Getting Applications to Move at the Speed of Flash

Matt Goldensohn

Director, WHIPTAIL Labs

More than 20 Years Ago This Was a Game Changer. WHY?

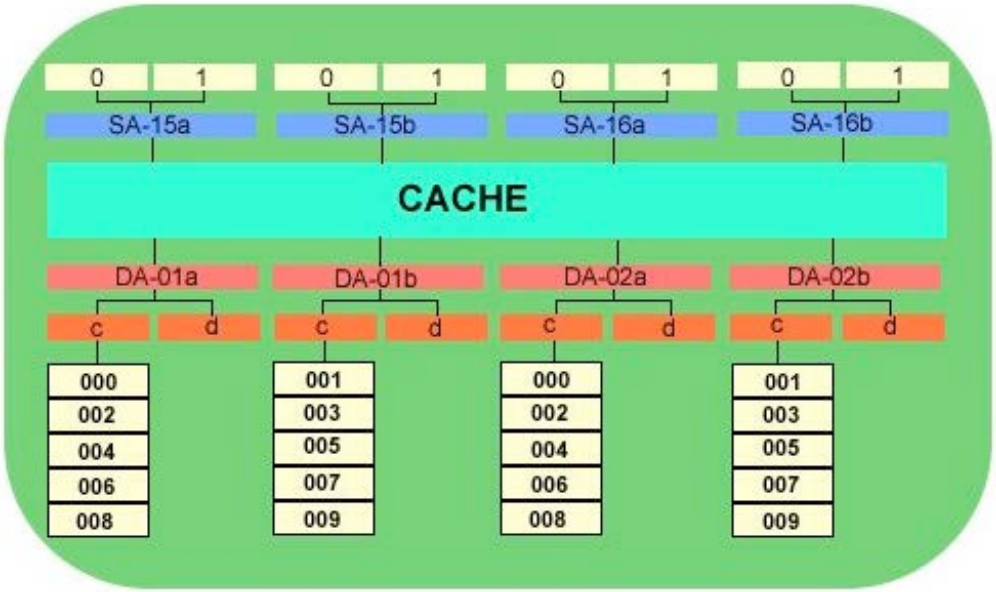

Cost/GB | The Storage Trade-Off

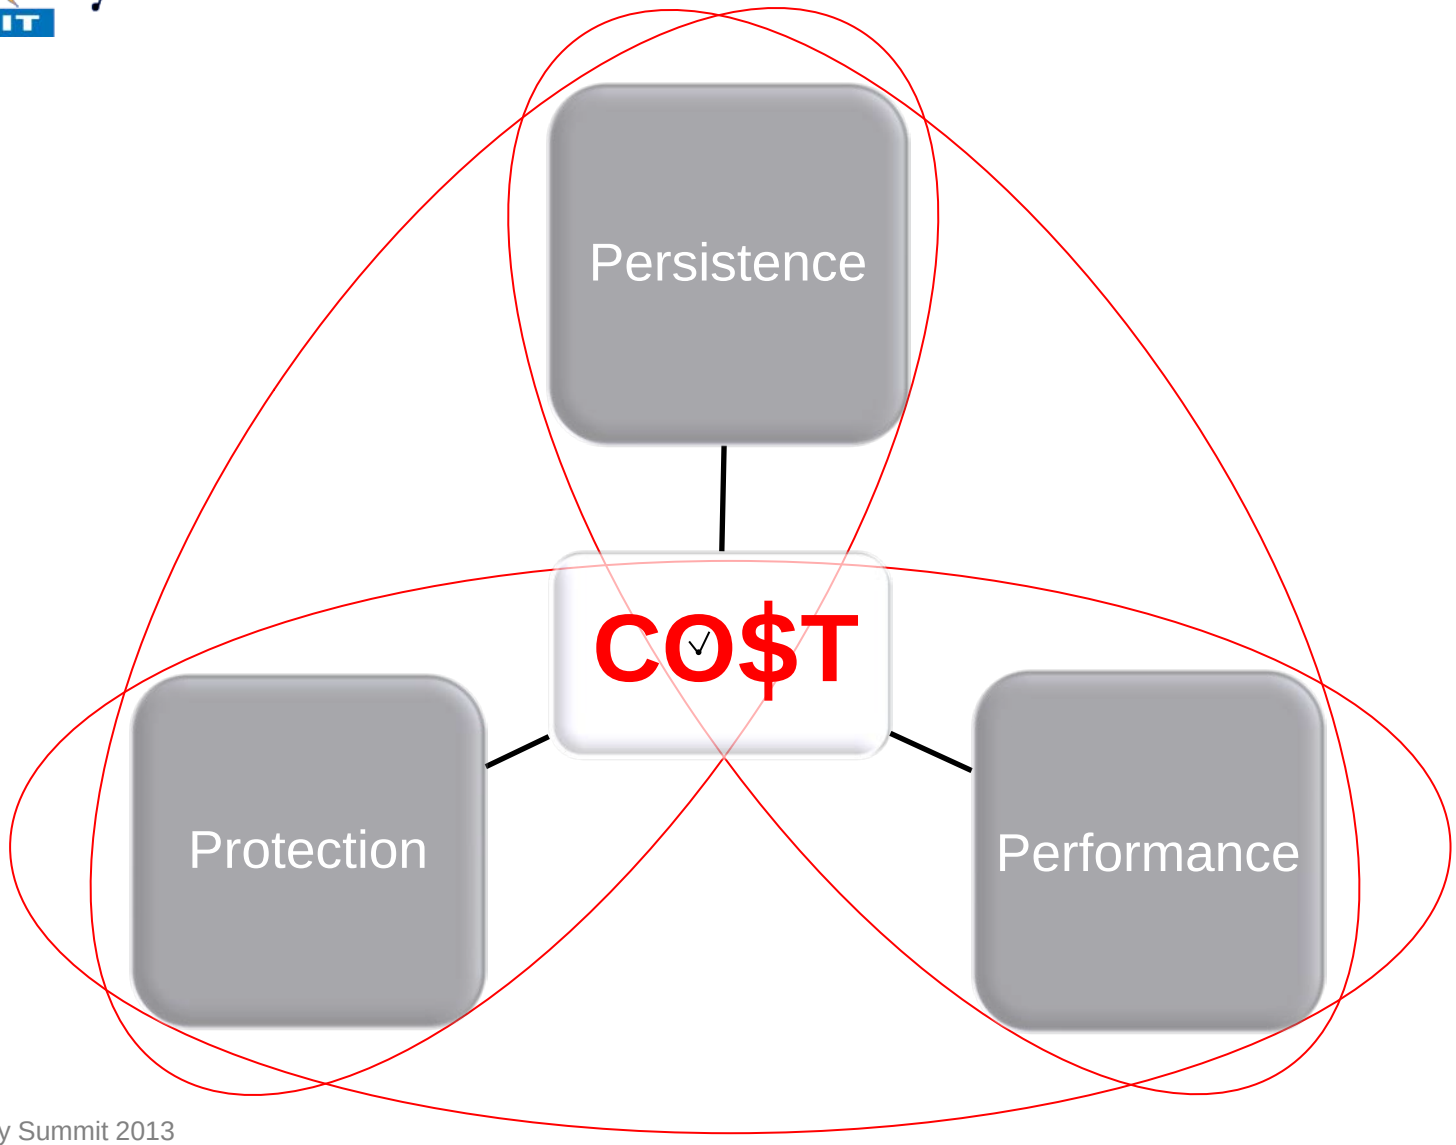

Flash is Today's Game Changer

- Traditional Storage Metrics *No Longer Apply*
  - Chasing IO, Writes & Reads
  - Measurement Based vs. Duration Based Testing
  - Use real data whenever possible
- SSD Moves Performance Constraints Up the Compute Stack
  - Operating System Configuration
  - Application
  - Network
  - Cache Management
- Expanding Pools of Flash Storage
  - Implications at the Application Level
    - Programs Have Been Focused on Mitigating Writes
    - Capacity to support an application is shrinking
    - With an abundance of capacity, *application design will be able to grow*
  - Meet tighter recovery time and point objectives

Flash Deployment will Affect Enterprise Application Behavior

- Enterprise Applications are Heavily WRITE Dependent
  - Lock, Group, & Join are All Writes
- A Flash Aware BTL is **Critical** to Sustain High WRITE Performance
  - Mitigate poor write behavior of SSDs
  - Garbage Collection
  - Minimize Write Amplification
  - Protection – RAID

Flash Accelerates Workloads Across All Industries

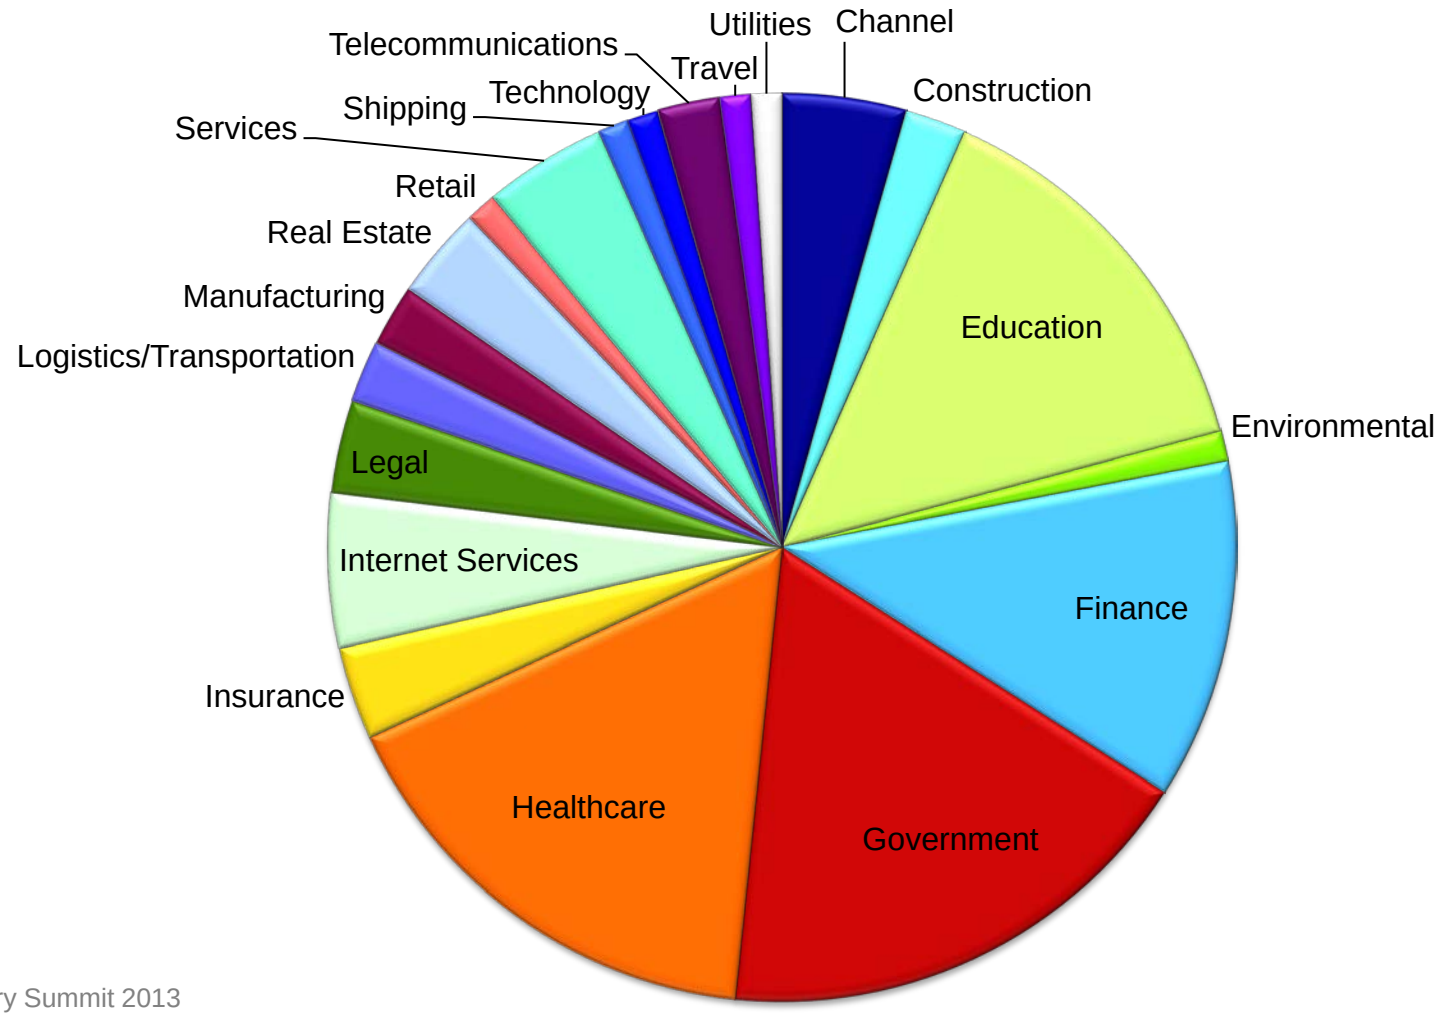

Enterprise Applications Deployed on Flash Platforms

High IOPS

Analytics & Intelligence

High IOPS, High Bandwidth

Batch Processing

High IOPS

Database Loads

High IOPS, Low Latency

Online Transaction Processing

High IOPS

Email

High Write IOPS

Virtual Desktops

High IOPS, High Bandwidth

Video Transcoding

Low Latency

High Performance Computing

Cost/Transaction | You CAN Have It All

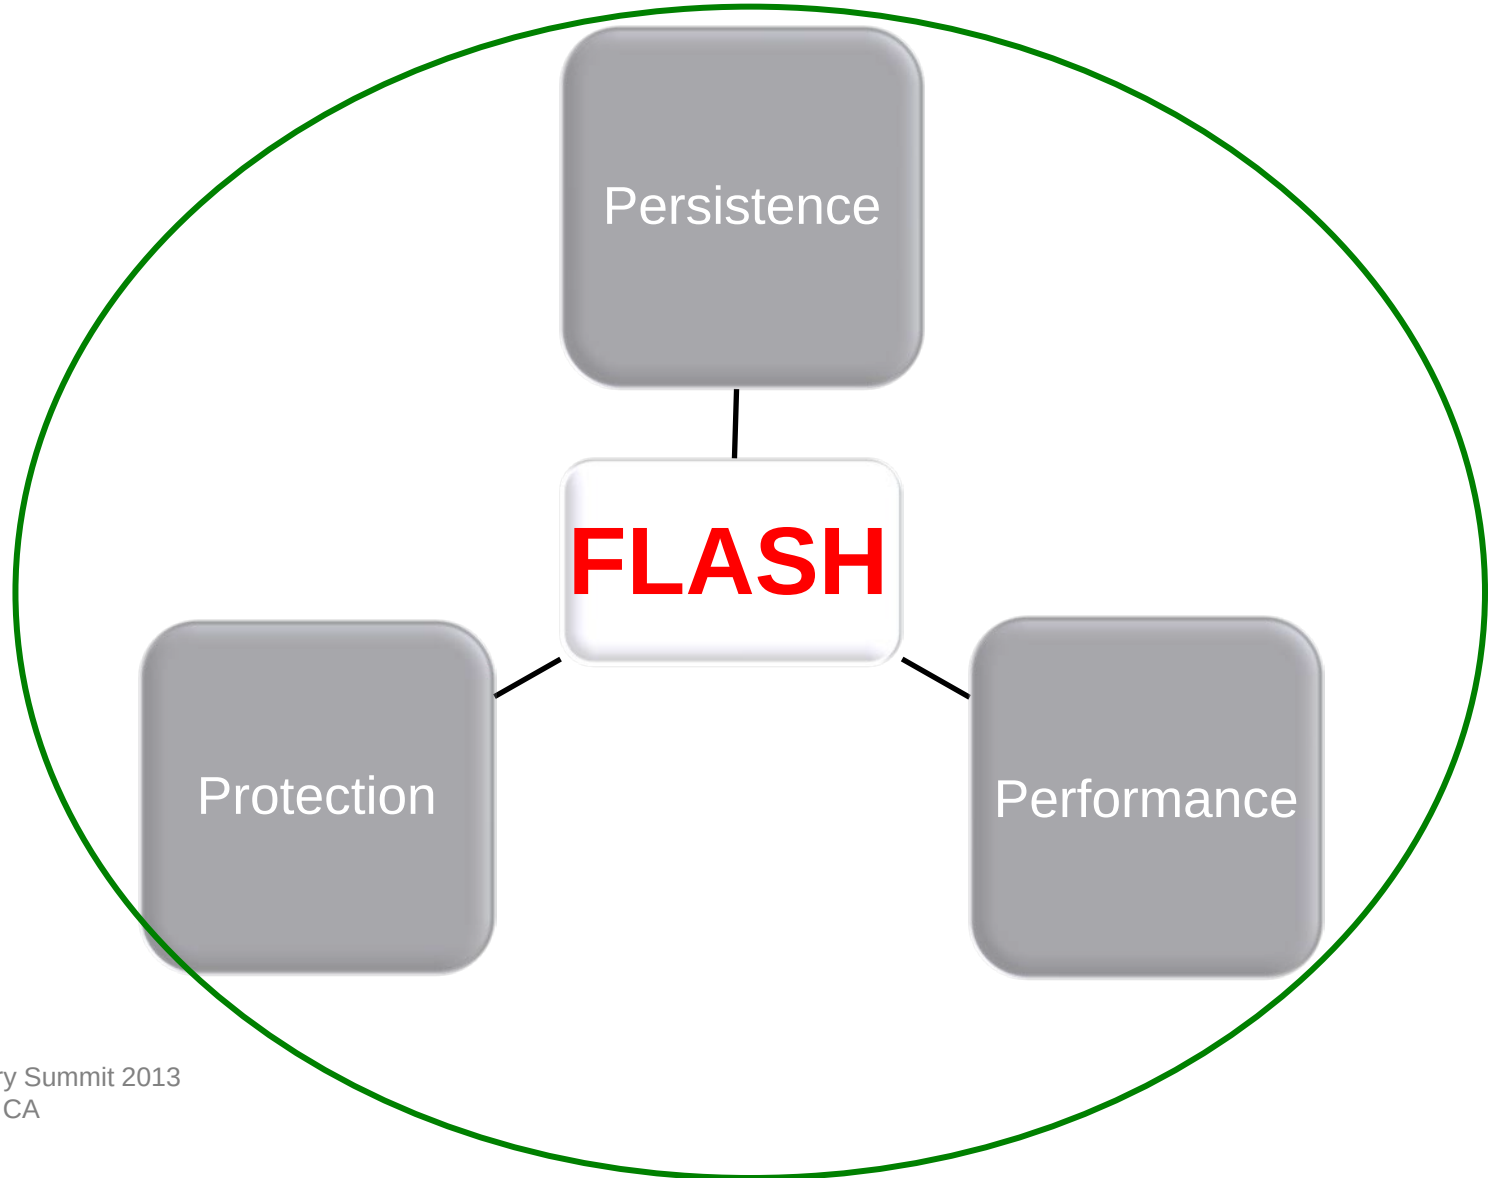

What does the Labs Do?

Rethink Deployments

Real-World Applications

Exploit Performance of INVICTA Flash Array

What We Have Learned In The Lab

- Applications must be Optimized for Flash
  - Buffer Less
  - Read Faster
  - Write Faster
- Consolidation
  - Databases
  - Process
- Retune
  - Indexing Strategy
  - Placement
  - Block Size
  - OS IO Scheduler

Application Optimization

- Write Faster
- Read Faster
- Buffer Less

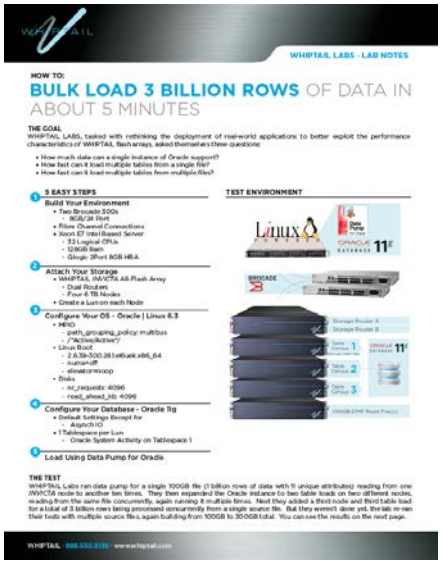

WHIPTAIL LABS - LAB NOTES

HOW TO:

BULK LOAD 3 BILLION ROWS OF DATA IN ABOUT 5 MINUTES

THE GOAL

5 EASY STEPS

TEST ENVIRONMENT

THE TEST

RESULTS

100GB / 1 Billion Rows

READ	WRITE	AVG. TIME
		00:02:55
		00:04:30
		00:05:32
		00:03:35
		00:05:32

Consolidation

- Databases
- Processes

WHIPTAIL LABS · LAB NOTES

HOW:

POSTGRESQL SCALES OUT TO **LOAD 15 BILLION RECORDS** AT ONCE

THE GOAL

5 EASY STEPS

TEST ENVIRONMENT

4 Node INVICTA Configuration

THE TESTS

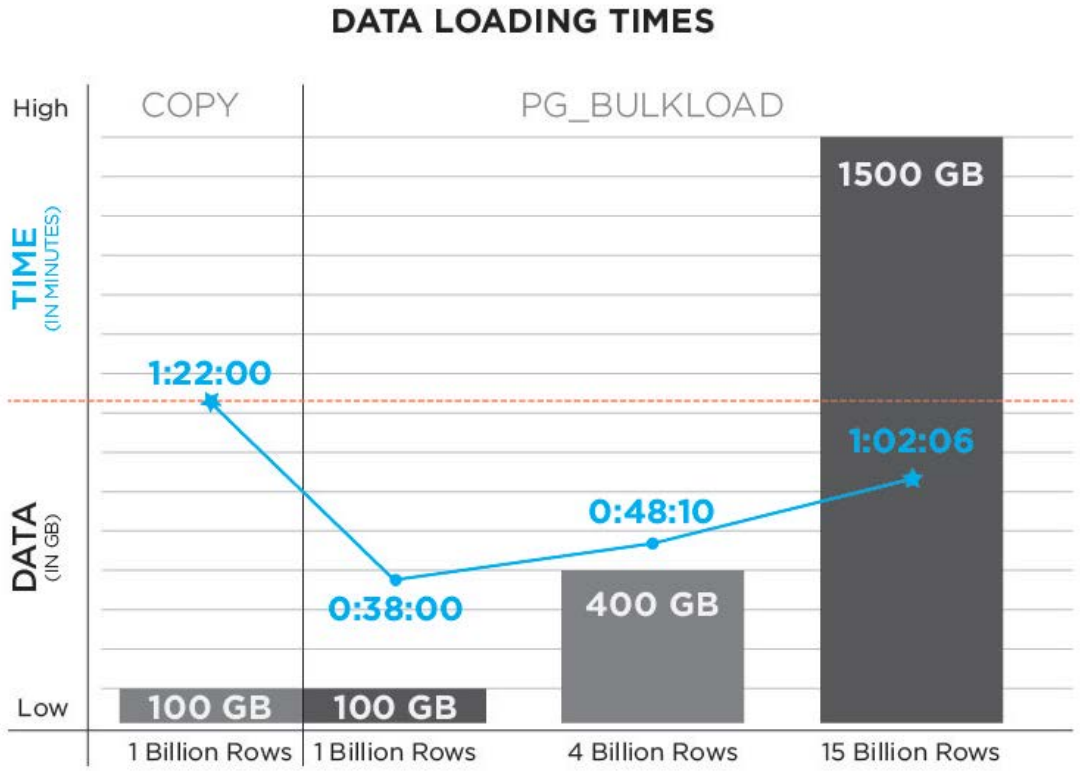

Rethink, Retune

- Indexing Strategy
- Placement
- Block Size
- OS IO Scheduler

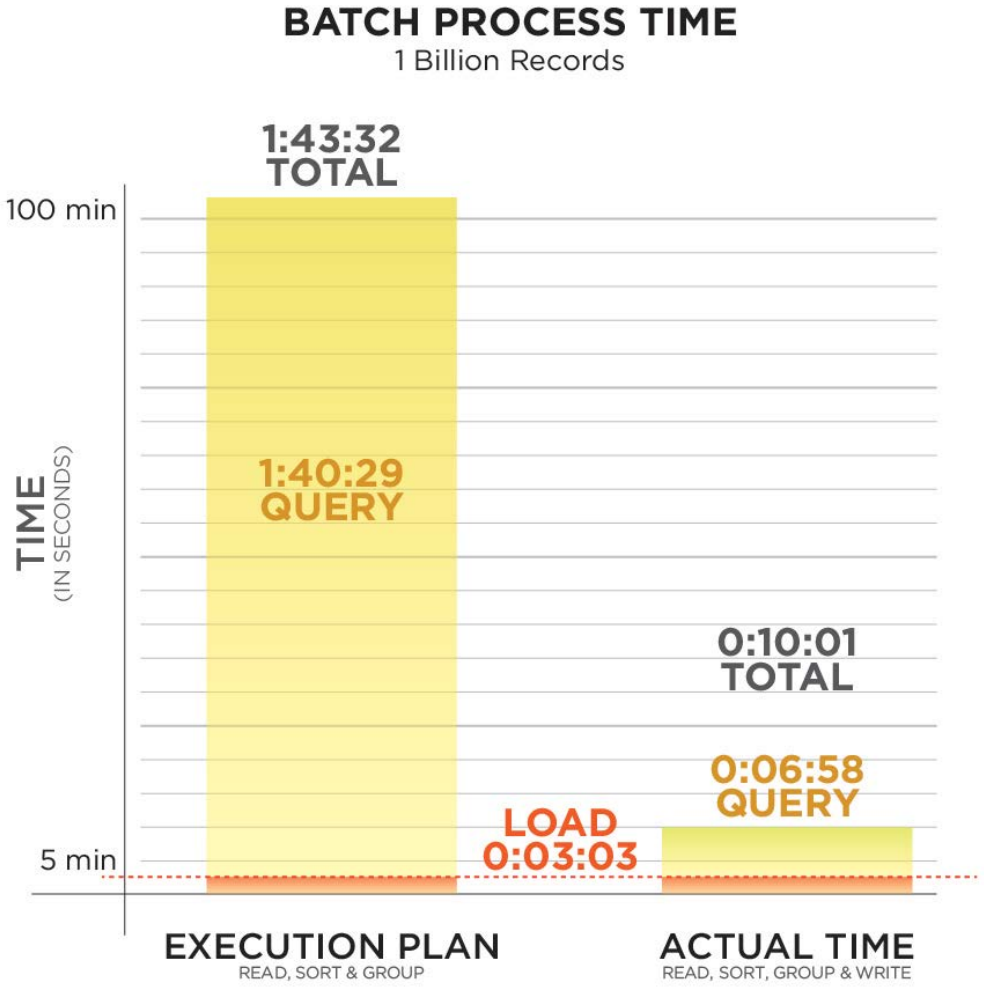

The Take Away

- Perform Ad Hoc Queries Quickly & Easily
- Consolidate Databases & Processes
  - Manager Fewer LUNs
  - Lower Costs
    - Licensing
    - Server Management
    - Maintenance
  - Efficient Utilization of Existing Resources
- Critical Enterprise Applications Move to Flash
  - Prepare for New Challenges
  - Unique to Each Application