


HITACHI
Inspire the Next

Hitachi Data Systems

Hitachi Dynamic Provisioning with HNAS v12.1

 Hitachi Data Systems

MK- 92HNAS057-01



© 2011-2015 Hitachi, Ltd. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written permission of Hitachi, Ltd.

Hitachi, Ltd., reserves the right to make changes to this document at any time without notice and assumes no responsibility for its use. This document contains the most current information available at the time of publication. When new or revised information becomes available, this entire document will be updated and distributed to all registered users.

Some of the features described in this document might not be currently available. Refer to the most recent product announcement for information about feature and product availability, or contact Hitachi Data Systems Corporation at <https://portal.hds.com>.

Notice: Hitachi, Ltd., products and services can be ordered only under the terms and conditions of the applicable Hitachi Data Systems Corporation agreements. The use of Hitachi, Ltd., products is governed by the terms of your agreements with Hitachi Data Systems Corporation.



Hitachi Data Systems products and services can be ordered only under the terms and conditions of Hitachi Data Systems' applicable agreements. The use of Hitachi Data Systems products is governed by the terms of your agreements with Hitachi Data Systems.

Hitachi is a registered trademark of Hitachi, Ltd., in the United States and other countries. Hitachi Data Systems is a registered trademark and service mark of Hitachi, Ltd., in the United States and other countries.

Archivas, Dynamic Provisioning, Essential NAS Platform, HiCommand, Hi-Track, ShadowImage, Tagmaserve, Tagmasoft, Tagmasolve, Tagmastore, TrueCopy, Universal Star Network, and Universal Storage Platform are registered trademarks of Hitachi Data Systems Corporation.

AIX, AS/400, DB2, Domino, DS8000, Enterprise Storage Server, ESCON, FICON, FlashCopy, IBM, Lotus, OS/390, RS6000, S/390, System z9, System z10, Tivoli, VM/ESA, z/OS, z9, zSeries, z/VM, z/VSE are registered trademarks and DS6000, MVS, and z10 are trademarks of International Business Machines Corporation.

All other trademarks, service marks, and company names in this document or website are properties of their respective owners.

Microsoft product screen shots are reprinted with permission from Microsoft Corporation.


This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>). Some parts of ADC use open source code from Network Appliance, Inc. and Traakan, Inc.

Part of the software embedded in this product is gSOAP software. Portions created by gSOAP are copyright 2001-2009 Robert A. Van Engelen, Genivia Inc. All rights reserved. The software in this product was in part provided by Genivia Inc. and any express or implied warranties, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose are disclaimed. In no event shall the author be liable for any direct, indirect, incidental, special, exemplary, or consequential damages (including, but not limited to, procurement of substitute goods or services; loss of use, data, or profits; or business interruption) however caused and on any theory of liability, whether in contract, strict liability, or tort (including negligence or otherwise) arising in any way out of the use of this software, even if advised of the possibility of such damage.

The product described in this guide may be protected by one or more U.S. patents, foreign patents, or pending applications.

Notice of Export Controls

Export of technical data contained in this document may require an export license from the United States government and/or the government of Japan. Contact the Hitachi Data Systems Legal Department for any export compliance questions.



Document Revision Level

<i>Revision</i>	<i>Date</i>	<i>Description</i>
Version 1	12/4/14	Document created by GSS FCS Services Engineering
Version 2	1/7/14	Typo fix

Contact

Hitachi Data Systems
2845 Lafayette Street
Santa Clara, California 95050-2627
<https://portal.hds.com>
North America: 1-800-446-0744



Table of Contents

Overview	1
Contacting Hitachi Data Systems Support.....	1
Frequently Asked Questions	1
HNAS Features Requiring Use of Hitachi Dynamic Provisioning.....	1
Benefits of HDP Over Standard Parity Groups	1
Mixing SDs from a Thick Provisioned HDP Pool to Form a Thinly Provisioned Pool.....	2
Creating HDP Pools from Hard Drive Sets	2
Minimum Number of SDs Required for an HDP Pool	2
Striping Considerations When Expanding the HDP Pool.....	3
Sizing Considerations for the HDP Pool	3
When Upgrading an Earlier Release HNAS Pool to v12.1.....	3
Maximum File System Size.....	4
HDP Pool Free Space	4
Returning Space in the HDP Pool.....	6
For Further Assistance.....	6

Overview

Contacting Hitachi Data Systems Support

The Hitachi Data Systems customer support staff is available 24 hours a day, seven days a week. If you need technical support, log on to the Hitachi Data Systems Portal for contact information: <https://portal.hds.com>.

Frequently Asked Questions

HNAS Features Requiring Use of Hitachi Dynamic Provisioning

Question: Which HNAS features in the future will require use of Hitachi Dynamic Provisioning (HDP)? For example, will a 1PiB file system will require HDP.

Answer: There are currently no features that will require the use of HDP. However, to scale out the HNAS 4100 with upwards of 32 x 1 PiB file systems will require 512 x 64 TiB SDs.

Creating a standard RAID 6 (14+2) Parity Group with currently available 4 TB drives does not support that capacity. Therefore a 64 TiB SD can only be created from a Thin Provisioned DP Vol from an HDP pool.

Benefits of HDP Over Standard Parity Groups

Question: What are the benefits of HDP over standard Parity Groups?

Answer: In terms of HNAS, there are several advantages to using thinly provisioned DP Vol's from an HDP pool:

1. If the real storage capacity of the thinly provisioned DP Vol is exhausted (SD is not full), you simply add more DP RAID Groups/Pool Volumes to the HDP pool and the array rebalances the data without HNAS knowledge. This is obviously the recommended choice.
2. If the thinly provisioned capacity of the DP Vol is fully exhausted (each SD is full from the perspective of the array), you can then add more DP RAID groups/LDEVs to the HDP pool and the array once again rebalances. However, this time you must add more DP Vols as a new stripe set. Again, there is no need to perform careful expansions using Dynamic Read Balancing (DRB). Effectively DRB is rendered obsolete; in fact, it cannot run by design.

Even adding a single Pool Volume (DP RAID Group) will allow the span to go faster as there are simply more disks to access. That is not the case with standard Parity Groups, when you expand a span.

You do not need to expand a file system to benefit from the newly added HDDs. These alone help to make HNAS storage administration simpler and easier.

From a storage perspective, you are not required to initially purchase all real capacity, thus simplifying implementation and set-up.

Additionally, workloads are distributed among the DP RAID groups by the array, helping to eliminate hot-spotting (albeit, from an HNAS perspective, this is not guaranteed because of data layout and access patterns). For additional information, please see the CLI's `span-expand` man page.

Mixing SDs from a Thick Provisioned HDP Pool to Form a Thinly Provisioned Pool

Question: Can you mix SDs from a thick provisioned HDP Pool (or SDs from Parity Groups) with SDs from a thinly provisioned HDP Pool?

Answer: Taking advantage of the new features in v12.1 for thinly provisioning an HDP pool is independent of clusters, meaning a cluster may contain spans that use HDP and spans that do not.

Using HDP is a span option, not a per-cluster option.

In releases prior to v12.1 it was possible for the following scenarios to occur:

- Mix Logical Units within a stripe set that resided on an HDP pool(s) and standard Parity Groups
- Mix Logical Units from different HDP pool(s) within the same stripe set
- Have separate stripe sets that comprised Logical Units from HDP Pools and standard Parity Groups (this is a use case for Tiered File Systems (TFS))

However, in v12.1 it is no longer possible to create such configurations. In fact, upon upgrade to v12.1, you will not be permitted to create\expand\delete file systems that reside on such a storage pool. However, in v12.3 you will be able to issue a CLI command, `span-hdp-thickly-provisioned`, thus overriding the storage pool rules and allowing file system creation. If a system has used this command ensure that no HDP Pool which contributes DP Vols to the HNAS storage pool ever runs out of real storage capacity. Otherwise, the HNAS storage pool will fail until space is freed on the affected HNAS storage pool (deleting file systems, vacating chunks, and so on) or additional Parity Groups/Pool Volumes are added to the HDP pool.

Creating HDP Pools from Hard Drive Sets

Question: I have sixty 600GB 15K RPM and sixteen 4TB NLSAS disks. Should I create separate HDP Pools here or can I create single HDP Pool using all the disks (both 600G and 4TB disks)?

Answer: This question is more of a Tiered File System (TFS) versus HDT pool consideration.

To create a TFS, you would create two separate HDP pools (one HDP pool for each disk type), placing the 600GB 15K disks in Tier 0 and the 4TB NLSAS disks in Tier 1.


Note: Never attempt to mix different types of disk media in the same span unless it is tiered.

However, if you have created an HDT pool comprising both disk types, a TFS could not be created, because all SDs from the HDT pool must reside in the same tier. You would simply create a regular span/file system(s).

In v12.3, HNAS will support pinning of HDP pages to higher tier on the HDT pool, thus allowing for Tiered File System creation.

Minimum Number of SDs Required for an HDP Pool

Question: What is the minimum SDs required for an HDP Pool? And what is the best practice, either two or four SDs?



Answer: In v12.1, there is no hard requirement for the number of SD's within a stripe set. However, HDS recommends a modest use of HDP thin provisioning (upwards of two to three times). The current guideline is to create eight thinly provisioned DP Vols. More details will be available in the updated HDP guide.

In v12.3, the server will lockdown the initial minimum number of SD's residing on an HDP pool to four. Every subsequent stripe set must contain at least four SD's and at least as many as the first stripe set.

For example, if the original stripe set contained eight SDs, every subsequent stripe set must contain eight SDs. If your system is large—for example, sixteen SDs—you should put all sixteen SDs into a single stripe set. Do not break the original amount of SDs into smaller stripe sets.

There are different considerations for standard parity groups. Again, more details will be available in the updated HDP guide.

Striping Considerations When Expanding the HDP Pool

Question: When you expand the HDP Pool, do you need to maintain the same striping across SDs?

Answer: See the previous sections regarding lock-down considerations.

Sizing Considerations for the HDP Pool

Question: Is there a sizing guide for the HDP pool, considering future capacity needs?

Answer: Technically, there is not, because there are too many variables to consider. While it would be extremely beneficial to know the customer's future capacity needs from the onset, it is doubtful this is a reasonable expectation.

As noted previously, one of the great advantages of using a thinly provisioned HDP pool is the ability to add a small number of additional disks to increase performance. This is very beneficial for those customers with limited budgets and invariably the HNAS storage pool performance will increase after array rebalance occurs.


Eight thinly provisioned DP Vols of 8TiB will create an initial storage pool capacity of 64TiB. This would allow for sixteen storage pool expansions in a single 1 PiB storage pool (if each expansion was kept at the minimum of eight). Obviously if the initial real capacity of the HDP pool is greater than 64TiB, the thinly provisioned DP Vol capacity will need to be larger.

Generally speaking, the thinly provisioned capacity of the DP Vols should not exceed two to three times the initial real capacity of the HDP pool. For example, if the real capacity of the HDP pool is 32TiB, do not create eight thinly provisioned volumes of 24TiB (initial storage pool capacity of 192TiB).

Always try to make the DP Vols (in total) as large as you can foresee the span becoming. Also try to follow the guideline where the maximum thinly provisioned capacity will be up to two to three times the beginning capacity. If the thinly provisioned capacity is underestimated, the net result is that the span will need to be expanded (as described previously).

When Upgrading an Earlier Release HNAS Pool to v12.1

Question: An HNAS storage pool was created in an earlier release, using thick provisioned DP Vols with Full Capacity mode enabled. What happens upon upgrade to v12.1?



Answer: In HNAS releases prior to v12.1, using Full Capacity mode was the default setting when thickly provisioning DP Vols on Modular Storage Arrays. However, due to the integration that has occurred in v12.1 (and higher), the server now knows the amount of free space in the HDP pool, regardless whether DP Vols are of thinly or thickly provisioned.

Upon upgrade to v12.1, the server will notice there is no free space left on the HDP pool. Therefore, after the upgrade, do the following:

1. In SNM2, go to **Groups > Volumes > Format Mode** and select **Edit Format Mode**.
2. Change the Format Priority Mode to “Host Access”. Lowering the priority will allow HNAS IO priority on the array, however zero initialization will take longer.
3. Select the desired DP Vols and disable Full Capacity Mode.
4. On the HNAS CLI, run the command:

```
span-unmap-vacated-chunks -exhaustive <span>
```

This process will vacate any chunks from the span, thus returning any free space to the HDP pool. If there is no space to return that means the span is already full and the only way to expand the underlying file systems would be to add more storage.

- It is at this point that the array will begin the process of zero-initializing the unused pages within the HDP pool. This process will take time depending on the amount of space that can be freed.
- To monitor the progress use the command `span-list -s` and watch the “HDP “free” space grow.

Note: If thickly provisioning HDP using v12.1 (not an upgrade), do not enable Full Capacity Mode on the DP Vols.

Maximum File System Size

Question: What is the maximum file system size that can be created initially using HDP in v12.1?

Answer: The maximum file system size that can be created initially on SDs from an HDP pool is 1TiB, regardless whether provisioning is thick or thin the DP Vols on the array.

If you attempt to create a file system in excess of 1TiB, the process will fail on both the CLI and in the SMU. However, initially creating a 1TiB file system is obviously not the final size. You can manually expand the file system as many times as necessary in order to reach the desired size. Confining the file system appropriately, you can let the file system grow on demand.

Limiting the size is due to pre-allocation of HDP pages when the file system is created; that is, the file system can only use “real” storage capacity at any given time. In v12.1, the internal HNAS span mechanism will write one block to each HDP page in the chunks it allocates. In turn, this process requires time and as a result the initial file system capacity is limited. Similar limits are applied to background auto-expansion (64GiB).

HDP Pool Free Space

Question: How much free space is in the HDP pool?

Answer: The following example illustrates how to determine the free space in an HDP pool. It contains two thinly provisioned HDP pools, and each HDP pool consists of 4 DP Vols of 8TB.

HNAS does not know the individual DP Vol's "Consumed Capacity" in the HDP pool; HNAS does know the cumulative free space available in the HDP pool (as reported by the array).

The highlighted information is as follows:

- The **YELLOW** value is the total amount of free space on all the HDP pools that comprise the HNAS storage pool.
- The **LIGHT BLUE** value is the total amount of free space on the individual HDP pool that comprise the single stripe set.
- The **GREEN** value is the total size of the unused chunks in the stripe set. This value can either be larger or smaller than the amount of free space on the HDP pool where the SDs are located.

Since the HNAS storage pool in these examples only contain a single stripe set, both the YELLOW and LIGHT BLUE values are the same.

```
span-list -fsv

Span instance name      OK? Free  Cap/GiB Chunks          Con  Snap  Cod fullness
Cluster  Engrav  Permanent ID      DWB  HDP free
-----  -----  -----
bar      Dual      Yes  96%   32768  3750 x   9382494749  90%          1%, 1%, 0%
This     Dual      925957BDDEF7D28D  Yes   37230


  On HDP pool 1 with 37230GiB free
  Set 0: 4 x 8192GiB = 32768GiB, of which 31746GiB is free, 0GiB is vacated
    SD 4 (rack '92210503', SD '0006')
    SD 5 (rack '92210503', SD '0007')
    SD 6 (rack '92210503', SD '0008')
    SD 7 (rack '92210503', SD '0009')

fs fs2      UnMnt, EVS  1, cap  1022, unconfined , perm
92595648DEC31EDD, dev 1026, 4KiB, WFS-2, DSBs: 128, DWB: Yes , Thin: No, AX: OK

foo      Dual      Yes  96%   32768  3750 x   9382494749  90%          1%, 1%, 0%
This     Dual      9259570877CFE856  Yes   31776

  On HDP pool 0 with 31776GiB free
  Set 0: 4 x 8192GiB = 32768GiB, of which 31746GiB is free, 0GiB is vacated
    SD 0 (rack '92210503', SD '0002')
    SD 1 (rack '92210503', SD '0003')
    SD 2 (rack '92210503', SD '0004')
    SD 3 (rack '92210503', SD '0005')

fs fs1      UnMnt, EVS  1, cap  1022, unconfined , perm
9259565F76C0C74E, dev 1025, 4KiB, WFS-2, DSBs: 128, DWB: Yes , Thin: No, AX: OK
```



Returning Space in the HDP Pool

Question: If my file system is 30TiB and it is 95% full, and a large file (3TiB) is deleted, can the pages in the HDP pool be reclaimed for use with other clients?

Answer: No, the pages in the pool cannot be reclaimed to use the space for other clients.

The space used by the file system is all real capacity in the HDP pool. Although the internal HNAS Span Manager infrastructure can support shrinking a file system, the file system infrastructure does not currently possess the capability to defragment (and thus shrink) a file system. Therefore it is not possible to return the used space to the HDP pool.

In this case, there are no HNAS chunks to un-map. The only way to return space to the HDP pool is to delete and recycle the file system, then un-map vacated chunks.

For Further Assistance

If you have any other questions, please contact Hitachi Data Systems Support.

Hitachi Data Systems

Corporate Headquarters

2845 Lafayette Street
Santa Clara, California 95050-2639
U.S.A.

www.hds.com

Regional Contact Information

Americas

+1 408 970 1000

info@hds.com

Europe, Middle East, and Africa

+44 (0) 1753 618000

info.emea@hds.com

Asia Pacific

+852 3189 7900

hds.marketing.apac@hds.com



MK-92HNAS057-01