


IBM DS8880 High-Performance Flash Enclosure Gen2

(DS8880 Release 8.5)

Stephen Manthorpe
Axel Westphal



Storage

IBM DS8880 High-Performance Flash Enclosure Gen2

The IBM® DS8880 High Performance Flash Enclosure Gen2 (HPFE Gen2) is a 2U storage enclosure that is installed in pairs.

The HPFE Gen2 pair provides two 2U storage enclosures with associated RAID controllers and cabling. This combination of components forms a high-performance, fully redundant flash storage array.

The HPFE Gen2 pair contains the following hardware components:

- ▶ Two 2U 24-slot Serial Attached SCSI (SAS) flash drive enclosures
- ▶ 16, 32, or 48 encryption-capable 2.5 inch flash drives
- ▶ Each enclosure of the pair contains the following components:
 - Two SAS expander modules with two SAS ports each
 - Two power supplies with integrated cooling fans
 - One midplane or backplane for plugging components that allows maintenance of flash drives, expander modules, and power supplies

Installed with the HPFE Gen2 pair are a pair of flash RAID adapters configured for redundant access to the flash enclosures. Each RAID adapter supports concurrent maintenance and includes the following components:

- ▶ High Performance ASIC RAID engine
- ▶ Four SAS ports and cables connected to the four SAS expander modules, providing fully-redundant access from each RAID adapter to both of the flash enclosures
- ▶ PCIe Gen3 x8 connectivity to the processor nodes via the I/O enclosures

The HPFE Gen2 is only available in DS8880 models. This product guide discusses only DS8880 models 983, 984, 985, 986, and 988.

First generation DS8880 model 980, 981, and 982 are not discussed in this product guide. For information about installing HPFE Gen2 enclosures in first generation models of DS8880, see Appendix E of the *IBM DS8880 Version 8 Release 5 Introduction and Planning Guide*, GC27-8525.

Figure 1 Shows the High Performance Flash Enclosure Gen2, front view.



Figure 1 High Performance Flash Enclosure Gen2 (front view)

High-Performance Flash Enclosure Gen2 highlights

HPFE Gen2 includes features and capabilities, as characterized in the following list:

- ▶ As implemented in the DS8880, the HPFEs Gen2 are directly attached to the PCIe Gen3 fabric with flash RAID adapters, enabling increased bandwidth compared to Fibre Channel attached standard drive enclosures.
- ▶ Flash drives 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.
- ▶ Flash drives offer a number of potential benefits over spinning drives, including higher IOPS, lower power consumption, less heat generation, and lower acoustical noise.
- ▶ Compared to the fibre-attached flash drives (SSDs) installed in the standard drive enclosures, flash drives in HPFE Gen-2 offer even higher throughput using the flash RAID adapters, which have a direct PCIe Gen3 connectivity to the processor complexes.
- ▶ High-performance flash drives are classed as *flash tier 0*. Available flash tier 0 drives capacities include the following options:
 - 400 GB
 - 800 GB
 - 1.6 TB
 - 3.2 TB
- ▶ High-capacity flash drives are classed as *flash tier 1* or *flash tier 2*. Available flash tier 1 or flash tier 2 drives capacities include the following options:
 - 3.84 TB (flash tier 1)
 - 7.68 TB (flash tier 2)
- ▶ Flash drives in the HPFE Gen2 support full drive encryption (FDE).
- ▶ Each HPFE Gen2 pair contains up to 48 flash drives allowing up to 153.6 TB of raw capacity based on the 3.2 TB high-performance flash drives or a 368.6 TB raw capacity based on the 7.68 TB high-capacity flash drives.
- ▶ Up to 16 HPFE Gen2 pairs per DS8888F with two expansion racks give you an impressive total of 5,898.2 TB of raw flash capacity.
- ▶ The IBM Easy Tier® intra-tiering auto-rebalance (micro-tiering) feature is used to distribute the workload among traditional flash drives (SSDs) and *flash tier 0, 1, and 2* flash drives according to their IOPS capacity within the storage tier.

High-Performance Flash Enclosure Gen2 components

The following section describes the components of the HPFE Gen2 enclosure.

Flash drives

The HPFE Gen2 pair provides two 2U flash enclosures. The HPFE Gen2 is available with a choices of drive sets. The following drives are available:

- ▶ 2.5-inch flash tier 0 drives
 - 400 GB
 - 800 GB
 - 1.6 TB
 - 3.2 TB
- ▶ 2.5-inch flash tier 1 drives
 - 3.84 TB

- ▶ 2.5-inch flash tier 2 drives
 - 7.68 TB

Note: Intermix of flash tier 0 and flash tier 1 and 2 is not supported.

Flash drive sets

Flash drives are ordered in sets of 16. The HPFE Gen2 pair can contain 16, 32, or 48 flash drives (1, 2, or 3 drive sets). All flash drives in an HPFE Gen2 pair must be the same type. Half the drive set is installed in each enclosure of the pair. Figure 2 shows the HPFE Gen2 flash drive set install order. Flash drives and fillers are installed from the front of the enclosure.



Figure 2 Flash drive set install order

Storage-enclosure fillers

Storage-enclosure fillers fill empty drive slots in the storage enclosures. The fillers ensure sufficient airflow across populated storage. For HPFE Gen2, one filler feature provides a set of 16 fillers (feature code 1699).

HPFE Gen2 flash drive features

Table 1 lists the available feature codes for flash drive sets for HPFE Gen2.

Table 1 Feature Codes for HPFE Gen2 flash drive sets

Feature code	Disk size	Drive type	RAID support
1610	400 GB	flash tier 0	5,6,10
1611	800 GB	flash tier 0	5,6,10
1612	1.6 TB	flash tier 0	6,10 ^{1,2}
1613	3.2 TB	flash tier 0	6,10 ^{1,2}
1623	3.84 TB	flash tier 1	6,10 ^{1,2}
1624	7.68 TB	flash tier 0	6 ^{1,2}

Feature code	Disk size	Drive type	RAID support
Note:			
1. RAID 5 is not supported for drives larger than 1 TB, and requires a request for price quote (RPQ).			
2. RAID 6 is the default and preferred RAID type for all drives larger than 1 TB, and it is the only supported RAID type for 7.68 TB drives.			
3. Within a High-Performance Flash Enclosure Gen2 pair, no intermix of flash tier 0, flash tier 1, or flash tier 2 is supported.			

RAID capacities for HPFE Gen2 drive sets

Use the following information to calculate the physical and effective capacity for the HPFE Gen2. The default RAID type for all drives larger than 1 TB is RAID 6, and it is the only RAID type supported for 7.68 TB drives. RAID 5 is not supported for drives larger than 1 TB, and requires a request for price quote (RPQ).

Note: A RAID intermix within an HPFE Gen2 pair is not permitted. After the first array is created, the following arrays must be created with the same RAID type.

Enclosure SAS expanders

The enclosure also includes two redundant SAS expanders (also known as *Electronic Control Modules* or ECMs). They provide SAS connectivity from the flash RAID adapters to the HPFE G2 enclosure. Enclosure SAS expanders and power supplies are installed from the rear of the enclosure. For more details about the connectivity of the HPFE Gen2, see [“Flash RAID adapter connectivity” on page 6](#).

Enclosure power supplies

Each HPFE Gen2 has a pair of fully redundant power supply units (PSU). Each PSU has its own integrated fan.

Figure 3 is a rear view of the HPFE Gen2 enclosure showing the redundant SAS expanders and power supplies.




Figure 3 HPFE Gen2 rear view showing SAS expanders and power supplies

Enclosure midplane

The enclosure midplane provides the connectivity for the two SAS expander modules, two power supplies, and 24 flash drive slots for each flash enclosure.

Flash RAID adapters

The flash RAID adapters are PCIe adapters that are either installed in the DS8880 I/O enclosures, or remotely connected to the I/O enclosures through PCIe cables.

The flash RAID adapters have a PCIe3 eight-lane connection to the I/O enclosures, which provide PCIe connectivity to the processor nodes of the DS8880.

The main processor is a RAID engine that provides RAID and sparing management to the flash drives in the HPFE Gen2 flash enclosures. Each flash RAID adapter has four SAS ports, which provide connectivity from the flash RAID adapters to the HPFE Gen2 enclosures.

The flash RAID adapters are installed as a pair, one in each of an I/O enclosure pair. This is known as a *device adapter pair* (DA pair). Logical configuration should be balanced across the DA pair for load balancing and the highest throughput. The redundant DA pair ensures continued availability in the event of a flash RAID adapter or a logical I/O enclosure failure.

The DS8880 flash RAID adapter is specifically designed for connectivity and management of the DS8880 HPFE Gen2.

The DS8880 flash RAID adapter is available in three different form factors, depending on the DS8880 model and location within that model. Internally, the three different form factors have the same core hardware and function.

To differentiate between the three form factors, they have unique naming and features:

- ▶ Microbay flash RAID adapter
 - Remotely connected to the I/O enclosures by a PCIe3 x8 cable to a standalone enclosure
 - Connects to HPFE Gen2 pairs in models 984, 985, 986, and the first eight HPFE Gen2 pairs for model 988
 - The microbay enclosure has its own power supplies and integrated cooling
 - Feature code 1600 is a pair of HPFE Gen2 storage enclosures, and a pair of Microbay flash RAID adapters and all associated cabling
- ▶ SAS flash RAID adapter
 - Installed directly into a PCIe3 x8 adapter slot in the I/O enclosure
 - Connects to HPFE Gen2 pairs nine to sixteen in model 988
 - Feature code 1602 is a pair of HPFE Gen2 storage enclosures (no flash RAID adapters)
 - Feature code 1604 is a pair of SAS flash RAID adapters and associated cabling
- ▶ Base I/O expander with flash RAID adapter
 - A PCIe3 x8 adapter that is physically imbedded into the base PCIe I/O expander, which is installed in the DS8882F 2U I/O enclosure
 - Exclusively available only in the DS8882F model 983

For more information about the locations of HPFE Gen2 storage enclosures, Microbay flash RAID adapters, and SAS flash RAID adapters, see “DS8880 models support for HPFE Gen2” on page 8.

For more information about the DS8882F and its HPFE Gen2 enclosures and associated flash RAID adapters, see *Introducing the IBM DS8882F Rack Mounted Storage System*, REDP-5505.

Flash RAID adapter connectivity

The flash RAID adapter pairs provide high bandwidth, redundant PCIe3 x8 connectivity to the DS8880 processor nodes through the I/O enclosures.

The flash RAID adapter pairs then provide redundant SAS connectivity to the HPFE Gen2 pairs.

Figure 4 on page 6, Figure 5 on page 7, and Figure 6 on page 7 show the connectivity of each of the three flash RAID adapter form factors.

Figure 4 shows PCIe and SAS connectivity of the Microbay flash RAID adapters. Microbay flash RAID adapters are installed with HPFE Gen2 pairs on DS8880 models 984, 985, 986, and the first eight HPFE Gen2 pairs for model 988.

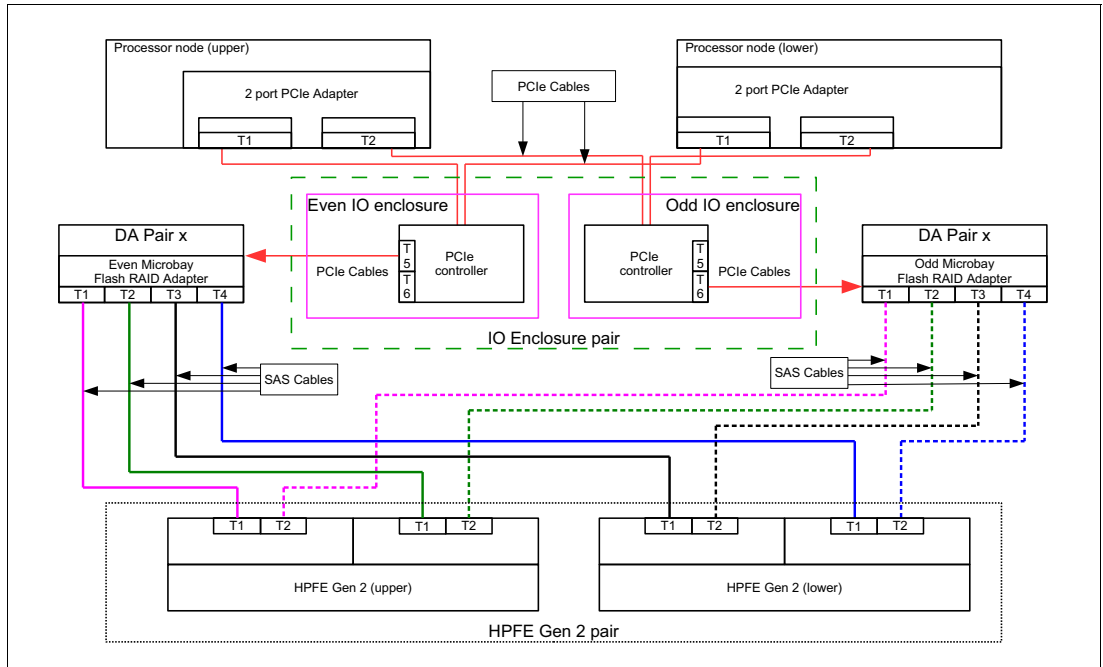


Figure 4 Microbay flash RAID adapters PCIe and SAS connectivity

Figure 5 Shows PCIe and SAS connectivity of the SAS flash RAID adapters. SAS flash RAID adapters are installed with HPFE Gen2 pairs nine to sixteen in model 988.

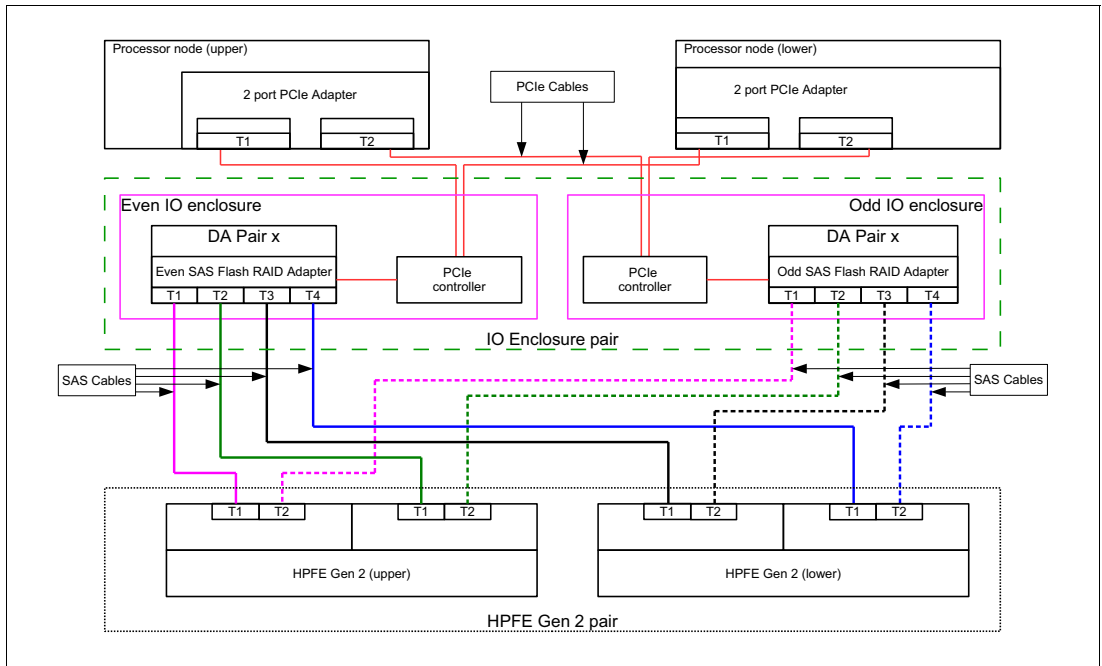


Figure 5 SAS flash RAID adapters PCIe and SAS connectivity

Figure 6 Shows PCIe and SAS connectivity of the base I/O expander with flash RAID adapter, exclusively available for connectivity to the HPFE Gen2 pair installed in DS8882F model 983.

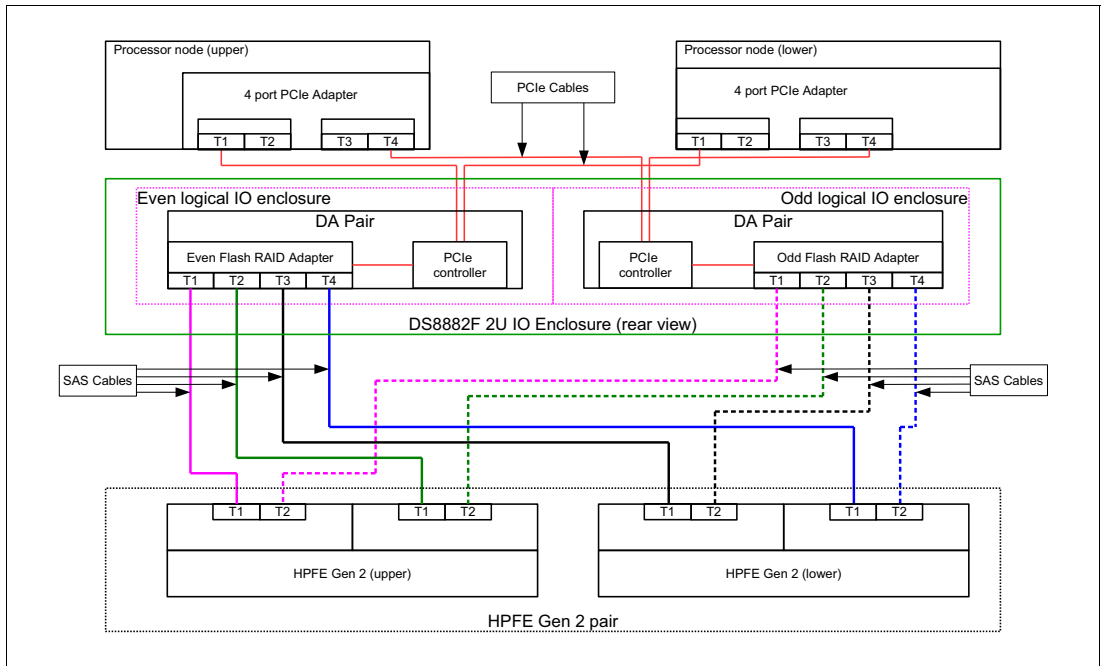


Figure 6 DS8882F flash RAID adapters PCIe and SAS connectivity

DS8880 models support for HPFE Gen2

The DS8880 family encompasses a total of 8 models that support HPFE Gen2 flash enclosures, associated flash RAID adapters and flash drives. The families are the all-flash 533x machine type and the hybrid 283x machine type. This section provides a summary of the HPFE Gen2 pairs the two machine types and associated models support.

Note: For information about HPFE Gen2 support for first generation DS8880 models 980, 981 and 982, go to Appendix E of the *IBM DS8880 Version 8 Release 5 Introduction and Planning Guide*, GC27-8525.

DS8880 all-flash models

An all-flash model means the system supports only flash drives installed in HPFE Gen-2 drive enclosures.

There are five DS8880 all-flash models belonging to the 533x machine type:

- ▶ DS8888F Analytic class (model 988)
- ▶ DS8886F Enterprise class 3 phase (model 986)
- ▶ DS8886F Enterprise class 1 phase (model 985)
- ▶ DS8884F Business class (model 984)
- ▶ DS8882F Rack Mounted (model 983)

Figure 7 Summarizes maximum HPFE Gen2 pairs and flash drives per frame of each all-flash model (machine type 533x) based on processor and memory configuration.

Models (machine type 533x)	Processor Cores	Total System Memory (GB)	Max # HPFEs Gen 2 pairs base frame	Max # HPFEs Gen 2 pairs 1st exp frame	Max # HPFEs Gen 2 pairs 2nd exp frame	Max # flash drives per system
DS8888F	24	1024	4	N/A	N/A	192
	48	2048		12	12	768
DS8886F	8	128	4	N/A	N/A	192
		256				
	16	256	4	8	N/A	384
		512				
		1024				
24	2048					
DS8884F	6	64	2	N/A	N/A	96
		128	4	N/A	N/A	192
		256				
	12	256				
DS8882F	6	64	1	N/A	N/A	48
		128				
		256				

Figure 7 HPFE Gen2 pairs and flash drives by processor and memory for machine type 533x

Diagrams in the following pages show the maximum HPFE Gen2 pairs supported by each all-flash model (machine type 283x), and the installed locations and DA pair IDs. Install order is from the bottom up in each frame, filling each frame to the maximum allowed pairs for that frame, then populating the next frame.

Figure 8 Shows DS8888F Analytic class (model 988, machine type 533x) HPFE Gen2 pairs:

- ▶ 16 HPFE Gen2 pairs maximum for the system
- ▶ 4 HPFE Gen2 pairs in the base frame
- ▶ 6 HPFE Gen2 pairs in first expansion frame
- ▶ 6 HPFE Gen2 pairs in second expansion frame

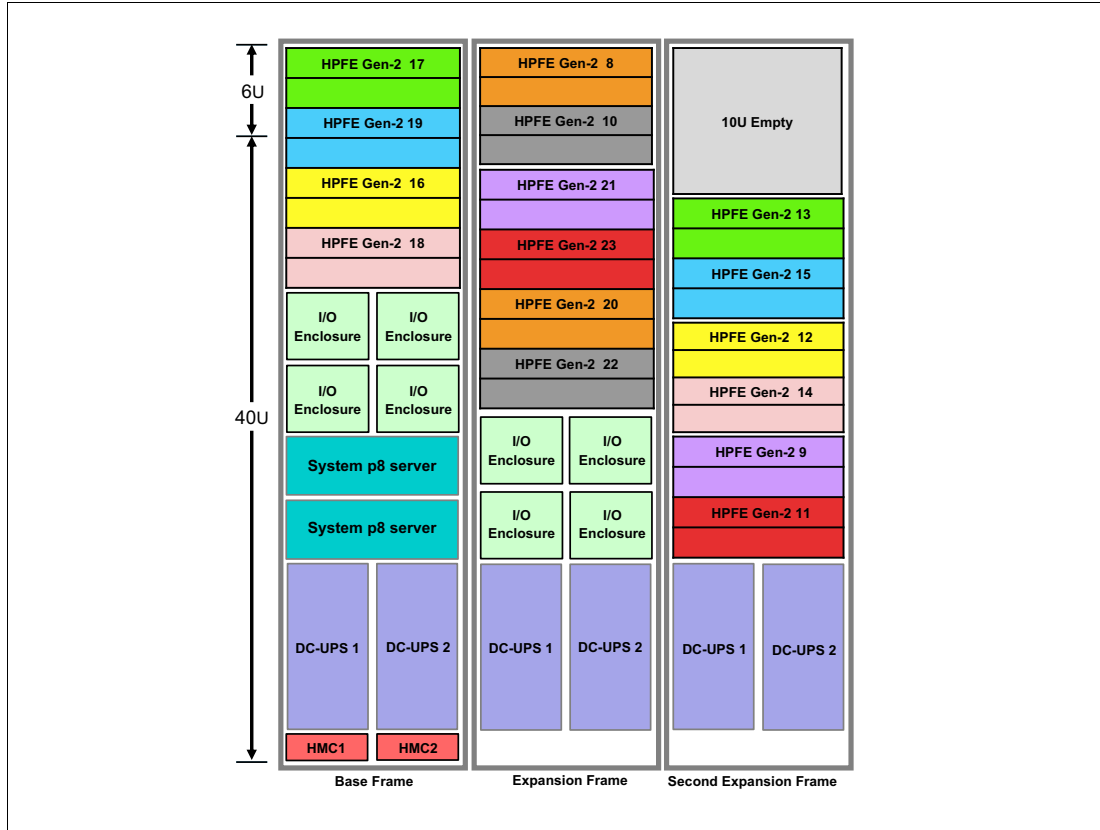


Figure 8 DS8888F (machine type 533x) HPFE Gen2 maximum 16 pairs

Figure 9 shows DS8886F Enterprise class (model 986, machine type 533x) HPFE Gen2 pairs. The model 986 (3 phase) and the model 985 (1 phase) support the same number of HPFE Gen2 pairs per frame and system total:

- ▶ 8 HPFE Gen2 pairs maximum for the system
- ▶ 4 HPFE Gen2 pairs in the base frame
- ▶ 4 HPFE Gen2 pairs in first expansion frame

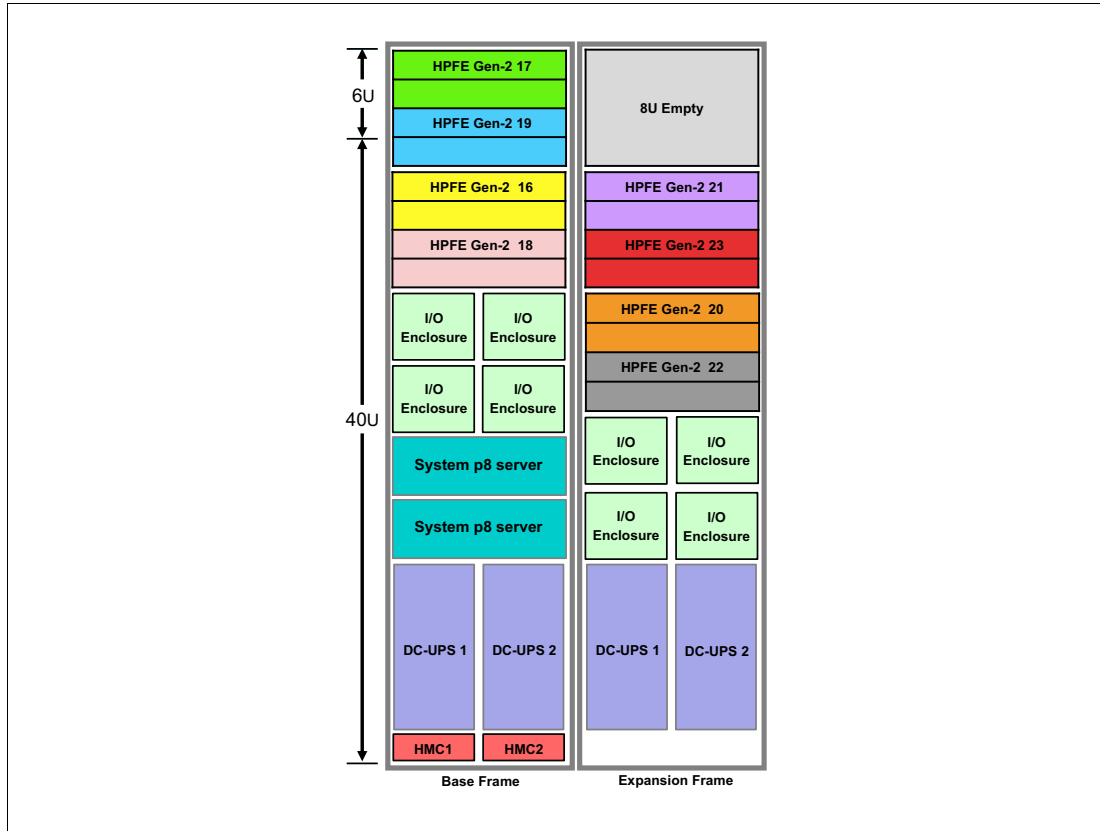


Figure 9 DS8886F (machine type 533x) HPFE Gen2 maximum 8 pairs

Figure 10 shows DS8884F Business class (model 984, machine type 533x) HPFE Gen2 pairs:

- ▶ 4 HPFE Gen2 pairs maximum for the system
- ▶ 4 HPFE Gen2 pairs in the base frame

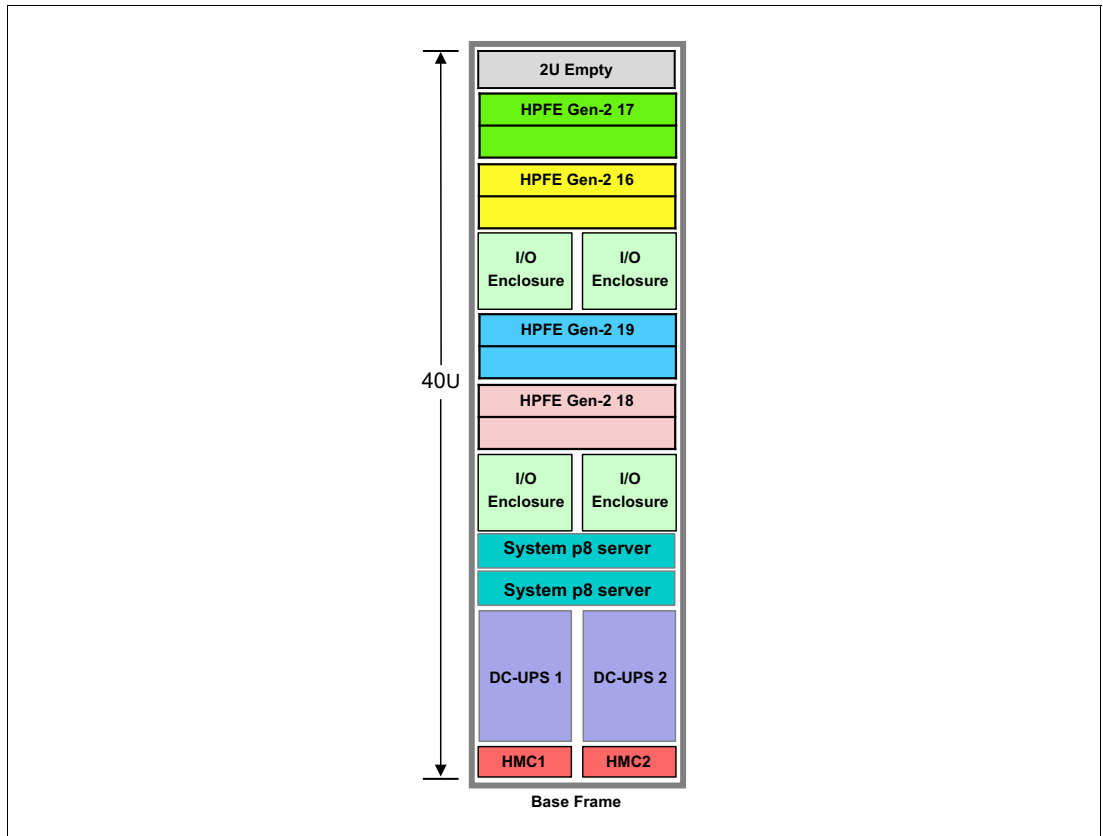


Figure 10 DS8884F (machine type 533x) HPFE Gen2 maximum 4 pairs

Figure 11 Shows DS8882F Rack Mounted (model 983, machine type 533x) HPFE Gen2 pairs:

- ▶ 1 HPFE Gen2 pairs maximum for the system

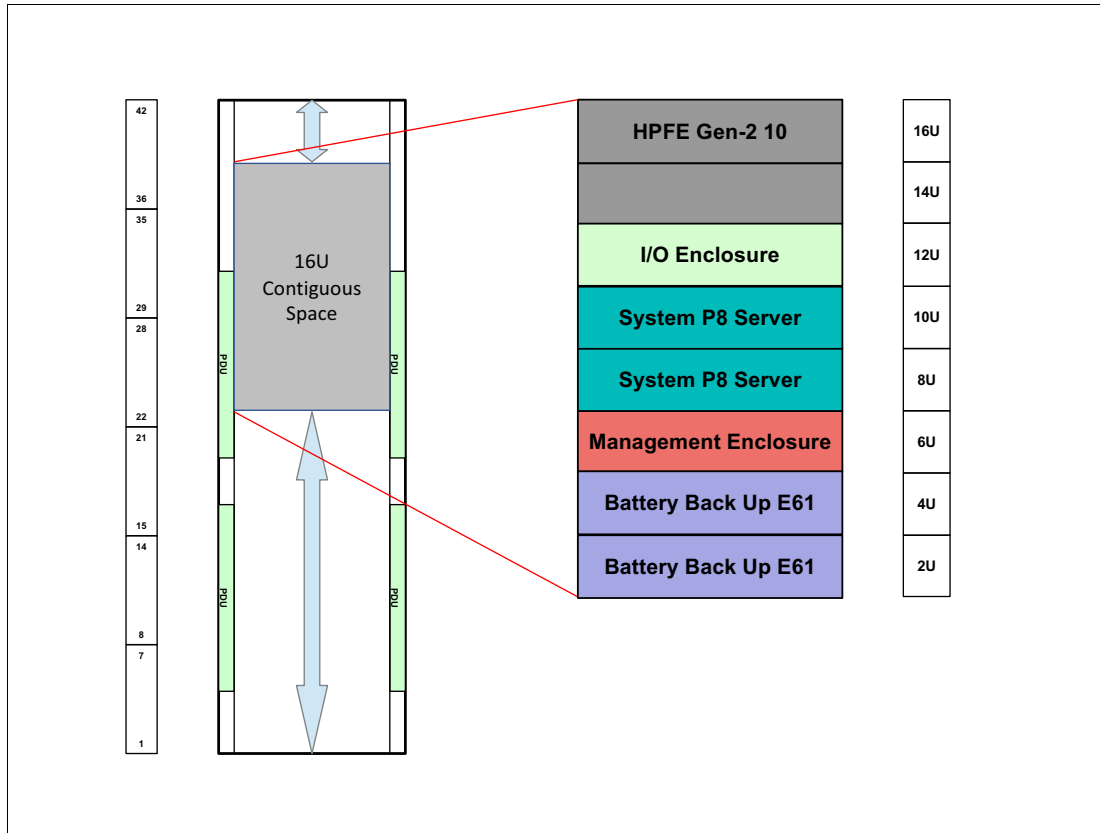


Figure 11 DS8882F (machine type 533x) HPFE Gen2 maximum 1 pair

DS8880 hybrid models

A hybrid model is one in which the system supports both flash drives installed in HPFE Gen-2 drive enclosures and standard drive enclosure pairs with spinning drives.

There are three DS8880 hybrid models belonging to the 283x machine type family that support HPFE Gen2 pairs and flash drives:

- ▶ DS8886 Enterprise class 3 phase (model 986)
- ▶ DS8886 Enterprise class 1 phase (model 985)
- ▶ DS8884 Business class (model 984)

Figure 12 Summarizes maximum HPFE Gen2 pairs and flash drives per frame of each hybrid model (machine type 283x) based on processor and memory configuration.

Models (machine type 283x)	Processor Cores	Total System Memory (GB)	Max # HPFEs Gen 2 pairs base frame	Max # HPFEs Gen 2 pairs 1st exp frame	Max # flash drives per system
DS8886	8	128	2	N/A	96
		256			
	16	256	2	2	192
		512			
		1024			
24	2048	2	2	192	
	2048				
DS8884	6	64	1	N/A	48
		128	1	1	96
	256				
	12	256	1	1	96

Figure 12 HPFE Gen2 pairs and flash drives by processor and memory for machine type 283x

The following graphics show the maximum HPFE Gen2 pairs supported by each hybrid model (machine type 283x), and the installed locations and DA pair IDs. Install order is from the bottom up in each frame, filling each frame to the maximum allowed pairs for that frame, then populating the next frame.

Figure 13 Shows DS8886 Enterprise class (model 986, machine type 283x) HPFE Gen2 pairs. The model 986 (3 phase) and the model 985 (1 phase) support the same number of HPFE Gen2 pairs per frame and system total:

- ▶ 4 HPFE Gen2 pairs maximum for the system
- ▶ 2 HPFE Gen2 pairs in the base frame
- ▶ 2 HPFE Gen2 pairs in first expansion frame

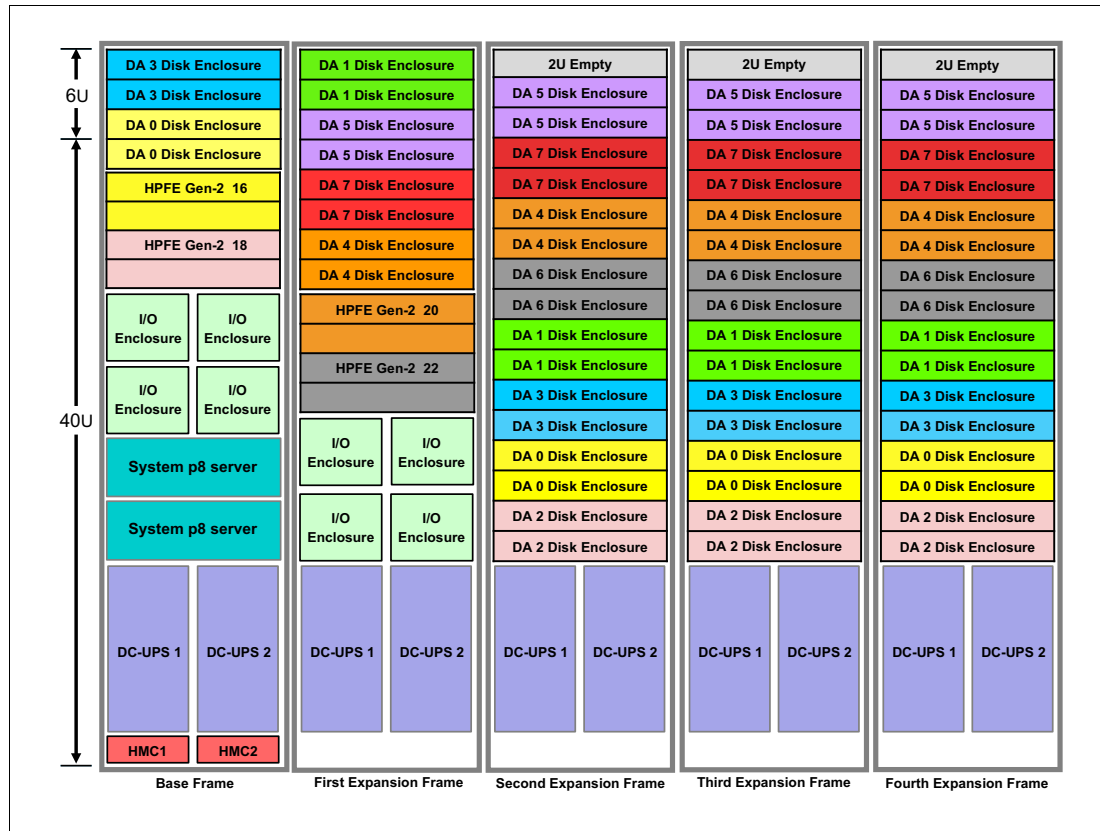


Figure 13 DS8886 (machine type 283x) HPFE Gen2 maximum 4 pairs

Figure 14 Shows DS8884 Business class (model 984, machine type 283x) HPFE Gen2 pairs:

- ▶ 2 HPFE Gen2 pairs maximum for the system
- ▶ 1 HPFE Gen2 pairs in the base frame
- ▶ 1 HPFE Gen2 pairs in first expansion frame

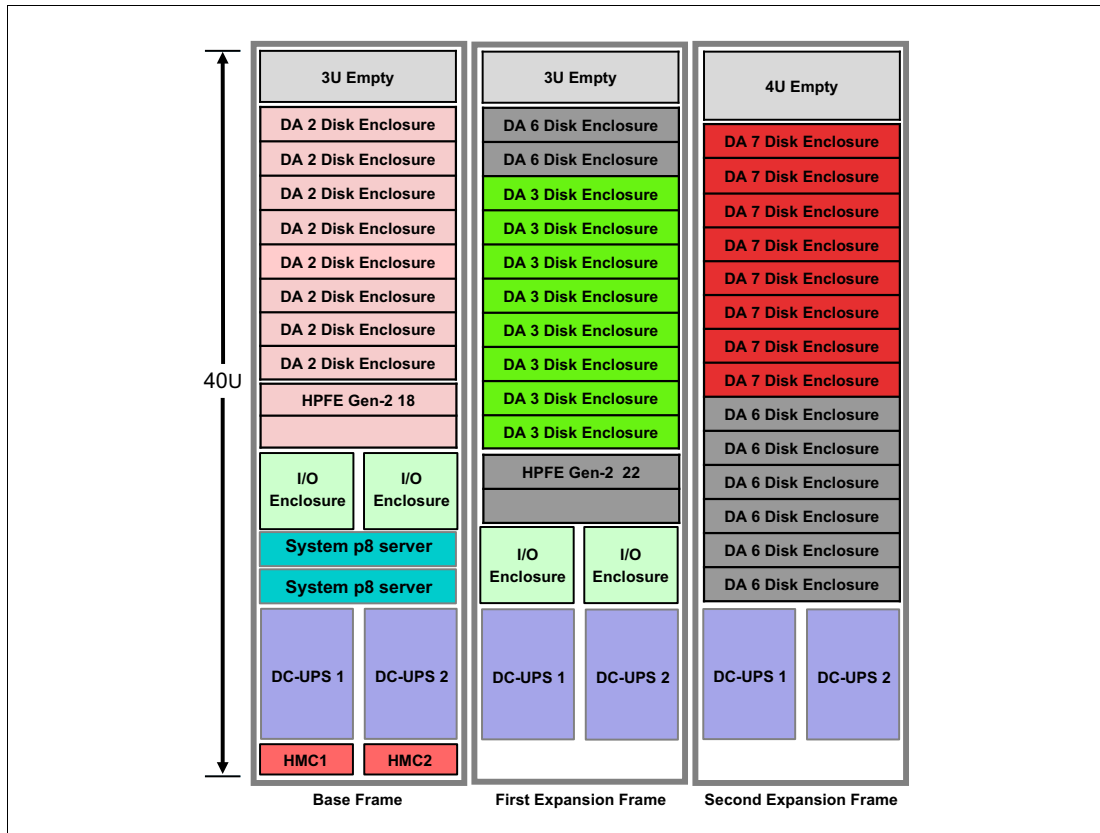


Figure 14 DS8884 (machine type 283x) HPFE Gen2 maximum 2 pairs

Upgrades

Adding HPFE Gen2 pairs (inclusive of associated flash RAID adapters pairs) to models DS8880 models 984, 985, 986, and 988 is supported. Adding drive sets to partially populated enclosure pairs to DS8880 models 984, 985, 986, and 988 is also supported.

All upgrades to add HPFEs Gen2 pairs or flash drive sets are non disruptive. However, upgrades might require co-requisite system memory and processor core upgrades. For DS8880 all-flash models (machine type 533x), see Figure 7 on page 8. For DS8880 hybrid models (machine type 283x) see Figure 12 on page 13.

For additional information about upgrades, see *IBM DS8880 Architecture and Implementation (Release 8.5)*, SG24-8323.

Easy Tier and flash drives

IBM Easy Tier dynamically optimizes performance for multi-tiered systems. It can also rebalance data within a single tier to help maintain optimal performance. Easy Tier offers full support for the High Performance Flash Enclosure Gen2, including Easy Tier Application and Easy Tier Heat Map Transfer.

Easy Tier Automatic Mode manages any combination of up to three tiers in a storage pool. For an HPFE Gen2 pair, the following drive classes are available, in order from highest to lowest performance:

- ▶ Flash tier 0 drives
 - High-performance flash drives
 - The highest performance drives, which provide high I/O throughput and low latency
- ▶ Flash tier 1 drives
 - The first tier of high capacity flash drives
- ▶ Flash tier 2 drives
 - The second tier of high capacity flash drives

Note: Intermix between flash tier 0 and flash tier 1 or 2 drives in the same HPFE Gen2 pair is not supported.

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

Disk Magic

Disk Magic is a Windows-based storage system 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 Gen2 in the DS8880 with Licensed Machine Code (LMC) R8.2 or later.

Related information

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this paper.

- ▶ *IBM DS8880 Architecture and Implementation (Release 8.3)*, SG24-8323
<http://www.redbooks.ibm.com/abstracts/sg248323.html>
- ▶ *Introducing the IBM DS8882F Rack Mounted Storage system*, REDP-5505
- ▶ *IBM DS8000 Easy Tier*, REDP-4667
<http://www.redbooks.ibm.com/abstracts/redp4667.html>
- ▶ IBM publication *IBM DS8880 Introduction and Planning Guide*, GC27-8525
<http://www.ibm.com/support/docview.wss?uid=ssg1S7005228>
- ▶ DS8880 support
https://www.ibm.com/support/entry/portal/product/system_storage/disk_systems/enterprise_storage_servers/ds8880
- ▶ IBM Knowledge Center
<http://www.ibm.com/support/knowledgecenter/>

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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, MD-NC119, Armonk, NY 10504-1785, US

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 jurisdictions 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

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

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

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

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 actual people or business enterprises 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

Redbooks (logo) ®
DS8000®

Easy Tier®
IBM®

The following terms are trademarks of other companies:

Other company, product, or service names may be trademarks or service marks of others.



REDP-5422-02

ISBN 073845723x

Printed in U.S.A.

Get connected

