



# The Elastic Storage Network for Elastic Stack

## The Elastic Stack

Elastic.co has created an open source software stack for analyzing machine-generated data in order to gain insights into customer behavior, system security, and other aspects of the digital business.



The Elastic Stack is based on the distributed search engine, Elasticsearch, the Kibana data visualization tool, and Logstash and Beats, which pipeline data into Elasticsearch for analysis.

## Elastic Stack satisfies a voracious appetite for insight

The beauty of Elasticsearch is the ease with which it can grow to handle more data or more analyses. You simply create new Elasticsearch servers (“nodes,” which can be virtual machines or containers) and connect them to the cluster of nodes you already have running. It’s analogous to opening more check-out lines in a grocery store to handle customers during peak shopping hours. This is in very sharp contrast to traditional SQL data warehouses, which are extremely costly and difficult to grow and change.

## Satisfying an elastic appetite for storage - “Goldilocks and the Three Storage Platforms”

As your Elasticsearch deployment grows, so must the storage space available to it. Storage for Elasticsearch should be fast, secure, and able to grow as agilely and economically as your Elastic Stack. You have three options,

which are not unlike the bowls of porridge Goldilocks encountered at the three bears’ house:

- **Too hot: On-premises storage** - Shared storage, like a SAN or NFS, is fast and secure, but most definitely not elastic. It’s a fixed resource that IT increases from time to time by forecasting how much storage you’ll need for the long term (say 18 months) and then procuring that. IT effectively over-provisions (and overpays for) your storage hardware to accommodate future growth. This process involves research, funding, acquisition, installation, and configuration of the storage systems, which can take weeks or months. This reduces the agility and cost effectiveness of your Elastic Stack-powered analytics initiatives.
- **Too cold: Public cloud storage** - Amazon Web Services and other public cloud platforms provide storage on demand, with usage-based pricing. It’s elastic and economical - grow or reduce the amount of storage services you consume, and your cloud bill rises or drops, respectively. The drawback is performance; network latency, the time it takes for data to make round-trips from the cloud to your Elastic Stack for processing, is too long, which can make the Elastic Stack experience for data scientists and developers prohibitively slow.
- **Just right: Hybrid cloud storage** - The answer is an entirely new generation of Elastic Stack-ready networked storage that extends the public cloud storage edge into your data center and into nearby storage hubs. The hybrid approach combines the performance and security of on-premises enterprise storage with the elasticity and economics of the cloud.

## The ClearSky Global Storage Network for machine data analytics

ClearSky Data's global storage network simplifies and reduces the cost of data management for the Elastic Stack by delivering enterprise storage that spans the entire data lifecycle as a fully managed service.

ClearSky reinvents data storage by securely combining the performance and reliability of your on-premises storage with the economics and scalability of public cloud storage. The result is a fully managed service that delivers low-latency data access and elasticity that data scientists and developers need, at about one-third the cost of alternatives.

The ClearSky service allows you to scale analytic data on demand, minimize your data center footprint and eliminate the complex management of primary, backup and disaster recovery storage infrastructure.

### Storage designed for analytics agility

- **Scalability on demand** – Simply scale your storage as needed, using the ClearSky dashboard to add unlimited capacity with a few clicks.
- **Agility without sacrificing performance** – Traditional cloud storage is economical, but can be slow and unpredictable. ClearSky's global network extends the cloud edge into your data center and nearby points of presence (PoPs), which results in fast, reliable data access speeds.
- **Simplified data lifecycle** – ClearSky's patent-pending Smart Tiered Caching™ continually categorizes data and places it in the right cache layer (hot, warm, cold) based on evolving usage requirements and customer policies.
- **Embedded data protection** – By building redundancy into all cache layers and network paths, ClearSky is able to provide an enterprise-class analytic data storage service with guaranteed high availability.
- **Enterprise-grade security** – ClearSky is a fully managed service, operated by enterprise storage experts who apply data encryption in flight and at rest, dedicated connections, and other security best practices to ensure customer data is always protected and isolated from the moment it enters the ClearSky network.
- **Fully managed service** – ClearSky is a service-level agreement (SLA)-based fully managed service.

## Dramatically improved agility, performance, and total cost of ownership (TCO)

The total cost of ownership (TCO) of storing analytic data using the ClearSky service is about 80% lower than the cost of managing it in a company's data center on physical storage arrays. It can be scaled up or down immediately, on-demand, to meet your workload requirements, which is strongly aligned with the scalability and agility provided by the Elastic Stack itself.

### Next Steps

To understand whether ClearSky can improve your analytic agility and reduce your costs, we invite you to visit our website and request a meeting or demo. Our data storage experts are happy to discuss your requirements, explain the ClearSky service in detail, and map out a provisioning timeline.

**For additional information, please visit [www.clearskydata.com/elastic](http://www.clearskydata.com/elastic)**