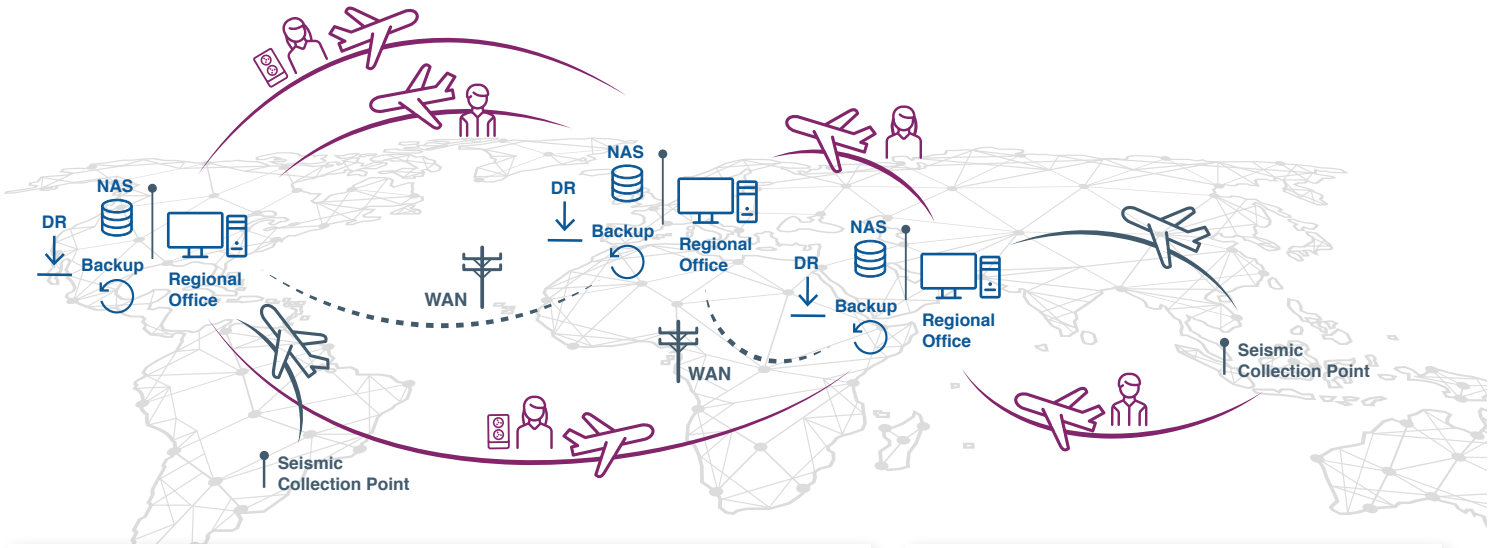

CHALLENGES

Storing data at the point of collection or in regional offices creates data silos.

Moving large files over WAN is slow and unreliable.

Moving seismic data to geoscientists or geoscientists to seismic data is inefficient and costly.

Archiving seismic data on legacy storage media is unreliable.

Recovering lost data is time-consuming and costly.

Multiple instances of data lead to versioning errors and security risks.

Physical transport of data and/or geoscientists

Geoscientists

Legacy storage media

High-performance workstations

The double-whammy of high costs and high risks

Deploying legacy infrastructure to capture, interpret and store massive seismic data files requires a significant CapEx investment. These CapEx expenditures are particularly challenging given the uncertainty of what computational needs will be. As a result, organizations tend to either intentionally overbuy or unintentionally underbuy storage and performance capabilities. And once new equipment is purchased, the enterprise has to maintain and protect the hardware, software and data – another significant expenditure and drain on resources.

Combined with the high costs, legacy systems carry high risks (availability, backup and security of the data) that make the expenditures increasingly hard to justify. To begin with, legacy systems do not lend themselves to reviewing and retrieving stored data assets. When data is stored on archaic media such as disks or tapes, finding relevant information can be slow and labor-intensive. Furthermore, the physical storage materials often deteriorate over time to the point of being faulty or missing critical information.

These outdated storage methods cause oil and gas companies numerous problems:

- There exists potential for errors when an organization needs to act on seismic interpretation that has not been reviewed by multiple geoscientists. They might drill in a place that the flawed seismic data or flawed analysis said would contain significant oil reserves, only to find less than expected. The impact of getting this wrong is simply too costly.
- The value of the stored data assets drops precipitously as the integrity of the data drops. Deterioration can diminish the value of data when the operator wants to sell the assets or if they want the assets to be part of the company's value during an acquisition.
- Operators can ensure that their most skilled (and highly compensated) geoscientists can review the data only by flying them to a location where the analysis is taking place or by transferring the data it to an office closer to the scientists, wasting time and money.

- If the original data has been stored on a medium that has degraded, the organization needs to retrieve a copy stored elsewhere – an additional expense and source of delay in time-to-oil.
- In addition to access, there are problems with data protection, backup and disaster recovery. If stored in the inhospitable environment where the data was accumulated, these data assets are vulnerable to whatever risks threaten the region – extreme weather, a natural disaster or even political unrest. This is why companies often want to fly the servers away from the location to begin with – hardly an ideal way of managing data in the 21st century.

Imagining a better way – a flexible, adaptable, faster and less expensive approach

Imagine a technology solution that overcomes all of the seismic interpretation lifecycle hazards and limitations:

- Operators don't have to contend with the logistics around getting highly valuable geoscientists and highly valuable data in the same place at the same time. Instead, they can empower their geoscientists to access the seismic data from anywhere in the world, all while maintaining compliance with local and regional regulations governing data protection and access.
- On-demand access to seismic, engineering, CAD, and other business file data from anywhere in the world means operators can maximize their geoscientists' productivity.
- Due to a superior method of distributing files that reduces time now spent copying and moving data from one location to another, the time from seismic collection to interpretation can be reduced by 70% to 80%.
- Multiple geoscientists can access massive seismic data files quickly and easily, allowing the most qualified experts to share and review the findings without the delays caused by slow processing speeds or high latency.
- Seismic data files – even those in storage for years – are easily retrieved as necessary for new interpretation or for sale to other parties.
- There are no limits on data volumes. Operators can store and access files within volumes that can be petabyte in size under a single file system and accessed globally.
- No longer does an organization need to keep spending on capital equipment to store and manage the data. Instead, an enterprise can shift to a more predictable OpEx model, with the flexibility to adjust spending as the dynamic oil market shifts up and down.
- The seismic data file storage infrastructure is a business enabler, not a hindrance. Because of the greater speed and accuracy it brings to the process, the new approach enhances an organization's ability to find the right place to drill.
- If the organization decides to sell seismic assets to another company, both buyer and seller can be confident that the data files are valid and the system has helped preserve their full value.

Seismic data in the cloud: The solution that makes the vision possible

All of the benefits of the system envisioned above can be achieved today by harnessing the power of Nasuni® cloud file services coupled with public cloud object storage from AWS, Azure, Google Cloud Platform or private cloud object storage from Dell EMC, Hitachi, IBM, NetApp and others. Together, Nasuni and cloud storage offer a modern file infrastructure that provides the necessary speed boost and unified controls that seismic interpretation requires.

Why centralized, cloud-based file storage with edge caching is essential

Until recently, working with file sizes in the terabyte range has required high-powered compute devices positioned alongside large network attached storage (NAS) devices or file servers. Oil and gas companies had no choice but to contend with multiple silos of compute, storage and data that were unable to access data effectively due to bandwidth and latency limitations.

Fortunately, Nasuni has overcome the silo problem with a solution that brings file infrastructure for oil and gas companies into the 21st century: *Seismic Data in the Cloud*. Nasuni's cloud-based solution allows operators to bypass risky IT installations at inhospitable seismic survey sites by consolidating their seismic data files to a public or private cloud.

Aside from simplifying the technology stack and improving the security posture for their data, the real promise of Seismic Data in the Cloud lies in allowing data scientists and other end users to access these seismic data files from wherever they are, eliminating the need to transport the data to the scientist or the scientist to the data. This means the organization's top subject matter experts can collaboratively apply their knowledge to the seismic interpretation process without having to be in the same place at the same time. Think about the implications for the seismic interpretation lifecycle!

But simply storing seismic data in the cloud does not eliminate data silos and does not automatically make data available to end users everywhere. The data files still need to be orchestrated across locations and made accessible in a user-friendly way. To avoid latency issues that would arise from accessing massive volumes of data across long distances, Nasuni has developed a "hub-and-spoke" architecture. This architecture places cloud storage in a hub and locates frequently accessed data near end users in regional spokes. These regional caches – called edge appliances – enable users to access seismic data at fast, local area network (LAN) speeds. Using this architecture, the Nasuni cloud file-services platform synchronizes data across any number of locations while providing the proverbial "single pane of glass" for managing an otherwise unwieldy global file infrastructure.

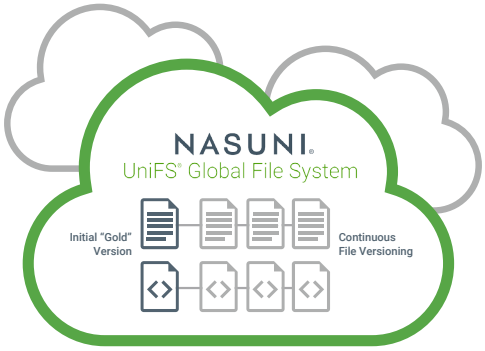

Nasuni provides continuous versioning of files to reflect only the modifications from the previous reading rather than replicating the entire file. Those files that are most frequently accessed are stored in the edge appliance cache, while the "gold copies" of all files remain in cloud storage until needed. Caching active data on lightweight edge appliance virtual machines (VMs) or physical machines and providing centralized orchestration and file synchronization allows operators to reap the benefits of digital transformation that have revolutionized IT processes in so many other industries.

Nasuni supports two desktop implementation models

On-premises workstation model – In this model, physical or virtual Nasuni edge appliances are deployed in conjunction with high-powered workstations in a corporate office. Data is stored in the cloud, but the Nasuni appliances cache the data locally and provide access through standard SMB and NFS file-sharing protocols. Doing so avoids the latency issues associated with transferring data to-and-from the cloud. The edge appliance speeds data transfer to the local workstations that otherwise would require retrieval from the cloud. Additionally, the Nasuni file-services platform orchestrates the synchronization of data across multiple offices, allowing geoscientists in different locations to share the same files.

File Services Platform Built for the Cloud—On Premises Workstation Model

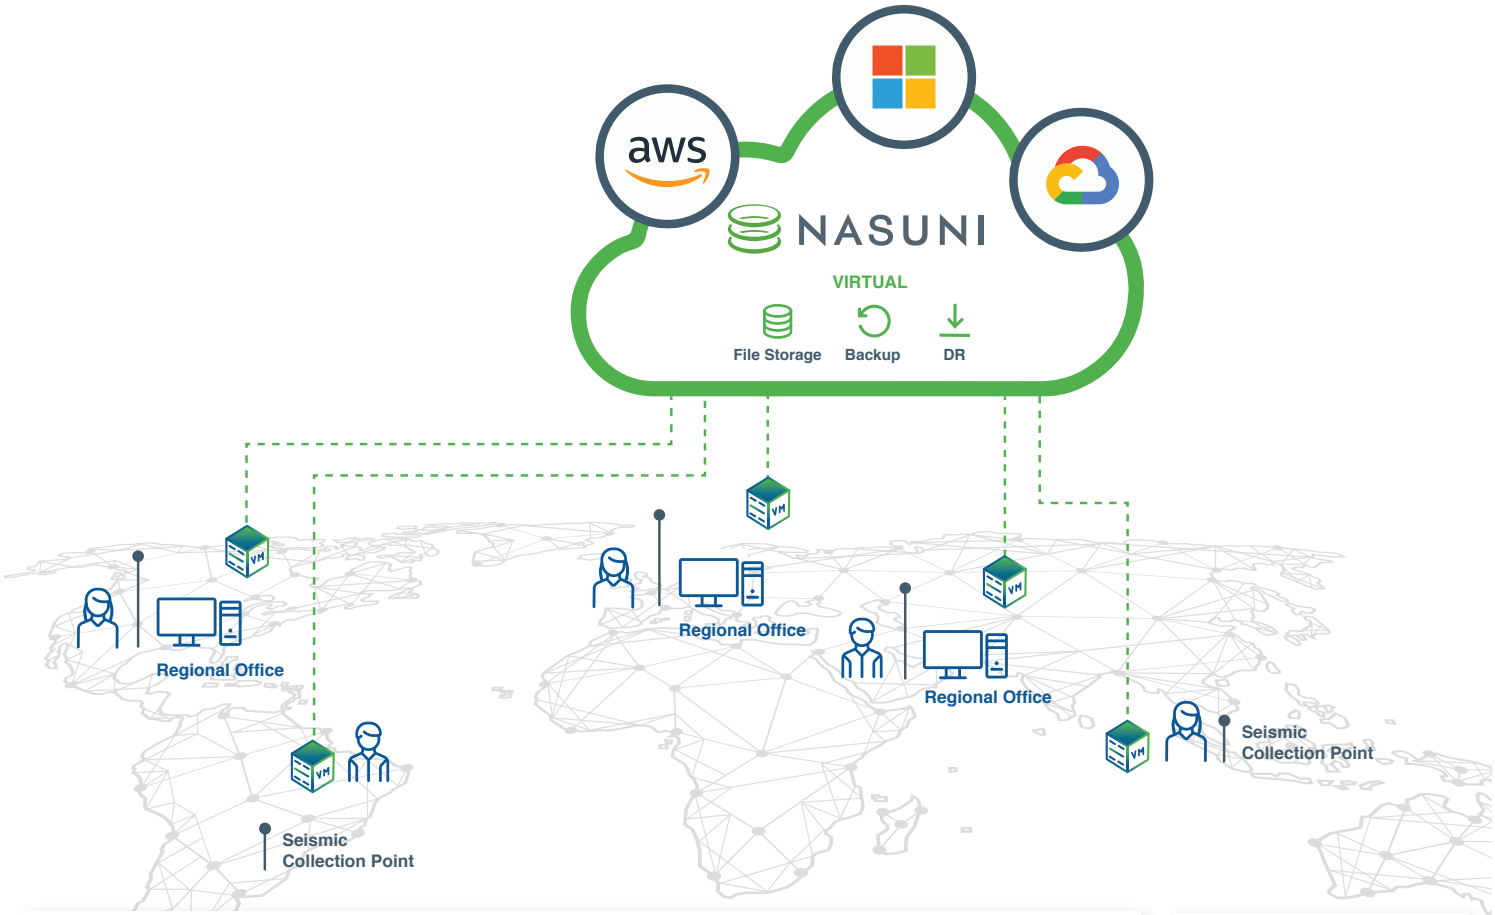

BENEFITS OF CLOUD STORAGE

Operational efficiency: Centralized file storage and data management minimizes seismic interpretation lifecycle and reduces time to oil.

Unlimited Capacity: UniFS global file system scales without limits; stateless Edge Appliances of any size provide fast local access to cached files.

Continuous Versioning: Tracking file changes with infinite versions delivers up-to-the-minute RPO while eliminating the cost and complexity of backup.

Global File Lock: Cloud architecture provides global file lock needed for multi-site file sharing without version conflict.

Flexibility: Simultaneously connect to multiple clouds to meet data security and residency requirements.

Enterprise NAS Features: Integrate QoS, cache pinning, audit, migration, SNMP alerting, antivirus, and other enterprise-class NAS features.

Workforce Productivity: Improve global productivity and automate creative file workflows.

Risk Reduction: Reduce risk of downtime following an outage/disaster and enable granular RTO/RPO.

Virtual desktop infrastructure (VDI) model – Here, both the data and the workstation are virtualized. Rather than utilizing locally based, high-powered workstations, operators work with the data through VMs that can be accessed from anywhere. All data is still stored by Nasuni in cloud storage, but now the Nasuni edge appliances are VMs deployed alongside the virtual desktops ensure high-performance file access. Data scientists in different regions working on thin-client workstations need to connect to a virtual desktop somewhere in their region to reduce latency. Nasuni's file synchronization will cache the active files from cloud storage on an edge appliance VM at each VDI site so the geoscientists can access their data from their virtual desktop without having to travel to a regional office. They can work from anywhere – even at home – and have a user experience on par with accessing the data from a high-powered workstation at the collection site or in a regional office.

Nasuni's file synchronization will cache the active files from cloud storage on an edge appliance VM at each VDI site so the geoscientists can access their data from their virtual desktop without having to travel to a regional office.

Switching to the VDI model provides significant advantages over the local workstation model. Costs are lowered by not having to replace high-powered workstations every three-to-five years. Instead, a thin client can be used, with an estimated lifespan of seven-to-10 years. Additionally, with centralized management of VDI, it is far easier and more cost effective to maintain and update a single master image rather than attending to each workstation individually. And from a security standpoint, VDI is unsurpassed. A thick-client workstation is a security risk because it contains the valuable data in its local memory. A thin-client device accessing a VM does not store the data, eliminating this security risk.

File Services Platform Built for the Cloud—Cloud VDI Model

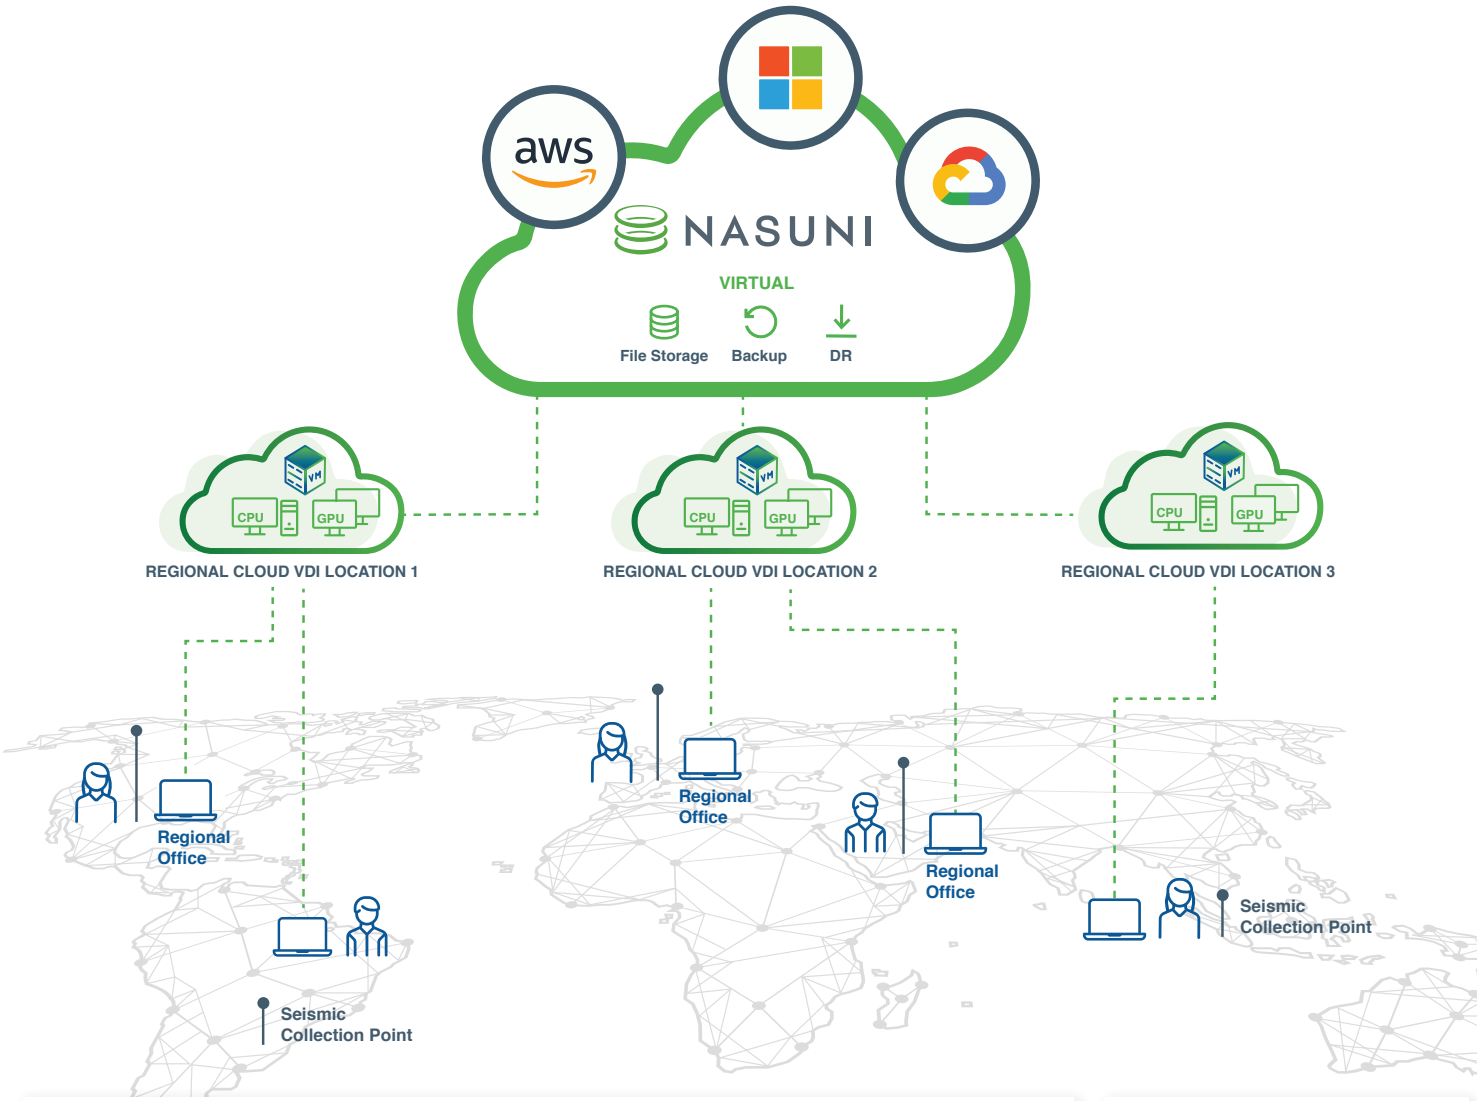

ADDITIONAL BENEFITS OF VDI WITH CLOUD STORAGE

Cloud-based file storage with VDI: Eliminate data silos and protect data.

Virtualized compute with edge caching: Reduce CapEx while offering centralized management of data and seismic interpretation software.

Anywhere Access: Geoscientists do not have to travel to regional offices and can work on thin-client devices from wherever they are.

Easy collaboration: Cloud storage with VDI allows multiple geoscientists to easily undertake seismic interpretation on the same data set.

- Geoscientists
- High-end compute and GPU
- Virtual Desktop Infrastructure
- Thin-client device
- Nasuni Edge Appliance

Nasuni also supports a variation on its VDI-based Seismic Interpretation in the Cloud solution. In this variation, the VDI solution is run in an on-premises data center to support devices in the office. In either VDI deployment model, only the screen and keyboard/mouse information is transmitted to the thin-client device.

Deploying a mixed model

For some distributed oil and gas enterprises, the ideal solution is a combination of the two deployment models: The on-premises workstation model in one or more regions and VDI in other locations. For some organizations, the mixed approach will serve as a transition to a fully cloud-based implementation in the future. For others, the mixed approach will remain the best long-term strategy. Nasuni supports all configurations and can help identify the optimal deployment for any enterprise.

Summary of benefits

The centralization of data means no more data silos thanks to the move to the cloud

No limits on data volumes because Nasuni's UniFS® global file system resides in limitless cloud object storage. Unlike legacy file systems that are constrained by the size of a block-based NAS device or cluster, UniFS has no limits on file, volume, or directory size; number of snapshots; or number of locations it can span.

No need for onsite evaluation of data and no problem with moving data across international borders

Reduced time to interpretation due to less burdensome administration of data, thus improving time-to-oil

Centralized file infrastructure for easier access to, and retrieval of, information

No concerns about compliance with local, national and regional data-management regulations

More reliable storage, free from worry about physical deterioration

Availability of data via VMs wherever geoscientists are located

More manageable, predictable expenditures because a cloud-based file infrastructure allows enterprises to buy what they need rather than guessing too high or too low

Enhanced data protection through encryption, continuous file versioning and an infinite version history for every file

Global recovery of files within minutes for any number of sites

Data optimization through deduplication and compression

Conclusion

The modernization journey for oil and gas companies starts with cloud-based file storage infrastructure. Whether you're an upstream operator or geosciences service provider, you need to be part of the effort to reduce time-to-oil by slashing the seismic interpretation lifecycle. Nasuni is your partner in your transformation journey to a better way of capturing, storing, analyzing and sharing seismic data.

Contact Nasuni to discuss the business case for the Nasuni Seismic Data in the Cloud solution and how it may help your organization improve productivity and operational efficiencies while reducing costs and risks.

800-208-3418
+44 (0)20 7788 8297 (EMEA)
sales@nasuni.com

About Nasuni

Nasuni is a file services platform built for the cloud, powered by the world's only global file system. Nasuni consolidates network attached storage (NAS) and file server silos in cloud storage, delivering infinite scale, built-in backup, global file sharing and local file server performance, all at half the cost of traditional file infrastructures. Leading companies from a wide array of industries rely on Nasuni to share and collaborate on files across multiple sites, enhance workforce productivity, reduce IT cost and complexity and maximize the business value of their file data. Sectors served by Nasuni include manufacturing, construction, creative services, technology, pharmaceuticals, consumer goods, oil and gas, financial services and public sector agencies.

One Marina Park Drive
Boston, Massachusetts 02210
United States
www.nasuni.com
SAL-0082 04/20