

Radically Improve ROI on Big Data Infrastructures



INCREASED EFFICIENCY CUTS INFRASTRUCTURE SPEND

"Cancun demonstrated the ability to do the same job with significantly less infrastructure, which can lead to a lot of efficiency and cost savings for us and our customers."

— David Vennergrund
Director, Data Science CSRA

Executive Summary

Each year, the amount of data that businesses need to manage grows at an astounding rate¹. The explosive growth of data presents a costly challenge to businesses who must continue to scale their big data infrastructures to keep pace with the speed of new data creation. Whether the data resides on-premises or in the cloud, the rising costs to maintain a robust infrastructure threaten to override the value derived from big data insights.

Cancun Systems offers customers an in-memory SDML (software-defined memory lake) platform that not only delivers massive speed and agility, but also allows companies to achieve new levels of infrastructure efficiency and cost savings, whether deployed on-premises or in the cloud.

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

Cancun's MemoryLake™ software platform delivers an SDML that enables applications to run up to 10X faster, allowing customers to accelerate time to insights and enjoy 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 can 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 having to purchase additional infrastructure. 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.

¹ [How Much Data is Produced Every Day?](#) Northeastern University, 5/13/16

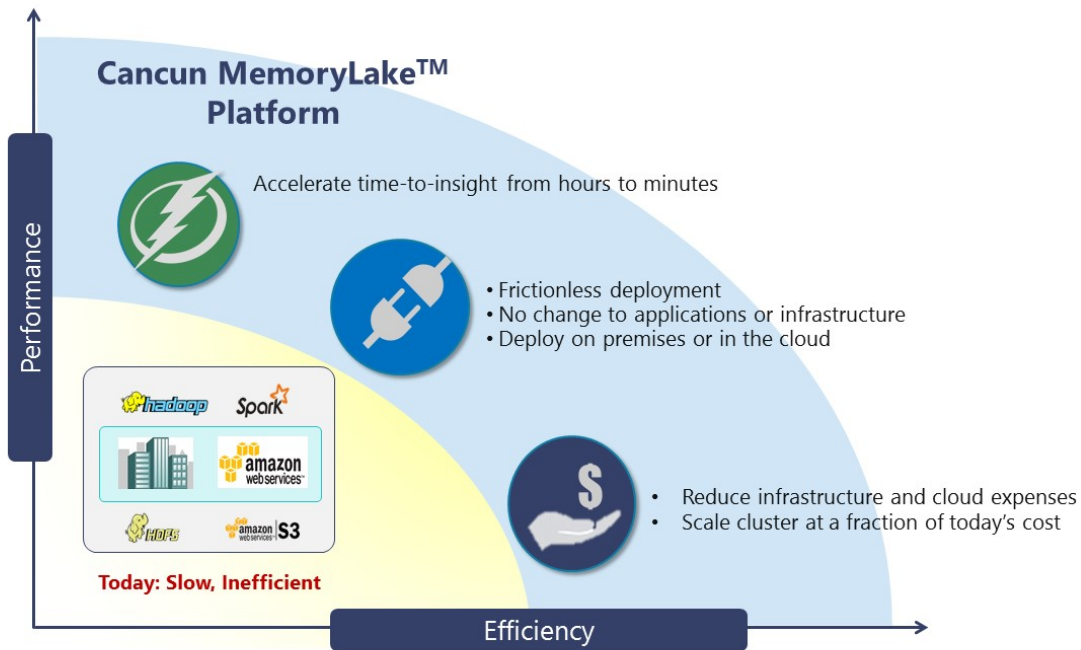


Figure 1 – Cancun MemoryLake delivers massive speed, agility, and cost savings to existing Big Data frameworks

On-Premise Deployments: Cancun MemoryLake™ Provides Nearly 500% ROI

Based on a conservative estimate of 2X data growth per year, companies that deploy the Cancun’s MemoryLake™ platform can achieve a 42 percent cost savings, resulting in an immediate 494 percent ROI. See Figure 2 – On-Prem Infrastructure TCO Based on 2X/year Data Growth.

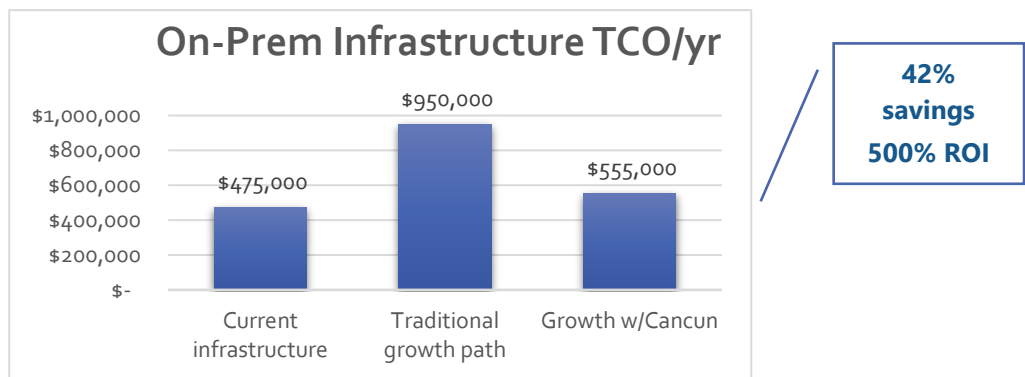


Figure 2 – On-Prem Infrastructure TCO Based on 2X/year Data Growth



The TCO calculations are based on a 50TB Spark job growing 2X to 100TB. It's important to note that the ROI savings derived from Cancun's MemoryLake™ are realized on day one, essentially allowing companies to still save 30-40 percent on infrastructure gains even after deploying Cancun.

Customers using Cancun's MemoryLake™ in new buildouts will realize immediate cost savings, since Cancun significantly reduces the amount of infrastructure required. Cancun is equally beneficial for use in existing environments, since customers need only deploy a fraction of the expected infrastructure since Cancun delivers unprecedented infrastructure efficiency. In addition, both new and existing environments benefit from the immediate boost in performance and ability to seamlessly ingest data from various sources (HDFS, NAS, cloud object stores, etc.) for richer insights.

Calculations are based on standard list pricing² as seen in Figure 3 – ROI On-Prem Calculations Based on 50TB Spark Job Growing 2X.

On-Prem: 50TB Spark job growing 2X to 100TB	Current infrastructure	Traditional growth path	Growth w/Cancun
1. # Nodes	25	50	25
2. Server cost/node/yr	\$3,333	\$3,333	\$3,333
3. OS cost/node	\$500	\$500	\$500
4. Hadoop license per node	\$2,500	\$2,500	\$2,500
5. OpEx/yr/node	\$12,666.67	\$12,666.67	\$12,666.67
6. Cancun ASP/node			\$3,200
Infrastructure TCO/yr	\$475,000	\$950,000	\$555,000
TCO Savings with Cancun			\$395,000
TCO Savings with Cancun (%)			42%
ROI from day1			500%

Figure 3– ROI On-Prem Calculations Based on 50TB Spark Job Growing 2X

² Pricing as of January 2017



Cloud Deployments: Cancun MemoryLake™ Provides 200% ROI

Deploying the Cancun's MemoryLake™ platform in the cloud yields similar benefits with 30 percent cost savings and an immediate 203 percent ROI. See Figure 4 – Cloud Infrastructure TCO Based on 2X/year Data Growth and Figure 5 – ROI Cloud Calculations Based on 50TB Spark Job Growing 2X.




Figure 4 – Cloud Infrastructure TCO Based on 2X/year Data Growth

Cloud: 50TB Spark job growing 2X to 100TB	Current infrastructure	Traditional growth path	Growth w/Cancun
1. Nbr of Nodes	25	50	25
2. Server cost	\$4,200	\$4,200	\$4,200
3. Storage cost	\$720	\$720	\$1,440
4. OS cost	\$84	\$84	\$84
5. Hadoop license	\$2,500	\$2,500	\$2,500
6. OpEx/yr/node	\$11,256	\$11,256	\$12,336
7. Cancun ASP			\$5,606
Infrastructure TCO/yr	\$469,000	\$938,000	\$654,150
TCO Savings with Cancun			\$283,850
TCO Savings with Cancun (%)			30%
ROI from day1			200%

Figure 5 – ROI Cloud Calculations Based on 50TB Spark Job Growing 2X

Conclusion

The Cancun's MemoryLake™ platform offers a unique solution for customers struggling to keep pace with big data demands. With Cancun's MemoryLake™, customers not only benefit from a massive boost in speed and agility, they are also able to add new initiatives and workloads to existing clusters for greater efficiency, or reduce their infrastructure footprint, on-premises or in the cloud, for immediate cost savings and significant ROI.

For more information or to request a demo, visit us at www.cancunsystems.net

Notes:

Figure 2 – On Prem ROI

- 2) Server cost: Xeon, 2 socket, 4 core/socket, 10G net, 64G ram, 12TB HDD, 500GB NVMe-SSD. \$10K on 3 yr. depreciation
- 5) Opex/yr/node=2X capex, (Professional Services is 1X plus operating expense of datacenter plus manpower)

Figure 3 – Cloud ROI

- 2) Server cost: 1xM4.2XL instance, 8vCPU, 32G Mem - \$4200/instance/yr.
- 3) Storage cost (EBS \$120/TB/yr.)
- 4) OS cost (\$84/yr.)
- 5) Hadoop license (\$2500/yr)
- 6) OpEx/yr/node = 1.5X capex, (Professional Services is 1X plus manpower)

