



Igneous Data Management with Pure FlashBlade

Within any industry, data-centric organizations are dealing with ongoing changes that drive innovation, create new revenue sources, and expand business opportunities, but which are also disrupting business as usual.

Data-centric teams need to manage their overall storage capacity, but they also need to protect data as fast as it's generated, as well as migrate older data to archive or cloud storage as it ages out of use. Through all this, they also need to track the location, usage, and movement of every version of billions of files across file systems, sites, and cloud platforms.

Enterprises who use Pure™ FlashBlade™ NAS systems for unstructured data services need a high-performance data-management solution that interfaces directly with their storage platform to index, protect, and move even massive amounts of data without impacting the Pure FlashBlade system's production workloads.

As-a-Service Simplicity with Any Platform

Igneous™ offers comprehensive unstructured data-management (UDM) that was engineered for today's rates of data creation and change. Consisting of Igneous DataProtect™ for backup and archive services, Igneous DataDiscover™ for index-and-search, and Igneous DataFlow™ for data movement between systems and sites, Igneous UDM was specifically designed to support the massive amounts of unstructured data that define the modern economy.

With native API-level integration with Pure FlashBlade, Igneous UDM automatically discovers system exports, manages its own snapshots for read-consistent backups, and balances data-movement traffic across storage blades to optimize throughput. Igneous UDM also includes native support for all Pure FlashBlade data formats, including NFS, SMB, and S3/Object storage – the only data-management solution capable of doing so.

Besides its ability to index, protect, and move data on local NAS systems, Igneous UDM includes native integration with all major public-cloud platforms – Amazon AWS, Google Cloud Storage, and Microsoft Azure – enabling simple, integrated access to online storage tiering and archive services as part of any UDM strategy, even at multi-petabyte scales.



Igneous UDM offers a modern approach to consolidated backup, archive, data movement, and data visibility. Its intuitive user interface provides at-a-glance data management with a few clicks. Igneous UDM is delivered as-a-Service, meaning that Igneous monitors service metrics and handles system maintenance and updates behind the scenes, so that data-centric teams can deliver comprehensive UDM across petabyte-sized data footprints and still focus on core business operations.

Data protection is as simple as assigning a policy to an export. For dataset movement operations, Igneous exposes its APIs directly, enabling simple command-line dataflow at scale between endpoints.

The Igneous Difference

Data movement at line speeds.

With its highly parallel data-movement engine, Igneous uses the full bandwidth of even a 10GbE network to move data rapidly between systems, sites, and the cloud, ensuring that data is always where it's needed, when it's needed.

Scan up to 1.8 billion files per hour.

The file-system crawler engines from a single Igneous instance can scan and index an entire NAS system in minutes, or an entire site in hours. Igneous delivers near-real-time visibility and actionable insights on hundreds of billions of files.

Time to value in days, not weeks.

The software-as-a-service deployment model means there's virtually no customer configuration or tuning required. High-performance data management can begin almost immediately after deployment.

API-enabled for integration and automation.

Engineered to integrate at an API level with Pure FlashBlade, Igneous offers native snapshot management for seamless backup operations, data-path optimization for maximum throughput, and file-system permissions for NFS, SMB, and S3/Object protocols.

Resource-efficiency via fully-

managed service. With a full suite of Unstructured Data Management as-a-Service (UDMaaS) solutions, Igneous monitors operational telemetry, hardware performance, and overall service levels remotely.



Data-Management Challenges in Today's Enterprises

New and faster data sources, such as drones, seismic sensors, and medical-imaging devices, are causing unstructured data footprints to double in only two years. For the owners of that data, the complexity that unchecked data growth brings into the environment presents unique challenges that must be actively and effectively managed, including:

- **Billions of files:** Effective management of unstructured data requires the creation and maintenance of a metadata index for all files in the portfolio, including filename, location, and last access date. A data footprint consisting of billions of files can overwhelm the scanning and index engines of legacy data management applications, leading to stale metadata and a loss of visibility into critical assets.
- **Hundreds of file systems:** Billions of files often require hundreds of file systems. Not only does this stress the IT team who must deploy and manage every one of these systems, but it strains data centers that are out of space, and it stresses the organization itself, which must budget for every file system in the enterprise.
- **Petabytes of data:** Petabyte-plus data footprints are painful to host, protect, and move. Single-threaded protocols, like NDMP-based backups and rsync-based data-copy operations, take too long to move the volumes of data that today's enterprises need to manage, leading to lapsed data protection and ultimately data loss.

Built specifically to address these challenges, Igneous UDM scales to manage billions of files across hundreds of file systems, providing backup and archive services through Igneous DataProtect, enterprise-wide data index and search analytics with Igneous DataDiscover, and hardware-agnostic data movement between both NAS and cloud endpoints with Igneous DataFlow.

Pure Storage FlashBlade: High-Performance at Petabyte Scale

Pure Storage FlashBlade is the industry's first cloud-era, flash-based NAS platform, purpose-built for modern analytics and delivering unprecedented performance for big-data applications. Its distributed architecture delivers consistent performance for all analytics applications, whether using NFS, S3/Object, SMB, or HTTP. Pure FlashBlade is as simple as it is powerful, and all driven by Purity//FB2™.

Pure FlashBlade is built on scale-out metadata architecture, capable of handling tens of billions of files and objects while delivering maximum performance. Storage services at every layer are elastic, thanks to the distributed transaction database built into the core of Purity for FlashBlade. Simply adding additional blades grows system capacity and performance, instantly.

Igneous DataProtect Features

- API integration with Pure FlashBlade for optimal data movement, snapshot management, and automated export discovery
- Native support for NFS, SMB and S3/Object on Pure FlashBlade
- Latency-aware data-movement engine with dynamic throttling to protect production workloads
- Highly parallel data-movement at line speed
- Automated, policy-based backup and archive workflows
- Archive data at the system, export, or directory level
- Point-and-click, policy-driven data replication
- Automated archive to Igneous and/or cloud storage endpoints
- Mount archived data via read-only NFS
- Search to restore via direct download or restore back to Pure FlashBlade
- Logging and audit trails for all data movement

Igneous DataDiscover Features

- Multi-threaded file-system crawler tracks billions of objects per day on all NAS systems and in cloud storage
- Scalable index engine delivers high performance even in multi-petabyte environments
- Deep analytics to provide end-to-end, actionable data insights
- Index all managed objects on all tiers, across all cloud environments

Igneous DataFlow Features

- Hardware-agnostic data movement between Igneous and any NAS or cloud endpoint
- Self-optimizing, able to move data in parallel streams to utilize line speed over any bandwidth
- Exposed APIs to initiate and manage dataset movement

About Pure Storage

Pure Storage (NYSE:PSTG) helps customers build a better world with data. The Pure Storage Data Platform, powered by all-flash storage, offers a simpler, more effective, and more flexible solution for cloud infrastructure and data-rich applications like artificial intelligence, machine learning and Big-Data analytics. With Satmetrix-certified NPS performance in the top 1% of B2B companies, Pure has an ever-expanding range of customers who are some of the happiest in the world.

About Igneous

Igneous delivers the only Unstructured Data Management (UDM) as-a-Service solution, giving data-centric enterprises visibility, protection, and data mobility at scale. Igneous' API-enabled, cloud-native solution combines all UDM functions so that organizations can tap the value of their unstructured data, while reducing risk and optimizing IT resource utilization.

Igneous: The right data, in the right place, at the right time.

Find out more at igneous.io.

Contact Igneous

To learn more about Igneous and about our data migration solutions, contact us:

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