



Pulse2011



Storage Virtualisation: Storwise V7000 and SVC

Anthony Vandewerdt
Storage Solutions Specialist

Is this new?

Six reasons to adopt storage virtualisation

1. Storage virtualisation reduces complexity
2. Storage virtualisation makes it easier to allocate storage
3. Better disaster recovery
4. Better tiered storage
5. Virtual storage improves server virtualisation
6. Virtual storage lets you take advantage of advanced virtualisation features





The screenshot shows a CNET News article page. At the top left is the CNET News logo. A navigation bar contains links for Home, Reviews, News (highlighted), Downloads, Video, and How To. Below this is a secondary navigation bar with links for Latest News, Webware, Crave, Business Tech (highlighted in red), Green Tech, and Wireless. The breadcrumb trail reads 'Home > News > Storage'. The article date is 'June 2, 2003 9:00 PM PDT'. The main headline is 'IBM enhances storage virtualization' in large red font. The author is 'By Ed Frauenheim', a Staff Writer at CNET News. A 'Related Stories' box is partially visible on the left. The main text snippet reads: 'IBM is poised to ship two data storage virtualization products in July, the latest step in the industry's push to squeeze more productivity out of storage systems.'


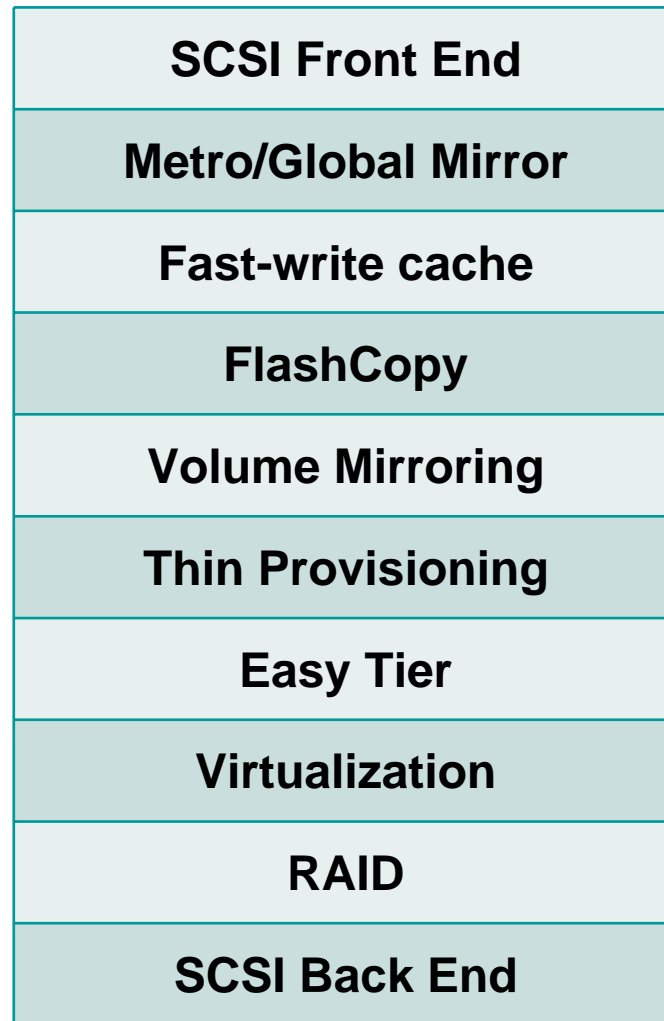
http://news.cnet.com/IBM-enhances-storage-virtualization/2100-1015_3-1012438.html?tag=mncol;txt



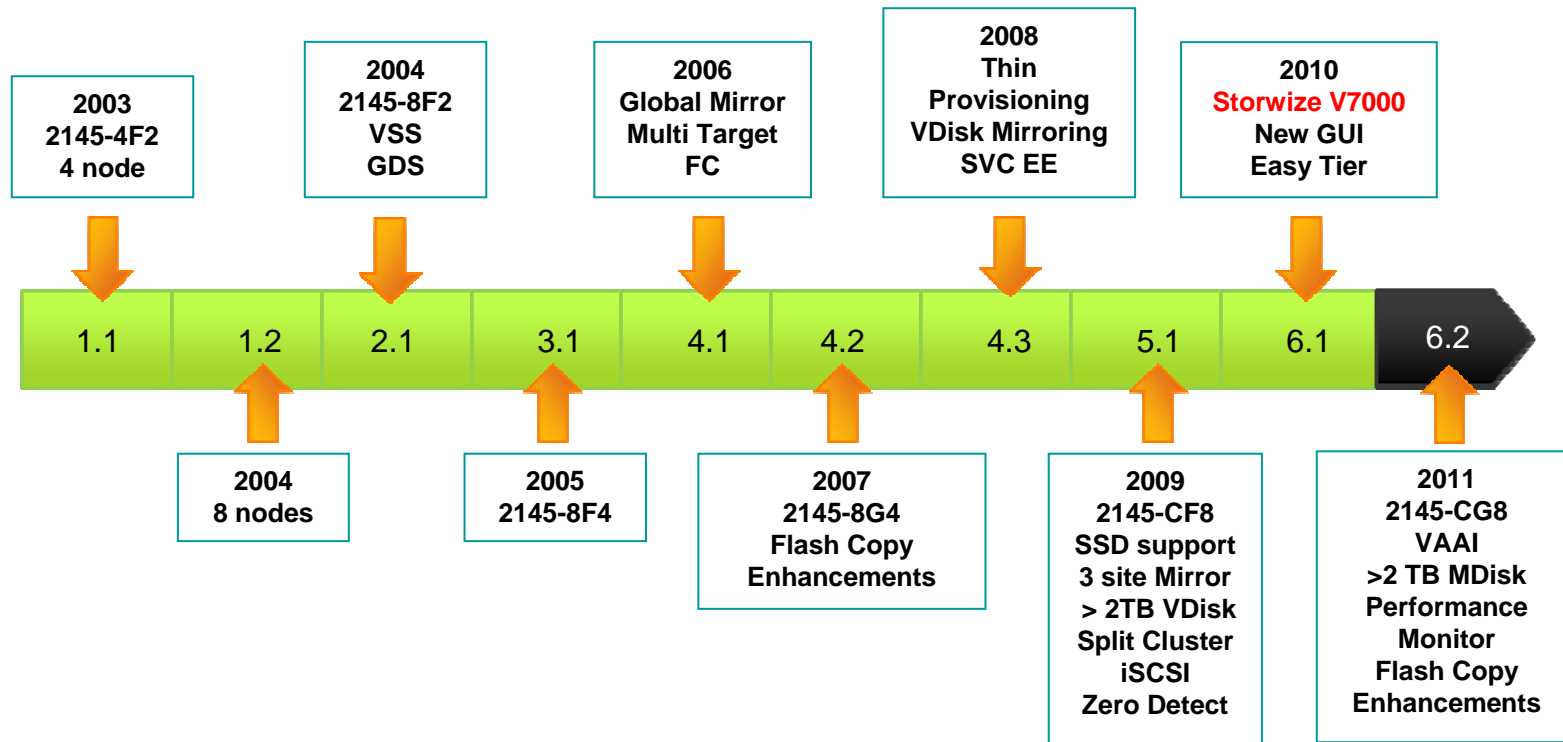
The hardware is not exciting




The software stack is far more exciting...



SVC release history is a continuous evolution



What was Project Quicksilver?




What do you get when you add these together?




A New Era in Midrange Disk Systems!




We start with a Storage Bridge Bay enclosure



We create a Control Enclosure






The Node Canister is an SVC Node



Each Control Enclosure Power Supply has a Battery



Then we add some disks



What are our disk choices?

Drive Type		Drive Sizes
2.5" SAS v2	e MLC SSD	300GB
2.5" SAS v2	15K RPM	146 GB
2.5" SAS v2	10K RPM	300 GB 450 GB 600GB
3.5" NL-SAS v2	7.2K RPM	2TB



No restrictions on mixing sizes or quantities in an enclosure

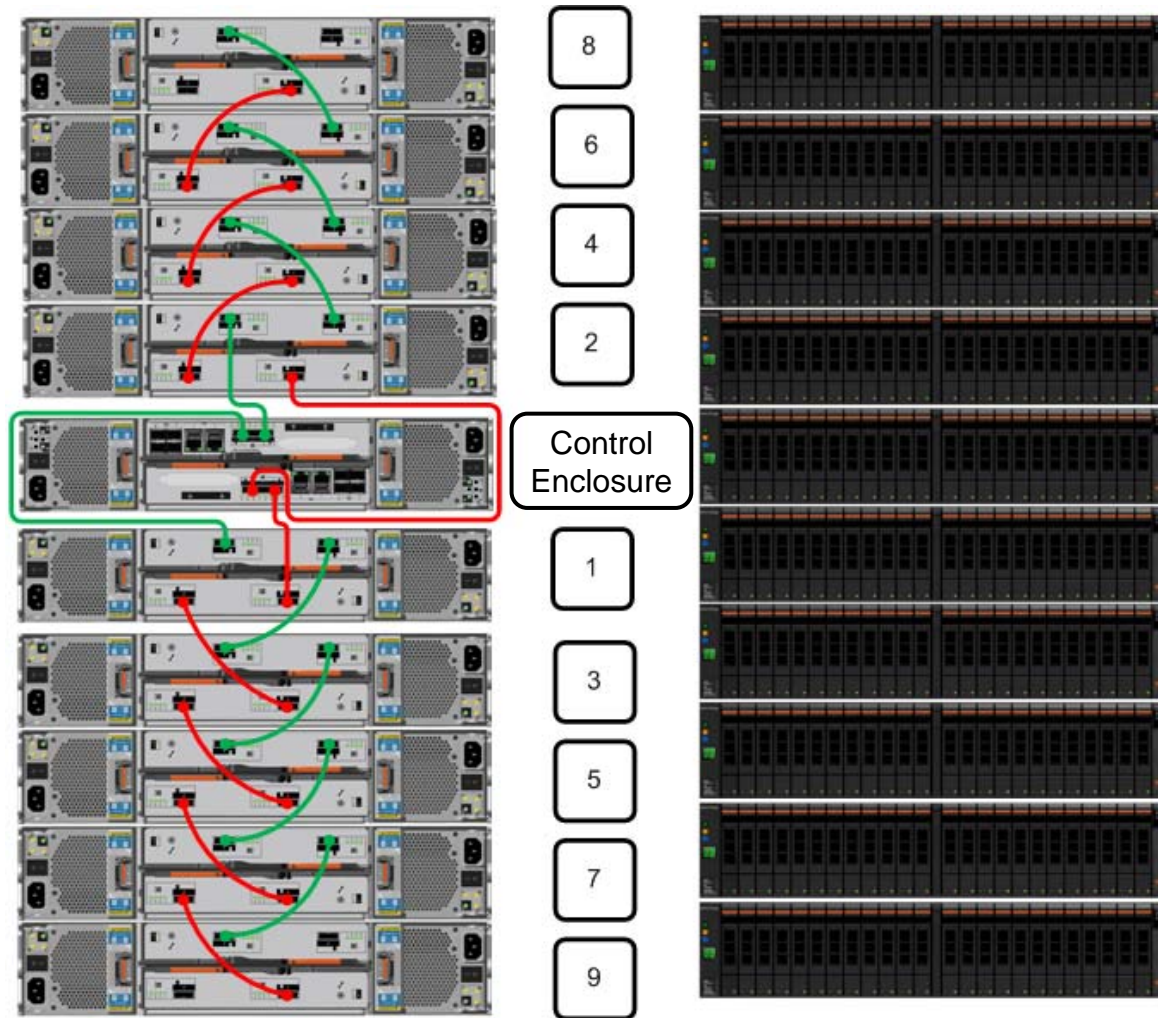




Why SAS Version 2.0?

Serviceability Feature	Support			Pure 6Gb/s SAS 2.0
	4Gb/s FC	3Gb/s SAS 1.0	3Gb/s SAS 2.0	
Drive Channel Health Monitoring	YES	YES	YES	YES
Collect All Port Statistics	YES	YES	YES	YES
Retrieve Drive Diagnostic Data	YES	YES	YES	YES
Cabling Error Mis-Wire Detection and Notification	YES	YES	YES	YES
Degraded Wide Port Notification	N/A	YES	YES	YES
Automatically Disable Faulty Port	YES	YES	YES	YES
Drive Channel Initialization Storm Protection	YES	NO	NO	YES
Maintain History of Port Statistics	YES	NO	NO	YES
Proactive Drive Health Monitoring	YES	NO	NO	YES
SATA Drive SMART Event Polling	YES	NO	NO	YES
Slow PHY Speed Detection	YES	NO	NO	YES
Drive Port Fault Protection	YES	NO	NO	YES
Controller FRU Fault Isolation Diagnostics	YES	NO	NO	YES
T10 PI and 520B Support with Recovered Error Limits	Yes	NO	NO	YES
Improved Redundant Drive Path Fault Tolerance	NO	NO	NO	YES
Controller to Controller Channel Fault Protection	NO	NO	NO	YES





SAS Cabling For Easy Expansion



Pulse2011



Configuration Concepts

RAID arrays




RAID Level	Description	Drive Count	Approximate Array Capacity	Redundancy	Striped Width Goal (with Preset)	Sparing Goal (With Preset)
0	Striped with no protection	1 - 8	$DC * DS$	None		0
1	Mirrored	2	DS	1	2	1
5	Striped with parity	3 - 16	$(DC - 1) * DS$	1	8	1
6	Striped with dual parity	5 - 16	$< ((DC - 2) * DS)$	2	12	1
10	Mirrored then striped	2 - 16 (even only)	$(DC/2) * DS$	1 - DC/2	8	1

DS = Drive Size - DC = Drive Count



Virtualisation Concepts & Limits



Volumes

- Up to 2,048 Volumes per cluster
- Max. Volume size of 256TiB
- Image Mode = Native Mode
- Managed Mode = Virtualized Mode

Hosts

- Up to 256 Hosts per cluster
- Up to 512 FC ports per cluster & Host
- Up to 256 iSCSI ports per cluster & Host

Storage Pools

- Up to 128 MDisks per Storage Pool
- Up to 128 storage pools per cluster

Internal Storage

1 to 8/16 drives (HDD) to form an Array
1 Array → Managed Disk (MD)

Up to 4,096 MDisks per cluster

Max. External MDisk Size of 1 PB with approved controllers

External Storage

Multiple Storage Subsystems
1 LUN → Managed Disk (MDdisks)

New GUI

Navigation, user indicator

IBM Storwize V7000 Welcome, **superuser** | About | Logout | Help | IBM

ATS_Oxford_1 > Home > **Getting Started** Getting Started System Status

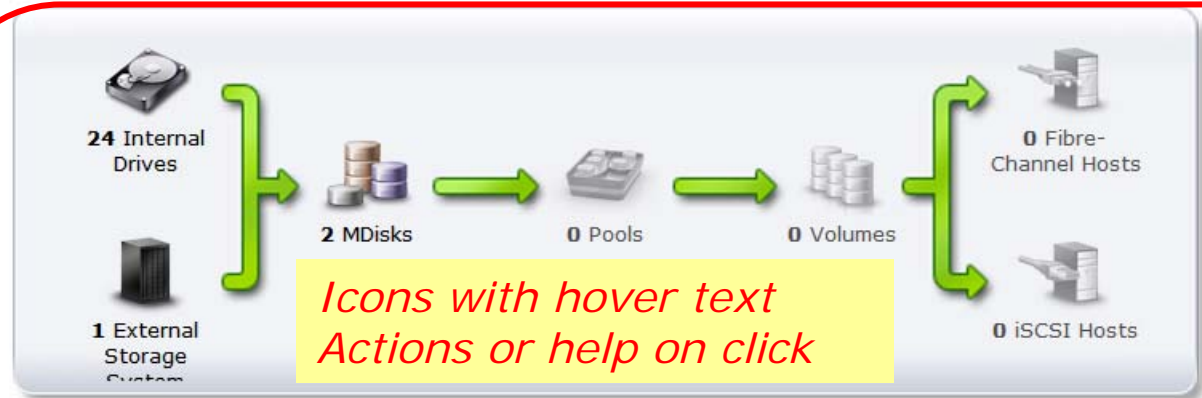
Quick navigation

Suggested Tasks

Action menus

-
-
-
-
-
-
-
-

Function icons



*Icons with hover text
Actions or help on click*

Overview

Watch e-Learning: [Overview](#)

Welcome!

The diagram represents all of the objects that need to be configured. To learn more about each object, click the icon in the diagram. The e-Learning modules include a tutorial of the steps that are required to complete the configuration of the objects, either select the associated task from Suggested Tasks or use the icons in the diagram.

Status details

Main display area with embedded help

- Connectivity
- Node Status
- External Storage
- Remote Partnerships

Overall Storage Allocation

0%

Free Capacity: 2.7 TB

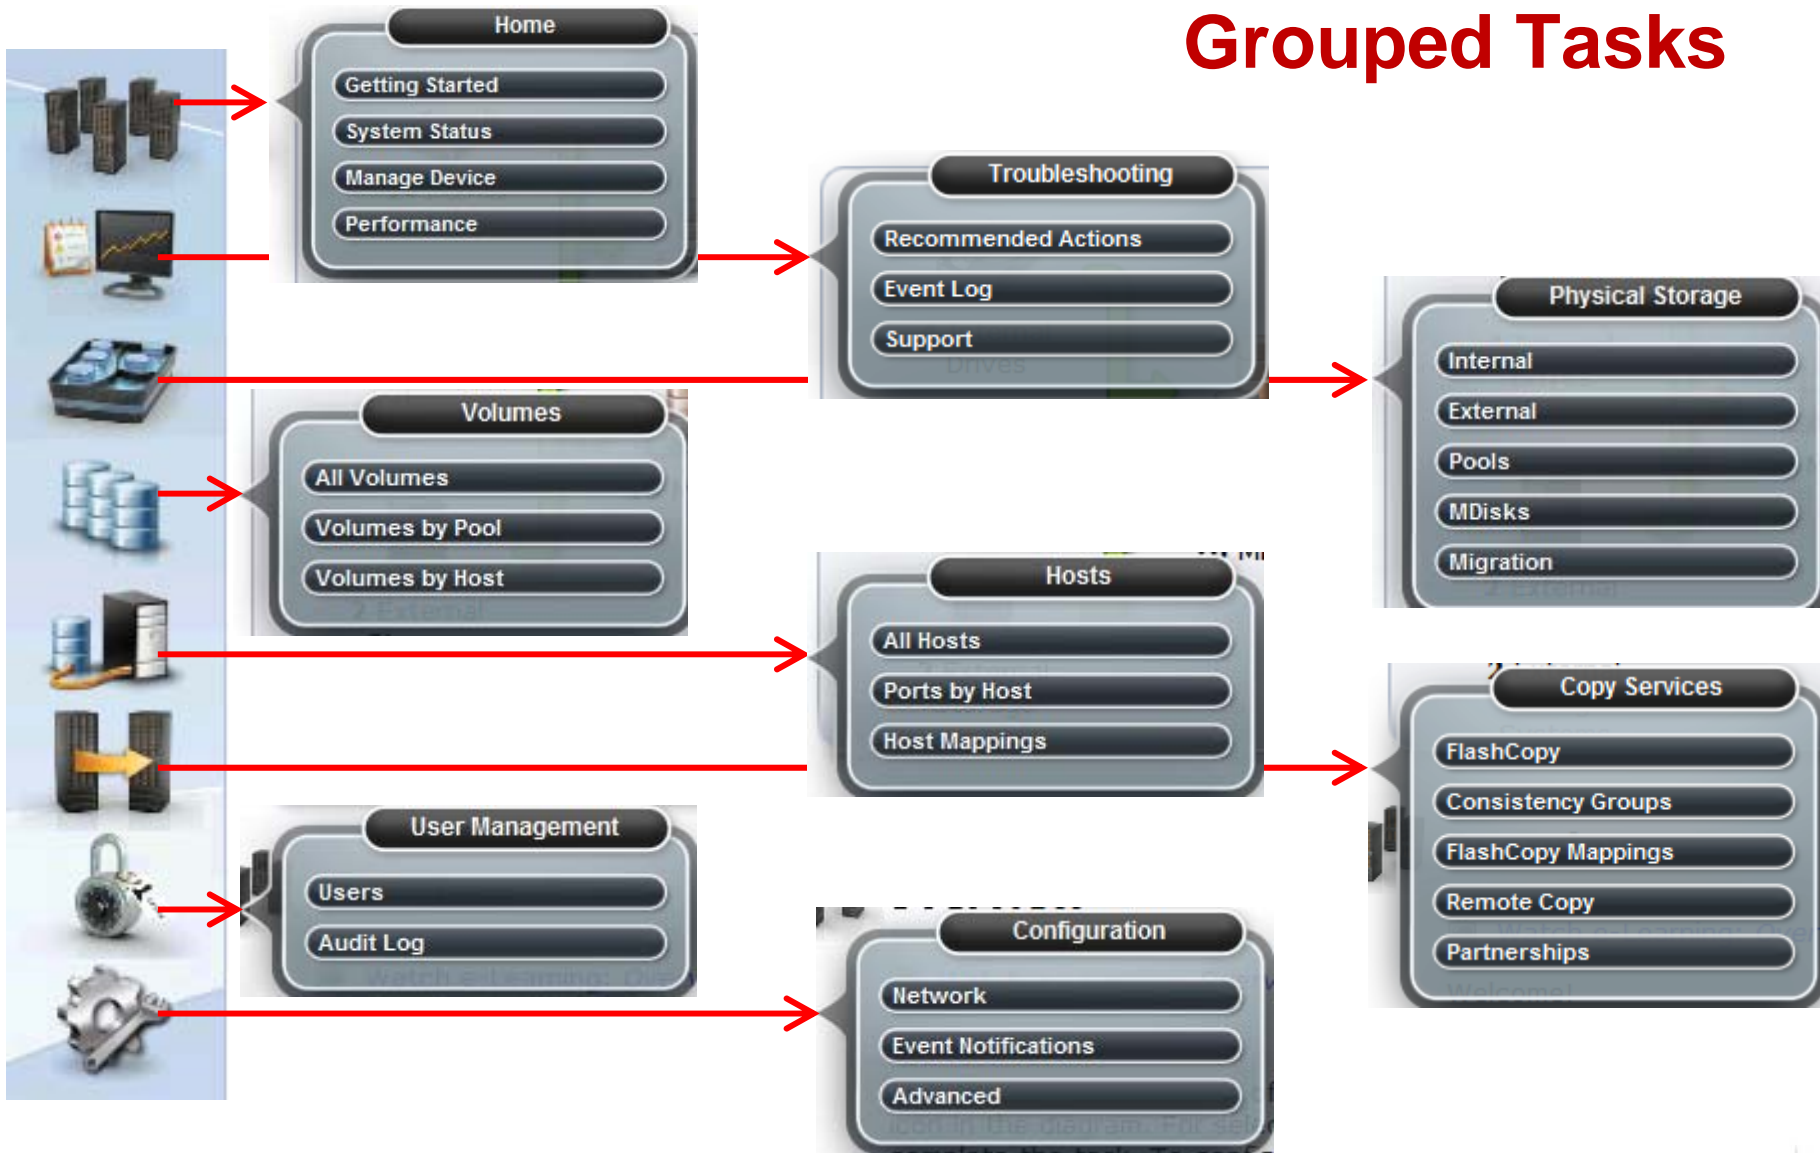
Physical Capacity: 2.7 TB

Status indicators

0 Running Tasks



Grouped Tasks



Status pods

The screenshot displays the IBM Storwize V7000 management interface. At the top, it says "IBM Storwize V7000" and "Welcome, superuser (3 users online)". The breadcrumb navigation shows "ATS_Oxford_1 > Home > System Status".

On the left, a vertical navigation menu contains icons for various system components. A yellow callout box labeled "Hover" points to this menu.

The main area shows a 3D rendering of a server rack labeled "ATS Oxford 1 (6.1.0.0)". A yellow callout box labeled "Click, move" points to the rack. A tooltip for the rack shows "Physical Capacity 200.0 GB".

To the right, a detailed view of "Enclosure 1" is shown. It is marked as "Online" with a green checkmark and has an "Identify" button. A yellow callout box labeled "Hover" points to this enclosure view. A tooltip for "Drive 16" shows it is "Online" with "278.9 GB SAS" and "10,000 RPM".

At the bottom, there are three status bars: "Connectivity" (checked), "0%" (with a progress bar), and "0 Running Tasks".



Persistent grids

The screenshot shows the IBM Storwize V7000 management console interface. At the top, it says "IBM Storwize V7000" and "Welcome, superuser (2 users online)". The main area is titled "ATS_Oxford_1 > Physical Storage > Internal".

On the left, there is a "Drive Type Filter" section with "All Internal" and "278.9 GB, SAS 10000 rpm" options. A yellow box labeled "Type filter" points to this section.

In the center, a "Configure Storage" button is visible above a drive icon labeled "278.9 GB, SAS 10000 rpm". A yellow box labeled "Drag-and-drop reorders columns" points to this area.

On the right, there is a "Capacity Allocation" section with a search filter and a table showing "MDisk Capacity", "Spare Capacity", and "Total Capacity". A yellow box labeled "Name filter" points to the search filter.

Below the drive details, there is a table of drives with columns for "Status", "MDisk Name", and "Drive". A context menu is open over the table, showing options like "Fix Error", "Take Offline", "Mark as...", "Identify", "Show Dependent V", and "Properties". A yellow box labeled "Context menu offers additional columns" points to this menu. The menu also shows a list of columns to display, including "Drive ID", "Capacity", "Use", "Status", "MDisk Name", "MDisk ID", "Member ID", "Enclosure ID", "Drive Slot", "Technology Type", "RPM", "Drive Type", "Show Select/Deselect All", and "Reset Grid Preferences".

At the bottom of the table, it says "Showing 24 drives | Selecting 3 drives". A yellow box labeled "Multi-select" points to this text.

At the bottom of the console, there are status indicators: "Connectivity" (checked), "0%", and "0 Running Tasks".



Presets reduce complexity

Viewing FlashCopy Mappings

Verify FlashCopy Mapping

Verify that the FlashCopy mapping attributes shown below are correct. If you change an attribute, click Back to return to the appropriate panel in the wizard. Click Finish to create the FlashCopy mapping.

Attribute	Value
FlashCopy Mapping Name	fcm-002
Source VDisk	vd-prod-0002
Target VDisk	vd-back-0002
Consistency Group Name	fcg-001
Background Copy Rate	NOCOPY
Cleaning Rate	50
Copy Type	Standard
Grain Size	256
IO Group	io_grp0
Delete when background copy completes	Disabled

< Back Next > Finish Cancel



STG_V7000 > Copy Services > FlashCopy ▾

Actions ▾

- New Snapshot
- New Clone
- New Backup

Status	Progress
--------	----------



Create MDisks

Configure Storage

All Internal

Configure Internal Storage *Step 1 of 1*

Use this wizard to allocate RAID arrays to storage pools. After this configuration wizard completes, you can create volumes from these storage pools.

Storage Found:
(0 drives) 278.9 GB, SSD, io_grp0
(7 drives) 418.7 GB, SAS, 10000 rpm, io_grp0

- Use the recommended configuration**
Select this option to configure all available drives based on recommended values for the RAID level and drive class. The recommended configuration uses all the drives to build arrays that are protected with the appropriate amount of spare drives.
- Select a different configuration**

Configuration Summary:

- 0 x **SSD RAID-1** (278.9 GB, SSD, io_grp0):
✖ There are not enough drives to satisfy the target number of drives that is needed for the configuration.
- 1 x **Basic RAID-5** (418.7 GB, SAS, 10000 rpm, io_grp0):
7 drives
0 Hot Spares
0 Unconfigured Drives

Finish **Cancel**

Configure Internal Storage *Step 1 of 2*

Use this wizard to allocate RAID arrays to storage pools. After this configuration wizard completes, you can create volumes from these storage pools.

Storage Found:
(0 drives) 278.9 GB, SSD, io_grp0
(7 drives) 418.7 GB, SAS, 10000 rpm, io_grp0

- Use the recommended configuration**
Select this option to configure all available drives based on recommended values for the RAID level and drive class. The recommended configuration uses all the drives to build arrays that are protected with the appropriate amount of spare drives.
- Select a different configuration**

Drive Class: (7) 418.7 GB, SAS, 10000 rpm, io_grp0

Preset: Basic RAID-5

- Automatically configure spares**
- Optimize for Performance**
- Optimize for Capacity**

7 **Number of drives to provision**

Configuration Summary:

- 1 x **Basic RAID-5** (418.7 GB, SAS, 10000 rpm, io_grp0):
7 drives
0 Hot Spares
0 Unconfigured Drives

Next > **Cancel**

Create Pools & Volumes



Create Storage Pool (auto-config)

New Pool Action

Welcome!

No storage pools have been defined on your system.

To create a storage pool, click New Pool.



Create Storage Pool Step 1 of 2

Choose Icon

Create Storage Pool Step 2 of 2

Pool Name (optional):

Select the MDisk to include in the storage pool (optional):

Name	Status	Capacity	Storage System	LUN
mdisk1	Online	23.6 GB	controller1	0000000000000001
mdisk2	Online	23.6 GB	controller1	0000000000000002
mdisk3	Online	23.6 GB	controller1	0000000000000003
mdisk4	Online	23.6 GB	controller1	0000000000000004

Showing 4 mdisks | Selecting 1 mdisk

Back Finish Cancel

Create Volumes

New Volume Action

Welcome!

It looks like you have no volumes defined yet.

To create new volumes in storage pools, click New Volume.



New Volume

Select a Preset

Generic Thin Provision Mirror Thin Mirror

Select a Pool

Primary Pool: mdiskgrp1 Edit

Select Names and Sizes

Volume Name	Size	Unit	Buttons
Volume1	5	GB	+ -
Volume2	5	GB	+ - X

Summary: 5 volumes, 25.0 GB total, 927.3 GB free in pool



Advanced... Create Create and Map to Host Cancel



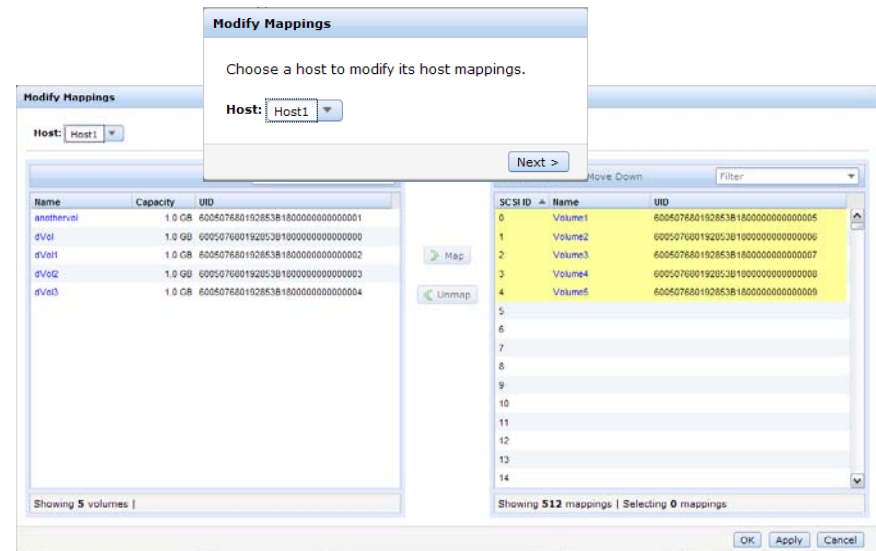
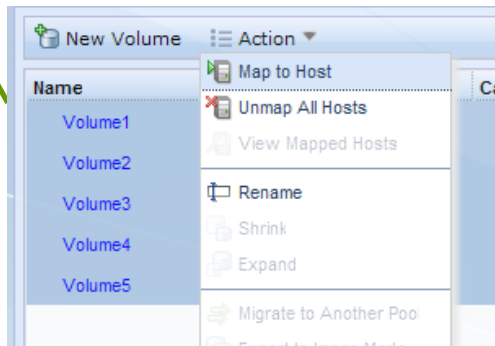
Create Hosts and Map Volumes



Create Host



Map Volumes to Host

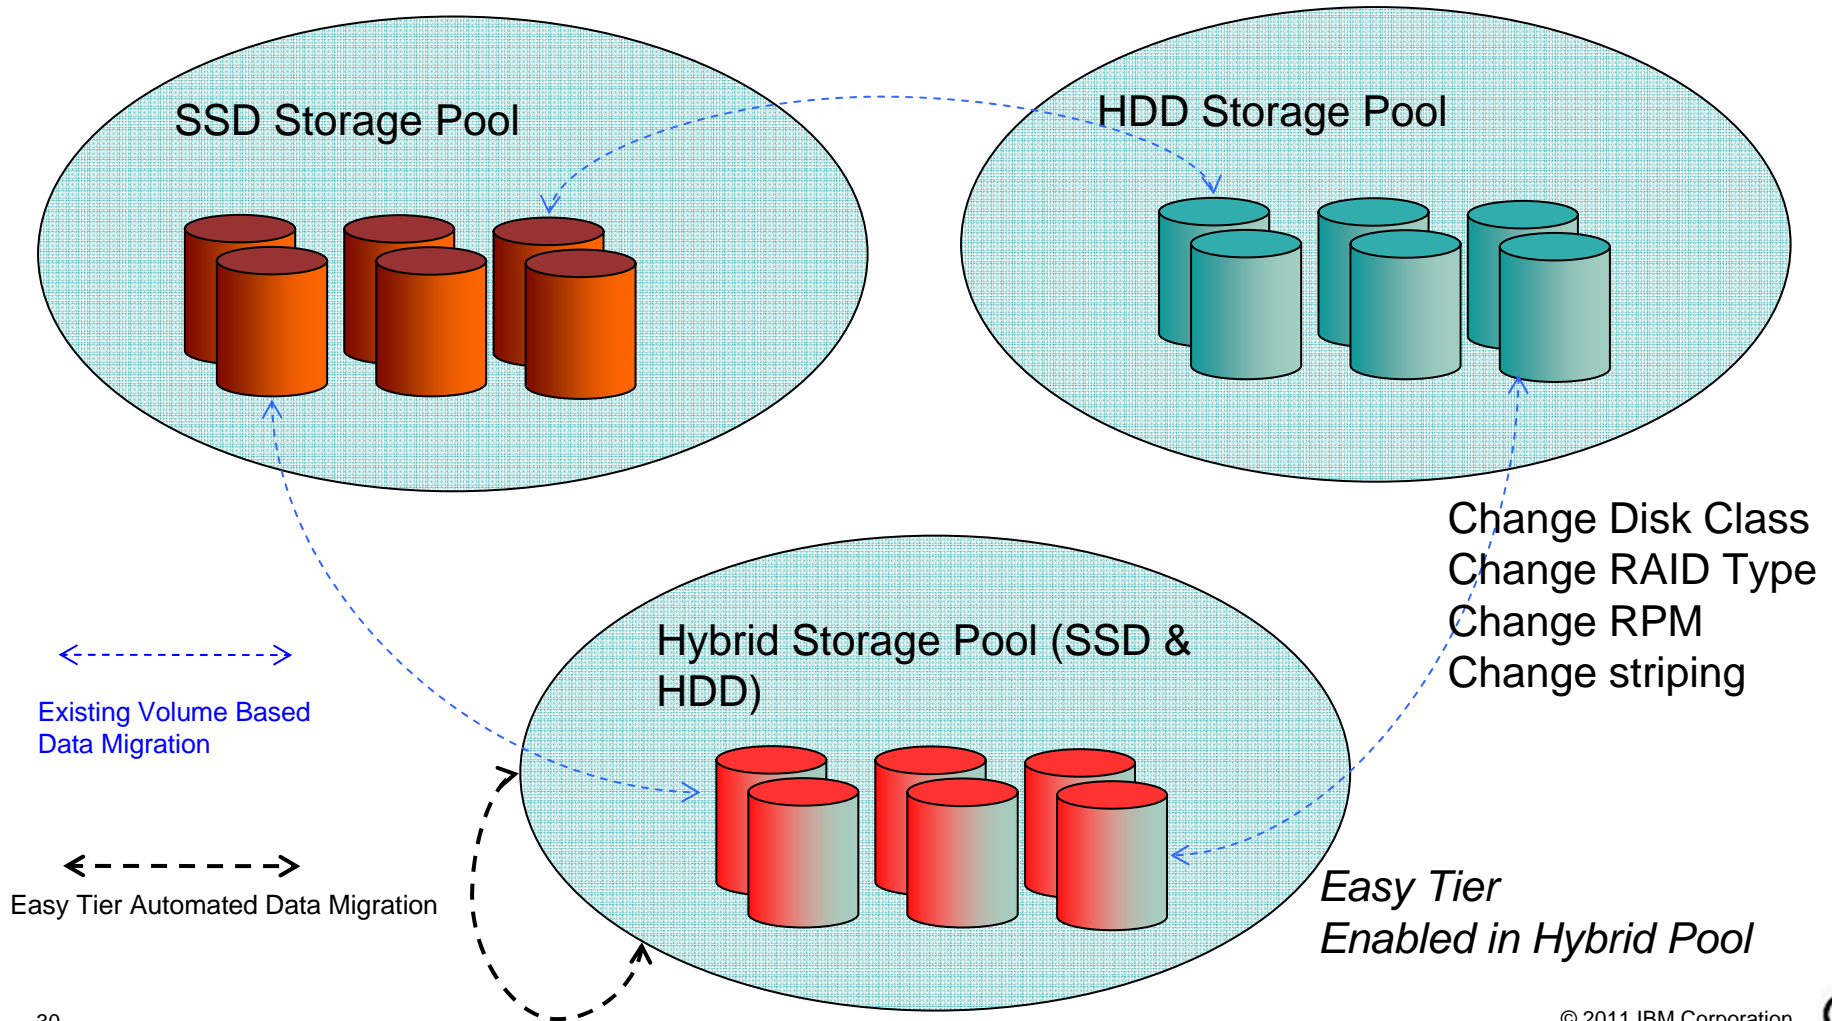


Easy Tier

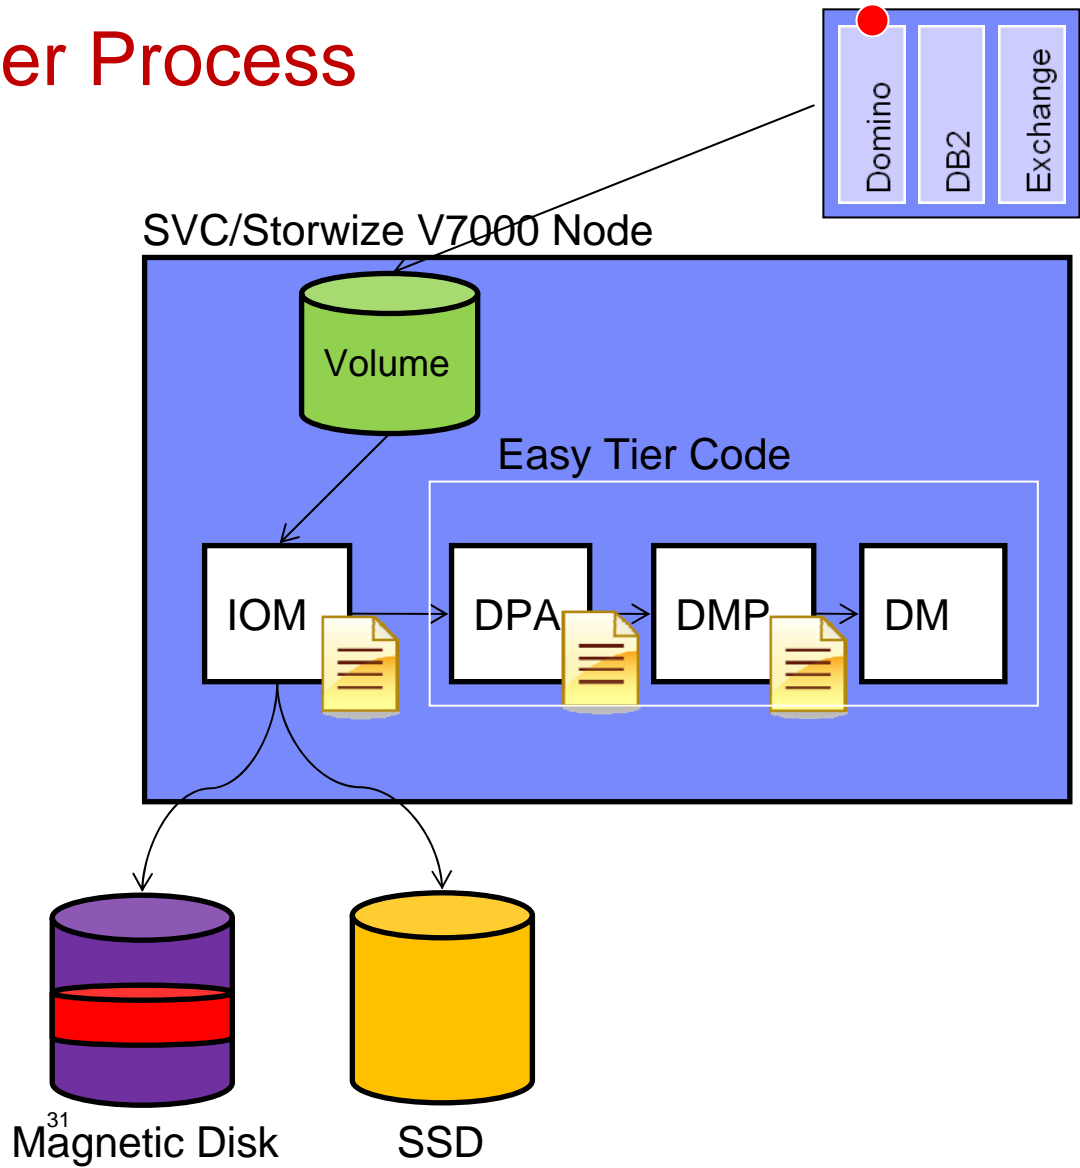
- SSDs (Solid State Drives) are still expensive
- Putting entire volumes on an SSD is wasteful



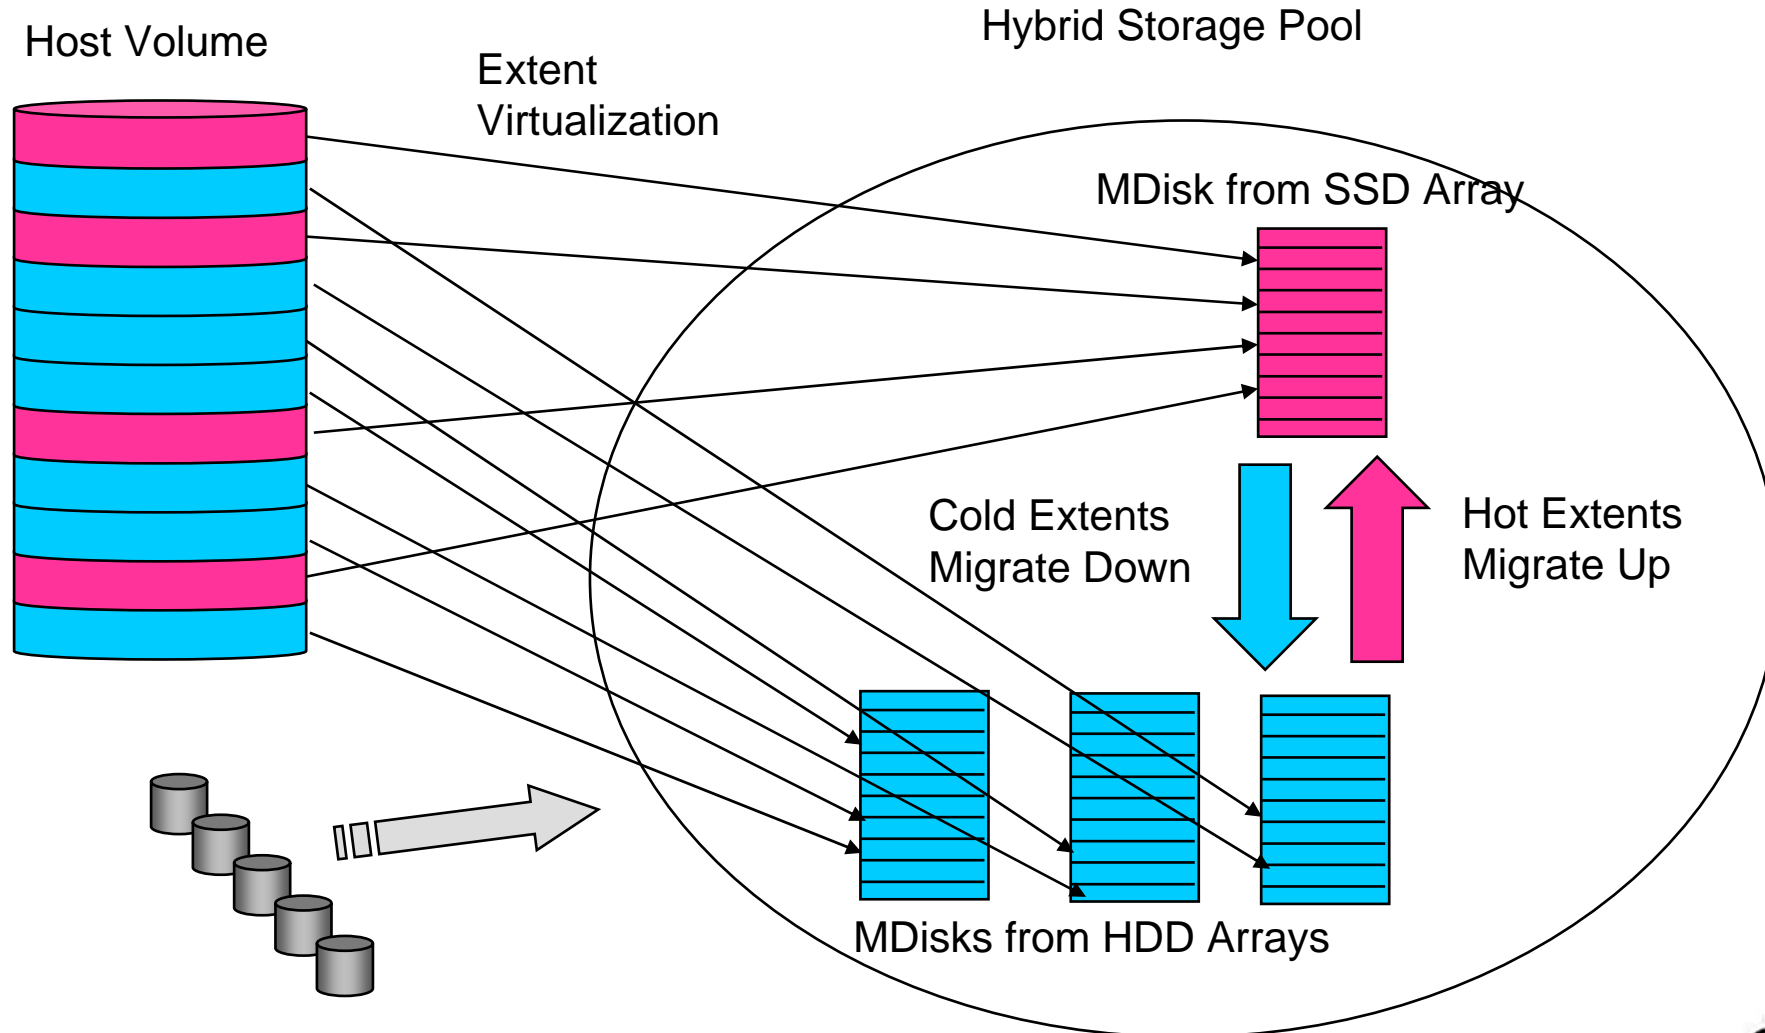
Data Migration by Volume and Sub-Volume



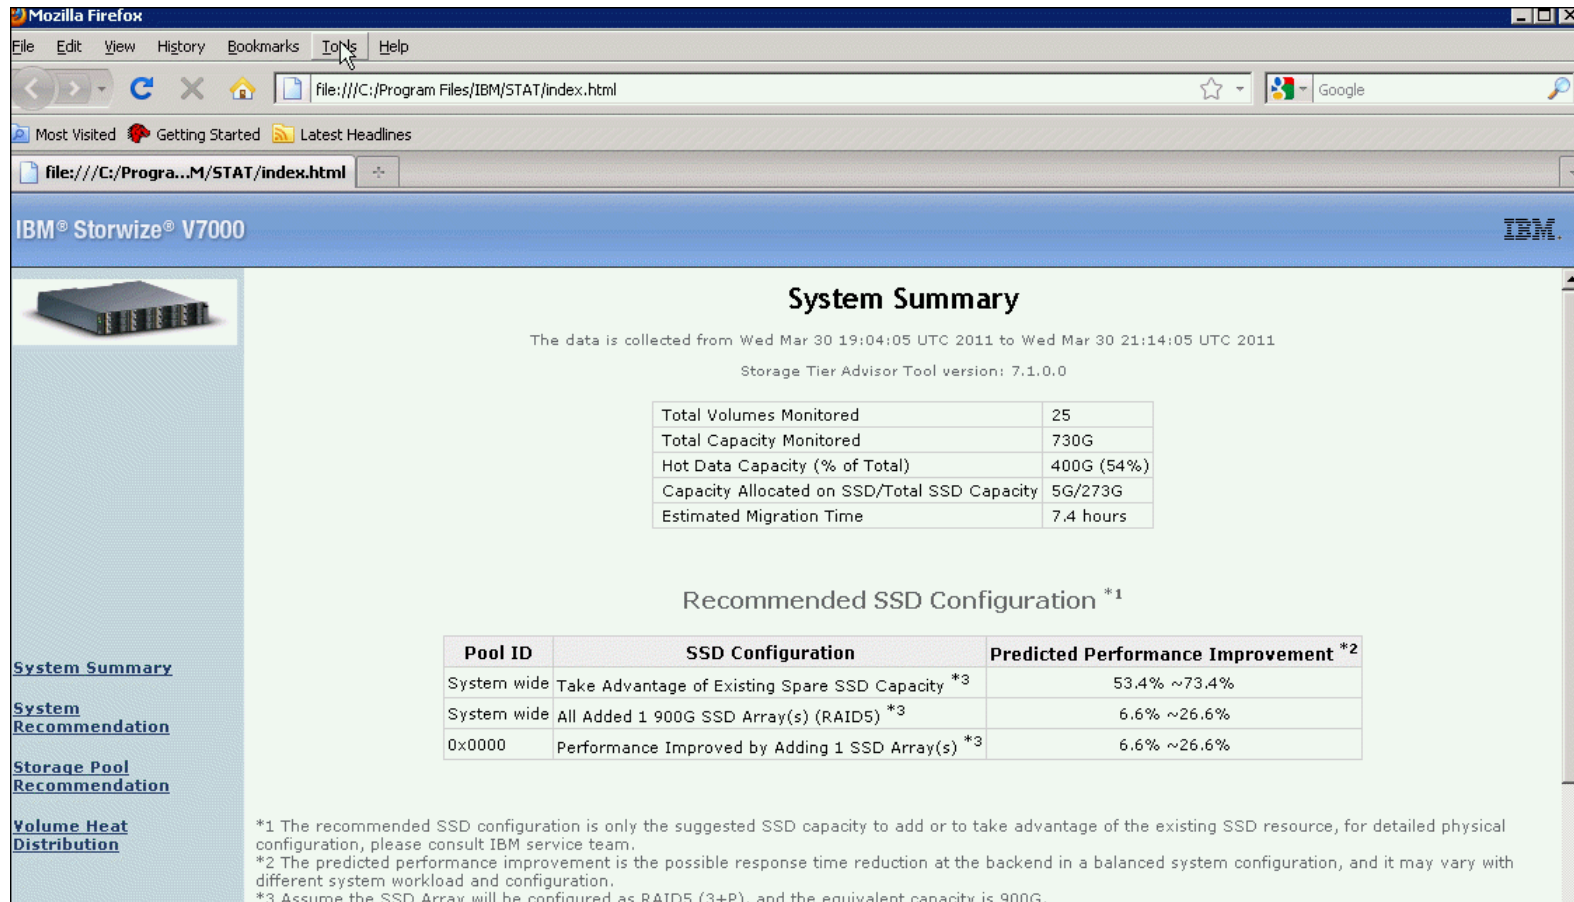
Easy Tier Process



Easy Tier – Automated Data Relocation



Easy Tier – Summary Report



IBM® Storwize® V7000

System Summary

The data is collected from Wed Mar 30 19:04:05 UTC 2011 to Wed Mar 30 21:14:05 UTC 2011
Storage Tier Advisor Tool version: 7.1.0.0

Total Volumes Monitored	25
Total Capacity Monitored	730G
Hot Data Capacity (% of Total)	400G (54%)
Capacity Allocated on SSD/Total SSD Capacity	5G/273G
Estimated Migration Time	7.4 hours

Recommended SSD Configuration ^{*1}


Pool ID	SSD Configuration	Predicted Performance Improvement ^{*2}
System wide	Take Advantage of Existing Spare SSD Capacity ^{*3}	53.4% ~73.4%
System wide	All Added 1 900G SSD Array(s) (RAID5) ^{*3}	6.6% ~26.6%
0x0000	Performance Improved by Adding 1 SSD Array(s) ^{*3}	6.6% ~26.6%

^{*1} The recommended SSD configuration is only the suggested SSD capacity to add or to take advantage of the existing SSD resource, for detailed physical configuration, please consult IBM service team.
^{*2} The predicted performance improvement is the possible response time reduction at the backend in a balanced system configuration, and it may vary with different system workload and configuration.
^{*3} Assume the SSD Array will be configured as RAID5 (3+P), and the equivalent capacity is 900G.



Easy Tier - Heat Distribution Report

IBM® System® Storage San Volume Controller
IBM



Volume Heat Distribution

The data is collected from Tue Aug 24 21:09:03 UTC 2010 to Tue Aug 24 22:14:03 UTC 2010

Volume ID *1	Copy ID	Pool ID	Configured Size *2	Capacity on SSD *3	Heat Distribution *4
0x0003	0x0000	0x0000	500.0G	0.0G	485.7G 14.3G
0x0002	0x0000	0x0000	500.0G	0.0G	485.2G 14.8G
0x0001	0x0000	0x0000	500.0G	0.0G	487.7G 12.3G
0x0000	0x0000	0x0000	500.0G	0.0G	488.2G 11.8G
0x0004	0x0000	0x0000	500.0G	0.0G	488.3G 11.7G
0x0005	0x0000	0x0000	500.0G	0.0G	487.9G 12.1G
0x0006	0x0000	0x0000	500.0G	0.0G	485.9G 14.1G
0x0007	0x0000	0x0000	500.0G	0.0G	486.6G 13.4G
0x0008	0x0000	0x0000	500.0G	0.0G	487.3G 12.7G
0x0009	0x0000	0x0000	500.0G	0.0G	487.0G 13.0G
0x000a	0x0000	0x0000	500.0G	0.0G	484.8G 15.2G
0x000b	0x0000	0x0000	500.0G	0.0G	485.2G 14.8G

20 Entries Per Page

|< << >> >|

Displaying Page 1 of 1

*1. Volume ID is the ID to represent the LUN.
 *2. The configured capacity of the volume.
 *3. The 'Capacity on SSD' column represents the amount of data that has already migrated to SSD drives.
 *4. The 'Heat Distribution' column shows the heat distribution of the data in this volume. The blue portion of the bar represents the capacity of the cold data, and the red portion represents the capacity of the hot data.

[System Summary](#)


[System Recommendation](#)

[Storage Pool Recommendation](#)

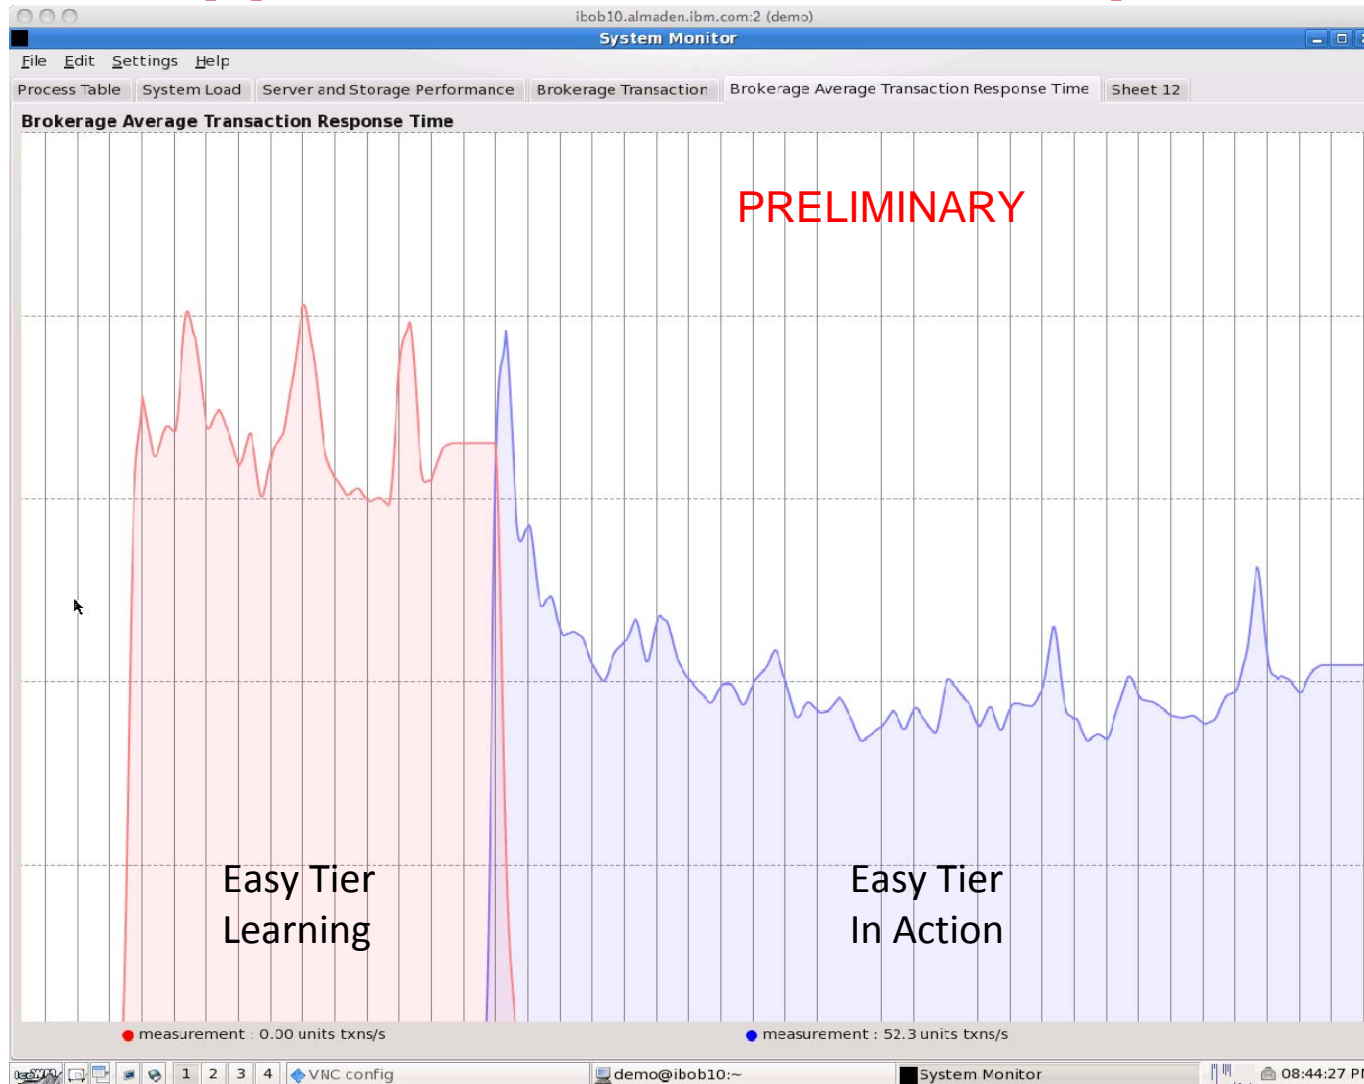
[Volume Heat Distribution](#)



Easy Tier Application Transaction Improvement



Easy Tier Application Transaction Improvement



Pulse2011



Enhancements v6.2

Real-time Performance Statistics




VAAI: A tale of three primitives

Block zero
Full Copy
Smarter locking

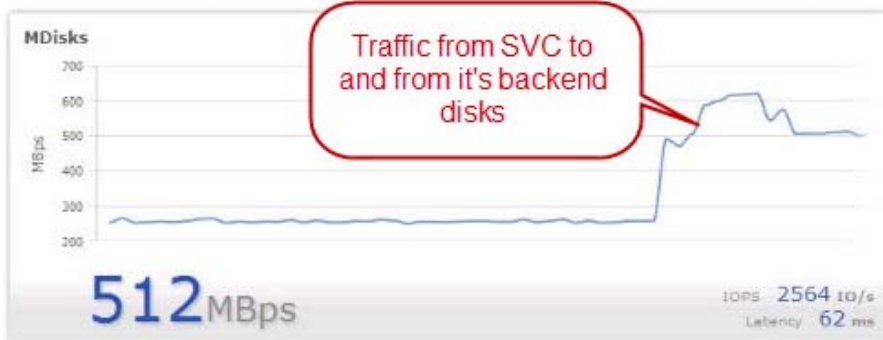
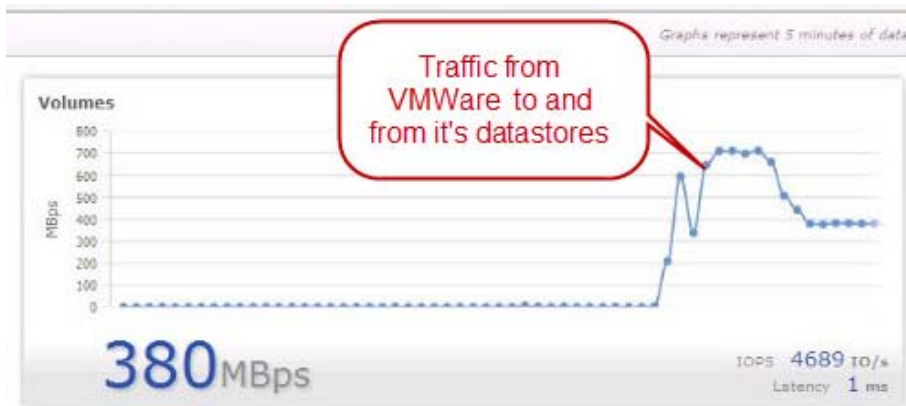


Storage VMotion done right

Before VAAI



After VAAI



Storwize V7000 Plug-In for VMware vCenter

The screenshot shows the vSphere Client interface for a Storwize V7000 storage system. The 'Storage Systems' table lists several systems, with 'Storwize_V7000' selected. The 'Storage Pools' section shows a pool named 'mdiskgrp0' with 1939 GB usage and a capacity of 1951 GB. The 'Details' section for the selected system shows the following information:

Property	Value
System ID	000002006440C4FC
System Name	StorwizeCluster
User Name	admin
Version	6.2.0.0 (build 35.2.1103250000)
System IP	9.119.41.113



Clustered Storwize V7000 Control Enclosures

Base Storwize V7000

Maximums

Hosts (FC): 256 FC ports: 8x8Gb
 Hosts (iSCSI) 64 Ethernet ports: 4x1Gb
 Capacity: 24TB (optional) 4x10Gb
 Cache: 16GB Volumes: 2048
 Drives: 12 LFF or 24 SFF
 Disk types: SAS, NL-SAS, SSD



Expand

Maximums

Hosts (FC): 512 FC ports: 16x8Gb
 Hosts (iSCSI) 128 Ethernet ports: 8x1Gb
 Capacity: 48TB (optional) 8x10Gb
 Cache: 32GB Volumes: 4096
 Drives: 24LFF or 48 SFF or mix
 Disk types: SAS, NL-SAS, SSD




Cluster

Expand

+ External Virtualization for more capacity

Maximums

Hosts (FC): 256 FC ports: 8x8Gb
 Hosts (iSCSI) 64 Ethernet ports: 4x1Gb
 Capacity: 240TB (optional) 4x10Gb
 Cache: 16GB Volumes: 2048
 Drives: 120 LFF or 240 SFF or mix
 Disk types: SAS, NL-SAS, SSD



Maximums

Hosts (FC): 512 FC ports: 16x8Gb
 Hosts (iSCSI) 128 Ethernet ports: 8x1Gb
 Capacity: 480TB (optional) 8x10Gb
 Cache: 32GB Volumes: 4096
 Drives: 240 LFF or 480 SFF or mix
 Disk types: SAS, NL-SAS, SSD




Cluster

History of SVC Node Development

Model	Announce Date	Cache per node	FC Port Speed	System x base	CPU
2145-4F2	June 2003	4 GB DDR1 1 Channel	2 Gbps	x335	2 x 2.8 Ghz Xeon
2145-8F2	25 Oct 2005	8 GB DDR2 2 Channels	2 Gbps	x336	2 x 3 Ghz Xeon
2145-8F4	23 May 2006	8 GB DDR2 2 Channels	4 Gbps	x336	2 x 3 Ghz Xeon
2145-8A4	28 Oct 2008	8 GB PC2-5300 2 Channels	4 Gbps	x3250	3 Ghz Dual Core Xeon E3110
2145-8G4	22 May 2007	8 GB DDR2 4 Channels	4 Gbps	X3550	2 x 2.33Ghz Dual Core Xeon 5160 2.5Ghz Quad Core Xeon 5420
2145-CF8	20 Oct 2009	24 GB DDR3-1333 6 Channels (6 x 4 GB)	8 Gbps	x3550 M2	2.4Ghz Quad Core Xeon 5500
2145-CG8	May 9 2011	24 GB DDR3-1333 (3 x 8 GB)	8 Gbps	x3550 M3	2.53Ghz Quad Core Xeon 5600 2.4Ghz Hex Core Xeon 5645

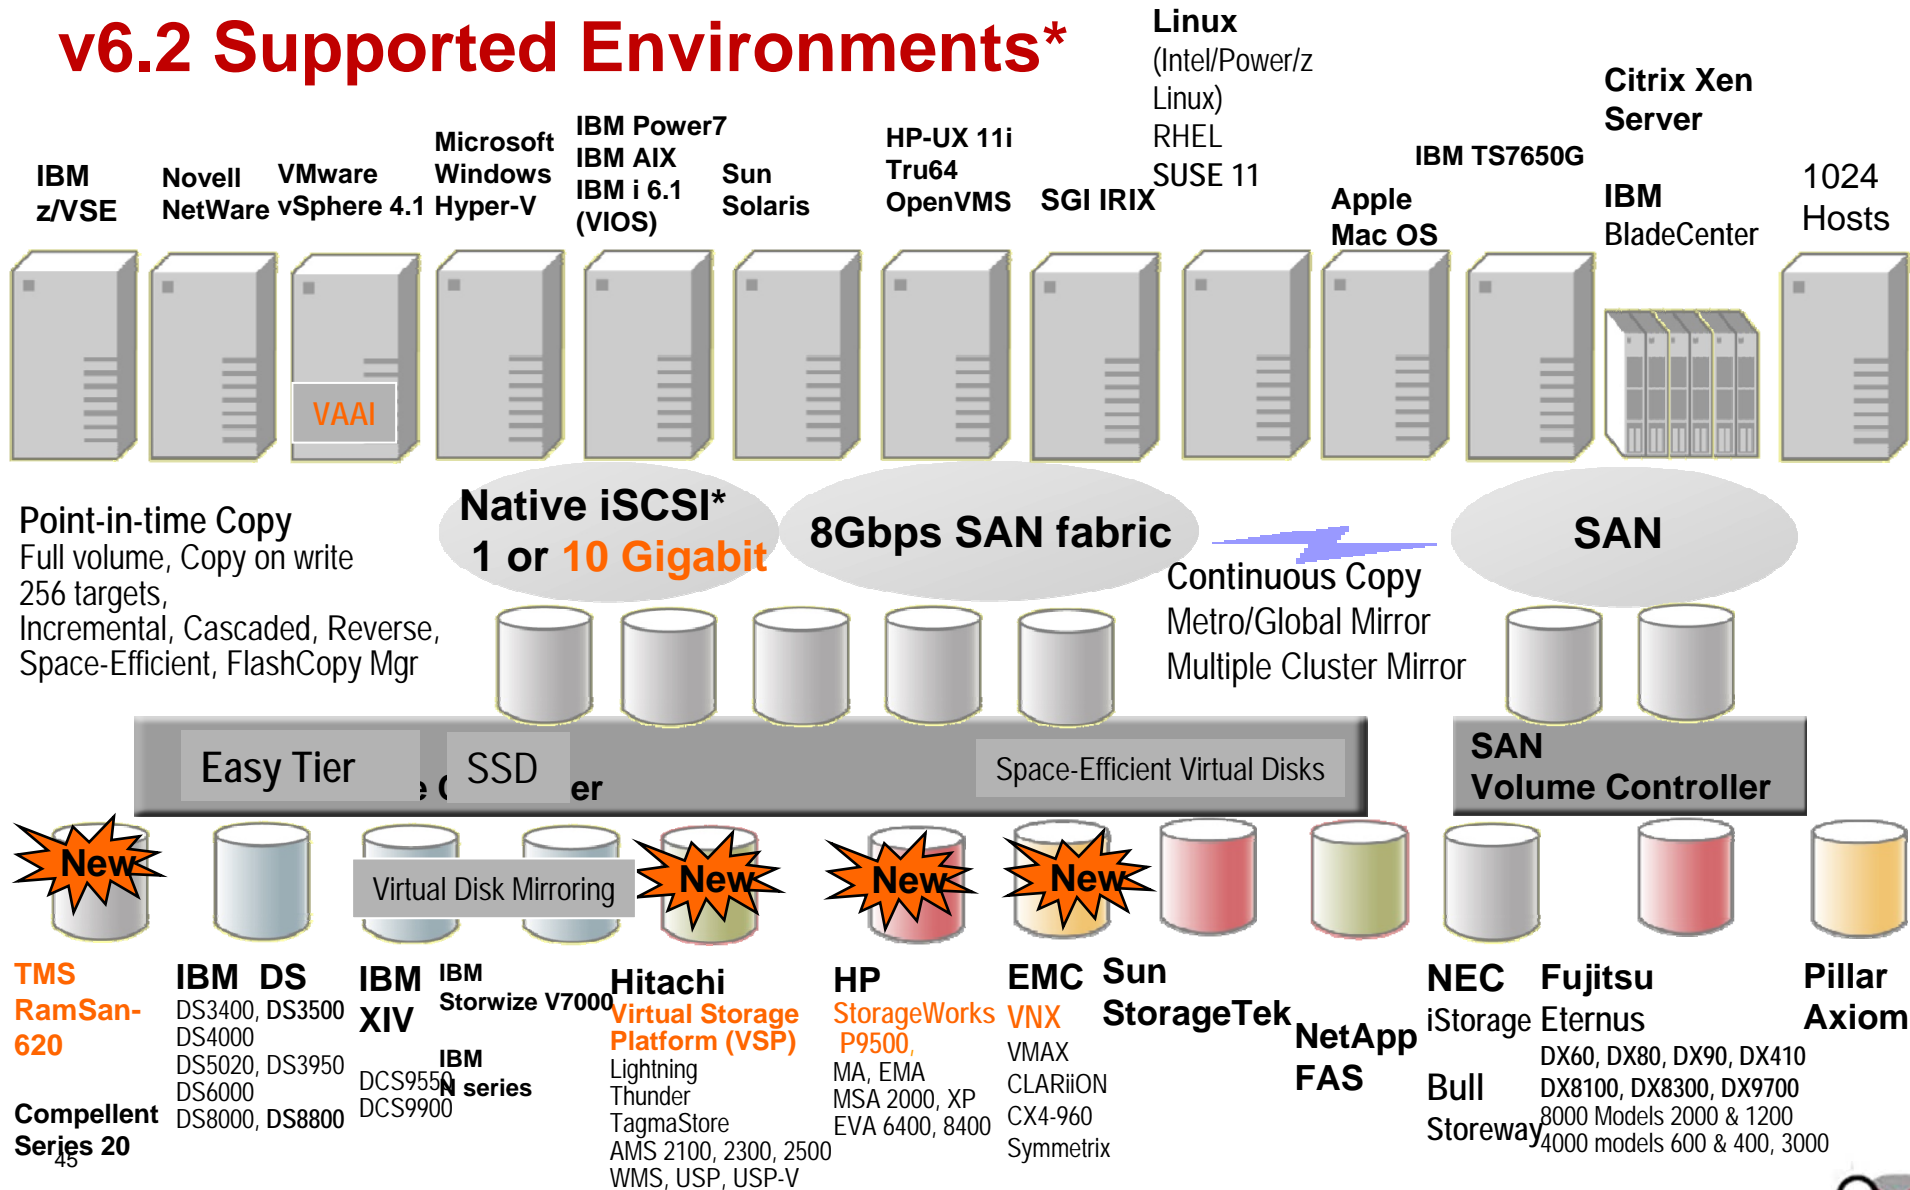
10 Gbps iSCSI Option



New Node CG8 with optional 10 Gbps iSCSI



v6.2 Supported Environments*



What about software integration?

- IBM Tivoli FlashCopy Manager (FCM)
- IBM Tivoli Storage Productivity Center (TPC)
- IBM TPC for Replication (TPC-R)
- IBM Tivoli Storage Manager FastBack
- IBM Systems Director
- Microsoft Operations Manager (SCOM)
- Microsoft VSS
- VMWare vCenter Plugin





So what does the future hold?



Thank you!

Email:

anthonyv@au.ibm.com

Twitter:

<http://twitter.com/#!/aussiestorblog>

Blog:

<https://aussiestorageblog.wordpress.com/>



Trademarks and disclaimers

© Copyright IBM Australia Limited 2011 ABN 79 000 024 733 © Copyright IBM Corporation 2011 All Rights Reserved.
TRADEMARKS: IBM, the IBM logos, ibm.com, Smarter Planet and the planet icon are trademarks of IBM Corp registered in many jurisdictions worldwide. Other company, product and services marks may be trademarks or services marks of others. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml

The customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer.

Information concerning non-IBM products was obtained from a supplier of these products, published announcement material, or other publicly available sources and does not constitute an endorsement of such products by IBM. Sources for non-IBM list prices and performance numbers are taken from publicly available information, including vendor announcements and vendor worldwide homepages. IBM has not tested these products and cannot confirm the accuracy of performance, capability, or any other claims related to non-IBM products. Questions on the capability of non-IBM products should be addressed to the supplier of those products.

All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

Some information addresses anticipated future capabilities. Such information is not intended as a definitive statement of a commitment to specific levels of performance, function or delivery schedules with respect to any future products. Such commitments are only made in IBM product announcements. The information is presented here to communicate IBM's current investment and development activities as a good faith effort to help with our customers' future planning.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

Prices are suggested U.S. list prices and are subject to change without notice. Starting price may not include a hard drive, operating system or other features. Contact your IBM representative or Business Partner for the most current pricing in your geography.

Photographs shown may be engineering prototypes. Changes may be incorporated in production models.

