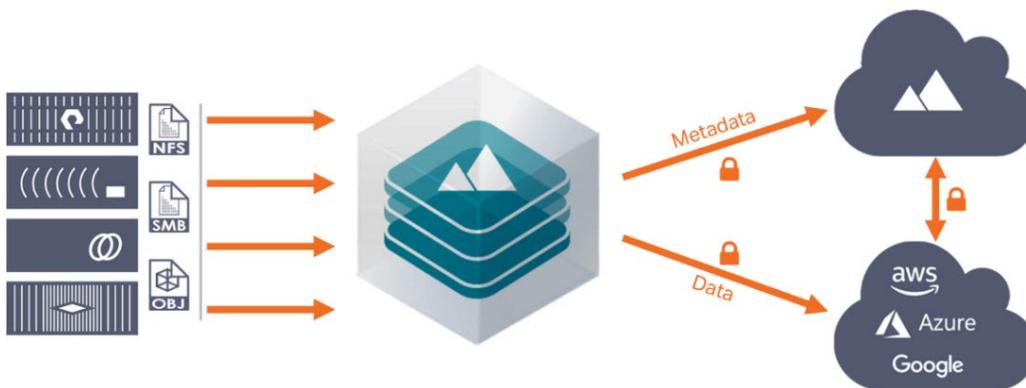


Igneous DataProtect Virtual Machine – Deployment Requirements and FAQs



Deployment Requirements

Virtual Infrastructure	
Hypervisor environment:	VMware ESXi v5.5 to 6.7 ¹ VMware vCenter v6.0 to 6.7
Virtual machine	OVA deployment
Compute capacity	4-8 vCPU
Memory capacity	32GB RAM
Storage	100GB disk space
Network	
Local connectivity	Full access to all discovered NAS systems
Outbound connectivity	HTTPS (port 443) to: <ul style="list-style-type: none"> - cloud.igneous.io - Customer's cloud account(s) – Amazon Web Services, Google Cloud Storage and/or Microsoft Azure - Customer's cloud endpoint(s)
IP addressing	1 x IPv4 address (static or DHCP)
NAS Systems	
NFS	Dell EMC Isilon, NetApp FAS, Qumulo QF2, Pure FlashBlade, Linux, GPFS, Lustre, Weka.io, Stornext, Gluster, BeGFS, Dell EMC Unity, Synology, any other NFS systems
SMB/CIFS	Dell EMC Isilon, NetApp FAS, Pure FlashBlade, Qumulo QF2
Administrative access (optional)	Dell EMC Isilon, NetApp FAS, Pure FlashBlade, Qumulo QF2

The Igneous Difference

Simple. Easy to deploy through the stateless virtual machine, Igneous leverages a software-as-a-service deployment model. This means there's virtually no customer configuration or tuning required, with same-day turnaround in nearly every instance.

Scale out. On-premises instances integrate highly scalable backup and archive software with scale-out index and object stores to manage billions of files.

High performance. Igneous performs highly parallel data movement, optimized for scale-out primary NAS systems from NetApp, Dell EMC Isilon, Pure Storage FlashBlade, Qumulo, and others. Dynamic throttling eliminates impact to users and applications accessing the primary storage. As such, backups can run continuously, eliminating the concept of a "backup window" where users and applications can't otherwise access data.

Search. A key problem with enterprise network attached storage (NAS) is not knowing what's there. The Igneous scale-out index store provides integrated search and discovery of all file data on primary, secondary, and cloud storage tiers.

Cloud integration. In a virtual machine-based deployment, Igneous DataProtect leverages your own public cloud capacity (any provider and any tier) for all backup and archive storage, eliminating any future concern about capacity expansion or backup storage refresh cycles.

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



Frequently Asked Questions

How is Igneous DataProtect deployed?

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

How does the virtual machine manage customer data?

During backup operations, backup data is uploaded to the customer's own cloud endpoint and storage tier using settings defined and provided by the customer. File metadata – information associated with each file such as create and modify times, file size, etc. – is uploaded by DataProtect to an Igneous-managed cloud bucket, where it is used to provide index and search capabilities.

How are data and metadata secured?

All data uploaded by DataProtect is compressed and encoded into a proprietary binary format, and is encrypted during upload via TLS over HTTP (HTTPS). Data is stored in buckets in the customer's account so the customer always retains control over their data. Each customer's metadata is stored in its own Igneous-managed bucket, which is single-tenant and customer-specific.

Backup and archive data are encrypted via HTTPS (TLS) while in transit, and are secured and protected by the cloud provider's own perimeter infrastructure while at rest on your cloud storage.

Which public cloud providers can DataProtect use as a storage target?

DataProtect can backup and archive unstructured data to AWS S3, Google Cloud Storage, and Microsoft Azure Blob Store. All S3 tiers on all three cloud platforms are supported.

What kind of performance can I expect with DataProtect?

Backup and restore performance are heavily dependent on the source NAS system, the file count and average size, directory depth and width, the speed of the connection to the customer's cloud provider, and the capabilities of the virtual infrastructure hosting the Igneous virtual machine.

For smaller files (<2KB), scan and compare rates can be up to 114,000 files per second, while compare-only can be more than 500,000 files per second. Larger files (average >400 kB) can approach 16,000 files per second. The Igneous DataProtect VM is capable of filling the network to full capacity.

Will DataProtect impact my system performance?

No. While DataProtect uses NAS system resources to scan and protect, the service monitors latency at all times, during all tasks, and will throttle back its resource consumption if NAS system latency starts to increase.

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 both the initial setup process and ongoing operations. With administrative access, Igneous DataProtect can leverage API-level integration to automate the discovery of new shares and exports, to protect export and share permissions, and manage file-system snapshots for point-in-time backup consistency.

For generic NAS systems, how is the setup experience different?

Explicit permissions will need to be granted for the Igneous Virtual Machine's IP address to scan, read, and modify file-system data on a generic NAS system's exports or shares. This may require more setup work on the NAS system to implement.

What happens if I discontinue Igneous DataProtect 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 metadata will be deleted.

Backup and archive data are stored in the customer's own cloud storage, so can only be accessed or deleted by the customer directly. To prevent data loss, it is recommended that customers recover any needed backup or archive data from their cloud storage before discontinuing their Igneous DataProtect service.

Can I use the same virtual machine that I have deployed for DataDiscover?

Yes!

Contact Igneous

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

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

About Igneous

Igneous offers the only Unstructured Data Management solution to be delivered as-a-Service, 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.