

Technology · Connections · Results

z/OS Basics: Migrating from HFS to zFS and things to watch out for

Jim Showalter IBM

March 3, 2011 Session 9037



Agenda





- Why you should convert from HFS to zFS
- How to convert HFS to zFS
- Things that are different
- Things to watch out for





Why you should convert from HFS to zFS



- zFS is the base for future z/OS UNIX file system development
 - HFS development is stabilized
 - HFS is still supported for service but support may be removed sometime in the future
- zFS generally performs better than HFS
- Recent significant performance improvements for zFS especially in the shared file system environment (z/OS V1R11 and z/OS V1R13)



How to convert HFS to zFS



- Use the HFS to zFS migration tool (BPXWH2Z)
 - It is an ISPF based tool executed from ISPF 6 (Command)
 - Documented in z/OS UNIX System Services Planning
- Use the pax command
 - pax –rwvCMX –p eW /etc/fromhfsmnpt /etc/tozfsmnpt
 - Documented in z/OS UNIX System Services Command Reference
- Use the copytree command
 - /samples/copytree /etc/fromhfsmnpt /etc/tozfsmnpt
 - Documented in z/OS UNIX System Services Command Reference



An Example of BPXWH2Z



- 4 filesystems
 - USSZFS.SHARE.F1.HFS
 - USSZFS.SHARE.F2.HFS
 - USSZFS.SHARE.F3.HFS
 - USSZFS.SHARE.FS.HFS
- USSZFS.SHARE.F1.HFS is mounted on /share and is a 14G Multi-Volume file system spanning 2 mod 9s
- USSZFS.SHARE.F2.HFS is mounted on /share/erahs
- USSZFS.SHARE.F3.HFS is not mounted
- USSZFS.SHARE.FS.HFS is mounted on /share/nomig and we will not be migrating it.



Start the tool by issuing BPXWH2Z –cv from TSO

PKSTHUB1.ws	Antiona 160 days 11-la				
		DATA SET	SPECIFICATIO	N	
Command ===	>				
Use the HEL	P command for	usage inform	ation on this	tool.	
Enter the n	ame of the HFS	6 file system	n to migrate t	o a zFS file s	system.
Note that a	ı data set patt	ern can be u	used.		
If the file	system is not	: currently m	nounted it wil	l be mounted o	on a
temporary d	lirectory durir	ng the migrat	ion.		
File system	name: <u>usszfs.</u>	share.*_			
		—			
F1=HELP	F2=SPLIT	F3=END	F4=RETURN	F5=RFIND	F6=RCHANGE
F7=UP	F8=DOWN	F9=SWAP	F10=LEFT	F11=RIGHT	F12=RETRIEVE
				,	11/034
Connected to remote server/ho	ost pksthub 1.pok.ibm.com using lu/	poor TCP00134 and port 23			

PKSTHUB1.ws	_ D 🗙
File Edit View Communication Actions Window Help	
CLASS AND VOLUME DEFAULTS	
Command ===>	

The volume or SMS classes for zFS allocations will default to the same names as the current HFS allocation. You can change these individually for each new allocation. If you want the default for each new allocation to be set to specific values other than that of the current HFS allocation, enter those values here and Press Enter.

Default	volume:	
Default	data class:	<u>zfsext</u>
Default	storage class:	
Default	management class :	

F1=HELP	F2=SPLIT	F3=END	F4=RETURN	F5=RFIND	F6=RCHANGE
F7=UP	F8=DOWN	F9=SWAP	F10=LEFT	F11=RIGHT	F12=RETRIEVE

<u>A</u>	a	11/035
ታ የ	Connected to remote server/host pksthub 1.pok.ibm.com using lu/pool TCP00134 and port 23	
		2011

PKSTHUB1.ws	
File Edit View Communication Actions Window Help	
-getting ds list for USSZFS.SHARE.*	
-getting ds info for USSZFS.SHARE.FS.HFS	
-getting ds info for USSZFS.SHARE.F1.HFS	
-getting ds info for USSZFS.SHARE.F2.HFS	
-getting ds info for USSZFS.SHARE.F3.HFS	
-getting ds info for USSZFS.SHARE.ZFS	
Skipping non-HFS data set USSZFS.SHARE.ZFS	
-getting ds info for USSZFS.SHARE.ZFS.DATA	
Skipping non-HFS data set USSZFS.SHARE.ZFS.DATA	

а

10/006

Connected to remote server/host pksthub1.pok.ibm.com using lu/pool TCP00134 and port 23

D PKSTHUB1.ws		
File Edit View Communication Actions Window Help		
• • • • • • • • • • • • • • • • • • •	Ø	
DAT Command ===>	TA SET LIST	Row 1 to 2 of 4
Use the HELP command for full Select items with D to delete Select items with A to alter a	usage information or items from this migr allocation parameters	n this tool ration list s for the items
Enter command FG or BG to beginstrate D US	in migration in fore SSZFS.SHARE.FS.HFS	ground or UNIX background
Save HFS as: USSZFS.SHARE.	FS.HFS.SAV	Utilized : 43%
Initial zFS: USSZFS.SHARE.	FS.HFS.TMP	Allocated: N
Final zFS: USSZFS.SHARE.	FS.HFS	
HFS space Primary : 50	Secondary: 5	Units: CYL
zFS space Primary : 50	Secondary: 5	Units: CYL
Dataclas : ZFSEXT	Mgmtclas :	Storclas:
MOUNTED Volume : 9SX902	Vol count: 1	
Us	SSZFS.SHARE.F1.HFS	
Save HFS as: USSZFS.SHARE.	.F1.HFS.SAV	Utilized : 87%
Initial zFS: USSZFS.SHARE.	.F1.HFS.TMP	Allocated: N
Final zFS: USSZFS.SHARE.	.F1.HFS	
HFS space Primary : 7000	Secondary: 1000	Units: CYL
zFS space Primary : 7700	Secondary: 1000	Units: CYL
F1=HELP F2=SPLIT F3=EN	ND F4=RETURN	F5=RFIND F6=RCHANGE
F7=UP F8=DOWN F9=SV	JAP F10=LEFT	F11=RIGHT F12=RETRIEVE
F1=HELP F2=SPLIT F3=EN F7=UP F8=DOWN F9=SW	ND F4=RETURN NAP F10=LEFT	F5=RFINDF6=RCHANGEF11=RIGHTF12=RETRIEVE

DKSTHUB1 ws							
Edit View Communication Actions	Window Help						
	a k k k	🖻 🔿 🔗					
						4	6.0
0)		UHIH :	SEI LISI -		R0	οω 1 το 2	отз
command ===> _							
Use the USLD	command fo	- full us		+:	this tool		
Soloct itoms	uith D to	r lull us delete it	age informa ome from th	ic migno	tion list		
Select Items	with D to	alter alle	ems from th	is migra	for the iter		
Entor common	d EC or PC	to bogin r	migration i	ameters n foroar	ound on UNIX	lis V backarou	nd
		119970 1	ES SHODE E1	HES		~ backgrou	
- Save HES as	• 11997F9	SHARE E1	HES SAV	. 111 3		Itilized ·	87%
Initial ZES	· 11997E9	SHARE F1	HES TMP		, I	Allocated:	N
Final ZES	· USSZES	SHARE F1	HES		•	nitocated.	
HFS space P	rimaru : 7	000	Secondaru:	1000	Units .	.: CYL	
zFS space P	rimaru : 7	700	Secondaru:	1000	Units .	.: CYL	
Da	ataclas : Z	FSEXT	Mamtclas :		Storclas	SMSZFS	
MOUNTED V	olume :		Vol count:	2			
		USSZR	FS.SHARE.F2	.HFS			
Save HFS as	: USSZFS	.SHARE.F2	.HFS.SAV		l	Utilized :	18%
Initial zFS	: USSZFS	.SHARE.F2	.HFS.TMP		í	Allocated:	Ν
Final zFS .	: USSZFS	.SHARE.F2	.HFS				
HFS space P	rimary : 4	2	Secondary:	5	Units .	.: CYL	
zFS space Pr	rimary : 4	2	Secondary:	5	Units .	.: CYL	
-				TUDN			CE.
F1=HELP	F2=SPLIT	F3=END	F4=RE	TURN	FD=RFIND	F6=RCHHN	GL

MA	а	A	02/015
J (onnected to remote server/host pksthub1.pok.ibm.com using	Ju/pool TCP00134 and port 23	
	10		2011

PKSTHUB1.ws							_ 7
File Edit View Communication Ac	tions Window Help						
9 BB 🛃 💻 🖩	🛋 🐚 🛃 👪	🗎 🔮 🤣					
		DATA S	ET LIST -		F	Row 3 to 3	of 3
Command ===>	_						
Use the HE	∟P command for	r full usa	qe informa	tion on	this tool		
			~				
Select ite	ns with D to a	delete ite	ms from th	is miar	ation list		
Select ite Select ite	ns with D to a ns with A to a	delete ite alter allo	ms from th cation par	is migr ameters	ation list for the ite	>ms	
Select ite Select ite Enter comm	ns with D to a ns with A to a and EG or BG of	delete ite alter allo to begin m	ms from th cation par igration i	is migr ameters p foreg	ation list for the ite	ems IX backgrou	und
Select ite Select ite Enter comma	ns with D to a ns with A to a and FG or BG	delete ite alter allo to begin m	ms from th cation par igration i s suppe ea	is migr ameters n foreg	ation list for the ite round or UN]	ems [X backgrou	und
Select iter Select iter Enter comm	ns with D to with A to wit	delete ite alter allo to begin m USSZF	ms from th cation par igration i S.SHARE.F3	is migr ameters n foreg .HFS	ation list for the ite round or UN] 	ems [X backgrou 	ind
Select iter Select iter Enter comm Save HFS as	ns with D to a ns with A to a and FG or BG s: USSZFS	delete iter alter allo to begin m USSZF .SHARE.F3.	ms from th cation par igration i S.SHARE.F3 HFS.SAV	is migr ameters n foreg .HFS	ation list for the ite round or UN] 	ems [X backgrou Utilized :	und 33%
Select iter Select iter Enter comm Save HFS as Initial zFS	ms with D to ms with A to and FG or BG s: USSZFS S: USSZFS	delete iter alter allo to begin m USSZF .SHARE.F3. .SHARE.F3.	ms from th cation par igration i S.SHARE.F3 HFS.SAV HFS.TMP	is migr ameters n foreg .HFS	ation list for the ite round or UN] 	ems [X backgrou Utilized : Allocated:	und 33% N
Select iter Select iter Enter comm Save HFS as Initial zFS	ms with D to ms with A to and FG or BG s: USSZFS S: USSZFS : USSZFS	delete ite alter allo to begin m USSZF .SHARE.F3. .SHARE.F3.	ms from th cation par igration i S.SHARE.F3 HFS.SAV HFS.TMP HFS	is migr ameters n foreg .HFS	ation list for the ite round or UN] 	ems [X backgrou Utilized : Allocated:	und 33% N
Select iter Select iter Enter comm Save HFS ar Initial zFS Final zFS HFS space	ms with D to ms with A to and FG or BG 5: USSZFS 5: USSZFS : USSZFS Primary : 42	delete iter alter allo to begin m USSZF .SHARE.F3. .SHARE.F3. 2	ms from th cation par igration i S.SHARE.F3 HFS.SAV HFS.TMP HFS Secondary:	is migr ameters n foreg .HFS 5	ation list for the ite round or UN] Units .	ems [X backgrou Utilized : Allocated: : CYL	und 33% N
Select iter Select iter Enter comm Save HFS ar Initial zFS Final zFS HFS space zFS space	ms with D to ms with A to and FG or BG : USSZFS S: USSZFS : USSZFS Primary : 42 Primary : 42	delete iter alter allo to begin m USSZF .SHARE.F3. .SHARE.F3. 2	ms from th cation par igration i S.SHARE.F3 HFS.SAV HFS.TMP HFS Secondary: Secondary:	is migr ameters n foreg .HFS 5 5	ation list for the ite round or UN] Units . Units .	ems [X backgrou Utilized : Allocated: : CYL	und 33% N
Select iter Select iter Enter comm Save HFS ar Initial zF Final zFS HFS space zFS space	ms with D to ms with A to and FG or BG s: USSZFS S: USSZFS : USSZFS Primary : 42 Primary : 42 Dataclas :	delete iter alter allo to begin m USSZF .SHARE.F3. .SHARE.F3. 2 2	ms from th cation par igration i S.SHARE.F3 HFS.SAV HFS.TMP HFS Secondary: Secondary: Mgmtclas :	is migr ameters n foreg .HFS 5 5	ation list for the ite round or UN] Units . Units . Storcla	ems [X backgrou Utilized : Allocated: : CYL : CYL as: SMSZFS	und 33% N
Select iter Select iter Enter comm Save HFS as Initial zFS Final zFS HFS space zFS space	ms with D to ms with A to and FG or BG : USSZFS S: USSZFS : USSZFS Primary : 42 Primary : 42 Dataclas : Volume : Of	delete iter alter allo to begin m USSZF .SHARE.F3. .SHARE.F3. 2 2 E1523	ms from th cation par igration i S.SHARE.F3 HFS.SAV HFS.TMP HFS Secondary: Secondary: Ngmtclas : Vol count:	is migr ameters n foreg .HFS 5 5	ation list for the ite round or UN] Units . Storcla	ems [X backgrou Utilized : Allocated: : CYL : CYL as: SMSZFS	und 33% N

F1=HELP	F2=SPLIT	F3=END	F4=RETURN	F5=RFIND	F6=RCHANGE
F7=UP	F8=DOWN	F9=SWAP	F10=LEFT	F11=RIGHT	F12=RETRIEVE
м <mark>А</mark> b					02/015
Connected to remote server/h	ost pksthub1.pok.ibm.com using lu	/pool TCP00182 and port 23			
44				,	
11					2011

Ele Edit View Communication Actions Window Help				
	CHANGE ALLOCA	TION ATTRIBUT	ES	
Command ===> _				
Use the HELP command for Change the allocation at If the file system is al name of the data set as	• usage inform tributes for ready allocat the Initial z	ation on this this new zFS ed enter Y for FS name. The a	tool. file system. r preallocate attributes wi	d and the ll not be used.
File system name Save HFS as <u>l</u> Initial zFS name <u>l</u>	JSSZFS.SHARE.F JSSZFS.SHARE.F JSSZFS.SHARE.F	1.HFS 1.HFS.SAV 1.HFS.TMP		
Final ZFS name Preallocated zFS Primary allocation Secondary allocation	<u>1552F5.SHARE.F</u> <u>1</u> <u>700</u> <u>.000</u> <u>.000</u> <u>.200</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u> <u>.000</u>	L or TRK)	F5=RFIND	E6=RCHANGE
F7=UP F8=D0WN	F9=SWAP	F10=LEFT	F11=RIGHT	F12=RETRIEVE
· · · · · · · · ·				

M <u>A</u> aa	02/015
Connected to remote server/host pksthub 1.pok.ibm.com using lu/pool TCP00134 and port 23	
	2011

PL PRSTHUBT.WS	
-getting ds info for USSZFS.SHARE.F1.HFS.SAV	
-getting ds info for USSZFS.SHARE.F1.HFS.TMP	
-getting ds info for USSZFS.SHARE.F2.HFS.SAV	
-getting ds info for USSZFS.SHARE.F2.HFS.TMP	
-getting ds info for USSZFS.SHARE.F3.HFS.SAV	
-getting ds info for USSZFS.SHARE.F3.HFS.TMP	
Migrating USSZFS.SHARE.F1.HFS 16:16:17	
creating zFS USSZFS.SHARE.F1.HFS.TMP	
copying USSZFS.SHARE.F1.HFS to USSZFS.SHARE.F1.HFS.TMP Blocks to copy: 22	03835
- cd /share; /bin/pax -rw -peW -XCM . /tmp/bpxwh2z.CHAIGHT.15:39:39.zfs	
ENTRY USSZFS.SHARE.F1.HFS.TMP ALTERED	
ENTRY USSZFS.SHARE.F1.HFS.TMP.DATA ALTERED	
mount /share USSZFS.SHARE.F1.HFS ZFS 0	
mount /share/nomig USSZFS.SHARE.FS.HFS HFS HFS	0
mount /share/erahs USSZFS.SHARE.F2.HFS HFS HFS	0
Migration complete for USSZFS.SHARE.F1.HFS 16:58:57	
Migrating USSZFS.SHARE.F2.HFS 16:58:57	
creating zFS USSZFS.SHARE.F2.HFS.TMP	
copying USSZFS.SHARE.F2.HFS to USSZFS.SHARE.F2.HFS.TMP Blocks to copy: 39	61
- cd /share/erahs; /bin/pax -rw -peW -XCM . /tmp/bpxwh2z.CHAIGHT.15:39:39	.zfs
ENTRY USSZFS.SHARE.F2.HFS.TMP ALTERED	
ENTRY USSZFS.SHARE.F2.HFS.TMP.DATA ALTERED	
mount /share/erahs USSZFS.SHARE.F2.HFS ZFS 0	

мA	a	24/006
പ	Connected to remote server/host pksthub 1.pok.ibm.com using lu/pool TCP00134 and port 23	
	13	2011

PKSTHUB1.ws
ile Edit View Communication Actions Window Help
Migration complete for USSZFS.SHARE.F2.HFS 17:03:06 Migrating USSZFS.SHARE.F3.HFS 17:03:06 creating zFS USSZFS.SHARE.F3.HFS.TMP
copying USSZFS.SHARE.F3.HFS to USSZFS.SHARE.F3.HFS.TMP Blocks to copy: 6204
- cd /tmp/bpxwh2z.CHAIGHT.15:39:39.hfs; /bin/pax -rw -peW -XCM . /tmp/bpxwh2z.C
ENTRY USSZFS.SHARE.F3.HFS.TMP ALTERED
ENTRY USSZFS.SHARE.F3.HFS.TMP.DATA ALTERED
Migration complete for USSZFS.SHARE.F3.HFS 17:04:49 ***



BPXWH2Z – Special Considerations

- Size of target zFS file system
- Size of target zFS file system log
- System Managed? greater than 4GB



Things that are different



SHARE Technology · Connections · Results

- HFS vs zFS (VSAM)
- zFS is a logging file system
- Growing a file system





HFS vs zFS (VSAM)

- HFS data set
 - DD with DSNTYPE=HFS

//USERIDA	JOB
//STEP1	EXEC PGM=IEFBR14
//HFS1 DD	DSN=OMVS.HFS.HOME,
//	SPACE=(CYL,(40,1,1)),
//	DSNTYPE=HFS,
//	DISP=(NEW,CATLG,DELETE),
//	STORCLAS=STANDARD

```
//USERIDA JOB
                                //STEP1
                                           EXEC PGM=IDCAMS
• zFS data set
                                //SYSPRINT DD SYSOUT=H
                                //SYSIN
                                           DD *

    define VSAM Linear

                                 DEFINE CLUSTER (NAME(OMVS.ZFS1) -
      Data Set (LDS)
                                        VOLUMES(PRV000) -

    format with

                                        LINEAR -
                                        CYL(40 1) -
      IOEAGFMT
                                        SHAREOPTIONS(3))

    (Can also use zfsadm)

                                /*
      commands to define
                                //STEP2
                                           EXEC PGM=IOEAGFMT, REGION=0M,
                                // PARM=('-aggregate OMVS.ZFS1 -compat')
      and format or pfsctl
                                //SYSPRINT DD SYSOUT=H
      APIs)
                                                                   in Anaheim
```



zFS is a logging file system



- zFS is a logging file system and has improved recovery from system failures
 - Metadata updates are transactional and are logged to maintain the structural integrity of the file system
- zFS has a utility called IOEAGSLV (Salvager) that can verify and optionally attempt to repair a zFS file system in the unlikely event that it becomes corrupted



Growing a file system



Dynamic grow

- HFS will grow when a write occurs and out of space if
 - A secondary allocation size was specified when created, and
 - there is space available on the volume(s)
- zFS will grow when a write occurs and out of space if
 - A secondary allocation size was specified when created, and
 - there is space available on the volume(s), and
 - aggrgrow=on is specified in the zFS IOEFSPRM configuration file or AGGRGROW is specified on the MOUNT PARM NOTE: zFS R13 is changing the default from aggrgrow=off to aggrgrow=on so zFS will act the same as HFS by default



Growing a file system ...



• Explicit grow with a command

- HFS can be explicitly grown with the z/OS UNIX command confights –x 10c pathname
- zFS can be explicitly grown with the z/OS UNIX zFS command
 zfsadm grow aggregate_name new_total_k_bytes
- Residing on the EAS (extended addressing space) portion of an EAV (extended address volume)
 - HFS is not eligible
 - zFS is eligible as of z/OS V1R10



Things to watch out for



Technology • Connections • Results

- Large directories
- DASD space usage
- Backup and Quiesce
- Plan to remove function
- Mount Parms
- Sysplex sharing
- Publications
- Backup slides





Large Directories

- zFS has a performance problem with large directories
 - As you approach 100,000 entries in a zFS directory, performance begins to suffer
 - If you can,
 - spread out entries among multiple directories, or
 - try to remove older files to keep directory from getting too large, or
 - use HFS for this directory
 - There is some guidance on this in the z/OS Distributed File Service zSeries File System Administration book (SC24-5989) in Chapter 4, "Minimum and maximum file system sizes".



DASD space usage



- HFS uses 4K blocks to store file data
- zFS uses 8K blocks but can store a small file (<53 bytes) in the inode
- zFS R13 does not use 1K fragments any longer
 - Simplifies zFS code making it less error prone
 - Necessary to support zFS R13 Direct I/O support
- This means that, in some cases, zFS R13 will use more DASD space than zFS R11
 - The worst case is files that are less than or equal to 1K (but larger than 52 bytes)
 - 1000 1K files could take (a maximum of) 10 cylinders more space in zFS R13 than zFS R11
- See z/OS Distributed File Service zSeries File System (SC24-5989), Chapter 4, zFS disk space allocation for more information



DASD space usage ...



- Fragmented files caused confusion about free space in zFS
 - df can report, for example, 20K of free space
 - but, if there are no free 8K blocks (that is, there are only free fragments), then you cannot, for example, create a 14K file
 - zfsadm aggrinfo aggregate_name –long shows detailed information including the number of free 8K blocks
- See z/OS Distributed File Service zSeries File System (SC24-5989), Chapter 4, zFS disk space allocation for more information



Space in a file system ...



SHARE Technology · Connections · Results

- An HFS that is multi-volume must be SMS managed (and cataloged)
- A zFS that is > 4GB must be extended addressability in the data class definition and therefore SMS managed (A VSAM LDS is always cataloged)





Backup and Quiesce

- DFSMSdss automatically quiesces a mounted file system on backup to ensure data integrity
 - For HFS, quiesce is a BPX1QSE call
 - df /hfsmntpoint displays Status of Quiesced
 - D OMVS, F, N=OMVS. HFS. FS1 displays Status of QUIESCED
 - D OMVS, F, E considers the HFS to be in an exception state
 - Message BPXF083I THE FOLLOWING FILE SYSTEM HAS BEEN QUIESCED FOR MORE THAN 10 MINUTES: OMVS.HFS.FS1 QUIESCING SYSTEM=DCEIMGVM JOB=SUIMGVM PID=67174418 LATCH=44.
 - For zFS, quiesce is a BPX1PCT call
 - Neither df, nor D OMVS,F show the ZFS to be quiesced
 - zfsadm aggrinfo omvs.zfs.fs1 displays status of quiesced
 - zfsadm lsaggr displays status of quiesce
 - Message IOEZ00581E There are quiesced zFS aggregates. After about 30 seconds



Plan to remove function



- In February 2011, IBM announced
 - z/OS V1.13 is planned to be the last release to support multifile system zSeries File System (zFS) aggregates, including zFS clones. Support for the zfsadm clone command and mount support for zFS file system data sets containing a cloned (.bak) file system will be removed. IBM recommends that you use copy functions such as pax and DFSMSdss to back up z/OS UNIX file systems to separate file systems. Support for zFS compatibility mode aggregates will remain.



MOUNT/automount Parms



- The MOUNT PARMs for HFS and zFS are different (the other options are the same – MOUNTPOINT, MODE, etc.)
 - HFS MOUNT PARMs (MOUNT TYPE(HFS))
 - PARM('FSFULL(threshold,increment)')
 - PARM('NOSPARSE')
 - PARM('NOWRITEPROTECT')
 - PARM('SYNC(sec)')
 - PARM('SYNCRESERVE(nn)')
 - zFS MOUNT PARMs (MOUNT TYPE(ZFS))
 - PARM('AGGRFULL(threshold,increment)')
 - PARM('AGGRGROW')
 - PARM('NBS')
 - PARM('RW')
 - PARM('RWSHARE')



MOUNT/automount Parms...



- Generic file system TYPE on MOUNT
 - If you specify TYPE(HFS) and the data set is not HFS or is not found, it is treated as ZFS (and you get a zFS reason code)
 - If you specify TYPE(ZFS) and the data set is HFS, it is treated as HFS
 - In each of these cases where the TYPE did not match the actual data set type, THE MOUNT PARMs ARE DISCARDED

(we don't want the mount to fail due to invalid PARMs)

 Once you have fully migrated a file system from HFS to zFS, you should specify TYPE(ZFS) so that MOUNT PARMs are effective



MOUNT/automount ...



- zFS is a logging file system (metadata is logged, not file data)
- Maintains file system consistency log is replayed on next mount if file system was not cleanly unmounted – requires a R/W mount
- Problem scenario can occur
 - 1. R/O file system (for example, version root) needs to be updated
 - 2. Remount R/O to R/W
 - 3. Update file system
 - 4. Before remount back to R/O, system is re-IPLd
 - 5. Mount of R/O version root fails because log needs to be replayed
- Should always do MODIFY OMVS,SHUTDOWN before a planned system shut down
- With R9 APAR OA20615, zFS provides new IOEFSPRM option (romount_recovery=on)



Sysplex sharing

- SHARE Technology · Connections · Results
- Both HFS and zFS support read-write sharing from multiple systems in a shared file system environment (BPXPRMxx SYSPLEX(YES))
 - HFS read-write file systems are always non-sysplex aware (z/OS UNIX always uses function shipping to a single owning system)
 - zFS read-write file systems can be sysplex-aware or nonsysplex aware

For sysplex-aware read-write, z/OS UNIX sends requests to the local zFS and then

- R11 zFS uses caching to sometimes avoid sending a read request to the owning system
- R13 zFS can do direct I/O for reading and writing





Technology · Connections · Results

RE

Sysplex sharing ...

Non-sysplex aware read-write

R11 Sysplex-aware read-write



Publications



echnology • Connections • Result

- z/OS UNIX System Services Planning (GA22-7800) General Administration of z/OS UNIX file systems
- z/OS UNIX Command Reference (SA22-7802) confights command for HFS
- z/OS MVS System Messages Volume 9 (IGF-IWM) (SA22-7639) IGWxxxt messages for HFS
- z/OS UNIX System Services Messages and Codes (SA22-7807) z/OS UNIX return codes, z/OS UNIX reason codes, X'5Bxxrrrr' reason codes for HFS
- z/OS Distributed File Service zSeries File System Administration (SC24-5989) was refreshed in April 2010 zFS Concepts and zfsadm command for zFS
- z/OS Distributed File Services Messages and Codes (SC24-5917) IOEZxxxt messages and X'EFxxrrrr' reason codes for zFS
- z/OS Distributed File Service zSeries File System Implementation (SG24-6580)
 - Redbook available (updated February 2010 to include z/OS V1R11)
 - <u>http://www.redbooks.ibm.com/abstracts/sg246580.html?Open</u>
- z/OS Version 1 Release 8 Implementation (SG24-7265)
 - Redbook available (contains zFS updates for z/OS V1R8)
 - <u>http://www.redbooks.ibm.com/abstracts/sg247265.html?Open</u>
- z/OS DFSMSTM Access Method Services for Catalogs (SC26-7394) IDCAMS utility
- z/OS DFSMSTM Storage Administration Reference (SC26-7402) ADRDSSU utility for backup







Backup



File access

- File access commands and APIs for zFS are the same as HFS except for reason codes on failures
 - z/OS UNIX reason codes X'0000rrrr' to X'20FFrrrr' Documented in z/OS UNIX System Services Messages and Codes (SA22-7807) – these are common to HFS and zFS
 - HFS specific reason codes X'5Bxxrrr' Documented in z/OS UNIX System Services Messages and Codes (SA22-7807)
 - zFS specific reason codes X'EFxxrrr' Documented in z/OS Distributed File Service Messages and Codes (SC24-5917)
- The bpxmtext shell command can be used to display the meaning of zFS reason codes (as of z/OS V1R8) and z/OS UNIX reason codes



Directory access



- Directory access commands and APIs for zFS are the same as HFS except for a few non-obvious situations
 - HFS returns names in a directory in (some) alphabetical order (using opendir, readdir, closedir APIs)
 DO NOT BECOME DEPENDENT ON THIS ORDER
 - This is not POSIX behavior
 - It is not controlled by localization envars (LC_COLLATE)
 - Is returns sorted names but that is because –C is the default
 - The order that HFS returns names is <u>not</u> the same as **Is** (Is -C returns uppercase characters first; HFS returns uppercase characters last.)
 - zFS returns names unsorted as every other UNIX file system does





Technology · Connections · Result

Directory access ...

- zFS directories can be read as a file; HFS directories return 0 bytes (using open, read, close APIs)
 - You can see this by using the strings command against a directory





Directory access ...

- SHARE Technology · Connections · Results
- zFS is limited to 64K subdirectories per directory
- HFS reports that it is limited to 64K (but may support more?)
 DO NOT CREATE MORE THAN 64K SUBDIRECTORIES IN A DIRECTORY

