



# Igneous Unstructured Data Management as-a-Service Optimizes Life Sciences Workflows

## Extend the value of high-performance computing

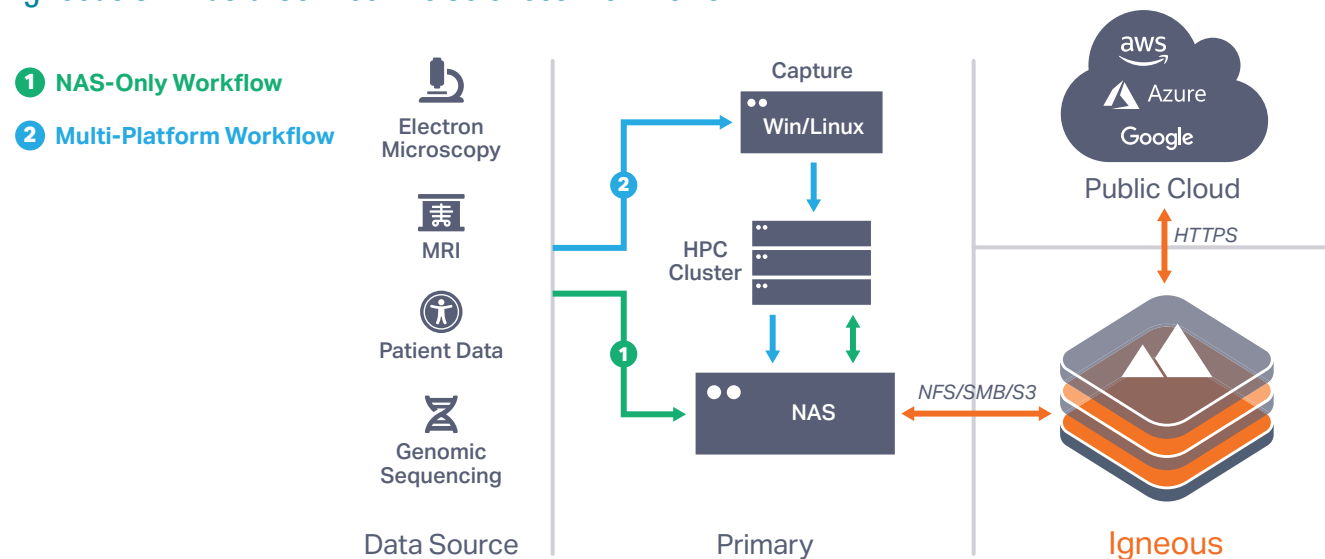
In today's high-performance computing environments, researchers must contend with the challenge of generating, analyzing, and storing enormous data sets for new insights, discoveries, and breakthroughs, all while keeping costs down.

For life-sciences solutions that use high-performance storage for hosting large amounts of raw and/or transformed data, Igneous Unstructured Data Management as-a-Service (UDMaaS) integrates directly with your primary NAS array to extend the value of your platform. Whether you use storage from Dell EMC Isilon for high capacity, Pure FlashBlade™ for high performance, Qumulo QF2™ for deep analytics, or NetApp for flexibility, Igneous UDM as-a-Service solution brings API-level integration to your NAS. Igneous' high-throughput data mover can quickly replicate data from any NAS system. You can even connect it directly to an NFS export or SMB share on your HPC systems.

Once your file system has been added to Igneous, any data written to the target export is automatically replicated to Igneous – where it can be used as a backup image, as an archive repository, or as a content tier – using Igneous' high-performance data-movement engine, IntelliMOVE™. Our patented data protection model offers efficiency and resiliency while our multi-threaded index service quickly discovers and enumerates billions of files at petabyte-scale, providing data protection and visibility at high performance for even your most demanding workloads.

Because Igneous is built on an as-a-Service delivery model, customers get high-performance, cost-effective unstructured data management without the administrative overhead of traditional solutions.

## Igneous UDM as-a-Service Life Sciences Workflows



## Protect Your Data, Extend Your Storage, Upload to the Cloud

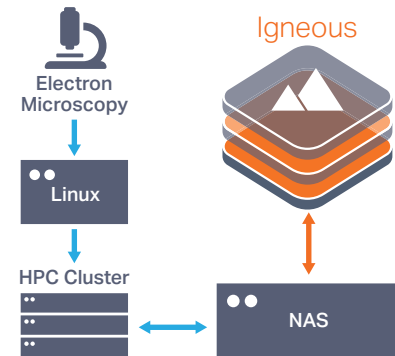
After your data has been offloaded to the Igneous UDM as-a-Service solution, simple user-defined policies let you determine how long it should be retained. Igneous can even upload it automatically to any of the major cloud providers – AWS (S3, IA, Glacier), Google (Regional, Nearline, Coldline), or Azure (Hot, Cool, Archive blob storage) – for longer-term retention, bringing capacity-on-demand and deep-storage archiving to your environment with just a few clicks.

### How Igneous enables an electron microscopy workflow to advance neuronal imaging at a large research institute

Sophisticated neuronal imaging requires advanced microscopy technology with high data output rates. In one real-world life sciences imaging use case, currently in place at a large research institute, an electron microscope kicks off the workflow when it scans a sub-matrix tile of a brain tissue image directly to a high-performance Linux server. Next, the first analysis on the data uses several lens-correction algorithms (perspective correction, distortions, chromatic aberrations) to normalize the data. Finally, the Linux server processes the tiles into high-resolution “Master” copies, which are written to a NAS platform for the next phase.

After this initial image acquisition, in a process that can take hours or days, systems in the HPC environment stitch the tiles into large 2D images, and write the results back to the NAS platform. These images are aligned with one another in layers to produce 3D models. Once the models are created, the original “Master” images are archived. Researchers at the institute then examine the reconstructed models to further our understanding of neuronal function.

Throughout this entire process, Igneous protects the 2D and 3D images against data loss. Igneous backs up between 75 and 90TB of data per day.



## Protect Your Data in any Environment


The following examples show how Igneous UDM as-a-Service can be easily integrated into any life-sciences solution stack.

With an Igneous UDM as-a-Service solution, you can easily and seamlessly manage data protection and data movement to make the most of your solution stack, protecting file data from any source, including the intermediary data on your compute nodes, as well as the result sets on the NAS. Igneous automates data replication and backup – from your primary storage to an archive tier, to a remote site, even to the public cloud - through simple policies that you define.



### NAS-Only

For use cases requiring NFS/SMB storage, shared-file or shared-workload access, or larger overall datasets. Igneous includes seamless integration with all NAS platforms.



### Multi-Platform

For scenarios requiring high-performance storage for ingest, but also needing support for large files or shared-storage access during data-analysis life cycle.

## 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.

To learn more about an Igneous Unstructured Data Management as-a-Service solution, and to see how it can integrate with your High-Performance Computing environment to accelerate and optimize your Life Sciences solution, contact us.

1-844-IGNEOUS / 206-504-3685 / sales@igneous.io