



# Taming Cloud Costs with Igneous

## Cloud efficiency and cost control at any scale

Today's enterprises are facing unprecedented challenges. Budget cuts and hiring freezes are forcing them to do more with less. Existing data centers are out of space for new infrastructure, and many can't be expanded any further. Through all of this, their data footprints are still growing – with millions of new files and terabytes of new data every day.

While data-center constraints are driving many enterprises to cloud storage, uncertainty about cloud costs are causing managers to hesitate to adopt cloud for backup and archive use cases.

Igneous enables effortless integration with all public-cloud providers and storage tiers, and its ability to compress and consolidate unstructured data before uploading it to the cloud minimizes both storage and transaction fees. As a backup and archive solution, Igneous delivers comprehensive unstructured data management that lets enterprises move data to and from the cloud, without the added worry of per-file storage and access costs.

### Proven Storage Efficiency

Igneous recently completed a series of backup tests to demonstrate and quantify our ability to minimize cloud storage and transaction costs in backup and archive use cases. We created a series of unique datasets, each characteristic of data found in different industry verticals, ran them through our backup engine, and compared the resulting backup images with their original datasets.

We found that for all data types and industry verticals, Igneous can lower consumption costs by up to 70%, and transaction costs by up to 99% compared to other backup and archive solutions.

Enterprises concerned about the high cost of "warm" cloud storage tiers (e.g. Azure Hot, Google Standard, or AWS Standard) will see their consumption costs minimized by Igneous' ability to compress data. Organizations who are concerned about the high transactional costs of "cold" tiers (Azure Archive, Google Ice Coldline, or AWS Glacier/Deep Archive) will see their transaction costs reduced by 95-99% by Igneous' data bundling. All enterprises will see cost reduction from Igneous' ability to write data directly to the cheapest cloud tiers.

**Consumption  
Cost Reduction: Up to 70%**

**Transaction  
Cost Reduction: 95-99%**

### Full Cloud Protection at Minimal Cost

Most backup solutions protect individual files by simply replicating them in their native state directly to cloud storage. If the backup/archive target is a warm cloud storage tier, then this approach maximizes consumption costs by not compressing the data first. Replicating data to a cold storage tier instead may offer reduced consumption costs, but it also incurs an ingress fee for every file uploaded. At large scales, these per-transaction charges quickly add up: a billion-file dataset – increasingly common in today's enterprises – could rack up \$50,000 in ingress fees alone!

Ingress fees also apply to incremental backups, in which only changed files are copied to cloud storage, meaning customers incur a transaction fee for every single file that changes. Restore operations, which pull files from cloud storage, incur similar egress fees. A single day's backup and restore tasks using other data protection solutions can cost thousands of dollars.

Igneous significantly reduces these costs – both consumption and transaction fees – by bundling individual files into "blobs" during the backup/archive process, and then compressing the resulting data before sending it to cloud storage.

Igneous' use of data blobs means that large numbers of files – in some cases, thousands of individual files – can be moved to cloud storage in a single transaction. Data compression improves backup performance by moving more data using the same bandwidth, and lowers costs by reducing the amount of cloud storage needed to host it.

The net effect, for customers challenged by both scale and cost considerations, is a backup solution that scales to any size while minimizing cloud costs.

### THE IGNEOUS DIFFERENCE

**Designed for Cloud**  
Igneous integrates fully with all public cloud platforms, letting enterprises backup, archive, and/or copy data to any cloud provider and any storage tier.

#### Engineered for Efficiency

For backup and archive workflows, Igneous bundles files into "blobs" and compresses data before writing to cloud storage.

#### Built for Cost Control

Igneous minimizes both storage costs and cloud transaction fees, even for petabyte scale backups with billions of individual files. For cold-tier storage, Igneous manages data retention settings to ensure data is deleted in the most cost-effective way.

#### Proven for Scale

Where other backup solutions fail at scale, Igneous was designed and built for data protection in any size environment. Combining high-speed data movement with seamless scalability, an Igneous deployment can easily grow to protect petabytes of data and billions of files every day.

## Test Data

The five unique datasets were standardized in size, but with object counts and average object sizes corresponding to those found in different vertical industries, with specific dataset parameters shown in Table 1 below:

1. **Electronic Design Automation (EDA)** – uses a very large number of very small files of different types
2. **Enterprise File Services (EFS)** – a broad mix of data types, but tending overall toward smaller numbers of larger files
3. **High-Tech Manufacturing (HTM)** – large number of small files, consisting of CAD/CAM designs, blueprints, test plans and results
4. **Life Sciences (LS)** – primarily imaging and analysis data, tending toward large-file datasets
5. **Media and Entertainment (M&E)** – widely varying file sizes and data types

	EDA replica	EFS replica	HTM replica	LS replica	M&E replica
<b>Dataset size</b>	4TB	4TB	4TB	4TB	4TB
<b>Object count</b>	80 million	425,000	25 million	1 million	30 million
<b>Avg. file size</b>	54KB	10MB	172KB	4MB	143KB

**Table 1: Dataset parameters**

Each dataset type was staged twice: once using non-compressible files, and again using data that was calibrated for 50% compression. We then ran a full, Level-0 (L0) backup of each dataset, including both compressible and non-compressible instances, and compared the overall results.

## Test Results

After running L0 backup tests for all datasets we compared their total size and object counts with the size and object count of their corresponding backup images. The results are itemized in Table 2 below:

	EDA	EFS	HTM	LS	M&E
<b>Dataset size</b>	4TB	4TB	4TB	4TB	4TB
<b>Object count</b>	80 million	425,000	25 million	1 million	30 million
<b>Compressible / non-compressible size</b>	2TB / 4TB	2TB / 4TB	2TB / 4TB	2TB / 4TB	2TB / 4TB
<b>Backup object count</b>	176,100	18,000	75,300	19,100	64,000
<b>Total object reduction</b>	99.78%	95.76%	99.70%	98.09%	99.79%

**Table 2: Consolidation results for all datasets**

## Interpreting the Results

While individual results may vary, and are dependent on each organization's particular data profiles, customers can expect to see the following benefits from an Igneous-based backup and archive solution:

Cloud storage fees will be minimized by data compression. Not all data is compressible, but even minimally compressible data will occupy a lower footprint on cloud storage than on local NAS. Most enterprises will see at least some compression, and some will see significant reduction of their data footprint during backup – anywhere from 10%-70% of their overall data. Compressed data results in lower costs on any platform, particularly on cloud storage, and even more so on hot/warm storage tiers.

Igneous reduces cloud transaction costs by up to 99% as data is consolidated for backup or archive. Even datasets with an average per-file size of 10MB saw object-count reduction of more than 95%.

Minimizing consumption costs is critical for enterprises who use hot or warm-tier cloud storage. Keeping transaction fees low is important to those who use cool or cold cloud storage. Igneous delivers cost savings for all cloud storage use cases – consolidating files to cut transaction costs, and compressing data to lower the amount of warm-tier storage needed.

With Igneous, the cost for an initial backup of the above-mentioned billion-file dataset would be reduced from \$50,000 to as little as \$110!

## ABOUT IGNEOUS

Igneous is the SaaS data management solution for file-intensive environments. Data-centric enterprises trust Igneous to provide visibility, protection, and movement at scale. The API-enabled, cloud-native solution combines all data management functions, allowing organizations to fully tap the value of their unstructured data while reducing risk and optimizing IT resources. Igneous unstructured data management is uniquely architected for the scale and performance of the modern enterprise environment.

## Contact Igneous

To learn more about Igneous and your unstructured data management workflows, contact us:

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