



Increase Savings with File Systems, Thin Provisioning and Storage Reclamation

John Harker, Senior Product Marketing Manager, Hitachi Data Systems

Chakri Avala, Principal Product Manager, Symantec Corporation

January 20, 2010

Increase Savings with File Systems, Thin Provisioning and Storage Reclamation

Learn about Write Same, a new feature of Hitachi Dynamic Provisioning that uses active, ongoing reclamation of file system storage. It is now available on the Hitachi Adaptable Modular Storage 2000 family and Hitachi Universal Storage Platform V and VM. Used in conjunction with Symantec Storage Foundation's VxFS file system and the Veritas Thin Provisioning API, Hitachi Dynamic Provisioning now further increases storage savings. Hitachi Data Systems is one of the first to release support for this significant technology advancement. Come to this WebTech and learn from Hitachi and Symantec experts how to take advantage of this feature.

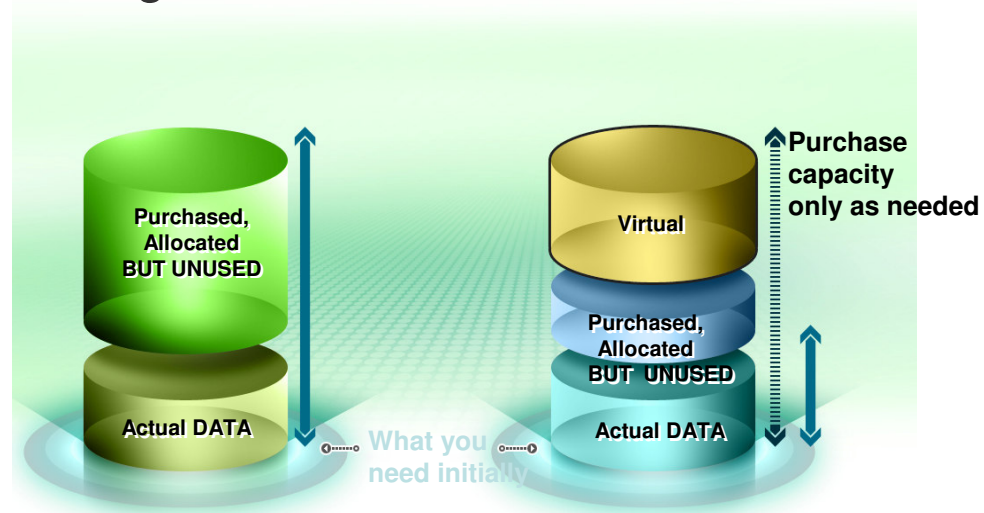
Attend this session and learn:

- When, why and how to integrate Symantec Storage Foundations server-based, thin provisioning capabilities with Hitachi storage-based thin provisioning.
- How the Veritas Thin Provisioning API and Write Same work in thin provisioned environments to enable active file system storage reclamation.
- About the evolution of open standards by SNIA and INCITS for the integration of file system and storage system reclamation.

Background: Why Dynamic Provisioning?

Storage Acquisition and Allocation Challenges


- High cost of storage
- Cumbersome provisioning
- Overprovisioning
- Difficult optimization
- Wasted storage



Dynamic Provisioning = Efficient Storage Allocation


Use Only What You Need Where You Need It When You Need It

HITACHI
Inspire the Next



Dynamic Provisioning Capabilities

- Simplify provisioning
- Provision only what is used
- Automate performance optimization
- Increase performance
- Automatic load balancing
- Storage Reclamation
- Replication Savings



Your Benefits

- Reduce storage expense
- Reduced operational expense
- IT agility



Major Retail Chain Reduces Storage Costs While Improving Performance

A large enterprise retail company reported that improved performance and deferred storage purchases were the key benefits of implementing Hitachi Dynamic Provisioning.

Source:

Senior IT Manager, Large Enterprise Retail Company

TechValidate
TVID: 37E-CB3-5B9

Background: Dynamic Provisioning and Storage Foundations make a good couple

Synergies:

- Storage Foundation VxFS is very thin friendly. Utilizes space efficiently and minimizes unnecessary allocation
- Storage Foundation SmartMove enables online migrations from thick to thin environments and ensures thin storage is allocated only for the used data blocks(mirroring doesn't cause underlying volumes to become fully populated).
- Dynamic Provisioning reclamation support enables Storage Foundation thin storage reclamation fully automated and fully transparent to the host configuration and the applications.
- Hitachi storage-system based volume mirroring offloads work from servers and offers further thin volume economic benefits for replication



Confidence in a connected world.

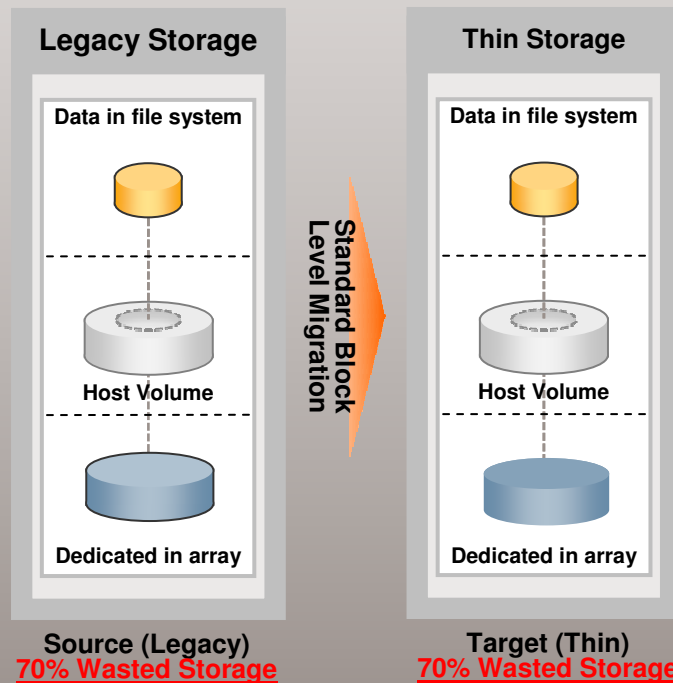
Get Thin with SmartMove™

Automatically Reclaim ALL Unused Storage with SF SmartMove™



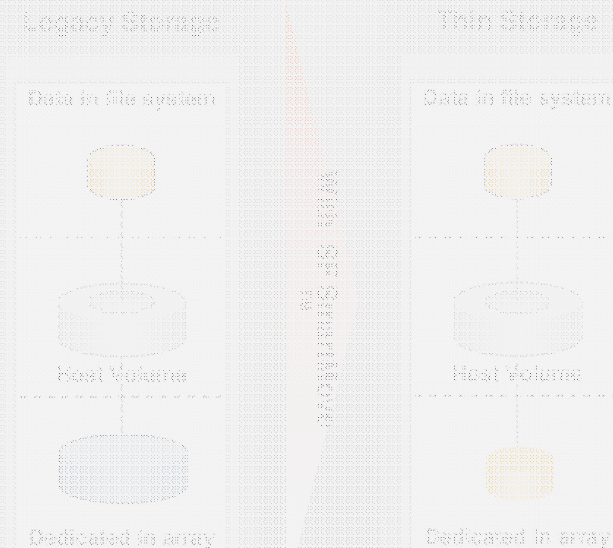
Standard Block Level Migration

- Migrating online typically involves FULL 'block level' copy to Thin Devices
- This applies to ALL block level copy methods (Host, Network, Appliance, Array)



With Storage Foundation SmartMove™

- Online reclamation as you migrate to Thin Storage
- Hardware independent and works with ALL thin storage vendors
- Industry's only solution for online thick to thin conversions




Source (Legacy) 70% Wasted Storage
Target (Thin) 70% Wasted Storage

SmartMove Makes Volume Migrations Smarter



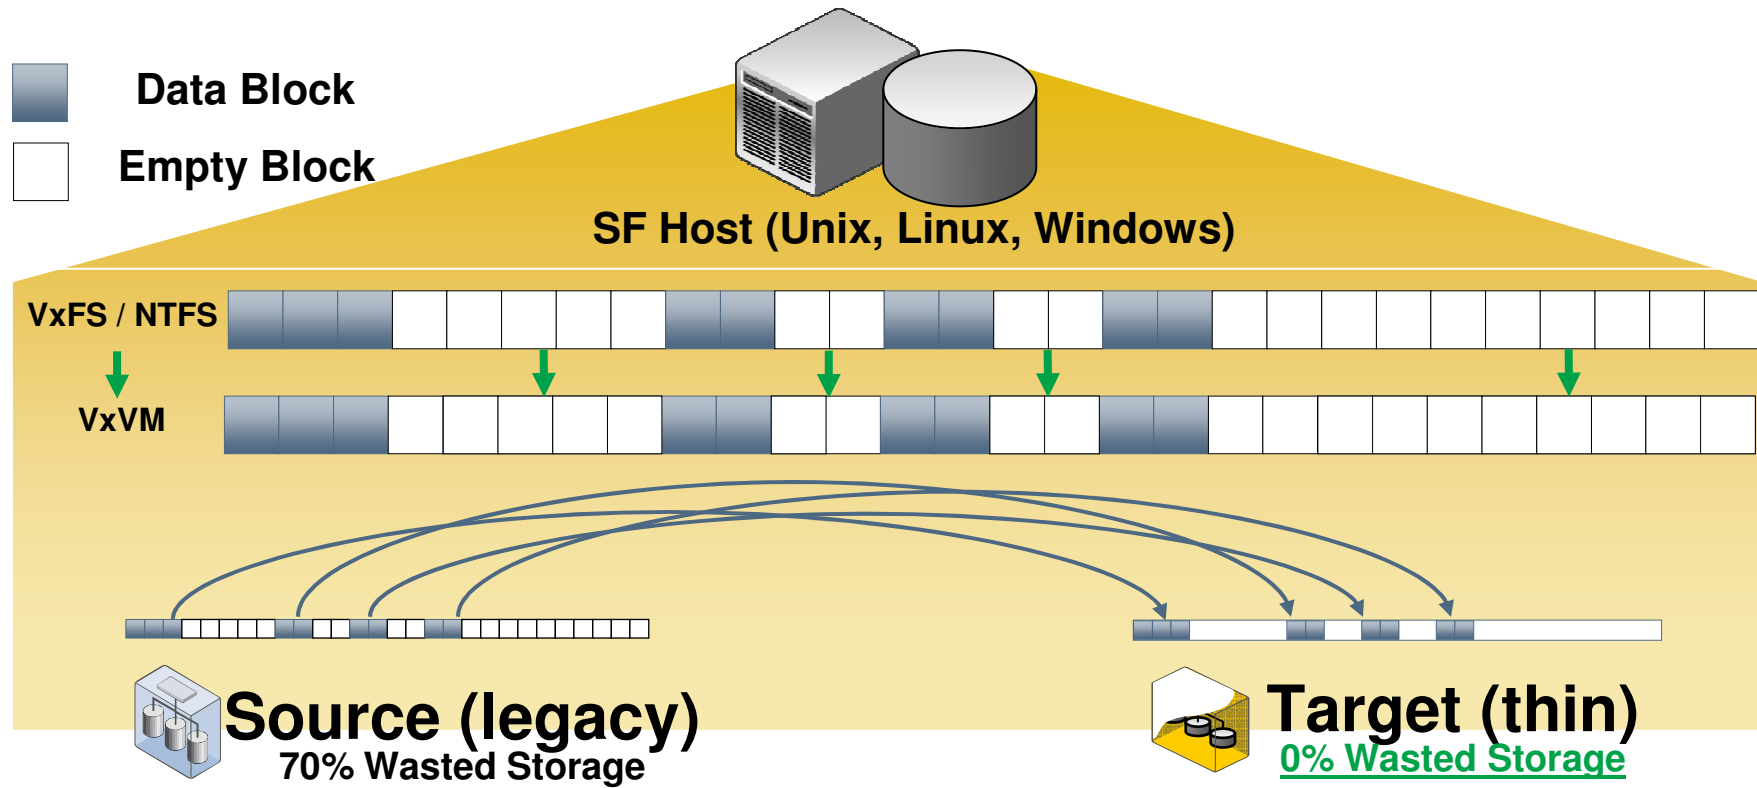
- SmartMove requires a mounted file system on the VxVM volume
 - VxFS on Unix/Linux and NTFS on Windows



SmartMove Makes Volume Migrations Smarter



- SmartMove requires a mounted file system on the VxVM volume
 - VxFS on Unix/Linux and NTFS on Windows
- VxVM only copies the blocks that contain user data
- Not writing unused data = No unnecessary allocation on the thin LUN



SmartMove & TP Array Support



Vendor	SmartMove
EMC CX	Yes
EMC DMX	Yes
HDS USP-V	Yes
HDS AMS	Yes
3PAR	Yes
IBM XIV	Yes
Netapp	Yes
HP EVA	Yes
Fujitsu	Yes
Dell EqualLogic	Yes



symantec™

Confidence in a connected world.

Storage Foundation Thin Reclamation

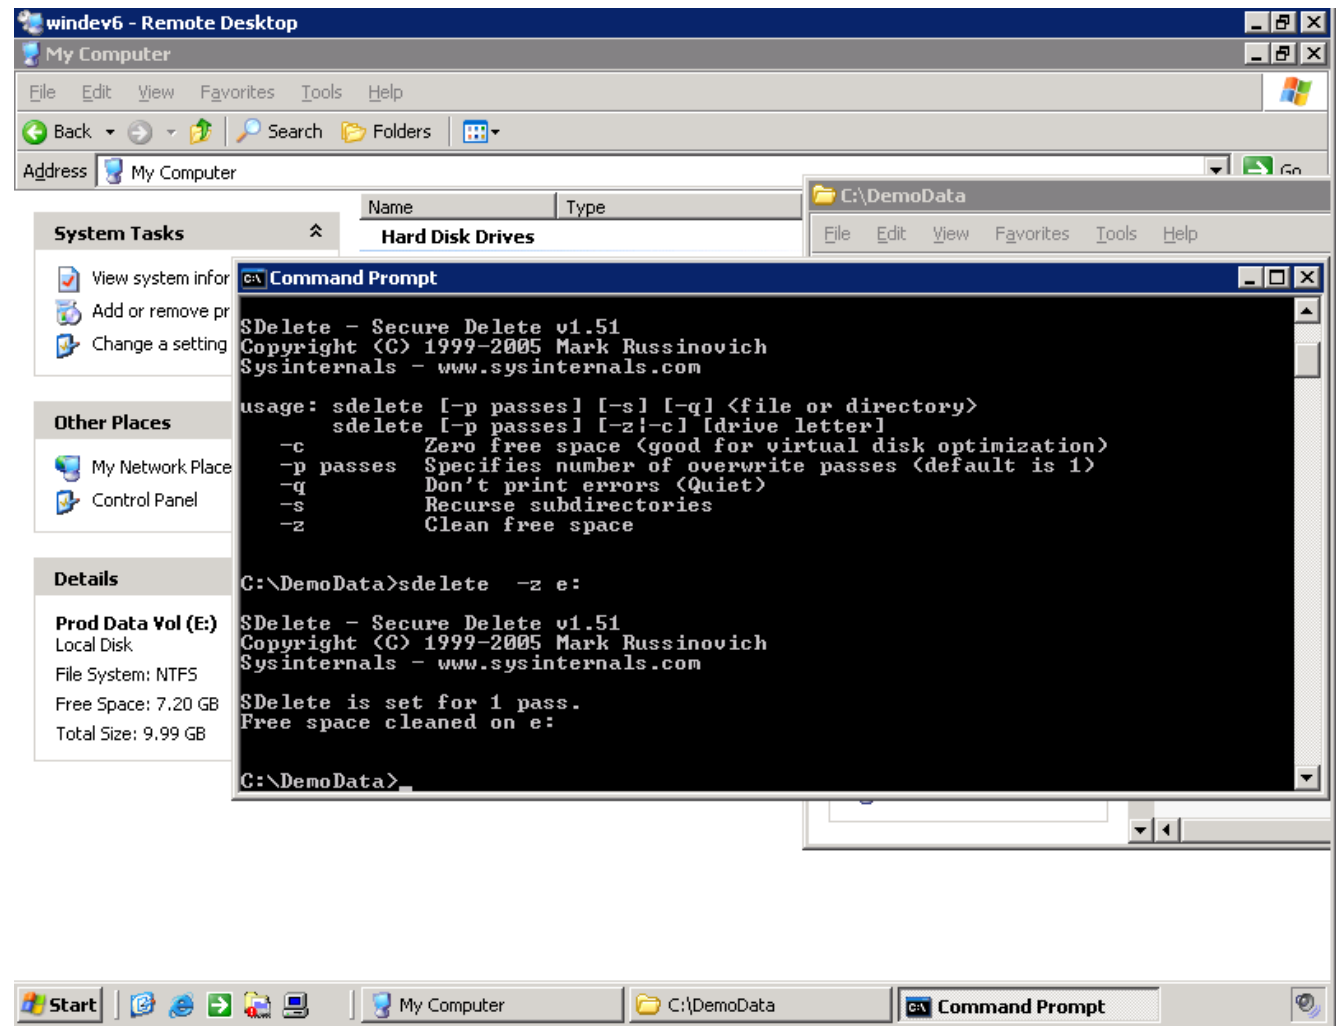
Background: Zero Page Reclaim Provides Limited Ability for Ongoing Storage Reclamation

To date only zero page reclaim and manual tools have been available to make reclamation useful on an ongoing basis

-SDelete

-vmkfstools

-HDS GSS
DPTools toolset
(best)




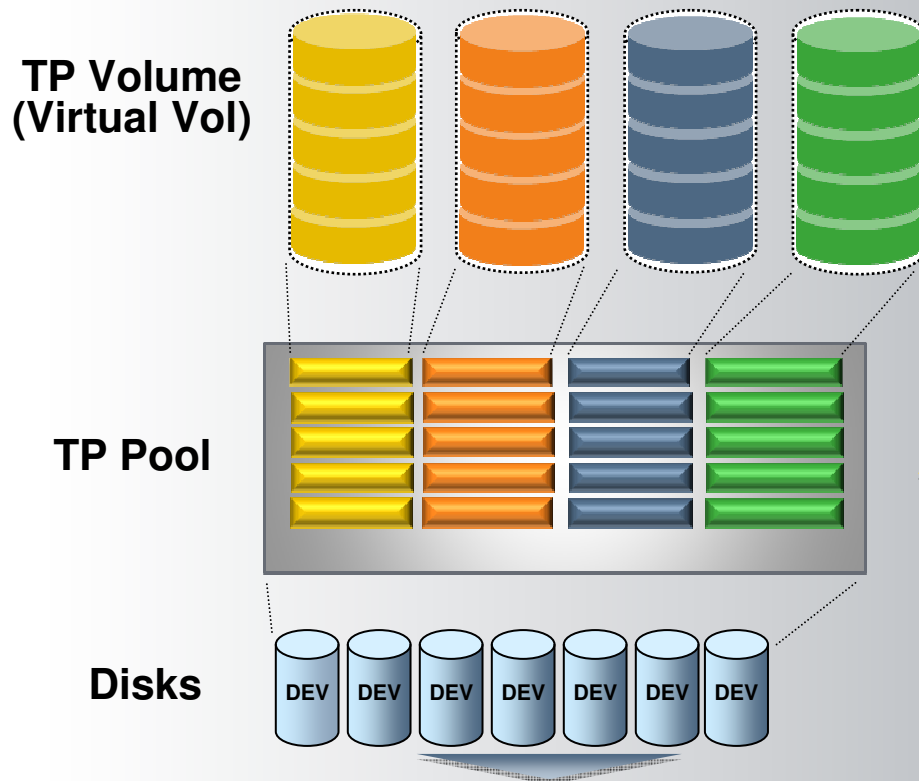
New **Write Same** command support for ongoing filesystem reclamation

- New in AMS2000 V8 and USP V/VM V6 microcode, Dynamic Provisioning can now be used in conjunction with Symantec Storage Foundation's VxFS filesystem to support ongoing filesystem storage reclamation to increase storage cost savings.
- This is a significant industry first “evolving standard interface” technology advancement. And Symantec on the server side and HDS, 3Par and IBM on the storage side are the first to support this capability.
- “Write Same” is a repurposed SCSI command that now provides a mechanism to allow the server-based filesystem such as Symantec VxFS to notify the array which pages on a thinly provisioned LUN or volume do not have meaningful data. The array can then safely disassociate the page from a thinly provisioned LUN/volume and make it available for re-use.
- The primary “Use Case” for this is to operationally keep filesystems thin. Through the use of Veritas Storage Foundation's VxFS filesystem and Hitachi Dynamic Provisioning the server and storage array now work together to continually automatically and transparently optimize storage utilization.

The Case for Thin Reclamation



Runaway Application in Action



- VxFS/SFW minimize unnecessary allocations
- BUT unnecessary allocations are caused by applications
- Think of the multiple temp files Word creates on Windows
- Think of an application running-away on a Thin Volume...

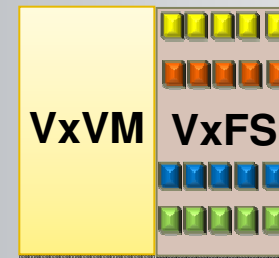

Storage Array

Host

The Case for Thin Reclamation



Runaway Application in Action



System Admin
Recovers and
deletes useless
data

- The file system fills up, triggering allocations in the array's shared TP Pool
- Deleting the data in the file system does not affect utilization in the TP Pool

Storage Array

Host

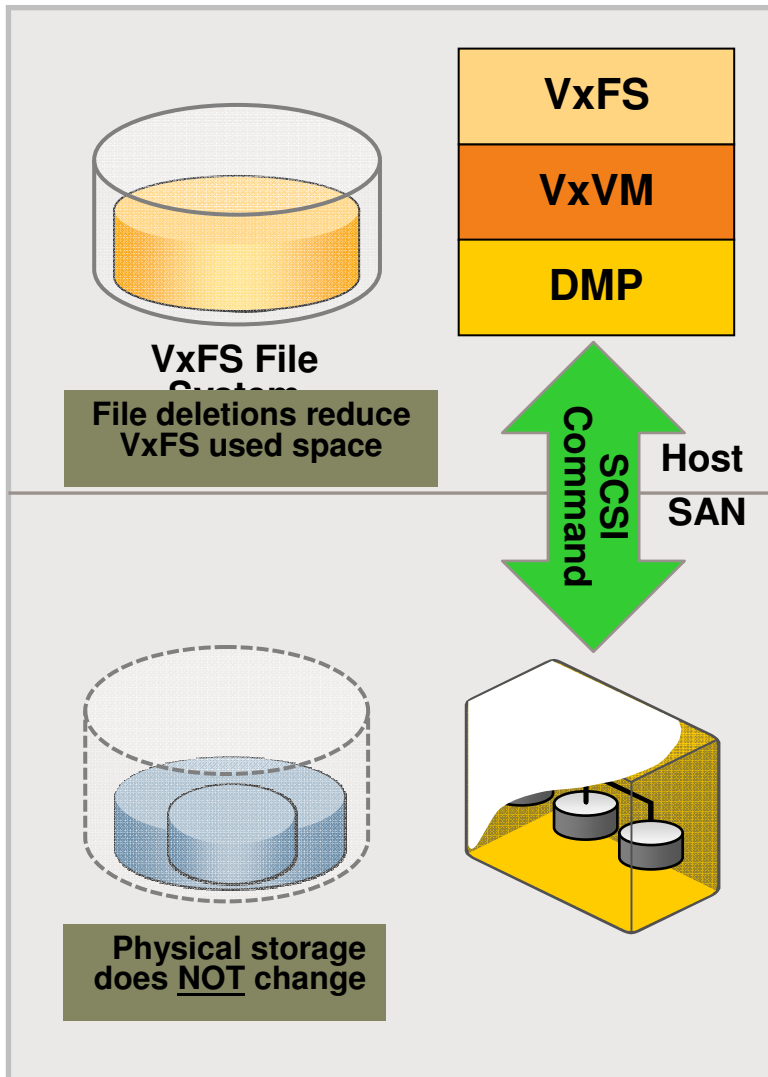
Host Stack Matters in Reclamation



- Array does not know what is a wasted & reclaimable space
 - Only the host FS can differentiate between free and used blocks
- VxFS can differentiate between free & used blocks
- Solution!! – Intelligence in the host, execution in the array
 - Leverage host intelligence to “right-size” physically allocated space

SF Thin Reclamation API

WRITE_SAME(16)



- Online & In-band Thin Reclamation API
 - WRITE_SAME(16) SCSI command based
 - Online reclamation - No impact to App I/O
 - File System, host volume or LUN granularity
 - Aligned with array chunks & Offset
 - Correctness & Simplicity is the key
 - Error Handling for Write_Same failures
- VxFS or VxVM based CLI

```
#>vxdisk reclaim 3par0_1,3par0_2
Reclaiming thin storage on:
Disk 3par0_1 : Done!
Disk 3par0_2 : Done!

#> fsadm -R
```

SF Thin Reclamation API Discovery & Reporting



- Thin Provisioning Device Discovery
 - Each array vendor has different values of TP “attributes”
 - ASL (DMP) discovers TP capabilities of any array
 - LUNS can be *thin* or *thinrclm*
 - Enables SmartMove & Allows SCSI Reclamation command
- Thin Reclamation Array Attribute Discovery
 - ASLs (DMP) discovers “Chunk Size” and other array specific TP attributes
 - Optimized handling of different storage vendor TP arrays
- Thin Reclamation Reporting
 - Reported via “vxdisk -o thin list” or via “vxdisk -e list”

DANAME	DISK SIZE	<i># vxdisk -e thin list</i> PHYS_ALLOC	DISK GROUP	ATTR
3par0_0	2000GB	50GB	dg1	thinrclm
3par0_1	200GB	50GB	dg2	thin

Thin Provisioning Challenges Solved



- Challenges

Efficient writes

Migrating to “Thin”

Staying “Thin”



- Symantec solution

VxFS extremely “thin friendly”

Storage Foundation SmartMove

Veritas Thin Reclamation API

Array Support for Veritas Thin Reclamation API

Today

End 2009

2010

- All arrays supported with SmartMove
- SF implements Thin Reclamation API
- Array vendors are working on it

HITACHI
Inspire the Next

EMC²
CLARIION

3PAR
Serving Information

XIV
Storage Reinvented

hp

FUJITSU

NetApp

EMC²
SYMMETRIX

IBM
DS8K

LSI
Sun
microsystems

Automated Storage Optimization With Thin Provisioning & Reclamation



- Thin Provisioning provides on-demand allocation
- Thin Reclamation automates returning storage to the pool
- Both operations are transparent to applications, host and SAN
 - Host file systems do not change
 - Lun configurations do not change
 - SAN configuration does not change
- Thin LUNs act as ‘access pipes’ to the array storage pools
- TP and TR together deliver automated storage management
- Enabling storage management at the pool level

New Write Same command support

What is the SCSI Write Same Command?

WRITE SAME(16) command as used to de-allocate and allow reclamation of pool capacity from allocated storage

Generally the WRITE SAME(16) command is a request that transfers a single logical block from the data-out buffer and writes the contents of that logical block, with modification based on the LBA DATA bit and the PBDATA bit, to the specified range of LBAs. Each logical block includes user data and may include protection information, based on the WRPROTECT field and the medium format.

This command fills up the specified segment on CDB from LBA to Number of Blocks with writing data of ONE block on Data-out.

CDB Format

Bit	7	6	5	4	3	2	1	0
0	Operation Code (0x93)							
1	WRPROTECT			Reserved	UNMAP	PBDATA	LBA DATA	Obsolete
2-9	LBA							
10-13	Number of Logical Blocks							
14	Reserved			GROUP NUMBER				
15	Control							
	Vendor Unique			Reserved		NACA	Flag	Link

Data-out Format

Bit	7	6	5	4	3	2	1	0
0-511	Writing Data							

New Write Same command support

Symantec Storage Foundations version support for HDS Dynamic Provisioning Write Same

Array Support	SF 5.0MP3, SF 5.0.1 Support	SF 5.1 Support	SF 5.1 MP1 Support	SFW 5.1 SP1 Support
HDS USP-V/VM	Solaris, SolarisX64, AIX, HPUX (SF5.0.1), Linux	Solaris, Linux RHEL5, Rest based on demand	Windows, Jan 2010	All, Jan 2010
HDS AMS2K	Solaris & Linux, Rest based on PAR request	Based on demand	Windows Jan 2010	All, Jan 2010

New Write Same command support

- The minimum Hitachi Adaptable Modular Storage 2000 systems microcode level for Write Same support is 0883/B and higher, initially released in November 2009.
 - Also referred to as “V8”
 - Write Same storage reclamation is fully automatic on the USP V/VM platforms
- The minimum USP V / VM microcode level is 60-06-05-00/00 and higher
 - Also referred to as “V6”
 - Write Same storage reclamation on the AMS2000 requires some script or administrator driven actions

Write Same Futures and Standards

- The Write Same interface can and will be used by other applications, filesystems and storage systems from an increasing number of vendors, increasing available support combinations
- But it is not yet a standard, and is still evolving. This means the relationships between vendors supporting the technology is key. HDS and Symantec have such a relationship
- Going forward the standards body responsible for the SCSI protocol - INCITS – has a T10 subcommittee that is including similar ideas into proposals for the SCSI standard which are likely to be adopted by many file systems, databases, and storage vendors. SNIA is also looking at it. Hitachi is monitoring this activity and will adopt as appropriate as it becomes a recognized industry approach.

Hitachi Storage Solutions

HDS Write Same Technical Slides

New Dynamic Provisioning Support in USP V and VM V6 microcode

- The following new Hitachi Dynamic Provisioning enhancements were delivered in the USP V and VM V6 microcode release:
 - **Write Same command support. This is handled automatically, no administrator intervention required.**
 - *Note Write Same supported on replicated volumes but only pvol space is immediately freed, svol zero page reclaim must be manually run to reclaim svol space.*
 - Improved pool monitoring through a new provisioning ratio threshold setting, which is percentage of the total size of all DP-VOLs with respect to the DP-Pool size (implemented only in Device Manager, both GUI and CLI)
 - Improved performance optimization - HDP Pool rebalance Phase II adds rebalance of pre V05 pools
 - HDP zero page reclamation improvement for copy PPs (TC/HUR defined volumes)

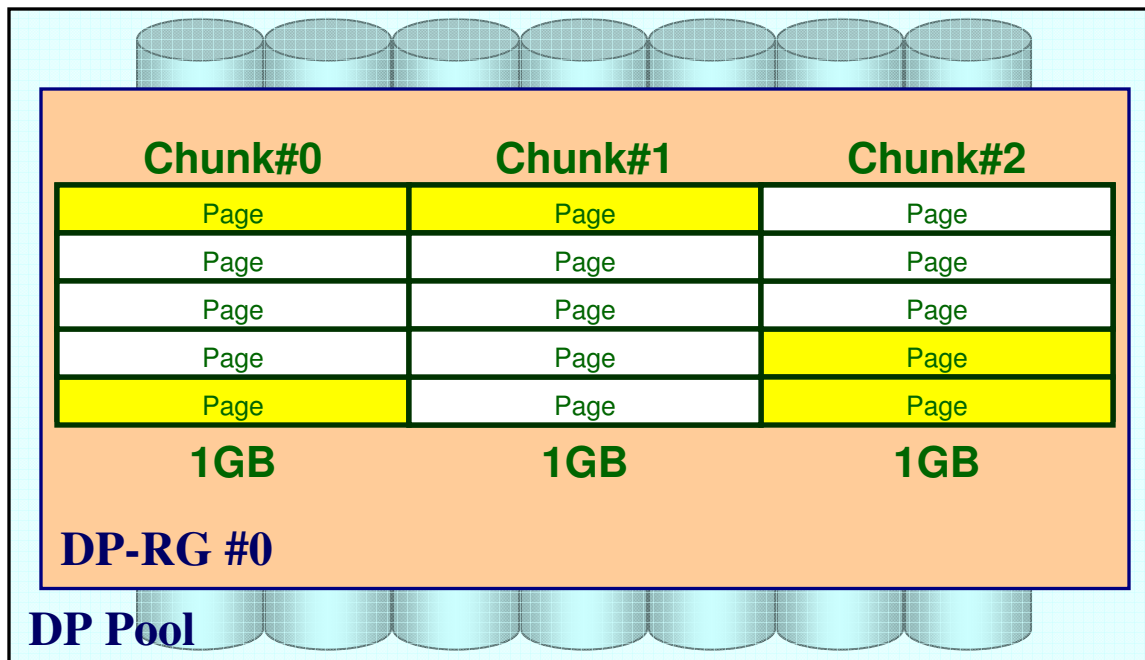
New in HDP on the AMS2000 with V8 microcode

#	New Item	Overview
1	HDP Pool rebalance	When RGs are added to a pool, the pool can be rebalanced
2	Reclaim Pages	Zero data in DP-Vol is checked for and deleted.
3	Write Same command support	This is a command from server to free up storage space for reclamation
4	Volume Migration Support	To migrate DP-Vol in the pool or between pools.
5	Data Retention Utility Volume support	LDEV guard function to DP-Vol
6	Pool Performance monitoring	Performance information for the unit of the pool is gathered.
7	HDP Stats to HiTrack	HDP information is offered from Hi-Truck.
8	Pool full write protect impacts all writes to the DP-Vol	When the state of the pool is full, all WR command is protected to DP-Vol.
9	SSD Support	SSD can be used by HDP volume.
10	General LU formatting with 0 data	LU is formatted by all 0 data as default.

HDP AMS Pool rebalance (DP Pool Optimization) with 0 page reclaim option

Chunks complicate reclaiming storage

On the AMS, first the system automatically does a zero page reclaim freeing pages in allocated chunks. To complete the reclamation you need to also need to consolidate the pages in order to free completely free and return the chunk to the pool for general reuse



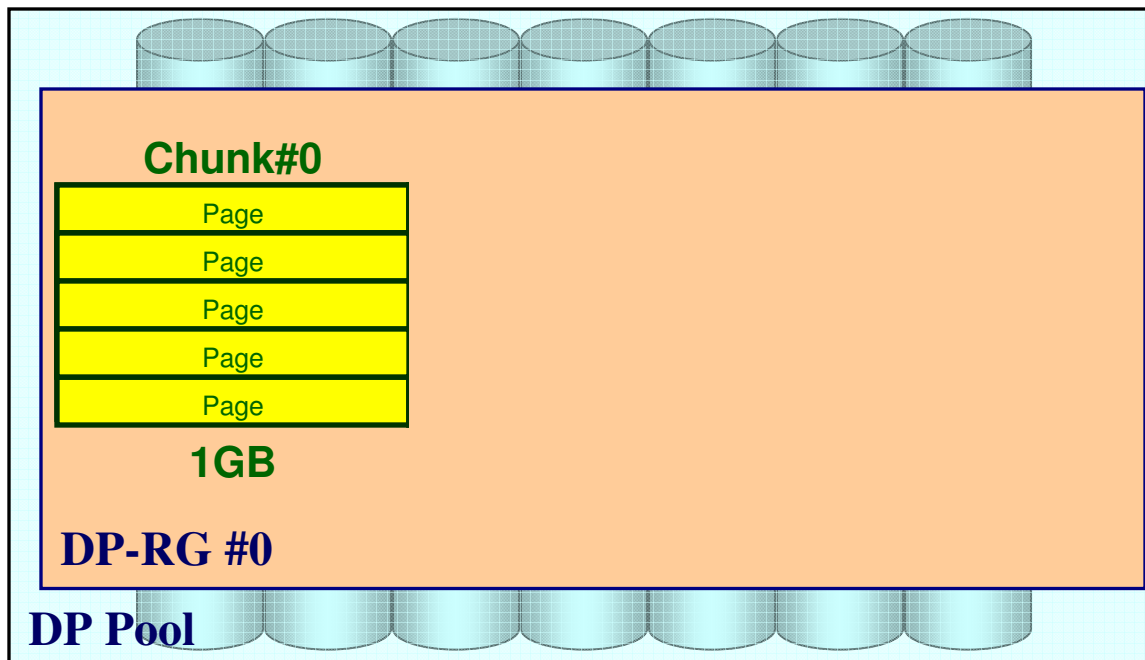
DP Pool
consumed capacity:
3GB

**Not fully reduced
Yet !!**

HDP AMS Pool rebalance (DP Pool Optimization) with 0 page reclaiming option

Step 2: Optimization Step defragments at the chunk level

Automatic Page de-fragmentation function !!



DP Pool
consumed capacity:

1GB

Reduced !!

HDP AMS Dynamic Provisioning Optimization

Use cases

#	Use cases			
	Case	Purpose		
		zero page reclaiming (option)	Optimization	
			Performance	Capacity
1	After expanding the DP Pool (RG addition to the DP Pool)	--	X	X
2	After shrinking the Virtual-LU(DPVOL) capacity	--	--	X
3	Reclaiming Zero Pages after migrating thick to thin or after restoring Host Backup data	X	--	X
4	After receiving the Write Same command	--	--	X

Note1: Automatic or manual execution after expanding or contracting a pool is configurable. On the AMS there is no automatic execution for Optimization after Zero Page Reclaim or after receiving write same, so user needs to create specific script by using the CLI or API if they want to fully automate.

HDP AMS Pool rebalance (DP Pool Optimization) with 0 page reclaiming option

Monitoring space to be reclaimed

The screenshot shows the Hitachi Storage Navigator Modular 2 web interface. The title bar reads "Hitachi Storage Navigator Modular 2". The top navigation bar includes "File Go Help" and "Logged in as: John Smith" with "Close" and "Logout" buttons. The interface is divided into three main sections: Explorer, Arrays, and Performance.

Explorer: Shows a tree view with "Resources" expanded to "Arrays".

Arrays: Shows "Array 1" expanded to "Performance". The "Performance" sub-menu is expanded, and "DP Optimization" is highlighted with a red box.

Performance: Shows a table with the following data:

Name	Description
Monitoring	View and output the monitored performance in the array
Tuning Parameter	Configure the parameter to performance in the array
DP Pool Trend	View and output the monitored DP pool trend in the array
DP Optimization	View and configure the DP Optimization

Monitoring space to be reclaimed

The screenshot shows the Hitachi Storage Navigator Modular 2 web interface. The browser address bar indicates the URL is `http://windev4:23015/HiCommand/frame/Console`. The page title is "Hitachi Storage Navigator Modular 2". The user is logged in as "null".

The left navigation pane shows the "Arrays" section expanded, with "DP Optimization" selected under the "Performance" category.

The main content area displays the "DP Optimization" configuration for array "DF800M_85010089". The "Summary" section shows the priority is "DP Optimization". Below this is a table titled "Logical Units in the DP Pools".

Logical Units in the DP Pools							
Rows/Page: 25 Page 1 of 1							
<input type="checkbox"/>	LUN Δ	DP Pool	RAID Level	Capacity		Reclaimed	Status
				Total	Consumed		
<input type="checkbox"/>	0000	000	RAID5(2D+1P)	50.0GB	18.0GB	0.0MB	Normal
<input type="checkbox"/>	0002	000	RAID5(2D+1P)	16.0GB	0.0MB	0.0MB	Normal
<input type="checkbox"/>	0008	000	RAID5(2D+1P)	50.0GB	0.0MB	0.0MB	Normal
<input checked="" type="checkbox"/>	0010	000	RAID5(2D+1P)	10.0GB	10.0GB	0.0MB	Normal
<input type="checkbox"/>	0024	000	RAID5(2D+1P)	15.0GB	15.0GB	1.0GB	Normal

At the bottom of the table, there are buttons for "Optimize DP", "Cancel Optimization", "Filter", and "Filter Off".

Optimize options

The screenshot displays the Hitachi Storage Navigator Modular 2 web interface. The main window shows the 'DP Optimization' page for array DF800M_85010089. A table lists logical units in DP pools, with LUN 0010 selected. An 'Optimize DP - 0010' dialog box is open, showing the following options:

- Optimize all logical units in DP pools including selected logical units: Yes
- Reclaim zero pages before optimizing DP: Yes

Buttons for 'OK', 'Cancel', 'Optimize DP', 'Cancel Optimization', 'Filter', and 'Filter Off' are visible at the bottom of the dialog.


LUN	DP Pool	RAID Level	Capacity			Status
			Total	Consumed	Reclaimed	
0000	000	RAID5(2D+1P)	50.0GB	18.0GB	0.0MB	Normal
0002	000	RAID5(2D+1P)	16.0GB	0.0MB	0.0MB	Normal
0008	000	RAID5(2D+1P)	50.0GB	0.0MB	0.0MB	Normal
0010	000	RAID5(2D+1P)	10.0GB	10.0GB	0.0MB	Normal
0011	000	RAID5(2D+1P)	15.0GB	15.0GB	1.0GB	Normal

HDP AMS Pool rebalance Case #1

DP Pool Optimization with 0 page reclaim

ShadowImage with DP Pool Optimization

Case1: PVOL Optimization, then execute ShadowImage backup




You do NOT need to execute the DP Pool Optimization for SVOL!!

HDP AMS Pool Rebalance Case #2

DP Pool Optimization with 0 page reclaiming option

DP Pool Optimization and ShadowImage

Case 2: If you execute a DP Pool Optimization on a paired PVOL you will not get a thinned SVOL. You still have to optimize the Svol to get the full benefits



You would then need to execute DP Pool Optimization for SVOL!!

Optimize Thin Storage with Veritas Storage Foundation



- Thin Provisioning will change how storage is managed
- Optimize Thin Storage on all OS with Storage Foundation
 - Get cross platform thin friendly host based storage management
 - Automatically reclaim space as you migrate online to thin storage
 - Be ready for online thin storage reclamation
- Standardize on Hitachi Dynamic Provisioning with SF 5.0 MP3rp2 and SFW 5.1sp1

Questions and Discussion

- Symantec / Hitachi Data Systems Thin Provisioning Whitepaper: <put name of paper>
- Available:
 - www.hds.com <need location>
 - www.symantec.com <need location> the site below does list the white paper or not that I could find.
- And on symantec.com at:
http://www.symantec.com/business/products/whitepapers.jsp?pcid=pcat_business_cont&pvid=203_1

- **Upcoming Webcasts**

- **VMware Support with Device Manager and Tuning Manager in Hitachi Storage Command Suite v6.3:** February 10, 9am Pacific
- **How to Manage Hitachi Dynamic Provisioning:** February 17, 9am Pacific
- **Storage Service Level Management with Hitachi Storage Command Portal:** February 24, 9am Pacific
- **Economic Impact of Storage Archive:** March 31, 9am Pacific

- **Please check www.hds.com/webtech for:**

- Link to the recording, the presentation and Q&A (available next week)
- Schedule and registration for upcoming WebTech sessions

Thank you!