

LAKE

Store Everything. At the Lowest Cost. On Any Object Storage

Introduction

Apica empowers you to effortlessly manage vast volumes of data through its unique data lake solution, LAKE. It seamlessly integrates with any object storage, offering comprehensive indexing and real-time access to all data. With limitless scalability and streamlined operations, it eliminates the need for additional log storage projects.

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.

Moreover, LAKE isn't only about lessening the storage tax; we're redefining data management. To achieve this, LAKE takes a unique approach. It treats any object store as the primary storage layer, enhancing its performance to match real-time storage capabilities. By decoupling storage and compute processes, LAKE simplifies the management of high-volume data requirements, making it easy to install, use, and scale within an organization.

Furthermore, in the data storage landscape, enterprises find it hard to escape the latent storage tax. Our objective is straightforward: we're here to liberate observability platforms from storage hassles. With LAKE's InstaStore, you get ZeroStorageTax from day one and beyond.

The storage tax refers to the hidden costs and operational overhead that organizations incur when managing their data storage infrastructure, which can significantly impact their overall efficiency and expenses.

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

HIGHLIGHTS

- LAKE is Apica's data lake solution for managing vast data volumes.
- Integrates seamlessly with any object storage for real-time data access.
- Offers limitless scalability, eliminating the need for additional log storage projects.
- Provides one-tier storage infrastructure, eliminating tiered storage complexities.
- Ensures compliance with petabyte-scale indexing and instant retrieval.
- Addresses data challenges related to volume, cost, integration, quality, scalability, security, and metadata.
- Key capabilities include centralizing data, historical data access, master data retention, data loss prevention, real-time storage, and compliance enhancement.
- LAKE empowers data management by eliminating storage tax and offering real-time performance.
- It's designed to simplify high-volume data requirements and optimize costs.

Addressing Data Challenges

In the modern landscape of distributed systems, dealing with a multitude of data challenges has become a critical aspect of maintaining system efficiency and reliability. These challenges encompass a range of issues related to data volume, quality, security, scalability, and more.

Handling Vast Data from Complex Systems Efficiently:

The proliferation of distributed systems has led to an exponential increase in data volumes generated from various sources. Managing and making sense of this vast amount of data can be overwhelming. LAKE's InstaStore technology offers a promising solution by streamlining the indexing and storage process for data from complex sources. This technology enables quick data organization and storage, significantly reducing the effort required to manage data streams.

Costly Log Storage Solutions:

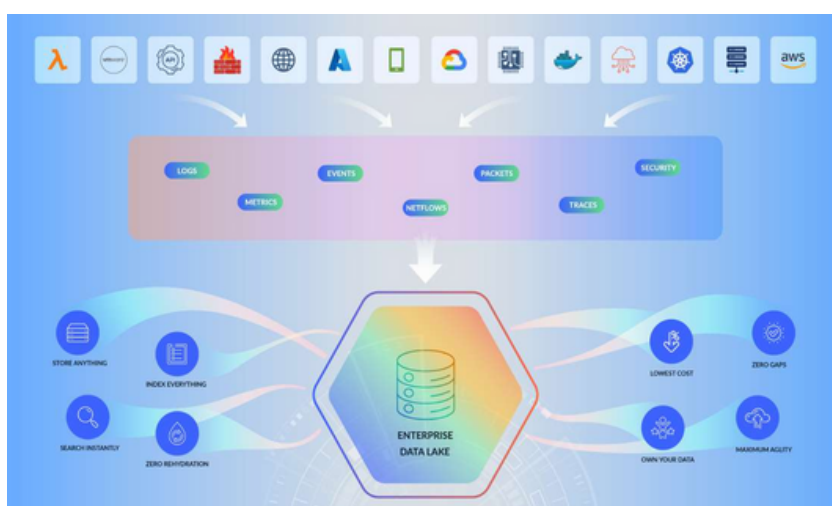
Traditional log storage solutions often lead to additional resource-intensive projects that strain both budgets and operational efficiency. LAKE addresses this concern by providing an efficient and user-friendly alternative. By eliminating the need for separate log storage projects, LAKE not only reduces costs but also simplifies the process of managing log data.

Scalability and Performance:

Scalability remains a concern as data demands continue to grow. The ability to scale systems while maintaining performance is crucial for avoiding bottlenecks and ensuring seamless operations. LAKE's technology offers features that enable efficient scaling, allowing organizations to accommodate increasing data loads without sacrificing performance.

Data Quality and Consistency:

Maintaining data quality and consistency across complex systems is a persistent challenge. Discrepancies in data formats, data corruption, and inconsistencies can negatively impact decision-making and analysis. LAKE includes mechanisms to validate and ensure data quality, contributing to more accurate insights and improved decision-making processes.



Data Ingestion and Integration:

Collecting and integrating data from diverse sources can be a complex endeavor. Inconsistent data formats, varying data speeds, and the need to accommodate different data sources can lead to integration challenges. LAKE facilitates efficient data ingestion and integration, enabling organizations to streamline this process and make the most of data resources.

Data Security and Compliance:

In distributed environments, ensuring data security and compliance with regulations is intricate. The complex nature of data flows and storage requires robust security measures and adherence to various compliance standards. LAKE provides functionalities that address these concerns, helping organizations navigate the complexities of data security and compliance within distributed systems.

Metadata Management:

Efficient metadata management is key to enhancing data discoverability and governance. Without proper metadata organization, finding relevant data becomes challenging, and maintaining data governance standards becomes more difficult. This includes features to efficiently manage metadata, enabling users to locate and utilize data resources effectively.

Cost-Effective Infrastructure:

Balancing the costs of infrastructure with the need to manage and analyze large datasets is a common concern. LAKE's solution optimizes infrastructure usage, allowing organizations to manage extensive data while minimizing operational costs. This can have a significant impact on the overall budget and resource allocation.



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.

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.

Retain master data from all your data sources:

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

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 your log data in any object storage as primary data with real-time storage capabilities using InstaStore. Thus, 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 to 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.

Key Features

- **Centralized Visibility:** Gain instant access to data at rest, effortlessly collecting logs, metrics, traces, and packets from all systems.
- **Time Travel Capabilities:** Query, retrieve, and replay historical data from any time-frame instantly.
- **Master Data Retention:** Ingest and store comprehensive observability data with uniform access and context.

- **Data Loss Prevention:** Elastic design prevents data loss, scaling to handle unexpected data spikes.
- **Real-Time Storage Capabilities:** Store log data in object storage with real-time access, eliminating tiering complexities.
- **Compliance Enhancement:** Centrally store and secure machine data, ensuring continuous compliance.

Benefits

- **Centralized Data at Scale:** LAKE consolidates vast data streams into a centralized repository, simplifying management.
- **Improved Data Availability:** Real-time access to data enhances operational efficiency and responsiveness.
- **Enhanced Analytics and Insights:** Easy data access leads to deeper insights and streamlined analytics.
- **Faster Compliance:** Simplified compliance through centralized log management and reporting.
- **Support for Diverse Data Types:** LAKE accommodates various data formats, ensuring compatibility.
- **Cost-Effective Storage:** Benefit from affordable storage solutions without compromising performance.

Empower Your Data Management

Just like there is Shadow IT, there is a hidden storage tax that is paid by every enterprise. Our mission is to free observability platforms from storage operations overheads. ZeroStorageTax from day 1 and beyond use is now possible with LAKE.

Apica's approach to solving this fundamental problem was to make any object store a primary storage layer and make it perform like real-time storage. We cleanly decoupled Storage and Compute and made high-volume data requirements simple to install, use, and scale.

Apica's LAKE redefines data management, offering limitless storage with efficient indexing and real-time access. With LAKE's advanced capabilities powered by InstaStore technology, you can streamline operations, enhance analytics, and ensure compliance, all while optimizing costs.

Experience the Power of LAKE today!

Let's Talk.
