

Cancun MemoryLake



Data Platform for Fast Insights

Business Benefits

Fast Time to Insights

- Accelerate time to insight by up to 10X.

Business Agility

- Elastically scale infrastructure.
- Seamlessly consume data from various sources (on-prem or cloud).

Infrastructure Efficiency

- Reduce hardware footprint by one-quarter or query four times more data on existing infrastructure.

High ROI

- Reduce costs of new deployments and increase ROI up to 500% on existing infrastructure.

Frictionless Deployment

- No change required to existing applications or infrastructure.

Use Case Examples

Financial Institutions

- Accelerate fraud detection and risk analysis from hours to minutes, while reducing infrastructure spend.

Retail/E-commerce

- Increase profits by improving customer sentiment analysis, 360 degree view of customers, and targeted recommendation engines.

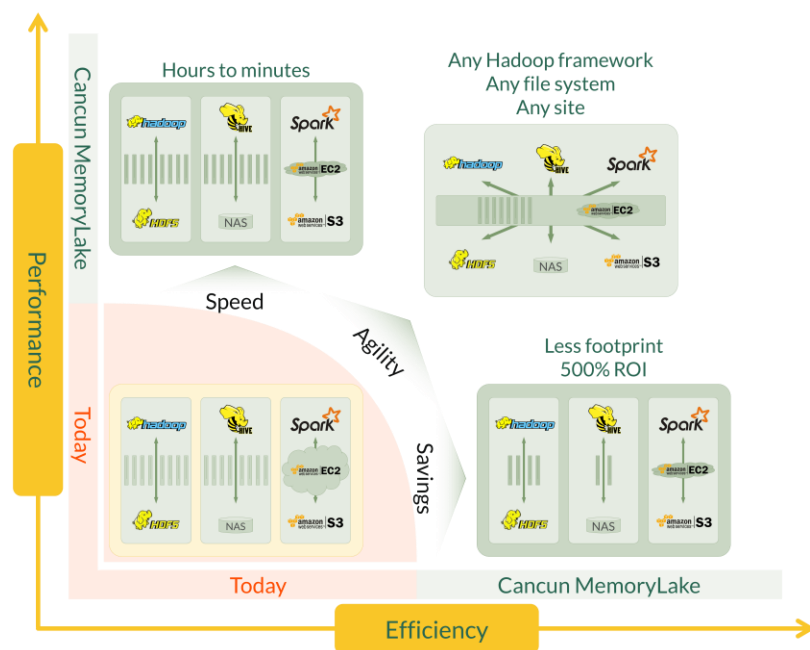
Government

- Accelerate insights for mission critical projects (national security, cyber defense, health services, etc.).

The Problem: Data Growth Leading to Infrastructure Sprawl

Data is growing faster than our ability to consume it, overwhelming existing Big Data infrastructures, driving up costs and slowing the pace of business. In order to remain competitive, businesses must be able to quickly extract insights from an avalanche of data. Yet today, most Big Data clusters are slow and inefficient. As data volume grows, so must infrastructure at a similar or faster pace. In addition, clusters are being set up in silos. These factors ultimately lead to poor application performance, inefficient architecture, and expensive scaling requirements.

Cancun MemoryLake: An In-Memory, Transparent, Data Platform



Cancun MemoryLake™ software platform enables applications to run up to 10X faster, allowing customers to get richer insights and save on infrastructure expenses leading up to 500% ROI. Cancun MemoryLake provides immediate benefits in three areas:

Faster time to Insights: By virtualizing and pooling available memory and storage resources within and across nodes, Cancun is able to present a much larger memory footprint to applications. In-memory applications like Spark can now run significantly faster by leveraging the larger and more-efficient memory enabled by Cancun MemoryLake.

Business Agility: Cancun enables businesses to run analytics applications in private, public, or hybrid cloud environments, and ingest data directly from various sources (e.g. HDFS, NAS, cloud object stores) for richer insights. Cancun MemoryLake technology makes it possible to separate compute from memory, allowing customers to scale these resources independently as needed; servers

Key Features

Performance

- Cluster-optimized memory virtualization
- Dynamic memory offload
- In-memory workflow pipelining

Agility

- Unified data access
- Scale as you grow, with decoupled compute, memory and storage
- API framework for 3rd party integration

Ease-of-use

- No change to applications
- Controller-based cluster management
- Hardware agnostic, support for on-prem or cloud deployments
- 1-click rollback

Product Information

Platform

- Any x86-based system, bare metal, VM, container, or cloud

Java

- Open JDK 1.7 and above

OS

- CentOS 6.6+, RHEL 6.6 +,
- *Ubuntu 14.04+*

Computation framework

- Apache Hadoop 2.7+, Apache MapReduce v1, v2, Apache Spark 1.6.1+, Hive1.2.1
- *Pig, HBase, Impala, Presto, Cassandra, Mongo DB, Splunk*

Storage Framework

- HDFS, Amazon S3, MapR-FS, EMR-FS, LocalFS, NFS, GCS
- *Azure, Cassandra*

Public Cloud support

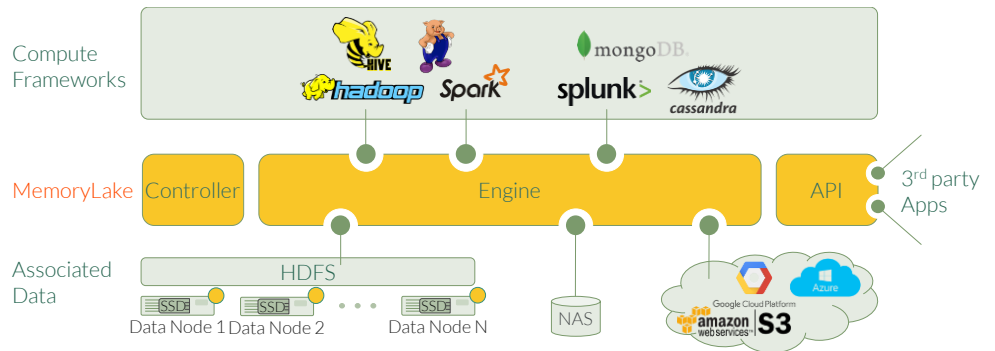
- Amazon EC2, Amazon EMR,
- *Google GCE, Microsoft Azure*

Roadmap items are italicized. Contact us for details.

with dense SSD/NVME configurations can be transparently pooled, increasing the performance of the entire cluster. And deployment is frictionless, requiring no changes to application code or underlying infrastructures.

Infrastructure Savings and ROI: In existing environments, customers using Cancun MemoryLake technology can run more jobs and query more data without having to purchase additional hardware. For new buildouts, customers need only deploy a fraction of the expected infrastructure since Cancun delivers unprecedented infrastructure efficiency. And in cloud deployments, Cancun MemoryLake allows customers to enjoy immediate savings by reducing the required number of compute instances.

Architected for Performance, Agility, and Efficiency



The Cancun MemoryLake Platform includes three elements:

MemoryLake Engine: The MemoryLake Engine is a transparent layer residing in each data node of the cluster offering virtually unlimited memory to northbound compute frameworks. Controller and APIs run in a controller node, either in a container or on a bare metal server.

MemoryLake Controller: The MemoryLake Controller provides a single pane for provisioning, monitoring, and maintenance operations. It includes functions for installation, rollback, upgrade, configuration, service scheduling, and metrics capture.

MemoryLake APIs: With a fully RESTful API, Cancun MemoryLake can be integrated into existing Big Data environments and allow third-party applications to integrate into the MemoryLake Platform to create higher value offerings.

Cancun MemoryLake: A Faster, More Agile & Efficient Infrastructure

Without any changes to existing applications or infrastructure, Cancun customers can immediately benefit from greater performance, increased agility and enhanced infrastructure efficiency resulting in significantly lower costs and a distinct competitive advantage. Visit us online to learn more.

