

DB2 for z/OS With EMC Storage Tiering: FAST VP

Paul Pendle
EMC Corporation

August 12, 2013
Session Number: 13988




Agenda

- The drivers for tiered storage
 - Technology changes
 - Workload skew
- FAST VP
 - Storage elements
 - Operating parameters
 - Lab testing and results
- Operational/host considerations
- Summary


Drivers Towards Storage Tiering

- Massive data growth
- Faster and faster processors
- Faster and faster channels
- Budgets are flat or decreasing
- Decline of the hard drive
- Arrival of SSD
 - Extremely high price and performance!
- Arrival of SATA
 - Extremely low price and performance!

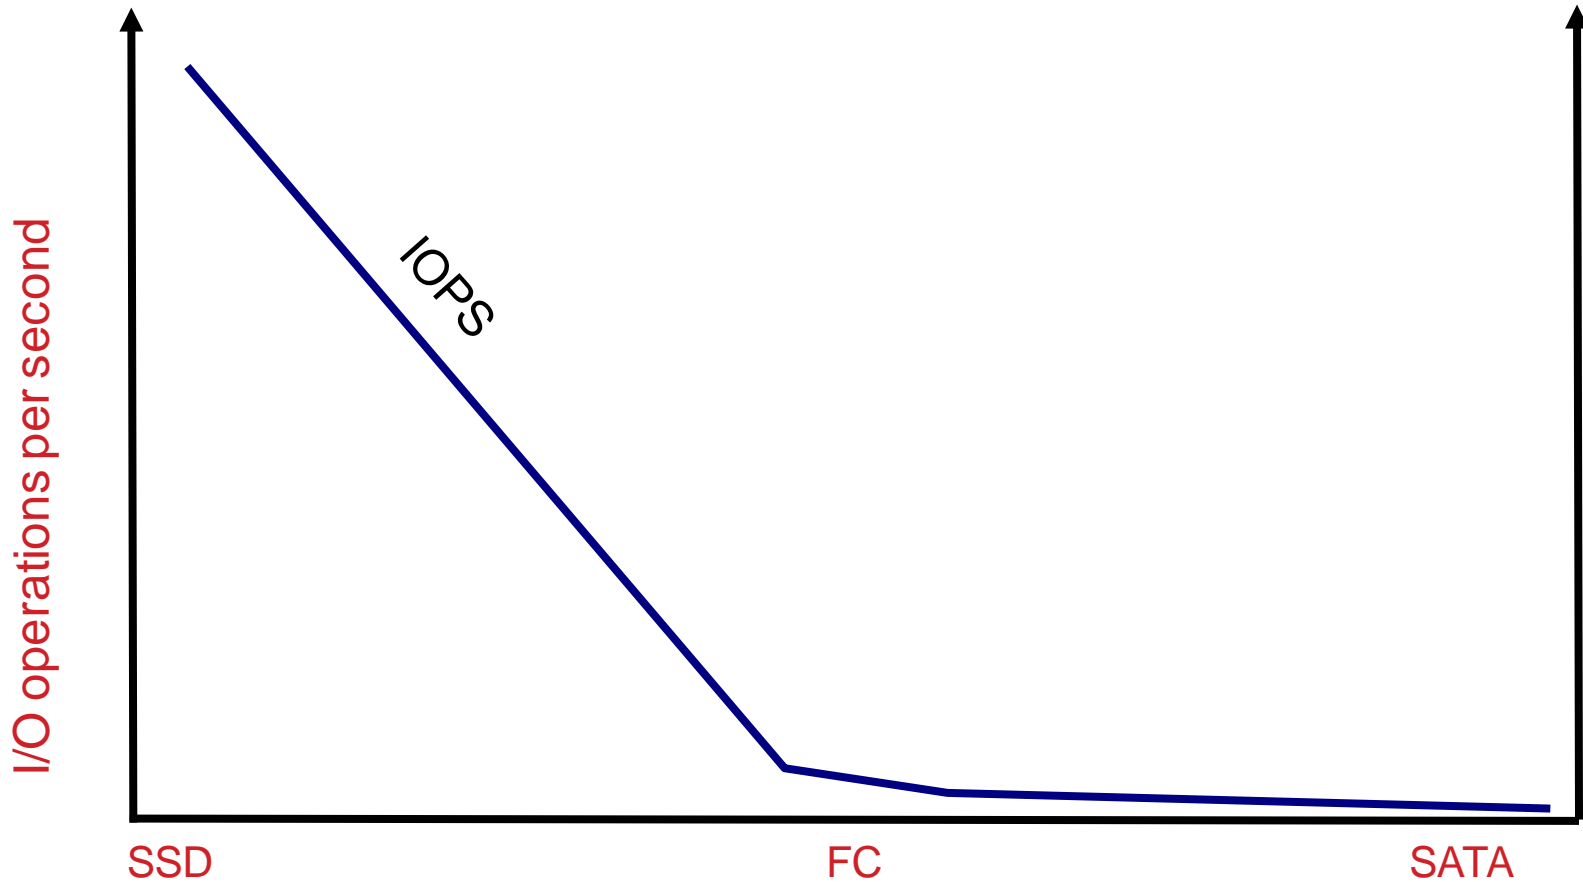
The Hard Drive RPM Story




IOPS Density Trend for Disk Drives



IOPS Comparisons of Drive Technology



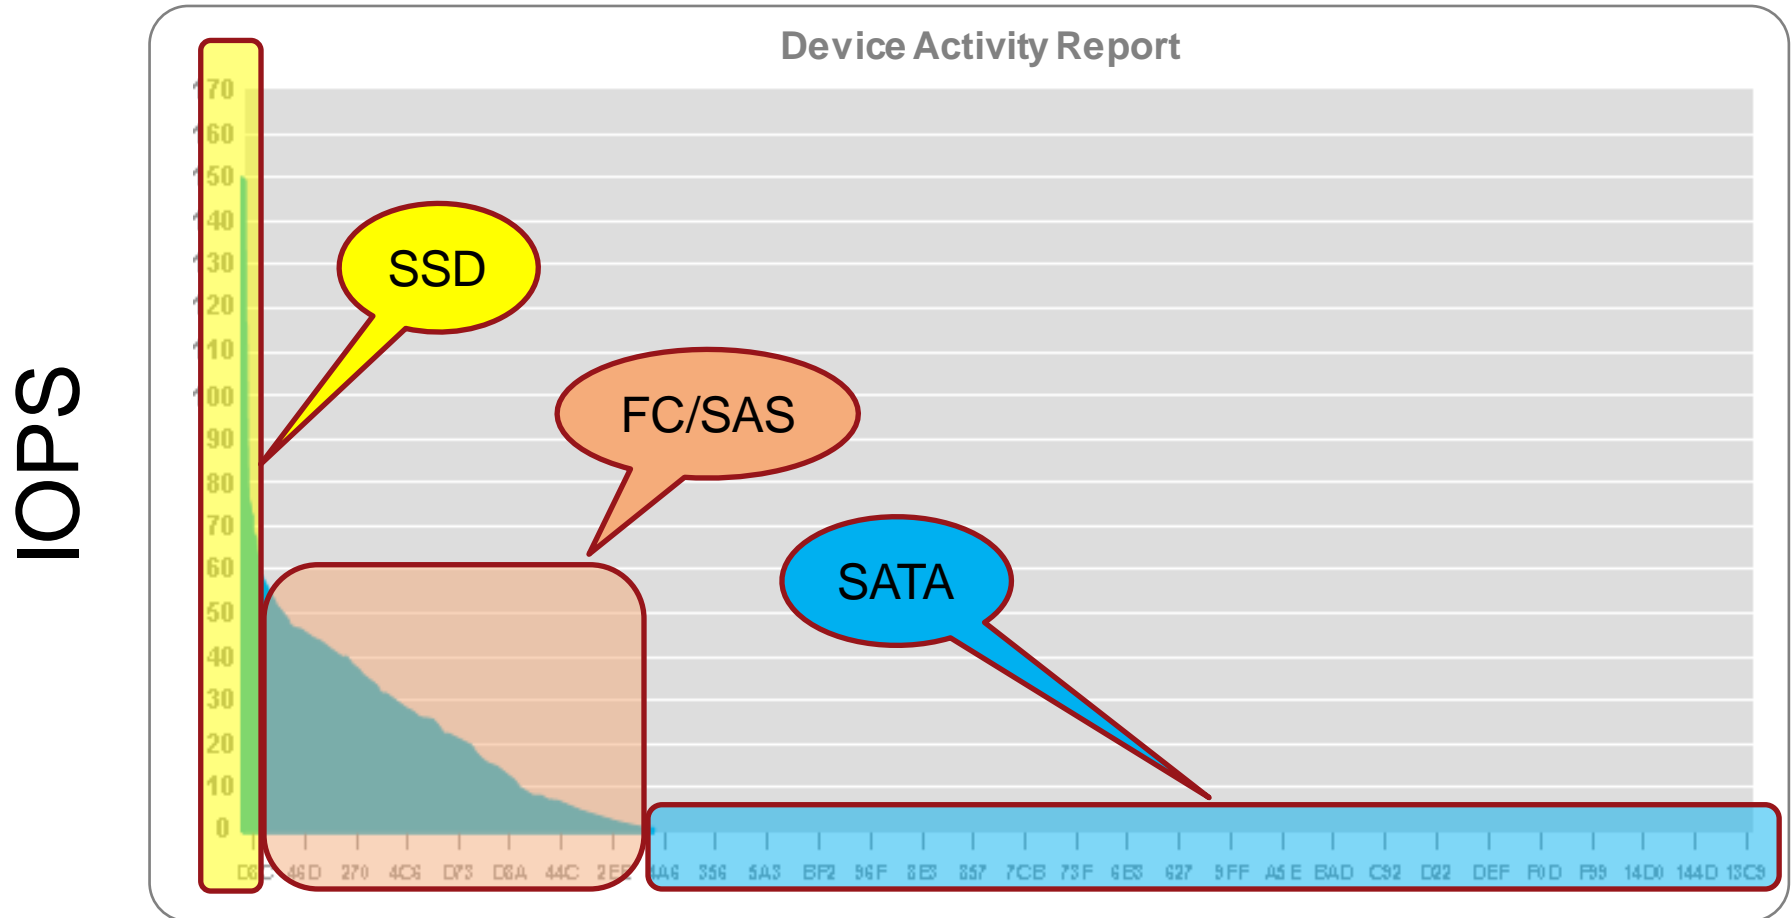
Data Life Cycle



Types of Workload Skew

- Persistent
 - To a large extent, historical activity is a good predictor of future activity
 - Good candidates for static tiering
- Non-persistent
 - Activity is mostly randomly skewed
 - Hot data today may not be hot tomorrow
 - Good candidates for automated tiering
- No skew (at full-volume level)
 - TPF, DB2 LUW DPF, Teradata

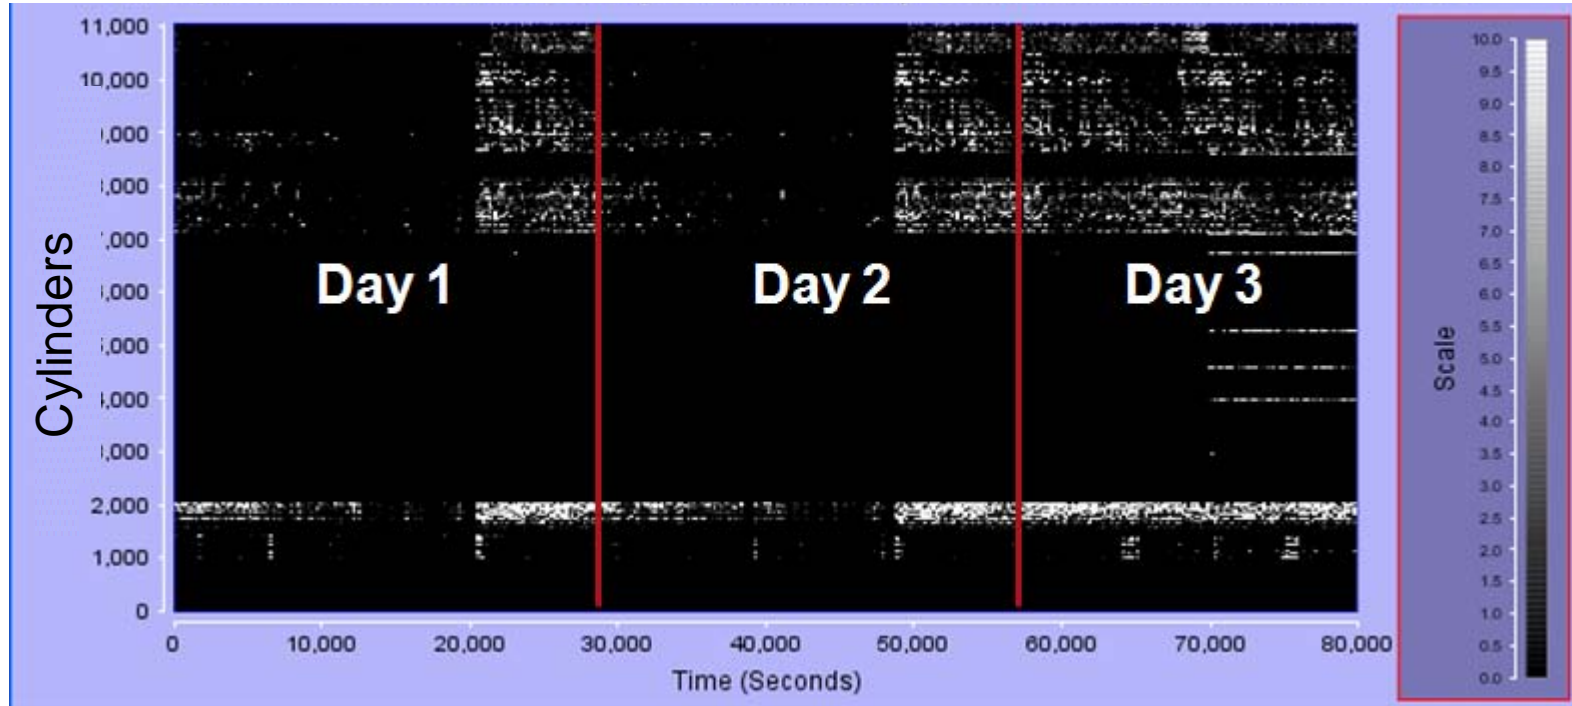
Workload Skew by Volume



The Static Tiering Challenge

- How do you know what database objects to place on each tier?
 - Largely, access patterns to an object change over time
 - The most frequently accessed objects are in the DB2 buffer pool or in the storage controller cache
 - The biggest objects are not good choices
- What about DB2 logs?
- High write table spaces?
- Sequentially accessed table spaces?
- The whole table space? Partition? Part of a table space?

Sub-Volume Skewing

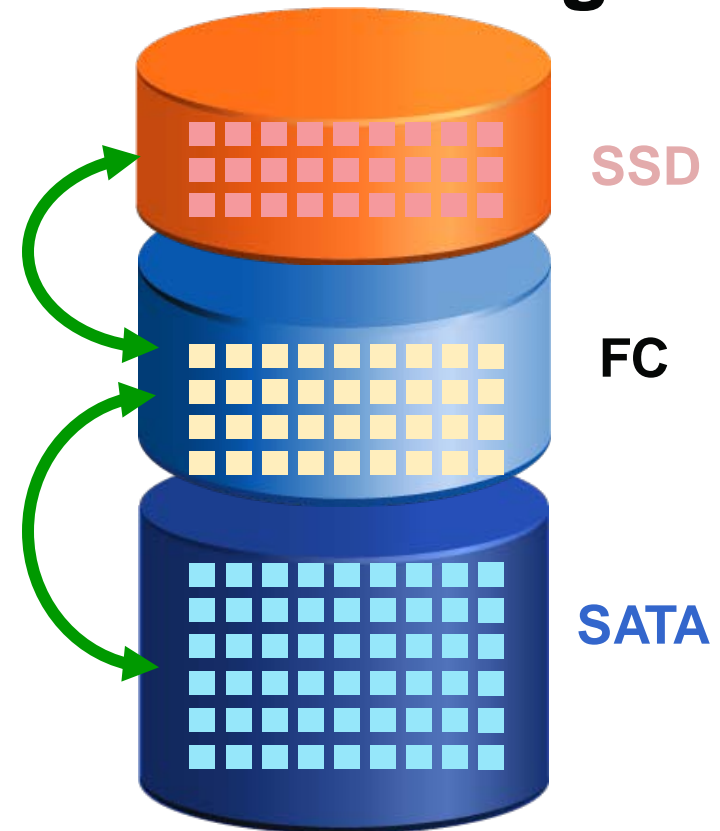


Typically, only parts of a volume are consistently “hot”

FAST VP: Automated Storage Tiering

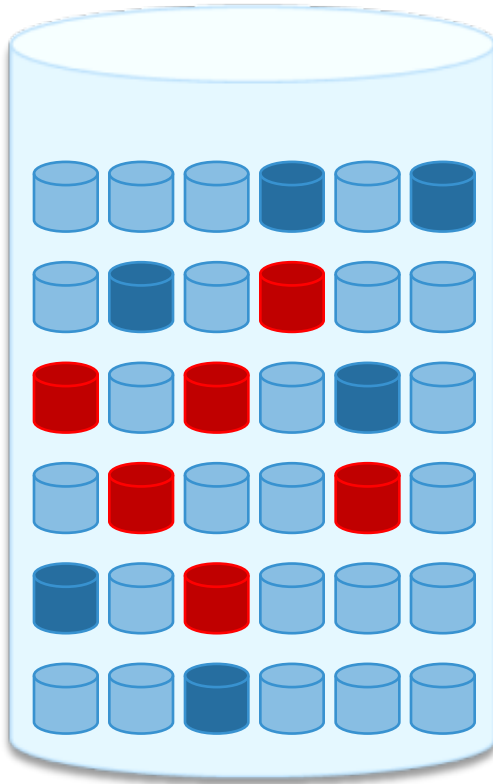
- **FASTVP**– Fully Automated Storage Tiering Virtual Pools
- FASTVP is a policy-based system that automatically promotes and demotes data across storage tiers to achieve performance objectives and cost targets
- Gets the **right data**, to the **right place**, at the **right time**

Tiered Storage

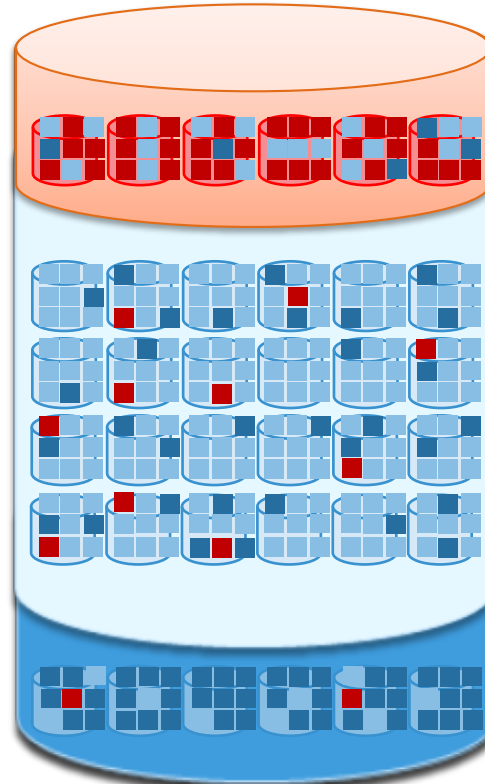


Evolution of FAST


Traditional




FAST



FAST VP



Tiered storage use case



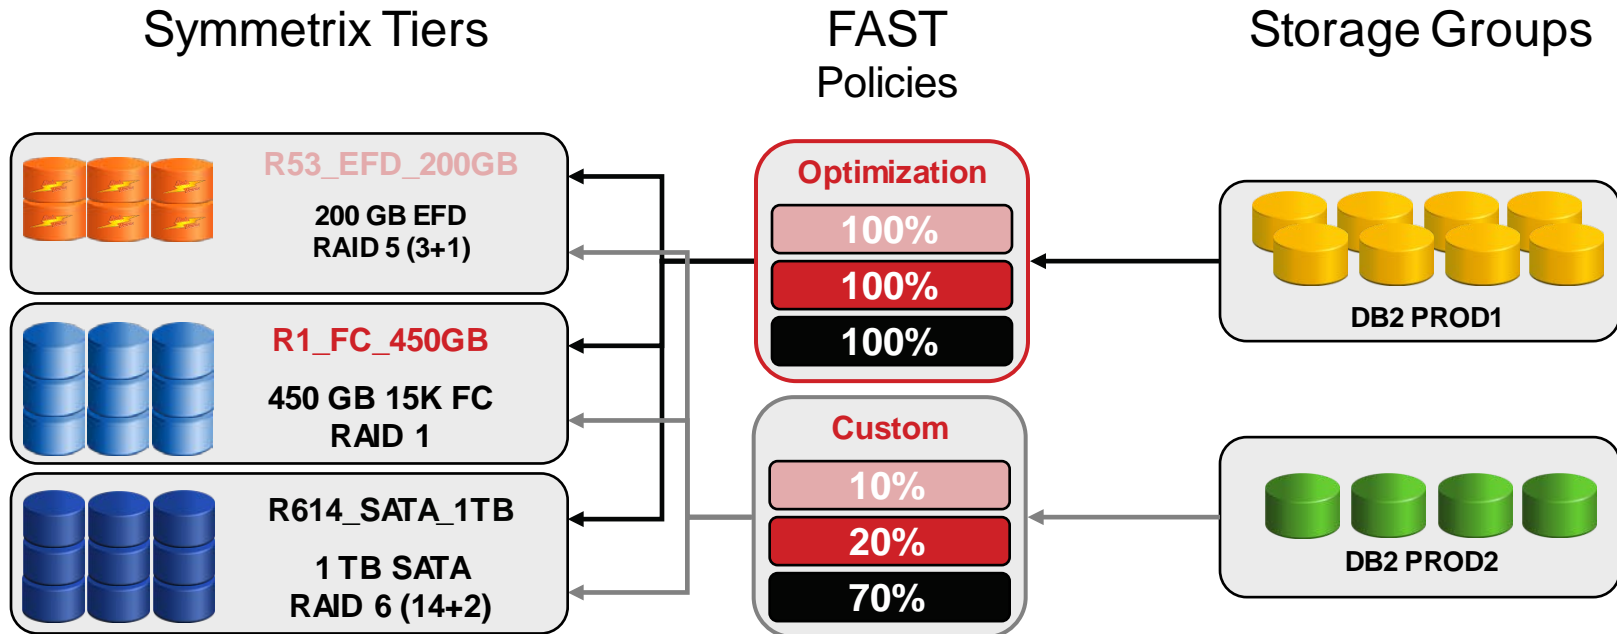
17% Lower Storage Costs
+ reduce maintenance & SW costs

45% More Disk IOPS
+ more aligned with workloads

32% Less Power & Cooling
+ more efficient use of space

37% Fewer Disk Drives
152 EFD+FC+SATA vs. 240 FC

FAST VP Elements



Symmetrix Tiers

- Tiers combine a drive technology and a RAID protection type, e.g.
 - Solid state
 - SAS/FC 15K RPM
 - SATA/SAS 7200 RPM

FAST Policies

- FAST Policies manage data placement and movement across Symmetrix Tiers for one or more Storage Groups
- Each Policy can contain up to three Symmetrix Tiers
 - Policies define the upper usage limit of each tier
- Each tier usage rule defines the maximum capacity of a storage group that can be moved to that tier
 - Each tier usage rule may be between 1% and 100%
 - Combined tier usage rules must total at least 100%, but may be greater than 100%
- Symmetrix Tiers may be shared among multiple FAST Policies

FAST Storage Groups

- Storage Groups logically group together devices for common management
- A Storage Group can have at most one policy associated with it
- Storage Groups may contain multiple device types
- Devices may be “pinned” to prevent FAST movement
 - Performance statistics will continue to be collected for pinned devices
 - Statistics included when generating a new performance movement policy

Time Windows

- Performance Window
 - Defines the times of the day, and the days of the week during which performance data is collected
 - Allows for “quiet” periods, or irregular workloads to be excluded from analysis
- Data Movement Window
 - Defines the times of the day, and the days of the week during which data movements will automatically be performed

Create FAST Policy

EMC Unisphere for VMAX - Windows Internet Explorer

https://10.12.160.254:8443/univmax/#0zszbfv0366

EMC Unisphere for VMAX V1.5.0.6

Home System **Storage** Hosts Data Protection Performance Support

000195700455 > Storage > FAST > Manage Policies

FAST Policies

Policy Name	Tier 1	Tier 1 %	Tier 2	Tier 2 %
FBA_Initial	Jim_EFDR53	3	Jim_FCR1	100
FBA_LOG	Jim_EFDR53	3	multipool	75
Gold	Jim_EFDR53	3	Jim_FCR1	20
paul_test	zOS_SD_R3	50	zOS_FC_2M	50
SAP	Jim_EFDR53	2	Jim_FCR1	27
test	multipool	100	N/A	0
zos_bronze	zOS_FC_2M	65	zOS_AT_R3	35
zos_gold	zOS_SD_R3	25	zOS_FC_2M	100
zOS_OPT	zOS_SD_R3	100	zOS_FC_2M	100
zos_silver	zOS_SD_R3	1	zOS_FC_2M	79

Create FAST Policy

* Policy Name: DB2FASTVP

* Emulation: CKD 3390

Tier *: zOS_SD_R3 10 %

Tier: zOS_FC_2M 20 %

Tier: zOS_AT_R6 70 %

Tier: N/A %

OK Cancel Help

Associated Storage Groups

Policy Name	# Associated Storage Groups
FBA_Initial	1
FBA_LOG	1
Gold	0
paul_test	0
SAP	0
test	0
zos_bronze	0
zos_gold	0
zOS_OPT	2
zos_silver	1

Common Tasks

- Create a new host
- Manage hosts
- Provision storage
- Create a storage group
- Expand a thin pool
- Create volumes
- Create a meta volume
- Replicate storage locally

FAST Policies with New Policy Listed

EMC Unisphere for VMAX V1.5.0.6

Home System **Storage** Hosts Data Protection Performance Support

000195700455 > Storage > FAST > Manage Policies

FAST Policies

Policy Name	Tier 1	Tier 1 %	Tier 2	Tier 2 %	Tier 3	Tier 3 %	Tier4	Tier4 %	# Associated Storage Groups
DB2FASTVP	zOS_SD_R3	10	zOS_FC_2M	20	zOS_AT_R6	70	N/A	0	0
FBA_Initial	Jim_EFDR53	3	Jim_FCR1	100	Jim_SATAR614	100	N/A	0	1
FBA_LOG	Jim_EFDR53	3	multipool	75	Jim_SATAR614	100	N/A	0	1
Gold	Jim_EFDR53	3	Jim_FCR1	20	Jim_SATAR614	100	N/A	0	0
paul_test	zOS_SD_R3	50	zOS_FC_2M	50	zOS_AT_R3	100	N/A	0	0
SAP	Jim_EFDR53	2	Jim_FCR1	27	Jim_SATAR614	100	N/A	0	0
test	multipool	100	N/A	0	N/A	0	N/A	0	0
zos_bronze	zOS_FC_2M	65	zOS_AT_R3	35	N/A	0	N/A	0	0
zos_gold	zOS_SD_R3	25	zOS_FC_2M	100	N/A	0	N/A	0	2
zOS_OPT	zOS_SD_R3	100	zOS_FC_2M	100	zOS_AT_R6	100	N/A	0	1
zos_silver	zOS_SD_R3	1	zOS_FC_2M	79	zOS_AT_R3	20	N/A	0	0

Common Tasks

- Create a new host
- Manage hosts
- Provision storage
- Create a storage group
- Expand a thin pool
- Create volumes
- Create a meta volume
- Replicate storage locally

Create New Time Window

EMC Unisphere for VMAX V1.5.0.6

Home System **Storage** Hosts Data Protection Performance Support

000195700455 > Storage > FAST

FAST Type FAST VP

FAST Performance Time Window

Show Week 4 Sunday January 20th - Saturday January 26th

Day / Time	00:00	01:00	02:00	03:00	04:00	05:00	06:00	07:00	08:00	09:00	10:00	11:00	12:00
Sunday	[Green bar]												
Monday	[Green bar]												
Tuesday	[Green bar]												
Wednesday	[Green bar]												
Thursday	[Green bar]												
Friday	[Green bar]												
Saturday	[Green bar]												

Open Time Windows (Inclusive)

Hide Advanced <<

Open Time Windows (Inclusive)

Closed Time Windows (Exclusive)

Manage Open Performance Time Windows

Define New Time Window

- Always Open.
- All Weekend (Fri:18:00 - Mon:00:00)
- 09:00 - 17:00, Monday - Friday
- 17:00 - 08:00, Monday - Friday
- Custom
 - Monday [00:00] to [00:00]

Existing Open Performance Time Windows

Day	Time
Monday	00:00 to 24:00
Tuesday	00:00 to 24:00
Wednesday	00:00 to 24:00
Thursday	00:00 to 24:00
Friday	00:00 to 24:00
Saturday	00:00 to 24:00
Sunday	00:00 to 24:00

Add Delete

OK Cancel Help

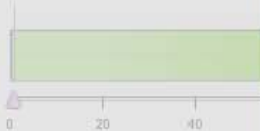
OK Cancel Help

Common Tasks

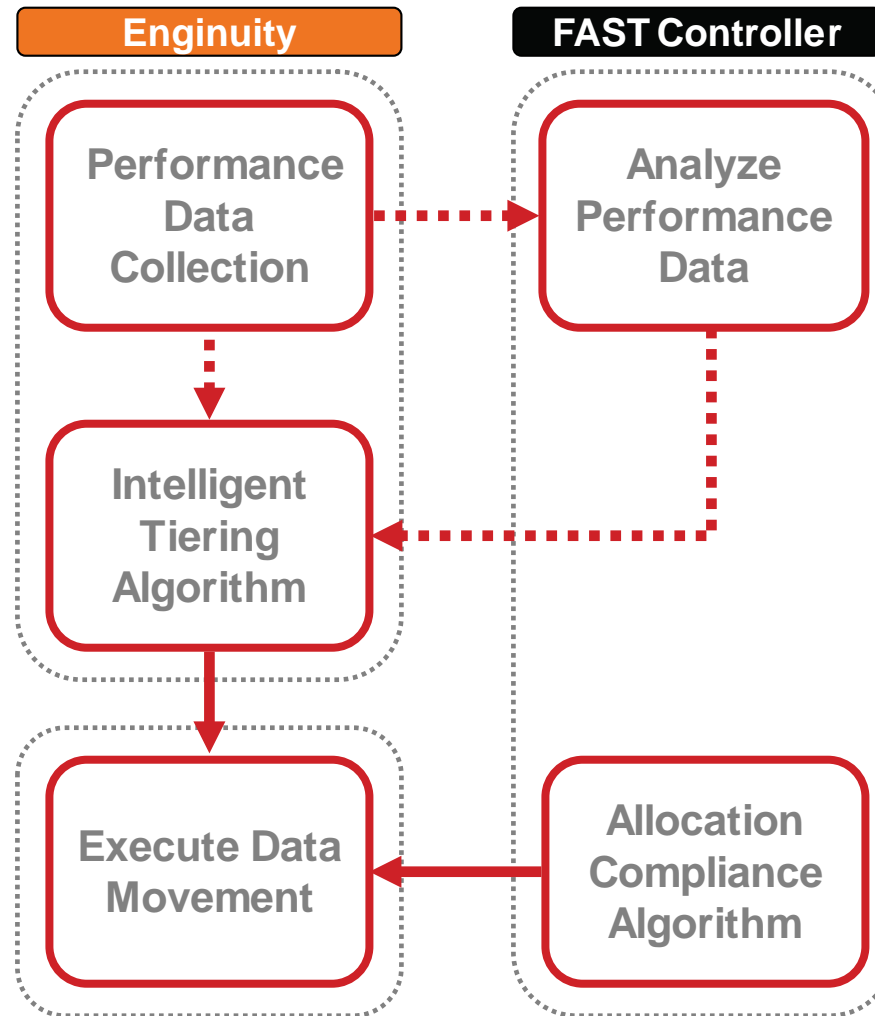
- Create a new host
- Manage hosts
- Provision storage
- Create a storage group
- Expand a thin pool
- Create volumes
- Create a meta volume
- Replicate storage locally

Tiers Demand Report

Jim_EFDR53



FAST VP Implementation

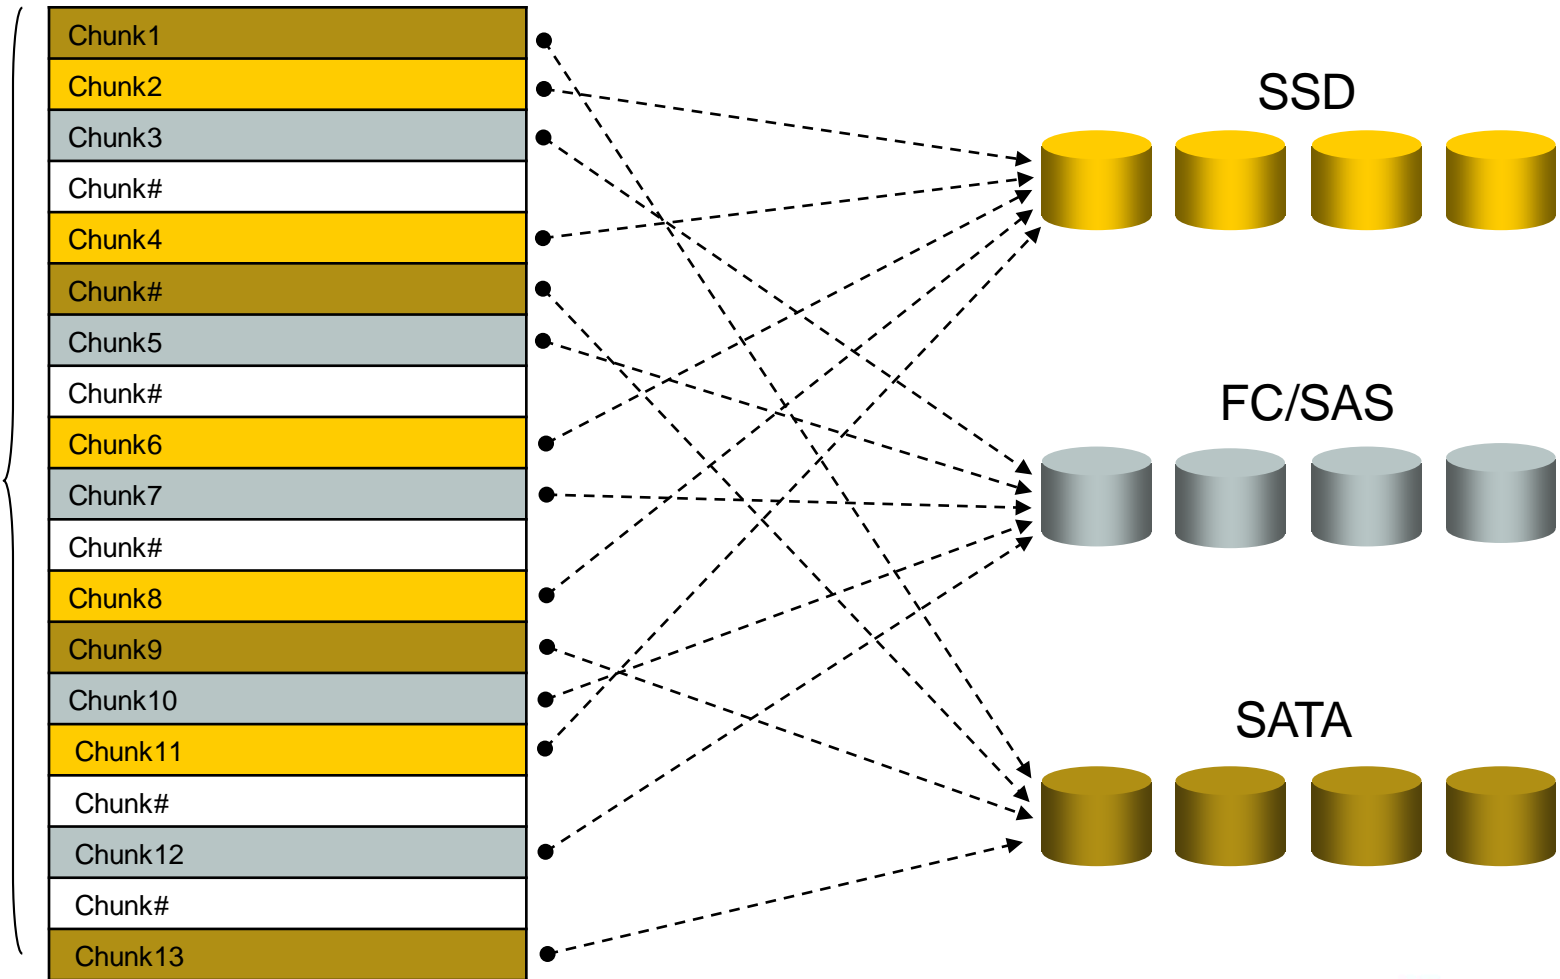


Data Movement

- Data chunks are moved using 6.8MB chunks
- Performance-based
 - Promotions due to high I/O rates
 - Demotions to free up space for promotions
- Compliance-based


Sub-volume Tiering

TABLE SPACE



Sequential Tablespace Scan

- Long Seek (costly)
- Rotational delay (costly)
- Read
- Read
- Read ...



FAST VP TESTING WITH DB2

What was tested

- DB2 V10
- Symmetrix VMAX SE (single Engine)
- 2x4Gb Channels
- Highly random OLTP workload driven by 32 batch jobs
- 4x200GB Solid State drives
- 32x300GB 15K FC drives
- 507GB DB2 subsystem
- FAST VP policy set to 10% SSD use

Policy display before workload

000195700455 > Storage > Storage Groups > ZOS1_0455_DB2_SG

Details : Storage Group : ZOS1_0455_DB2_SG

Properties	
Name	ZOS1_0455_DB2_SG
FAST Policy	zos_gold
FAST Priority	2
Total Capacity (GB)	507.5
Host Name	N/A
Volumes	64
Masking Views	0

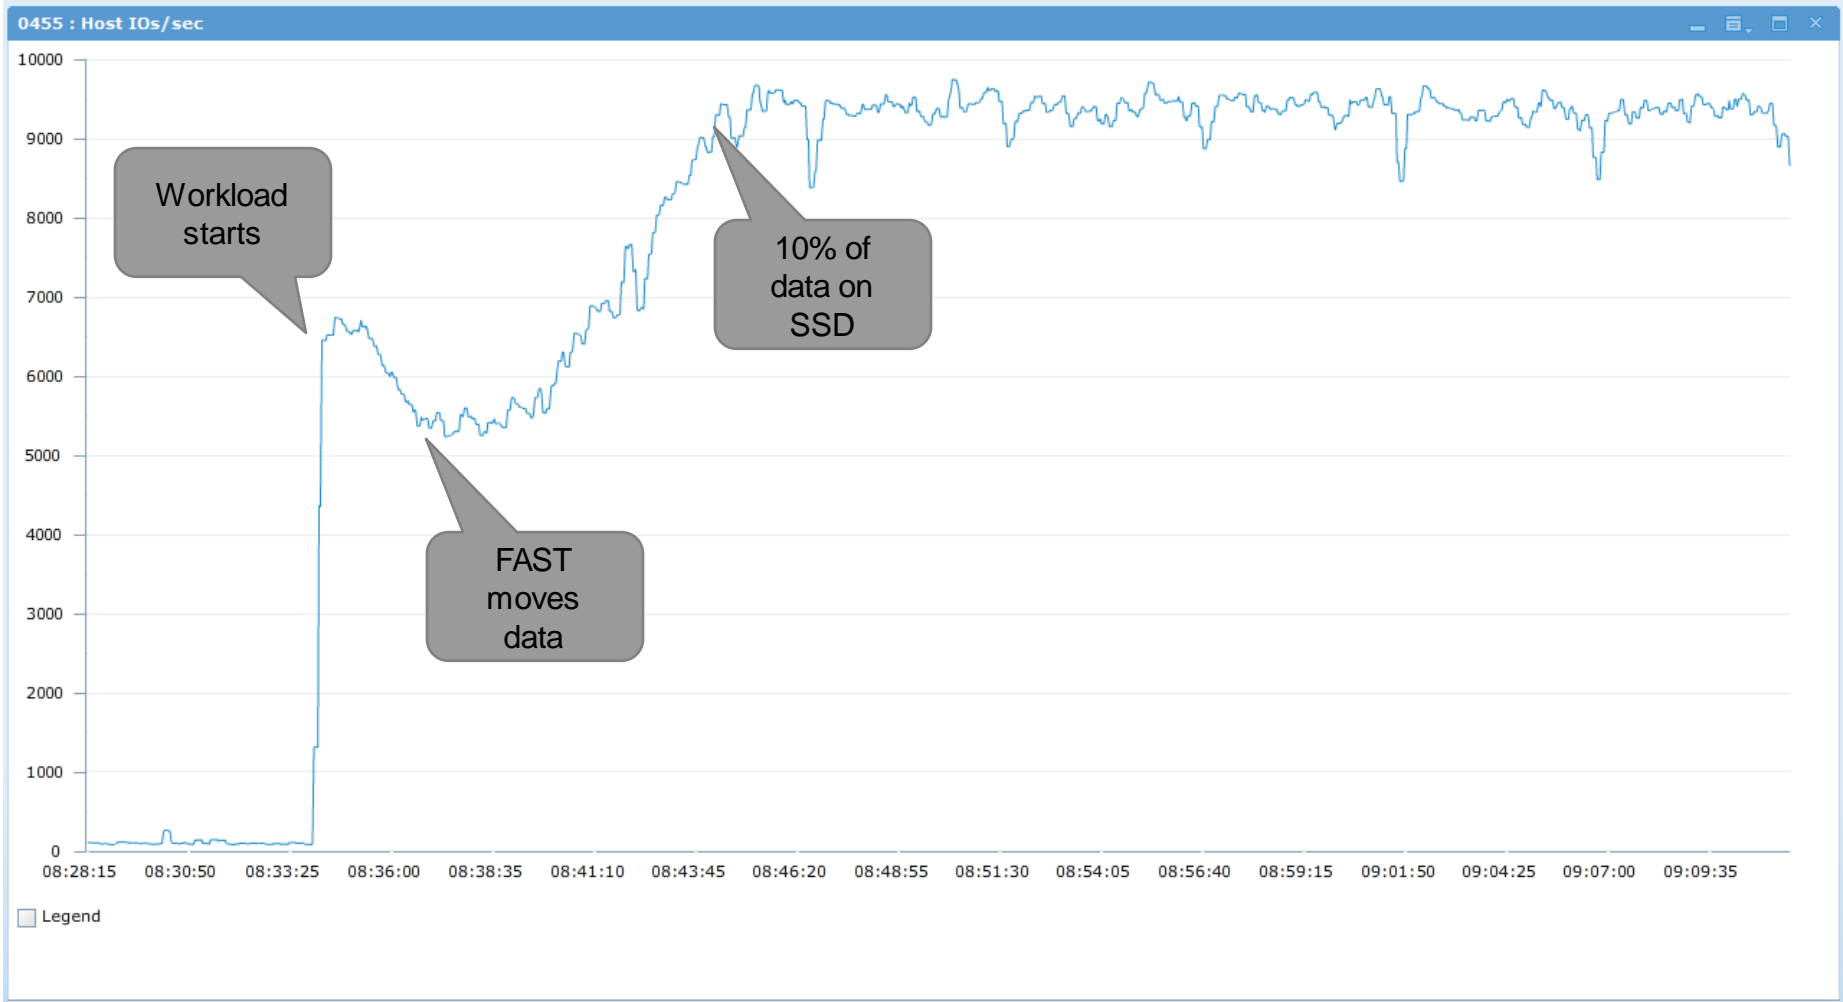
Related Objects

Contains : [Volumes - 64](#)
Associated With : [FAST Policy - 1](#)

FAST Compliance Report

Tier	Protection	Technology	Max SG Demand (%)	Limit (GB)	Fast SG Used (GB)	Growth (GB)
zOS_SD_R3	RAID-5 (3+1)	EFD	10	+50.75	0	+50.75
zOS_FC_2M	RAID-1	FC	100	+507.48	+507.95	--0.48

IOPS Measured: Unisphere for VMAX Performance View



Policy display after workload

000195700455 > Storage > Storage Groups > ZOS1_0455_DB2_SG

Details : Storage Group : ZOS1_0455_DB2_SG ?

Properties

Name	ZOS1_0455_DB2_SG
FAST Policy	zos_gold ▼
FAST Priority	2 ▼
Total Capacity (GB)	507.5
Host Name	N/A
Volumes	64
Masking Views	0

Related Objects

Contains : [Volumes - 64](#)

Associated With : [FAST Policy - 1](#)


FAST Compliance Report

Tier	Protection	Technology	Max SG Demand (%)	Limit (GB)	Fast SG Used (GB)	Growth (GB)
zOS_SD_R3	RAID-5 (3+1)	EFD	10	+50.75	+50.71	+0.04
zOS_FC_2M	RAID-1	FC	100	+507.48	+457.25	+50.23


HOST CONSIDERATIONS

Automation at all Layers

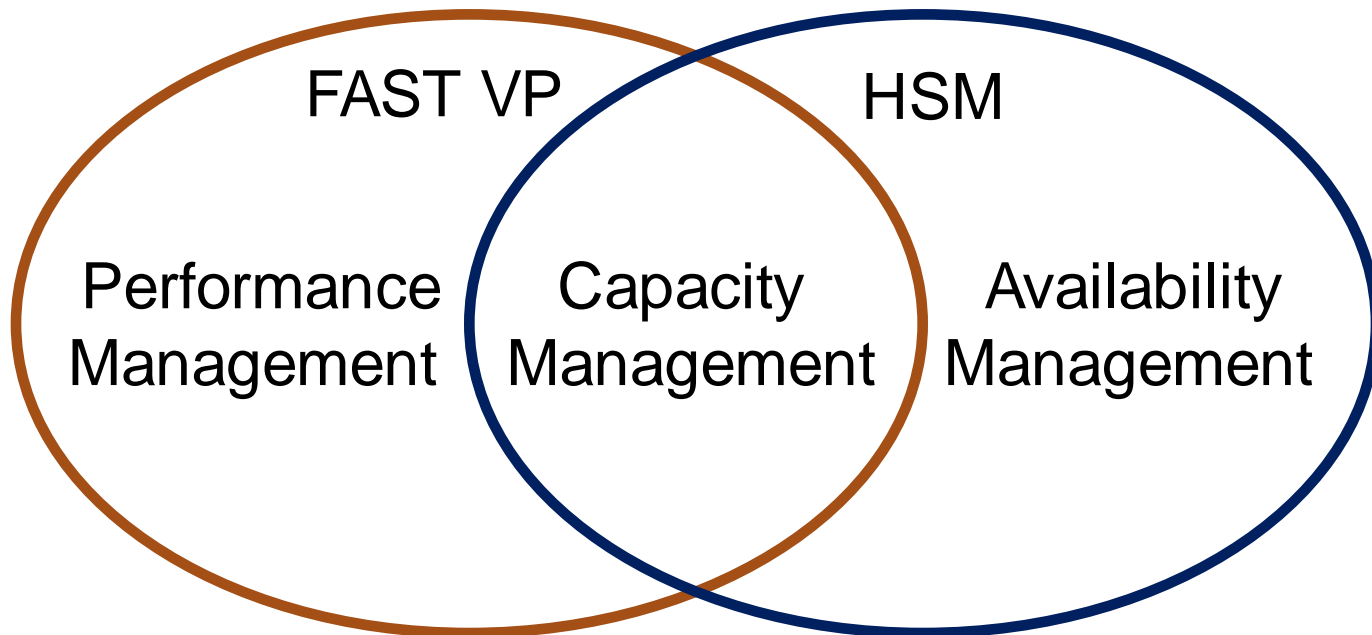
SMS &
HSM



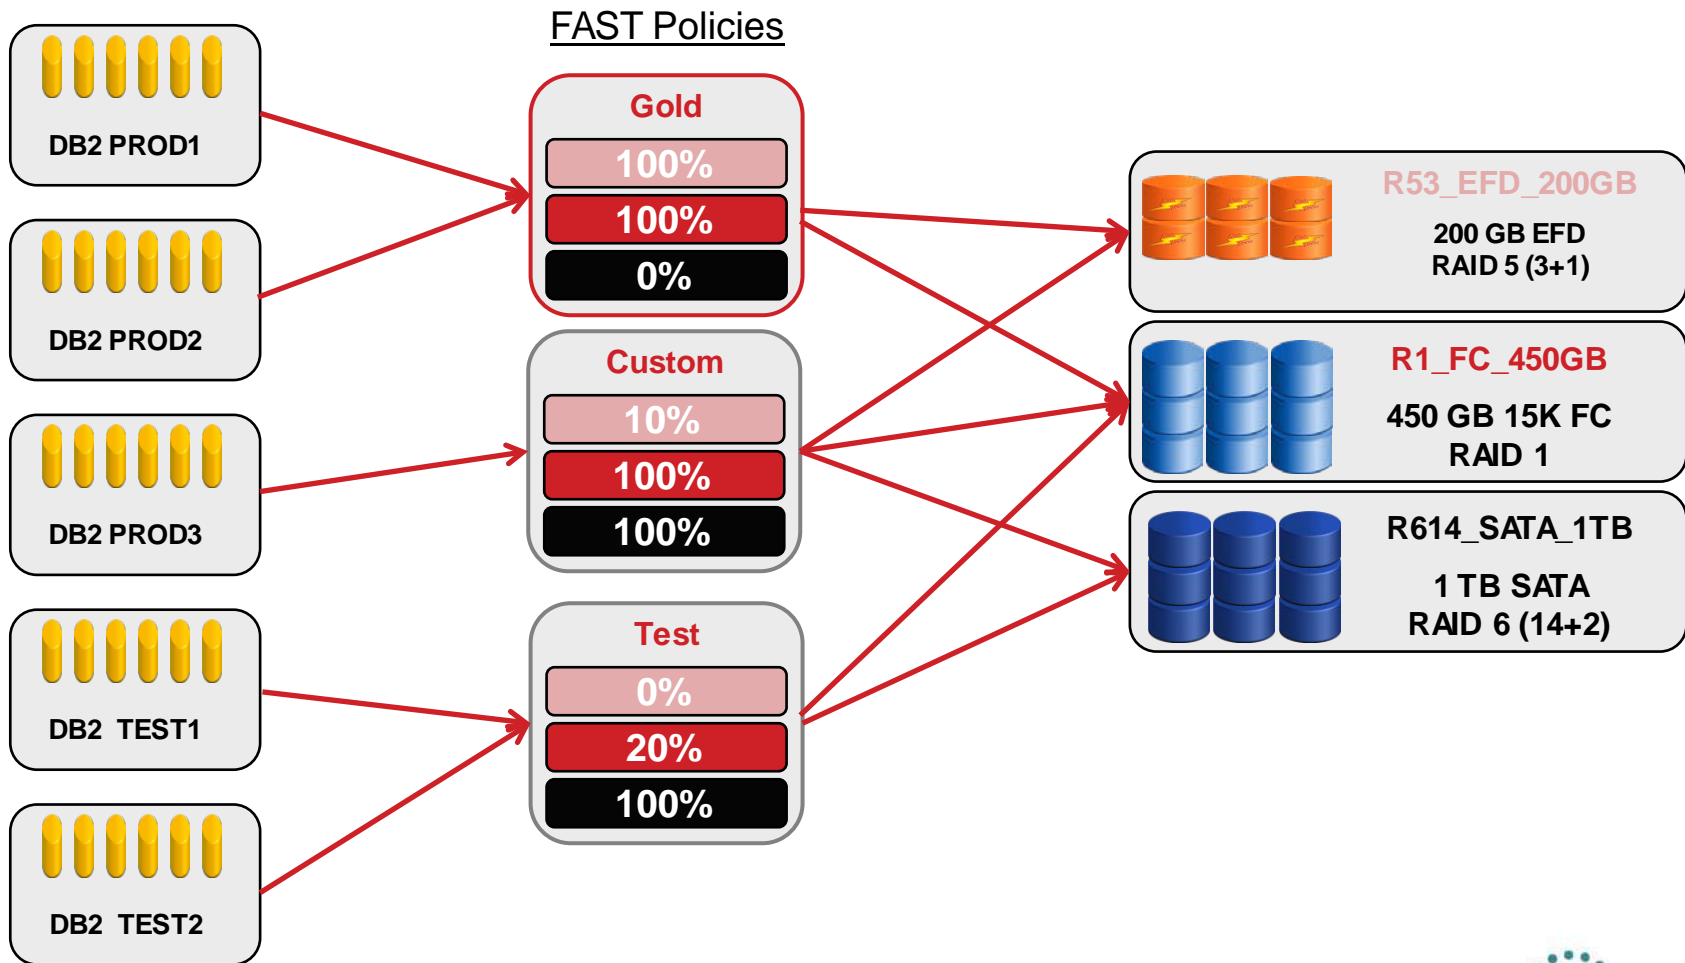
FAST VP



HSM and FAST VP Intersect




SMS Storage Groups

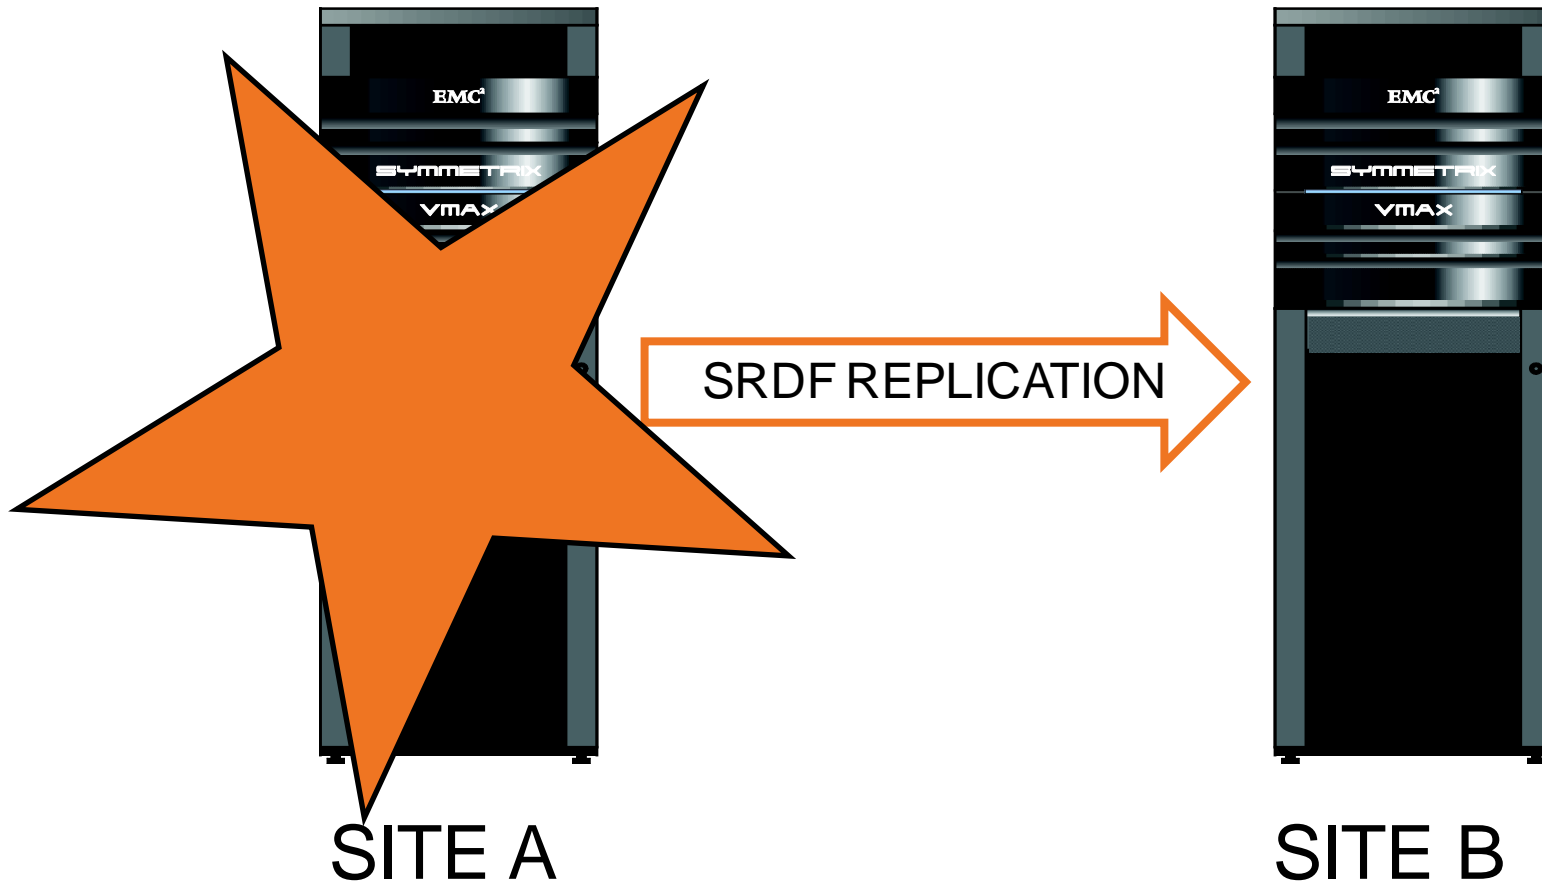


Operational Considerations

- Storage tiering interactions with z/OS
- SMS
 - Performance-based allocations
 - DIRECT MSR Values (what do they mean now?)
- Thrashing
 - DB2 REORGs
 - HRECALLs
 - Dataset moves
 - Volume restore



FAST VP with Remote Replication Integration



Operational Considerations (contd)

- How to manage charge back
- How to influence decisions in the performance engine
- How to determine where everything is/was

Benefits of Storage Tiering

- Autonomic/automatic operation
- Optimized performance
- Reduced cost (power and cooling)
- Reduced footprint
- Better capacity utilization
- Ease of management

DB2 for z/OS With EMC Storage Tiering: FAST VP

Paul Pendle
EMC Corporation

August 12, 2013
Session Number: 13988

