

Optimizing Performance and Cost Savings for Elastic on Pure Storage

This article originally appeared on [InsideBigData.com](https://insidebigdata.com).

There's no denying that data is growing at an unprecedented rate and access to this data is crucial for enterprises seeking insights for business growth. As data continues to expand, so does the demand for fast query performance and quick insights.

Across [Pure Storage](#)'s install base, organizations use Elastic as a tool to store, search, and analyze data at scale, across a broad range of use cases. While some use it within their applications for search and indexing, others use it for log analysis where they analyze application, infrastructure, or [security logs](#) to trace problems and find root causes to issues.

Whatever the use case, the proliferation of modern applications, IoT, DevOps, and [microservices](#) have resulted in an explosive growth in infrastructure, and ultimately growth in machine-generated [log data](#), which is critical in preserving system health, troubleshooting issues, and deriving valuable insights. As a result, Observability and Security Operations teams are not only challenged with rapid growth in data and increased infrastructure complexity, but also with keeping their systems running and doing faster root cause analysis of issues.

Challenge: Balancing Performance and Cost as Elastic Deployment Scales

Generally, Elastic is used as a central repository to analyze logs originating from distributed infrastructures. As an organization's usage of Elastic grows, enterprises bring more and more data and logs into Elastic. Over time, the log files are becoming larger. In addition, malware may be within an enterprise's firewalls for weeks before it is detected. It's no surprise that organizations want to retain security-related logs much longer.

What happens when data volumes begin to reach hundreds of terabytes? As applications scale and the demands on infrastructure increase, often operational challenges arise in maintaining the overall system. These may include unpredictable search performance, infrastructure complexity, underutilized

resources, and increase in operational overhead. Organizations considering Elastic Cloud see their costs increase significantly as well: Enterprises pay Elastic for the infrastructure resources consumed in the cloud, the number of days the data is retained, and for maintaining their Elastic cluster.

Not all organizations utilize the Elastic Cloud offering. Data residency, [data sovereignty](#), regulatory and compliance mandates, and rising cloud costs are some of the reasons why enterprises explore an on-premises deployment of Elastic. Even when organizations use Elastic on-premises, they do want a cloud-like experience and get the most performance from Elastic while optimizing costs.

Use the Right Storage Choices with Elastic Data Tiers

The first step organizations can take to balance performance with cost is to use the Data Tier capability in Elastic and pair each Tier effectively with a storage solution that best matches its performance and capacity requirements.

In some instances, data gets searched less often as time goes by. Logs, metrics, and transactions, for example, may be of less interest to users as days/weeks/years go by. Elastic allows users to effectively manage growing volumes of data through [Data Tiers](#). With Elastic Data Tiers, users can select between hot, warm, cold, and frozen tiers to optimize storage cost and query performance based on the needs of their use case.

Pure Storage® provides storage solutions suited to get the most out of each Data Tier. Organizations can store data that is accessed the most in the expensive hot tier and search it in milliseconds. When the Elastic cluster is relatively small with extremely low latency needs, Pure [FlashArray//X™](#) is a commonly used solution. When the Elastic cluster is relatively large with high throughput needs, FlashBlade//S® on NFS is most commonly used for the hot tier.

All infrequently accessed data for which query time of seconds or minutes is acceptable, can be stored in the cheaper cold, or frozen Data Tier.

Use Cold and Frozen Data Tiers with Capacity-Optimized Storage

Use of [cold and frozen tiers](#) can help reduce infrastructure costs with minimal performance impact. Elastic cold and frozen Data Tiers use searchable snapshots. Searchable snapshots rely on the same snapshot mechanism an organization

already uses for backups and allows for automatic retrieval of frozen data from an S3 bucket. Organizations can search across data in snapshots stored on low-cost object storage. Generally this is the data that is accessed infrequently, is read-only, with query performance of seconds or minutes.

Data stored in cold or frozen tiers eliminate the need for replica shards in Elastic as these are now stored in the low cost storage as snapshots. This reduces the storage consumed by almost 50% and as a result Elastic storage costs. By offloading historical data to cold and frozen tiers and storing less data in the hot tier, organizations reduce the amount of data processing that has to be done in the hot nodes. This in turn significantly reduces the number of hot/warm nodes required and subsequently compute costs. While all of this sounds compelling, the use of cold/frozen data tiers has not gained as much traction as organizations still need storage performance from the low cost object storage supporting cold and frozen tiers. Pure Storage uniquely provides an all-flash, capacity-optimized object storage technology that provides the needed performance for the cold and frozen tier at disk economics. Organizations can augment their existing hot tier, and offload historical data to FlashBlade//E – a capacity-optimized, yet performant S3 bucket.

Pure Storage and Elastic Deployments

As enterprises grapple with escalating volumes of data and the challenges of balancing performance and cost in Elastic deployments, Pure Storage is a strategic partner providing tailored solutions. Pure Storage's FlashArray and FlashBlade offerings address diverse needs, from small clusters with low-latency requirements to large clusters demanding high throughput. The introduction of FlashBlade//E in 2023 further revolutionizes the landscape by offering all-flash, capacity-optimized object storage, addressing the longstanding performance concerns associated with low-cost disk-based solutions.

Organizations can now confidently embrace Elastic, enhance their hot tier storage, and seamlessly manage historical data with cost-efficient capacity-optimized storage. Pure Storage not only meets the demands of the modern data landscape but also empowers organizations to simplify their Elastic architecture, reflecting the industry trend towards a more streamlined and efficient approach.

[Learn how to enhance log analytics for infrastructure, applications, and security operations on Pure Storage.](#)