




Igneous Unstructured Data Management as-a-Service

In the modern, data-driven economy, IT organizations are being squeezed from all sides, facing business pressure to launch new applications, expand services, integrate new cloud services, and meet the ongoing demands of runaway data growth.

Igneous offers the first comprehensive data management as-a-Service solution for IT teams and Data Owners managing unstructured data at scale. Purpose-built to handle billions of files, hundreds of file systems, and petabytes of data, Igneous' solutions help organizations manage their unstructured data at scale, anywhere data lives. Igneous' offerings include **Igneous DataDiscover**, **Igneous DataProtect**, and **Igneous DataFlow**.



See, analyze and act on all unstructured data assets for maximum value with the only unstructured data management as-a-service solution — **At Scale, Anywhere, For Everyone.**

Architecture Matters in Unstructured Data Management

The Igneous architecture includes unique capabilities that deliver these services for customers with billions of files and petabytes of unstructured data:

AdaptiveSCAN™

File system scans that use 75% fewer IOPS than Linux, moving data at line speed through any file structure.

InfiniteINDEX™

A global, infinitely scalable metadata index that maximizes file query efficiency while remaining file system agnostic.

IntelliMOVE™

Highly parallelized, latency-aware file movement optimized for massive unstructured data libraries from any NAS storage platform.

Supporting Your Unique Infrastructure Today and Tomorrow

The importance of choice cannot be underscored enough when it comes to managing unstructured data throughout the lifecycle. Igneous gives choice back to IT teams.

- Key Protocols: NFS, SMB, S3
- NAS API-based Integrations:
 - Dell/EMC Isilon OneFS including multiprotocol
 - NetApp Cluster Mode and NetApp 7 Mode
 - Pure Storage FlashBlade - File and Object
 - Qumulo QF2
- Clouds:
 - MS Azure Hot, Cold, and Archive Blob
 - AWS S3, Infrequent Access, Glacier
 - Google Multi-regional, Regional, Nearline and Coldline



Igneous DataDiscover

Organizations are experiencing explosive growth of unstructured data footprints, reaching the point where a single storage platform, single site, and manual management processes are no longer sustainable.

Igneous DataDiscover provides high-performance, protocol and platform agnostic file system scans by utilizing the Igneous AdaptiveSCAN™. Complete system scans that use 75% fewer IOPS than Linux allows organizations to understand their entire unstructured infrastructure, regardless of scale.

Igneous DataDiscover collects and aggregates file-system metadata from any source into the Igneous InfiniteINDEX™. This data catalog with limitless scalability simplifies locating datasets, files and objects and also enables organizations to answer meaningful questions that lead to effective unstructured data management strategies.

Igneous DataDiscover shows you...	So you can...
Datasets that have not been accessed recently	Continually optimize storage costs by moving cold data to less expensive storage.
Datasets that have experienced significant growth and change	Forecast when and where to expect data generation to accurately manage capacity growth
Data owners that produce and store significant amounts of unstructured data	Develop growth management strategies optimized for the consumption patterns of end users



Igneous DataProtect

Large enterprises using legacy backup solutions find that the rate of unstructured data growth is outpacing the rate at which those legacy solutions can protect it. Further complicating the problem, organizations are now forced to choose between continuing to expand their Tier-1 NAS footprint indefinitely, or migrating data off the primary tier at the same rate at which new data is created.

Any File and Object Protocol, Any NAS platform, Any Public Cloud

Igneous DataProtect simplifies and automates data protection and archive, giving IT teams the ability to offer self-serve data protection services across the enterprise. Igneous DataProtect offers high-performance, protocol and platform agnostic file discovery, which collects and aggregates file-system metadata from any source into the Igneous InfiniteINDEX™. With support for any file and object protocol, any NAS platform, and for any public cloud platform, Igneous DataProtect enables IT and end users to consolidate their data protection and archive operations without vendor lock-in.

Igneous DataProtect is protocol agnostic, centralizing backup and archive data from every NAS instance in the enterprise. Organizations that deploy a mix of NAS architectures can use Igneous to consolidate data from all systems, either onsite, offsite or in the public cloud.

Igneous is the only backup and archive solution that includes full API integration with Dell EMC Isilon, NetApp, Pure Storage FlashBlade, and Qumulo QF2 storage. Using platform-specific programming commands, Igneous simplifies backups by:

- Automating discovery of NAS exports
- Automating permissions provisioning for the entire export structure
- Monitoring overall system latency to ensure backups do not impact business-critical applications
- Managing snapshots for read-consistent backups
- Indexing and backing up both SMB and NFS file permissions for all storage types

For all other NAS platforms, Igneous DataProtect offers parallel, latency-aware backup capabilities and share/export discovery.

To ensure source system performance isn't impacted, the Igneous multi-threaded, platform-independent data movement engine – Igneous IntelliMOVE™ – intelligently deploys or scales back on threads in response to changing system performance requirements. This eliminates any impact to production services and allows file and object discovery to operate in the background at any time of the day.

The Igneous Difference

Latency-Awareness: All data movement is highly parallel and latency-aware, delivering fast throughput without disrupting business-critical applications that depend on the primary storage tier. If Igneous identifies storage performance changes on the primary NAS system during file-system crawl or backup operations, it automatically scales back the workload as appropriate to protect production applications and users.

Faster Backup: After the initial full backup operation, subsequent jobs are incremental only, called "virtual fulls", which capture new and changed files, ensure faster time to completion, and simplify the restore process by providing point-in-time versions of the entire file system.

Simplified cloud-tiering: Through policy-based data movement, datasets can be tiered to any tier of public cloud providers, such as Amazon Web Services, Google Cloud Platform, or Microsoft Azure, allowing enterprises to constrain capacity growth on-premises, and optimize their storage infrastructure.



Igneous
DataFlow

As data usage needs increase, so does the need for Data Owners – not just IT– to be able to copy, move, and sync datasets, at scale, and on their schedule. The process of moving large unstructured datasets, however, is time consuming, resource intensive, and highly complex. As a result, these active datasets often become extremely stationary, severely limiting their value to Data Owners and the overall organization.

To ensure source system performance isn't impacted, the Igneous multi-threaded, platform-independent data movement engine – Igneous IntelliMOVE™ – intelligently deploys or scales back on threads in response to changing system performance requirements. This eliminates any impact to production services and allows file and object discovery to operate in the background at any time of the day.

Igneous DataFlow helps IT teams give their customers the ability to mobilize their own datasets for analysis, simulations and collaboration. By getting the right data to the right place at the right time, data-centric organizations can accelerate business agility, time to insight, and the realization of value from the data.

Outcomes Achieved with Igneous DataFlow

- Empower Data Owners with secure access to datasets to automate a workflow
- Copy, move, sync datasets at petabyte scale
- Automate data workflows

About Igneous:

We deliver the only UDM as-a-Service solution enabling data-centric organizations with visibility, protection and data mobility at scale, wherever datasets and workflows live. Our customers see, organize and understand all of their unstructured data – anywhere. Our customers protect petabytes of data on a single cloud-native platform – at scale. Our customers automate movement of datasets – for everyone needing them. We combine all UDM functions into a single, API-enabled, cloud native solution.

The right data, in the right place at the right time. Find out more at [igneous.io](https://www.igneous.io)