



Data Sheet

Apica Architecture

**Built for infinite scale.
Industry-lowest TCO**

At Apica, we architect and build our software for infinite scale. Infinite scalability as part of our Day 0 thinking. Our system is also designed for elasticity. Any modern technology that processes data needs to be architected as such.

Most, if not all, observability systems Today are built on scale-up models which need reconfiguration and in some cases, reconstruction to ingest and process higher volumes of data. Sudden surges of data volumes due to a special or unforeseen event incapacitate the system.

Apica's design incorporates architectural components and a philosophy that allows infinite scale.

We built the architecture on the following tenets.

- Auto scale-out architecture
- Decoupled storage and compute
- Infinite compute
- Infinite storage
- Lowest TCO
- Real-time performance



Auto scale-out architecture

Data growth rates today demand the use of scale-out architectures. Scale-up architectures cannot meet the demands exerted by the sprawl and proliferation of data sources.

The Apica architecture is growth-focused and allows the addition of compute and storage resources instead of merely increasing the capacity or specifications of existing resources. Unlike a scale-up architecture, our architecture built on a native K8S (Kubernetes) architecture automatically adjusts to increased data streams or data rates by proportionally adding bandwidth, compute, storage, and throughput.

We build our computing systems on Kubernetes containers and our storage systems leverage object storage. Both of these technologies allow for auto-scaling.

Compute and storage decoupled (truly)

Since observability involves the two dimensions of data: volume and retention, it is highly imperative for any modern architecture to completely decouple storage from compute. This decoupling allows for scenarios where data generation rates remain steady, but retention requirements increase. While claiming to employ decoupling, most systems like Splunk, Datadog, QRadar, ArcSight, etc., truly don't. The (expensive) storage needed for indexing and processing is still tightly coupled to their compute resources and only the storage used for long-term retention is decoupled. This leads to complexity, lack of flexibility, and increased cost.

However, Apica has true decoupling where indexes and retention are both decoupled from compute. The ingest capacity is 100% decoupled from storage capacity!

In the following sections, you will see how Apica's architecture enables compute and storage expansion using simple API calls. Yes, it is that simple!

Infinite compute

Using Kubernetes containers as our compute layer offers the inherent benefits of autoscale. Apica software administrators can dynamically/automatically increase or decrease the number of running pods as data rates change. Increases in data load, either in short bursts or gradually over time, are automatically handled; they are no longer cumbersome and time-consuming. Kubernetes Autoscaler addresses situations where the system needs to accommodate surging data rates using a scaling out of pods and/or nodes.

Using this scale-out architecture, the Apica system can ingest any scale seamlessly. The architecture to support GBs, TBs, PBs of data ingestion per second is essentially the same.

With the addition of intelligent algorithms using AI/ML techniques that take care of capacity planning, infinite scale capability is a fit-it-and-forget-it scenario with Apica.



Infinite storage

Apica uses any object store or S3-compatible object store as its primary storage layer. You read that right! This capability of using object stores as primary storage means that our system can handle any growth in data volumes (even theoretically infinite) due to ingestion or long-term retention requirements using simple API calls. The operational agility that comes with Apica using object stores as primary storage is unparalleled by any other system in the industry today. Any volume of data retention, be it TBs or PBs of data, works precisely the same way with the Apica architecture.

Apica is the first real-time platform to bring together benefits of object store like scalability, one-hop lookup, faster retrieval, ease of use, identity management, lifecycle policies, data archival, and other capabilities.

Most scaled-out self-service log analytics solutions require costly management of volumes at scale! Apica abstracts it as an S3 API.

Lowest TCO

Apart from the costs of licenses and management, infrastructure costs significantly influence the TCO of any observability system. Let's talk about the infrastructure costs involved in running Apica, namely, compute and storage systems. Containers cost remarkably less than virtual machines and bare-metal servers. Kubernetes deployments effectively fund themselves - not just in direct dollar terms but also in terms of simple management, efficient resource utilization, high availability, fault-tolerance, and end-user productivity.

Object storage is our primary storage. This has similar architectural benefits as a Kubernetes system provides for compute. Simple management, efficient resource management, high availability, fault-tolerance, and enhanced end-user productivity. This directly translates to a Zero Storage tax model for the end-user. We not only simplify Day 0 operations but eliminate any and all storage overheads with Day 1 and beyond operations. Thus, Apica offers the lowest possible TCO on the planet with respect to scaling storage operations with growing data needs.

Real-time performance

A distributed Kubernetes-based compute architecture enables our software to keep performing at the best performance-to-scale ratio and continuously ingest and process high-volume data streams in real-time. We previously mentioned that we use S3-compatible storage systems as our primary storage layer. You might be wondering, "How do you manage real-time query performance on what everyone widely considers and use as secondary or cold storage?" Thanks to our innovative engineering, that's a reality. While other use-cases are possible down the line, this real-time performance extraction from object storage serves the observability use-case perfectly.

Let's summarize.

Apica is an INFINITE SCALE OBSERVABILITY SYSTEM offering REAL-TIME PERFORMANCE at the LOWEST TCO. If you'd like to know more about how Apica can accelerate your business transformation through unmatched data pipeline control and total observability, do reach out.

