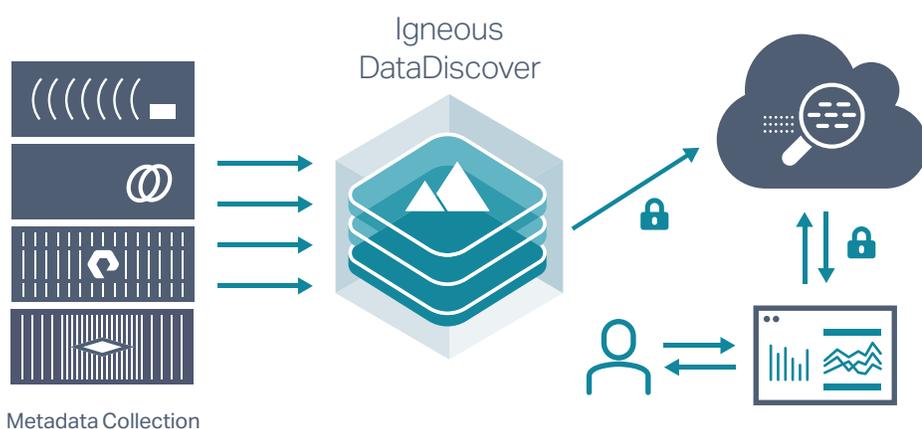


## Igneous™ DataDiscover™ – Deployment Requirements and Frequently Asked Questions



### Deployment Requirements

Virtual Infrastructure	
VMware™ vSphere™	v6.0 or higher <sup>1</sup>
Host resources	
Compute capacity	4-8 vCPU
Memory capacity	16-32GB
Storage	100GB
Network	
Connectivity	Full access to all discovered NAS systems
Port access	443 to cloud.igneous.io and to customer-specific cloud endpoint
IP addressing	1 x IPv4 address (static or DHCP)
NAS Systems	
NFS	Dell EMC™ Isilon™, NetApp™ FAS™, Qumulo™ QF2™, Pure™ FlashBlade™, generic NAS
SMB/CIFS	Dell EMC Isilon, NetApp FAS, Pure FlashBlade

<sup>1</sup> The Igneous DataDiscover virtual machine may be compatible with earlier versions of VMware vSphere, but full functionality has not been tested or verified.

### THE IGNEOUS DIFFERENCE

**Simple.** Easy to deploy via a stateless virtual machine, Igneous is delivered operated as-a-service. Once deployed, Igneous DataDiscover is self-optimizing, with no further customer configuration needed.

**Fast.** Capable of scanning up to 17 billion files per day with little to no impact on NAS performance, Igneous' file-system crawlers can scan and index an entire NAS system in minutes-to-hours, not days to weeks. Igneous delivers near-real-time visibility and actionable information on hundreds of billions of files.

**Scalable.** Built for the most demanding file environment, scanning hundreds of thousands to hundreds of billions of files.

**Action-oriented.** A global view of data allows fact-based conversations with data owners to drive outcomes such as deletion or archive (through Igneous DataProtect™) of stale datasets.

**Resource-efficient.** With a full suite of Unstructured Data Management as-a-Service (UDMaaS) solutions, Igneous monitors operational telemetry, software performance, and overall service levels remotely, dramatically reducing overhead for IT teams.

**Try Igneous for Free**

Visit:

[www.igneous.io/datadiscover-free-trial](http://www.igneous.io/datadiscover-free-trial)

## Frequently Asked Questions

### How is Igneous DataDiscover deployed?

Igneous DataDiscover is deployed using a stateless virtual machine, which is hosted on the customer's own virtual infrastructure. After the initial setup and configuration are complete, the virtual machine's operational status and performance are monitored as-a-Service by Igneous. The overall process requires very little on-premises infrastructure to deploy, with fast time to results.

### How does Igneous DataDiscover work?

The on-premises Igneous DataDiscover virtual machine intelligently scans the metadata of all files on all exports hosted on all registered NAS systems. The collected metadata is then securely uploaded to a dedicated, single-tenant, Igneous-managed cloud instance to be processed and aggregated into a summarized, interactive view.

### How does the virtual machine communicate with the Igneous cloud instance?

There are two endpoints that need to be accessible to the virtual machine. One is [cloud.igneous.io](https://cloud.igneous.io) for all service monitoring and software updates, the other endpoint is the customer's specific cloud repository for metadata upload and processing. The virtual machine communicates with each endpoint via outbound-only access over port 443.

### How is metadata secured?

All metadata scanned by the Igneous DataDiscover virtual machine is compressed and encoded into a proprietary binary format. Data in transit is encrypted as it is uploaded to the Igneous cloud instance via TLS over HTTP (HTTPS). This uploaded data is sent to a customer specific endpoint where all metadata is secure and isolated. Each provisioned Igneous cloud instance is single-tenant and customer-specific.

### How fast can the virtual machine scan?

A single 4-vCPU virtual machine can scan up to 200,000 files per second if source systems do not throttle performance. Actual scan rate can vary depending on NAS system configuration and other active workloads.

### Will Igneous DataDiscover scanning impact my system performance?

No. While Igneous DataDiscover uses NAS system resources to scan metadata, the scan service monitors latency as it's scanning and will throttle back if NAS system latency starts to increase.

### What is needed to deploy and configure the Igneous DataDiscover virtual machine?

Once deployed into the local vSphere infrastructure, the virtual machine will need to be configured with a single IPv4 address that provides visibility to all in-scope NAS systems, as well as TCP443 access to the public Internet. Both static and DHCP-assigned addresses are supported.

### Why is administrative access requested for Dell/EMC Isilon, NetApp FAS, Qumulo QF2 and Pure FlashBlade?

While administrative access is not required, it greatly simplifies the scanning process. With administrative access, Igneous DataDiscover will leverage API access to manage share permissions, take snapshots to scan, and release snapshots when complete. There will be fewer instances of "Access Denied" errors to shares, exports and files.

### For generic NAS systems, how is the scanning experience different?

Explicit permissions will need to be granted for the Igneous DataDiscover virtual machine's IP address to have administrative access to scan data on a generic NAS system's exports or shares. This may require more setup work.



## Frequently Asked Questions (continued)

### When does scanning start and how often does it run?

Scanning starts on first import of a system. End-to-end scanning repeats weekly on Sunday.

### What does Igneous DataDiscover report off?

Reporting is based on the collected metadata from all files and directories that have been scanned by Igneous DataDiscover. Size as reported is based on file size; file count includes all files and directories, and age views are based on each file's mtime (since atime is often not enabled for large-scale NAS systems and is therefore not a reliable metric).

### What happens if I discontinue Igneous DataDiscover service?

The virtual machine can be terminated and deleted from the on-premises vSphere cluster. The customer's cloud-based Igneous instance will also be terminated, and all data will be deleted.

### How does Igneous compare to other file analytics tools?

The biggest differences between Igneous DataDiscover and any other packaged file analytics tool is our simplicity, speed, and scale. Igneous is designed for billions of files, petabytes of data, and deep/wide directory structures across all NAS systems. Unlike other file-analytics tools, there is no software to update, no database or elastic search clusters to manage, and no tuning or optimization required.

## Try Igneous for Free

Visit: [www.igneous.io/datadiscover-free-trial](http://www.igneous.io/datadiscover-free-trial)

### 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](http://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](mailto:info@igneous.io)