

# LAKE

## Compliance & Search

### INTRODUCTION

Apica's LAKE is a novel storage paradigm built to meet the demands of modern high-volume business environments. The platform, InstaStore, runs on top of any object storage and solves traditional storage operation challenges such as running out of disk space, node failures, and storage recovery.

LAKE is a one-tier storage infrastructure that fully indexes incoming data and allows uniform, on-demand, and real-time access to all data. This eliminates the need for tiered storage and allows teams to focus on their core expertise and business problems.

LAKE is compliant with petabyte-scale indexing and instant retrieval, unifying, storing, and securing all machine data centrally on LAKE. The platform also enables low-cost longer-term retention with superfast querying, real-time access, and extensive auditing and reporting capabilities that enable companies to maintain and demonstrate compliance. With centralized log management, proactive analytics, and monitoring, granular reporting, and audit trails, and built-in security controls, LAKE features and the Apica platform help maintain continuous compliance at any scale.

### CAPABILITIES

Apica's LAKE capabilities include:

**Works on any object store:** Turn your favorite object storage system, on-premises or cloud, into your primary and only storage layer for observability data. Get hot storage like speed on object storage to easily access and ship your data across public or private clouds to build applications, perform analytics, and speed time to insight

**Gain instant access to data at rest:** LAKE indexes 100% of all incoming observability data from all your data sources. Data at rest uses Apache Parquet format that ensures data availability at any scale, simplified retrieval, and faster time to insight.

On LAKE, repository and processing come together to enable data access that's 2x faster and columnar data retrieval.

**Retain master data from all your data sources:** Ingest and store every byte of observability data on InstaStore at object storage prices. LAKE enables uniform access to unfiltered master data and data optimized for your target systems in real time.

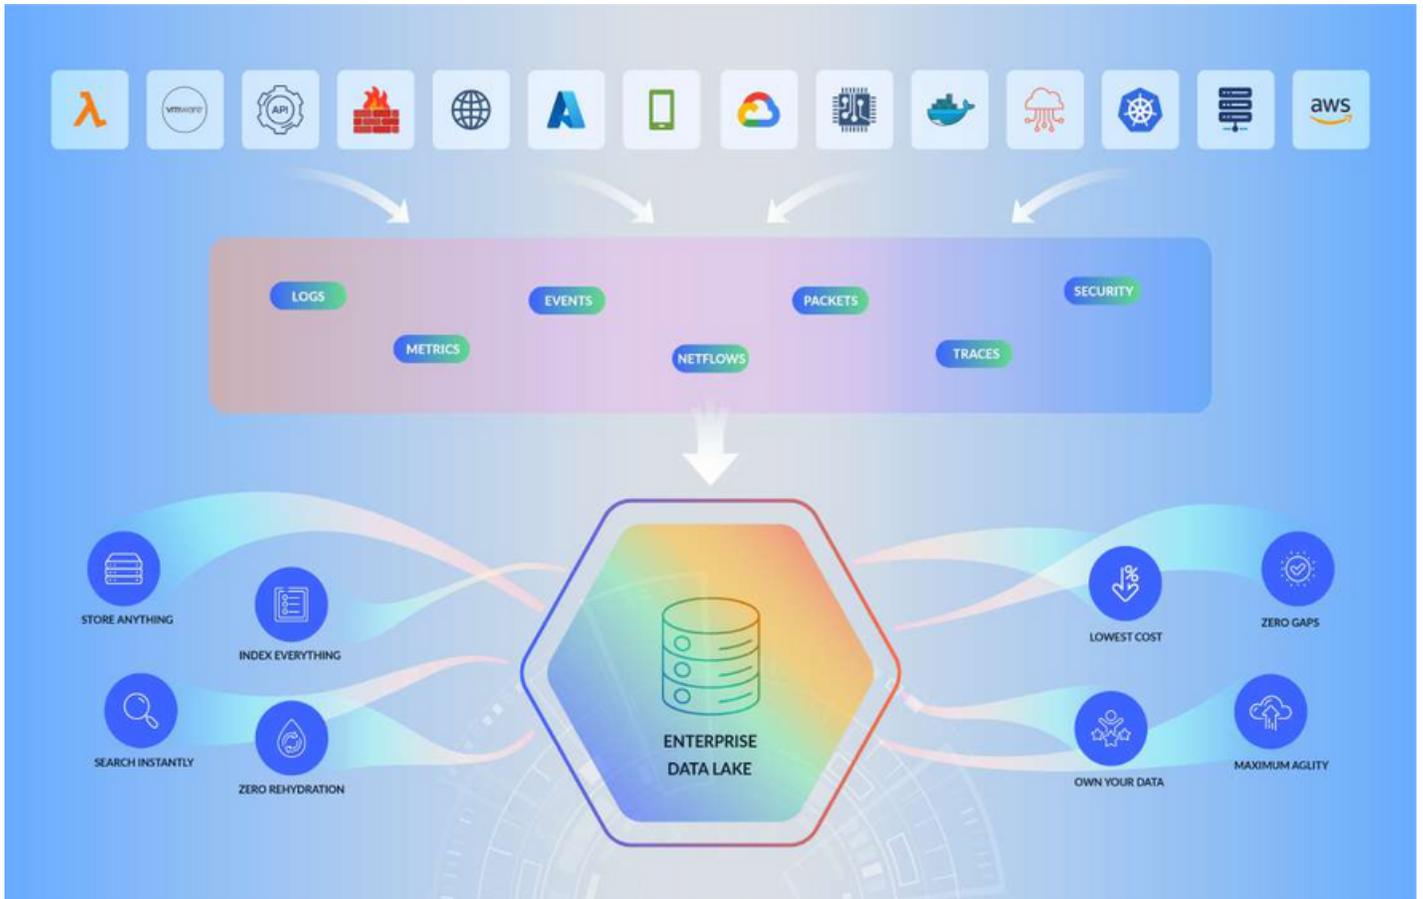
### HIGHLIGHTS

- Apica's LAKE is a novel storage paradigm designed to meet the demands of modern high-volume business environments.
- InstaStore, which runs on top of any object storage, solves traditional storage operation challenges such as running out of disk space and node failures.
- LAKE enables petabyte-scale indexing and instant retrieval, unifying, storing, and securing all machine data centrally on one platform.
- It allows for low-cost longer term retention with superfast querying capabilities in real time along with extensive auditing reporting features that enable companies to maintain compliance at scale.
- With centralized log management analytics monitoring granular reporting audit trails built in security controls InstaStore helps maintain continuous compliance at any scale while providing access to unfiltered master data optimized for target systems in real time via forwarders making it available downstream applications for additional context.

Use built-in forwarders to make master data available in downstream applications for additional context.

**Never lose a single byte:** The elastic design of LAKE ensures high storage availability in the most voluminous data environments. LAKE enables your data pipelines to scale horizontally to handle any unexpected data spikes at endpoints, avoid data backlogs, and prevent data loss at scale.

With LAKE, you'll never have leaks in your observability data pipeline.



### Search and replay historical data instantly:

Since LAKE indexes all your machine data, you can query, retrieve, and replay historical data from any timeframe in an instant.

Whether you're retrieving one log line from several billion or a batch of logs from a year ago, LAKE gets them ready to be replayed to any target system of choice in real time.

### Limitless storage, minus the overhead:

Store ALL of your log data in any object storage as primary data with real-time storage capabilities using InstaStore. With InstaStore, you can get rid of tiering, the complexities it introduces, and the storage operations overheads forever.

Your log data on InstaStore is fully indexed and searchable, mineable, comparable, and replayable to any target system in real-time.

### Unlock better compliance at scale

Unify, store, and secure all your machine data centrally on LAKE.

With object storage at its core, LAKE enables low-cost longer term retention with superfast querying, real-time access, and extensive auditing and reporting capabilities that enable your company maintain and demonstrate compliance.

With centralized log management, proactive analytics and monitoring, granular reporting and audit trails, and built-in security controls, LAKE and the Apica platform help you maintain continuous compliance at any scale.

## Compliance

Petabyte-scale indexing and instant retrieval

Unify, store, and secure all your machine data centrally on LAKE. With object storage at its core, LAKE enables low-cost longer-term retention with superfast querying, real-time access, and extensive auditing and reporting capabilities that enable your company to maintain and demonstrate compliance.

## BENEFITS

The Key Benefits of LAKE include:

**SIMPLICITY:** Stores 100% of all data in object storage as primary data - no secondary tier

**QUERY READY:** 100% Indexed, searchable, comparable & minable in real-time!

**SPEED:** Old Data or New Data - Same Time to First Byte (TTFB)  
**TIME TRAVEL:** It is possible to go back in time! Go back in time and replay any old data to any target, anytime

**DATA OWNERSHIP:** 100% of data belongs to you, not Apica

**EFFICIENCY:** No storage management projects.

**OPENNESS:** Open data format eliminating vendor lock-in and guaranteeing future-readiness

**FLEXIBILITY:** Works with environments of any size - small to large

**SECURITY:** Built-in security & encryption, protecting data from unauthorized access or modification.

**SCALABILITY:** Ability to scale and add more storage as the need increases.

**RELIABILITY:** Uptime guarantees of 99.999%, ensuring no downtime and consistent performance.

**COST EFFECTIVENESS:** Low cost of ownership, eliminating the need for costly maintenance and upgrades.

**DATA VISUALIZATION:** Easily visualize data using pre-built tools or custom-tailored solutions.

With centralized log management, proactive analytics and monitoring, granular reporting and audit trails, and built-in security controls, LAKE and the Apica platform help you maintain continuous compliance at any scale.

### Search

Instantly search and visualize at petabyte-scale

All of your log data may be stored in any object storage as primary data with real-time capabilities using LAKE. You can get rid of tiering, the issues it causes, and the storage operations overheads for good with LAKE. Your log data on LAKE is fully indexed, searchable, mineable (in some cases), comparable, and replayable to a target system in real-time.

Ingest and store every byte of observability data on LAKE at object storage prices. LAKE enables uniform access to unfiltered master data and data optimized for your target systems in real-time.

Use built-in forwarders to make master data available in downstream applications for additional context.

### Replay

Instantly replay historical data to any target

LAKE indexes all of your machine data and gives you access to logs from any time period you need. With our advanced search and filtering capabilities, you can quickly identify the exact data point or series of points you're looking for. And since LAKE intelligently compiles all the necessary information to recreate your exact environment, it's easier than ever for you to troubleshoot problems and ensure quality control across all your systems.

You can query, retrieve, and replay historical data from any time period in an instant since LAKE indexes all of your machine data. Whether you're looking for one log line or a whole batch of logs dating back to a year ago, LAKE gets them ready for real-time playback on any target system of choice.

Use our powerful search and filtering tools to find exactly the data point or series of points you're looking for. LAKE intelligently compiles all the data necessary to recreate your exact environment, so you can quickly identify issues, validate fixes, and ensure quality control.

---

## SUMMARY

With LAKE, you can instantly search and visualize at a petabyte scale, getting rid of tiering, the issues it causes, and the storage operations overheads for good. Your log data on LAKE is fully indexed, search-able, mineable, comparable, and replayable to a target system in real-time.

Furthermore, LAKE capabilities allow you to instantly replay historical data to any target, enabling the query, retrieve and replay of historical data from any time period in an instant since LAKE indexes all of your machine data.

Whether you're looking for one log line or a whole batch of logs dating back to a year ago, LAKE gets them ready for real-time playback on any target system of choice.