IBM Systems & Technology Group

IBM ® TotalStorage ® DS8000

Rosemary McCutchen
Storage Systems Advanced Technical Support

ON DEMAND BUSINESS™

Topics

- DS8000 Enhancements
- DS8000 Logical Configuration
- z/OS Software Support
- z/OS HCD Considerations
- Advanced Functions
- Performance Monitoring
- Benefits
- References

DS8000 Enhancements

- 73GB, 146GB, 300GB switched bunch of disk (SBOD) disk drives (up to 640)
- pSeries Power5
- 4-port FC/FICON Host Adapter cards (up to 128 ports)
- Device Adapters scale with capacity
- 64K devices; 255 Logical Control Units per logical DS8000 (FICON)
- Smaller footprint
- Two logical DS8000s in one physical DS8000 (9A2 'LPAR')
- Metro and Global Mirror compatibility with ESS and DS6000

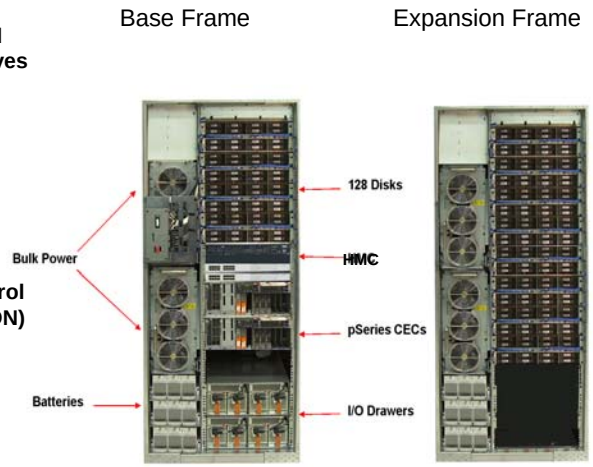

DS8000 Physical Overview

	DS8100 921 (Base Frame Only)	DS8100 921 + Expansion Frame	DS8300 922 or 9A2 (Base Frame Only)	DS8300 922 or 9A2 + Expansion Frame	DS8300 922 or 9A2 + 2 Expansion Frames
Processor - pSeries Power5	2-way 1.5GHz	2-way 1.5GHz	4-way 1.9GHz	4-way 1.9GHz	4-way 1.9GHz
Cache up to 256GB	16 to 128 GB	16 to 128 GB	32 to 256 GB	32 to 256 GB	32 to 256 GB
Expansion Rack	Yes (1)	--	Yes (1 or 2)	--	--
Host Adapters up to 128 ports - per HA: 4-port FC / FICON (2 Gb) - per HA: 2-port ESCON	2 to 16 (e.g., 8 to 64 FC / FICON ports)	2 to 16	2 to 16	2 to 32	2 to 32
Device Adapter Pairs	1 to 4	1 to 4	1 to 4	1 to 8	1 to 8
Drives up to 640 disks - 73 GB (15K rpm) - 146 GB (10K rpm) - 300 GB (10K rpm)	16 to 128 (Increments of 16)	16 to 384 (increments of 16)	16 to 128 (Increments of 16)	16 to 384 (increments of 16)	16 to 640 (Increments of 16)
Physical Capacity up to 186 TB	1.1 to 37.2 TB	1.1 to 111.6 TB	1.1 to 37.2 TB	1.1 to 111.6 TB	1.1 to 186.3 TB
Power	Three-Phase	Three-Phase	Three-phase	Three-Phase	Three-phase
Dimensions - Height x Width x Depth - Footprint	76 x 33.25 x 43 in 9.93 sq. ft.	76 x 66.5 x 43 in 19.86 sq. ft.	76 x 33.25 x 43 in 9.93 sq. ft.	76 x 66.5 x 43 in 19.86 sq. ft.	76 x 99.75 x 43 in 29.79 sq. ft.

DS8000 Performance Enhancers

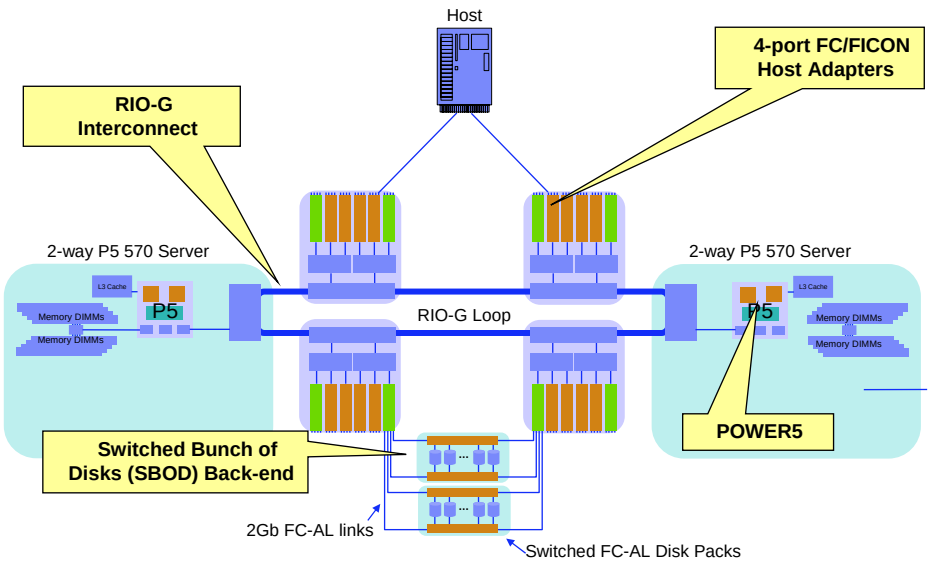

ESS 800 Arrays and Device Adapters

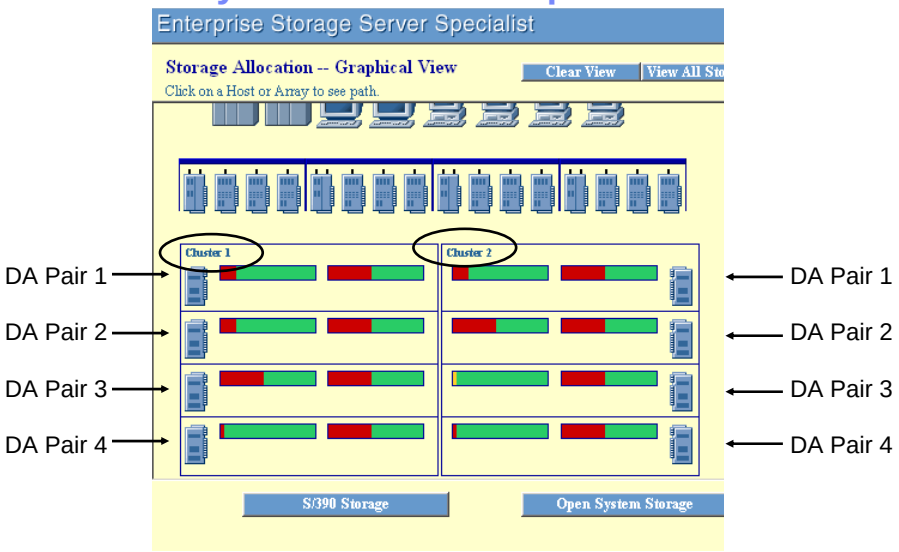

DS8100 Disks and Device Adapters (DAs)

- **8 arrays (64 disks) on DA0**
- **DA Pair 0**
- **Up to 4 DA pairs**
  - DA pairs 0 to 3
  - 2 DA pairs (DA0 and DA2) used for disks in full base frame
- **Up to 16 Host Adapters**
  - Up to 64 host ports

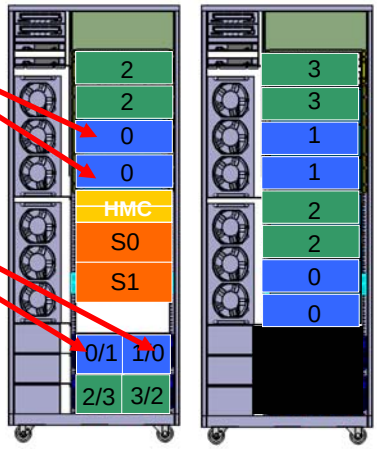

DS8300 Disks and Device Adapters (DAs)

- **Up to 8 DA pairs**
  - DA pairs 0 to 7
  - 2 DA pairs for disks in full base frame
  - 6 DA pairs for disks in full A+B frames
  - 8 DA pairs for disks in full A+B + half full C frame
- **Up to 32 Host Adapters**
  - Up to 128 FICON/FCP host ports
  - 64 ports in base frame
  - 64 ports in expansion frame

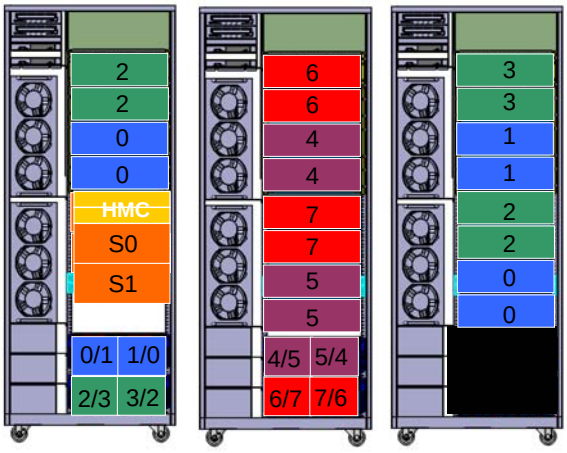

DS8300 2107-9A2 LPAR – Dual Storage Images

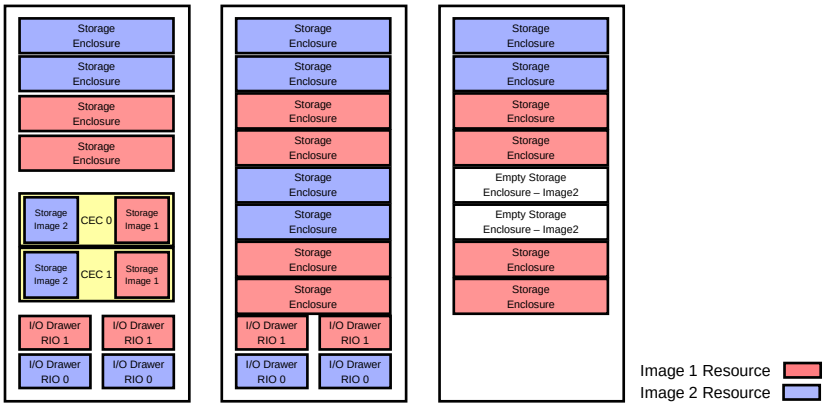

- Current implementation is 2 Storage Images (logical DS8000s)
  - 50% of total host adapters, device adapters and disks dedicated to each image

DS8000 Enhanced Addressing Capability

	ESS 800	DS8000	DS8000 w/LPAR
Max Logical Subsystems	32	255	510
Max Logical Devices	8K	64K	128K
Max Logical CKD Devices	4K	64K	128K
Max Logical FB Devices	4K	64K	128K
Max N-Port Logins/Port	128	509	509
Max N-Port Logins	512	8K	16K
Max Logical Paths/FC Port	256	2K	2K
Max Logical Paths/CU Image	256	512	512
Max Path Groups/CU Image	128	256	256

DS8000 Logical Configuration Enhancements

DS8000 Array Site

- **Logical grouping of disks**
  - Same capacity (73 GB, 146 GB, 300 GB)
  - Same speed (10K rpm, 15K rpm)
  - Determined by software during installation
  - Assigned to Device Adapter (DA) Pair
    - DA assignment determined by software during installation
- **DS8000**
  - 8 disks (DDMs)

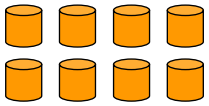

DS8000 array site

DS8000 Array and Rank

- **RAID type and storage type now defined in separate steps**
  - Array defines RAID type (RAID5 or RAID10)
  - RANK defines storage type (CKD or FB)
- **RAID array**
  - One DS8000 array site becomes one array
  - RAID5 6+P, 7+P
  - RAID10 3+3, 4+4
  - Parity is striped across all disks in array but consumes capacity equivalent to one disk
- **RANK**
  - One DS8000 array becomes one rank
  - CKD or FB
  - Divided into extents
    - 3390 Mod1 (1113 cylinders/.94GB)

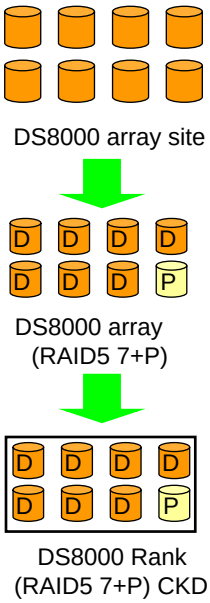

ESS 800 Diskgroups and Logical Control Units

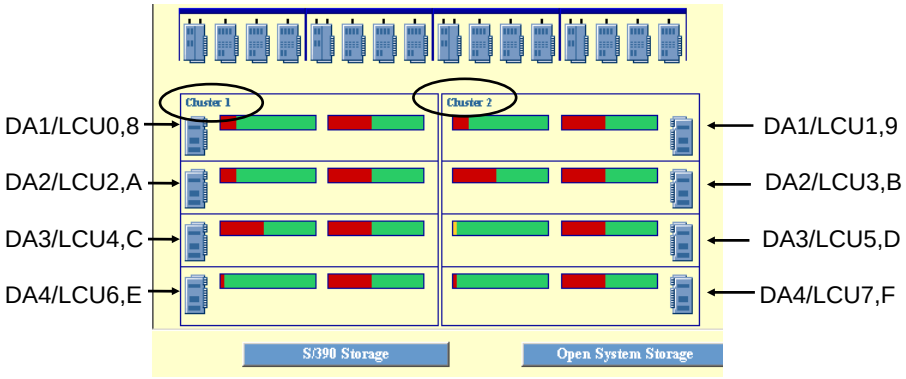

- **Unlike ESS diskgroups, DS8000 ranks have no pre-determined or fixed relation to Server0/Server1 or Logical Control Units (LCUs)**
- **DS8000 ranks are associated with Server0/Server1 when assigned to an extent pool**

DS8000 Logical Control Unit and Address Group

- **Similar to ESS:**
  - LCU has a maximum of 256 addresses
  - Aliases are shared within an LCU
  - LCU is the basis for Copy Services paths and consistency groups
  - Even LCUs are associated with Server0 and odd LCUs are associated with Server1
- **New:**
  - LCU does not have a pre-determined relationship to rank/DA pair
  - Up to 255 LCUs are available
  - Set of 16 Logical Subsystems (LSSs) is called an Address Group
    - LSS 00-0F, 10-1F, 20-2F, 30-3F, etc.
    - Storage type for entire Address Group (16 LSSs) is set to CKD or Fixed Block by the first LSS defined
    - ESCON devices must be in Address Group 0 (LCU 00-0F)

DS8000 Extent Pool

- **Logical grouping of one or more ranks which have been divided into extents from which volumes will be created**
- **User creates Extent pool & assigns it to Server 0 or 1**
- **User puts one or more Ranks in the extent pool**
- **User creates LCU**
- **User creates volume from ONE extent pool**
  - User specifies a 4-digit hex volume ID which includes address group, LCU and device ID:
    - Example: Volume ID x'2A10'
      - 2=Address Group 2 (group of 16 LCUs 20-2F)
      - 2A=LCU 2A (2A is 'even'; may only be used with Server0 pools)
      - 10=device ID
  - Even LCUs are available for Server0 extent pools
  - Odd LCUs are available for Server1 extent pools

Extent Pool P0
Assigned to Server0
'Even' LCUs available

DS8000 Flexible LCU Assignment

- User has created extent pools for Server0 and Server1 and placed one or more ranks in each extent pool
- Volume creation associates LCU with rank(s)
- One LCU per rank
  - Simplifies management
- Multiple LCUs on one rank
  - Provides additional addresses for rank
- One LCU across multiple ranks
  - Enables cross-rank Copy Services consistency group

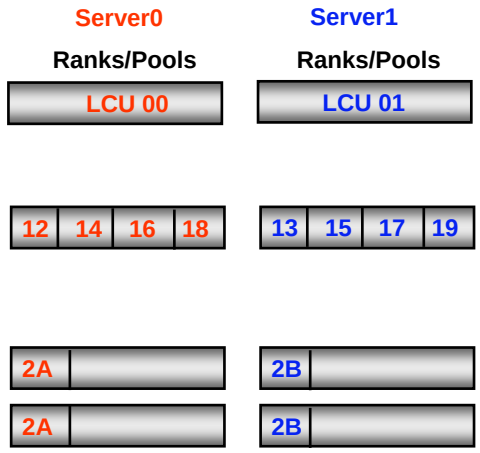

DS8000 Volumes

- CKD standard volumes
  - 3380, 3390M3, 3390M9
- CKD customer volumes
  - Minimum specification is 1 cylinder
    - Minimum space allocation is 1 extent (1113 cylinders)
  - Maximum volume size is 64K cylinders/56GB
    - With appropriate z/OS software support
- Maximum number of volumes is 64K per logical DS8000 FICON access (4K max ESCON access)
- Volumes may be dynamically deleted and extents reused
- Current implementation does not provide volume striping across multiple ranks in extent pool
  - Volume may 'spill' across ranks in pool

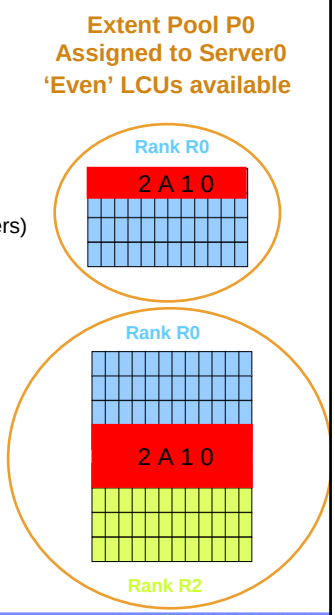

DS8000 Configuration Planning Considerations

- **One extent pool per rank is a good starting point**
  - Balance extent pools and ranks across servers
- **One LCU per rank is a good starting point**
  - Unless more addresses are needed
- **Use a limited number of device types for ease of management**
- **Use customer volumes that are even multiple of 1113 cylinder extents**
  - 3390M3
  - 3390M9
  - 30051 cylinders
  - 60102 cylinders
- **Use dynamic Parallel Access Volumes (PAVs) if possible**
  - ESS users may use the same PAV ratios as for ESS
  - Define as many PAVs as you can afford
- **New DS8000 addressing capability may require zSeries device number and SSID planning**

DS8000 Storage Manager

- **Easy-to-use, powerful and flexible User Interface**
  - Wizards, Filters, Sorts, Hyperlinks, Animation, Copy/Paste
- **Includes optional automated methods**
- **Runs on DS8000 integrated HMC; accessed via Browser**

DS8000 Command Line Interface (DSCLI)

```
#6+P
#Array/Rank/Pool2
#mkckdvol -extpool p2 -cap 3339 -name csfb_m3_#h 0000-0023
#mkckdvol -extpool p2 -cap 3339 -name csfb_m3_#h 0200-0223
#mkckdvol -extpool p2 -cap 3339 -name csfb_m3_#h 0400-0423
#mkckdvol -extpool p2 -cap 3339 -name csfb_m3_#h 0600-0623
#mkckdvol -extpool p2 -cap 3339 -name csfb_m3_#h 0800-0823
#mkckdvol -extpool p2 -cap 3339 -name csfb_m3_#h 3e00-3e0b

#7+P
#Array/Rank/Pool4
#mkckdvol -extpool p4 -cap 3339 -name csfb_m3_#h 0a00-0a23
#mkckdvol -extpool p4 -cap 3339 -name csfb_m3_#h 0c00-0c23
#mkckdvol -extpool p4 -cap 3339 -name csfb_m3_#h 0e00-0e23
#mkckdvol -extpool p4 -cap 3339 -name csfb_m3_#h 3000-3023
mkckdvol -extpool p4 -cap 3339 -name csfb_m3_#h 3200-3223
```

- **Powerful tool for automating configuration tasks and collecting configuration information**
  - Can be used in conjunction with PDF
- **Same DSCLI for DS6000 and for ESS 800 Copy Services**

Disk Storage Feature Authorization (DSFA)

- **Operating Environment (base microcode) and advanced functions require license activation**
- **Machine type, model, serial and signature are used to retrieve authorization keys from IBM website**
  - www.ibm.com/storage/dsfa
- **Keys are activated on DS8000 through the DS Storage Manager or command line interface (DSCLI)**
- **Configuration will be limited by advanced function key with lowest capacity for that storage type**

DS8000 z/OS Software Support

- **z/OS 1.4 and higher**
- **Device support facility (ICKDSF) Release 17**
- **DFSORT**
- **Data Facilities Storage Management Subsystem (DFSMS) SPE**
- **Resource Management Facility (RMF) SPE**
  - APAR OA06476, PTF UA90079, UA90080
- **DS8000 supported as 2105 without DFSMS and z/OS SPE installed**
- **DEVSERV QDASD and DEVSERV PATHS show device type 2107**
- **Check the 2107device PSP bucket for maintenance required**

DS8000 z/OS HCD Considerations

- **New device support for D/T2107**
- **DS8000 supports up to 16 Address Groups**
  - 65K logical volumes
  - For IOCP and HCD, the CU addresses are hex 00 – FE
  - Logical subsystem (LSS) addresses are 0 – 254
  - LCU / LSS do not have to be contiguous
- **ESCON addressing only supported on Address Group 0 (LCU 00-0F)**
- **If running in 2105 mode, do not have extended addressing**

Address Group 2 (LCU 20-2F)

```
CNTLUNIT CUNUMBR=A000,PATH=(52,53,54,55),
        UNITADD=((00,256)),LINK=(24,34,25,35),
        CUADD=20,UNIT=2107,DESC='N150 LCU20'
CNTLUNIT CUNUMBR=A100,PATH=(52,53,54,55),
        UNITADD=((00,256)),LINK=(24,34,25,35),
        CUADD=21,UNIT=2107,DESC='N150 LCU21'
CNTLUNIT CUNUMBR=A200,PATH=(52,53,54,55),
        UNITADD=((00,256)),LINK=(24,34,25,35),
        CUADD=22,UNIT=2107,DESC='N150 LCU22'
```

Examples provided at: http://w3-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD102127

DS8000 Advanced Functions

- **Parallel Access Volumes**
  - Shared within an LCU
  - Same considerations as ESS
- **Copy Services**
  - Similar to ESS
  - New with DS8000:
    - **New command line interface (DSCLI) provides enhanced support**
    - **9A2 ('LPAR') Storage Images (Point in Time Copy within 1 image)**
    - **TSO commands and ANTRQST support for Global Mirror**
    - **Exploitation of Parallel Access Volumes on secondary devices for Remote Mirror for z/OS (APAR OA09238)**
    - **PPRC links are FCP only (no ESCON PPRC support)**
    - **No Metro/Global Copy**

DS8000 Performance Monitoring

- **New counters reported in RMF**
  - OA06476 APAR must be installed for RMF
  - Cache reports
  - Link statistics
  - Extent Pool statistics
  - Rank statistics

RMF Cache Reports

```
0------------------------------------------------------------------------------------------------------------------------------
                                                CACHE SUBSYSTEM OVERVIEW
-------------------------------------------------------------------------------------------------------------------------------
OTOTAL I/O    29627   CACHE I/O    29627   CACHE OFFLINE        0
 TOTAL H/R    0.813   CACHE H/R    0.813
-CACHE I/O   ------------READ I/O REQUESTS-------------   ----------------------WRITE I/O REQUESTS----------------------          %
 REQUESTS      COUNT     RATE     HITS     RATE    H/R      COUNT     RATE     FAST     RATE     HITS     RATE    H/R       READ
ONORMAL        22168    375.7    16641    282.1  0.751       7459    126.4     7459    126.4     7459    126.4  1.000       74.8
 SEQUENTIAL        0      0.0        0      0.0    N/A          0      0.0        0      0.0        0      0.0    N/A        N/A
 CFW DATA          0      0.0        0      0.0    N/A          0      0.0        0      0.0        0      0.0    N/A        N/A
OTOTAL         22168    375.7    16641    282.1  0.751       7459    126.4     7459    126.4     7459    126.4  1.000       74.8
-------------------------CACHE MISSES-------------------------   -----------MISC------------   ------NON-CACHE I/O------
 REQUESTS       READ     RATE    WRITE     RATE  TRACKS    RATE                      COUNT   RATE                    COUNT   RATE
                                                                  DFW BYPASS            0    0.0          ICL           0    0.0
 NORMAL         5527     93.7        0      0.0    5527    93.7  CFW BYPASS            0    0.0          BYPASS        0    0.0
 SEQUENTIAL        0      0.0        0      0.0       0     0.0  DFW INHIBIT           0    0.0          TOTAL         0    0.0
 CFW DATA          0      0.0        0      0.0             ASYNC (TRKS)       5830   98.8
OTOTAL         5527     RATE    93.7
----CKD STATISTICS---      ---RECORD CACHING---                  ---HOST ADAPTER ACTIVITY---      --------DISK ACTIVITY-------
                                                                          BYTES  BYTES                    RESP   BYTES  BYTES
 WRITE             0     READ MISSES        0                             /REQ   /SEC                     TIME    /REQ   /SEC
 WRITE HITS        0     WRITE PROM      5082          READ              5.5K   2.1M      READ     5.456   52.9K   5.0M
                                                       WRITE             4.1K 519.8K      WRITE    9.235    4.9K 482.1K
1                                  C A C H E   S U B S Y S T E M   A C T I V I T Y                                  PAGE    2

            z/OS V1R6              SYSTEM ID GDP2          START 06/07/2005-21.03.00  INTERVAL 000.00.59
                                   RPT VERSION V1R5 RMF      END   06/07/2005-21.04.00
OSUBSYSTEM  2107-01     CU-ID  8F01     SSID 177F     CDATE  06/07/2005     CTIME 21.03.02     CINT  00.00.59
 TYPE-MODEL 2107-922    MANUF  IBM      PLANT  75     SERIAL 0000000020331
0------------------------------------------------------------------------------------------------------------------------------
                                              CACHE SUBSYSTEM DEVICE OVERVIEW
-------------------------------------------------------------------------------------------------------------------------------
OVOLUME      DEV  XTNT    %    I/O   ---CACHE HIT RATE--   ----------DASD I/O RATE----------    ASYNC   TOTAL    READ   WRITE      %
 SERIAL      NUM  POOL   I/O   RATE    READ    DFW    CFW   STAGE  DFWBP    ICL    BYP  OTHER    RATE     H/R     H/R     H/R   READ
0*ALL                  100.0  502.2   282.1  126.4    0.0    93.7    0.0    0.0    0.0    0.0    98.8   0.813   0.751   1.000   74.8
 *CACHE-OFF              0.0    0.0
 *CACHE                100.0  502.2   282.1  126.4    0.0    93.7    0.0    0.0    0.0    0.0    98.8   0.813   0.751   1.000   74.8
 PR8F00     8F00  001F  22.3  112.2    65.1   28.0    0.0    19.1    0.0    0.0    0.0    0.0    19.1   0.830   0.773   1.000   75.1
 PR8F01     8F01  001F  11.3   56.6    31.9   14.2    0.0    10.5    0.0    0.0    0.0    0.0    10.6   0.815   0.753   1.000   75.0
 PR8F02     8F02  001F  11.0   55.2    30.3   14.3    0.0    10.5    0.0    0.0    0.0    0.0    10.4   0.810   0.743   1.000   74.0
 PR8F03     8F03  001F  11.2   56.2    31.8   14.0    0.0    10.4    0.0    0.0    0.0    0.0    10.6   0.815   0.754   1.000   75.1
 PR8F04     8F04  001F   3.6   18.2    10.5    4.6    0.0     3.1    0.0    0.0    0.0    0.0     3.9   0.828   0.770   1.000   74.7
 PR8F05     8F05  001F   3.5   17.8     9.5    4.5    0.0     3.8    0.0    0.0    0.0    0.0     3.9   0.789   0.717   1.000   74.6
 PR8F06     8F06  001F   3.6   18.1     9.8    4.5    0.0     3.7    0.0    0.0    0.0    0.0     4.1   0.793   0.723   1.000   74.9
 PR8F07     8F07  001F   3.9   19.8    11.1    5.0    0.0     3.7    0.0    0.0    0.0    0.0     4.3   0.812   0.749   1.000   75.0
 PR8F08     8F08  001F   3.5   17.7    10.2    4.2    0.0     3.3    0.0    0.0    0.0    0.0     3.6   0.814   0.756   1.000   76.2
 PR8F09     8F09  001F   3.7   18.8    10.4    4.7    0.0     3.6    0.0    0.0    0.0    0.0     4.1   0.806   0.741   1.000   74.9
```


RMF Link Statistics

```
                                                                                                                      150.4
1                                                E S S   L I N K   S T A T I S T I C S

            z/OS V1R6                 SYSTEM ID GDP2            START 06/07/2005-21.03.00  INTERVAL 000.00.59
                                      RPT VERSION V1R5 RMF      END   06/07/2005-21.04.00  CYCLE 1.000 SECONDS
OSERIAL NUMBER  0000020331     TYPE-MODEL  002107-922     CDATE  06/07/2005     CTIME  21.03.01     CINT  00.01.00
0------ADAPTER------    --LINK TYPE--      BYTES       BYTES          OPERATIONS       RESP TIME          I/O
 SAID  TYPE                               /SEC        /OPERATION     /SEC             /OPERATION         INTENSITY
00401  FIBRE 2Gb        NO DATA TO REPORT OR ZERO
00530  FIBRE 2Gb        ECKD READ             4.1M        4.1K          1436.4             0.0              62.4
                        ECKD WRITE            2.1M        4.1K           504.1             0.2              88.0
                                                                                                          ------
                                                                                                           150.4
00531  FIBRE 2Gb        NO DATA TO REPORT OR ZERO
00600  FIBRE 2Gb        ECKD READ             6.2M        4.1K          1503.3             0.0              62.7
                        ECKD WRITE            2.0M        4.1K           498.0             0.2              86.9
                                                                                                          ------
                                                                                                           149.6
00601  FIBRE 2Gb        NO DATA TO REPORT OR ZERO
00730  FIBRE 2Gb        ECKD READ             6.2M        4.1K          1505.8             0.0              62.7
                        ECKD WRITE            2.0M        4.1K           498.6             0.2              86.9
                                                                                                          ------
                                                                                                           149.6
00731  FIBRE 2Gb        NO DATA TO REPORT OR ZERO
```

RMF Extent Pool Statistics

```
1                                E S S   E X T E N T   P O O L   S T A T I S T I C S

            z/OS V1R6                SYSTEM ID GDP2            START 06/07/2005-21.03.00  INTERVAL 000.00.59
                                     RPT VERSION V1R5 RMF      END   06/07/2005-21.04.00  CYCLE 1.000 SECONDS
0SERIAL NUMBER  0000020331     TYPE-MODEL  002107-922     CDATE  06/07/2005     CTIME  21.03.01     CINT  00.01.00
0--EXTENT POOL--          ------- REAL -------
  ID    TYPE              CAPACITY      EXTENTS
00000  CKD 1Gb                382          434
 0001  CKD 1Gb                382          434
 0002  CKD 1Gb                382          434
 0003  CKD 1Gb                382          434
 0004  CKD 1Gb                446          507
 0005  CKD 1Gb                446          507
 0006  CKD 1Gb                446          507
 0007  CKD 1Gb                446          507
 0008  CKD 1Gb                382          434
 0009  CKD 1Gb                382          434
 000A  CKD 1Gb                382          434
 000B  CKD 1Gb                382          434
 000C  CKD 1Gb                446          507
 000D  CKD 1Gb                446          507
 000E  CKD 1Gb                446          507
```


DS8000 Rank Statistics

```
1                                E S S   R A N K   S T A T I S T I C S                                                PAG

            z/OS V1R6                SYSTEM ID GDP2            START 06/07/2005-21.03.00  INTERVAL 000.00.59
                                     RPT VERSION V1R5 RMF      END   06/07/2005-21.04.00  CYCLE 1.000 SECONDS
0SERIAL NUMBER  0000020331     TYPE-MODEL  002107-922     CDATE  06/07/2005     CTIME  21.03.01     CINT  00.01.00
0                        ------ READ OPERATIONS -------      ------ WRITE OPERATIONS ------
 --EXTENT POOL--              OPS    BYTES    BYTES   RTIME       OPS    BYTES    BYTES   RTIME     -ARRAY--  MIN     RANK  RAID
   ID    TYPE       RRID     /SEC     /OP      /SEC    /OP       /SEC     /OP      /SEC    /OP     NUM  WDTH  RPM     CAP  TYPE
00000  CKD 1Gb    0000      94.2    53.3K     5.0M     5.5       96.3     5.0K   480.6K     9.5      1     6  N/A     438G  RAID 5
00001  CKD 1Gb    0001      96.1    53.1K     5.1M     6.2       95.8     5.0K   478.4K    11.2      1     6  N/A     438G  RAID 5
00002  CKD 1Gb    0002      95.3    53.3K     5.1M     6.4       95.1     4.9K   469.7K    11.6      1     6  N/A     438G  RAID 5
00003  CKD 1Gb    0003      93.7    53.0K     5.0M     5.5       95.6     4.9K   467.5K     9.6      1     6  N/A     438G  RAID 5
00004  CKD 1Gb    0004      93.3    53.0K     4.9M     6.0       95.3     4.9K   469.7K    10.3      1     7  N/A     511G  RAID 5
00005  CKD 1Gb    0005      93.6    53.2K     5.0M     6.1       95.8     4.9K   471.9K    11.5      1     7  N/A     511G  RAID 5
00006  CKD 1Gb    0006      91.9    53.3K     4.9M     5.9       95.7     4.9K   467.5K     9.8      1     7  N/A     511G  RAID 5
00007  CKD 1Gb    0007      95.0    53.5K     5.1M     5.9       95.0     4.9K   469.7K    10.8      1     7  N/A     511G  RAID 5
00008  CKD 1Gb    0008      94.3    53.1K     5.0M     5.5       96.2     5.0K   478.4K     9.2      1     6  N/A     438G  RAID 5
00009  CKD 1Gb    0009      95.0    53.1K     5.0M     5.6       94.7     4.9K   467.5K     9.3      1     6  N/A     438G  RAID 5
0000A  CKD 1Gb    000A      92.7    53.1K     4.9M     5.5       94.1     4.9K   465.3K     8.9      1     6  N/A     438G  RAID 5
0000B  CKD 1Gb    000B      94.0    53.2K     5.0M     5.6       95.2     5.0K   471.9K     9.7      1     6  N/A     438G  RAID 5
```

DS8000 Benefits

- **More capacity**
- **More host ports**
- **More performance**

- **More LCUs**
- **Flexible LCU alignment with ranks**
  - Ability to maintain small volumes on ranks of large DDMs
  - Ability to create application-based LCUs and consistency groups

- **Larger volumes**
- **Dynamic volume deletion**
- **User-specified volume IDs**

- **Enhanced command line interface**
- **Copy Services compatibility with ESS and DS6000**
- **Dual Storage Images in a single physical storage unit**

DS8000 References

IBM TotalStorage DS8000 Series: Concepts and Architecture	SG24-6452
IBM TotalStorage DS8000 Series: Implementation	SG24-6786
IBM TotalStorage DS8000 Series: Copy Services with IBM eServer zSeries	SG24-6787
IBM TotalStorage DS8000 Series: Copy Services in Open Environments	SG24-6788
IBM TotalStorage DS8000 Introduction and Planning Guide	GC35-0495
IBM TotalStorage DS8000 Installation Guide	SY27-7641
IBM TotalStorage DS8000 User's Guide	SC26-7623
IBM TotalStorage DS8000 Host Systems Attachment Guide	SC26-7628
Multipath Subsystem Device Driver User's Guide	SC30-4096

zSeries – HCD Updates with z/VM

- **Support of ATTR=1750 for a VM FBASCSI device:**
  - z/VM 5.1 supports the emulation of an FBA disk on SCSI disks defined on the DS6000 D/T1750.
  - When defining the ATTR=1750 parameter on an FBASCSI device type in HCD, for each given path, it has to be specified whether the path is preferred or not using the new PREFPATH parameter
  - APARs OA09114 and OA07873 provide support
- **Support of ATTR=ESS for a VM FBASCSI device:**
  - z/VM 5.1 supports the emulation of an FBA disk on SCSI disks defined on the DS8000 D/T2107.
  - ATTR=ESS parameter supports an FBASCSI device type for both 2105 and 2107
  - APAR OA06830 provides support

Trademarks

Disclaimer

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This information could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or programs(s) at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectually property rights, may be used instead. It is the user's responsibility to evaluate and verify the operation of any on-IBM product, program or service.

THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted according to the terms and conditions of the agreements (e.g., IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. IBM is not responsible for the performance or interoperability of any non-IBM products discussed herein.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.