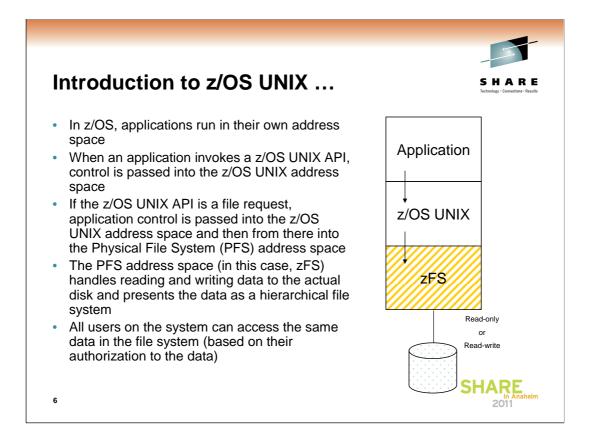


The application programming interface of z/OS UNIX supports C callable services (documented in the z/OS XL C/C++ Run-Time Library Reference) and Assembler callable services (documented in the z/OS UNIX System Services Programming: Assembler Callable Services Reference).

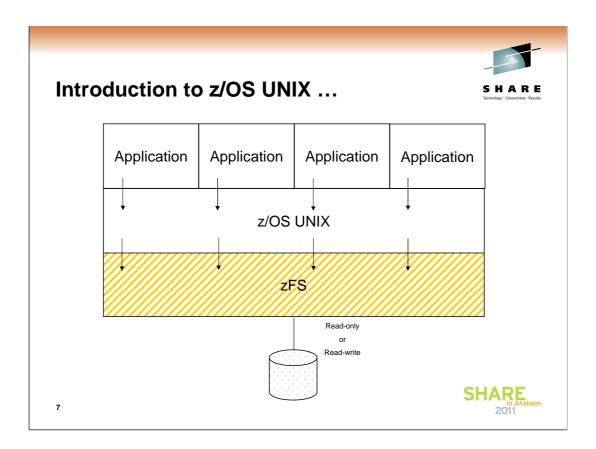
The interactive shell interface provides a UNIX command interface (commands are documented in the z/OS UNIX System Services Command Reference).

The TSO/E commands are also documented in the same book.

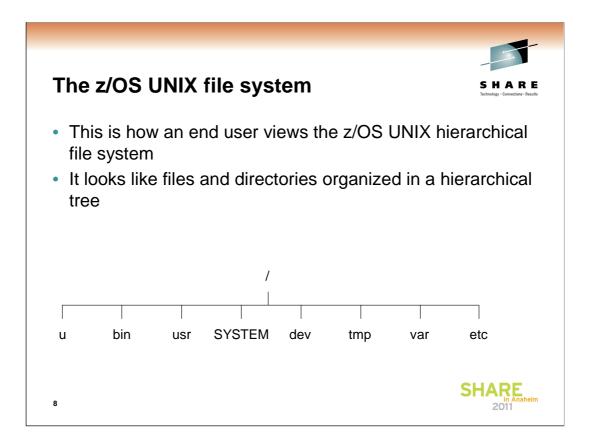
The ISPF shell is documented in the z/OS UNIX System Services User's Guide.

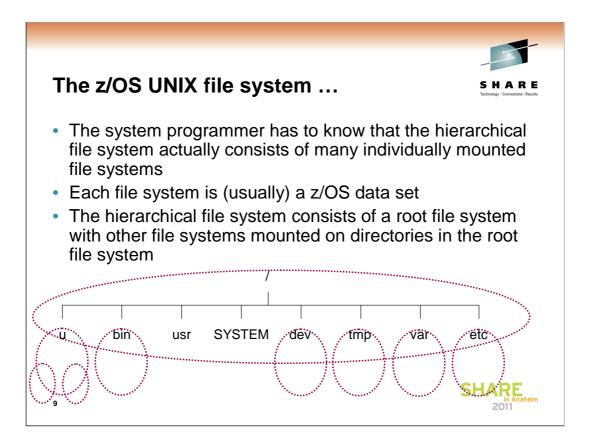


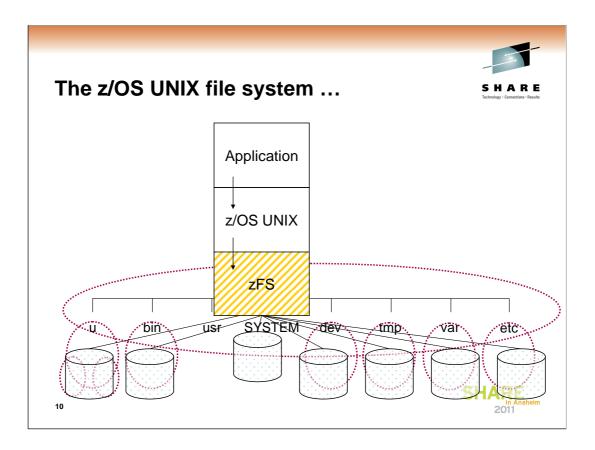
zFS runs in its own address space (sometimes referred to as a Colony Address Space). z/OS UNIX application requests program call into the z/OS UNIX address space and when z/OS UNIX determines that the request is for a zFS file system, z/OS UNIX program calls into the zFS address space. The zFS address space handles all I/O to the file system on DASD.

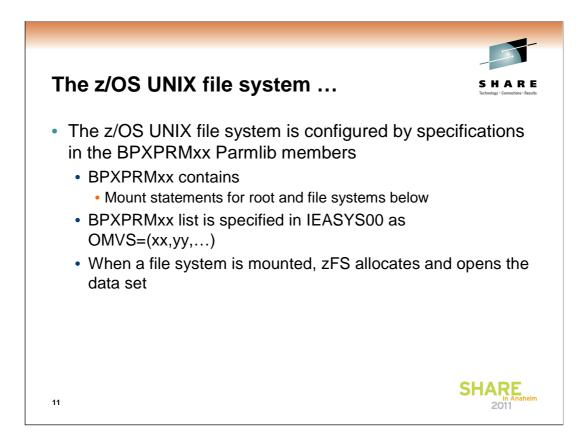


Of course, there can be multiple applications accessing z/OS UNIX and zFS at the same time, so this is a bit more accurate. For simplicity, I will show only a single application in the remaining pictures, but you should understand that there can be many applications.





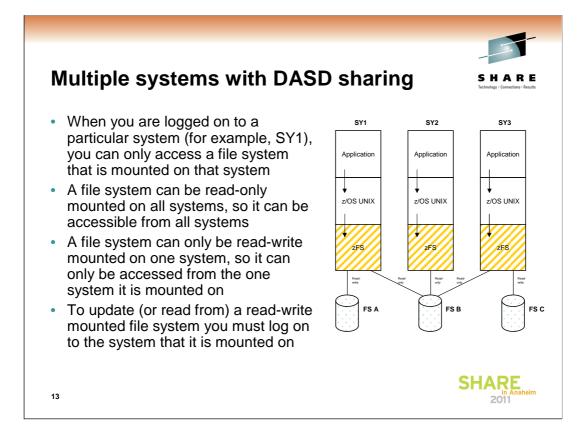




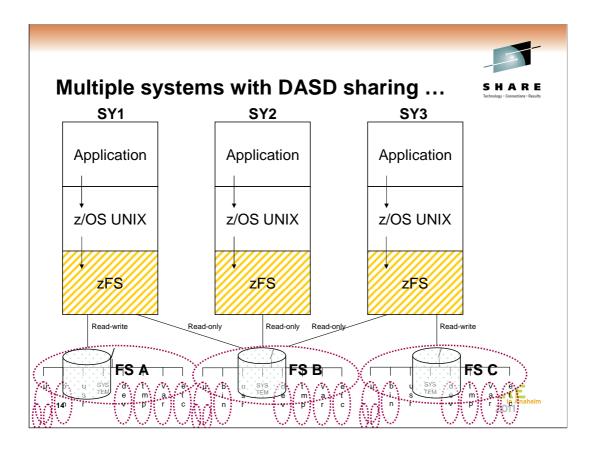
The logical parmlib concatenation is a set of up to ten partitioned data sets defined by PARMLIB statements in the LOADxx member of either SYSn.IPLPARM or SYS1.PARMLIB (quote from z/OS MVS Initialization and Tuning Reference).

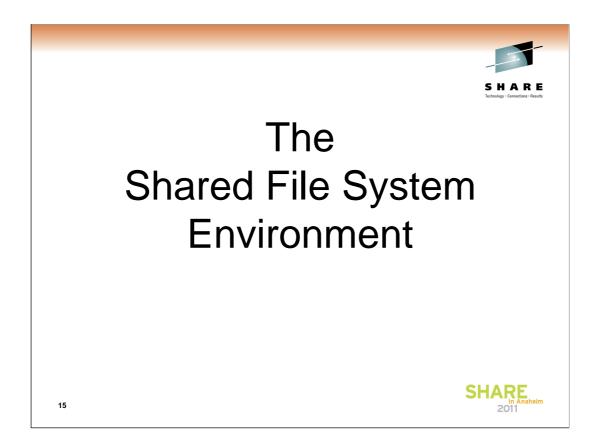
BPXPRMxx also contains the FILESYSTYPE statement for zFS and other PFSes – defines the TYPE (for example, ZFS) and the load module that contains the zFS code.

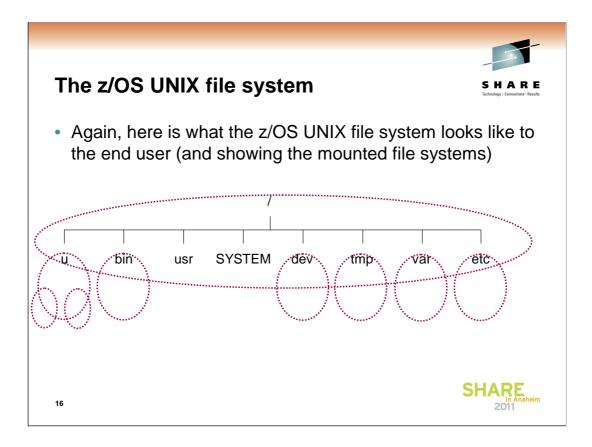


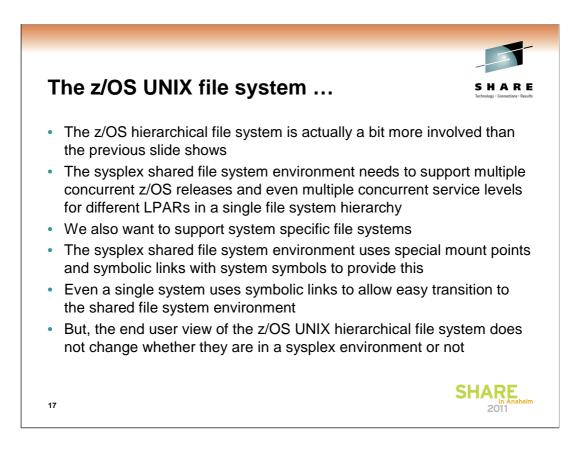


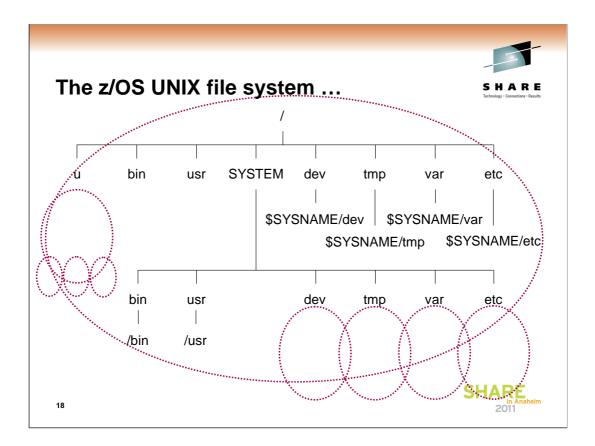
If you try to mount a zFS read-write file system on multiple systems, zFS will deny the second mount.





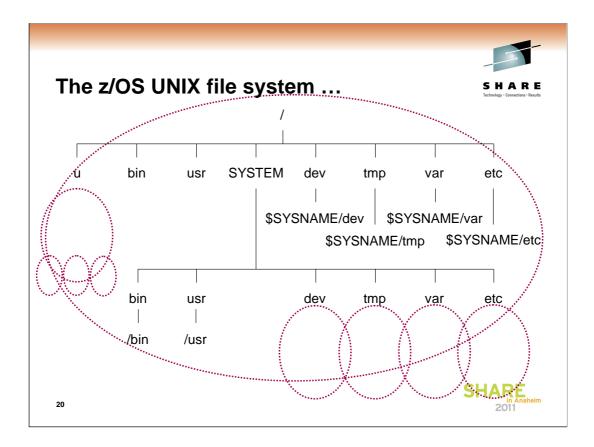


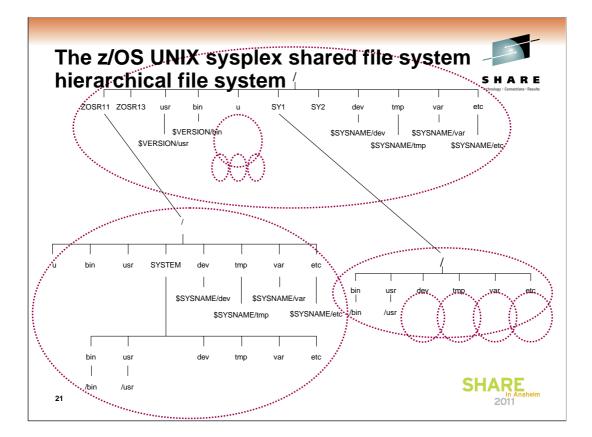


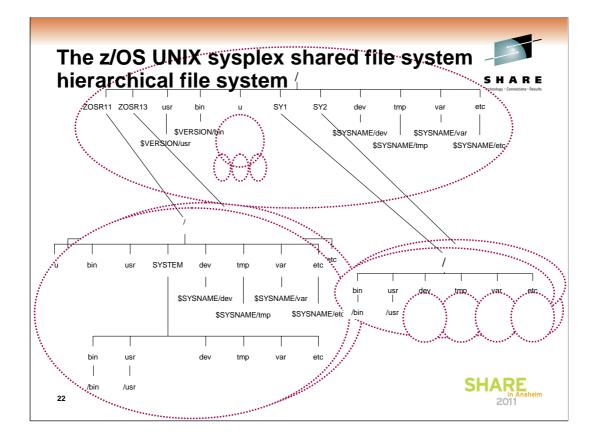


If the content of the symbolic link begins with \$SYSNAME and SYSPLEX is specified NO, then \$SYSNAME is replaced with /SYSTEM when the symbolic link is resolved.









From SY1 (running z/OS V1R11)

# df /bin

Mounted onFilesystemAvail/TotalFilesStatus/ZOSR11(OMVS.MNT.ZOSR11.ZD1111.ZFS)26236/45792004294950685Available

From SY2 (running z/OS V1R13)

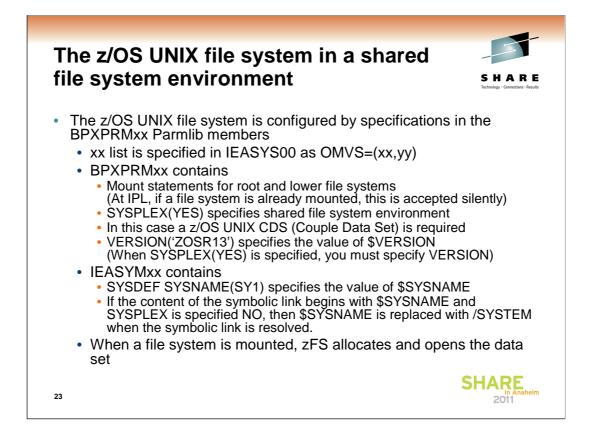
# df /bin

Mounted on Filesystem Avail/Total Files Status

/ZOSR13 (OMVS.MNT.ZOSR13.ZD1131.ZFS) 1113638/5760000 4294951449 Available

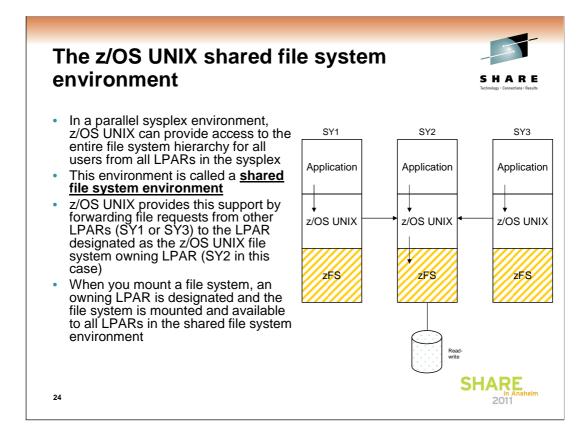
Anyone can access a different system's /etc by using a full pathname such as /SY1/etc

Anyone can access a different release of /bin by using a full pathname such as /ZOSR13/bin



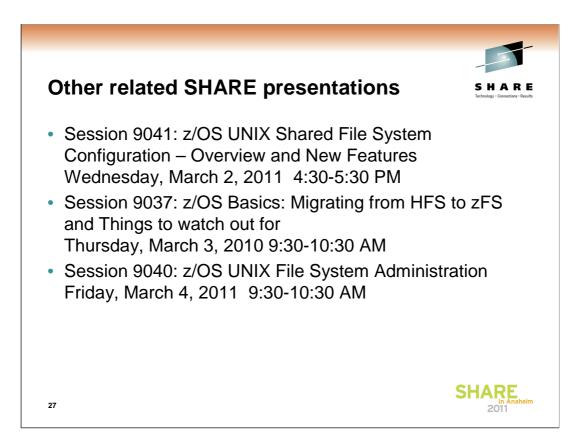
You can use System Symbols in Parmlib members. For example,

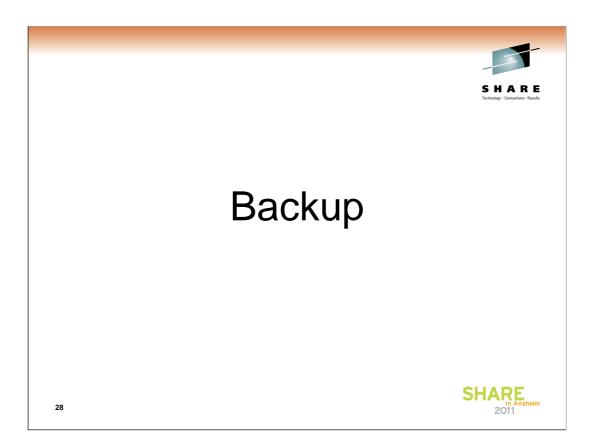
In IEASYS00, you can say OMVS=(&SYSCLONE.,01)

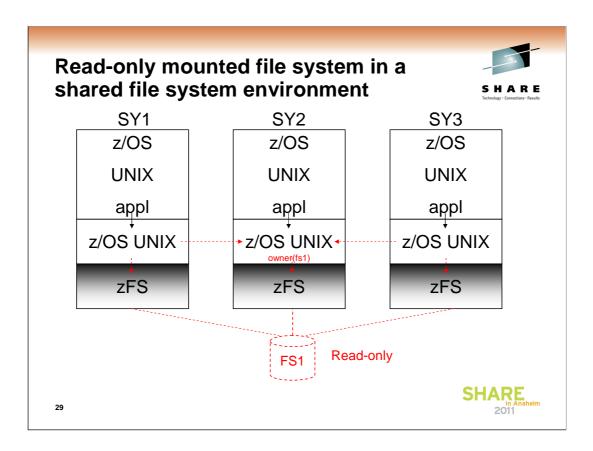


Benefits of the shared file system environment	SHARE Technology - Connections - Results
<ul> <li>System independence         All z/OS UNIX data can be accessed from any system in the sysplex     </li> <li>Availability         If a system is shutdown or if it abnormally goes down, file system             ownership movement occurs automatically and file systems remain             accessible from the other systems (although temporary failures may             be visible to applications during abnormal shutdown)     </li> </ul>	
<ul> <li>Flexibility         General users and applications automatically access the co (system specific and/or version/release) file systems while administrators can access any file system     </li> </ul>	rrect
<ul> <li>Transparency Users and applications do not need to change to run in a sh system environment (except to possibly handle and recover temporary failures)</li> </ul>	
25	SHARE 2011

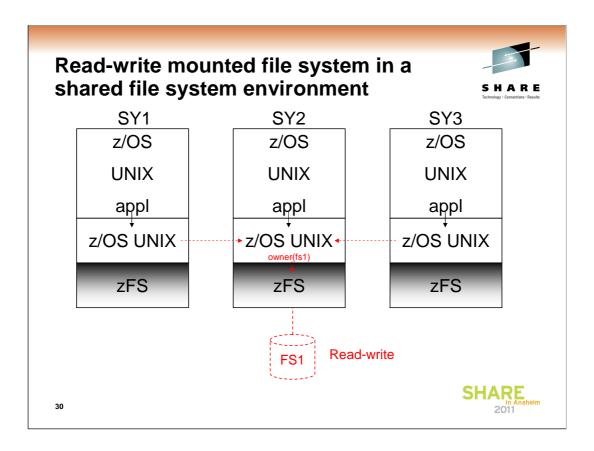




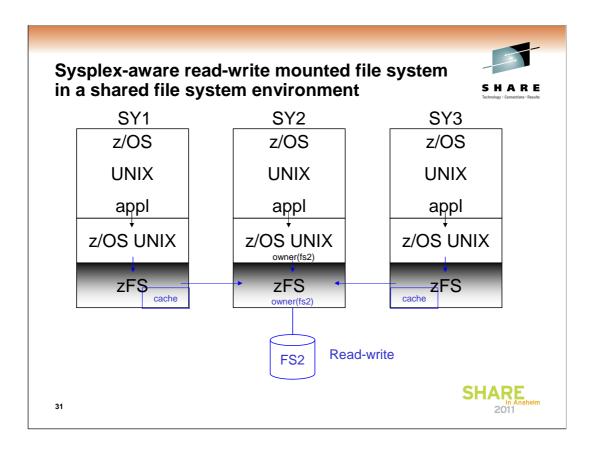




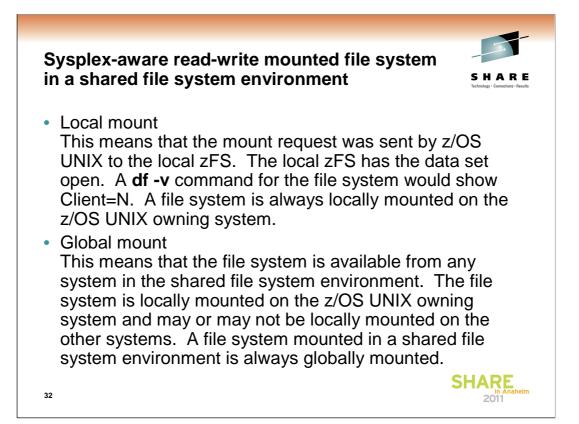
This shows a shared file system environment in a sysplex (z/OS UNIX BPXPRMxx specifies SYSPLEX(YES)). When a zFS file system is mounted read-only, it is locally mounted on each system. Read-only mounted file systems are always sysplex-aware. This has always been the case for shared file system environments. There is no communications between sysplex members to access a read-only mounted file system.

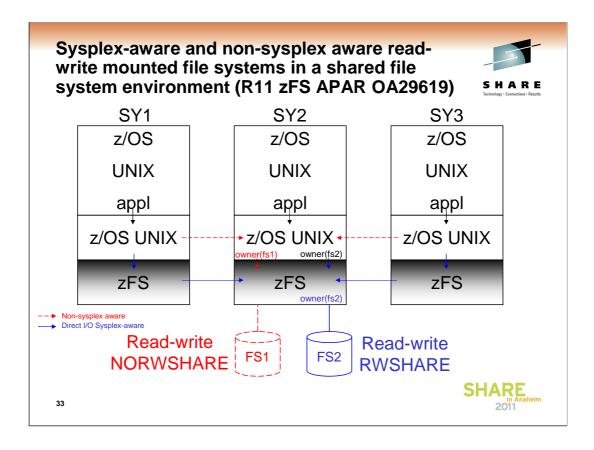


In contrast to a read-only mounted file system, this is a read-write mounted file system. One system is the owning system and the other systems are "clients". The file system is only locally mounted on the owning system. (It is still externally mounted and available on all systems.) Applications running on the owning system (SY2) access the file system locally (without any XCF communications). But applications that run on the other systems (SY1 and SY3), access the file system through the use of z/OS UNIX function shipping (using XCF communications). That is, the request is forwarded by the z/OS UNIX on that system (SY1 or SY3) to the z/OS UNIX running on the owning system (SY2) which then program calls the local zFS. The response goes back along the same path. So, access to the file system from systems other than the owner (that is, from the client systems) involves XCF communications. This makes it important to have the z/OS UNIX owning system, be the system that is doing the most accesses.



Here is a picture of the new R11 zFS sysplex-aware for read-write support. When zFS runs sysplex-aware (for read-write) on all systems, a read-write mounted sysplex-aware file system is locally mounted on all systems. There is still a z/OS UNIX owning system but there is no z/OS UNIX function shipping to the owner. Rather, requests from applications on any system are sent directly to the local zFS on each system. This means it is now the responsibility of the local zFS to determine how to access the file system. One of the systems is known as the zFS owning system. This is the system where all I/O to the file system is done. zFS uses function shipping to the zFS owning system to access the file system. (If this was all that zFS did, it would be essentially the same as the z/OS UNIX function shipping as shown in the previous slide.) However, each zFS client system has a local cache where it keeps the most recently read file system information. So, in many cases (when the data is still in the cache), zFS can avoid the zFS function shipping (and the XCF communications) and satisfy the request locally.

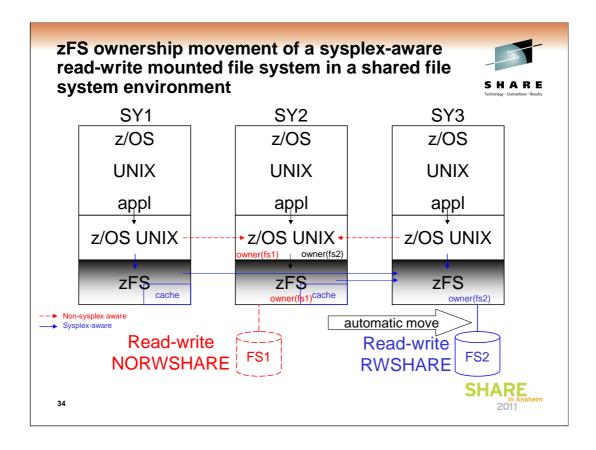




Here are two zFS read-write file systems on a sysplex running zFS sysplex-aware on file system basis on all members. One file system (FS1) is mounted NORWSHARE and the other (FS2) is mounted RWSHARE. They are both z/OS UNIX owned on SY2. The NORWSHARE file system (FS1) acts like a non-sysplex aware file system (it is only locally mounted on the z/OS UNIX owner (SY2) and requests from z/OS UNIX clients (SY1 and SY2) are function shipped to the z/OS UNIX owner (SY2) by z/OS UNIX). A df –v command for the NORWSHARE file system (FS1) from SY1 would display Client=Y.

The other file system (FS2) is mounted RWSHARE. It acts like a sysplexaware file system (it is locally mounted on all systems and z/OS UNIX never function ships requests to the z/OS UNIX owner (SY2)). A df –v command for the RWSHARE file system (FS2) from SY1 would display Client=N.

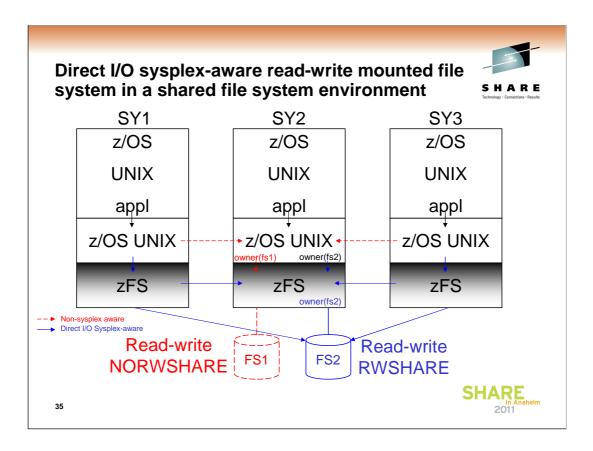
When you run zFS sysplex-aware on a file system basis on all your members, the zFS Physical File System initializes as sysplex-aware but it can individually determine which file system is sysplex-aware and which is not based on the RWSHARE/NORWSHARE mount parm.



In addition, zFS can move its ownership of zFS read-write sysplex-aware file systems dynamically based on system usage among zFS systems where the file systems are locally mounted. If one sysplex member is a better fit for ownership (because more requests to the file system are being done there than any other sysplex member), then zFS will dynamically move zFS ownership to that system. zFS will also move ownership when the owning system is shutdown, or an abnormal outage occurs on the zFS owning system. Thus the z/OS UNIX owner and the zFS owner can be two entirely different systems depending on the sequence of events. This is normal, and for a sysplex that is using the zFS sysplex=filesys support on all sysplex members this should have no negative consequences.

zFS does not have any external commands to move zFS ownership between systems. It does provide query commands to show the zFS owner, such as **zfsadm Isaggr**.

Note that zFS ownership movement based on performance criteria is nondisruptive.



In zFS R13, zFS supports direct I/O for sysplex-aware read-write file systems. This means that when all systems are running z/OS V1R13, zFS can directly read and write user data to a sysplex-aware (RWSHARE) read-write file system (for example, FS2). This generally improves the performance of client system access to the files (from SY1 and SY3) since the data does not need to pass through SY2 with XCF communications. Only metadata updates are sent to the zFS owning system.



## Adding a new system or version root

- In general, the contents of the sysplex root should only change when you need a new version root or system-specific root file system directory for your shared file system configuration. When a system is IPLed (initialized), the mount processing for the sysplex root file system will include defining the appropriate \$SYSNAME or \$VERSION directory in the sysplex root file system if the sysplex root is mounted as read/write. Assuming that you have the sysplex root file system mounted as read-only, the procedure to use to create a new version root or system-specific file system directory is as follows:
  - · Remount the sysplex root file system to read-write
  - IPL the new system
  - · Remount the sysplex root file system to read-only

36

