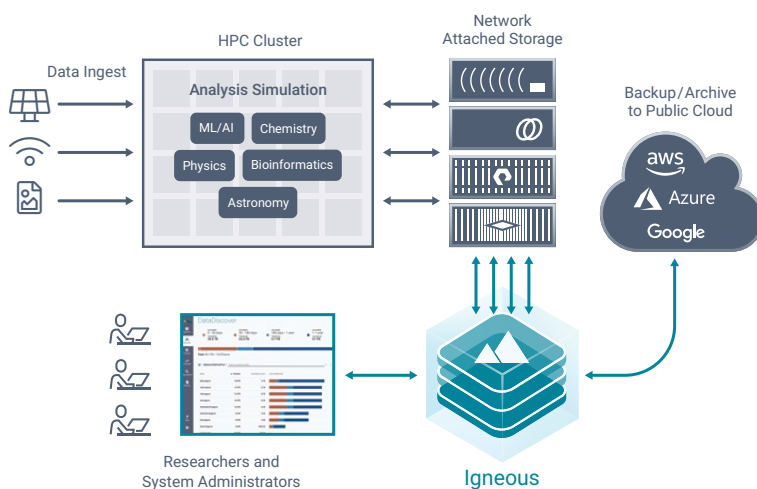


Finding, Moving, and Protecting Data in High-Performance Computing (HPC) Environments

High-Performance Computing (HPC) is quickly becoming a critical platform for many organizations, including research, analytics, finance, life sciences, and more. HPC is being used to solve complex problems, track real time trends, simulate test scenarios, and develop cures for disease. All that progress depends on data, and the data needed to power HPC is growing at double-digit rates every year, coming from more sensors and more software, which in turn are creating more files, which themselves get bigger with each successive update. Today's HPC workloads need consistent high performance and extreme scalability to keep up with massive data growth. Data management in any enterprise can be a challenge, but with petabytes of data that needs to be tracked, delivered, protected and archived at massive scale, older data management solutions are no longer viable in HPC environments.

At the multi-petabyte scale needed to drive HPC workflows, data visibility and management comes with unique challenges. Legacy tools for finding, moving, and protecting data are overwhelmed by the scale and growth rate of unstructured data in today's environments. Organizations that use HPC workflows need scalable, fast, easy-to-use visibility, mobility, protection and archive solutions to manage their data through every stage of its lifecycle. From ingest to simulation to archive, Igneous addresses the challenges that accompany high-performance workflows, enabling organizations to reliably discover, move, and protect billions of new files.

Igneous Unstructured Data Management for HPC Workflows



Visualize, protect and move data through all HPC workflows, including ML/AI, Chemistry, Physics, Bioinformatics and Astronomy

Requirements of Unstructured Data Management for High-Performance Workflows

Consolidated, up-to-date data visibility

Data footprints that can rapidly scale to billions of files pose a particular challenge for legacy solutions that can quickly lose track of new and changed files, leading to lost data and stale analytics. Enterprises need to be able to see and search through all their data, wherever it sits, in order to manage it effectively.

Simple, comprehensive data mobility

HPC environments need to move massive datasets quickly and reliably through all stages of the pipeline, from ingest to analysis to archive. They need solutions that can automate data movement to and from devices, storage platforms, and cloud targets, and which take advantage of today's 10GbE networks to move data at scale.

Scalable, flexible data protection and archive

Petabytes of file data can stress legacy backup and archive solutions that aren't engineered for today's level of data proliferation. To protect their valuable digital assets, organizations need effective data protection that supports the rate of data creation and change in petabyte-scale environments with file-dense workloads.

Performance and throughput

At every stage of their data-management lifecycle, high-performance workflows need high-throughput solutions that protect and move data without impacting production workloads or services, and which can seamlessly move data to and from cloud storage.

Business Benefits

Business benefits for organizations with high-performance workloads using Igneous SaaS data management solutions include:

- **Throughput at scale** that can protect and move billions of files and petabytes of data. See, move, and protect more data, faster than anyone else.
- **Scalable architecture** that can easily grow as needed to support any amount of unstructured data, even in exabyte-size environments.
- **Fast and flexible deployment** on a virtual machine, Igneous can move, archive and protect data from anywhere to anywhere – local NAS, private S3, or public cloud.

Igneous Overview

Igneous provides data management for any NAS, any cloud, and any workflow. Igneous consolidates all data management functions, empowering organizations to fully tap the value of their unstructured data while lowering costs, reducing risk, and optimizing IT resources.

Igneous DataProtect simplifies and automates data protection and archive for HPC workflows. With support for any file and object protocol, any NAS platform, and any public cloud, Igneous DataProtect lets organizations easily consolidate their backup workflows and eliminate vendor lock-in. Even in heterogeneous environments with a mix of NAS architectures, Igneous can protect data from any system, and use any available storage – NFS, S3, or any public-cloud platform and tier – for backup and archive capacity.

Igneous DataDiscover is the fastest way to see all data across all NAS systems – including data age, usage patterns, total capacity, and object counts at any level – all through an intuitive global view showing where data is, how old it is, and the last time it was accessed. DataDiscover can search and analyze data, by name, file type, path, or owner, with results that can be organized into collections or shared with data owners or other stakeholders to help maximize data value, as well as manage the explosive data growth common in HPC environments.

Igneous DataFlow can easily move any amount of data from any NAS system to any other - including the cloud. Moving large amounts of data has traditionally required lots of labor, time, and patience. Igneous eliminates the bottlenecks of legacy tools that were designed for one-and-done data movement, automating the movement of any amount of data and simplifying even the largest migration tasks – all without impacting production systems.

Purpose-built to handle unstructured data at any scale, including hundreds of file systems and petabytes of data, Igneous delivers high-speed data protection, data visibility and data mobility in even the largest environments. To protect business-critical workloads, Igneous monitors system latency during discovery and backup operations, and automatically throttles back on resource consumption as necessary during busy periods to ensure system availability.

Key Features of Igneous Unstructured Data Management:

- **SaaS delivery.** For simplified deployment and streamlined operations, Igneous remotely monitors, troubleshoots, and performs all software updates.
- **High performance.** Igneous moves data in parallel streams to maximize network bandwidth and manage data at scale from any NAS system.
- **Search.** Igneous' scale-out index store provides integrated search and discovery of all file data on primary, secondary, and cloud tiers.
- **Cloud integration.** Igneous integrates seamlessly with all cloud providers, locations, and tiers. Igneous also offers data bundling of small files, and data compression in-flight and at rest, minimizing cloud consumption and transaction costs.

The Igneous Difference



Fast file movement:
up to 100TB per day



Simple solution delivered,
managed, and consumed
as-a-Service



Support for any NAS system,
file protocol and cloud service



Support for on-premises,
cloud, or hybrid backup and
archive workflows



Time to Value in minutes and
hours, vs. days and weeks as
with other solutions

Before Igneous

- Legacy backup infrastructure and processes
- Disk-to-Disk (D2D) replication silos
- Offsite tape rotation

With Igneous

- High-performance, latency-aware data protection
- Consolidated, platform-agnostic data movement
- Backup and archive to and from cloud storage

"That deep involvement of Igneous engineers and experts in actually developing our data infrastructure— not only storage but also the management systems—was absolutely key to building a high performance compute infrastructure."

Thomas Fuchs,
Chief Scientific Officer at Paige

About Igneous

Igneous is the SaaS data management solution for file intensive environments. Data-centric enterprises trust Igneous to provide visibility, protection, and movement at scale. Find anything. Protect everything. Move anywhere.

Igneous: Petabytes today, exabytes tomorrow.

To learn more about Igneous unstructured data management and your HPC workflows, contact us:

1-844-IGNEOUS or 206-504-3685 or info@igneous.io