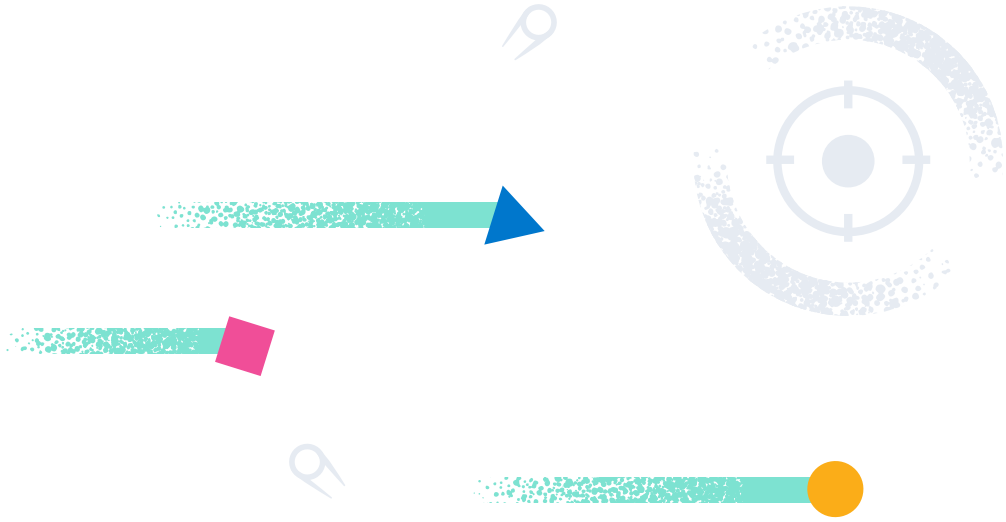

elastic | The Search AI Company

Optimizing SIEM performance

Understanding Elastic's data tiers

Comparing performance and cost figures for SIEM evaluations is rarely simple. Apples to apples comparisons of multiple vendor's storage/performance tier nomenclatures is almost impossible, as terms like HOT, WARM, COLD, FROZEN, and OBJECT don't mean the same thing across technologies.

Put simply, not all vendors' "Hot" tiers provide the same data search speed.

Data tiering has long provided various mechanisms for migrating data from the fastest, most expensive tier of storage for frequently accessed read/write data (Hot), all the way down to slower spinning disk and even object storage for low cost read-only archives (Frozen).

Here, we define our own requirements in the **Elastic Security** solution to help you compare our technology against other vendors in as true of a comparison as possible.

Security data tier overview

Let's start by defining what we need in various "security data tiers" to be able to effectively balance performance, retention, and budget in a modern security operation.

The table below represents a breakdown of Elastic Cloud storage tier options for Elastic Security for SIEM. Note that not every vendor will provide each of these tiers, nor do we expect every client will need to implement each tier to meet their security goals.

Storage Tier	Used for	Performance / Response	Cost	Elastic Cloud Tier
Ingest	**Read-Write** Log ingest, normalization, & enrichment, ML inference and anomaly detection, rule processing, case and investigation management	**Fastest** < seconds		Hot
Interactive	**Read Only** Near and mid term investigations, analytics, threat hunting, data aggregations, weekly metrics	**Fast** seconds		Cold
Historical	**Read Only** Long term historical analysis, threat hunting, KPI analysis, compliance live data retention requirements	**Baseline** 10 seconds +		Frozen
Offline	**Read Only** Retention requirements and compliance concerns	**Recovery Time Object driven**		Snapshot

*Elastic Cloud also has a warm tier available, but in a security context it is typically only required to support high volume concurrent investigations in MSSP or enterprise environments with large security analyst teams.

**Self-managed deployment sizing can be customized to your environment. We recommend working with your Elastic Solution Architect for optimal sizing.

Elastic Security for SIEM performance

While it's certainly possible to size an Elastic Security for SIEM implementation with a bulk of the storage allocated to the Hot tier, in practice, leveraging the appropriate tiers for your use case and security teams' requirements will result in a much more cost-effective deployment that meets or exceeds other vendors "hot" tier performance.

In the table below, we've gathered metrics to showcase real-world use cases such as threat hunting queries, SIEM dashboards, visualizations, and dashboard performance across Elastic Cloud storage tiers. All tests were performed multiple times against 7 days of data (~275 million events, ~ 3TB data) on a shared demo system, with performance speeds denominated in seconds (s)

Test	**Ingest** (Hot)	**Interactive** (Warm)	**Interactive** (Cold)	**Historical** (Frozen) uncached / cached
Queries				
Simple Query (all events with user name root)	.4 s	.45 s	3 s	45 / 6.5 s
Event Correlation Query Hosts with >3 failed ssh admin logins	6 s	8 s	9.2 s	18 / 10.1 s
SIEM Dashboards				
Detections & Response	.35 s	.55 s	1.62 s	16 / 1.8 s
Kubernetes	4.7 s	8.2 s	9.2 s	34 / 13.5 s
Entity Analytics	1.6 s	2.2 s	2.5 s	12 / 8.5 s
Visualizations and Dashboards				
Simple Visualization (Root user authentication events over time)	4.1 s	5.6 s	5.9 s	24.5 / 17 s
DNS Threat Hunting Dashboard (12 visualizations, multiple aggregations)	9.25 s	10.5 s	12.7 s	35.3 / 24 s

SIEM performance on the Elastic Cloud ingest/hot tier is second to none. The interactive/cold tier shows a 29% average performance hit compared to the ingest/hot, which is still more than fast enough (and more economical) for most security team needs. Our historical/frozen tier leverages object storage and adaptive caching on each cluster node to enable the best possible performance for fully live searching and analysis — at the absolute best price.

Once you've aged your data out to offline storage, you'll still be able to quickly access that data by re-mounting it on your frozen nodes, providing access to offline data in minutes. This contrasts with other vendors, who require additional hot nodes and 24 hours to thaw their frozen data.

Case studies

So how does Elastic's data tiering stack up in real-world applications? Let's take a look at companies that leverage the efficiency of Elastic data tiers:

THG

With Elastic snapshots, THG slashed their physical storage costs by 60%. And it didn't come at the expense of performance, either — the team also reduced their mean time to respond to security incidents by 60% with Elastic.

> Elastic allows us to store large volumes of data. We're able to store that data for long periods of time in frozen, which reduces the size and the cost of having to store it long-term.
>
> **Ryan Kennedy**, Head of Security Engineering, THG

personal CAPITAL

Personal Capital's deployment of Elastic Security has dramatically decreased their mean time to respond and has given them more control over their security data lifecycle.

> We can configure Elastic Security to our exact needs, especially when it comes to data conservation. We use Elastic's [Historical] Tier and Searchable Snapshots for more cost-effective hot and cold data management.
>
> **Eric Sekercan**, Security Team Lead, Personal Capital

The Elastic pricing model is much more cost-efficient. Personal Capital now pays for how data is used, not the volume of data consumed. Instead of paying for features that aren't needed, the security team can fine-tune Elastic Security for greater control and visibility over their annual spend.

OP SYS

Managed security provider OpSys enables their customer base to take advantage of Elastic Cloud as part of OpSys' business model. This allows customers to access, manage, and retain their security data for regulatory and business purposes. Storing data in Elastic's Frozen Tier minimizes compute costs without teams taking too big of a performance hit.

In addition, Elastic's **cross cluster search** allows OpSys to secure, log, and search their customer's security sources from their central SOC, even if the customer's data is outside the region.

Effective data tiering

Speed skate on your frozen data

With Elastic, data tiering is available for all security and observability data, providing greater flexibility in how you store, search, and analyze data. Unlike other solutions that require their "Hot" tier just to respond to queries, Elastic queries, analytics, and machine learning capabilities can all run fast — even in the "Active" and "Frozen" tiers.

Now, real-time search queries take milliseconds, not seconds. And historical queries take minutes, not hours.

Other vendors:

Accessing historical data can be painfully slow. Data in frozen tiers often has to be restored before searching, and users may have to wait up to 24 hours for the data to be searchable – time that can have serious consequences when you're facing issues that impact your customers and revenue.

Sort, get set, go

Pre-indexing data at ingestion time with Elastic means users don't need their data stored in the fastest storage tier to provide the same levels of interactivity. Users experience better price/performance, without the need to re-hydrate data to act on it.

Other vendors:

Require users to keep all their data in their fastest tier for interactive performance. Transfering data into these vendors' "Hot" tiers comes with additional costs and manual input for analysts.

Predict and control costs easily

Elastic's entire platform offering is a single, unified product, priced via a transparent, resource-based **consumption model**. Unlike other vendors, whose SIEMs often come with high costs and complex pricing and licensing structures, Elastic's simplified approach can save you money on both licensing and infrastructure through efficient data tiering.

If you're trying to rein in your hot storage costs, Elastic's data lifecycle management feature automatically stores your data in the most cost-effective tier, so you don't need to worry about slowing your operations teams down.

Next steps

Feeling like it might be time to explore a new SIEM technology that can deliver on your desired outcomes in the most cost-effective manner? You're not alone. **44% of organizations** are looking to migrate or augment their existing SIEM. See whether it might be time to **replace your SIEM**.