Service Level Objectives for Storage Solutions with Different Applications

Dr. M. K. Jibbe

Director of Quality Architect Team, NetApp APG

mahmoudj@netapp.com

K. H. Tan,

MTS QA PTI, NetApp APG

KuokHoe.Tan@netapp.com

Abstract

- Application dictates what other components are used to determine the solution.
- Solutions will determine the SLO that we need to target.
- QA based on SLO for each Application.
- Current problems.
  - Very little time measurement
  - Performance measurement
  - Redundancy among teams
  - Injection timing measurement

Component Overview

Applications	Platforms	Connectivity & Protocol	Storage
A		Fiber Channel	
B		SAS	
		iSCSI	
C		Infiniband	

End customer solutions from applications to our storage systems.
Applications: Analytics, Big Data and Content
Platforms: OS, File systems, I/O stack incl. Failover, Connectivity Devices
Connectivity & Protocol: Switches and Protocols
Storage: E-Series systems

Component Parameters

SLO → **Criteria**

Applications	Platforms	Connectivity & Protocol	Storage
I/O Profile IOPs / MBs Latency Connectivity Profile	Emu. Tools I/O Tools Mea. Tools Tuning	Link Speed Bandwidth Latency Tuning	Front/Back IOPs/MBs Latency System Utilization Concurrent Multi-Protocol Support

A

B

Fiber Channel

SAS

iSCSI

C

Infiniband

High Level Design Flow (1 of 2)

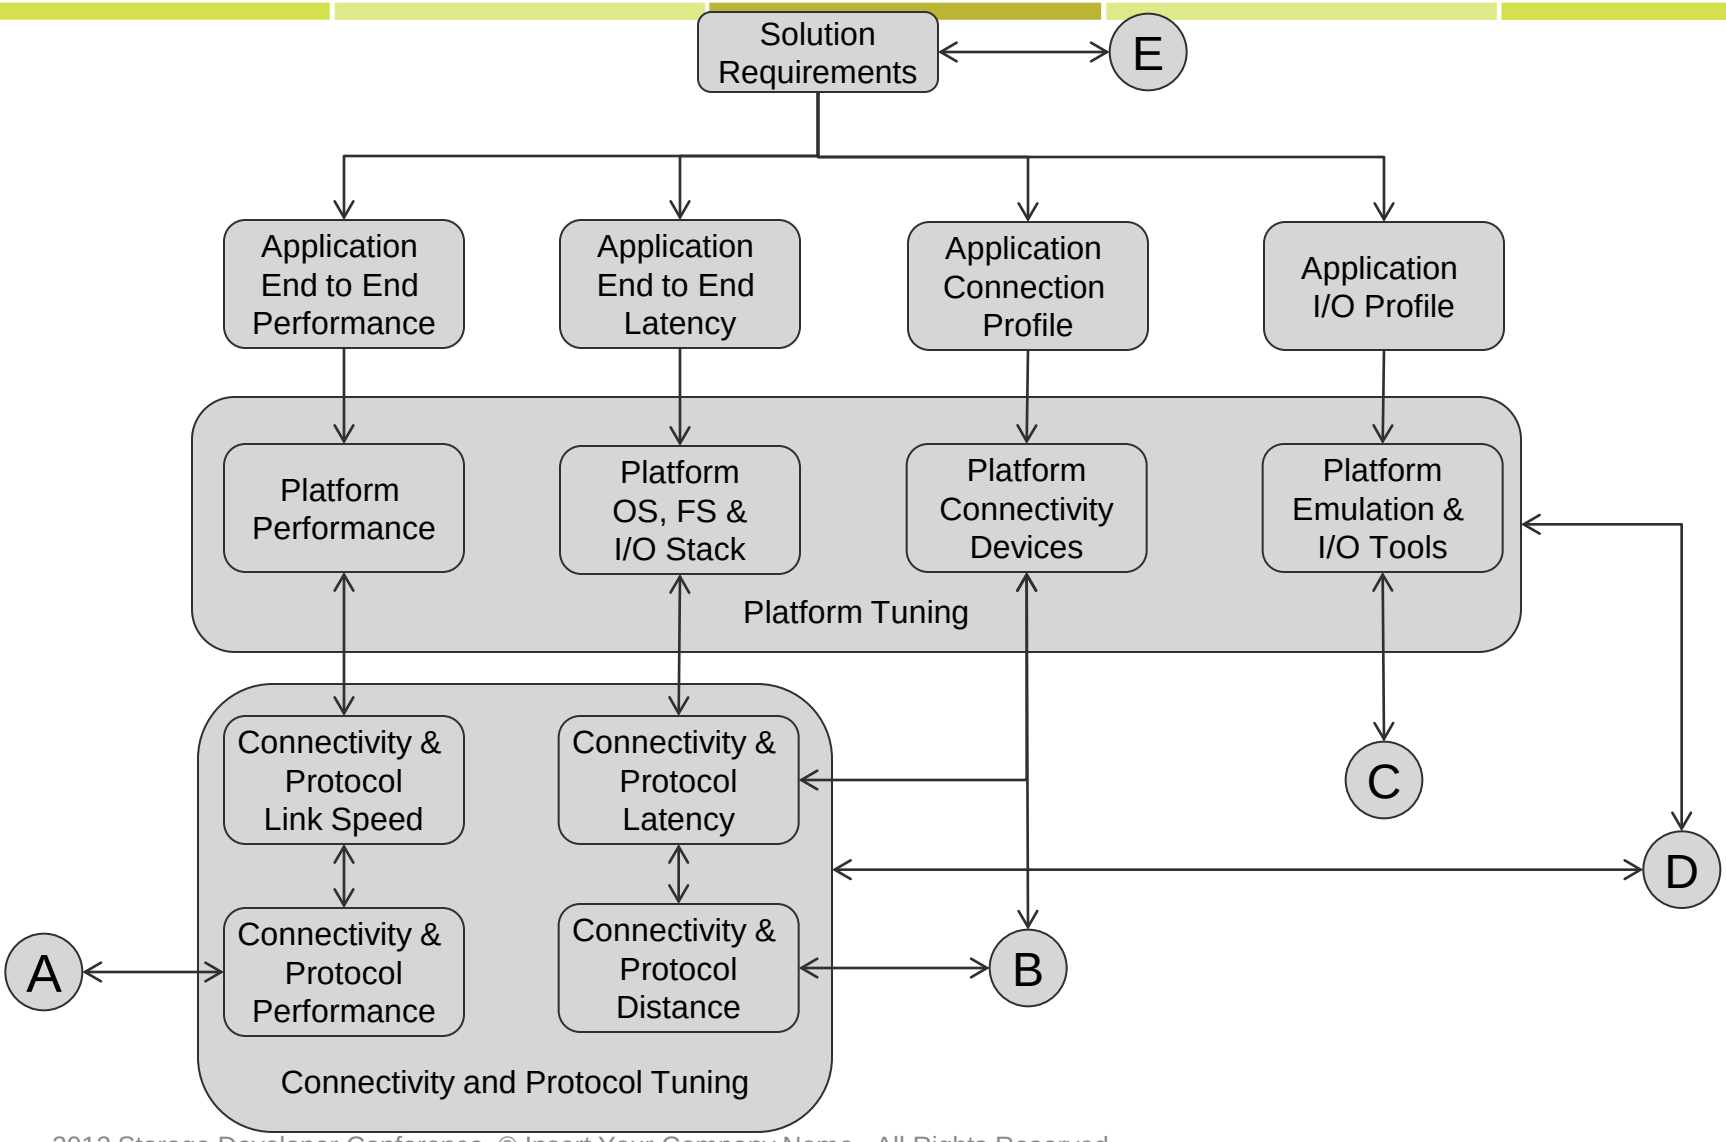

High Level Design Flow (2 of 2)

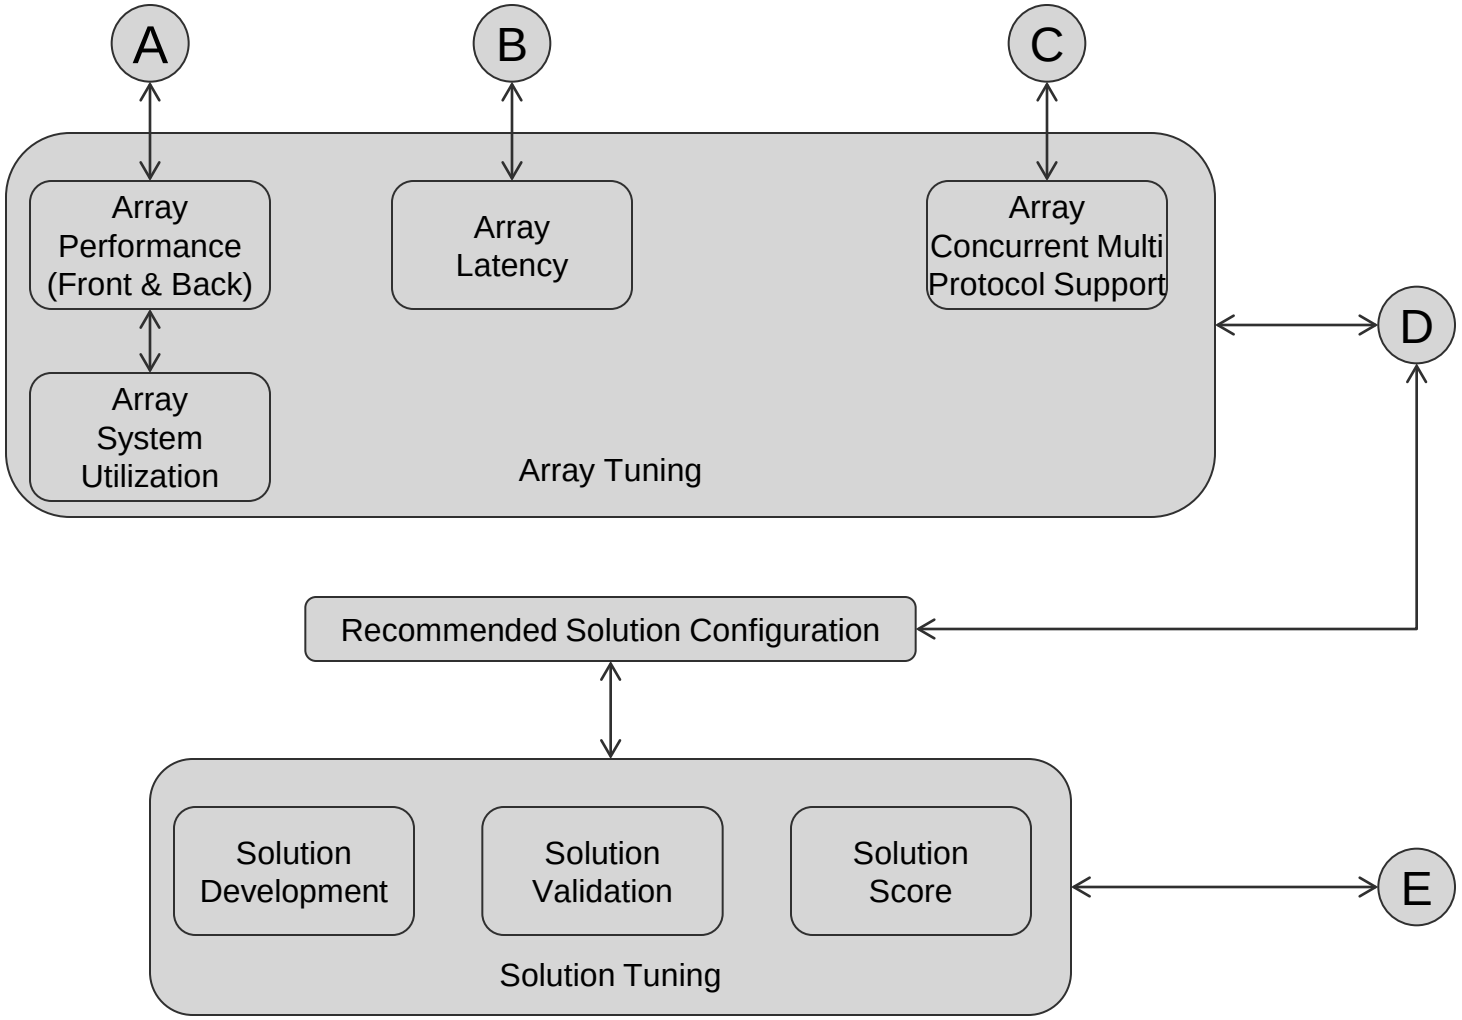

Discrete and Deterministic System Testing

- Service Level Objective (SLO) focus for various system states
  - Optimal system state
  - Transitional system state
  - Degraded system state
- Discrete criteria focus
  - I/O and Access Profile
  - IOPs/MBs
  - Latency (Optimal and Fault Tolerant Latency)
- Deterministic (bounded) operational behavior
  - Controlled and repeatable system criteria
  - Controlled and repeatable system measurement points
  - Controlled and repeatable system injection points
- Division of validation effort
  - SLO serves as overarching criteria for all validation efforts

System Event Diagram - Optimal

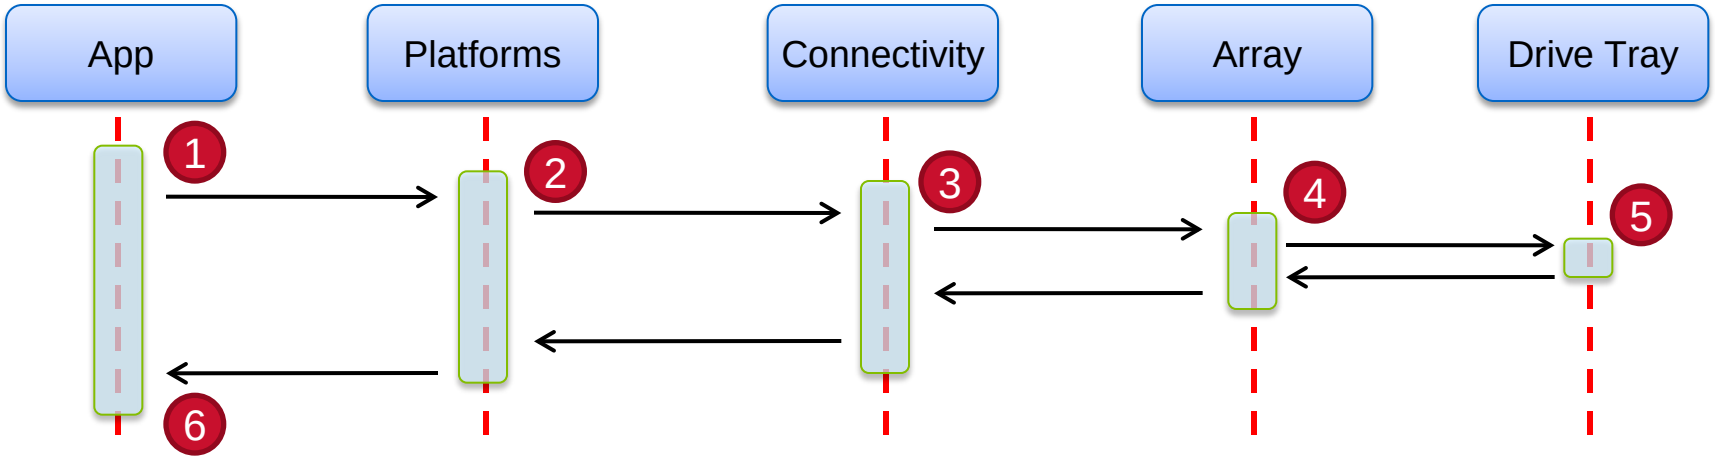

1. Application initiates I/O requests
2. Platform processes I/O requests
3. Interconnectivity (SAN, Ethernet, etc.) routes I/O requests to end devices
4. Array processes I/O requests
5. Drives processes R/W requests
6. Application completes I/O requests within target SLO

System Event Diagram - Exception

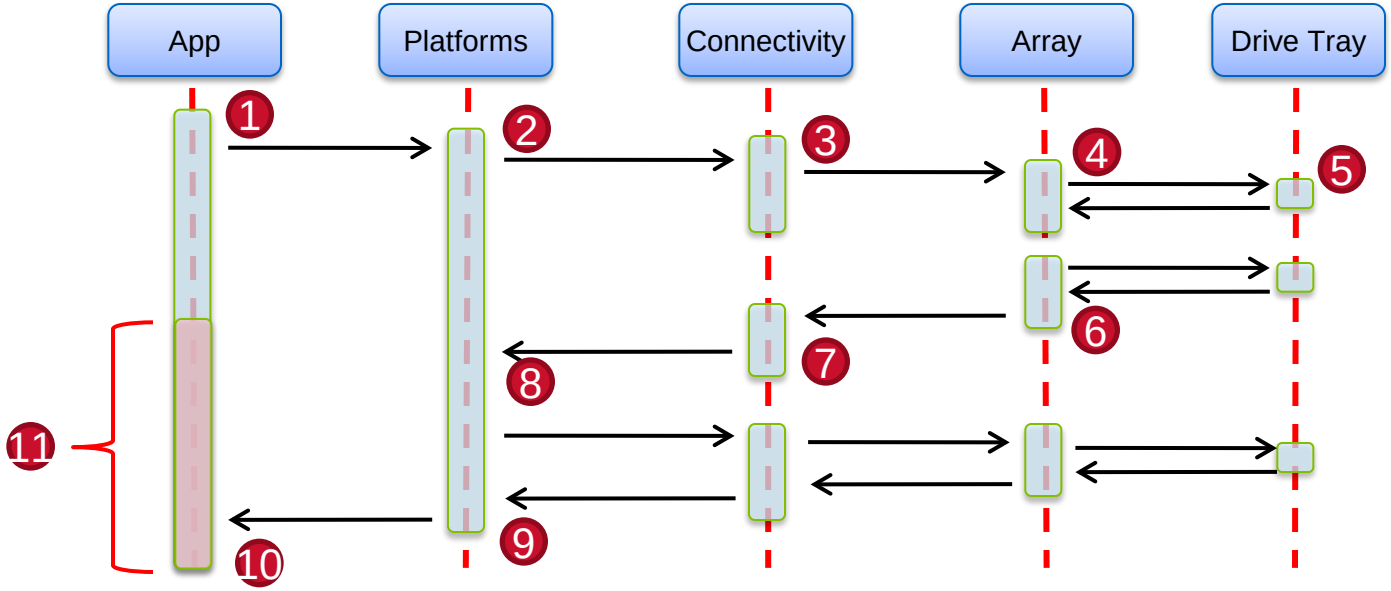

1 - 5 – Optimal operation

6 – Array rebooted

7 – Rejoins connectivity infrastructure

8 – Platforms detects link event and retries application I/O requests

9 – Optimal operation to complete retried I/O requests

10 – Application completes I/O requests within extend SLO targets

11 – Exception induced extended SLO

Validation Focus

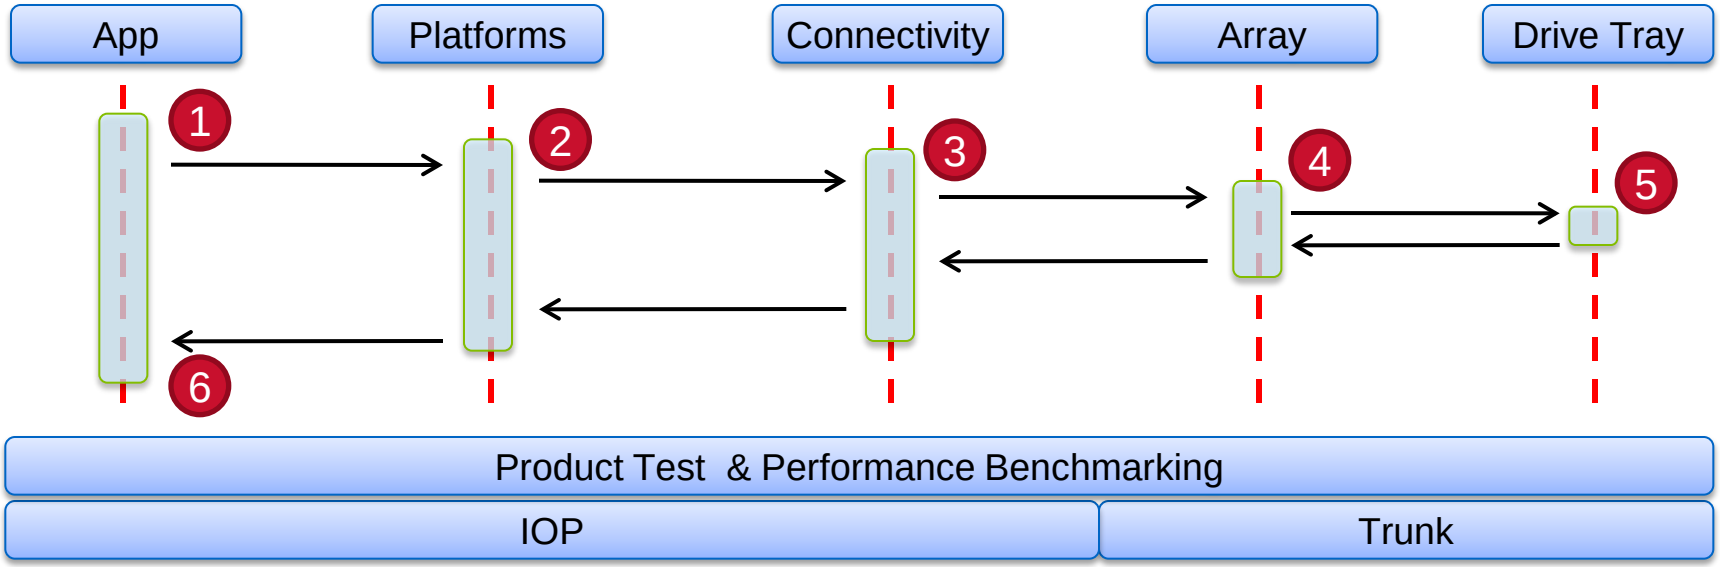

- All validation focus on key SLO
  - Clear areas of focus for each validation team
  - Streamline effort so that we can build on top of each other efforts
  - End to end focus on key SLO
- Pass/fail criteria determined by key SLO and end user experience
  - Multi faceted criteria to satisfy key SLO and user expectations
  - Criteria for each component must satisfy key SLO and user experience

SLO – Example Overall Flow

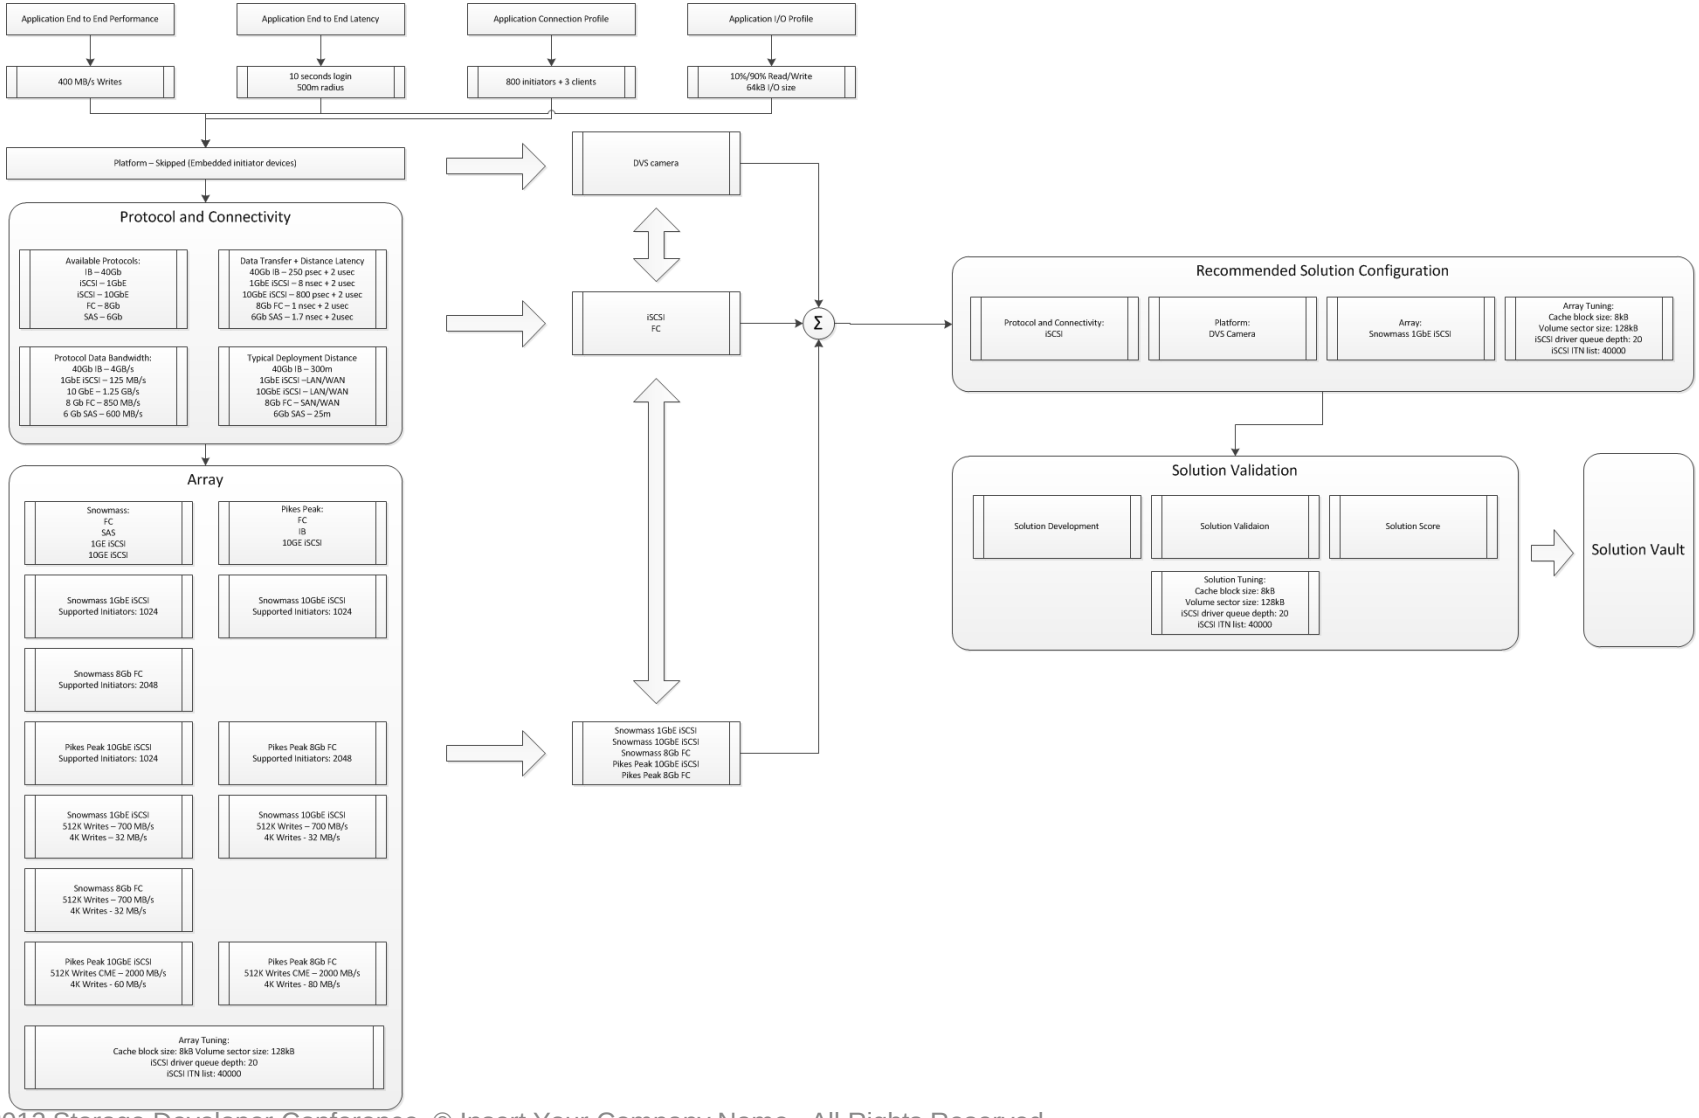

Example Overall Flow – Zoom In (1 of 3)

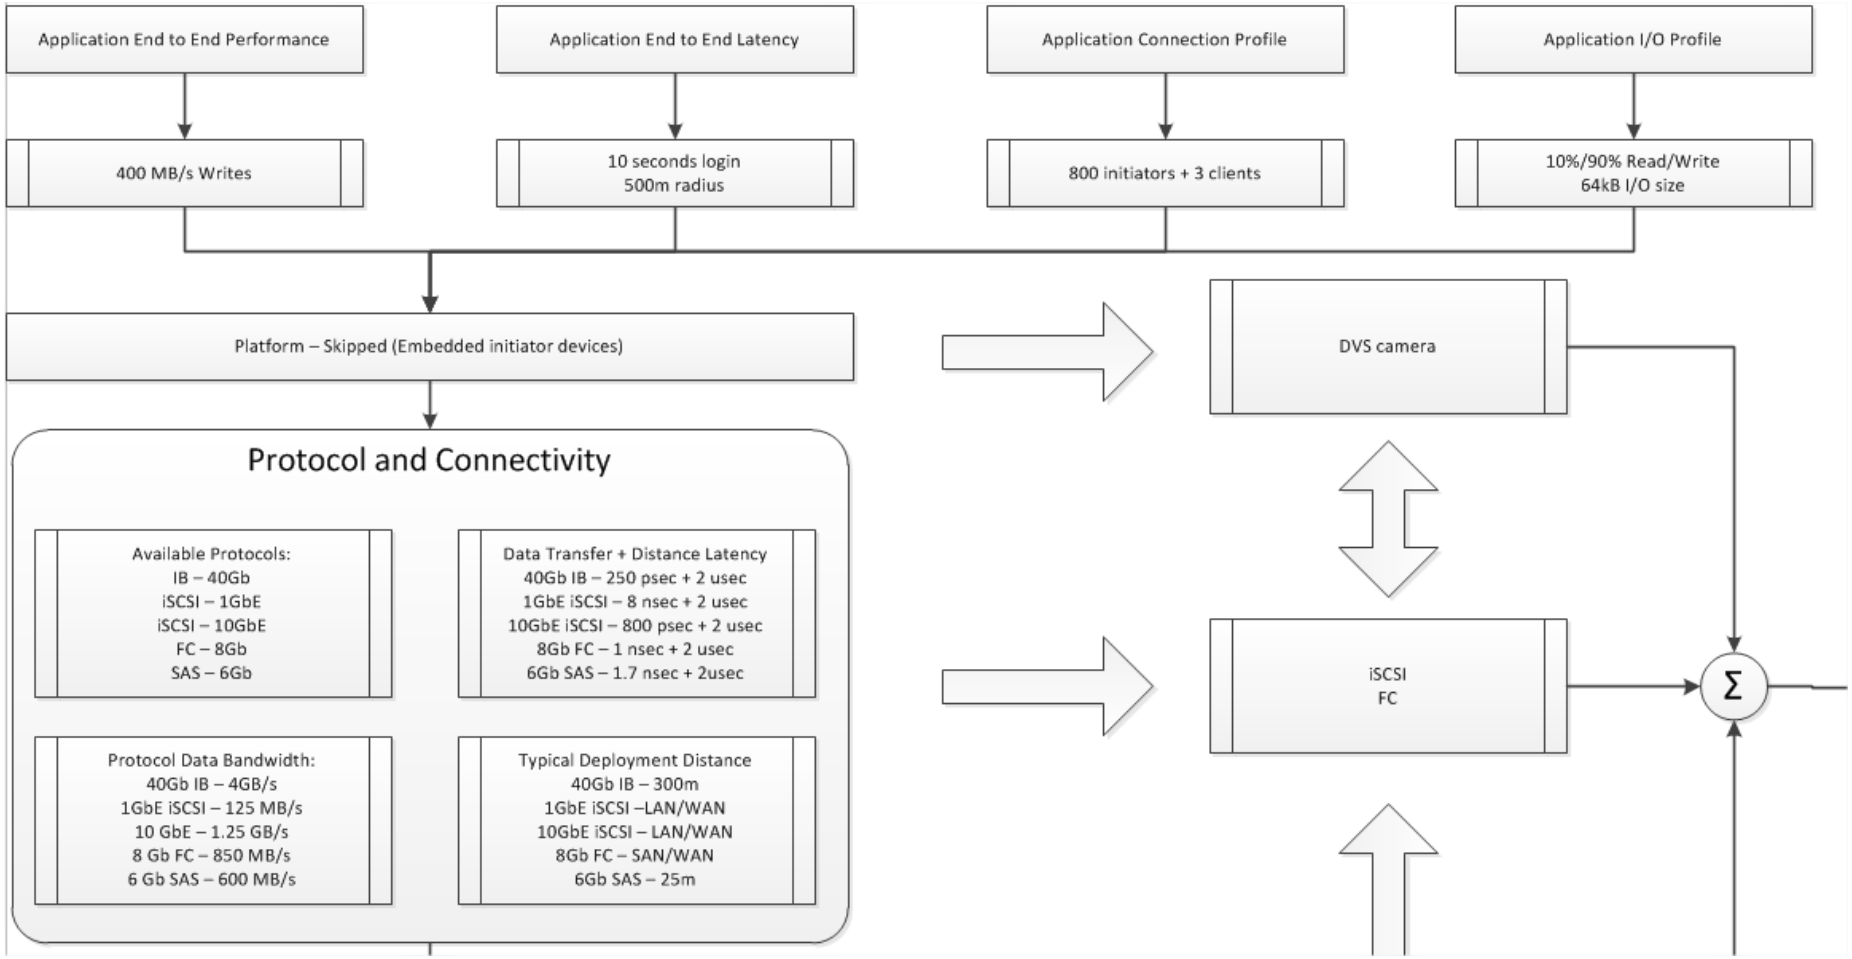

Example Overall Flow – Zoom In (2 of 3)

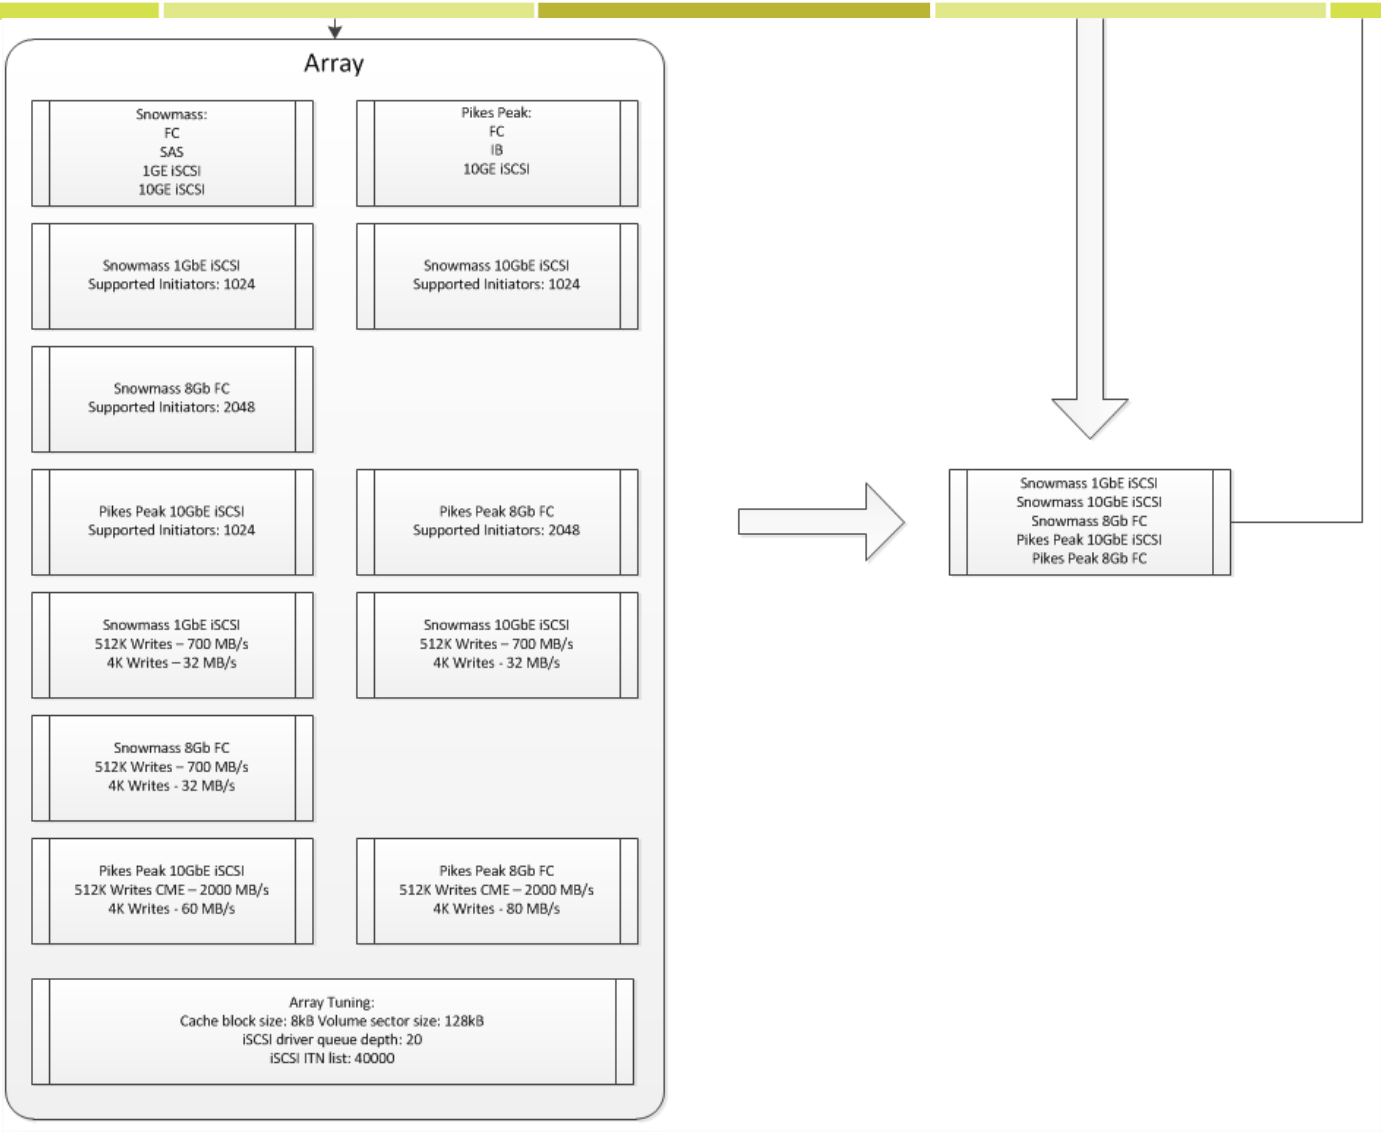

Example Overall Flow – Zoom In (3 of 3)

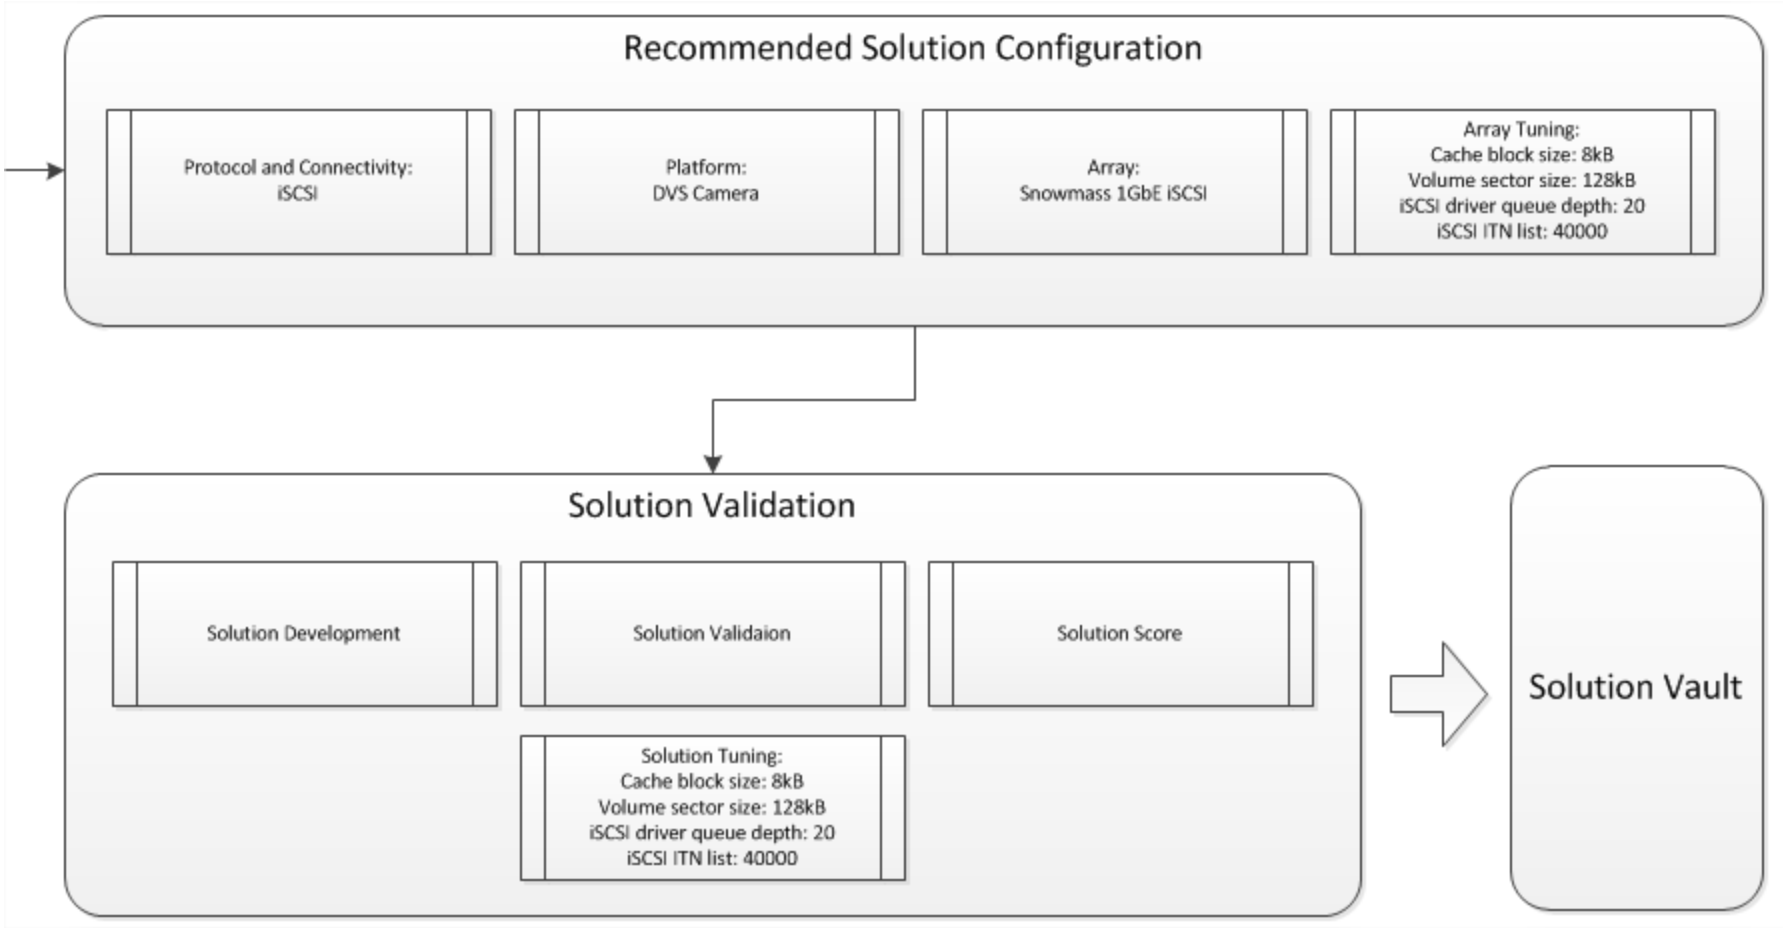

SLO To Platform Parameter Mapping

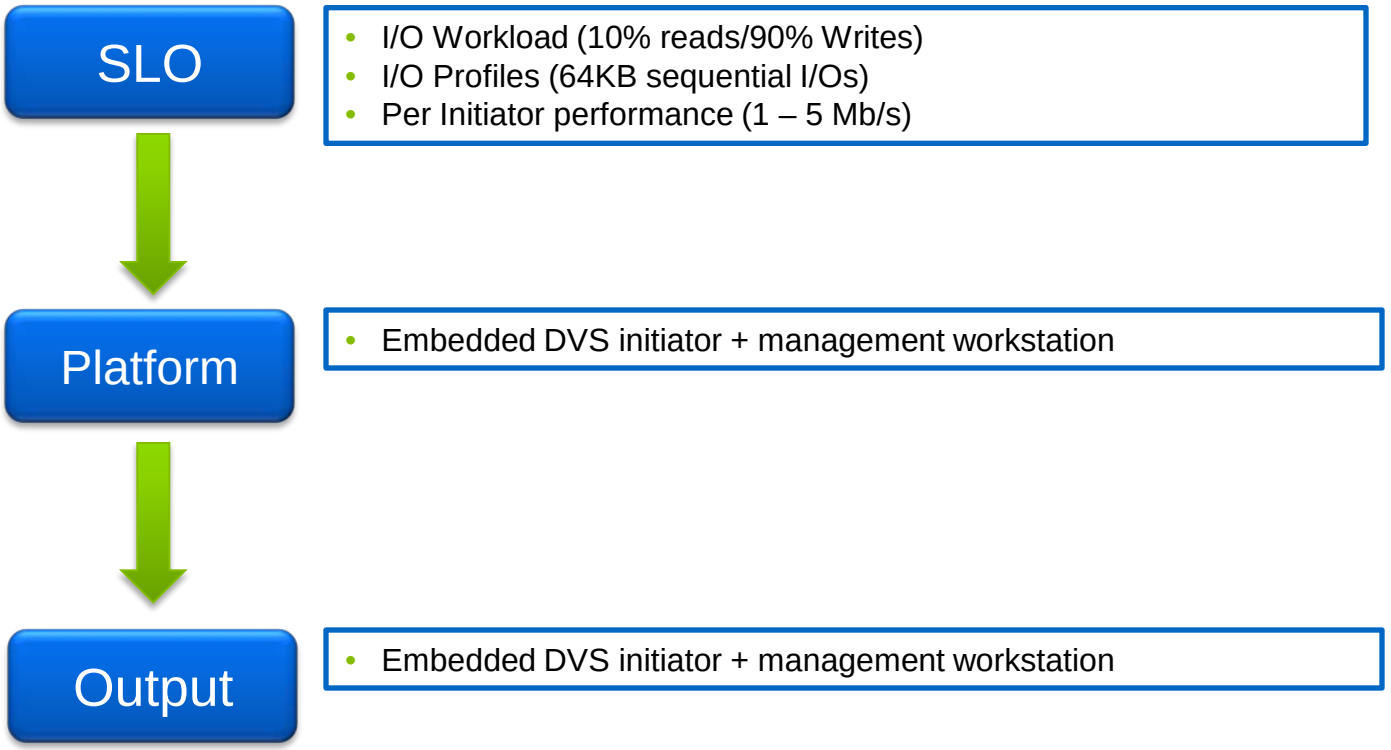

SLO To P & C Parameter Mapping

SLO

- Device Timeouts (< 10 sec logins)
- Product Boundaries (Unlimited)
- QoS (Average < 1 frame @ 30 fps)

Protocol & Connectivity

- IB 40Gb
  - Data Transfer + Distance Latency – 250 psec + 2 usec
  - Data Bandwidth – 4 GB/s
  - Typical Deployment Distance – 300m
- iSCSI 1GbE
  - Data Transfer + Distance Latency – 8 nsec + 2 usec
  - Data Bandwidth – 125 MB/s
  - Typical Deployment Distance – LAN/WAN
- iSCSI 10GbE
  - Data Transfer + Distance Latency – 800 psec + 2 usec
  - Data Bandwidth – 1.25 GB/s
  - Typical Deployment Distance – LAN/WAN
- FC 8Gb
  - Data Transfer + Distance Latency – 1 nsec + 2 usec
  - Data Bandwidth – 850 MB/s
  - Typical Deployment Distance – SAN/WAN
- SAS 6Gb
  - Data Transfer + Distance Latency – 1.7 nsec + 2 usec
  - Data Bandwidth – 600 MB/s
  - Typical Deployment Distance – 25m

Output

- iSCSI (1 & 10 GbE)
- FC

SLO To Array Parameter Mapping

- I/O Workload (10% reads/90% Writes)
- Login Rate (800 Initiators < 40 secs)
- Device Timeout (< 10 secs logins)
- I/O Profiles (64KB sequential I/Os)
- Write BW (1600 – 1800 MB/s)
- Read BW (50% of Write BW)
- Connection Profile (800 init + 3 clients)

- SM (FC, ~~SAS~~, 1GbE iSCSI, 10GbE iSCSI)
  - FC (Supp. Init: 2048, 512KB Writes: 700 MB/s)
  - iSCSI (Supp. Init: 1024, 512KB Writes: 700 MBs)
- PP (FC, iSCSI, ~~IB~~)
  - FC (Supported Initiator: 4096, 512KB Writes: 2000 MB/s)
  - iSCSI (Supported Initiator: 1024, 512KB Writes: 2000 MB/s)
- Tuning
  - Cache Block Size: 8KB
  - Volume Sector Size: 128KB
  - iSCSI driver queue depth: 20
  - iSCSI ITN list: 40000

- SM (8Gb FC; 1GbE iSCSI; 10GbE iSCSI)
- PP (8Gb FC; 10GbE iSCSI)

SLO To Solution & Validation Result

- I/O Workload (10% reads/90% Writes)
- Initiator Speed (1 – 5 Mb/s)
- Login Rate (800 Initiators < 40 secs)
- Device Timeout (< 10 secs logins)
- Product Boundary (Unlimited)
- QoS (Avg < 1 dropped frame @ 30fps)
- I/O Profiles (64KB sequential I/Os)
- Write BW (1600 -1800 MB/s)
- Read BW (50% of Write BW)
- Connection Profile (800 init + 3 clients)
- Budget & Training ($15k - $50k)

- Platform – DVS camera
- Protocol and Connectivity – iSCSI
- Array – SM 1GbE iSCSI
- Array Tuning:
  - Cache block size – 8kB
  - Volume sector size – 128kB
  - iSCSI driver queue depth: 20
  - iSCSI ITN list: 40000

- I/O Workload (10% reads/90% Writes)
- Initiator Speed (1 – 5 Mb/s)
- Login Rate (400 Initiators < 40 secs)
- Device Timeout (8 - 10 secs logins)
- Product Boundary (Unlimited)
- QoS (0 dropped frames)
- I/O Profiles (64KB sequential I/Os)
- Write BW (1848 MB/s)
- Read BW (52% of Write BW)
- Conn. Profile (800 init on 2 separate ports)
- Budget & Training ($16k - $65k)

Thank you