

Software-Defined MemoryLake™

Software Platform for Fast, Efficient Insights



Business Benefits

Fast Time to Insights

- Accelerate time to insight by up to 10X

Infrastructure Efficiency

- Reduce cloud expenses by completing jobs and decommissioning clusters faster
- Reduce infrastructure and cost of new build-outs.

Run more jobs with current infrastructure

Simple, Frictionless Deployment

- Hassle-free cloud simplicity – provisioned in minutes
- 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

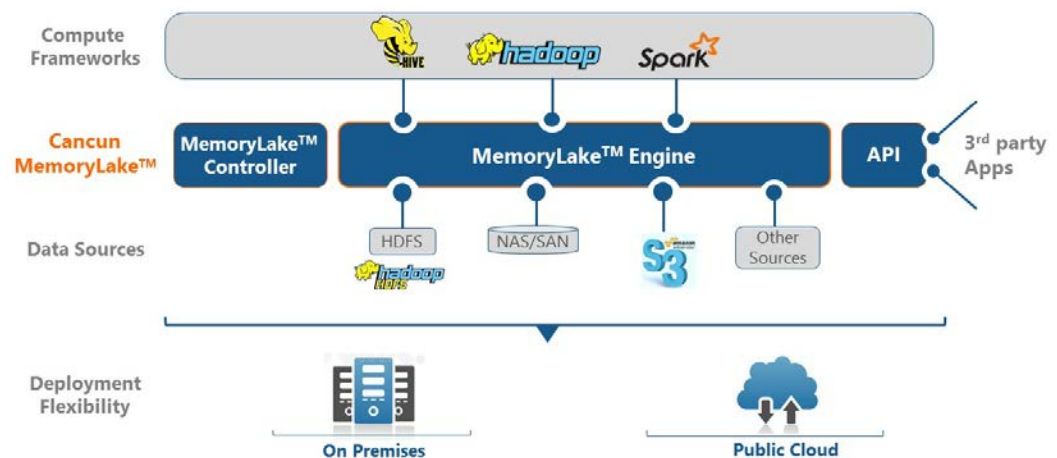
- Accelerate insights for mission critical projects (national security, cyber defense, health services, etc.).

The Problem: Data growth and slow data access is impeding time to insights and leading to infrastructure sprawl

The exploding growth of data combined with the limitations of today's infrastructures are outpacing data-centric businesses' ability to effectively gain meaningful insights that help them operationalize their businesses and drive competitive advantages. Most big data workloads spend the majority of time performing data access to and from slow media, and less time in computation. Slow data IO cripples the analytics process and dramatically increases time to insight - and processing growing data sets with the same infrastructure just exacerbates these outcomes. To counter these challenges, corporations are deploying larger infrastructures to parallelize processing, which leads to infrastructure sprawl and escalating costs.

Software-Defined Memory Lake: A faster and more efficient solution

Consider a solution that allows applications to access data from disk at memory speed, one that makes intelligent usage of available resources across memory and storage and delivers the speed of memory at the cost efficiency of disk. To accomplish this, a software-defined memory lake (SDML) approach is needed.



Cancun's MemoryLake™ Platform includes three elements:

MemoryLake™ Engine: The MemoryLake™ Engine is a transparent software layer residing in each data node of the cluster, and delivers memory-speed access to data. The MemoryLake™ Engine can leverage memory and storage resources within or across nodes in the cluster.

MemoryLake™ Controller: The MemoryLake™ Controller provides a cluster-wide management console for provisioning, monitoring, and maintenance operations. It includes functions for installation, rollback, upgrade, configuration, service scheduling, and metrics capture. The MemoryLake™ Controller can run either in a container or on a bare metal server, and can be configured for redundancy.

Key Features

Performance

- Memory-speed data IO
- In-memory workflow pipelining
- Off-node IO scaling

Efficiency

- Optimal memory and storage tiering across RAM/SSD/NVMe
- Unique in-memory IO scale options, independent from compute
- Unified data access

Simplicity

- Deployment flexibility: cloud and on-premises
- No change to applications
- Provisioned in minutes
- Controller-based cluster management
- 1-click rollback
- API framework for 3rd party integration

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, Kafka*

Storage Framework

- HDFS, Amazon S3, MapR-FS, EMR-FS, LocalFS, NFS, GCS

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

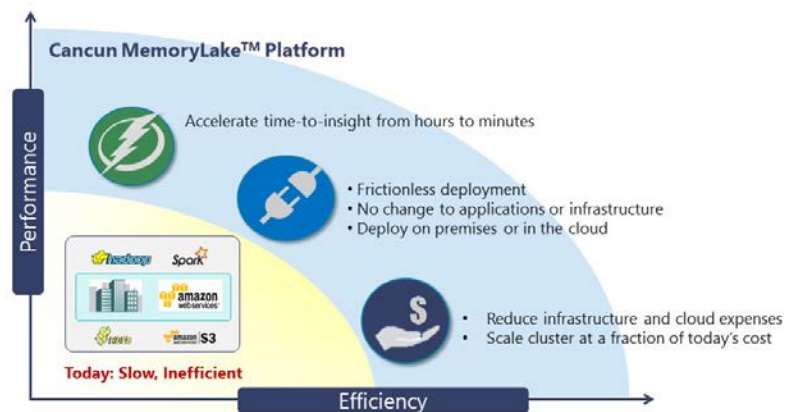
Cancun MemoryLake™: In-Memory Software Platform for Accelerated Insights

Cancun's MemoryLake™ delivers an SDML that enables applications to run up to 10X faster, accelerating time to insights and tremendous infrastructure efficiencies. Cancun's MemoryLake™ provides immediate benefits in three areas:

Faster Time to Insights: By pooling and virtualizing available memory and storage resources within or across nodes, Cancun is able to create a software-defined memory lake. In-memory applications like Spark can now run significantly faster by accelerating and pipelining applications at memory speed, enabling workflows to complete in a fraction of the time.

Infrastructure Efficiency and Savings: Whether deployed on-premises or in the cloud, Cancun's MemoryLake™ software delivers unprecedented infrastructure efficiency. Existing build-outs can run more jobs and query more data without additional infrastructure purchases. New build-outs only require a fraction of the expected infrastructure. For cloud deployments, customers can experience both faster insights and immediate savings because they are able to complete jobs and decommission clusters much faster.

Deployment Simplicity and Flexibility: Cancun enables businesses to deploy MemoryLake™ software in private, public, or hybrid cloud environments, and ingest data directly from various sources (e.g. HDFS, NAS, cloud object stores) for richer insights. Installation is simple and takes only minutes. And deployment is frictionless, requiring no changes to application code or underlying infrastructures.



Cancun's MemoryLake™: Delivering Faster and More Efficient Insights

Without any changes to existing applications or infrastructure, Cancun customers can gain a distinct competitive advantage and immediately benefit from accelerated time-to-insights at significantly lower costs, and with the simplicity and flexibility to deploy either on-premises or in the cloud.

Visit us online to learn more.



©2017 Cancun Systems, Inc.
www.CancunSystems.net