Increasing Velocity in Financial Services with the Power of Dell Technologies

Stack Summit – Spring 2022

Joe Steiner
CTO

DELLTechnologies

Maximizing the value of data in your strategy

The data value equation

Downward pressures

Lacking employee skillsets

Difficult workflows

Data silos

Technical debt

Cyber threats

Vendor lock-in

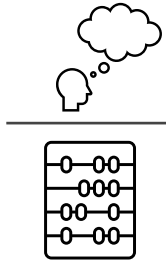Resources & budget

$$\text{Data value} = \frac{\text{VELOCITY}}{\text{FRICTION}} \times \text{USE CASES}$$

DELL Technologies

Dell Technologies Solution

From archives to ultra low latency

P or V Applications

Buffer Cache

SUT Host
(also the Client Host)

FC

Brocade 300 SAN Switch

NetApp EF550 Flash Array
(firmware ver 08.10.05.00)
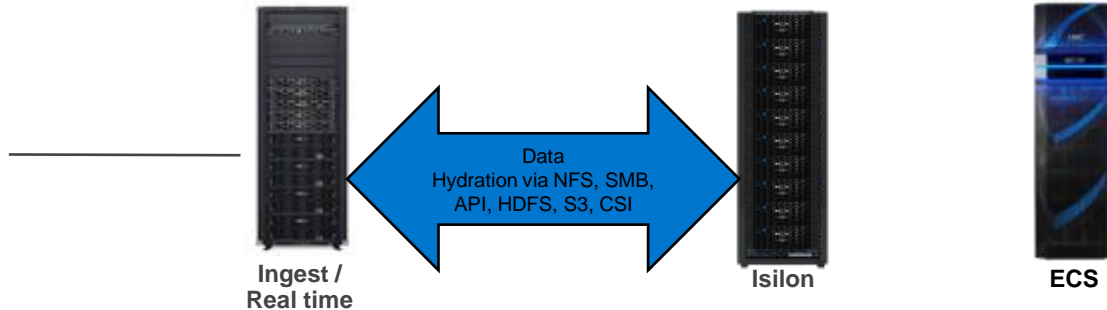
kdb+ 3.1 / RHEL 6.4 with xfs /
Intel White-box server/
4 x Intel® Xeon® E7-4890 v2 CPU @ 2.80GHz /
2x QLogic QLE2562 – PCI-E Dual Channel 8Gb

Centralized File / Object

Market /
Exchange

Ingest /
Real time

Data
Hydration via NFS, SMB,
API, HDFS, S3, CSI

Isilon

ECS

D&LL Technologies

Overcoming Friction by Increasing Velocity

Data explosion with constrained budgets

Unpredictable cloud data costs due to egress and access fees.

Data's value continues past the storage appliance lifecycle.

Disruptive upgrade paths create churn and create barriers to innovation.

Cost of storage administration on the rise due to skilled labor gap.

We made the investment to better serve

Data In-lining

25% Streaming performance gain when upgrading to OneFS 9.5 (Not a STAC benchmark)

Lower storage costs with new 15 and 30TB QLC drives.

Improved SmartPools tiering enables better cost/TB using blended nodes.

Benefit from PowerEdge 16G based portfolio.

Data Reduction Services

Improved cloud economics with new Public cloud and APEX aaS offers.

Network Acceleration with RDMA

DELLTechnologies

Case Study | Major Financial Services

Extreme concurrency and speed are king

Advanced Algorithmic Trading- The OMG Workload

This organization regularly runs upwards of 750 concurrent quant jobs which span years worth of data. These jobs keep getting bigger and more complex but the majority them focus on just 10 days worth of data, the same 10 days worth of data. This is why concurrency matters.

Today's solutions try to solve this embarrassingly parallel concurrency problem with higher performance NAS or by creating copies across filers. **But** these solutions don't scale linearly and have diminishing returns on both performance and investment.

Business Challenge: The financial services industry needs infrastructure that can increase the agility and capabilities of their data science departments to differentiate their business.

Isilon Impact

Up to **5x better performance** at high concurrency enbables organziations to get answers and make decisions faster with:

- **Shorter Model Development Time**

- **Score Models Faster**

- **Support 100's of Concurrent Jobs**

- **Support Larger Data Sets**

- **Shrink the footprint; Lower TCO**

Not a STAC benchmark   es