Redbooks
ibm.com/redbooks

IBM DS8000
High-Performance Flash Enclosure

Jeff Cook

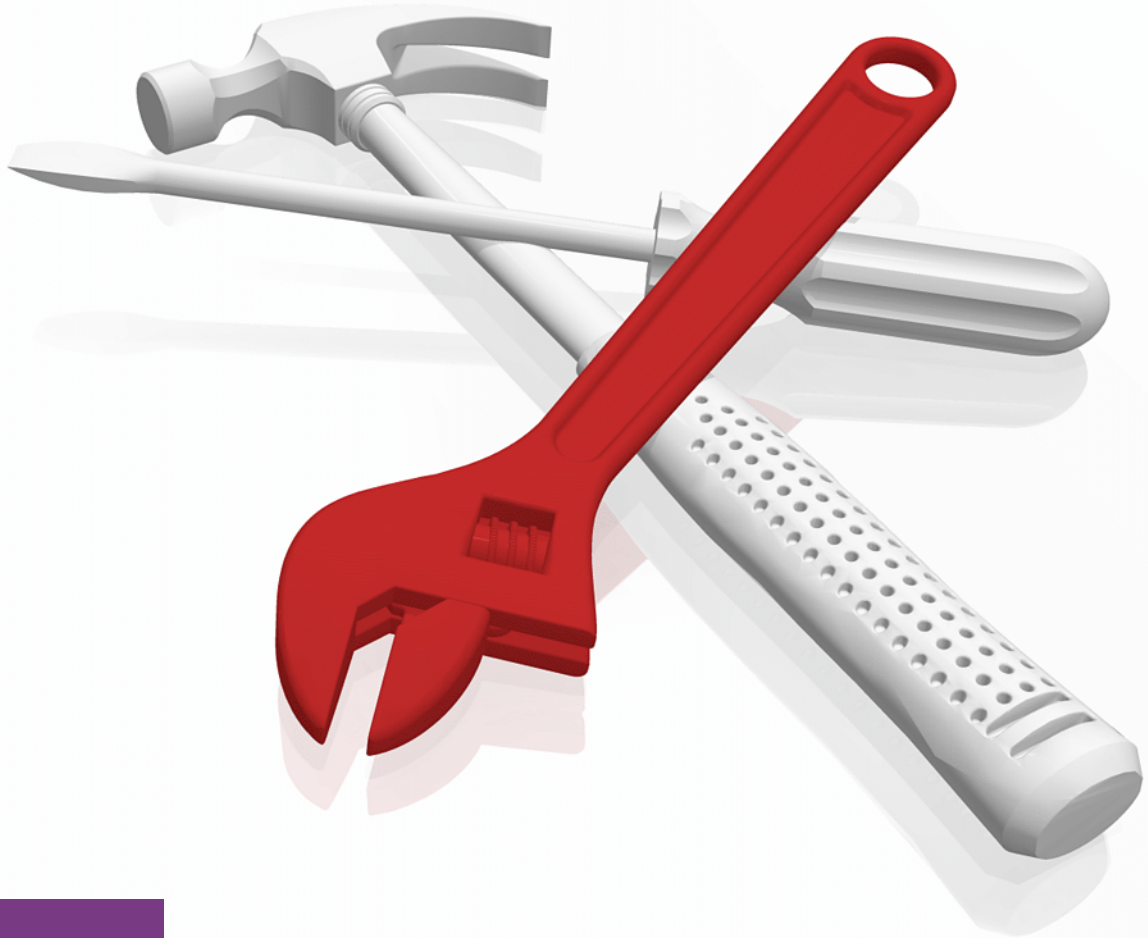

Storage

IBM

Product Guide

IBM DS8000: High-Performance Flash Enclosure

The high-performance flash enclosure (HPFE) is a Redundant Array of Independent Disks (RAID) storage enclosure that can support sixteen or thirty 400 GB encryption capable flash cards (1.8-inch, 46 mm form factor) in a 1U rack space. This IBM® Redbooks® Product Guide describes the IBM DS8000® high-performance flash enclosure.

HPFEs can be installed in the IBM DS8870 and IBM DS8880 storage systems:

- Up to four HPFEs can be installed in the DS8886 base frame, and up to four additional HPFEs in the first expansion frame.
- One or two HPFEs can be installed in the DS8884 base frame, and up to two additional HPFEs in the first expansion frame.
- One to eight HPFEs can be installed in the DS8870 high-performance all-flash configuration.
- Up to four HPFEs can be installed in the DS8870 enterprise class and business class configuration base frame, and up to four HPFEs in the first expansion frame.

Figure 1 shows the IBM high-performance flash enclosure.

Figure 1 IBM high-performance flash enclosure

Did you know?

Consider these facts:

- Compared to flash drives (also known as solid-state drives, or SSDs), the HPFE provides a higher standard of flash performance. As implemented in the DS8870 and DS8880, the HPFEs are directly attached to the PCIe fabric, enabling increased bandwidth compared to Fibre Channel attached standard drive enclosures.
- Flash cards are enterprise class storage devices that are targeted at I/O-intensive workload applications that can benefit from a high level of fast-access storage.
- High-performance flash cards, along with flash drives, are considered as Tier 0 storage. Both offer a number of potential benefits over spinning drives, including higher IOPS, lower power consumption, less heat generation, and lower acoustical noise. However, flash cards offer even higher throughput using the flash RAID adapters in the high-performance

flash enclosure and direct PCIe connectivity to the processor complexes, compared to the fiber-attached flash drives installed in the standard drive enclosures.

- The IBM Easy Tier® intra-tiering auto-rebalance (micro-tiering) feature is used to distribute the workload among traditional flash drives and flash cards according to their IOPS capacity within the storage tier

High-performance flash enclosure highlights

The high-performance flash enclosure includes various performance and other capabilities, as characterized in the following list:

- In DS8870, the high-performance flash enclosure delivers up to 250,000 IOPS and up to 3.4 GBps bandwidth in just one EIA (1U) of rack space. DS8880 doubles the number of PCIe lanes to each HPFE.
- Database performance can be accelerated by up to four times.
- Batch times can shrink by up to 10%, therefore reducing impact on production applications.
- Each High-Performance Flash Enclosures each contain up to 12 TB of raw capacity.
- Up to 8 HPFEs per DS8870 or DS8886, for a total of 96 TB raw capacity.
- Flash cards in the High-Performance Flash Enclosure support full drive encryption.

Architecture and key components

The high-performance flash enclosure is a 1U enclosure that integrates a fully redundant pair of flash RAID adapters, dual power supplies with integrated cooling fans and sixteen or thirty 1.8-inch encryption-capable flash cards.

The flash RAID adapters are directly connected over a PCIe fabric to the I/O enclosures, with a pair of Gen2 PCIe cables. In DS8870, each cable provides a four-lane 2 GBps full duplex connection. In DS8880, each cable provides an eight-lane, 4 GBps full duplex connection.

Figure 2 shows the major components of the HPFE, with the cover removed from the top, and from one of the Flash RAID adapters.

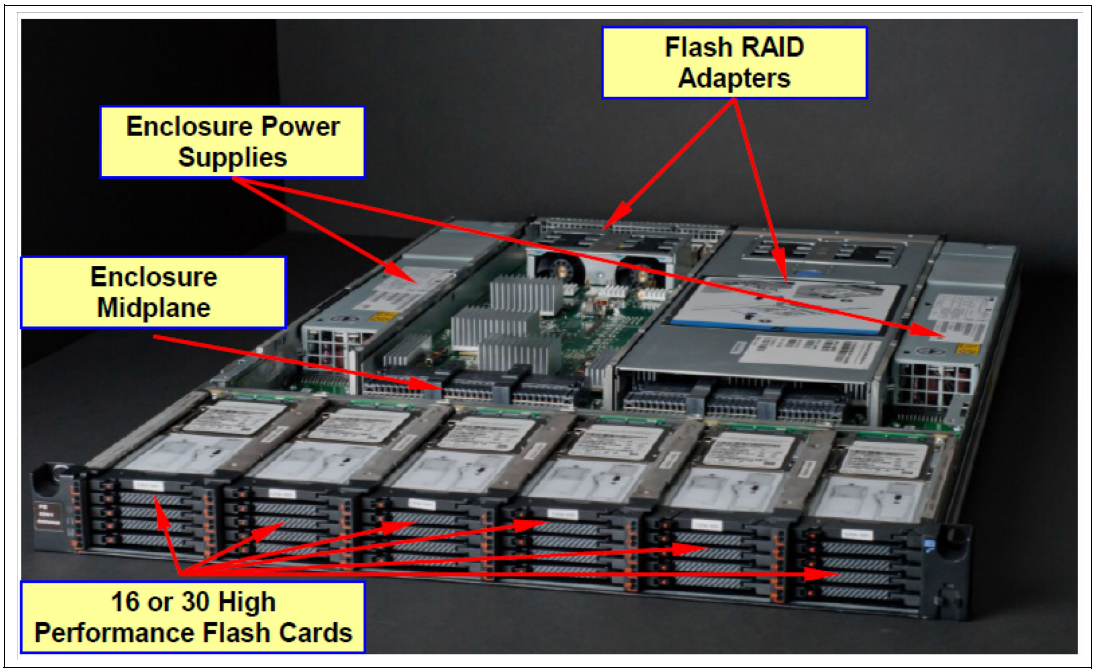

Figure 2 High-performance flash enclosure: key components

Flash RAID Adapter

Each HPFE contains two flash RAID adapters that provide redundant data paths to the flash cards in the enclosure. These adapters also provide enclosure services to control power, cooling, and other non-data functions. Each flash RAID adapter has a redundant pair of cooling fans.

- In DS8870, each flash RAID adapter connects to an I/O enclosure using a 2 GB PCIe Gen 2 four-lane cable.
- In DS8880, these cables are 4 GBps, eight-lane.

The PCIe connectivity has none of the protocol overhead associated with Fibre Channel architecture. Each I/O enclosure pair supports up to two HPFEs.

Enclosure power supplies

Each HPFE has a pair of fully redundant power supply units (PSU). Each power supply unit has its own integrated fans. Cooling fans from the PSUs and the flash RAID adapters provide cooling for the entire enclosure.

Enclosure midplane

The enclosure midplane provides the connectivity for the flash RAID adapters, PSUs, and flash cards.

High-performance flash cards

Each HPFE is installed with a minimum of sixteen 1.8 inch high-performance flash cards. An optional second set of 14 flash cards can be installed for a maximum total of 30 flash cards. If the second set is not installed, then a set of fillers are installed in the empty slots.

The flash cards are installed into the HPFE in rows. Figure 3 shows flash card locations, when HPFE is installed horizontally and vertically.

Note: Flash cards are not available as capacity on demand (CoD) features.

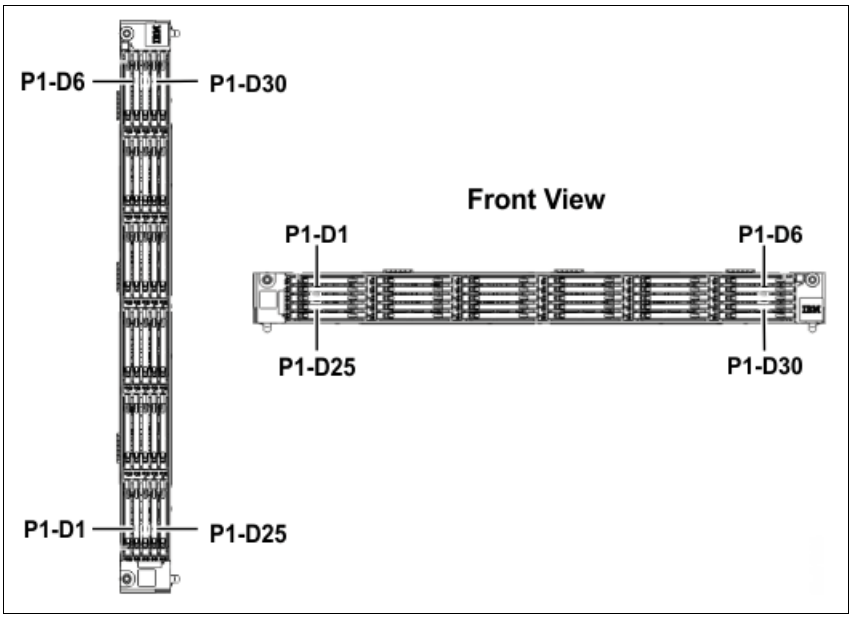

Figure 3 Flash card locations within the HPFE

Virtualization

The HPFE contains 16 or 30 flash cards. Two of the initial 16 flash cards are allocated as spares. HPFE supports only RAID 5 arrays. The arrays are configured as follows:

- First 16 flash cards installed
  - arraySite 1 = 6+P+S RAID 5
  - arraySite 2 = 6+P+S RAID 5
- Optional second set of 14 flash cards installed
  - arraySite 3 = 6+P RAID 5
  - arraySite 4 = 6+P RAID 5

Each HPFE can have a maximum of 12 TB of raw capacity, or 9.125 TB of usable capacity when configured with RAID 5. With the supported maximum of eight fully populated HPFEs in a DS8870 or DS8880 storage system, maximum usable RAID 5 capacity is 73 TB of high performance flash.

High-performance flash enclosures in the DS8880

The DS8880 is available in two models:

- DS8886
  - Supports up to four HPFEs in the base frame
  - Supports up to four HPFEs in the first expansion frame

  Figure 4 on page 5 shows HPFE locations and associated device adapter (DA) Pair numbering for the DS8886.
- DS8884
  - Supports up to two HPFEs in the base frame
  - Supports up to two HPFEs in the first expansion frame

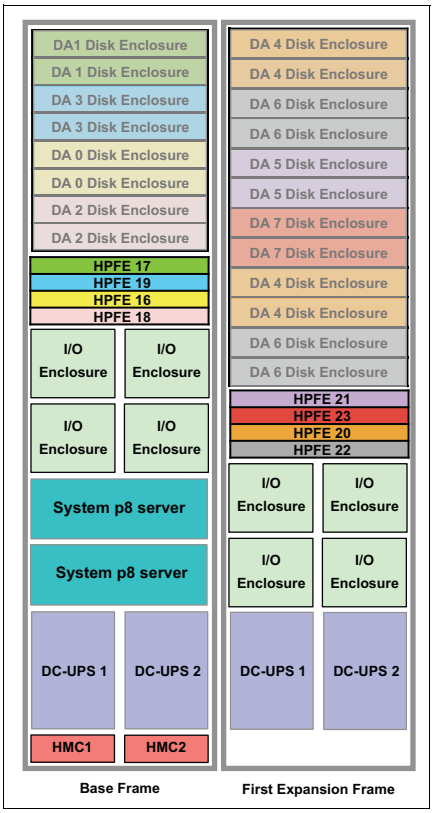

Figure 4 DS8886 with eight HPFEs installed

The DS8886 supports zero to eight HPFEs: four installed in the base frame and four installed in the first expansion frame. The installation order is as follows:

- R1-B04 - DA pair 18
- R1-B03 - DA pair 16
- R1-B02 - DA pair 19
- R1-B01 - DA pair 17
- R2-B04 - DA pair 22
- R2-B03 - DA pair 20
- R2-B02 - DA pair 23
- R2-B01 - DA pair 22

The DS8884 supports zero to four HPFEs: two installed in the base frame and two installed in the first expansion frame. The installation order is as follows:

- R1-B04 - DA pair 18
- R1-B03 - DA pair 19
- R2-B04 - DA pair 22
- R2-B03 - DA pair 23

Figure 5 shows HPFE locations and associated DA Pair numbering for the DS8884.

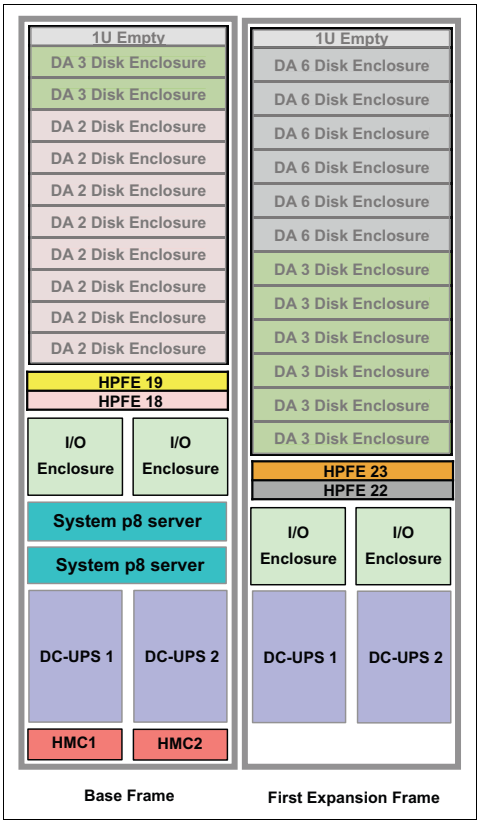

Figure 5 DS8884 with four HPFEs installed

Figure 6 provides a summary of HPFE and Flash cards supported in DS8880 configurations. For more information about DS8870 configurations and supported components, see *IBM DS8880 Architecture and Implementation,* SG24-8323.

Processor Cores	System Memory (GB)[1]	Max HPFEs Base Frame	Max HPFEs 1st Exp Frame	Max Flash Cards
DS8886 Configuration				
8-core	128	4	N/A	120
8-core	256	4	N/A	120
16-core	256	4	4	240
16-core	512	4	4	240
24-core	1024	4	4	240
24-core	2048	4	4	240
DS8884 Configuration				
6-core	64	2	N/A	60
6-core	128	2	2	120
6-core	256	2	2	120
Notes: 1 - System memory is a total of memory from both CPCs 2 - Number of HPFEs is independent of number of standard drive enclosures				

Figure 6 DS8880 HPFE supported configurations

Connectivity

The high-performance flash enclosures are directly attached to the I/O enclosure using PCIe cabling, which increases bandwidth and transaction-processing capability compared to standard drive enclosures connectivity (FC-AL).

Figure 7 is a block diagram that shows a simplified view of the PCIe cabling topology.

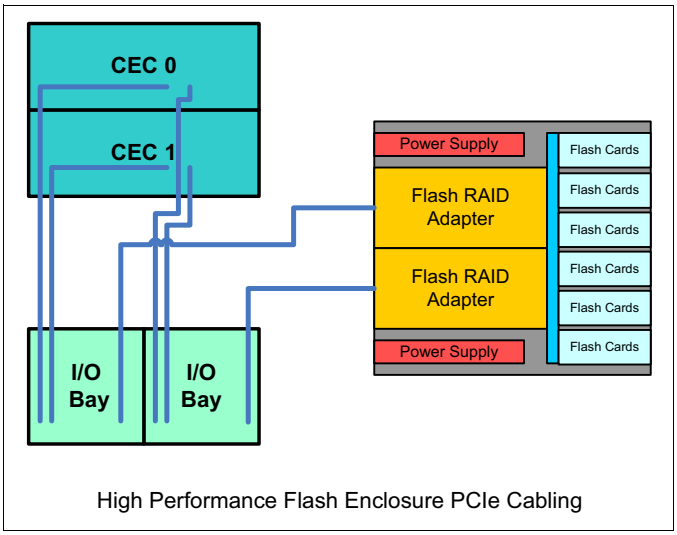

Figure 7 DS8880 HPFE PCIe cabling block diagram

Figure 8 shows the PCIe connections in the DS8880 I/O enclosure. Host adapters (HA) and device adapters (DAs) are also shown.

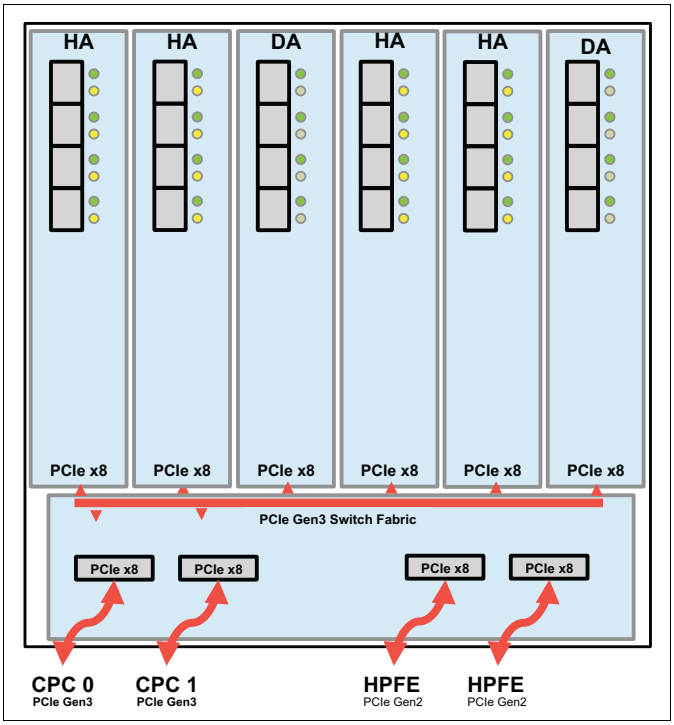

Figure 8 DS8880 HPFE PCIe cable connections

High-performance flash enclosures in the DS8870

The DS8870 has three configurations:

- The DS8870 high-performance all-flash configuration
  - Supports one to eight HPFEs in a single frame
- The DS8870 enterprise class configuration
  - Supports up to four HPFEs in the base frame
  - Supports up to four HPFEs in the first expansion frame
- The DS8870 business class configuration
  - Supports up to four HPFEs in the base frame
  - Supports up to four HPFEs in the first expansion frame

The DS8870 high-performance all-flash configuration supports a minimum of one HPFE installed in location R1-B04. The installation order is as follows:

- R1-B04 - DA pair 10
- R1-B03 - DA pair 8
- R1-B02 - DA pair 11
- R1-B01 - DA pair 9
- R1-B14 - DA pair 14
- R1-B13 - DA pair 12
- R1-B12 - DA pair 15
- R1-B11 - DA pair 13

Figure 9 shows HPFE locations and associated DA Pair numbering, for the high-performance all-flash configuration.

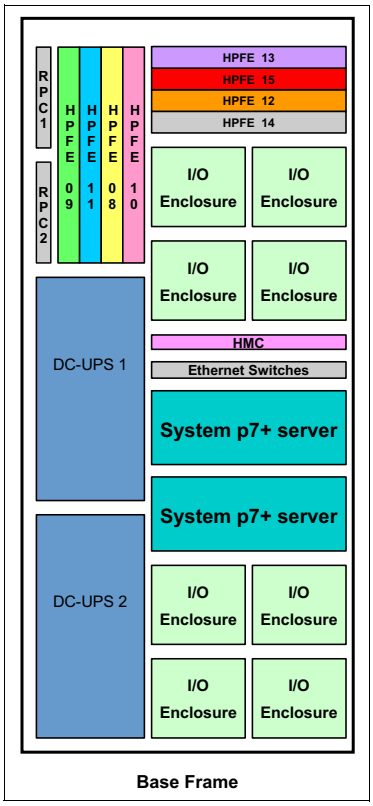

Figure 9 DS8870 high-performance all-flash with 8 HPFEs installed

The DS8870 enterprise class and the business class configurations support zero to eight HPFEs: four installed in the base frame and four installed in the first expansion frame. The installation order is as follows:

- R1-B04 - DA pair 10
- R1-B03 - DA pair 8
- R1-B02 - DA pair 11
- R1-B01 - DA pair 9
- R2-B04 - DA pair 14
- R2-B03 - DA pair 12
- R2-B02 - DA pair 15
- R2-B01 - DA pair 13

Figure 10 shows HPFE locations and associated DA Pair numbering for the enterprise class configuration. The HPFE installation order and locations are the same in the business class configuration.

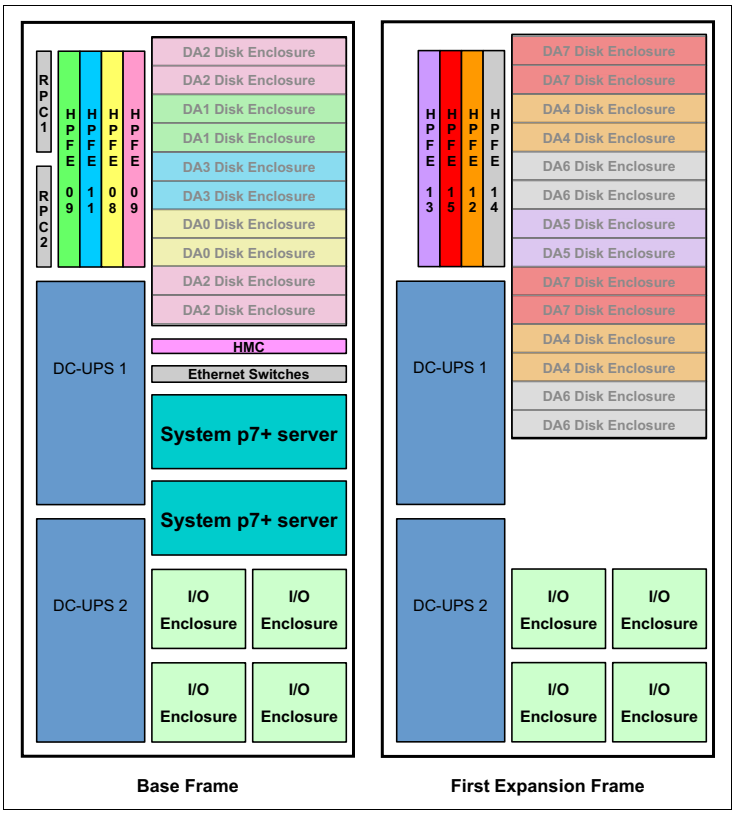

Figure 10 DS8870 enterprise class configuration with 8 HPFEs installed

Figure 11 provides a summary of HPFE and Flash cards supported in DS8870 configurations. For more information about DS8870 configurations and supported components, see *IBM DS8870 Architecture and Implementation,* SG24-8085.

Processor Cores	System Memory (GB)[1]	Max HPFEs Base Frame	Max HPFEs 1st Exp Frame	Max Flash Cards
High Performance All-Flash				
8-core	256	8	N/A	240
16-core	512	8	N/A	240
24-core	1024	8	N/A	240
Enterprise Class and Business Class Configurations				
2-core	16	0	N/A	0
2.core	32	2	N/A	60
4-core	64	4	N/A	120
8-core	128	4	4	240
8-core	256	4	4	240
16-core	512	4	4	240
16-core	1024	4	4	240
Notes: 1 - System memory is a total of memory from both CPCs 2 - Number of HPFEs is independent of number of standard drive enclosures				

Figure 11 DS8870 HPFE supported configurations

Connectivity

The high-performance flash enclosures are directly attached to the I/O enclosure using PCIe cabling, which increases bandwidth and transaction-processing capability compared to standard drive enclosures connectivity (FC-AL). The flash interface cards extend the I/O enclosure PCIe connections to the HPFE. Each flash interface card is a PCIe redrive card that is located in the I/O enclosure.

Figure 12 is a block diagram that shows a simplified view of the PCIe cabling topology.

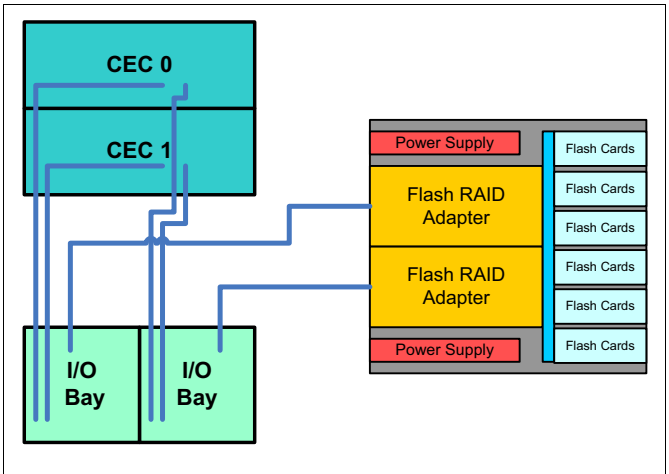

Figure 12 DS8870 HPFE PCIe cabling block diagram

Figure 13 shows the locations of flash interface cards in the DS8870 I/O enclosure. Host adapters (HA) and device adapters (DA) are also shown.

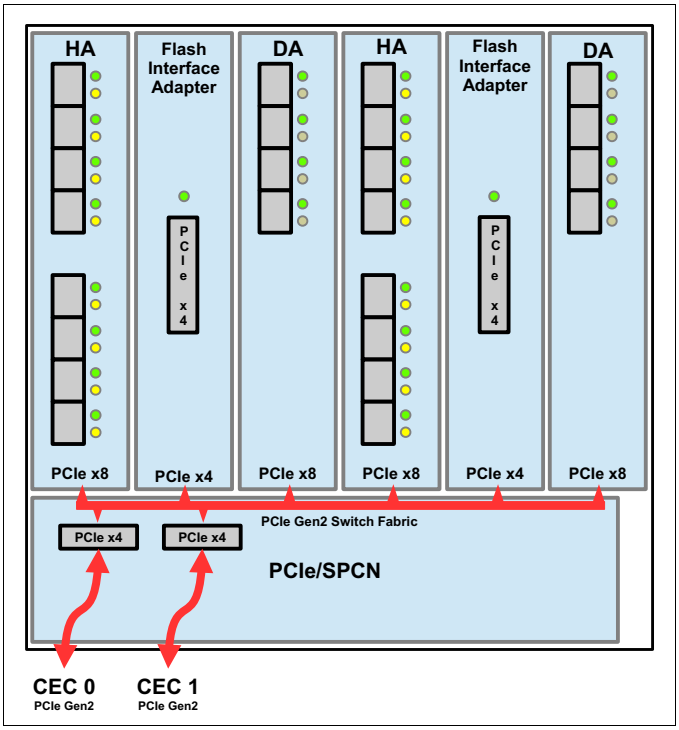

Figure 13 Locations of flash interface cards in I/O enclosure

Easy Tier and flash cards

IBM Easy Tier dynamically optimizes performance for multitiered systems. It can also rebalance data within a single tier to help maintain optimal performance. Now in its seventh generation, Easy Tier offers full support for the high-performance flash enclosure, including Easy Tier Application, and Easy Tier Heat Map Transfer.

All previously announced Easy Tier functions apply, including manual volume migration, automated sub-volume data relocation, automated performance rebalancing within drive tiers (in both single tier and multi-tier drive pools), hot spot management, and rank depopulation.

Currently, IBM supports up to three drive classes or storage tiers that can be configured in the same IBM DS8000:

- Tier 0 or Flash tier: Contains flash cards and flash drives (SSDs). Although flash cards and flash drives are in the same tier, the Easy Tier intra-tier auto-rebalance function will recognize the higher IOPS capability of the HPFE, and migrate hotter extents accordingly.
- Tier 1 or Enterprise tier: Contains the Enterprise drives (SAS 15 k or 10 k rpm).
- Tier 2 or Nearline tier: Contains the Nearline drives (SAS 7.2 k rpm).

For more information, see *IBM DS8000 Easy Tier*, REDP-4667.

Disk Magic

Disk Magic is a Windows based storage subsystem performance modeling tool that is used by IBM and IBM Business Partners to model storage subsystem performance. It supports disk systems from multiple vendors and offers detailed support for IBM storage systems. Contact your IBM Representative or IBM Business Partner to evaluate a Disk Magic study.

Disk Magic supports the high-performance flash enclosure in the DS8880, and in DS8870 with Licensed Machine Code (LMC) R7.3 or later.

Upgrades

High-performance flash enclosure upgrades are supported in the DS8880 base and first expansion frames. For DS8870, HPFEs are supported in the base frame with Licensed Machine Code (LMC) R7.3 or later. HPFEs are supported in the first DS8870 expansion frame for Enterprise and Business Class, with Licensed Machine Code (LMC) R7.4 or later.

All upgrades to add HPFEs or flash card sets are nondisruptive, but might require co-requisite system memory and processor core upgrades. For DS8880, see Figure 6 on page 6. For DS8870, see Figure 11 on page 10.

Related information

For more information, see the following documents:

- *IBM DS8870 Architecture and Implementation*, SG24-8085:
  http://www.redbooks.ibm.com/abstracts/sg248085.html
- *IBM DS8880 Architecture and Implementation*, SG24-8323:
  http://www.redbooks.ibm.com/abstracts/sg248323.html
- *IBM System Storage DS8000: Host Attachment and Interoperability*, SG24-8887:
  http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/sg248887.html
- *DS8000 Performance Monitoring and Tuning,* SG24-8318:
  http://www.redbooks.ibm.com/abstracts/sg248318.html
- *IBM DS8870 Copy Services for Open Systems,* SG24-6788:
  http://www.redbooks.ibm.com/abstracts/sg246788.html
- *IBM DS8870 Copy Services for IBM z Systems,* SG24-6787:
  http://www.redbooks.ibm.com/abstracts/sg246787.html
- *IBM System Storage DS8000 Copy Services Scope Management and Resource Groups*, REDP-4758:
  http://www.redbooks.ibm.com/abstracts/redp4758.html
- *IBM DS8870 Multiple Target Peer-to-Peer Remote Copy*, REDP-5151:
  http://www.redbooks.ibm.com/abstracts/redp5151.html
- *IBM DS8000 Easy Tier*, REDP-4667:
  http://www.redbooks.ibm.com/abstracts/redp4667.html
- *DS8870 Easy Tier Application*, REDP-5014:
  http://www.redbooks.ibm.com/abstracts/redp5014.html

- *IBM DS8870 Easy Tier Heat Map Transfer,* REDP-5015:

  http://www.redbooks.ibm.com/abstracts/redp5015.html
- *DS8000 Thin Provisioning*, REDP-4554:

  http://www.redbooks.ibm.com/abstracts/redp4554.html
- *IBM DS8870 Disk Encryption*, REDP-4500:

  http://www.redbooks.ibm.com/abstracts/redp4500.html
- *DS8000 I/O Priority Manager*, REDP-4760:

  http://www.redbooks.ibm.com/abstracts/redp4760.html
- *Introduction to IBM Assist On-site Software for Storage*, REDP-4889:

  http://www.redbooks.ibm.com/abstracts/redp4889.html
- *LDAP Authentication for IBM DS8000 Storage*, REDP-4505:

  http://www.redbooks.ibm.com/abstracts/redp4505.html
- *Using IBM DS8870 in an OpenStack Environment,* REDP-5220:

  http://www.redbooks.ibm.com/abstracts/redp5220.html
- IBM publication *IBM DS8880 Introduction and Planning Guide*, GC27-8525:

  http://www.ibm.com/support/docview.wss?uid=ssg1S7005228
- IBM publication *IBM DS8870 Introduction and Planning Guide*, GC27-4209:

  http://www.ibm.com/support/docview.wss?uid=ssg1S7004088
- IBM System Storage DS8884 and DS8886 (Machine type 2831) Models 980 and 981 with one-year warranty - Product Announcement (20 October 2015):

  http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/9/897/ENUS115-139/index.html
- IBM System Storage DS8884 and DS8886 (Machine type 2832) Models 980 and 981 with two-year warranty - Product Announcement (20 October 2015):

  http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/2/897/ENUS115-142/index.html
- IBM System Storage DS8884 and DS8886 (Machine type 2833) Models 980 and 981 with three-year warranty - Product Announcement (20 October 2015):

  http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/0/897/ENUS115-140/index.html
- IBM System Storage DS8884 and DS8886 (Machine type 2834) Models 980 and 981 with four-year warranty - Product Announcement (20 October 2015):

  http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/1/897/ENUS115-141/index.html
- IBM DS8880 (Machine type 283x) high-performance flagship Function Authorizations - Product Announcement (20 October 2015):

  http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/8/897/ENUS115-138/index.html
- IBM Assist On-Site:

  http://www.ibm.com/support/assistonsite/
- IBM Tivoli® Assist On-site Remote Support Utility. User's Guide:

  https://aos.us.ihost.com/AssistOnSiteAdmin/docs/AOS_Utility_User_Guide.pdf

- *IBM DS8000 Series Command-Line Interface User's Guide*, SC27-8526:

  http://www.ibm.com/support/docview.wss?uid=ssg1S7002620
- IBM Knowledge Center:

  http://www.ibm.com/support/knowledgecenter/
- DS8880 support:

  https://www.ibm.com/support/entry/portal/product/system_storage/disk_systems/enterprise_storage_servers/ds8880
- DS8870 support:

  https://www.ibm.com/support/entry/portal/product/system_storage/disk_systems/enterprise_storage_servers/ds8870
- IBM System Storage Interoperation Center (SSIC):

  http://www.ibm.com/systems/support/storage/ssic
- IBM data storage feature activation (DSFA):

  http://www.ibm.com/storage/dsfa

Authors

This Product Guide was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.

Jeff Cook joined IBM in 1979 as a large systems Customer Engineer. Prior to joining the DS8000 Product Engineering team in 2011, he worked as a storage systems regional Top Gun for the western United States. He has developed and taught DS8000 and IBM XIV® education in the United States and South America.

Thanks to the following people for their contributions to this project:

Andrew Butcher, Bert Dufrasne, Peter Kimmel, Stephen Manthorpe, Michael Stenson.
IBM

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document, REDP-5318-00, was created or updated on December 9, 2015.

Trademarks

REDP-5318-00

ISBN 0738454761

Printed in U.S.A.