CirrusData

DSI 2015 Session

Flash Cache in the Data center
The Future is NOW

Wayne Lam
Founder
Cirrus Data Solutions, Inc.

Cirrus Data Solutions (CDS) at a Glance

Company Snapshot	• Founded August 2011 by a team of data storage experts each with over 20 years of experience. Founders have track record for multiple successful startups. • Core team of seasoned engineers in storage virtualization, protection/recovery, I/O acceleration, and connectivity protocols (iSCSI, FCoE, Infiniband, FC, CIFS, NFS). Together, they hold 18 US patents. • Specializes in supplying storage solutions to enterprise and cloud data centers to lower their cost of operations.
Target Clients	• Enterprise and Midmarket end-customers across all market segments who need to reduce IT TCO while improving efficiency.
Partners	IBM, Cognizant, DELL, FUJITSU, datalink
Financials	• 2011/2014 Private Funding.

Quick Tutorial on FC Zoning

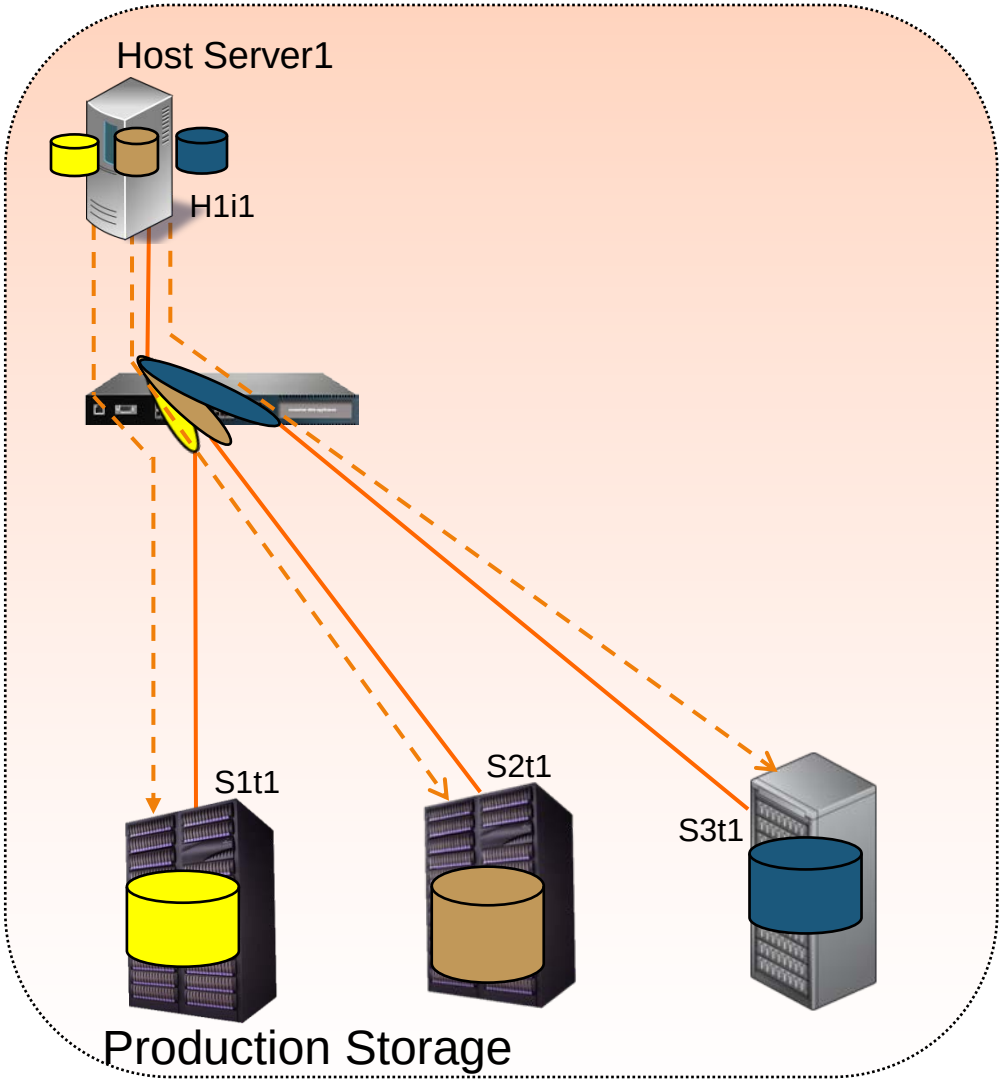

WWNs

Host1 Initiator1:	H1i1
Storage1 Target1:	S1t1
Storage2 Target1:	S2t1
Storage3 Target1:	S3t1

FC Zones

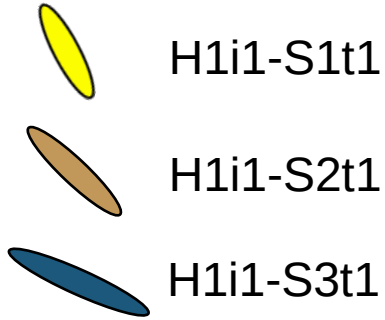

Basically:
Zones are logical connectors

FC Port Spoofing

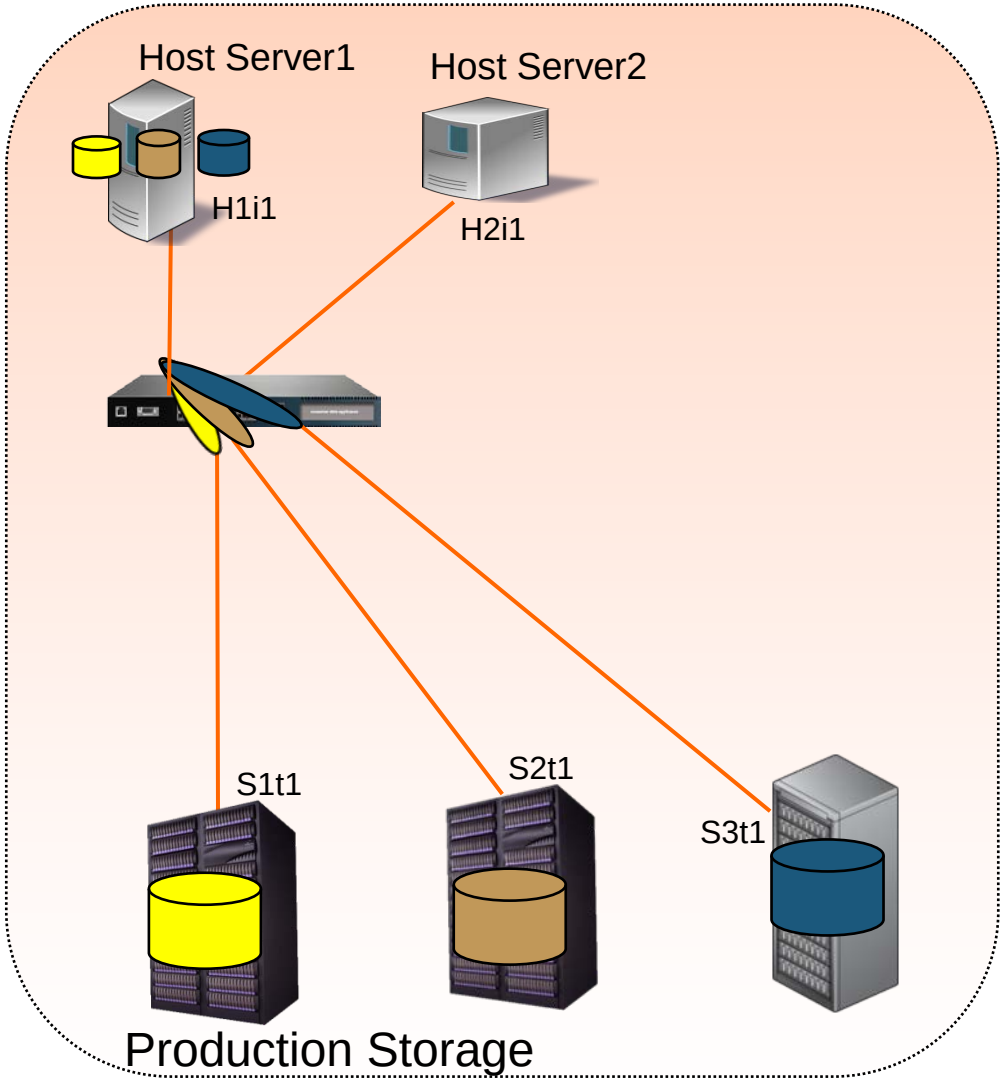

What if a different Host takes over the H1i1 WWN?

FC Port Spoofing

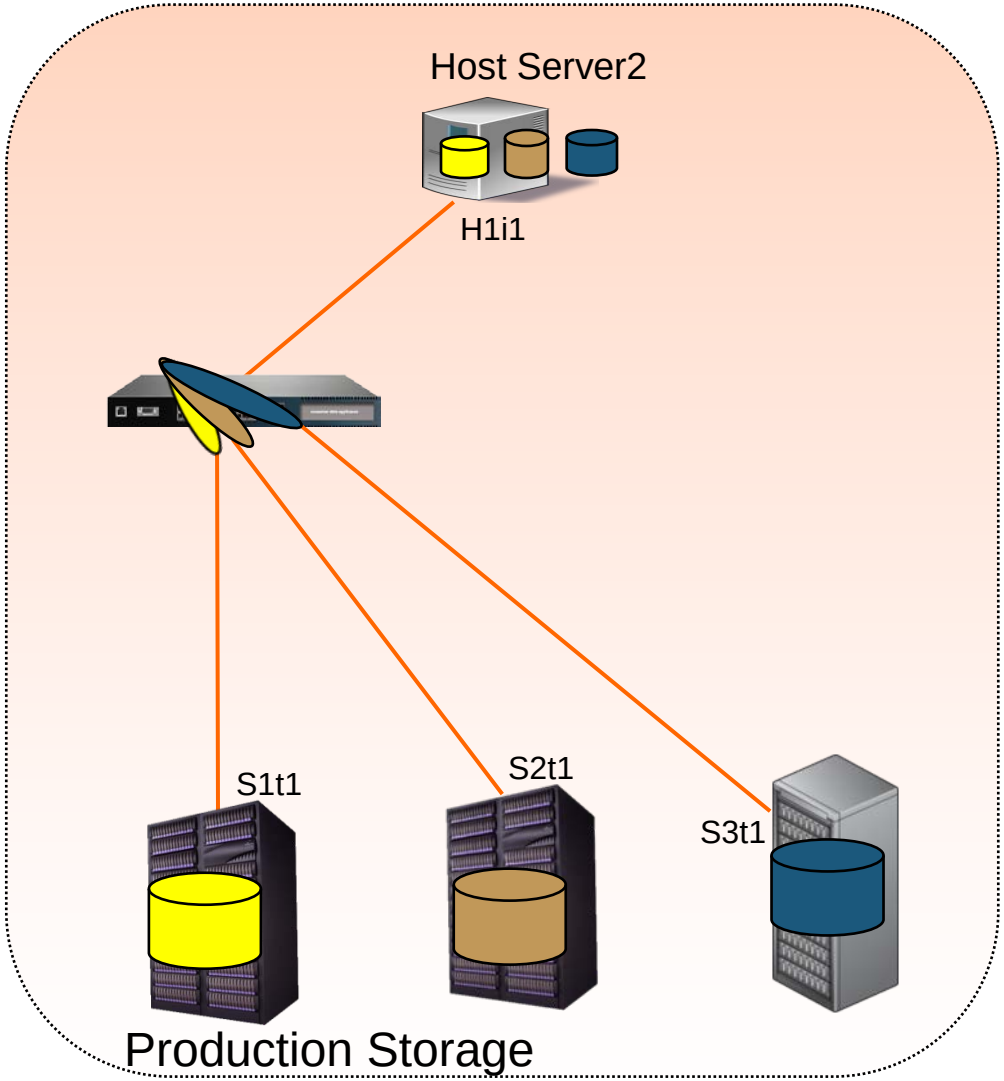

What if a different Host takes over the H1i1 WWN?

Host Server2 WILL have access to all the disks of Host Server1!

FC Port Spoofing: The Bad and the Good

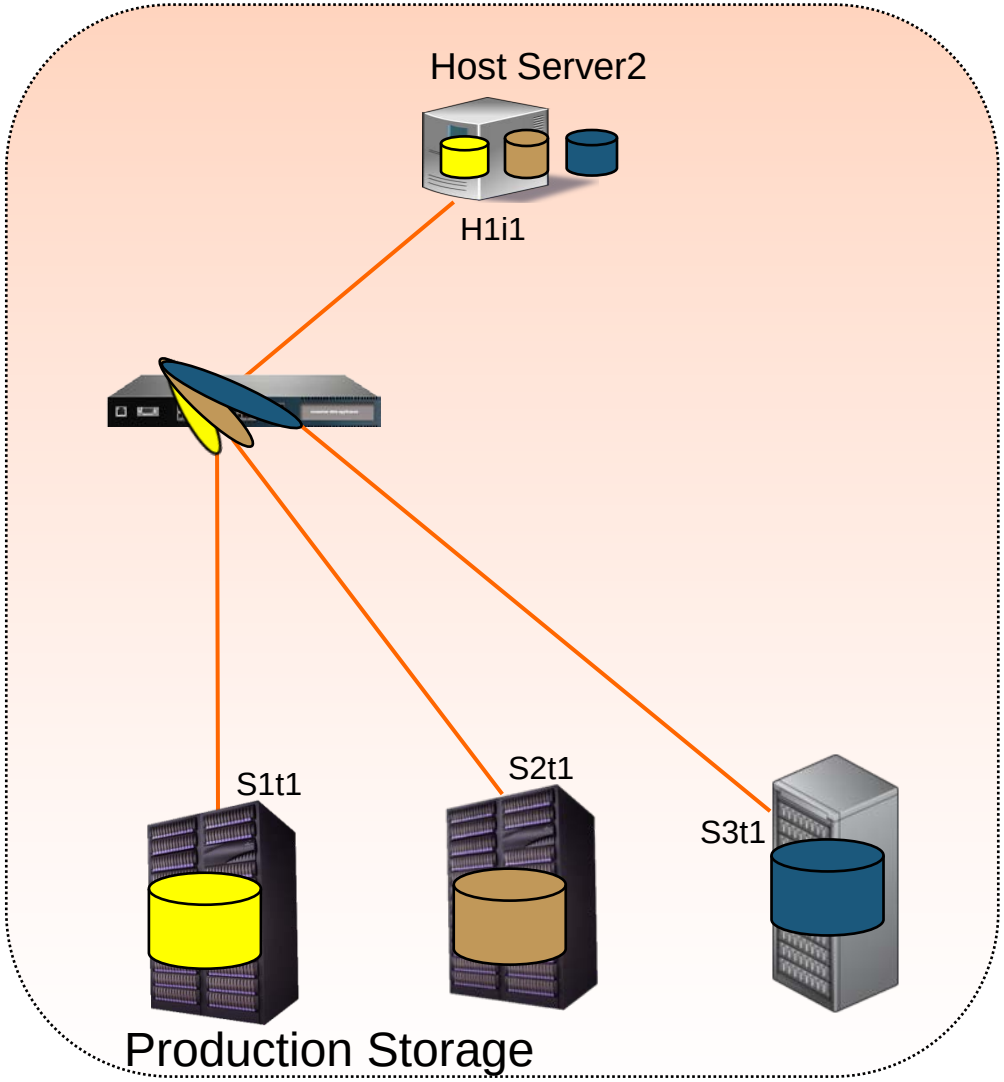

Bad: If you don't physically secure your hosts, switches, cable panels, storage, then You are vulnerable to Port Spoofing attacks.

FC Port Spoofing: The Bad and the Good

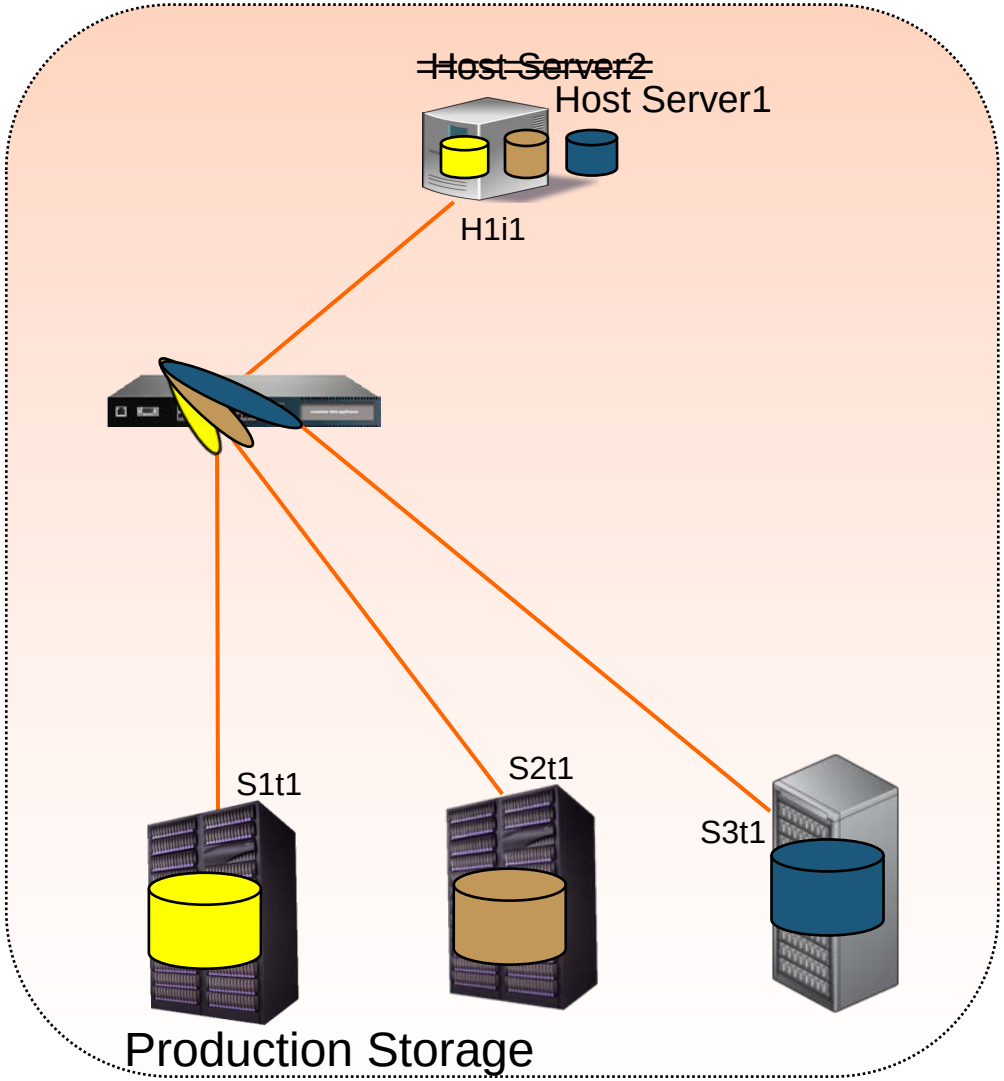

Bad: If you don't physically secure your hosts, switches, cable panels, storage, then You are vulnerable to Port Spoofing attacks.

Good: Port Spoofing allows for minimum work in replacing a defective Host Server or FC HBA.

It also enables creative ways to physically intercept the data path without requiring any changes at the host, the switch, or the storage!

For example: Transparent Datapath Intercept

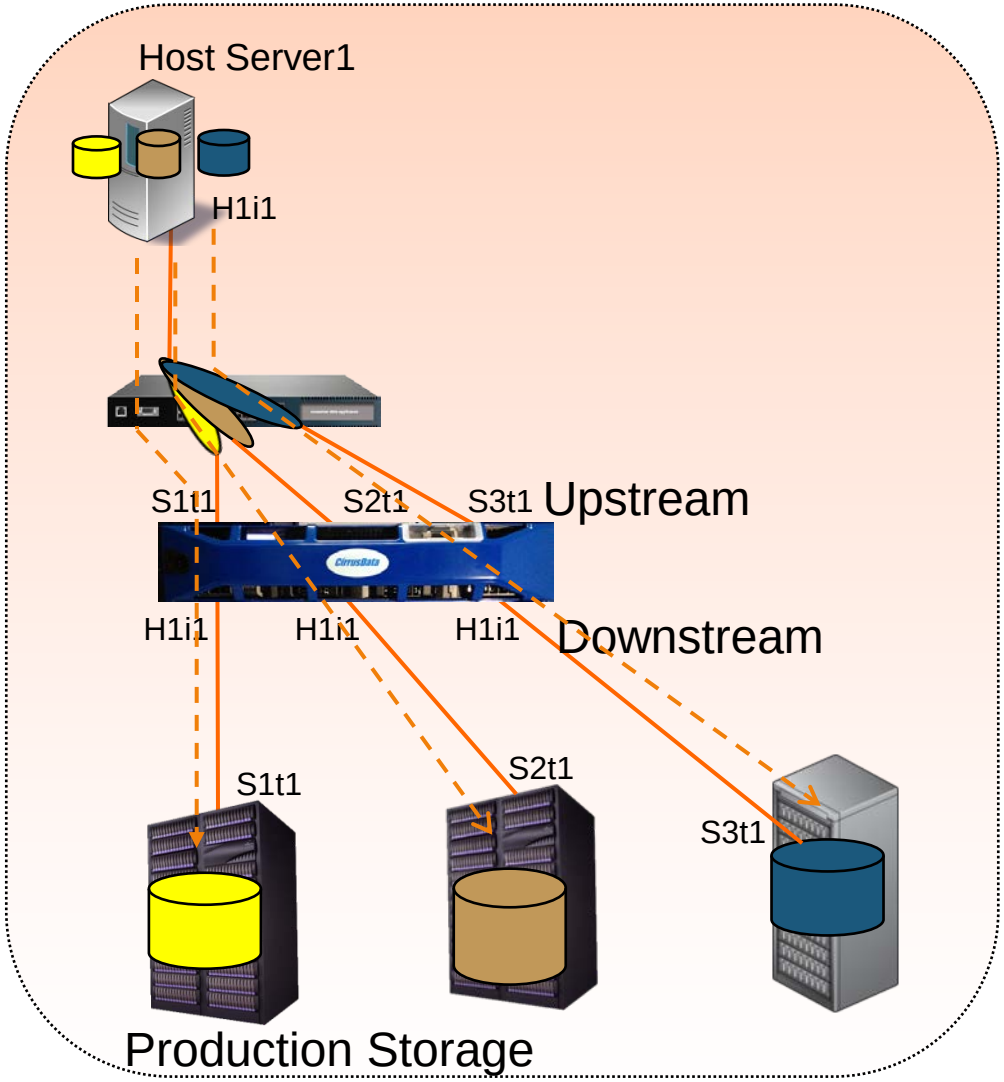

What if we can take advantage of WWN Spoofing to do something good:

An appliance in the middle with an Upstream and a Downstream FC port.

Double spoofing:

Upstream spoofs Storage
Downstream spoofs Hosts

FC Commands are sent through the appliance. Result: Total Transparency!

For example: Transparent Datapath Intercept

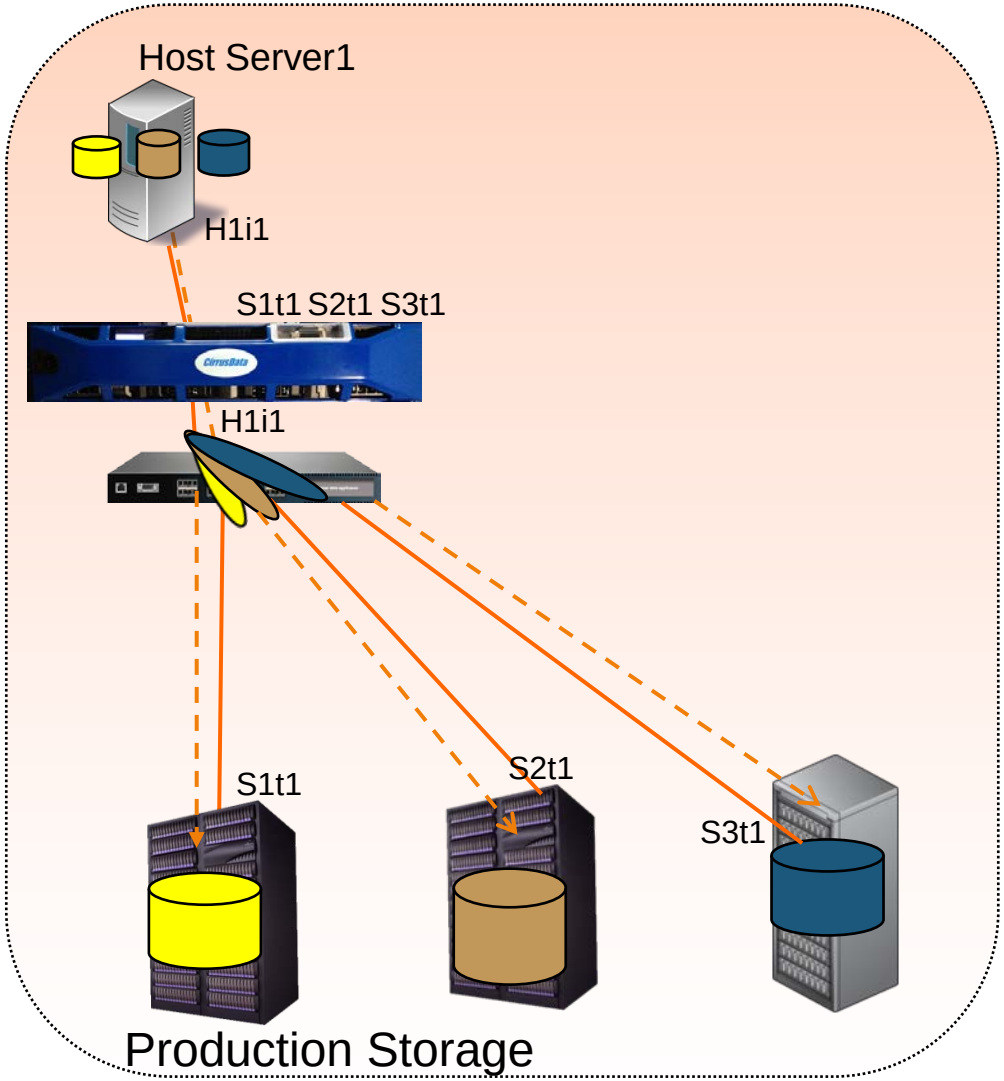

Variation:

Host-side Intercept

Typical Data Center SAN (Storage Area Network)

Very complex, delicate configuration with lots of parts: best left untouched.

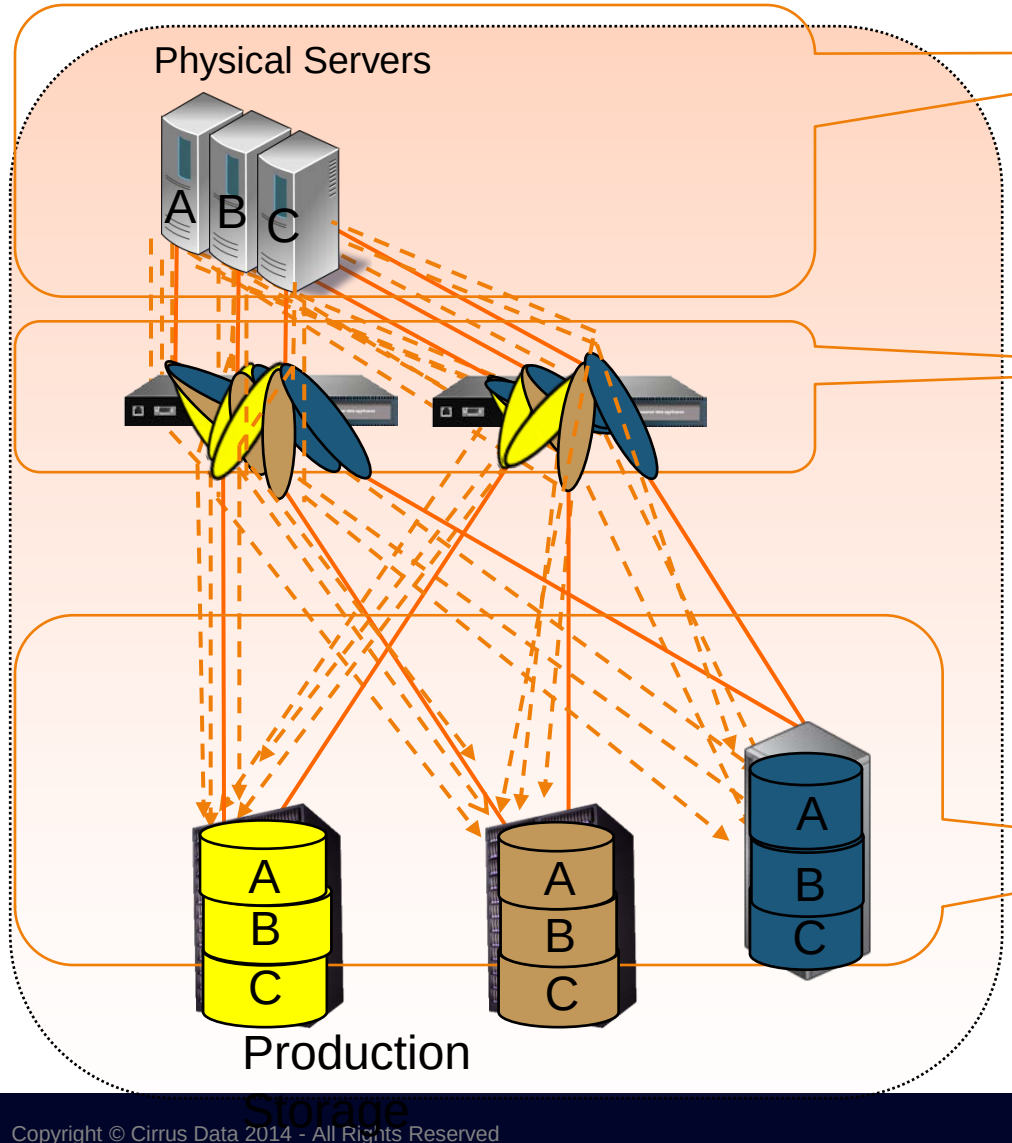

Existing environment has multi-path drivers suitable only for current storage.

Host is zoned to see WWN of current storage target ports.

For M servers and N storage, there is a total of 2xMxN zones.

Each LUN is then LUN-Masked to the intented host for access.

Traditional Data Migration: The Dreadful Changes Insertion by Zoning Changes

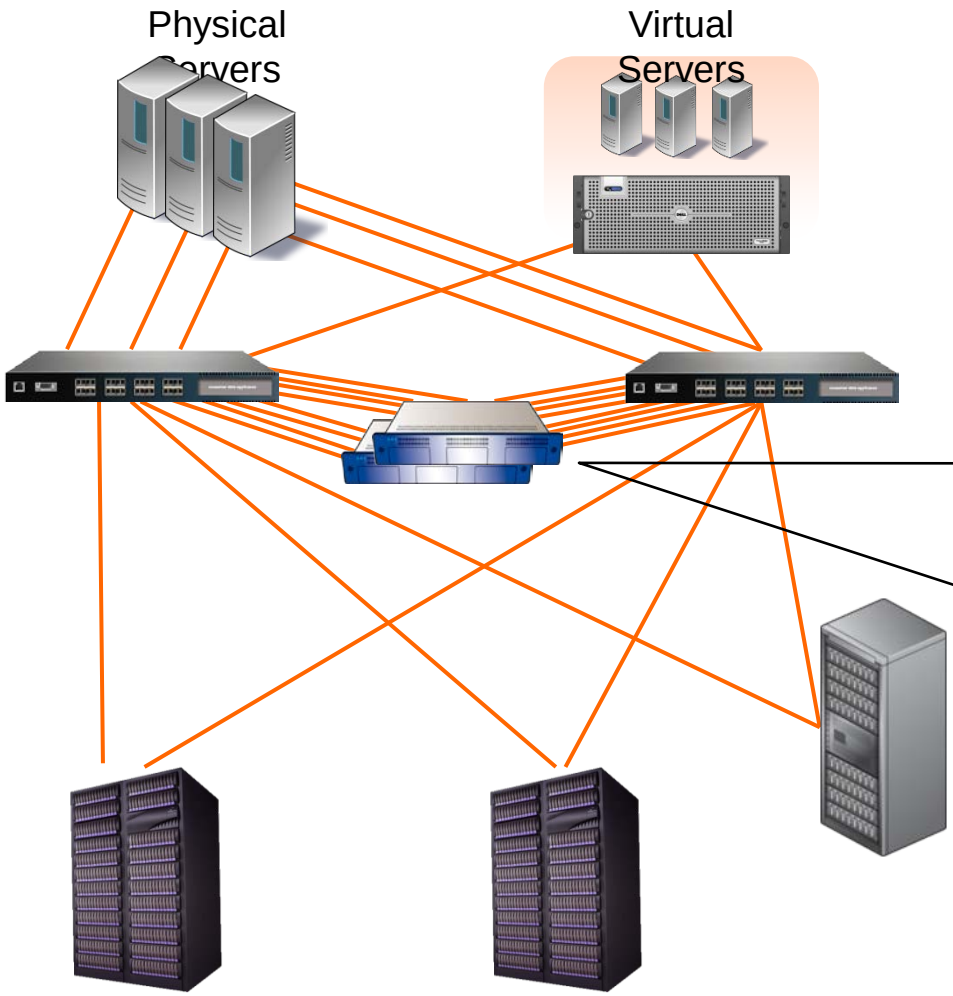

Traditional approach involves inserting a pair of migration appliances into the data path by zone changes and LUN masking changes.

Traditional Data Migration: The Dreadful Changes
Look what happens to just ONE of the zones

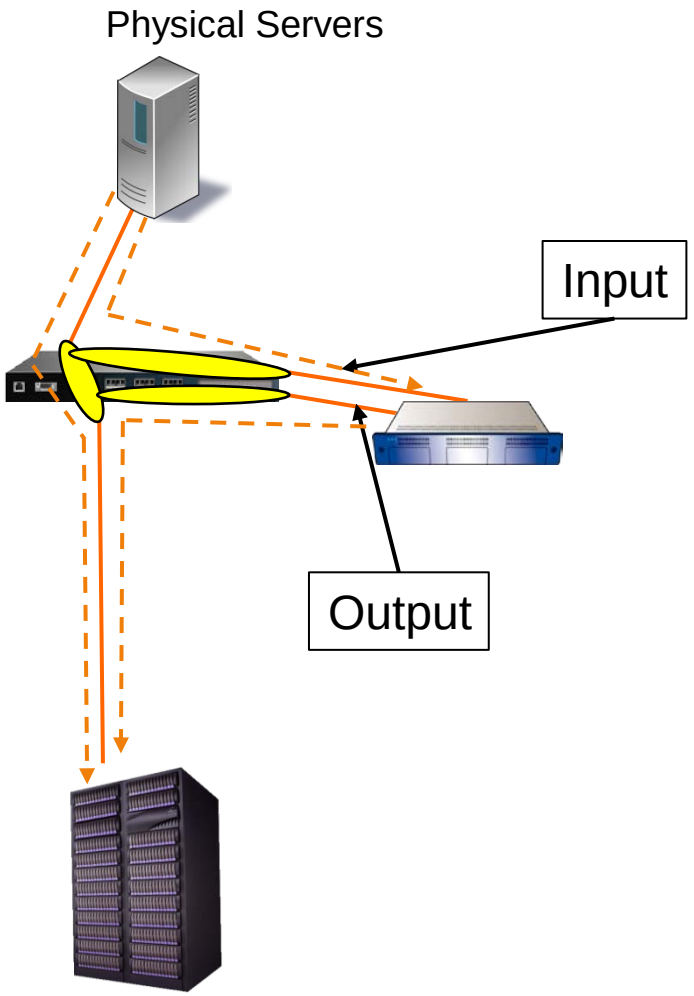

Rezone Tasks

Step 1: Remove the existing zone

Step 2: Create a zone to connect the host to the “Input” port of both appliances

Step 3: Create a zone to connect the Storage to the “Output” port for both appliances

Traditional Data Migration: The Dreadful Changes

Pain points of Appliance Insertion by Zoning

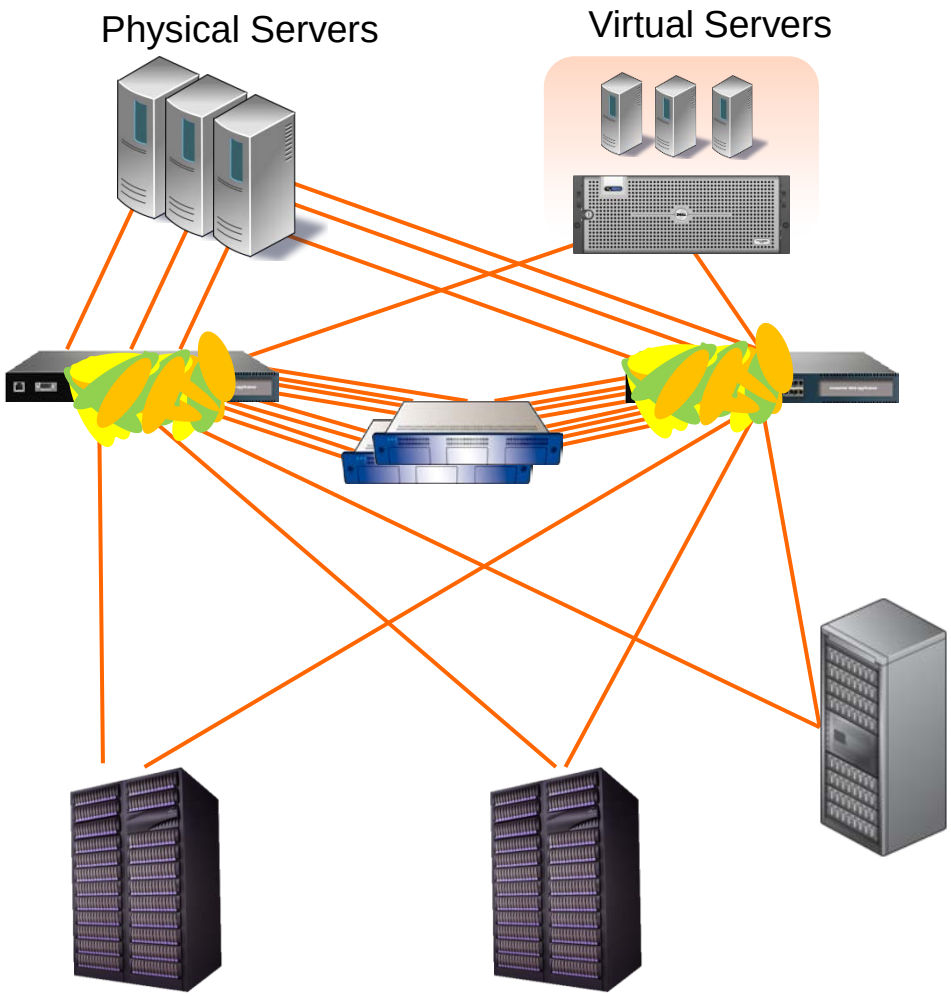

Work and Risks

Multi-path software must be reconfigured at each host: reboots=outage.

FC Switch Rezone: For a 4 Server x 3 Storage SAN, needs to delete 24 zones, create 56 new zones! Mistakes can cause data loss!

All LUN Masks must be changed to be presented to the virtualization appliance.

Extremely disruptive, risky and complex procedures!

TDI: Zero changes, just plug it in!

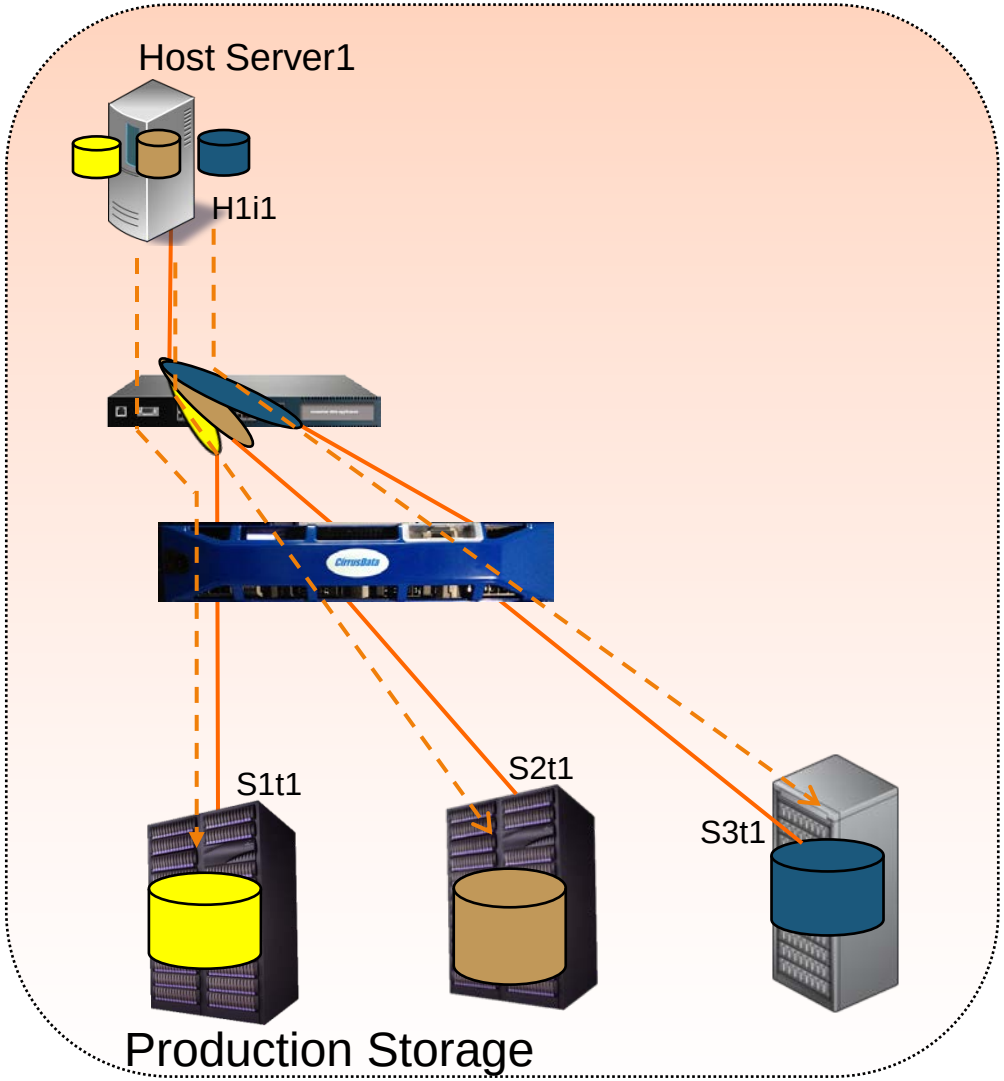

Very Coooool Usage Cases for TDI:

Zero change, zero-downtime deployment of appliances for: *Data Migration*, *Data Acceleration*, and *Data Protection* solutions in legacy and cloud based datacenters.

Data Acceleration: Cache! You know you need it!

But where?

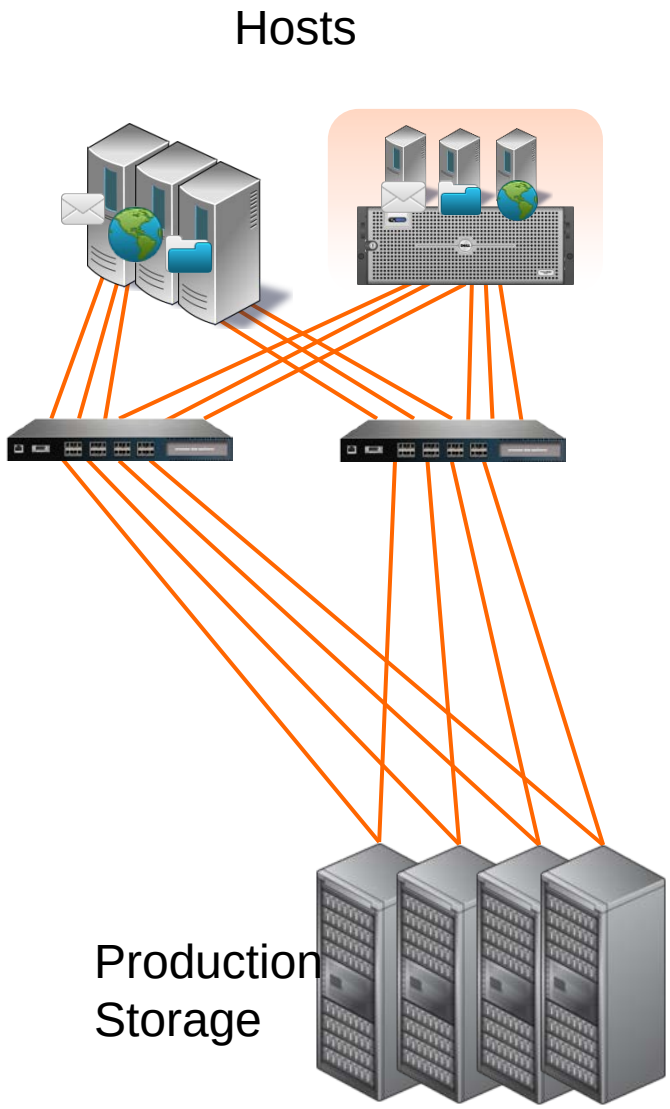

Three Layers in the SAN for Cache

Host-based

SAN-based

Storage-based

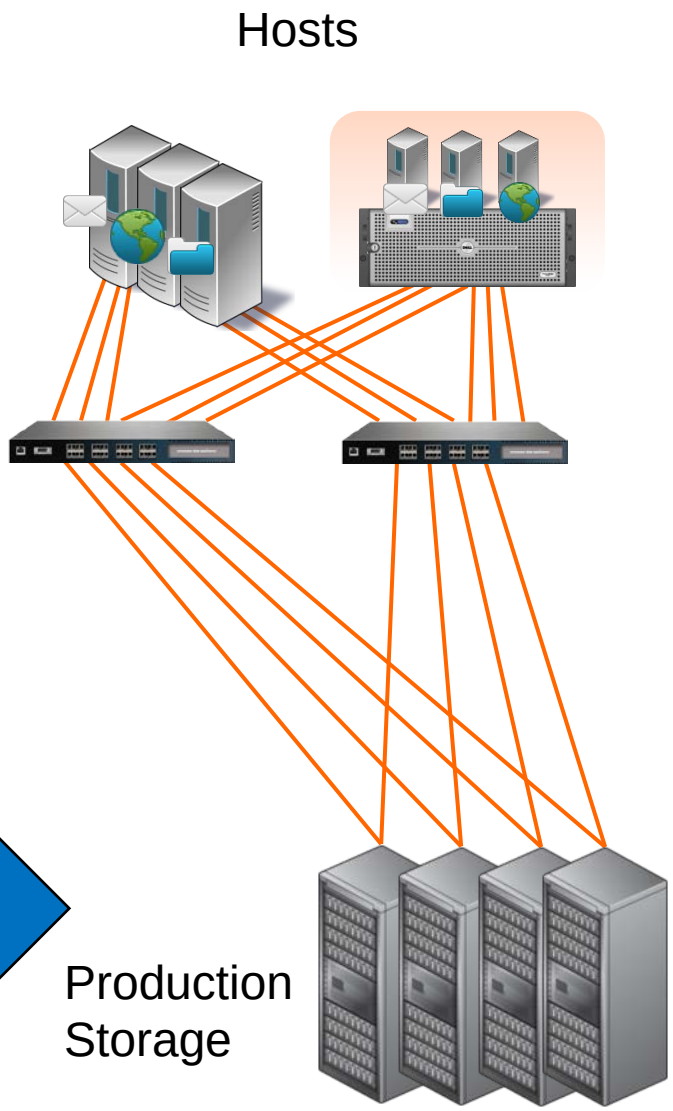

Host-based Cache

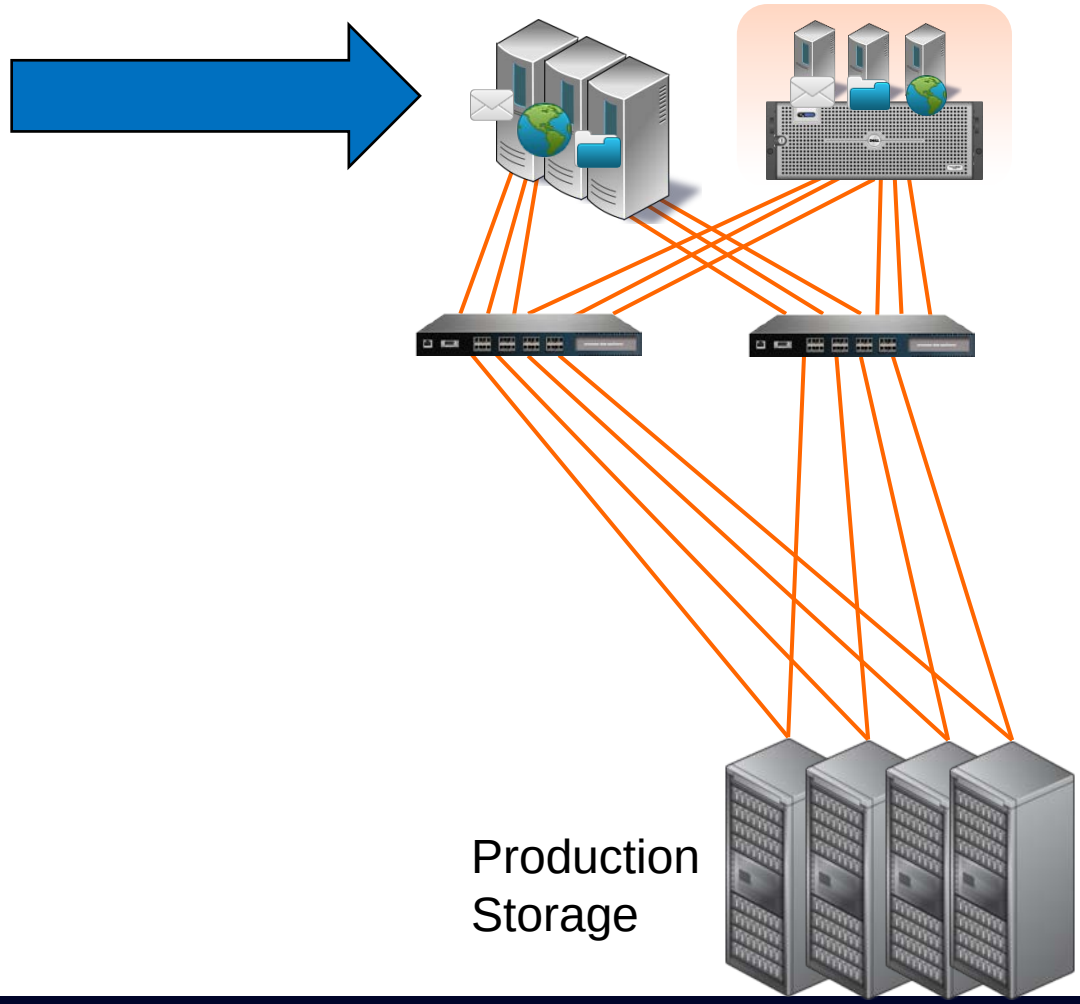

Host-based Cache

- The best place to cache in terms of performance
  - local PCIe bus speed delivers best IOPS and MB/s
- Can cache all downstream storage
- There are cases where it may be impractical:
  - No down time available
  - Hardware Compatibility issues
  - Does not work with Clusters
  - Wasted Cache due to inefficient allocation

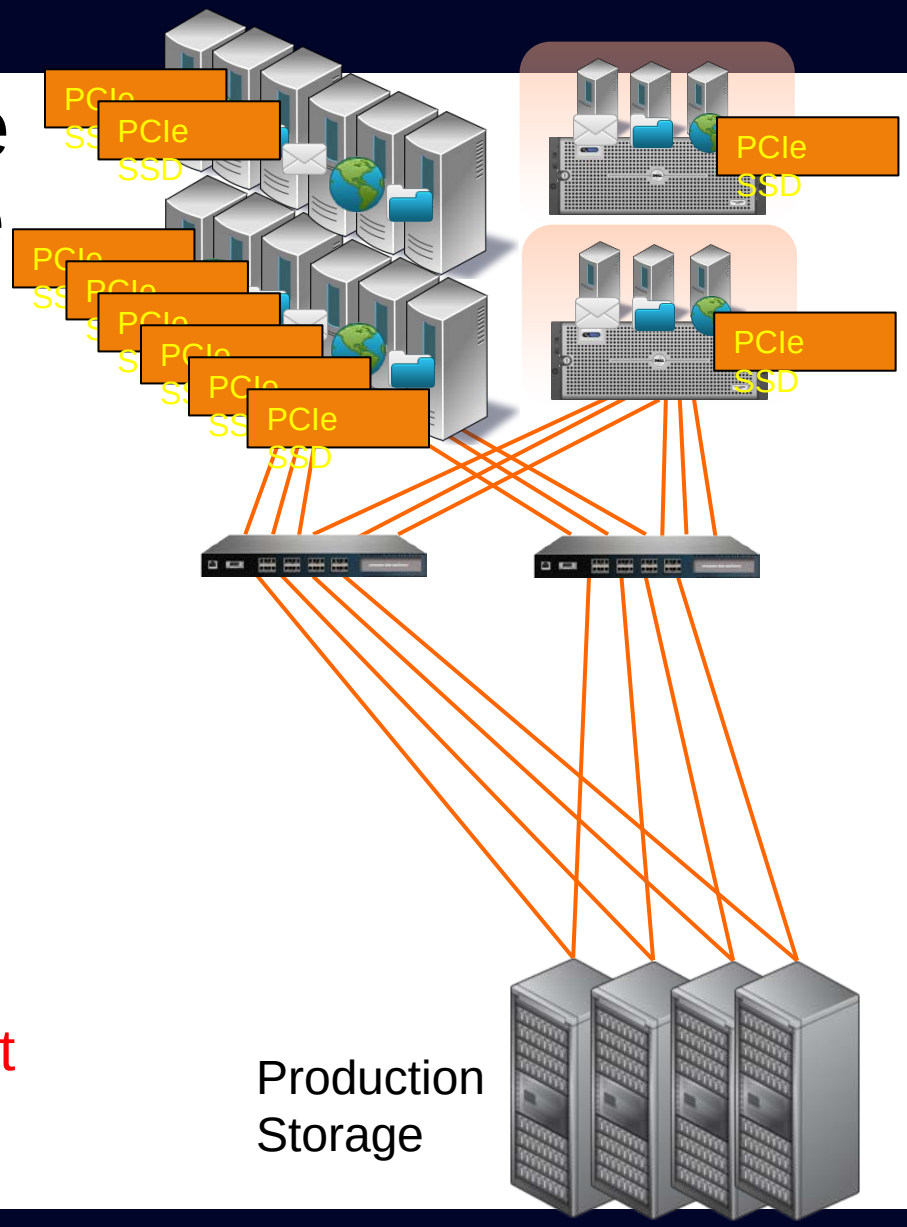

Storage-side Cache

- Shared by all hosts
  - Cluster friendly
- No work at the host
- No work at the FC switch

- Lock-in by storage vendor! $$
- Not shareable by all storage

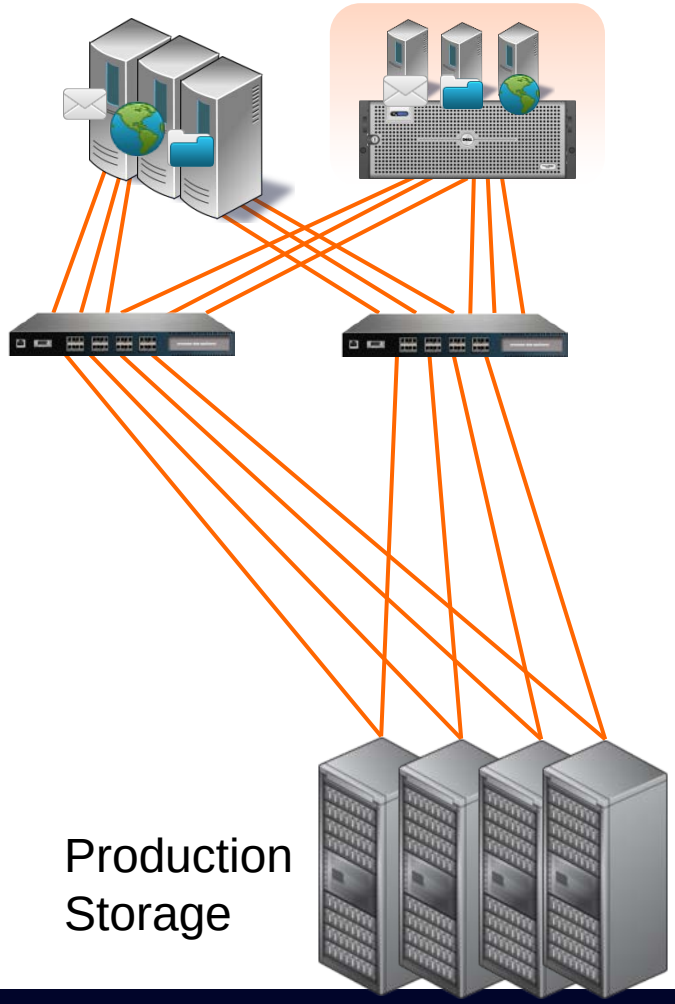

SAN-based Cache

- Best of both world!
- Shared by all hosts and storage
- No work at the Host side
- No work at the Storage side
- Cluster Friendly
- Most strategic place for Cache

- How to address the Complexity and Risks of Changes to SAN?

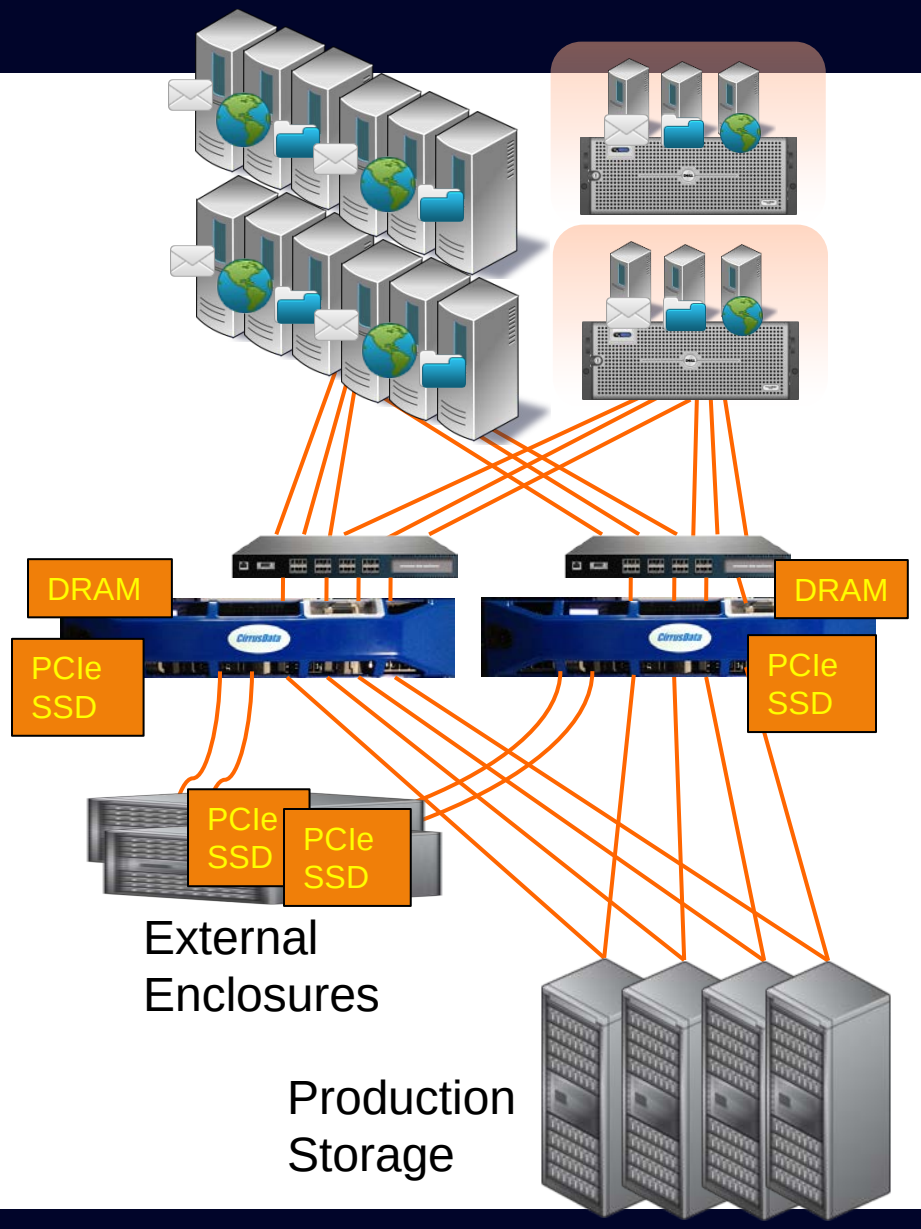

SAN-based Data Caching Server (DCS)

- Transparent Datapath Intercept (TDI) allows for Zero Change, plug-and-play deployment of DCS appliances in a live SAN
- Paths are inserted one at a time, 5 seconds per path, just like inserting an answering machine
- DCS Appliances auto-detects everything.
- Eliminates all the work and risk associated with In-band deployment.

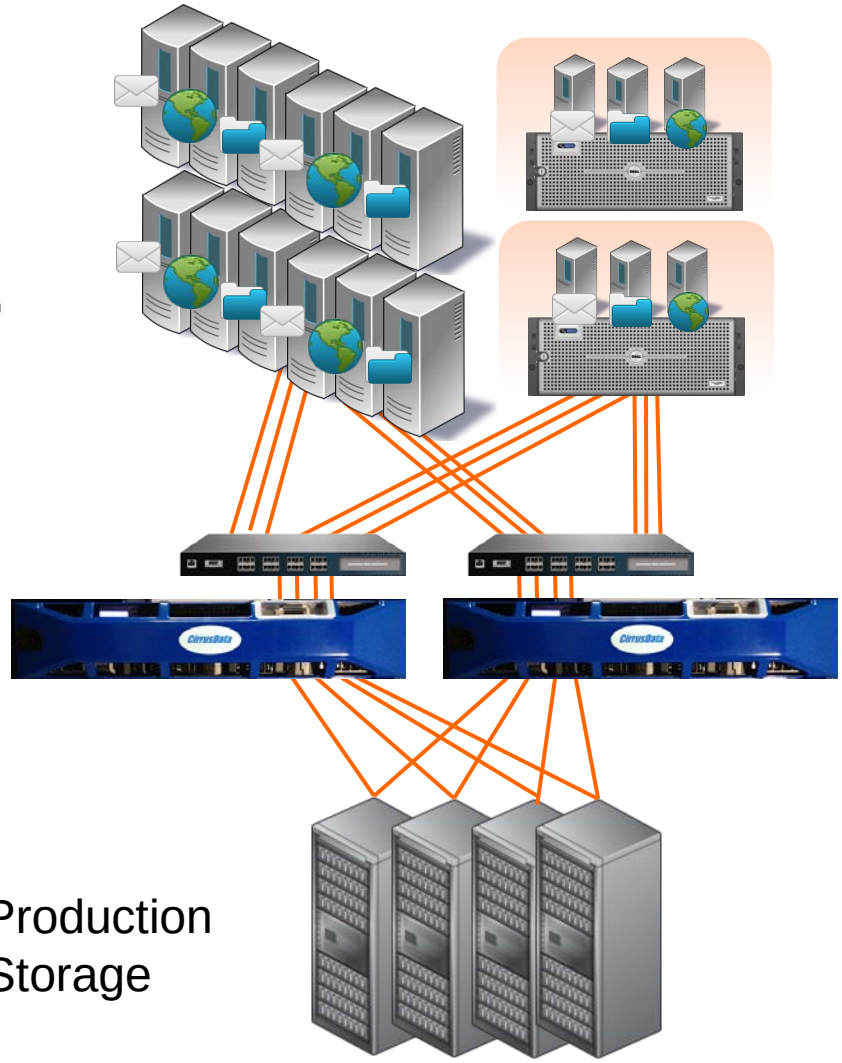

DCS Deployment Without Touching Original Physical Connections

- For Large Environment Mandating “No Touch” policy on existing ports and paths
- Connect to unused ports on the Switch and Production Storage
- For each set of LUNs being intercepted:
  - Create new zones
  - Present LUNs and activate new path
  - Unpresent LUNs from original path
  - Repeat for another path

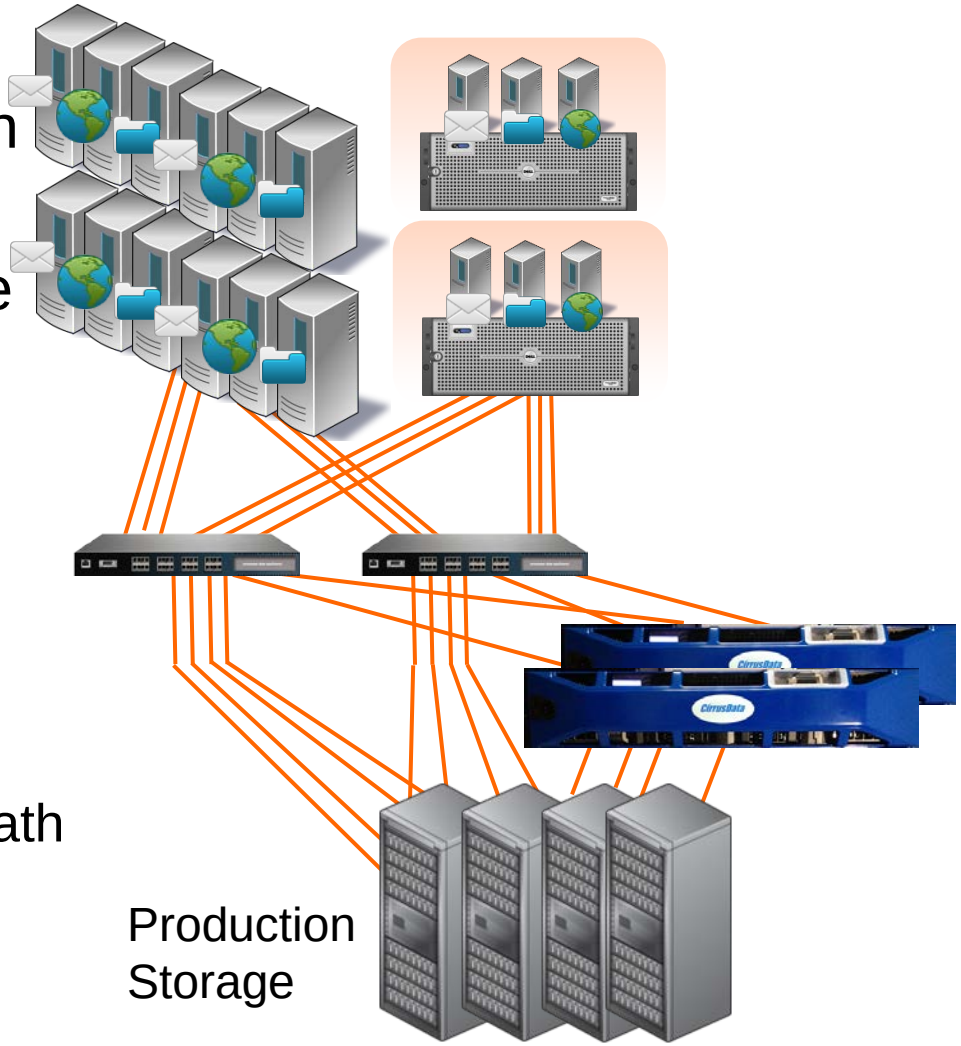

Data Caching Server (DCS) Appliance

The DCS Appliances Advance Features:

- **Zero downtime** data path insertion: just plug it in and cache
- **Zero reconfiguration** of existing hosts, FC Switch zones, and Storage
- **Plug and Play integration** to production environment, live
- **Auto Discovery:** Hosts, Storage Targets, Disks, FC Switches, and the entire SAN topology.
- **I/O Tracking:** All I/O activities and performance data are tracked, with reports made available for analysis.

Actual photo of a DCS-4000 with: 8x8Gbps FC Ports, Seagate Nytro PCIE SSD (up to 20TB per cluster pair)

DCS: The Plug-and-Cache Appliance

- Easy: Plug-and-play appliances deployed in live environment without downtime. Uses DRAM and Internal/External SSD
- Simple: No re-configuration of Hosts, FC Switches, Storage. Just use Web-GUI to create/run Cache Policies
- Effective: Increases IOPS thousands of times. Analyzes I/O of all discovered LUN (latency, queue) and applies appropriate caching scheme for best results.
- Safe: Clustered operation, enterprise-class hardware, no single-point of failure.
- Optimal: Centralized, shared cache. Best place to have all the I/O intelligence for optimized caching.

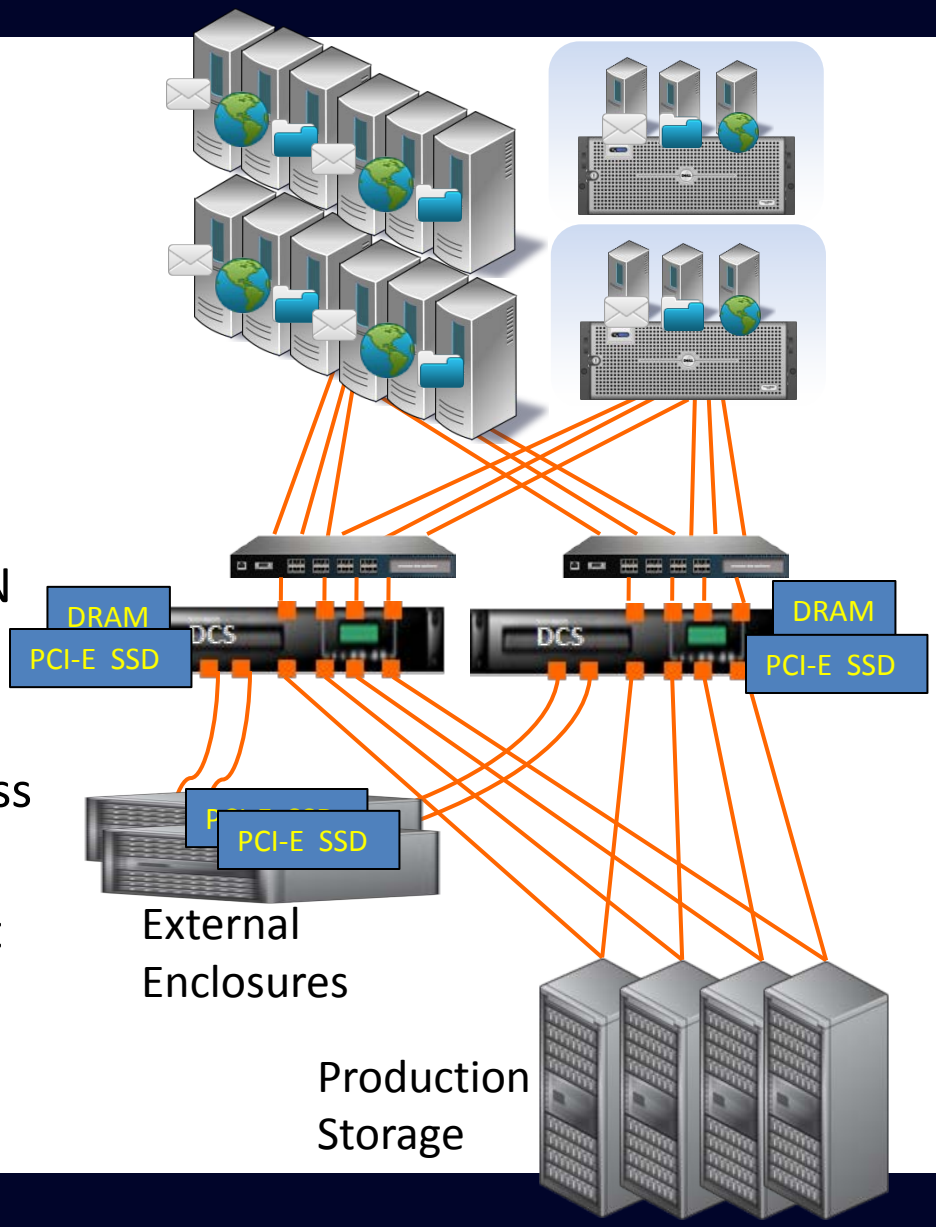

CDS DCS Cache Resources

- Manages one or more Pools of:
  - PCI-E SSD
  - FC Attached SSD
  - SAS-Attached SSD
  - RAID with 2.5”/3.5” SSD
- Built-in DRAM is always used as first-tier cache resource

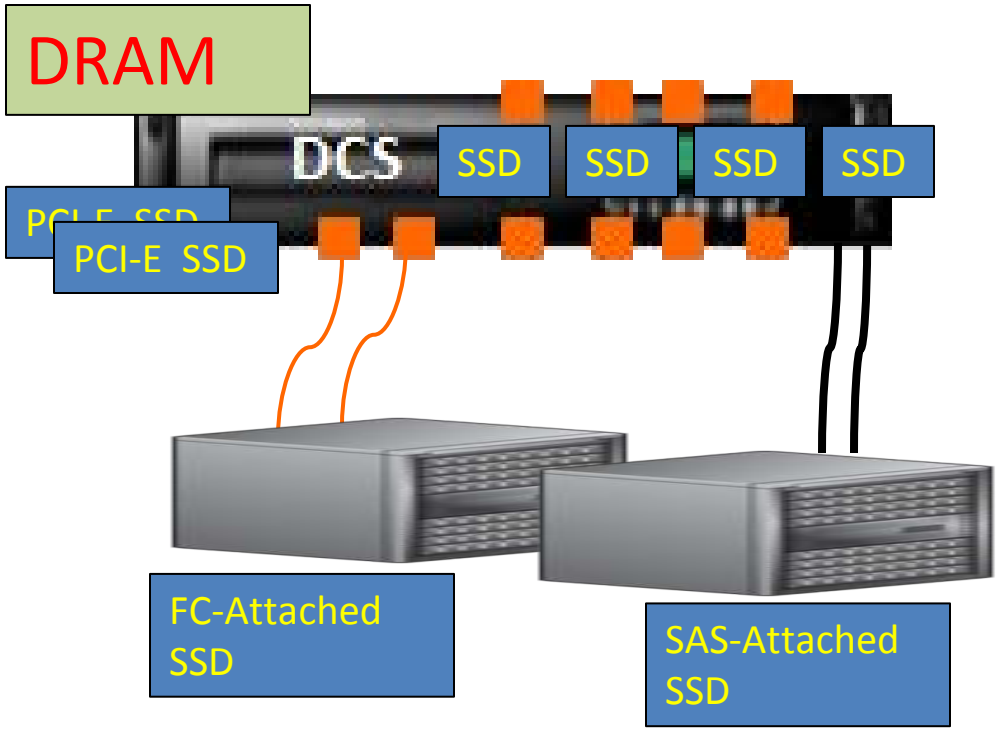

CDS DCS Cache Resources

- Manages one or more Pools
  - PC
  - FC
  - SA
  - RA
- Built-
  used
  resou

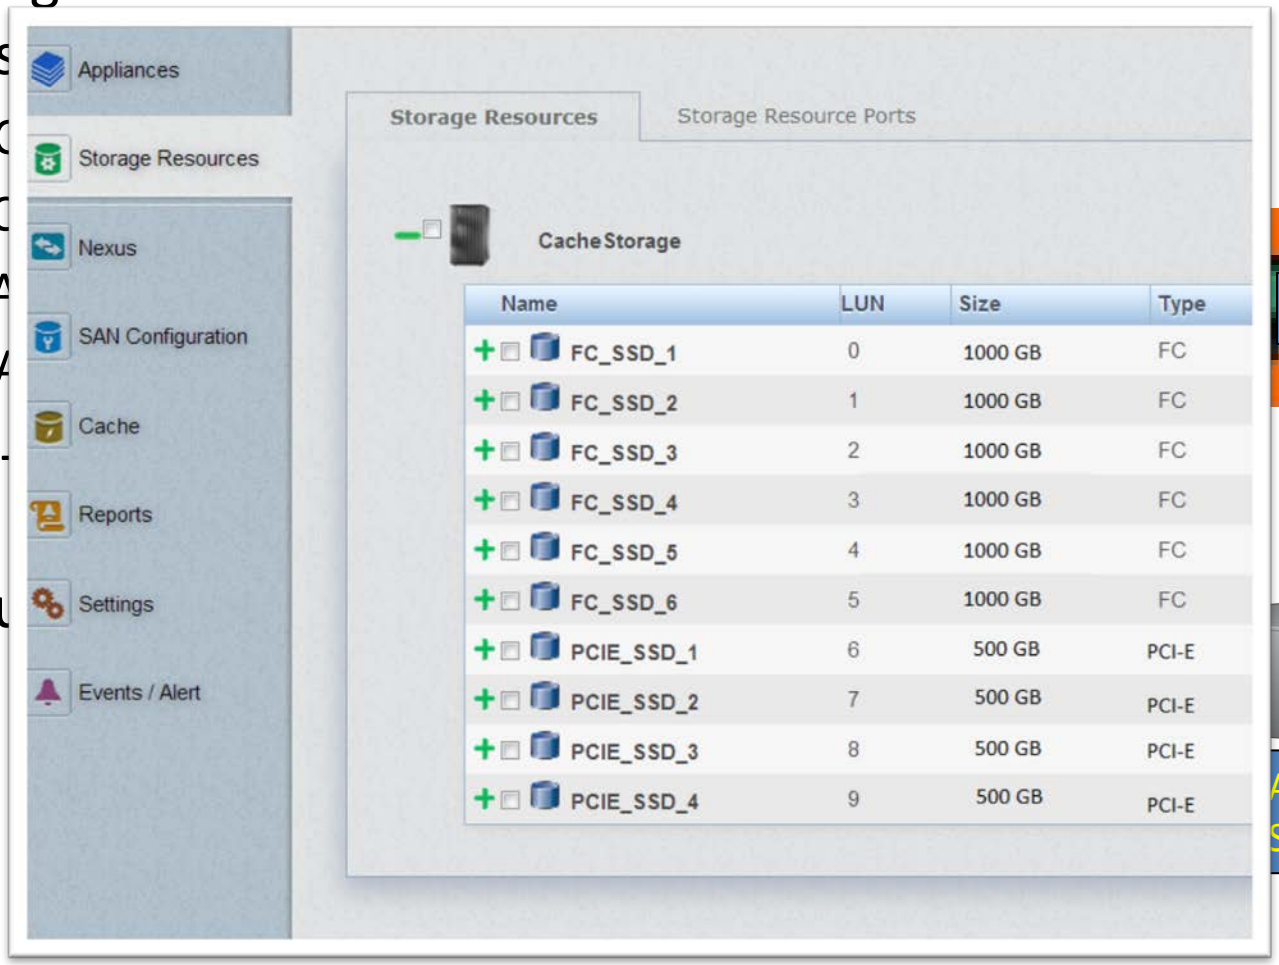

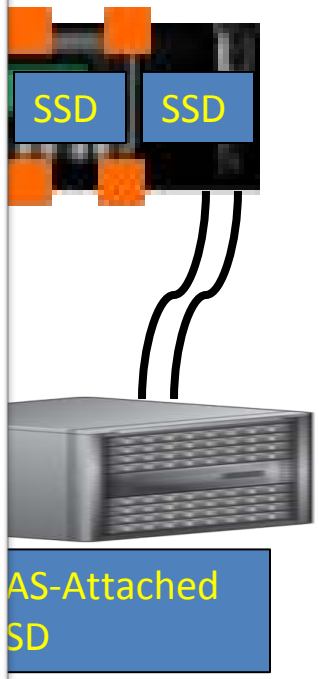

Pool-based Resource Management

- Availabe Cache Resources are organized into Pools (Resource Groups)
  - By performance characteristics
  - By hosts
  - By storage
  - By application

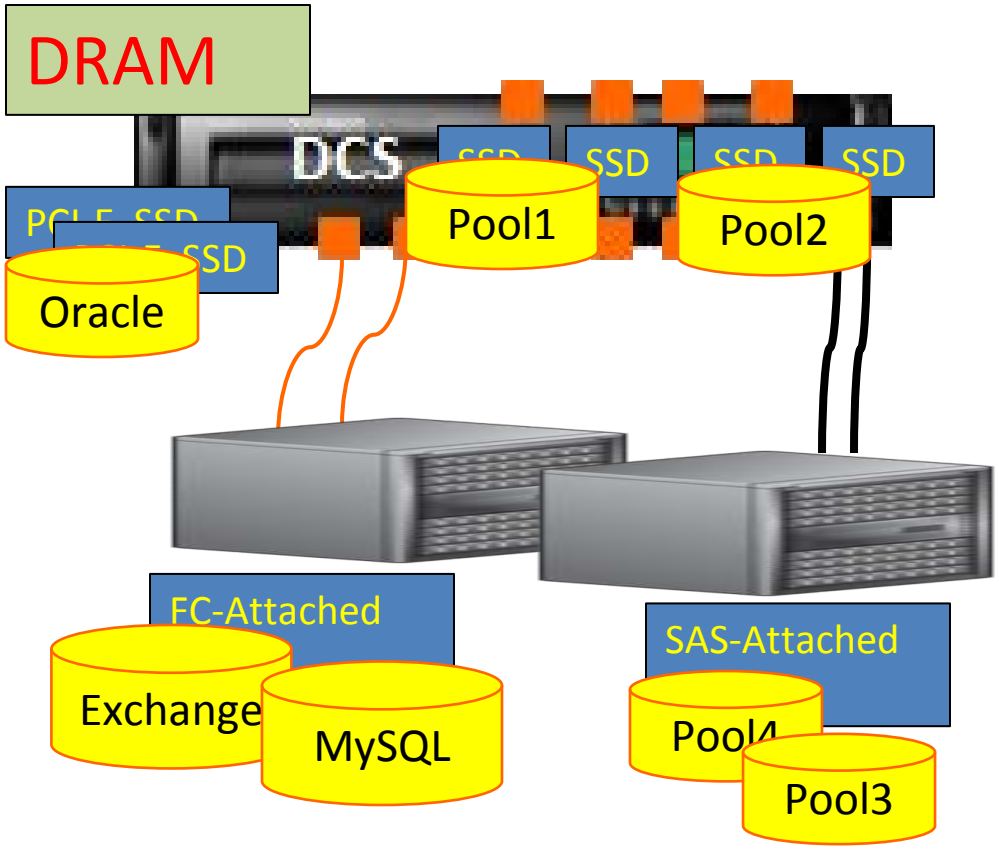

Pool-based Resource Management

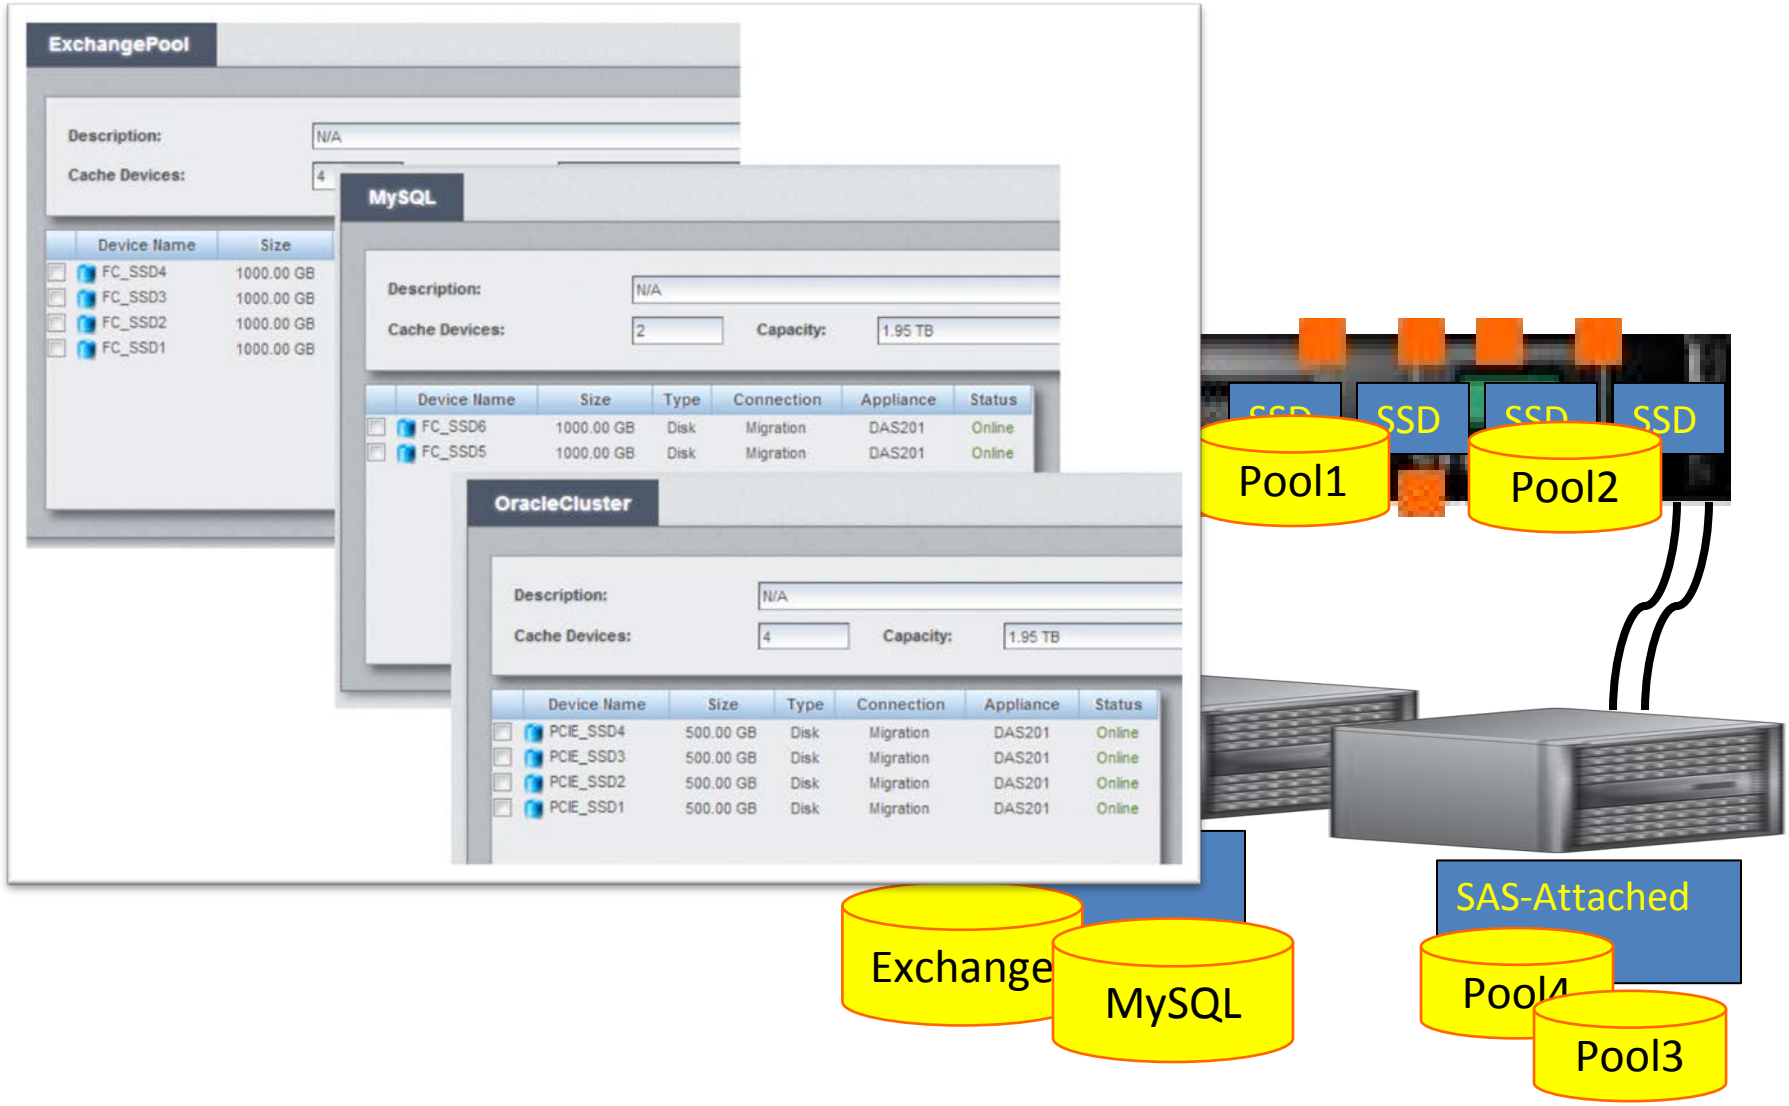

DCS Cache Policy

- Create Cache Policies for one or multiple application LUNS
  - Assign a set of Application LUNs
  - To a Resource Group (Pool)

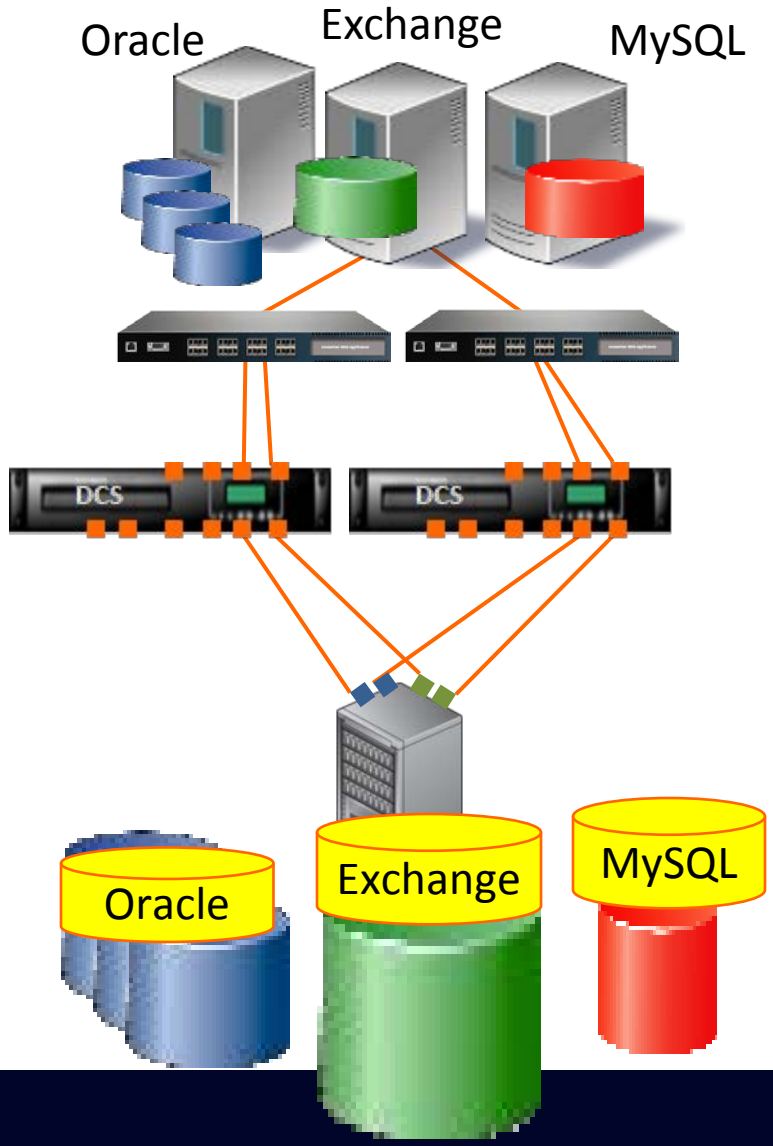

DCS Cache Policy

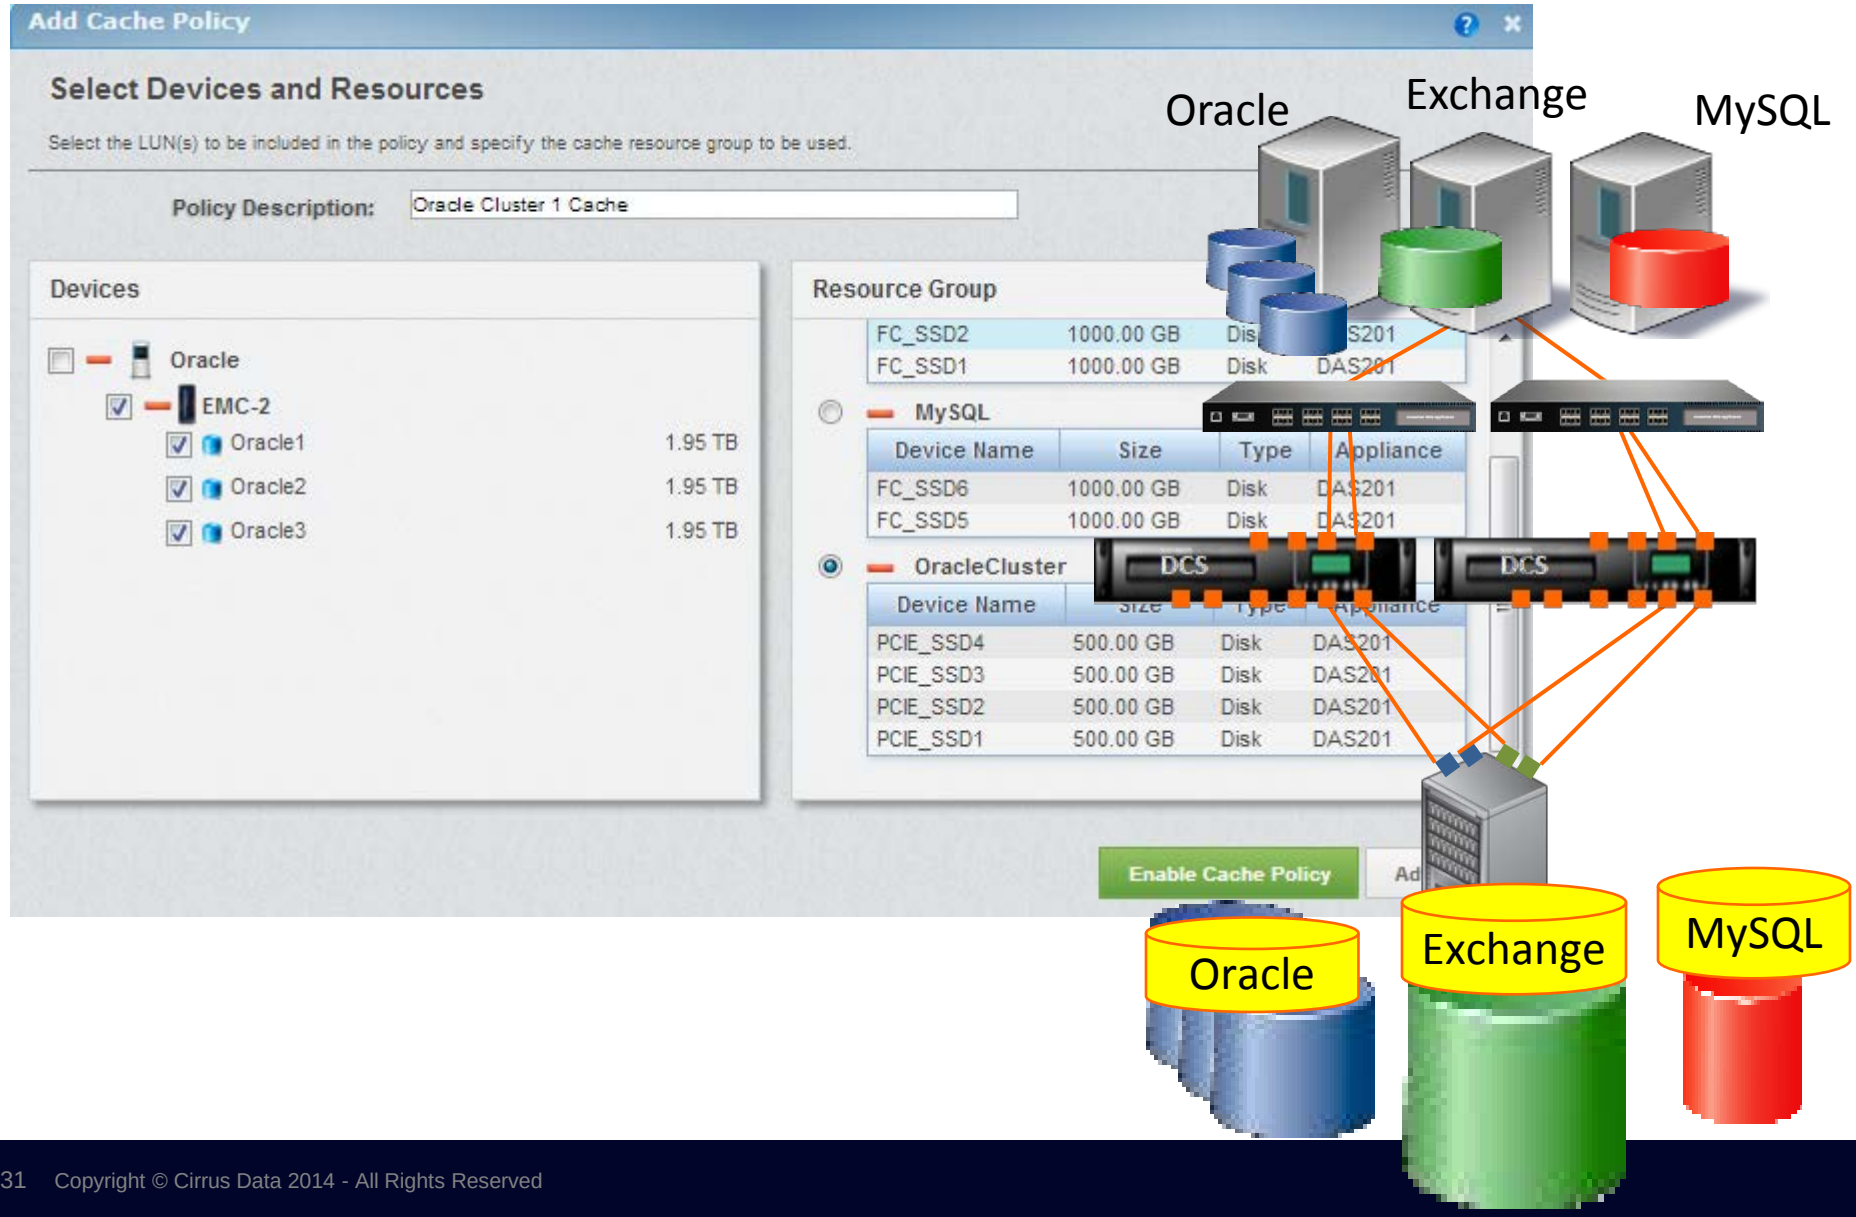

Balanced LFU/LRU Cache Algorithm Optimizes Hit-Ratio

- Each Cache Policy configured with optimized LFU/LRU settings
- “Evolution” algorithm automates the fine-tuning to dynamically optimize settings

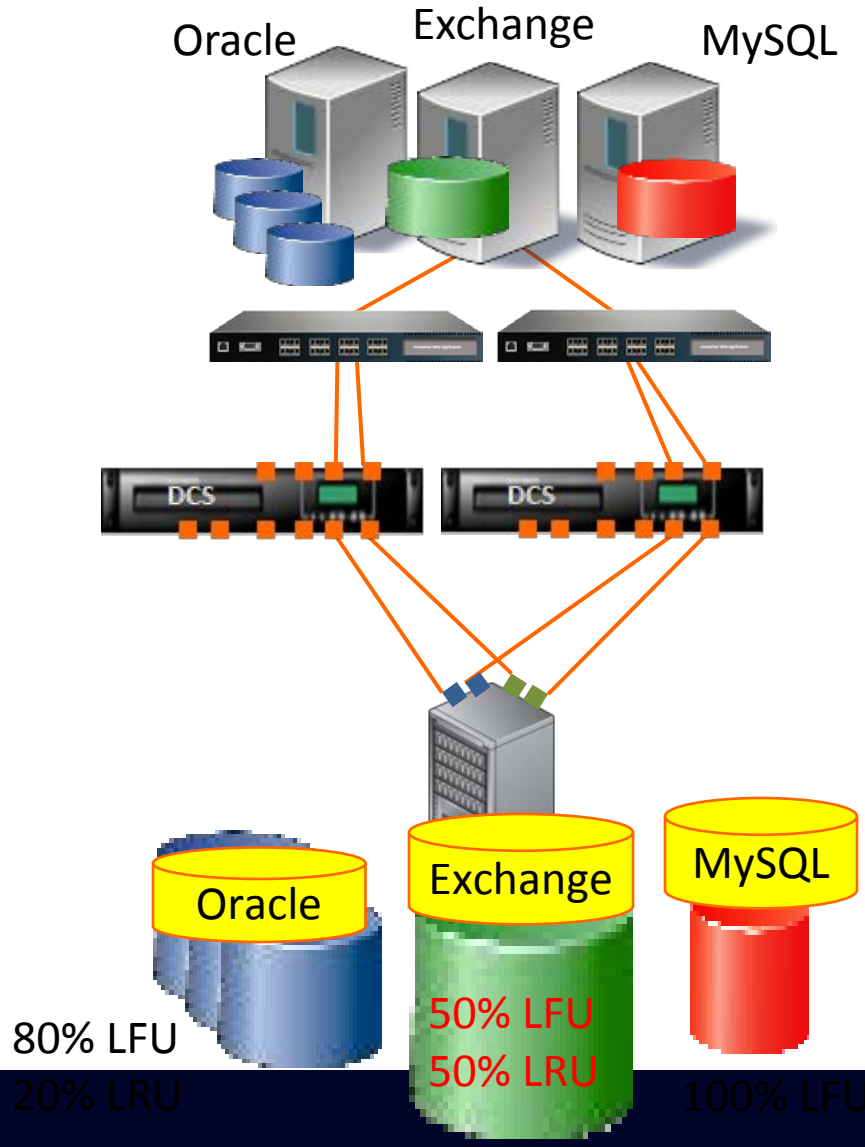

Balanced LFU/LRU Cache Algorithm Optimizes Hit Ratio

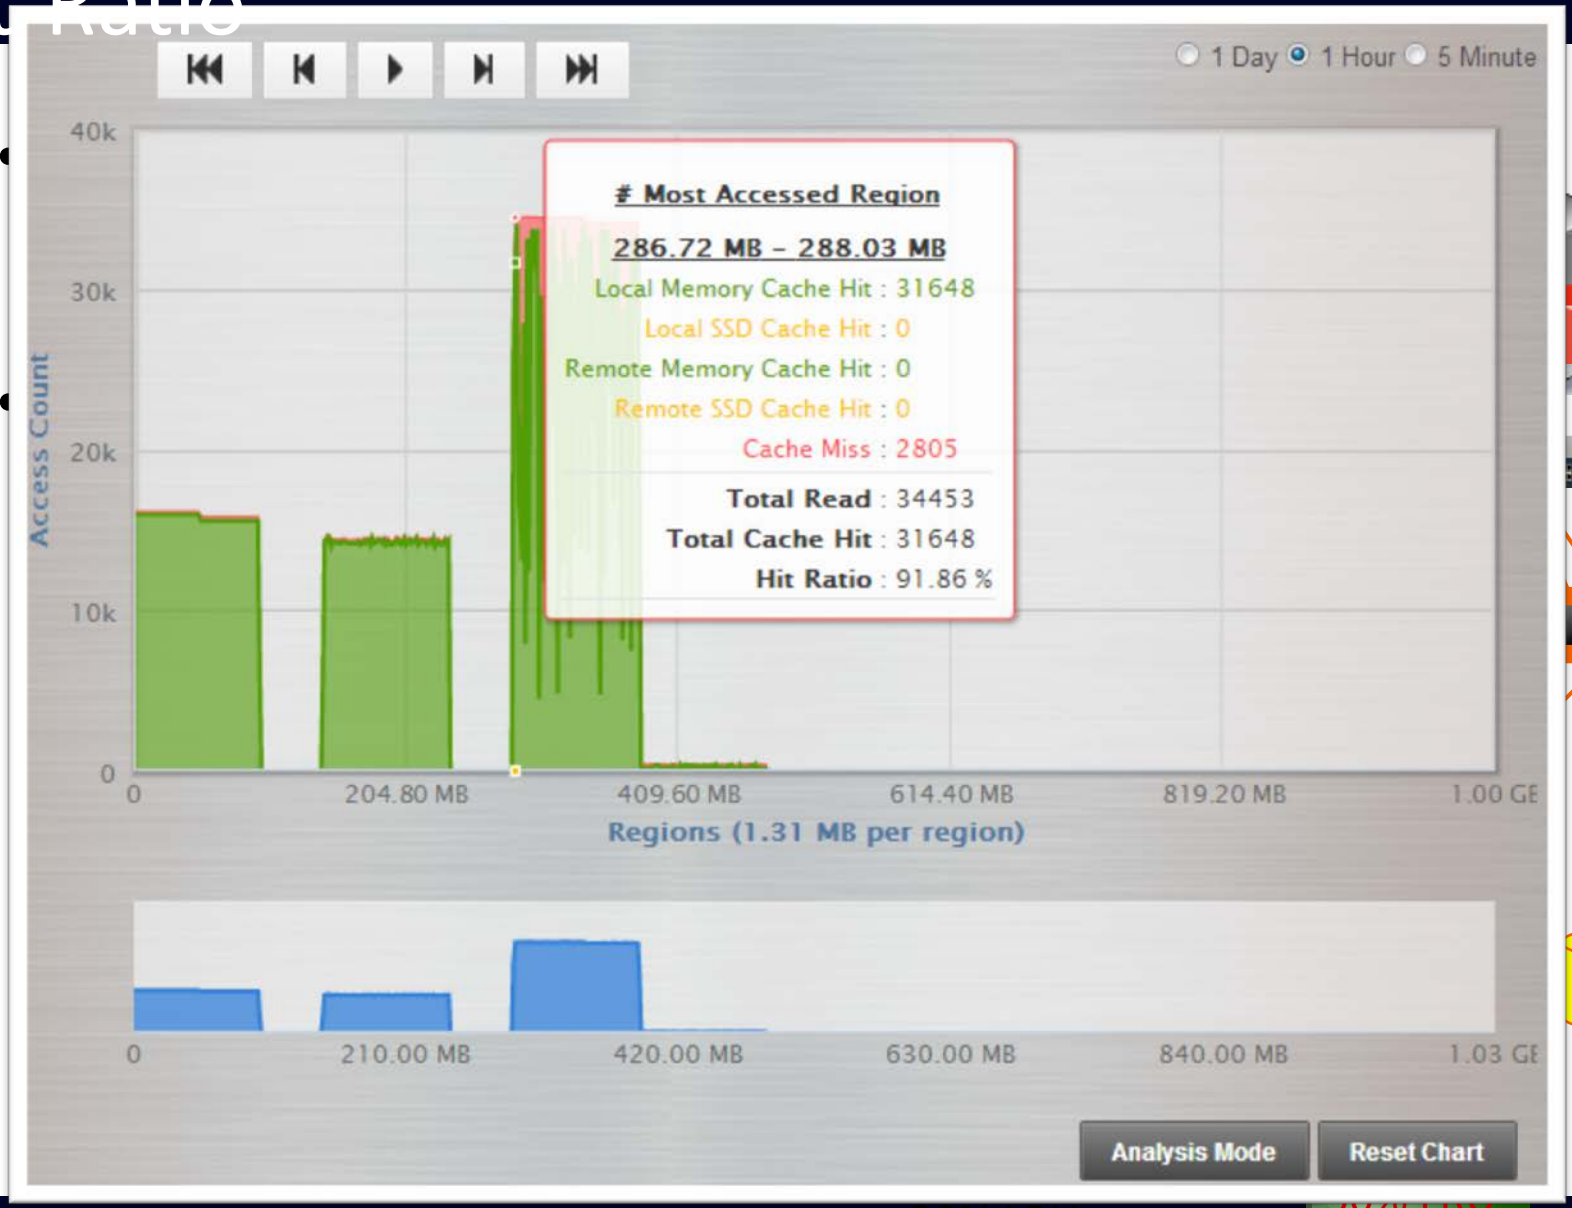

20% LRU

100% LFU

Built-in Analyzer Quantifies Cache Requirements

- Granular tracking of all disk access
- Sort by frequency of access
- Accurate estimation of amount of Cache Storage needed to hold x% of frequently hot blocks

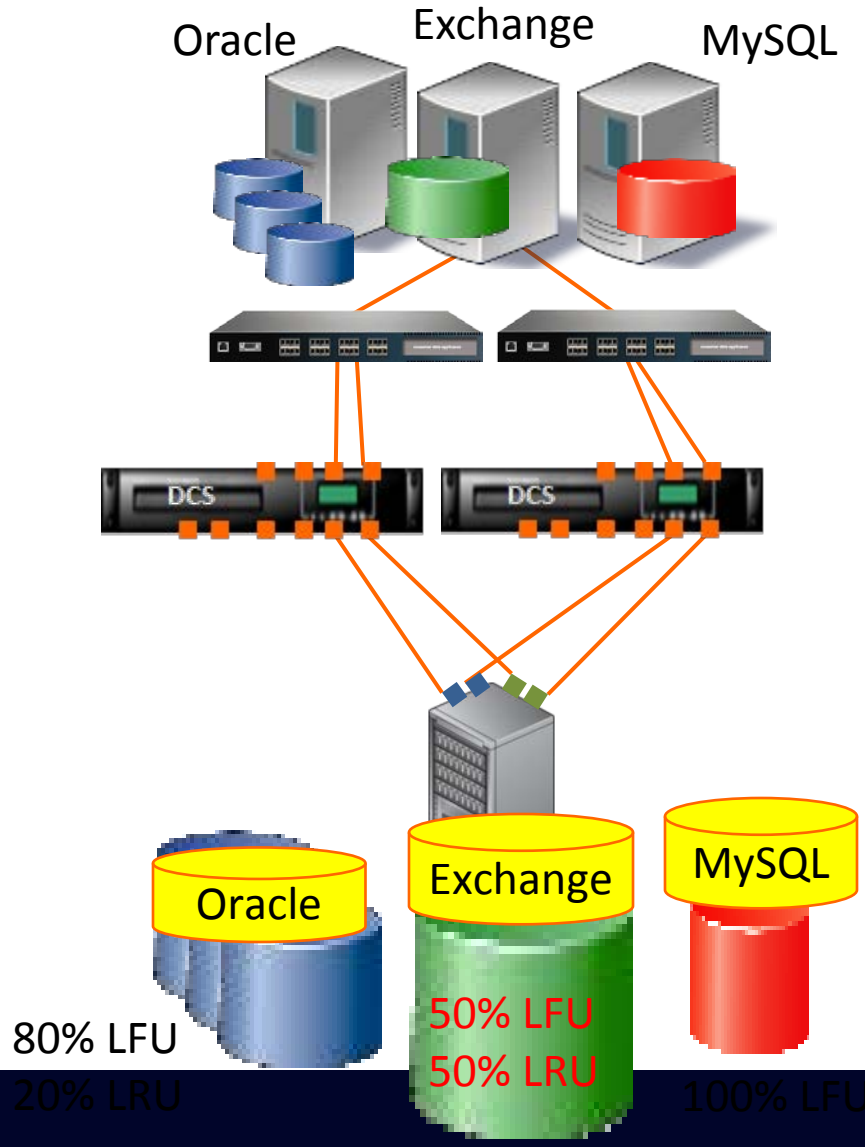

Built-in Analyzer Quantifies Cache Requirements

- G
  a
- S
- A
  a
  n
  fr

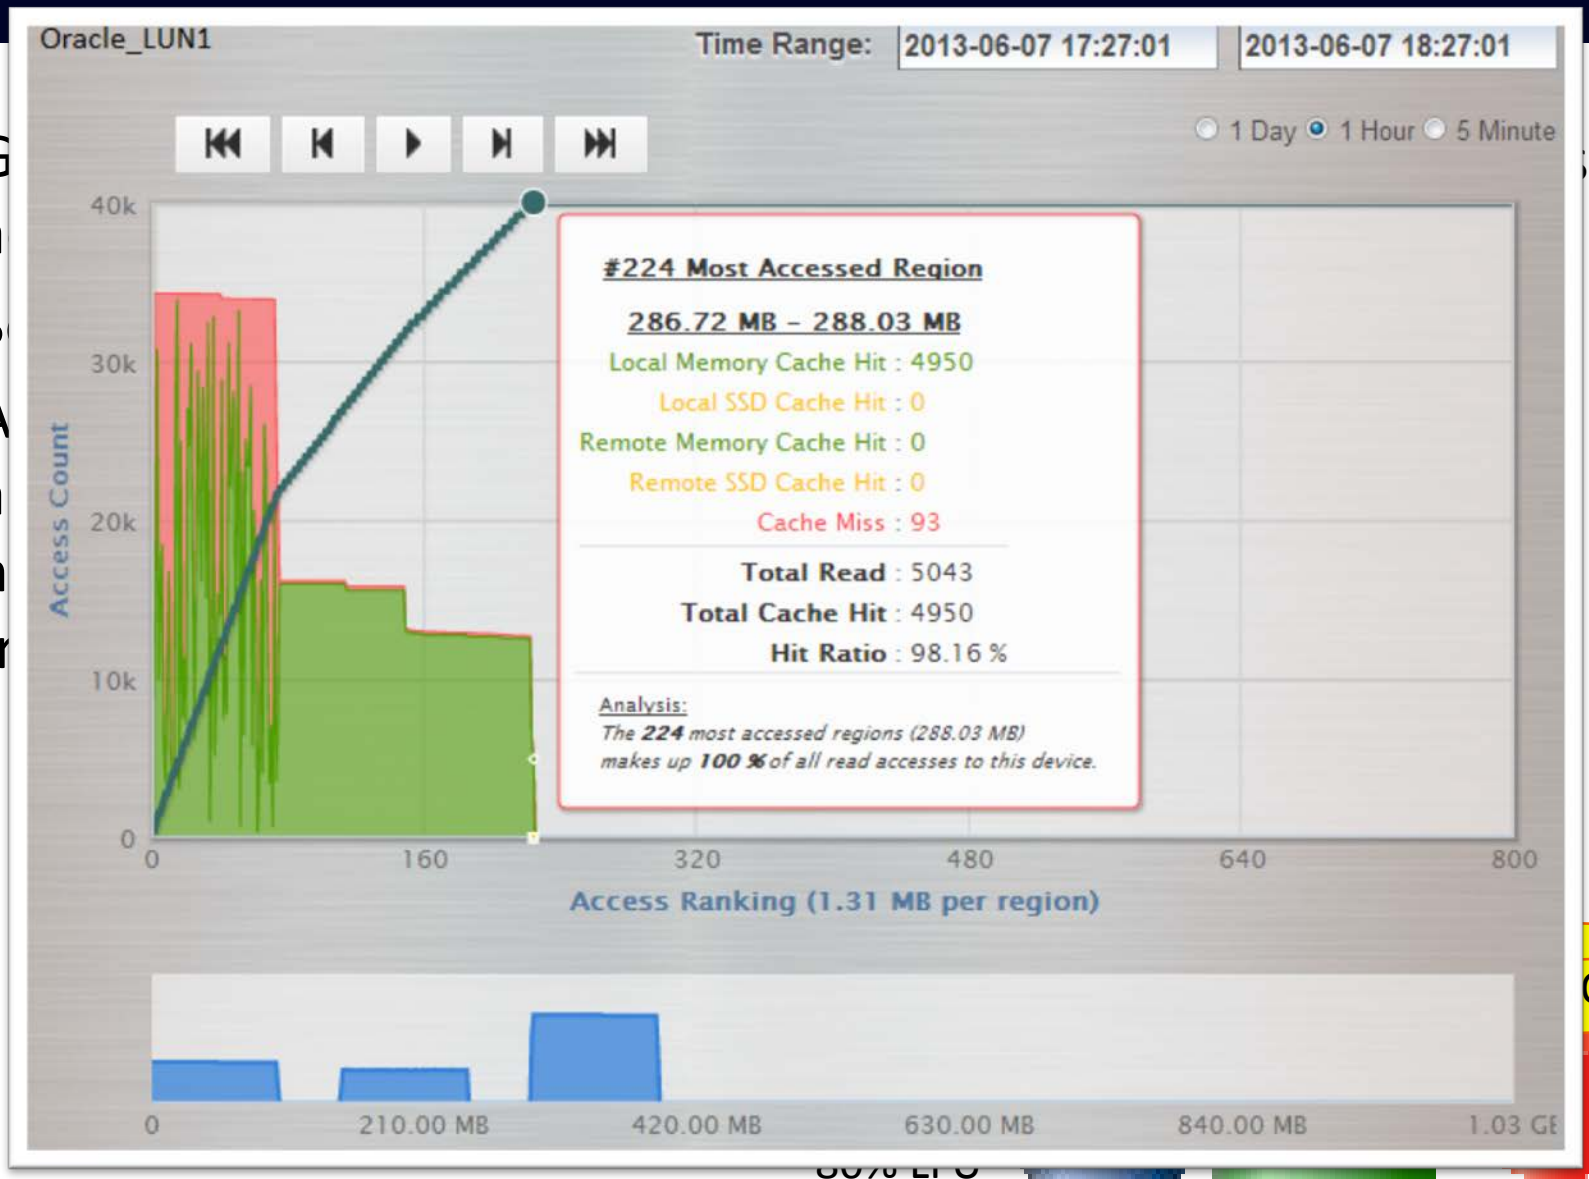

QL

QL

20% LRU

50% LRU

100% LFU

DCS Selling Points Summary

- DCS can be inserted into a production environment with no interruption to clients or storage, and with no user input!
- Insertion is plug and play, based on Patented “Transparent Data Intercept”. No software changes at hosts, no zoning changes at switches, and no changes at storage. Just plug in DCS between the switch and the storage port (or between the Host and switch).
- DCS has built-in I/O analyzer to help decide which disk needs caching, and how much cache storage is desirable
- Active-active DCS clustering eliminates single-point-of-failure. If either DCS fails, I/O continues on the other unit
- DCS automatically discovers and monitors the entire fibre channel topology, and generates I/O activity reports (latency, throughput, IOPS, queue depth).
- Deploy DCS as a permanent cache server, or use it as an emergency cache resource during I/O intensive data crunching projects.

Watch a recorded demo for Plug-and-Cache

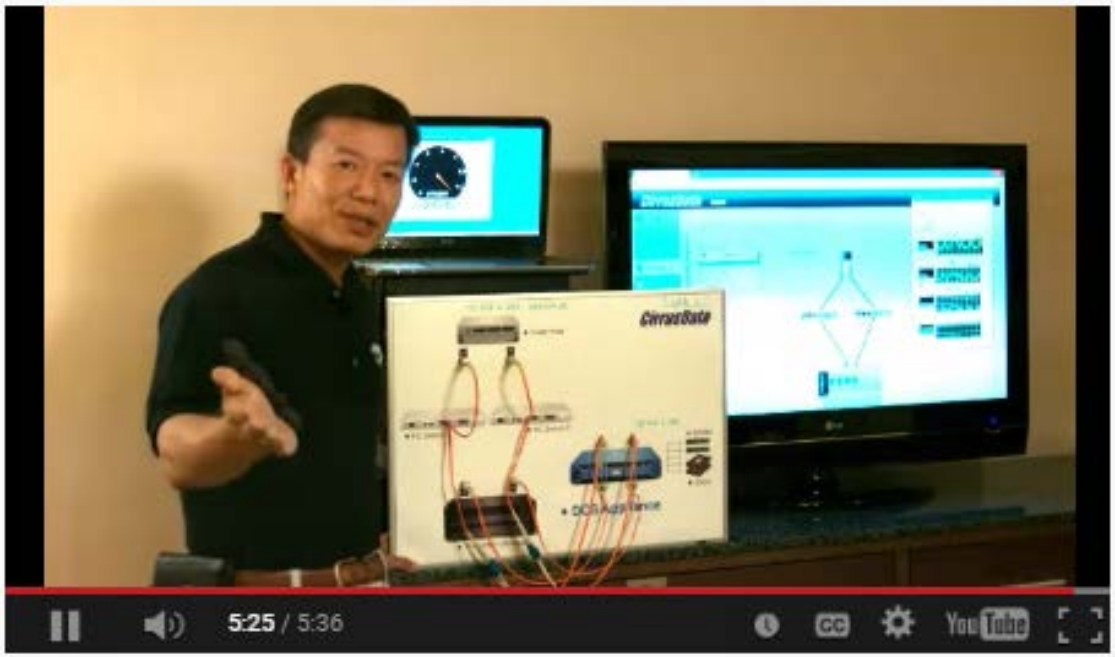

http://www.cdsi.us.com/data-caching-server/

Thank You!

CirrusData

Thank you

www.cdsi.us.com

Corporate Headquarters
America's
6800 Jericho Turnpike
Suite 213W
Syosset, NY 11791 USA

EMEA Headquarters
2 place de Touraine
78000 Versailles – France

APAC Headquarters
739 ShengAn Avenue, Suite 218
Nanjing Jiangsu, 211199
China