A 5-point blueprint for the sanctity of unstructured data for the energy sector

NASUNI

It was in 2012 when I realized that the current technology solutions of the time would not be sustainable, given the scale of data growth. This is what prompted me to specialize in data security.

Malcolm Brown, Manager, Energy Solutions

Table of contents

Introduction

Approximately 80% of a company's data is unstructured, making it an ideal target for a malware or ransomware attack. If your company was hit by malware today, what are the odds that you'd be able to recover your datain full?

In the energy sector, we hear all too often, of companies that are unable to recover their data after a breach. Or in the case of a ransomware attack, we hear that after the ransom has been paid, the data is unusable.

Data is one of the fastest-growing areas in the energy industry, and today's IT managers are faced with a range of challenges, least of all sharing seismic file data across the globe. They need to manage complex tasks such as connecting remote oil rigs and construction sites to the same cloud file system as the users in the office, all the while being prepared for ransomware attacks.

Compromises to intellectual property can make or break a company's fortunes. File data such as engineering and CAD files, financials, legal documents, seismic interpretation, or reserves audits are open to attack from threat actors. Given the large volumes of data in use today, trying to manage file data on complex legacy file storage and archive systems is no longer a scalable, secure solution - particularly when it comes to providing business continuity, and disaster recovery.

"Approximately 80% of a company's data is unstructured, making it an ideal target for a malware or ransomware attack.

As the energy industry’s first blueprint for unstructured data, this guide aims to help energy IT leaders meet their company’s growing storage needs and prepare for what lies ahead in the years to come.

As such, this five-point blueprint aims to help IT leaders meet their responsibility for sanctifying their company’s data. By ‘sanctifying’, we mean putting data first, and acknowledging that if your company’s intellectual property (which contains years of file data) is compromised, you’re putting your entire organization at risk.

This guide has been created specifically for the energy sector and the unique challenges that those IT departments face. The step-by-step actions that we’ve outlined are based on real-world best practices, drawn from the decades of experience of successful IT leaders.

We’ll look at what it means for IT leaders to be responsible for the sanctity of data, why traditional file infrastructure and backups are inadequate for energy companies today, and what actions can be taken to protect your company’s data.

About the author

Malcolm Brown has over 30 years' experience as an IT professional in the energy industry, primarily working across the exploration and production sub-sector. He joined Nasuni from Ithaca Energy, a leading independent oil and gas company, where he was the IT Operations Manager. During his time at Ithaca, Malcolm took a hands-on role to modernize the company's IT landscape. He also led the transition of the IT assets when the business acquired Chevron North Sea in 2019.

Malcolm works with Energy firms to introduce new security models and strategies to counter threats posed by today's geo-political and increasingly interconnected environments.

Throughout his career, Malcolm has worked in several senior management positions within various organizations. He also set up his own consultancies, including Clearscape, where he specialized in subsurface geoscience application deployment on multi-platform technologies.

Notably, he became Head of Information Services at Sterling Resources in 2005, having initially joined as an IT consultant. He redesigned the company's IT infrastructure to incorporate a Wide Area Network, using site-to-site Virtual Private Networks, remote desktop services, remote 3D modeling, and adopting cloud services. He was responsible for information and data management, merger and acquisition data governance, and security operations and compliance.

Malcolm's forward-thinking approach enables him to stay up-to-date with the latest developments while ensuring that any adoption is well researched and planned appropriately, so that any deployment is as seamless as possible. His ability to evaluate how new ideas and technologies are used to incorporate into business models to improve organizational workflows.

Malcolm works with Energy firms to introduce new security models and strategies to counter threats posed by today's geo-political and increasingly interconnected environments. Whether it be physical or cyber, he strongly believes that security needs to underpin all activities related to an organization's assets and intellectual property.

In today's ever-changing energy landscape, data security is very much at the forefront of Malcolm's work. As the leading factor in all the decisions he makes, Malcolm sees his mission at Nasuni as working to protect the sanctity of unstructured data in the energy sector.

Point 01

Be ready to support business growth initiatives

In order to plan for the protection of your company's data, you need to know how the company will grow in the future.

It's likely that your company's current IT requirements won't be the same in five years' time. In communicating with company executives, and those responsible for business development, you'll gain an overview of future plans. These can include organic business growth, future service offerings, possible mergers and acquisitions, and the data acquisition that comes with these movements.

To allow for growth, you need to use systems that can be expanded in an instant. They need to be capable of ingesting data from new sources as quickly as possible so that you always have a single, well-defined version of all the data entities in your company's ecosystem. In other words, you need to have a "golden master" of every data file.

In the case of acquisitions, you need to have a robust process in place that reduces the turnaround time from implementing the acquisition to using the data that you acquire. The sooner you can include new data in the golden master, the safer and more accurate it will be.

It's also beneficial to enhance integrations into existing data sets. The more fluid the integration, the quicker the data can be used, and the less time you'll spend making manual adjustments. During such times, I highly recommend using a dedicated data management team to carry out these tasks.

What to look for:

Scalable - without capacity or file size limitations

Flexible - data instantly available

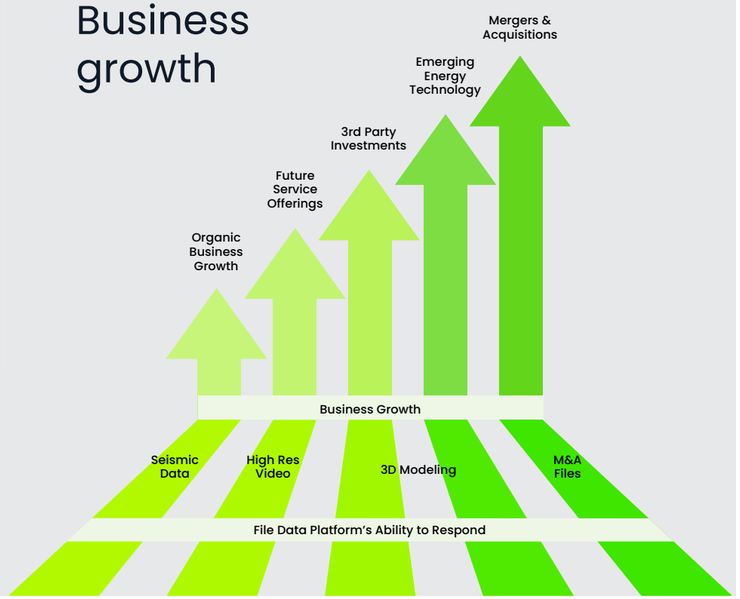

> You need to use systems that can be expanded in an instant.

Point 02

Implement the shortest possible path to final storage from all locations

As IT leaders, we should be aiming to use solutions that have been designed natively for the cloud, taking advantage of the hyperscale capabilities that object storage can offer. In simplifying the whole process, we ensure that our data is inherently safe by design.

We need to reduce the number of times we copy data to and from different types of media in order to prevent data loss and corruption. We should only ever have to ingest a piece of data once. By identifying the shortest path to storage, regardless of how remote the location, we achieve peace of mind that our data hasn't been altered, guaranteeing data authenticity.

Consider the outcome if you were to take data directly from its source, storing it centrally, and making it accessible from a location within a few hours. From survey vessels out at sea, to hydroelectric plants in isolated locations, with files and folders stored in the same place, collaborations are enhanced, lifecycle times are shorter, and IT costs are reduced.

What to look for:

Accessible - anywhere

Secure - data instantly protected

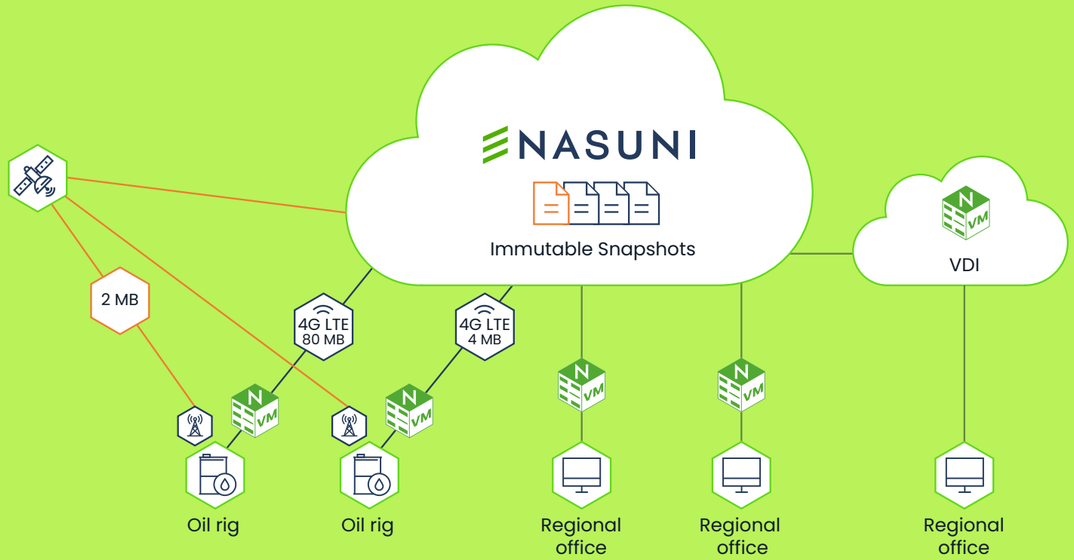

We should only ever have to ingest a piece of data once.

Point 03
Adopt a cloud-first data storage strategy

Most companies are now cloud-first and are going through data center transformations, working with the large public cloud providers.

File storage infrastructure, however, is often one of the last considerations in a company's digital transformation journey. After having moved computing, networking, security and other elements, file storage is often neglected.

With very large data sets, traditional storage methods are time-consuming, and expensive to manage and maintain. Companies in the energy sector are often hampered by a lengthy seismic interpretation lifecycle. Collecting, storing, copying, sharing, analyzing and managing seismic data with legacy file storage infrastructure can often take 12 to 18 months, or more. These older file storage systems keep their valuable information highly siloed, preventing companies from using their full expertise to make critical decisions.

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. They offer greater flexibility and the option of immediate expansion. By breaking the cycle of copying data to new volumes, the instant recovery of data sets is also entirely achievable.

> Collecting, storing, copying, sharing, analyzing and managing seismic data with legacy, file-storage infrastructure can often take 12 to 18 months, or more.

Object storage is the most cost-effective type of storage available, and the cloud makes it the most scalable and durable. There's no need for hardware refreshes every 3-5 years, and you won't run out of storage and have to provision more. Likewise, with cloud storage, you won't be burdened with having to manage remote office location equipment, and you won't have to pay for separate technologies for backup, disaster recovery, and file synchronizing.

What to look for:

Fit for Purpose - built on the cloud, for the cloud

Consolidated - only one master data store

Durable - for today's and tomorrow's requirements

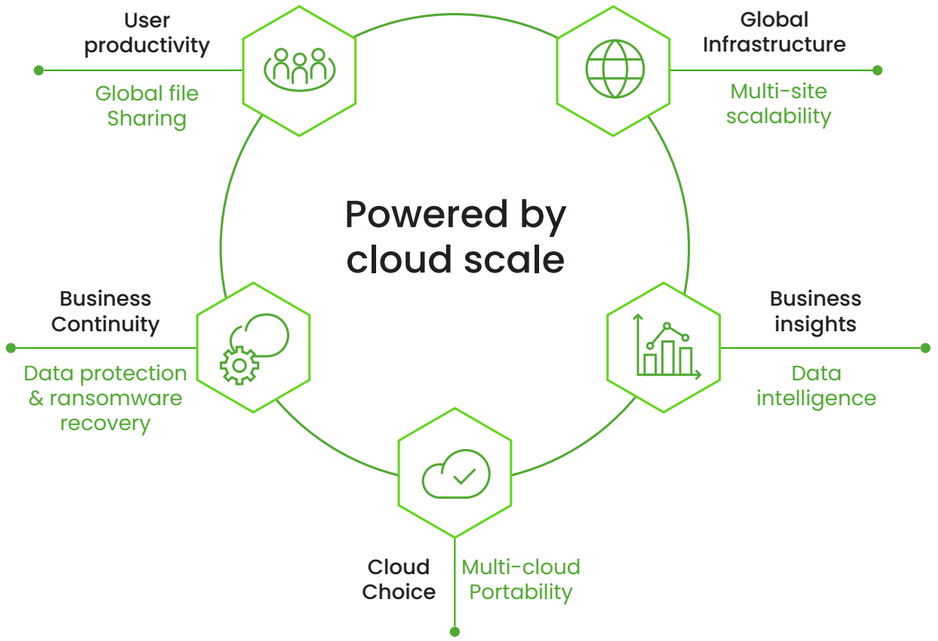

Point 04

Create a golden master of all data from day one, using immutable objects

Create a golden master of all data from day one, using immutable objects.

The original copy of any piece of data, no matter the size, is where the intellectual property exists. We spend millions of dollars acquiring data, interpreting and analyzing it, but we often give less credence to the small spreadsheets and reports, which is where our IP actually exists.

As IT leaders, it is our responsibility to be the custodians of this data and to ensure that it's protected immutably. As such, it should be readily accessible by analysts and engineers from anywhere in the world within hours, if not minutes, of being loaded.

By employing continuous file versioning, we create immutable objects of every subsequent update that is made. Granularity is reduced down to seconds and minutes, vastly improving the time it takes to restore previous versions, if required.

What to look for :

Immutable - file data objects that cannot be modified

Accessible - all versions of any file at anytime

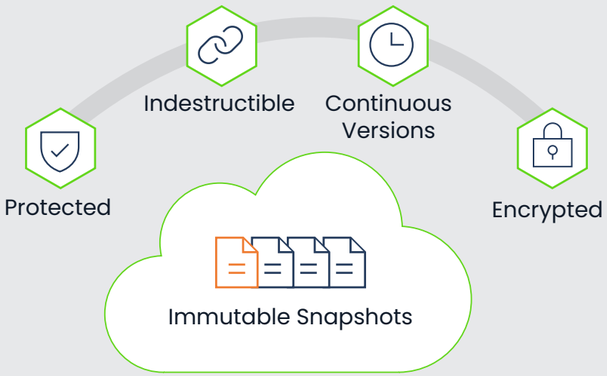

> As IT leaders, it is our responsibility to be the custodians of this data and to ensure that it's protected immutably.

Point 05
Protect, detect, recover

The final stage in protecting the sanctity of data is to adopt a "zero trust everywhere" approach.

This means being extra vigilant in every aspect of our data storage, including cloud infrastructure. To guarantee the safekeeping of our data, inspect firewalls and control traffic in both directions to prevent malicious content from compromising our systems.

In the network itself, consider using AI technologies with autonomous response built in to detect, prevent, and respond to the spread of threats without relying on human intervention. In the middle ground, to triage email and media in a sandbox and quarantine, or remove malicious content before it can be used against the company.

In the cloud, we need to identify and prevent data loss, particularly in our SaaS and file sharing services.

And for end user devices, it's important to be just as vigilant in monitoring and preventing the spread of malicious content. Additionally, we should implement robust user awareness campaigns to educate our end users on what to look out for, should they identify anything suspicious.

Today's sophisticated malware and ransomware technology mean that threats can easily go undetected and sit quietly in the background, learning user behavior and the file systems they have access to. If malware is resident on an authenticated end-user device, it can simulate the actions of a user and gain access to file systems and all the privileges that the user has been granted.

To counter this type of threat, we need file data platforms that have real-time protection, detection, and recovery against attacks. This integrated protection can mitigate the disruption of a threat, and in the worst-case scenario, minimize its impact and get the business back in operation as quickly as possible.

As we mentioned earlier, “traditional” storage solutions can no longer cope with the scale of our file data environments. A backup that needs to be restored from an archive can take many days or weeks to be recovered. It then needs to be validated, and the intermediate backups need to be applied to recover as close to the event time as possible. All data sets then need to be re-validated to ensure they’re clean.

Recovery times for hyperscale data sets of petabytes need to be achievable to within minutes for both RTOs and RPOs. In doing so, we can return the business back to operation efficiently within minutes as opposed to having to wait days/weeks for a full restore (RTO) without the likelihood that up to of 24hrs data could be lost (RPO). This is where continuous file versioning with immutable objects comes into its own.

What to look for:

Scalable - without capacity or file size limitations

Flexible - data instantly available

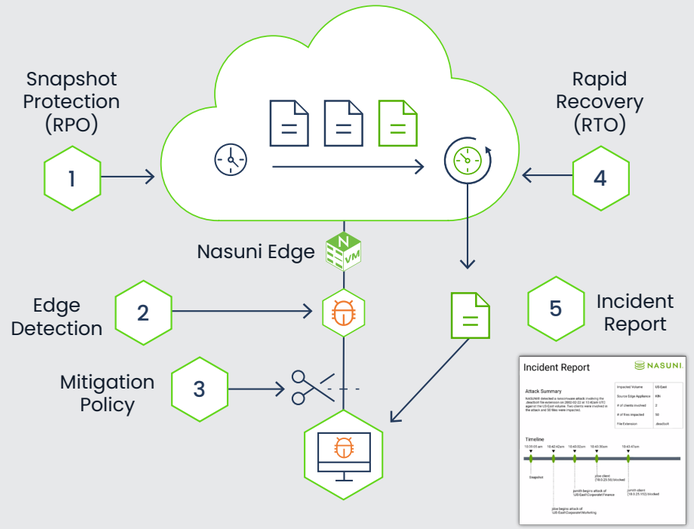

Edge Detection & Rapid Recovery

1. Recovery Points (RPO) as often as 1 minute with infinite snapshots
2. In-line edge detection
3. Mitigation policy to reduce business impact
4. Recovery Times (RTO) measured in minutes
5. Incident reporting for auditing

> Your file data cloud provider should offer an efficient way to move all that data, even at the multi-PB scale.

Let’s talk

Want to find out more about how Nasuni can provide your business with a fluid data infrastructure designed for the hybrid cloud world?

Nasuni’s hybrid cloud platform unifies file and object data storage to deliver effortless scale and control at the network edge.

Learn more

NASUNI