

ibm.com



e-business



Parallel Sysplex Update 2004 Usability, Scalability, Availability

Frank Kyne - ITSO Poughkeepsie
kyne@us.ibm.com



Redbooks

General

Agenda:

- Start 09:00 (ish!)
- Lunch about 12:00 for 1 hour
- Finish about 17:00.
- Breaks every 90 minutes or so, for **15** minutes



General

My background.....

Your handouts vary slightly from the presentation as I am constantly changing the presentation as new information becomes available - sorry!

The latest handouts will be available to IBM'ers in ZIPped PDF format (pw ITSO2004) - go to w3.itso.ibm.com, then Redbooks Online, then Additional Materials, then ITSO Materials Repository. Customers can get the latest handouts (with the same password) from:

ftp://www.redbooks.ibm.com/redbooks/itso_zseries_workshops_2004/2004fk.zip

PLEASE complete the evaluation forms.

Questions?? Please ask as I go along. Also, if you can't understand my strange accent, please let me know!



Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

eServer™	Encina®	OS/390®
ibm.com®	ESCON®	Parallel Sysplex®
z/OS®	FlashCopy®	Perform™
z/VM®	Footprint®	PAL®
zSeries®	FICON™	PC 100®
AnyNet®	GDPS®	PR/SM™
Balance®	HyperSwap™	Redbooks™
CICS/ESA®	IBM®	RACF®
CICS®	IMS™	RMF™
CICSplex®	Language Environment®	S/390®
Domino™	Multiprise®	Sysplex Timer®
DB2®	MQSeries®	SLC™
DFS™	MVS™	VTAM®
DFSMSshm™	MVS/ESA™	WebSphere®
DFSORT™	NetView®	
Encina®		

The following terms are trademarks of other companies:

Intel, Intel Inside (logos), MMX, and Pentium are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

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



Agenda

Topics:

- Sysplex and availability-related functions in recent releases of z/OS (1.5, 1.6, 1.7 preview, z990 GA3)
- Availability checklist - latest news and views
- Parallel Sysplex configuration options and recommendations
- Planning for second data centers
- Update on z/OS HealthChecker - latest news and future directions
- Bits and bytes



Summary of z/OS 1.5 and 1.6 new features, 1.7 preview, z990 GA3



z/OS 1.5 enhancements

Base Control Program (BCP)

- [Console restructure](#) (previously available on top of 1.4)
 - There have been ZERO console-related outages in customers with the new support installed!
- [System-Managed Duplexing](#) support rolled into base
- [Service task monitoring](#) in System Logger provides ability to bypass or redrive hung requests
- New option [STOPGTF](#) on SLIP command, allows you to automatically stop GTF tracing when SLIP trap is disabled
- [IPCS performance](#) improved by keeping multiple data sets open



z/OS 1.5 enhancements

DFSMS

- [JOB/CAT/STEP/CAT](#) support removed by default - use F CATALOG,ENABLE(JOBSTEP/CAT) to enable
- Automatically enable [AUTOADD](#) for ECS following an IPL - SYSCAT keyword in LOADxx
- [BUILDIX](#) no longer requires volume to be taken offline to sharing systems
- [REFORMAT](#) can be used to non-disruptively extend VTOC and VTOCIX
- Ability to recreate a lost SMS SCDS ([SETSCDS](#) command)
- [Enhanced dataset integrity](#) for sequential data sets
- New manual specifically to describe new features:
 - DFSMS: Using DFSMSdfp in the z/OS V1Rx Environment



z/OS 1.5 enhancements

Comms Server:

- Sysplex Distributor has new keyword (**DISTMethod ROUNDROBIN**) to provide another option for routing requests
- Sysplex Distributor **Application Server Affinity** feature allows for coexistence of workload balancing and persistent connections
- New **MOVEABLE (IMMED|WHENIDLE)** keyword on VIPABACKUP allows a VIPA to be started on the backup if not already active in the plex, and subsequently automatically moved to its "normal" stack
- New **SERVICEMGR** keyword on VIPABACKUP tells TCP/IP it is to act as the service manager for CISCO's MNLB
- Limit on number of DVIPAs increased from 256 to 1024 and VIPADISTRIBUTE ports increased from 4 to 64
- Performance of **SYSPLEXPORTS** enhanced
- Resolver setup file and TCPIP.DATA support system symbols



z/OS 1.6 enhancements

z/OS 1.6 will only run on processors that support Architecture Level Set 2 (zSeries for IBM - no 9672 support)

BCP:

- z/OS 1.6 will support up to 24 CP's per Lpar.
 - On z990 you are able to define 32 CP's per Lpar but T&Cs restrict you to 24 until some time in 2005.
- IRD - new keyword in IEAOPTxx (**VARYCPUMIN**)
 - Specifies the minimum number of CPs which must stay online during WLM LPAR management.
 - This will be rolled back to z/OS 1.4 via APAR OA09047
- SETXCF **REALLOCATE** command to place structures based on PREFLIST
- SYMUPDTE-like program shipped (in OBJ form) in SYS1.SAMPLIB member **IEASYMUP**



z/OS 1.6 enhancements

BCP (cont):

- GRS SYNCHRES default changed to YES
- You can delete bound and unbound PAV alias devices via ACTIVATE
- With Security Server in z/OS 1.6, it is no longer necessary to IPL to pick up a new Class Descriptor Table
 - Tool called CDT2DYN will be available from <http://www.ibm.com/servers/eserver/zseries/zos/racf/goodies.html>
- With z/OS 1.6 Security Server, most installation-defined classes will not require a change to the RACF router table (ICHRFROx)
- RRS enhanced so that Resource Managers can be restarted on another system even if RRS is still up on the original system (restart anytime anywhere)



z/OS 1.6 enhancements

DFSMS

- Restartable PDSE address space (SMSPDSE1)
 - Eliminates need to re-IPL a system due to a hang, deadlock condition, or out of storage condition
 - SMSPDSE still not restartable
- SMS Volume Selection based on PAV
 - Allows allocations to be automatically allocated to high performance devices
- DFSMSrmm can now run on a system which does not have direct access to the DASD containing the DFSMSrmm CDS
 - I/O requests to the CDS are handled over TCP/IP
 - Allows multiple sysplexes to have a single tape inventory without shared DASD



z/OS 1.6 enhancements

Comms Server:

- Sysplex functions supported by IPV6
 - Dynamic VIPA and Dynamic VIPA Takeover
 - Sysplex Distributor
- Automatic VIPA takeover if stack is not healthy
- Autonomic functions to monitor sysplex health added
 - Monitor storage usage
 - CSM, TCPIP Private & ECSA
- Monitor dependent networking functions
 - VTAM availability
 - XCF routes available
 - OMPROUTE availability
- TN3270 can now be in its own address space, so can be restarted independently, and given separate WLM goals



z/OS 1.6 enhancements

USS:

- Automove System List
 - Wildcard('*') character support is added to use in the Automove include system list.
- PFS Termination in a shared HFS
 - If zFS terminates, ownership will be moved to another member if possible. If not, it is globally unmounted and unavailable.
 - If possible, former owner of filesystem will function ship filesystem request to new owner. Both for R/O and R/W.
 - When PFS restarts, any R/O filesystems will go back to local mounts instead of function shipping.
- Latch contention
 - New ability to abend a task causing latch contention for extended period of time (F BPXOINIT,RECOVER=LATCHES command)



z/OS 1.6 enhancements

Growth constraint relief:

- Maximum number of SMF buffers increased
 - 2 new keywords in SMFPRM (BUFSIZMAX, BUFUSEWARN)
- Maximum number of members in an XCF group doubled to 2047 rolled back to z/OS 1.4 with APAR OA04034
- The maximum number of structures in the CFRM policy is increased to 1024 via APAR OA06240
- Number of linkage indexes increased from 2K to 32K
- Number of Dynamic VIPAs increased from 256 to 1024
- Number of VIPADISTRIBUTE ports increased from 4 to 64
- Maximum number of file descriptors per task doubled to 128K
- Control blocks used by EXCP moved above the line
- Max number of CPs increased to 24, then 32 in 2005



z/OS 1.7 Preview

XRC+ - Will allow the installation to specify whether I/Os to the Logger staging data sets should be synchronous or not.

Support for >64K tracks for sequential data sets

Support for >255 extents for VSAM files

Enhancements to make zFS more resilient to problems - zFS can be used in place of HFS in all cases

TCP/IP enhancements in areas of load balancing, optimizing available paths within the sysplex, and recovery

z/OS HealthChecker becomes part of base operating system

SDSF enhanced to let you look at HealthChecker messages



z990 GA3 announce

Number of ISC3 links increased from 32 to 48

Support added for spanned external CF links (ISC, ICB2, ICB3, ICB4). Spanned IC links were already supported

CFLevel 13 provides performance enhancements

Non-disruptive apply of new CFLevels

On/Off Capacity on Demand for ICFs

- **It is possible now to have both OOCoD and CBU installed - no longer mutually exclusive**

Support for up to 24 CPs per LPAR when used with z/OS 1.6 or z/VM 5.1

Ability to dynamically re-assign PUs (change to CPs, ICFs, or IFLs)

Support for dynamically renaming LPARs (requires z/OS 1.6)



Parallel Sysplex Availability Checklist



Objective

All mainframe shops have (at least!) two challenges:

- Provide better availability
- Do it with fewer staff!

To help, this presentation provides **SOME** recommended ways of improving the availability of your applications in a sysplex environment. Checklists are provided in an appendix at the end of the presentation, to make it easy to identify features that you are not currently using.



Parallel Sysplex availability

There are a number of facets to availability

- You can have *planned or unplanned* outages
 - Vast majority of outages are planned
 - As machines and systems get larger, planned outages are getting harder and harder to schedule
 - Planned outage avoidance consists of exploiting capabilities in the software and hardware to make dynamic changes, plus good planning
 - Unplanned outage avoidance consists of good configuration planning, keeping up-to-date on service, careful system monitoring, and so on.



Parallel Sysplex availability

Facets to availability.....

- You want to avoid outages, but for unavoidable ones, you want to minimize (mitigate) the impact of the outage
 - In the case of an outage, you want your users to still be able to use their applications.
 - The only way to achieve this is to exploit Parallel Sysplex data sharing and dynamic workload balancing
 - You also want to get back up and running as quickly as possible - this requires education, automation, planning, and testing



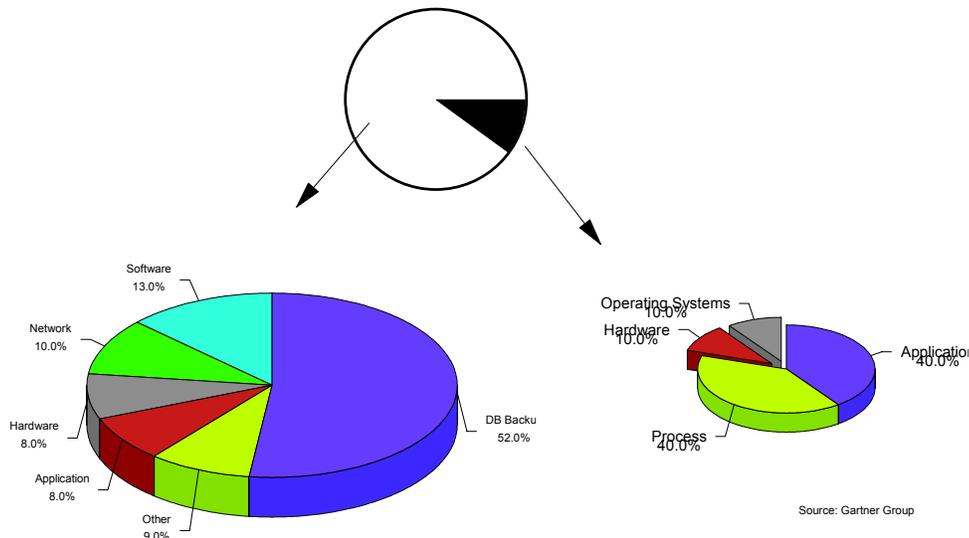
Parallel Sysplex availability

Facets to availability.....

- Performance
 - There comes a point when performance is *so* bad that the application must be viewed as being unavailable
 - Do you have agreed thresholds? Do you have monitoring to:
 - ▶ Warn you when the threshold is reached
 - ▶ Tell you *what part* of the application is causing the problem? There is no point re-IPLing z/OS if the cause of the problem is a faulty network router.....
 - Do you have agreed actions to take immediately when those thresholds are reached? Or do you extend the "outage" by arguing over what action should be taken?



Continuous Application Availability



Parallel Sysplex Availability

To get the availability benefits enabled by a Parallel Sysplex, you *must* make the effort to exploit the capabilities it provides

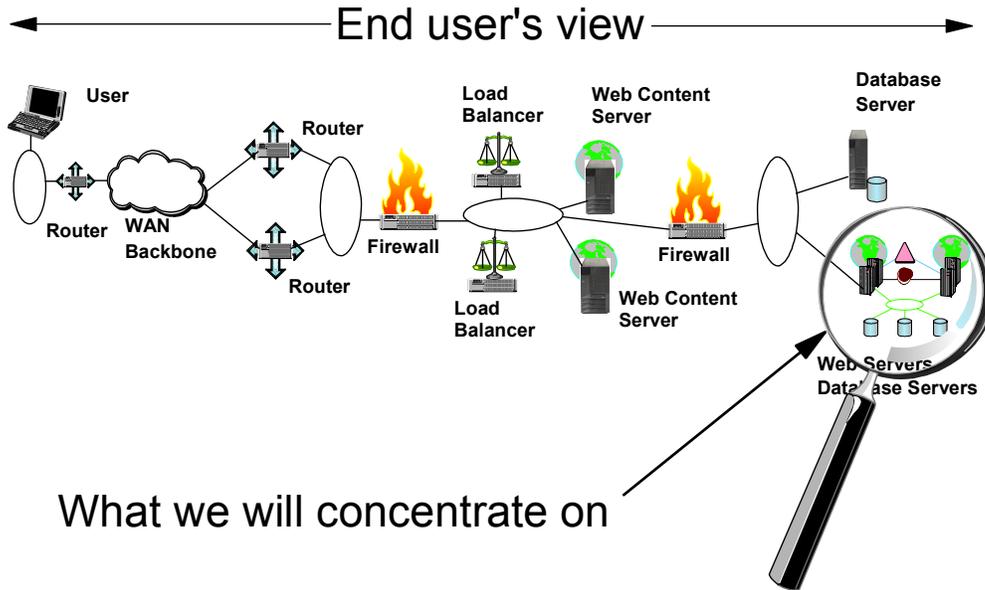
Installing a Coupling Facility does not automatically eliminate all outages

There is NO magic silver bullet to provide continuous availability

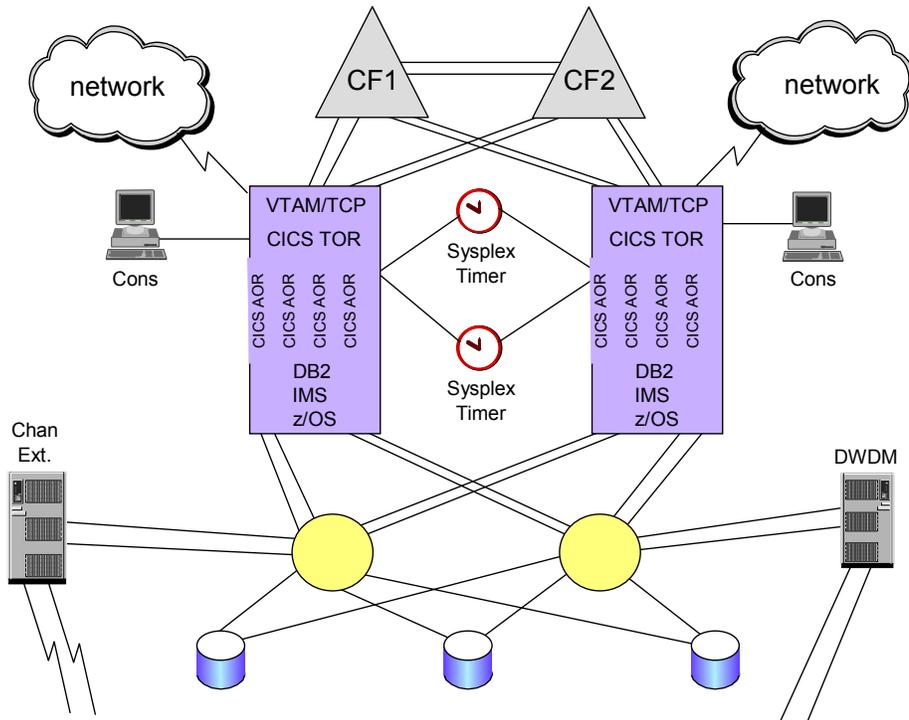
BUT... there are many small changes you can make - when added together, these can make a significant improvement to availability



Parallel Sysplex Availability



Application service delivery components



Management considerations

Does the spending on IT applications reflect their revenue contribution to the business?

Does the executive responsible for availability have the power to stop applications going into production? Does he have the budget to address availability issues?

Availability relies on Ops, Tech Support, and Application Development - do all these groups have availability as part of their objectives?

Is education provided when new HW or SW is installed?
Are the features of new releases being exploited?



Systems Management considerations

What is your current availability? If you don't know where you are now, how will you know if changes have the hoped-for impact?

Where is your pain? You have to know where your problems are so you can address them

How many problems are caused by problems where the fix is already available? Would more aggressive service policy have any impact?

Automation. Who writes it? What happens if it breaks?
Does it raise or reduce operator skill levels?



CPU Considerations

Where possible, use the latest available technology.
Ensure you are aware of, and exploit, features like CBU, OOCOD, Concurrent PU conversion, I/O Upgrade planahead, etc

Implement at least 2 HMCs and 2 HMC phone servers

Predefine spare LPARs on z990 using LPAR name of *

Define all LPARs with RESERVED CPs

ALWAYS enable Automatic I/O Interface Reset in Reset profile unless you are using ARF

LPARs should be successfully IPLed or DEACTIVATED - don't have LPARs that are ACTIVATED but not running

Apply APAR OA09001 to avoid boxed CTCs when upgrading 9672 to 2064 or 2064 to 2084



Switches, directors, and timer considerations

Sysplex Timer mod 1 out of support since 2003 - upgrade to mod 2 is disruptive

ESCON Director Mod 5 AND ALL MESes withdrawn from marketing effective end 2004

Next generation HMC will not support Sysplex Timer or ESCON Director consoles, so address now while you can still order an MES for ESCD

A patch will be made available for z900 and z990 to address high FICON CTC response times in some cases

Ensure that fibre for FICON and FCP links is within IBM specifications



CF Considerations

Keep on current CF Levels

CFLevel 13 adds ability to upgrade non-disruptively to future CFLevels on z890/z990

If using large DB2 GBPs, or have data and indexes in the same GBP, move to CFLevel 13 on z800/z890/z900/z990

Make sure structures are in the CF you intended them to be in - see SETXCF START,REALLOCATE command (OA03481 on z.OS 1.4 and later)

If using n-way CFs, or lots of XCF, move to CF Level 14



CF Considerations

Always have at least 2 CFs

CFs should not be run > 50% busy

If using z990 or have long links, carefully monitor CF Link utilization - 2 links may no longer be sufficient

- **Path busy conditions should not exceed 20%**

Always use highest speed link possible - IC, then ICB, then ISC

Avoid use of shared engines for production or response-time sensitive CFs

If dynamically changing CF Links, use SOFTWARE with VALIDATE option on ACTIVATE (See APAR OA09012)



z/OS HealthChecker

Pull the HealthChecker from

<http://www14.software.ibm.com/webapp/download/search.jsp?go=y&rs=hchk>

Modify your environment and/or the HealthChecker until you get a clean run

Consider using WTO feature to connect automation to HealthChecker findings

Then, run it every day, to warn you of changes to your configuration

Futures...



z/OS considerations

Configure paging subsystem to handle maximum paging rate

- **Can place multiple page data sets on a single (large) volume if you use Dynamic PAV on that volume**

Use extended format stand alone dump data sets (requires APAR OA04140) and stripe them over 4 volumes using AMDSADDD exec

Check GRS Planning book for recommended values for RACF ENQs and SYSZVVDS and SYSIGGV2

z/OS 1.5 removes requirement for IPL to implement new templates

z/OS 1.6 adds support for dynamic updates to Class Descriptor Table and removes need for router table entries (ICHRFRxx)



z/OS considerations

z/OS 1.4 adds JES2 health monitor (no relation to z/OS HealthChecker!). Runs in separate address space, delivers \$J command.

Exploit Console Restructure in z/OS 1.5 or feature on top of 1.4

APAR OA07335 (or z/OS 1.6) dramatically reduces time to build DPS arrays during master scheduler initialization

z/OS 1.6 (and 1.5 with APAR OW56001) provides ability to increase SMF buffers up to 1GB



XCF considerations

Should have 1 XCF structure per CF per transport class

Every transport class should have at least 2 paths defined

Ensure XCF structures don't all end up in the same CF

- Ensure **POPCF** or **REALLOCATE** is used after every CF outage

Ensure primary and alternate CDSs are on different devices and preferably different CUs

ALWAYS enable SFM, with **ISOLATETIME(0)**



CFRM considerations

Ensure CFRM CDS is configured with (at least) SMREBLD and SMDUPLEX (if all systems > z/OS 1.2 with APAR OW41617)

ALL structures should have at least 2 CFs in PREFLIST
NEVER specify a REBUILDPERCENT value - allow it to default

Use appropriate structure sizes. Check CF Sizer and RMF reports. Check CFRM definitions against actual sizes - many exploiters will automatically increase structure size



System Logger considerations

Refer to recent ITSO Redbook on System Logger (z/OS Systems Programmers Guide to: System Logger (SG24-6898)) for detailed information on Logger usage and tuning

Ensure LOGR CDS is formatted with SMDUPLEX option (use D XCF,C,TYPE=LOGR to display)

z/OS 1.3 added ability to change most log stream attributes while the log stream is still connected

All logger structures should contain at least two log streams, and be connected to by more than one system



DFSMS Features

Use Multiple Address Space HSM support - allows you to specify higher WLM priority for recall tasks

Exploit Common Recall Queue - customers very pleased with results - faster recalls, better workload balancing, reduced tape mounts

If HSM is very busy, enable Fast Subsequent Migration - will increase CDS sizem but decrease amount of data movement HSM has to do

Use VSAM RLS for CDS sharing - major improvement for CDS-intensive processing



DFSMS features

Many enhancements in SMS ("Reduce space by %", Extend Storage Groups, Overflow Storage Groups, etc) to help reduce the chance of space abends

- See also [Tivoli Allocation Manager \(5698-A30\)](#)

z/OS 1.5 adds ability to issue a WTO if storage group reaches high threshold

z/OS 1.6 adds ability to request a PAV volume when allocating a data set

OS/390 2.10 added ability to rename an ENQed duplicate data set - see STGADMIN.DPDSRN RACF profile

Enhanced Data Set Integrity added in z/OS 1.5. Controlled via SYS1.PARMLIB IFGPSEDI member

Specify PDS_RESTARTABLE_AS in IGDSMSxx to get restartable SMSPDSE1 address space



CICS considerations

Ensure all CICS applications are compiled and linked for 31-bit. Single largest cause of CICS outages is storage exhaustion below the line

Define CICS TORs as VTAM generic resources and have two cloned TORs on each system

Move newer code (Java, Sockets) into its own address space and have at least two of them

If using DB2 group attach support in CICS TS 2.2, specify RESYNCMEMBER(YES) so CICS will reconnect to the same DB2 after a failure



DB2 considerations

Use DB2 data sharing

ALWAYS use duplexing for production DB2 GBPs

Use RESTART LIGHT to restart DB2 on another image following a failure. V8 enhances this so that DB2 will connect to other Resource Managers to resolve in doubt UOWs

Use partitioned databases

Use CFLevel 13 if you have data and index in same GBP and use online utilities

V8 adds ability to ALTER partition boundaries dynamically

V8 lets you alter more DSNZPARM values dynamically

V8 lets you add columns to a table dynamically



IMS considerations

Use cloned IMS control regions, data sharing, and FDBR

If using Shared Queues, consider Common Service Layer and Sysplex Terminal Management

If using VTAM GR, do NOT use IMS Rapid Network Reconnect

Define more than one RACF TCB - RCFTCB in IMS Parmlib

Exploit High Availability Large Databases (HAL DBs) - similar concept to DB2 partitioned DB2

IMS V9 removes need for an IPL to add Type 4 SVC or Resource Cleanup modules

To restore IMS DBs to any point in time (as you can with DB2 DBs), use Database Recovery Facility



Parallel Sysplex Availability

Other sources of information:

- [SG24-6061 Sysplex Availability Guide \(currently in development\)](#)
- [SG24-2086 Continuous Availability S/390 Technology Guide](#)
- [SG24-4502 Parallel Sysplex Continuous Availability Presentation Guide](#)
- [SG24-4503 Parallel Sysplex Continuous Availability SE Guide](#)
- [SG24-4593 Planning for CICS Continuous Availability in an MVS/ESA Environment](#)
- [SG24-5680 Implementing ESS Copy Services on S/390](#)
- [Parallel Sysplex Training Environment, available from IBM Learning Services](#)
- [Parallel Sysplex Availability Checklist, available at:](#)
 - http://www.ibm.com/servers/eserver/zseries/library/whitepapers/pdf/availchk_parsys.pdf
- [Other Parallel Sysplex availability documents available at:](#)
 - <http://www.ibm.com/servers/eserver/zseries/psa/>



Parallel Sysplex Configuration Options



Configuration options

In this session we will discuss:

- Coupling Facility introduction
- Coupling Facility attributes - *what is available*
 - Function levels
 - Link types
 - Availability characteristics
 - Speed
 - Cost
- CF exploiter requirements - *what you need*
 - 3 categories
- Matching requirements to attributes
 - Getting maximum value for your money

Coupling Facility quick review

Some general points about CFs:

- ALL CFs run the SAME CFCC, regardless of whether they run in standalone boxes (9674, 2064-100, etc), in CF LPARs on general purpose CPCs, use dedicated or shared engines, use ICMF, or even run under VM.
- CFCC normally runs in a polling loop, ensuring good response times, but also consuming all MIPS available to that LPAR.
- Availability is mainly a factor of the lack of single points of failure between a CF and the Operating System LPARs connected to it - not of anything within the CF itself - and of having a second CF to rebuild into in case of an outage.
- CF storage requirements are dictated by the structures that will reside in each CF, and the CFLevel. Unlike MVS, providing more storage than necessary does not affect performance.



CF levels

The CFLevel (Level of the Coupling Facility Control Code (CFCC)) is important because it determines the functionality available in that CF

CFCC is Licenced Internal Code, and is generally related to a CPC Driver level - you cannot order a CFLevel upgrade independently

The current CFLevel for each CF can be determined using the D CF MVS command on a connected system

New CFLevels generally deliver additional function and often affect structure sizes. Old software levels can coexist with new CFLevels, but to exploit new functions may require current software levels

With CFLevel 13 (or later) and z990/z890, you can have CF LPARs with two different CFLevels on the same CPC



CF Function levels

CFLevel 14 - CFCC Dispatcher rewrite

CFLevel 13 - Non-disruptive patch apply, Castout enhancements

CFLevel 12 - (zSeries only) z800 support, 64-bit, System-Managed Duplexing (zSeries), support for 48 CF tasks, new batched IXL requests, Message Time Ordering

CFLevel 11 - (9672 G5/G6 only) System-Managed Duplexing (G5/G6)



Mixed CFLevel coexistence considerations

Generally speaking, it is OK to have two different CFLevels in the same sysplex, however:

- When you rebuild a structure into a CF, XES may adjust the size of the structure based on the CFLevel and the number of entries and elements in the structure.
- Make sure that the structure size defined in the CFRM policy is appropriate for the CF that the structure will normally reside in (normally the first CF in the PREFLIST).
- Make sure that you don't exploit any functions that are not supported in ALL attached CFs.



CF Link types

The types of links you will use to connect the CFs to the operating system LPARs is important because of the impact of link performance on response times, and, therefore, coupling overheads

For configurations covering large distances, the time spent on the link can be the largest part of the response time

Generally speaking, newer link types provide reduced latency, and better long-distance performance

Not all processor types support all link types



Link types

Fibre links

- **ISC links (also known as ISC2 or compatibility mode)**
 - Fibre
 - Supports up to 20 km (with RPQ), 40 km with DWDM
 - Up to 50MB/sec with multimode, 100MB/sec with single mode
- **ISC3 links (Peer mode)**
 - Fibre
 - Supports up to 20km (with RPQ), 100 km with DWDM
 - Single mode only, up to 200 MB/sec (100 MB/sec if using >10 km RPQ card (8P2197))
 - ISC3 and ISC connectors are compatible - upgrading from ISC to ISC3 on a zSeries is a HCD definition change
 - ▶ **NO ONE should be using compat mode ISC links if both ends of the link are zSeries processors**



Link types

Copper links

- ICB links (also known as ICB2 or compatability mode)
 - Copper
 - Max distance of 7 metres between CPCs (cable is 10 metres)
 - Up to 333MB/sec
- ICB3 links (Peer mode)
 - Copper
 - Max distance of 7 metres between CPCs (cable is 10 metres)
 - Up to 1000MB/sec
 - ICB3 use different cards and a different connector than ICB2 - ICB2 to ICB3 upgrade is a hardware change
- ICB4 links (peer mode)
 - Copper
 - Max 7 metres
 - Up to 2000MB/sec
 - Different cards again - ICB3 to ICB4 is hardware change



Link types

Internal links

- ICMF links - should never be used any more, not available on zSeries
- IC links
 - Microcode
 - Can only connect LPARs in the same CPC
 - "Free" (but they do take up a CHPID)
 - Up to 750MB/sec
- ICP link (Peer mode)
 - Equivalent to IC, but only available on zSeries
 - Also "Free" (also take up a CHPID)
 - Up to 1250MB/sec on z900, up to 3500MB/sec on z990!



Links

So, what's so great about peer links:

- Higher speeds than equivalent non-peer links
 - Link cards are faster, so better response times
 - Higher bandwidth, so better response time due to less contention
- Single physical link can be both Sender and Receiver. Sender can be shared between multiple OS LPARs and one CF LPAR. Receiver must be dedicated to one CF LPAR.
- More subchannels - 7 per link/LPAR instead of 2 per link/LPAR
 - Especially important with System-Managed CF Structure Duplexing or if CF is a significant distance from z/OS system
- More buffer-space per subchannel - 32K vs. 4K
 - Especially important for long distances as it reduces "handshaking" for large data transfers



Links

To summarize:

	ISC	ISC3	ICB	ICB3	ICB4	Total Ext Links	IC	ICP	ICMF	Total All Links
9672-R06	32	0	18	0	0	32	32	0	N/A	64
2066-0CF	0	24	0	6	0	26	0	32	0	58
2064-100	0	32/42 ⁽¹⁾	16	16	0	32	0	32	0	64
2086-0CF	0	48	0	16	8	58	0	32	0	64
2084-300	0	48 ⁽²⁾	8	16	16	64	0	32	0	64

1) up to 42 links with RPQ 8P2248

2) Increased from 32 to 48 links with GA3



CF Availability characteristics

Most companies implement Parallel Sysplex because of the higher availability it can provide

So, because the Coupling Facility is the heart of the Parallel Sysplex, the availability of the Coupling Facility is a vital consideration in configuring a Parallel Sysplex



CF Availability characteristics

Non-volatile - Internal Battery Facility (IBF) and/or UPS

- Some exploiters prefer a non-volatile CF
- IBM recommends both IBF and UPS
- Remember to tell the CF if it is attached to a UPS!!

Failure-isolated - no single point of failure in common with connected operating systems

- Some structures should be allocated in failure-isolated CFs (or be duplexed)

Duplexing

- User Managed - for DB2 Group Buffer Pools
- System Managed (z/OS 1.2 and later)
- Requires CF Level 12 (zSeries) or 11 (9672 G5/G6)



CF Availability characteristics

Isolation from OS CPCs:

- If a disruptive microcode change is required by the CF, OSs will be unaffected if the CF is in a standalone CF OR if the CF is in z990 at GA3 or later.
- If a disruptive microcode or configuration change is required by the OSs, the CF will be unaffected if the CF is in a standalone CF.

Dynamic reconfiguration support for CF Receiver and peer Links

- Available on z900 GA3 or later with APAR OW54538
- Only if CF LPAR is in same CPC as z/OS...
- See WSC Flash 10191



Speed

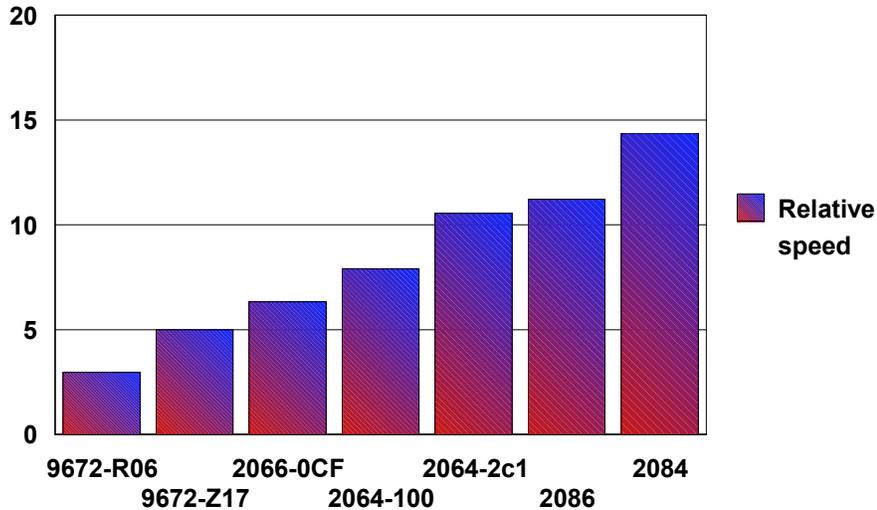
CF speed is important for two reasons:

- It is a significant factor in CF response times, thereby affecting response and elapsed times for CF users
- It directly affects the 'data sharing overhead' or cost of using the CF



CF Speeds

Relative CF Speeds



CF Speeds

And how relative speed of CF and CPC affects the overhead....

Host CF	G3	G4	G5	G6	z800	z900 1xx	z900 2xx	z890	z990
C04-SM	10%	11%	16%	19%	21%	22%	25%	---	---
C05-HL	9%	10%	14%	16%	18%	19%	22%	26%	30%
R06-HL	9%	9%	12%	14%	16%	17%	19%	22%	26%
R06-ICB	---	---	9%	10%	---	13%	14%	17%	20%
G5/6-IC	---	---	8%	8%	---	---	---	---	---
z800 ISC	9%	9%	11%	12%	11%	12%	13%	15%	18%
z800 ICB/IC	---	---	---	---	9%	10%	11%	12%	14%
z900 ISC	8%	9%	11%	12%	10%	11%	12%	14%	16%
z900 ICB/IC	---	---	8%	9%	8%	9%	10%	11%	12%
z890 ISC	7%	8%	8%	9%	9%	10%	11%	13%	15%
z890 ICB/IC	---	---	8%	8%	7%	8%	8%	9%	10%
z990 ISC	7%	8%	8%	9%	9%	10%	11%	13%	14%
z990 ICB/IC	---	---	8%	8%	7%	8%	8%	9%	9%

Based on roughly 9 CF Requests/MIP/second and dedicated engines for CFs
With z/OS 1.2 and later, CF Overhead is capped at about 15%



CF and Link Speed comparisons

The response times you can expect are largely a factor of the CF type and CF Link type - the following table gives an indication of relative times for different configurations

CF and Link type	Lock request	Cache/List request with 4K data transfer
z890-ISC3-z890	3.33	6.88
z890-ICB3-z890	1.50	2.50
z890-ICB4-z890	1.42	1.75
z990-ISC3-z990	3.17	6.38
z990-ICB2-G5	3.00	5.75
z990-ICB3-z900	2.00	3.50
z990-ICB4-z990	1.33	1.75
z990-IC-z990	1.00	1.00

* Numbers courtesy of Poughkeepsie zSeries performance team



Cost considerations

Use of ICFs vs. CPs (SW vs HW cost)

Standalone vs CF LPARs on general purpose CPC (availability)

Dedicated vs shared engines (direct cost vs overhead cost)

- Defining an engine as shared rather than dedicated, but with same capacity, increases overhead by roughly 1.2 to 2x
- Dynamic CF Dispatching
 - Is automatically turned on initially for any CF LPAR defined with shared engines
 - Should ALWAYS be turned OFF for production CFs
- Dynamic ICF expansion
 - Into shared CPs or shared ICFs

Target for CF utilization is below 50%



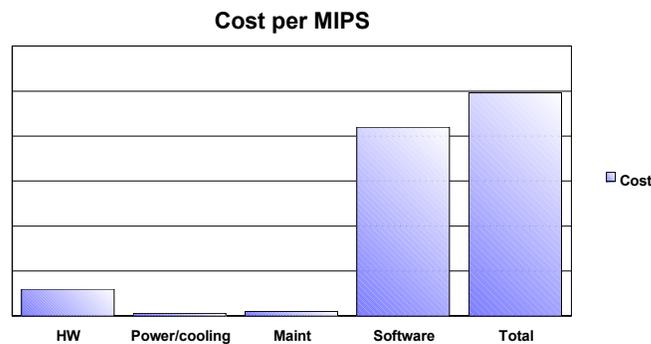
Coupling overhead cost considerations

Cost of consumed MIPS in connected CPCs (Overhead)

Cost of software associated with those MIPS

MIPS-related maintenance costs

Many customers use a figure of 10 times the actual hardware MIPS cost to determine overall cost per MIPS.



How many MIPS is it costing YOU

Synchronous operation to CF

- Software CPU time (exploiter + XES)
 - varies by structure type and exploiter
 - lock: average 26 mics on 9672-RX6
 - cache/list: average 44 mics on 9672-RX6
 - must scale to your processor's "per CP" speed
- Hardware dwell CPU time (host hdw + CF link + CF hdw)
 - reported in RMF as SYNC SERV TIME

Asynchronous operation to CF

- Software CPU time (exploiter + XES + task switch)
 - any type: average 140 mics on 9672-RX6
 - Hardware dwell CPU time - none
- In z/OS 1.2 and later, response time dictates Synch or Asynch

Lock contention event

- Software CPU time: average 800 mics on 9672-RX6



Sample overhead calculation

SYSA (9672 Y86 - per CP ratio to RX6 = 1.28)

Type	Frequency (per sec)	Software (mics)	Hardware (mics)	Total (secs)
Lock_1	3305 ¹ x	(20 +	22 ²)	= .1388
Lock_1 cont	53 ³ x	(625)		= .0331
GBP_1	1181 x	(34 +	44)	= .0921
GBP_2	2755 x	(34 +	42)	= .2094
GBP_3	4724 x	(34 +	43)	= .3637
Tot async (405+174) (from both CFs)	⁴ x	(110)		= .0637
			TOTAL	.9008

Total Host Effect = .9008 * 116 mips/cp = 104 mips



Where to get the information

RMF CF Activity reports...

```

CF NAME = CF1                                INTERVAL 015.00.851
      COUPLING FACILITY STRUCTURE ACTIVITY
STRUCTURE NAME = LOCK_1                      TYPE = LOCK

SYSTEM  TOTAL          #    -SERV TIME(MIC) -
NAME    AVG/SEC        REQ  AVG      ...  CONTENTIONS

SYSA    2977K          SYNC 2977K 22.1 2          REQ TOTAL 3969K
        3305 1        ASYNC 0    0.0          REQ DEFERRED 48K
                                   -CONT 3 48K

      SUBCHANNEL ACTIVITY

SYSTEM          #    -SERV TIME(MIC)
NAME            REQ  AVG

SYSA    ...    SYNC    3542k  42.9
        ASYNC   364771 4 212.8
        CHANGED    61
    
```



Exploiter requirements

Category 1 - no failure isolation requirement, suitable for placement in an ICF in same failure domain as connected systems. Recovery time is not an issue.

Category 2 - no failure isolation requirement, may be suitable for ICF. Recovery takes some time and will impact service during that time.

Category 3 - requires failure independence as double failure could cause service outage. Must be in standalone CF, or structure must be duplexed.



Exploiter requirements

Exploiter	Category 1	Category 2	Category 3	CFLevel
CICS Logger			✓	1
CICS Temp Stor		✓		
DB2 GBP		✓		5
DB2 Lock			✓	1
DB2 SCA			✓	1
Enhanced Cat Shr	✓			
GRS Star		✓		
HSM CRQ	✓			9
IEFAUTOS (2)	✓			
IMS Cache		✓		
IMS Lock			✓	1
IMS Logger			✓	1
IMS Shared Msg Q		✓		
IMS VSO		✓		
IMS Resource		✓		8
JES2	✓			8
LOGREC	✓			
MQSeries			✓	1

1) Category 2 if System-Managed CF Duplexing used for this structure

2) IEFAUTOS goes away with z/OS 1.2



Exploiter requirements

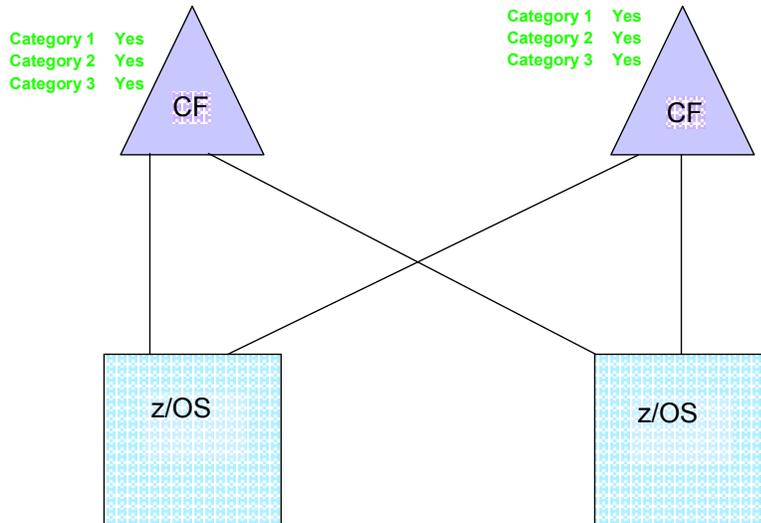
Exploiter	Category 1	Category 2	Category 3	CFLevel
OPERLOG	✓			
RACF		✓		
RRS Logger			✓	1
VSAM/RLS Cache		✓		
VSAM/RLS Lock			✓	1
VTAM GR		✓	2	1
VTAM MNPS	✓	3	✓	1
WLM Enclaves	✓			9
WLM LPAR Clust.	✓			9

- 1) Category 2 if System-Managed CF Duplexing used for this structure
- 2) If no VTAM applications in the same failure domain as the CF use the MNPS function
- 3) If no applications in the same failure domain as the CF use LU 6.2 sync level 2 conversations



Config options

2 Standalone CFs with dedicated engines



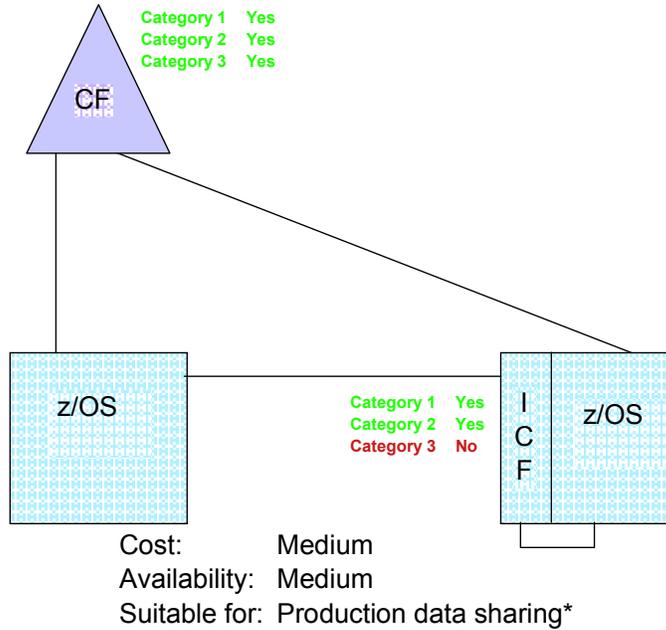
Category 1 Yes
 Category 2 Yes
 Category 3 Yes

Category 1 Yes
 Category 2 Yes
 Category 3 Yes



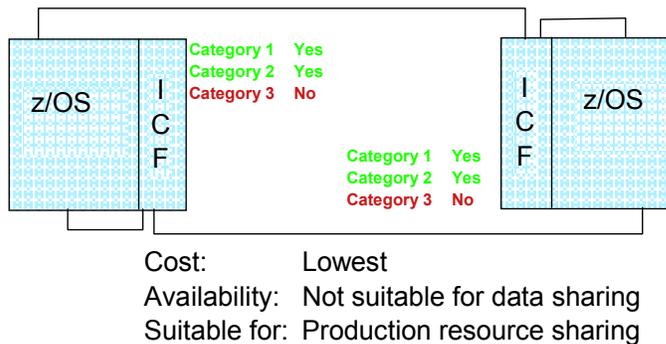
Config options

1 Standalone CF, 1 ICF, both with dedicated engines



Config options

2 ICF, both with dedicated engines, or shared engines with Dynamic CF Dispatching disabled

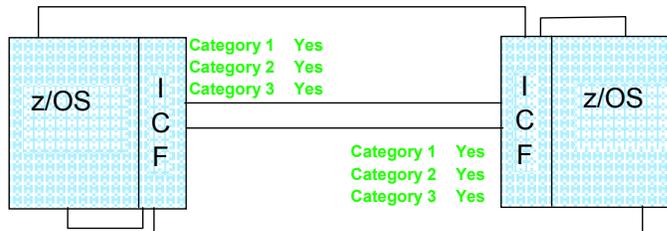


It is NOT recommended to share a CP between a CF and an operating system LP, ESPECIALLY on a CPC with just 1 engine. See WSC FLASH 10169 for details.



Config options

2 ICFs, both with dedicated engines, and System-Managed CF Duplexing



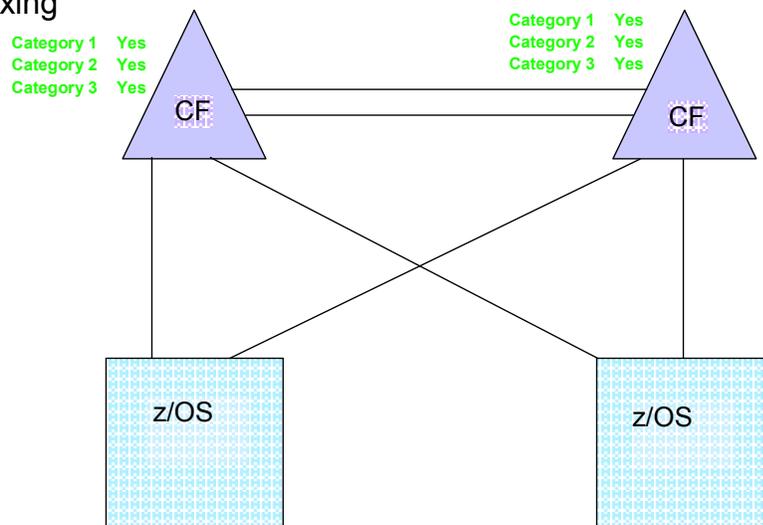
Cost: Low initial (depends on duplexing cost)
Availability: High (potential pause during CPC failure recovery)
Suitable for: Production data sharing

Note: If you have structures that require failure independence, but you don't want to duplex them, you will still require a standalone CF for those structures.



Config options

2 Standalone CFs with dedicated engines and System-Managed Duplexing



Cost: Highest
Availability: Highest
Suitable for: Production data sharing



CF Models

z800 Model 0CF



2066-OCF

- Highest supported CF level: 13
- CF links: Supports ISC, ISC3, ICB3, and ICP
- Capacity: Up to 4 engines
- Acquisition cost: Higher than R06
- Ongoing cost: Low to Moderate, depends on request rates from fast CPCs
- Failure isolation: **Yes**
- System-Manged CF Duplexing support: **Yes**
- For more info, see:

–<http://www.ibm.com/servers/eserver/zseries/library/refguides/pdf/gm130117.pdf>



CF Models

z900 Model 100



2064-100

- Highest supported CF level: 13
- CF links: Support ISC, ISC3, ICB2, ICB3, and ICP links
- Capacity: Up to 9 engines
- Acquisition cost: Similar price/MIPS to z800-OCF, but more MIPS
- Ongoing cost: Low - depends on request rates from fast CPCs
- Upgradable from 9672-R06
- Failure isolation: **Yes**
- System-Manged CF Duplexing support: **Yes**
- For more info, see:

–<http://www.ibm.com/servers/eserver/zseries/library/refguides/pdf/g3263092.pdf>



CF Models

2086-A04

- Highest supported CF level: 14 - so far
- CF links: Supports ISC, ISC3, ICB3, ICB4 and ICP links
- Capacity: Up to 4 engines
- Acquisition cost: ??
- Ongoing cost: Low - depends on request rates from fast CPCs
- Upgradable from z800
- Failure isolation: **Yes**
- System-Manged CF Duplexing support: **Yes**
- For more info, see:

–<http://www.ibm.com/servers/eserver/zseries/library/refguides/pdf/gm130522.pdf>



CF Models

2084-300

- Highest supported CF level: 14 - so far
- CF links: Supports ISC, ISC3, ICB2, ICB3, ICB4 and ICP links
- Capacity: Up to 16 engines
- Acquisition cost: ??
- Ongoing cost: Depends on request rates from fast CPCs
- Upgradable from z900, z890
- Failure isolation: **Yes**
- System-Manged CF Duplexing support: **Yes**
- For more info, see:

–<http://www.ibm.com/servers/eserver/zseries/library/refguides/pdf/gm130229.pdf>



CF Models

Next generation (z990 follow-on)

- Compatability mode ISC will NOT be supported
- This means that the next generation of CPCs will not support connection to a pre-zSeries CF



All general purpose models

Same CF level support and functionality as corresponding dedicated CF box

Generally support fewer CF Links than dedicated CFs

Capacity roughly equivalent to corresponding dedicated CF

Acquisition cost: lower because infrastructure is already paid for

Ongoing cost: Depends on types of CPCs connecting to it

Failure isolation: No (unless none of the LPARs in the CPC are connected to that CF)

Availability: not as good as dedicated CF due to impact of disruptive microcode upgrades or disruptive OS-related changes

Advantage is that CF speed gets upgraded in line with speed of connected LPARs



Choosing your configuration

To pick the best configuration for your environment, you need:

- CFs that will deliver the function (i.e. CFLevel) you need (e.g. IRD, Duplexing)
- A configuration that provides the availability you need
 - ICF, or stand alone, or ICF with System Managed Duplexing
- CFs that provide value for money in terms of comparing acquisition cost of CFs against overhead on connected CPCs.
- When selecting a CF, remember:
 - CF utilization in normal operation should not exceed 50%
 - CF subchannel utilization should not exceed 30%
 - Should configure sufficient storage, allowing for growth due to CFLevel changes
 - Overhead when using shared CF engines is significantly higher than for dedicated CF engines



Additional information

Coupling Facility Configuration Options: A Positioning Paper, GF22-5042-06

Parallel Sysplex Cluster Technology: The IBM Advantage, GF22-5015-10

System Managed CF Structure Duplexing, GM13-0103-05

Parallel Sysplex Configuration Planning for Availability (WSC Flash W98029)

Parallel Sysplex Configuration Vol 2: Cookbook, SG24-5638

- There should be an updated Vol 1, including Config info, in 1Q05

CFLevel info: <http://www.ibm.com/servers/eserver/zseries/pso/cftable.html>

CF Sizer: <http://www.ibm.com/servers/eserver/zseries/cfsizer/>

Parallel Sysplex Quicksizer - See your IBM rep



Second data center considerations



Objective

The objective of this session is to help you identify the considerations for running systems spanning multiple data centers, and especially multi-site sysplexes. The aim is to help you ensure that the project will deliver the expected benefits and meet your SLAs.

We will NOT give you all the answers - every situation is different, meaning that the right answer is different for everyone.

But we *will* highlight the things you need to worry about, and tell you where you can get more information about each.

Topics

WHY are you planning a second data center?

Remote copy considerations

Supported distances

Connectivity considerations

Performance considerations

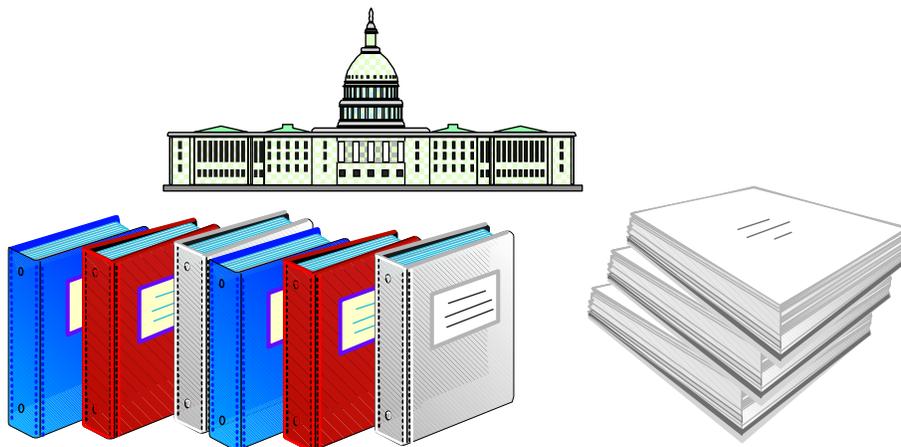
Systems Management considerations

Other sources of information



WHY do you want a second data center?

Disaster recovery



Government regulations for financial institutions



WHY do you want a second data center?

Continuous availability



A recent survey put data center outage costs at:

\$125,000/hr for pay per view TV service

\$2,600,000/hr for credit card validation

\$6,400,000/hr for online retail brokerage



WHY do you want a second data center?



You merged with another company and want to benefit from the data center investments



WHY do you want a second data center?



Software cost savings through aggregation of systems
in both sites into a smaller number of sysplexes



Remote copy considerations



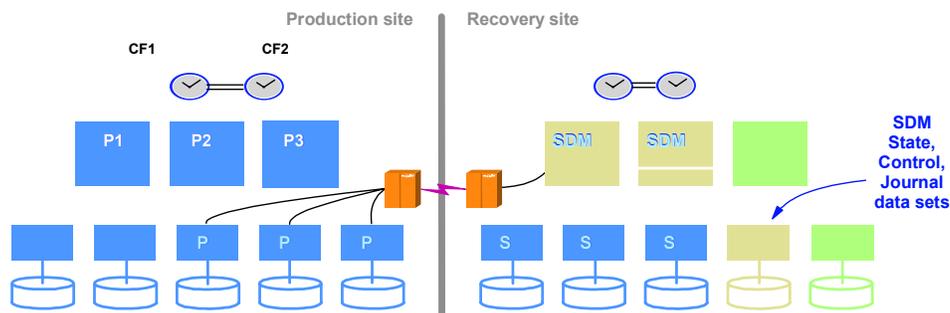
Remote copy considerations

The type of remote copy you can use (synchronous or asynchronous) depends on:

- **Recovery point objective:** how much data can you afford to recreate
- **Recovery time objective:** how long can you afford to be without your systems
- **Likelihood of a regional disaster** - is there a minimum distance that the second site must be from the first one
- **Is your objective solely disaster recovery, or do you need continuous availability as well?**



Asynchronous Remote Copy - XRC



Recovery Point Objective:

- Depends on many factors - distance, bandwidth, SDM capacity, CU type, etc
- Generally data in recovery site is less than one minute behind production site

Recovery Time Objective:

- Depends on level of automation, tested procedures, etc
- *Potentially* as low as less than 2 hours



Asynchronous Remote Copy - XRC

Pros:

- Negligible impact on response time
- Unlimited distance between primary and secondary sites
- XRC guarantees time consistency of remote DASD
- Supports very large configurations (1000s of volumes)

Cons:

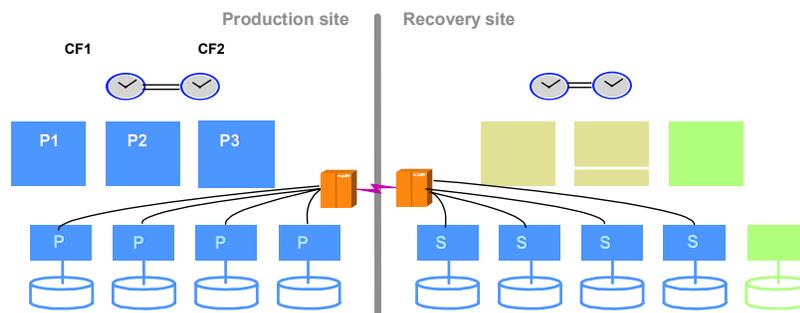
- Requires z/OS capacity to drive data movement (System Data Mover (SDM))
- Guaranteed that you will have to recreate some data

More information:

- XRC Planning and Installation Guide, GC35-0481
- XRC Reference Information for Advanced Users, GC35-0482



PPRC Global Mirror - Asynchronous PPRC



Recovery Point Objective:

- Depends on a number of factors - distance, bandwidth
- Data in recovery site might be as little as 3-5 seconds behind production site

Recovery Time Objective:

- Depends on level of automation, tested procedures, etc
- Potentially as low as less than 2 hours



PPRC Global Mirror - Asynchronous PPRC

Pros:

- Negligible impact on response time
- Unlimited distance between primary and secondary sites
- Guaranteed time consistency of remote DASD
- Removes requirement (and cost!) of an SDM

Cons:

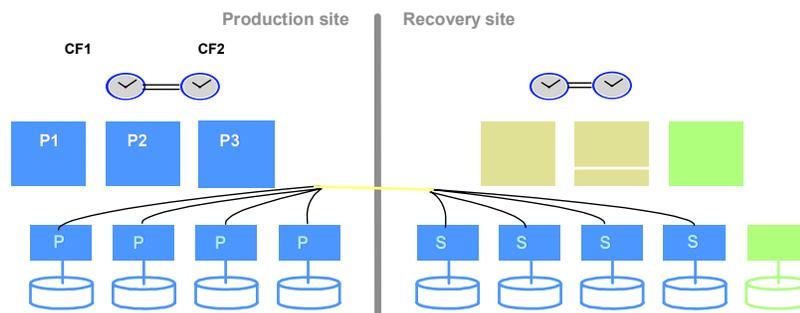
- Currently supports a maximum of 8 ESSs in total (Primary and secondary)
- Guaranteed that you will have to recreate some data
- Requires 3 copies of data - Primary, Secondary, FlashCopy

More information:

- Implementing ESS Copy Services on S/390 or zSeries Hosts, SG24-5680-04



PPRC Metro Mirror - Synchronous PPRC



Recovery Point Objective:

- As low as zero data loss

Recovery Time Objective:

- Depends on level of automation, tested procedures, etc
- Potentially as low as less than 1 hour



PPRC Metro Mirror - Synchronous PPRC

Pros:

- Potential for zero data loss
- Potential for continuous availability (HyperSwap)
- Supports very large configurations
- No requirement for an SDM

Cons:

- Will have an impact on response times
- Limited distance compared to asynch - up to 303km with ESS, less if both sites are to be in the same sysplex
- Possibly needs more bandwidth than asynch options

More information:

- Implementing ESS Copy Services on S/390 or zSeries Hosts, SG24-5680



Remote copy considerations

Question - how will you guarantee consistency of secondary DASD across multiple CUs?

- With XRC, the System Data Mover (SDM) ensures that data is applied to secondary DASD in consistent manner
- With Asynch PPRC, the PPRC microcode manages the data to ensure the secondary DASD are updated in a consistent manner
- With Synch PPRC..... you need something on top of PPRC, such as Freeze and automation



Remote copy considerations

Question - how will you decide that you have a disaster?

- If you continue to update the primaries after the remote copy relationship is broken, and this turns out to be a real disaster, you will have to recreate data, regardless of whether you are using PPRC or XRC
- This requires automation AND *prior* management decision on how to react to failure situations



Remote copy considerations

And don't forget that every time you break the relationship between the primary and secondary devices, the secondary devices contain inconsistent data while you resynchronize.....

Answer is to take a copy of the secondaries BEFORE you start to resync, so at least you have a consistent set of secondary DASD, even if they are aged. And if you want to be able to resynch in both directions, you need FlashCopy devices at both ends

AND..... don't forget that every PPRC secondary and every FlashCopy target device takes up a subchannel.....



Supported distances

OK, so what is the supported distance for synchronous remote copy and for asynchronous remote copy?

Answer: It depends (remember, this IS an IBM presentation!)



Supported distances

The maximum supported distance depends on what you want to do:

- Do you just want asynchronous remote copy or offsite tape vault?
– Is the moon far enough? See <http://www.transorbital.net/>



TRAILBLAZER®

THE FIRST COMMERCIAL VENTURE TO THE MOON: Transorbital® is the only private company to be authorized by the US State Department and NOAA for commercial flights to the Moon. The TrailBlazer® lunar orbiter will be the first delivery service to the Moon. Delivered to the Moon surface in a special capsule will be your certificates, business cards, cremated remains, jewelry, artwork and many other items of choice. The Trailblazer® satellite will deliver commercial and scientific projects and experiments to lunar orbit, as well as conduct lunar exploration and mapping.



Supported distances

Do you want synchronous remote copy, with all production CPCs in one site, and only backup CPCs in the remote location?

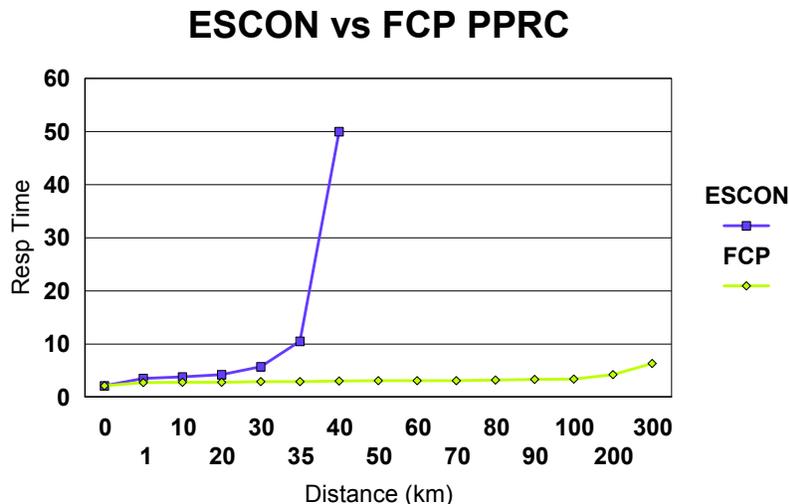
- Maximum distance is about 103 kms using ESCON sync PPRC, 303km with FCP sync PPRC
- Must carefully consider impact of long distances on primary DASD response time, especially if using ESCON for PPRC connections
- Does not provide continuous availability capability

Supported vs realistic

- Just because something is *supported* does not necessarily mean you can successfully implement it in your environment - you must check its validity for your configuration and workload



Comparison of ESCON and FCP PPRC Links



Sample configuration using 2105-800 doing 2100 I/O per sec and 2 ESCON PPRC or 2 FCP PPRC links, no PAV



Supported distances

Do you want a multi-site sysplex?

- Max distance is 100 kms (sort of)

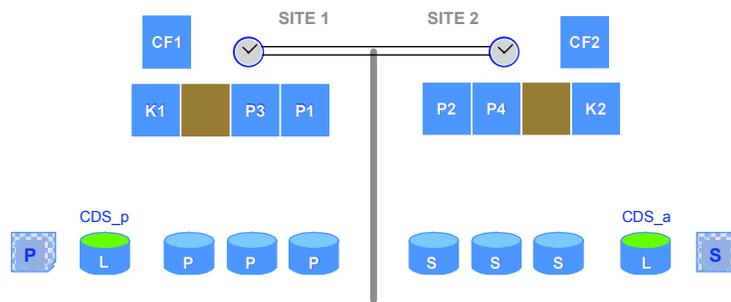
At very large distances, need to consider impact of distance on response times, even for "shamplexes". Some things (like GRS and XCF) are sysplex-wide and all systems will be impacted by the response time impact of large distances



Supported distances

Do you want to do multi-site sysplex data sharing?

- Realistic limit is *about* 10 km.
- THIS IS ONLY A RULE OF THUMB
 - The actual realistic distance depends on your workloads and how they are impacted by increased response times
 - Remember batch jobs as well as online - large batch update jobs are more likely to be impacted by longer response times

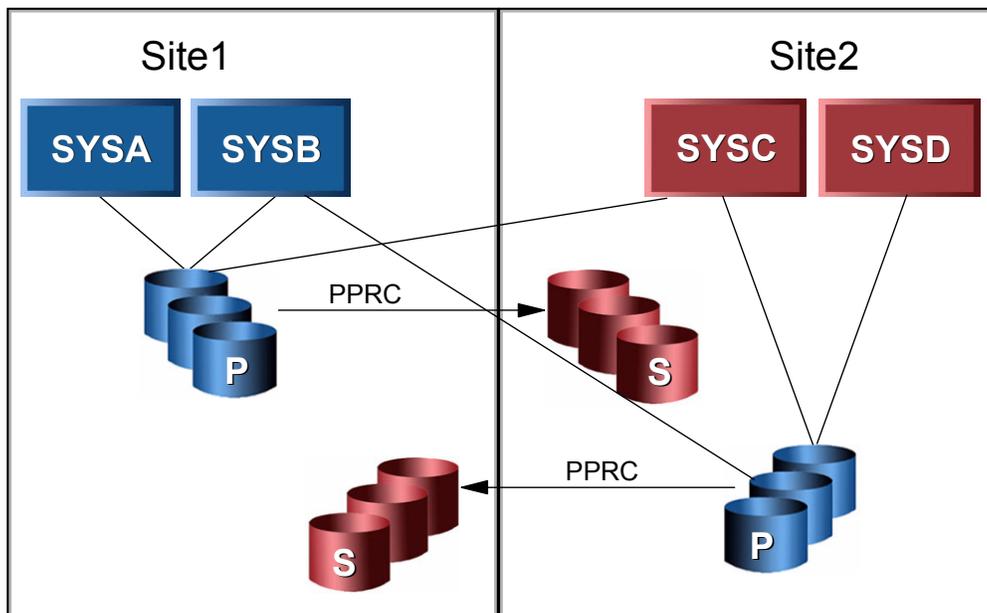


Supported distances

AND.... for any synchronous remote copy implementation, you need to consider where all your primary DASD will be.....

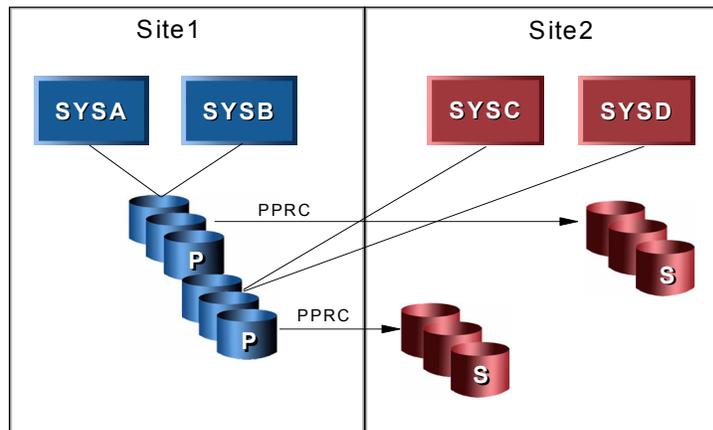


What's wrong with this picture?



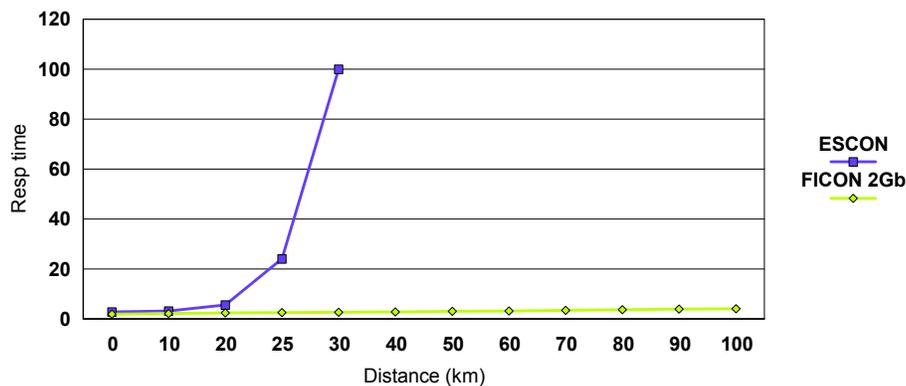
Supported distances

In a multi-site sysplex, and especially when doing multi-site DASD sharing, ALL the primary DASD must be in the same site. This means that some systems will have the cost of long connectivity *in addition* to the PPRC cost.



Comparison of ESCON and FICON CU connectivity

ESCON vs FICON CU Connections



Sample configuration using 2105-800 doing 2100 I/O per sec and 8 ESCON or 4 FICON channels



Bandwidth

The bandwidth **REQUIRED** depends on:

- The type of channel - FICON provides much more bandwidth than ESCON, for example
- FCP PPRC links provide more bandwidth AND more efficient use of that bandwidth compared to ESCON
 - ESS provides more efficient use of PPRC links (that is, better performance) than pre-ESS CUs.
 - PPRC links between an SSID pair must be ALL ESCON or ALL FCP
 - ESCON PPRC links can only operate in one direction, but FCP PPRC links can operate in both direction concurrently
 - In nearly all cases, two FCP links should provide acceptable performance and availability
- Use RMF Magic and Disk Magic to calculate DASD Write Rate MB/sec and estimate impact of various channel types and numbers



Bandwidth

More bandwidth considerations.....

- Are you going to use P/DAS or HyperSwap or similar function? Do you plan on swapping between DASD in the two sites?
 - If so, you need to plan for sufficient CPU to CU bandwidth, not just PPRC bandwidth
 - If you are going to use HyperSwap, that assumes that the systems will continue to run in Site1, requiring full connectivity to secondary DASD from Site1 CPCs.
- If you have to use the secondary DASD, where will the systems run - in the normal site or that site?
 - If in their normal site, you need production level bandwidth between the two sites
 - If in the second site, you need sufficient spare capacity AND connectivity from those CPCs to the secondary DASD
- Will you run systems in Site 2 off DASD in Site1?
 - If so, you need to plan for sufficient CPU to CU bandwidth from Site 2 back to Site 1



Bandwidth

More bandwidth considerations.....

- I/O Rate (how many per second)
- I/O Size (how long is the channel busy for each request)
- Performance characteristics of connected device
 - Tape can provide acceptable performance at higher channel utilizations than DASD
- If link is used for remote copy, the write intensiveness of the primary DASD is critical
- Need to consider single points of failure
 - Diverse routing
 - Two DWDMs
 - If you lose one DWDM/Path, can you continue operating on half capacity?
- Is the length of both paths similar and within related limits?



Bandwidth

The bandwidth AVAILABLE depends on:

- The connectivity you provide
 - Dark fiber:
 - ▶ Just ESCON and FICON directors?
 - ▶ DWDM
 - Telecoms lines
 - ▶ Channel extenders
- The cost!
 - Varies hugely from country to country.
 - Availability also varies - in some countries, dark fibre not available



Connectivity options



Connectivity considerations

Discuss **HOW** you can connect to remote devices (which technology can be used):

- Connectivity options
- CPU to DASD, Tape, Printers, 3x74, etc.
- CPU to Coupling Facility and Coupling Facility to Coupling Facility
- Sysplex Timer to Sysplex Timer, Sysplex Timer to CPU
- DASD to DASD and Tape to Tape (PtPVTs)
- XCF Signalling
- HMC LAN
- Consoles
- Network

I presume Parallel channels are no longer being used.....

(but if they are, remember that converters can't be used with FICON Bridge (FCV) channels)



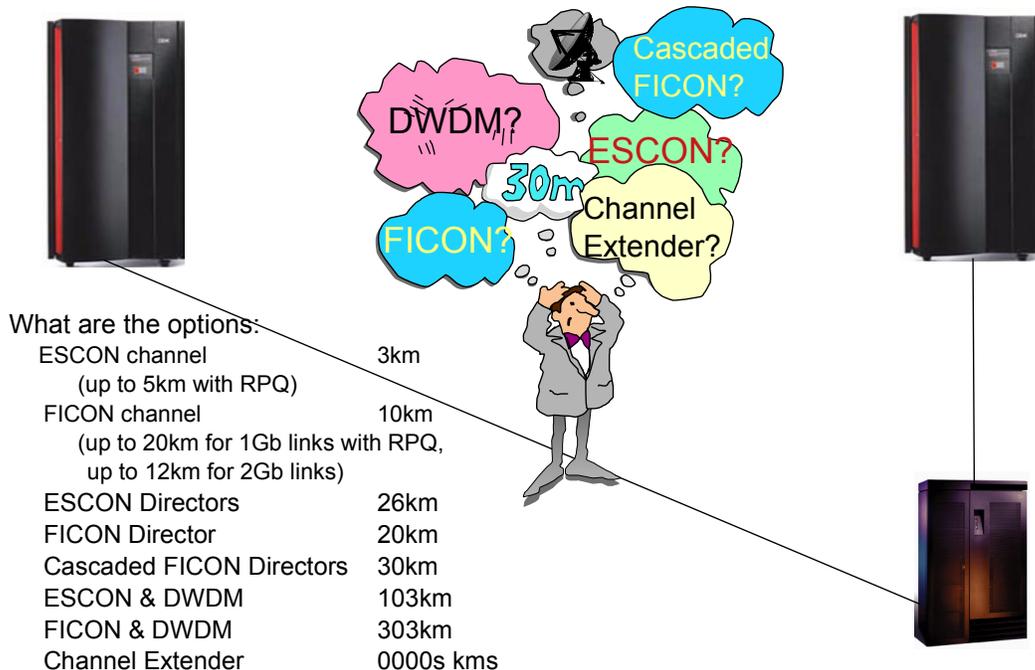
Connectivity considerations

Connectivity options

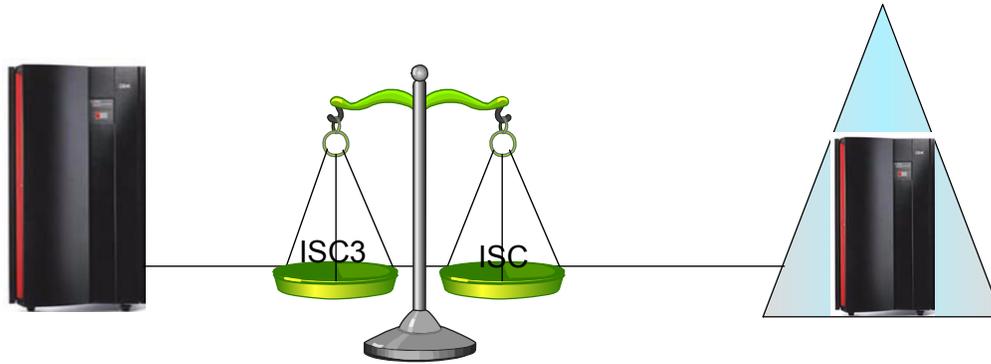
- There are a variety of ways to connect to a "device" over distances
 - Direct channel attach (point-to-point)
 - ESCON and FICON directors and switches
 - Use of repeaters such as IBM 9036
 - DWDMs
 - Channel extenders
- The distance supported for each option varies - by manufacturer, by device at the other end of the link, and especially by link quality
- You must contact the vendors to get the latest information and capabilities - this varies a lot by time and vendor. For example, one DWDM supports between 50 and 175km, depending on which features you select
- Not all devices perform acceptably over long distance
 - For example, ESCON CTC is terrible at 100 km, but ESCON VTS is fine at this distance



CPU to DASD, Tape, Cons controller, etc. connectivity



CPU to Coupling Facility



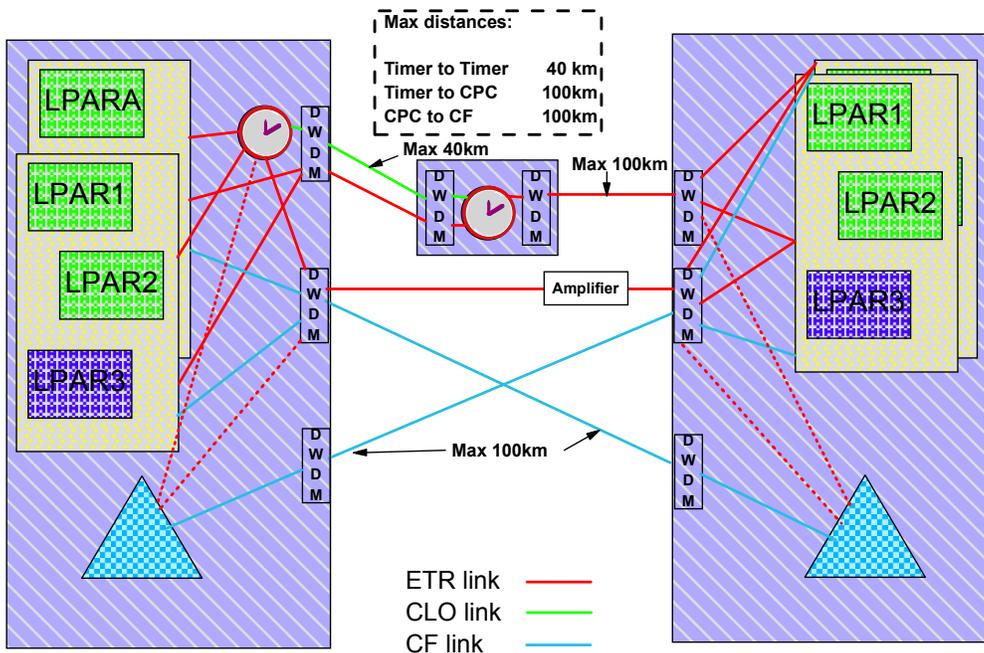
What are your options:

ISC3 peer mode (200MB)	10km (no repeater)
ISC3 peer mode RPQ 8P2197 (100MB)	20km (no repeater)
ISC compat mode (100MB) + DWDM + RPQ	40km
ISC3 peer mode (200MB) + DWDM + RPQ	100km

*See item RTA000174949 in viewblue for more information



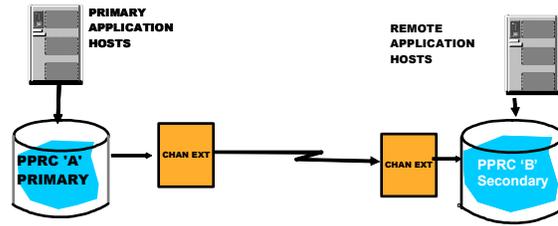
Sysplex Timer connectivity



*9036 Model 3 supports extending ETR signals, however this device is no longer marketed.



Remote copy links



Local Site

With PPRC, you can use:

- Switches, directors
- DWDMs
- Channel extenders (recent models make this more viable)

Remote Site

With XRC, you can use:

- Channel extenders (Most common)
- Switches, directors
- DWDMs

With VTS, you can use:

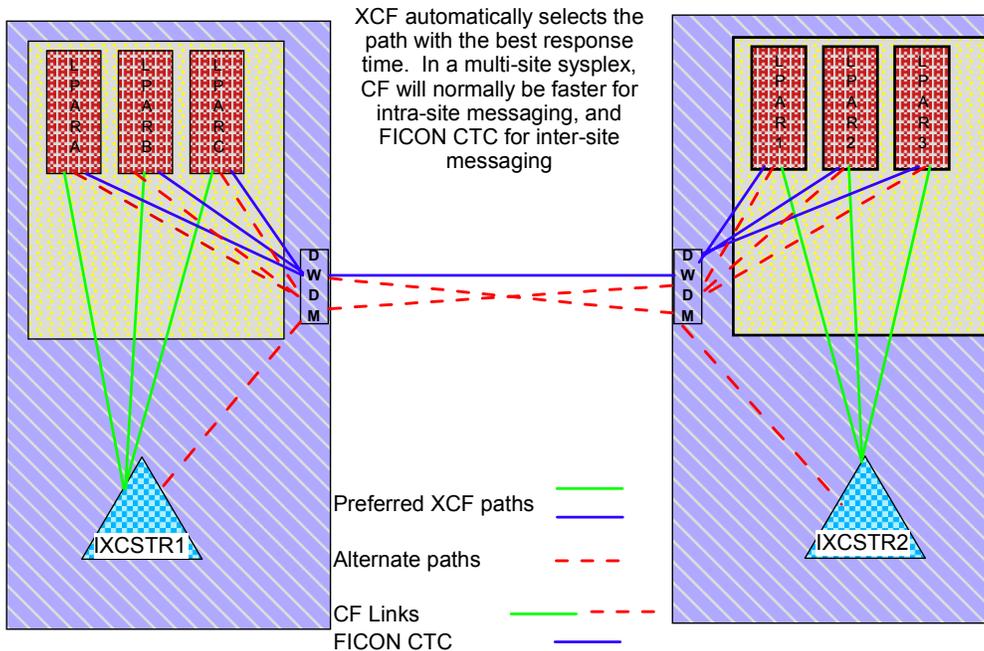
- Channel extenders
- DWDMs
- Switches, directors

DWDM requires dedicated fiber (dark fiber)

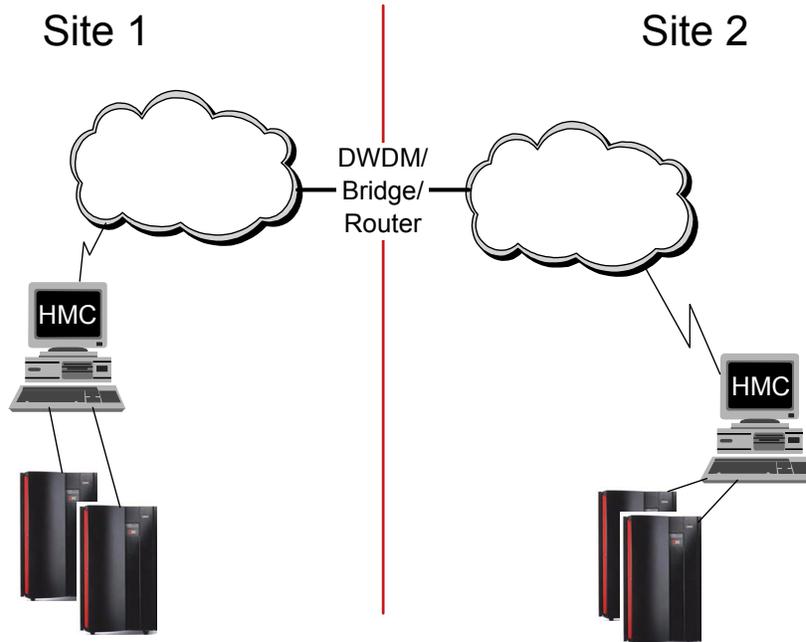
Channel extenders use leased lines and SONET



XCF Signalling



HMC LAN interconnect



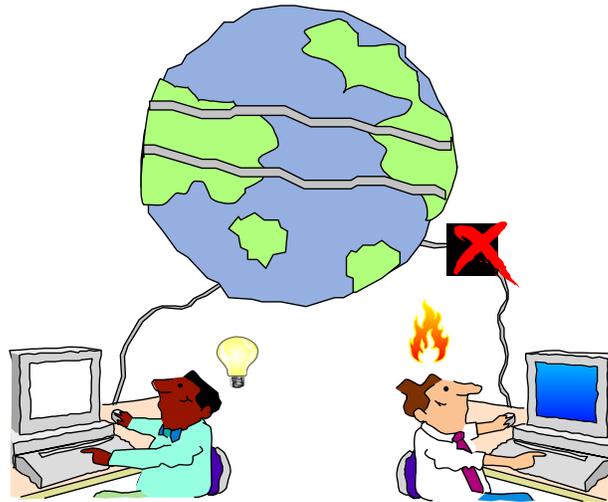
Console connectivity

Console connectivity:

- 2074
 - 2074s should be used to replace 3x74 and similar
 - 2074s are ESCON attached, so can attach via point to point ESCON, ESCON Director, or FICON Bridge
- OSA-ICC
 - Other alternative is OSA-Integrated Console Controller (OSA-ICC), available on z990, z890
 - Simply connect OSA-ICC port to LAN, and as long as you can access that LAN, you can bring up a console on any PC
- SNA Consoles
 - Real MVS console can be brought up on any PC that can access VTAM services
 - Cannot be used for NIP processing - only accessible after VTAM starts



Network connectivity



There is little point in having all the systems recovered and running in Site 2, if all the users are still connected to Site 1.....
In addition to connectivity, you also need automation and/or intelligent routing to switch your users over....



Connectivity summary

Device / Connection method	Direct (1)	Director / Switch (1)	DWDM (1)	Channel Extender (1)
CF Links	Yes	No	Yes	No
Sysplex Timer	Yes	No	Yes	No
CPU to CU	Yes	Yes	Yes	Yes
XRC	Yes	Yes	Yes	Yes
PPRC	Yes	Yes	Yes	Yes (2)
CTCs	Yes	Yes	Yes	Yes
HMC LAN	Bridge/Router	No	Yes	No
Console	Yes	Yes	Yes	Yes

1) Distances supported vary by device type

2) Only on newer channel extenders, need to carefully investigate performance



Performance considerations



Performance considerations

Performance is dependant on:

- **Distance.**
 - Degradation is not linear
- **Technology.**
 - ESS uses ESCON for PPRC much more efficiently than RVA
 - FICON is hugely better over distance than ESCON
- **I/O Rate.**
 - High channel utilization is much more painful at long distance, AND, long distance (via long response times) drives up channel and UCB utilization
- **Bandwidth.**
 - You may not be able to afford to provide as much bandwidth to a remote site as you would within the computer room
- **Remote copy.**
 - Synchronous remote copy may require more bandwidth than asynchronous



Performance considerations

Just because something is technically possible, does NOT mean that it can be made to work in your environment at an affordable cost:

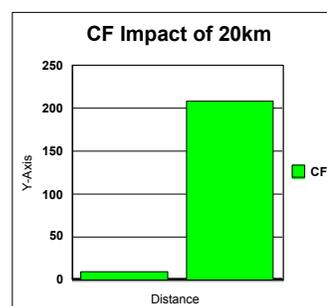
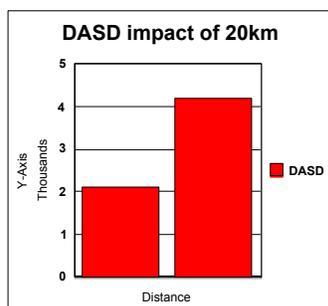
- Long distance between CPU and connected control units impacts response times
 - In the example earlier, adding 20 km doubled ESCON response time
 - You MUST use a tool like Disk Magic to project actual anticipated response times
- Long distance between primary and secondary CU impacts response times (for synch PPRC) or amount of data that needs to be recreated (for async PPRC or XRC)
 - In the example earlier, adding 20 km doubled response time when using ESCON PPRC links
 - Once again, use a tool like Disk Magic to project response times
- Newer technology (FICON and FCP), new features (PAV), and more bandwidth/adaptors can reduce impact of distance to some extent



Performance considerations

Also have to consider Coupling Facility impact:

- Consider that "good" CF response times are about 100 times faster than good DASD response times
- Distance adds a fixed amount - 10 microseconds per km - to response times. Therefore, impact of distance on CF is relatively much higher than on DASD



Performance considerations

Consider:

- Who is using the CF
- Synch/asynch algorithm in z/OS 1.2 will cause long sync requests to become asynch, significantly increasing response time for those requests. This change limits the CPU cost of high response times, but this does not help applications using the affected structures
- Peer mode links should be considered a must for longer distances
- It *is* possible to have many requests running against a given structure at the same time - access is not serial
 - One customer is doing 60,000 lock requests per second at 5km distance and about 150 mics. response time
- What is the change in lock response time compared to lock hold times?



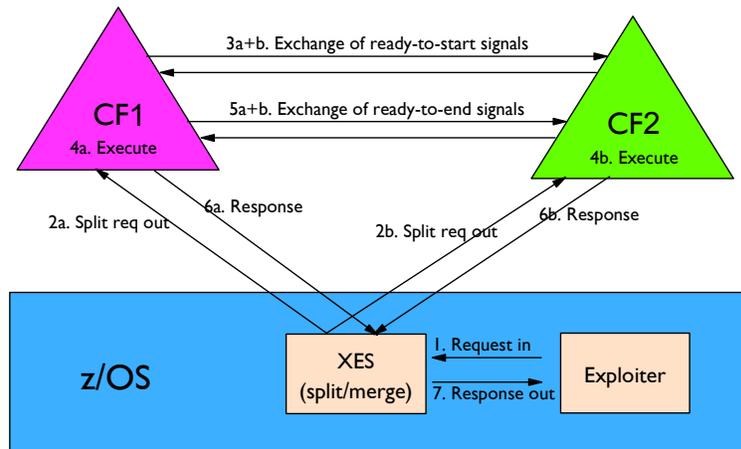
Performance considerations

CF Links:

- Subchannels are busy for the whole time of the CF request. For asynch requests, it is still for the whole of the response time
- As distance increases, response time goes up, driving up subchannel utilization, and making response times even worse
- To offset this, add CF links and carefully monitor subchannel utilization
 - Use Parallel Sysplex Quicksizer (SPSSZR) to project impact of distance on CF Link utilizations



Performance considerations



System Managed CF Duplexing especially does not like long distances

- Every km travelled adds 10 microseconds to response time
 - ▶ Distance from CPC to furthest CF * 10 microseconds PLUS
 - ▶ Distance between CFs * 4 * 10 microseconds
 - ▶ If one CF is 20 km away, impact would be approx 1000 microseconds....



Automation, processes, and testing



Systems Management considerations

Planning for the hardware and software is only the start

- MUST wrap this in automation, to get the speedy response and consistency that is required
- Ideally automation will handle everything from identifying the disaster trigger through to getting all your systems up and running in the recovery site
 - Automation should also manage and provide operations interface to Remote Copy mechanism and status
- You must test until you get all this working. Then keep on testing to ensure it continues working.
- Give consideration to placement of operations - some installations place ops in a 3rd site.
- If building a new site, don't forget to allow for fitting-out time after the building is complete, but before you can use it.



Systems Management considerations

What else.....

- Expect the unexpected - problems often arise from "trivial" things, like cabling, power supplies, physical access, and so on.
- Calculate your worst case elapsed time to set everything up - then double it! Many of the things you encounter will be outside your control and may take time to rectify.
- The devil is in the details.... for example, while direct attachment from DWDM to CU is possible, some CUs require attachment to a switch because the CU does not provide sufficient credits for the distance to the CPU



More information

IBM Redbooks:

- [IBM eServer zSeries Connectivity Handbook, SG24-5444 \(new version due this October\)](#)
- [Implementing ESS Copy Services on S/390, SG24-5680-04](#)
- [Planning for IBM Remote Copy, SG24-2595](#)
- [IBM Enterprise Storage Server, SG24-5465](#)
- [A Disaster Recovery Solution Selection Methodology, REDP3847](#)

IBM Product Manuals:

- [Advanced Copy Services](#)
- [XRC Planning and Installation Guide](#)
- [XRC Reference Information for Advanced Users](#)



More information

IBM Service offerings

- [Bandwidth studies - contact local storage specialist](#)
- [IBM Disk Magic tool projects impact of distance, remote copy, and connectivity type](#)
- [eRCMF - Does not support zSeries](#)
- [RCMF - Remote Copy Management Facility](#)
- [GDPS - Search IBM Web site for GDPS for White Papers, plus more detailed GDPS Introduction Redbook coming out later this year](#)



More information

GDPS-qualified DWDMs and Channel Extenders

- Adva
- Cisco
- Nortel
- InRange
- CNT
- Ciena
- See individual vendor's Web site for information on these devices.
 - Later this year, there will be a series of RedPapers on qualified DWDMs, with links to all the vendor's Web sites.
- See <http://www.storage.ibm.com/disk/ess/supserver.htm> for more information on devices that are tested with ESS



Summary

This is a very complex subject, requiring skills in many areas

Technology is changing daily

Before anything else, need to clearly identify the objective:

- Is it Continuous Availability?
- Is it Disaster Recovery?
- Is it SW cost savings?

Involve the experts from the very beginning

Results can be stunning, IF the project is properly planned and managed



z/OS HealthChecker



z/OS HealthChecker

In order to ensure high availability and optimum performance, it is vital that "best practices" are adhered to AND that a properly configured system remains so.

The z/OS HealthChecker helps you identify aspects of your configuration that do not conform to what IBM believe are the best ways to configure your system/sysplex. The list of features the HealthChecker looks for is based on analysis of multi-system outages.

Currently, the HealthChecker is available as a Web download from:

<http://www.ibm.com/servers/eserver/zseries/zos/downloads/>



z/OS HealthChecker

So, what is the current HealthChecker?

- ▶ Code that checks about 30 aspects of system operation - based on actual outage analysis data
- ▶ Highly parameterized, allows easy updates of "best practices" values
- ▶ Users can override best practices values, but must provide a reason and timestamp - if corresponding IBM best practice is updated, user must revalidate his override
- ▶ Runs on OS/390 2.10 and later
- ▶ Runs as a batch job.
- ▶ The latest available version is 3.01 and was released in October 2003. To date, there have been over 2700 downloads.
- ▶ Customer feedback is universally positive



z/OS HealthChecker

Version 1 was made available on February 6, 2003

- Coupling facility structure attributes
- XCF transport class verification and XCF cleanup values
- CF Structure locations compared against defined preferences
- Sysplex, CFRM, LOGR couple data set separation
- Sysplex, EMCS, and MVS console definitions
- z/OS UNIX Automove definitions
- Available frame queue thresholds
- Real storage settings
- Reconfigurable storage settings

See WSC Flash 10213



z/OS HealthChecker

Version 2 was made available on April 29, 2003

- **Additional support and new checks**
 - Write-to-operator (WTO) message support
 - Severity designation for checks
 - Virtual Storage checks and mapping
 - New data set (HC DATA) required □

See WSC Flash 10225



z/OS HealthChecker

Version 3 was made available October 4, 2004

- **Linklist checks**
 - Data sets in the link list have secondary space allocated
 - New secondary extent created in a current linklist data set data set
- **APF data sets exist on the volume defined in the APF list**
- **Synchronous reserve processing enabled**
- **Duplicate member names in linklist and LPA list data sets**
 - Disabled by default
- **Customer control over highlighting of exception messages**

See WSC Flash 10261



z/OS HealthChecker

HealthChecker mini-survey

- ▶ Is there anyone that has NOT downloaded the current version?
- ▶ Is there anyone that does NOT run the HealthChecker on all systems in your installation?
- ▶ Is there anyone that does NOT run this on a regular (at least weekly) basis?



z/OS HealthChecker

So, where do we go from here:

- Customers requested a more formal product:
 - We are making the HealthChecker part of the operating system
- Customers requested checks from more z/OS components and more IBM products:
 - We are adding more IBM-provided checks
- Customers requested ability to add their own checks:
 - We are opening up HealthChecker services so the HealthChecker infrastructure can be used by customer and vendor-written checks
- Change the HealthChecker from a batch job to be a constantly running started task
- Provide an open interface that can view the output from any check
- Provide documentation to assist in setting up and running the HealthChecker and in writing your own checks



z/OS HealthChecker

Integrating HealthChecker into the operating system:

- Preview for z/OS 1.7 states that HealthChecker will be shipped as part of 1.7
- HealthChecker will have its own FMID in 1.7
- Prior to general availability of z/OS 1.7, it is hoped to make the new HealthChecker available via Web download. This version will run on z/OS 1.4 and later*

* Console checks will require z/OS 1.4.1



z/OS HealthChecker

Providing more IBM-written checks:

- All the checks in the existing HealthChecker will be carried forward to the new version and updated as appropriate
- Ownership for identifying, developing, and maintaining the checks is assumed by the developers of the respective components
- Examples of the type of checks that are being investigated for the future include:
 - JES2
 - Virtual and real storage management
 - More USS-related checks
 - SMF checks
 - RACF
 - XCF and XES
 - More console checks



z/OS HealthChecker

Opening up HealthChecker infrastructure:

- The next version of the HealthChecker is broken into a Framework part and individual checks.
- The Framework part:
 - Will run as a started task
 - Provides callable services to schedule individual checks
 - Provides messaging, logging, and reporting services for checks
 - Provides an operator interface to display and control individual checks
 - The Framework includes documentation to help check writers



z/OS HealthChecker

Opening up HealthChecker infrastructure:

- And the individual checks:
 - Checks are provided independently of the Framework
 - Checks will be delivered via PTF, meaning that new checks can be delivered without having to wait for a new z/OS release
 - ▶ Currently investigating ways to make it easy to identify the PTFs that deliver new checks
 - Checks can be added dynamically
 - Checks can be grouped, so you can enable or disable groups, depending on your requirements
 - Support is provided to override the default values provided with the checks
 - ▶ Defaults can be overridden via a Parmlib member, Modify command, or via the CK command in SDSF



z/OS HealthChecker

Running the HealthChecker as a started task:

- In order to provide an operator interface, and scheduling services, the HealthChecker services must be available continually. Therefore, HealthChecker will now run as started task rather than a batch job.
- Setup work is trivial - objective is to keep installation as simple as it is for the current HealthChecker
- Resource used by started task will be minimal - depends on the checks and how frequently they will be scheduled



z/OS HealthChecker

Providing more flexible interface for viewing output:

- Output from the first version of the HealthChecker could be viewed in the spool, or through msys for Ops or SA/390
 - However, the SA/390 interface was based on specific knowledge about each check.
- In order to provide the ability to view the output from more checks without the viewer having to understand the individual checks, HealthChecker has been changed
- SDSF in z/OS 1.7 will exploit this new capability, allowing the output from any HealthChecker check to be viewed from a new CK panel in SDSF.



z/OS HealthChecker

Using new CK command in SDSF to get a list of checks...

```
SDSF HEALTH CHECKER DISPLAY (ALL) LINE 1-3 (3)
COMMAND INPUT ==> SCROLL ==> CSR
NP OWNER Status Interval Severity Start-Time
XCF XCF Initializing 48:00 MEDIUM 22:00:00
XCF XCF Initializing 48:00 MEDIUM 22:00:00
GRS GRS Running 48:00 MEDIUM 22:00:00
```



z/OS HealthChecker

Displaying the results of a specific check...

```
SDSF OUTPUT DISPLAY ALTERNATE_CONSOLE_GROUPS LINE 0 COLUMNS 02- 81
COMMAND INPUT ==> SCROLL ==> CSR
***** TOP OF DATA *****
*Medium severity Exception: IBM Criteria not met*
ALTERNATE_CONSOLE_GROUPS
The following consoles have no alternate group (ALTGRP) defined:
Console Console Console Active
ID Name Type System
1 POSIXCON MCS SY1
2 PLEXSY2 MCS (Inactive)
3 BARCON1 MCS (Inactive)
IBM suggests that alternate groups be defined to increase availability if
there is a console failure. MVS can then switch to another console. MVS
searches for the first available console based on the order of the
console members defined for the alternate console group. Alternate
groups help to avoid single points of failure. Note that IBM does NOT
suggest use of the ALTCONS facility.

Action: To define an alternate group, use the ALTGRP keyword of the
```



z/OS HealthChecker

HealthChecker documentation:

- The new version of HealthChecker will be provided with two manuals:
 - z/OS HealthChecker Developers Guide. Provides information about planning and developing the check, processing messages, controlling via the HZSPRMxx member, and testing the checks
 - z/OS HealthChecker User's Guide. Provides information to set up the HealthChecker started task, and a guide to messages issued by the Framework
 - ▶ For z/OS 1.4-1.6 and later, this guide will be the only document describing the checks. In 1.7 and later, checks will be described both here and in the component's own manuals (GRS Planning, for example).



z/OS HealthChecker

HealthChecker summary:

- There is significant value in running the current version of the HealthChecker:
 - Especially, once you get it running cleanly, schedule it to run regularly and only notify you of non-zero return codes
- The new HealthChecker will build on the value and ease-of-installation delivered by the current one, create a flexible base for further exploitation, and deliver more checks based on IBM's ongoing experiences
- The new HealthChecker will be available in 1.7, and hopefully in advance as a Web download. Future additional checks will be SMP-installable and may also be available as downloads. The new HealthChecker will run on z/OS 1.4 and later



Bits and bytes



GDPS Update



GDPS Update

Quick review of GDPS Offerings

Review of HyperSwap

Review of Open LUN support

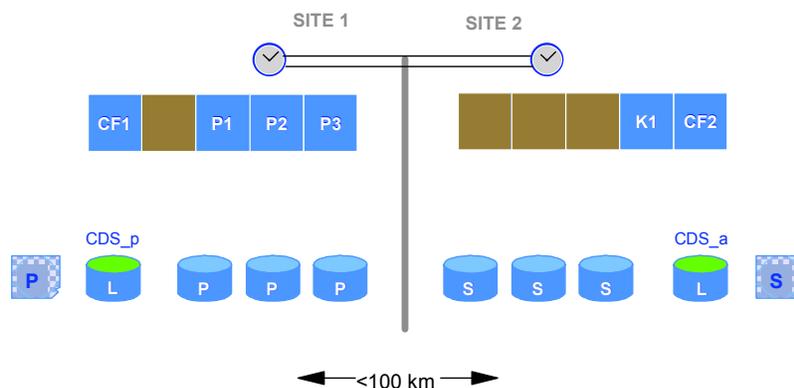
Review of enhancements to GDPS in 2004

Future directions



GDPS/PPRC - Single site

Single Site Workload (cross-site sysplex)

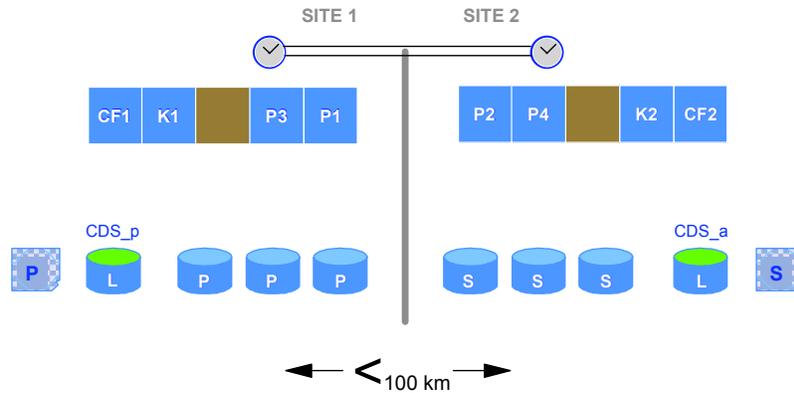


- ◆ One Sysplex, all production systems in site 1
- ◆ Controlling System in site 2
- ◆ Expendable workload and/or CBU capable processors in site 2
- ◆ Primary disks / tape in site 1; secondary disks / tape in site 2
- ◆ Site recovery by restarting failed system images



GDPS/PPRC - Multi-site

Multiple Site Workload (cross-site sysplex) Continuous Availability Configuration

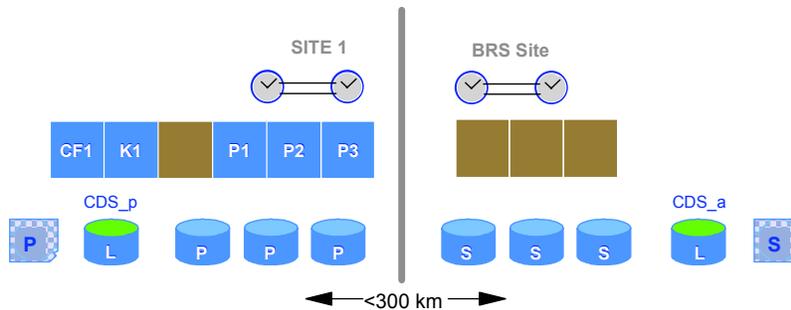


- ◆ One Sysplex, production systems in **both** sites, sharing data
- ◆ Controlling System in site 2 (at least, one in each site is preferable)
- ◆ Expendable workload and/or CBU capable processors in both sites (optional) or site 2
- ◆ Primary disks / tape in site 1; secondary disks / tape in site 2
- ◆ Continuous Availability configuration
 - ◆ Applications cloned and exploiting data sharing across 2 sites



GDPS/PPRC - BRS Config

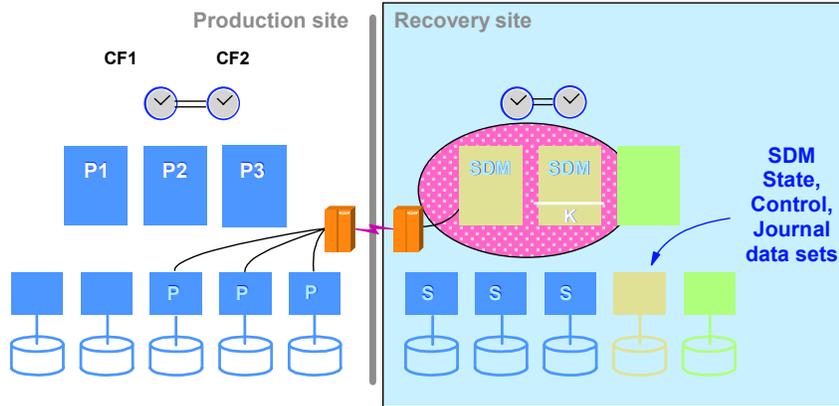
Single Site sysplex (BRS configuration)



- ◆ Configuration can be used if:
 - ◆ Do not own a second site, but a D/R facility exists within maximum PPRC distance
 - ◆ Cannot install dark fiber between sites
 - ◆ Need to minimize cross-site link connectivity
 - ◆ Increased RTO is acceptable
- ◆ One Sysplex with production systems and controlling system in site 1
- ◆ Site 2 can be customer site or third part D/R facility
- ◆ Primary disks / tape in site 1; secondary disks / tape in site 2
- ◆ Increased RTO compared to cross-site sysplex configurations
- ◆ K1 needs to be IPLed in site 2 when site 1 has a failure



GDPS/XRC



- ◆ Production system can be no, Base or Parallel Sysplex
- ◆ Common time reference required in site 1
- ◆ Primary data in Production site, secondary data in Recovery site
- ◆ System Data Mover(s) and Controlling System in site 2
- ◆ System Data Mover(s) must run in Base or Parallel Sysplex
- ◆ Expendable workload in site 2 and/or CBU capable processors



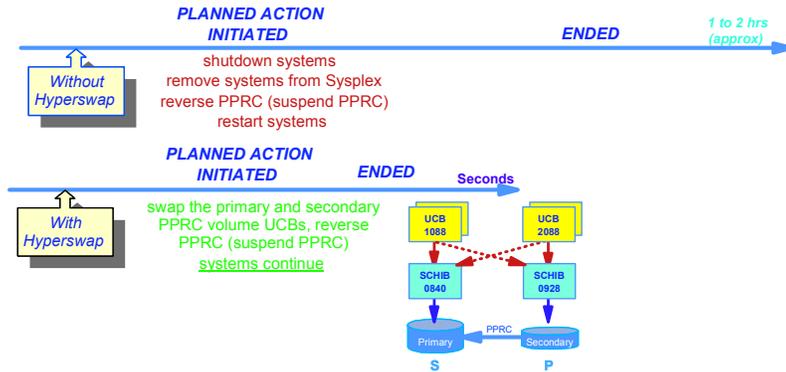
GDPS Offering comparison

Offering / Feat	GDPS/PPRC	GDPS/SM	GDPS/XRC	RCMF/PPRC	RCMF/XRC
Cont Avail	Y (with HS)	No	No	No	No
Disaster Recovery	Yes (Freeze)	Yes (Freeze)	Yes (SDM)	No	Yes (SDM)
FlashCopy Support	Yes	Yes	Yes	No	No
Prod Sysplex Automation	Yes	No	No	No	No
Supported Distance (in km)	100 (300 in BRS config)	100 (300 in BRS config)	1000s	300	1000s
PtPVTS	Yes	No	Yes	No	No
Span of control	Both sites	Both sites (DASD only)	Recovery site	Remote Copy only	Remote Copy only
Open LUN	Yes	Yes	No	No	No
Heterogeneous support	Yes	No	No	No	No
Monitoring and alerting	Yes	Yes	Yes	No	No
Automation scripts	Yes	No	Yes	No	No



GDPS/PPRC HyperSwap

Planned Disk Reconfiguration with HyperSwap

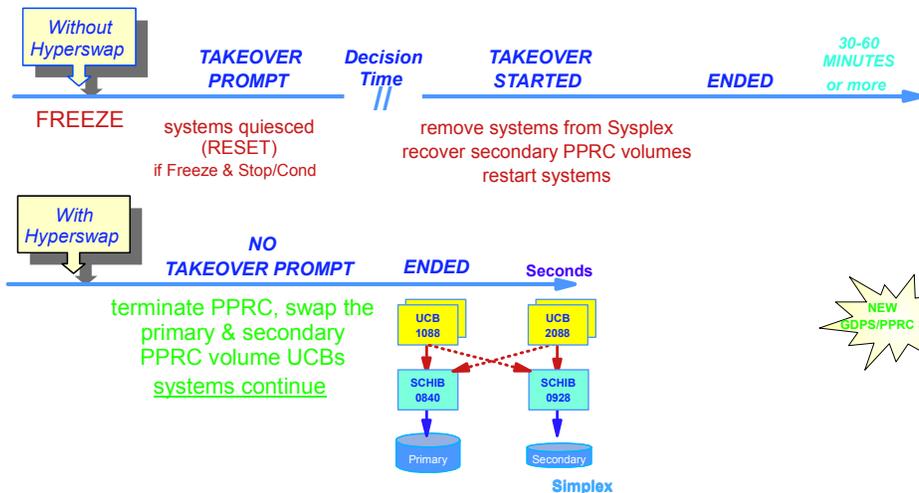


Reference Customer	n-way Sysplex	Disk configuration	Switch Time (note 1)
ARZ (Austria)	6	2300 vol pairs (14 TB)	82-84 secs
Postbank (Germany)	8	1800 vol pairs (32 TB**)	80-84 secs
iT Austria (Austria)	3	650 vol pairs (12 TB**)	32-36 secs
iT Austria (Austria)	5	4200 vol pairs (24 TB)	75 secs
iT Austria (Austria)	11	4500 vol pairs (76 TB**)	75 secs
A N Other	10	4000 vol pairs	70-87 secs
		** 3390-9 device type volumes	Note 1: Failover/Failback expected to reduce this time by 50% - expected end '04



GDPS/PPRC HyperSwap

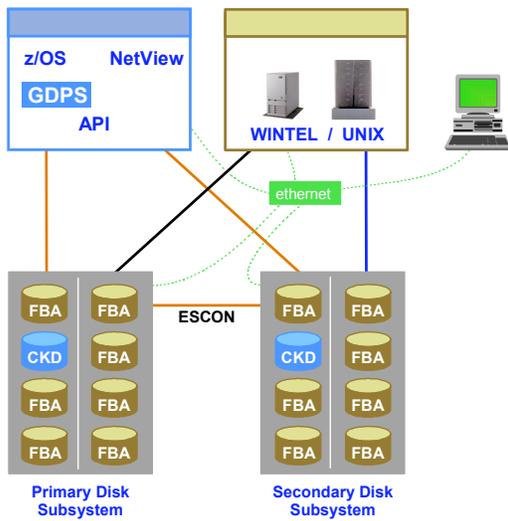
Unplanned Disk Reconfiguration with HyperSwap



Tests at GDPS Solution Center, Montpellier France:
 Configuration of ESS disks; 6000 volume pairs; 40 LSSs
 Unplanned disk reconfiguration 23 to 25 secs



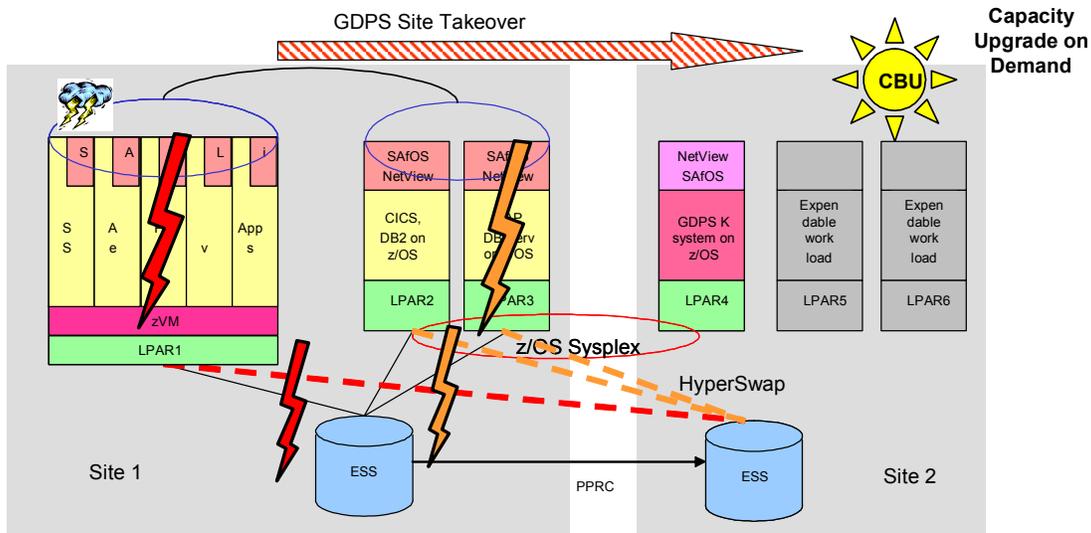
GDPS/PPRC Open LUN support



- ◆ Extends GDPS/PPRC technology to manage distributed applications across multiple platforms
 - ◆ z/OS and open systems data (Unix, NT, Linux)
- ◆ GDPS/PPRC running in a z/OS system manages the PPRC status of devices that belong to the other platforms
- ◆ Provides data consistency across both z/OS and/or open systems data when failures occur
- ◆ Requires
 - ◆ Some CKD capacity in disk subsystem
 - ◆ PPRC level 4
- ◆ Support details
 - ◆ Supports x-platform or platform level Freeze
 - ◆ FlashCopy not supported for Open
 - ◆ No GDPS Code running on Open Systems host - suspend reported through SNMP alert
 - ◆ Manual restart of Open systems required



GDPS Multiplatform Resiliency for zSeries



- ◆ Support for Linux for zSeries running under z/VM
- ◆ PPRC and HyperSwap support
- ◆ Ability to automatically IPL z/VM or Linux guest following a failure



GDPS Update

What's new for GDPS in 2004:

- Support for FCP PPRC links
- Support for FlashCopy II ability to have source and target in different LSSs
- Complete replacement of AO Manager with BCP Internal Interface, even for duplicate volser support
 - AO Manager product no longer supported, but interface IS supported for other products
- PPRC Link monitoring
- Now possible to use HyperSwap with CA-MIM
- GDPS now pre-reqs SA/390 V2.1 or later*
- GDPS now includes a Message Flood Automation tool to address the huge number of msgs that can arise during failover processing

* SA/390 2.1 support ends 10/31/05



GDPS Update

What's in the crystal ball:

- Support for PPRC Global Mirror
 - Significant new offering, with aspects of both the existing GDPS/PPRC and GDPS/XRC offerings
 - ▶ Does not provide an CA capability (no HyperSwap)
 - ▶ Does not require an SDM in the recovery site
 - ▶ Does support unlimited distances
 - ▶ Does support Open LUN
 - ▶ Is asynchronous like XRC, so does not offer zero data loss capability
 - ▶ Not as scalable as GDPS/XRC
- Support for continuous data availability within a single site/campus, exploiting Parallel Sysplex and HyperSwap
- Ability to have a PPRC primary device act as a FlashCopy target



z990 considerations



z990 migration - sysplex considerations

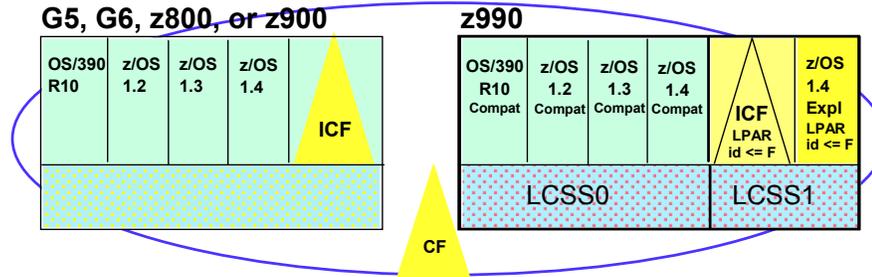
The following are the **ONLY** releases that can *coexist* in a sysplex with a z990:

- OS/390 R10 (only if running on 9672 G5 or later) (EOS 9/30/04)
- z/OS 1.2 (EOS 10/31/04)
- z/OS 1.3 and z/OS.e 1.3
- z/OS 1.4 and z/OS.e 1.4
- z/OS 1.5 and z/OS.e 1.5
- z/OS 1.6 and z/OS.e 1.6 (zSeries only)
- z/OS 1.1 is **NOT** supported ANYWHERE in the sysplex if the sysplex contains any image on a z990.

z990 migration - sysplex considerations

Compatibility NOT required

Compatibility (or Exploitation)
required on all z990 images



- Compatibility not required on non-z990 images in Parallel Sysplex when:
 - ▶ Coexisting z990 Compatibility (or Exploitation) operating system images have LPAR ID ≤ 15 (x'F')
 - ▶ z990 CF LPAR ID ≤ 15 (x'F')
 - LCSS of CF LPAR irrelevant

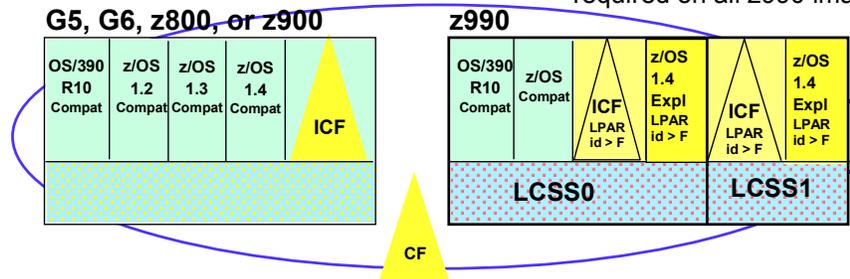
Recommendation: Rollout Compatibility level software to all images in the Sysplex as soon as possible



z990 migration - sysplex considerations

Compatibility IS required

Compatibility (or Exploitation)
required on all z990 images



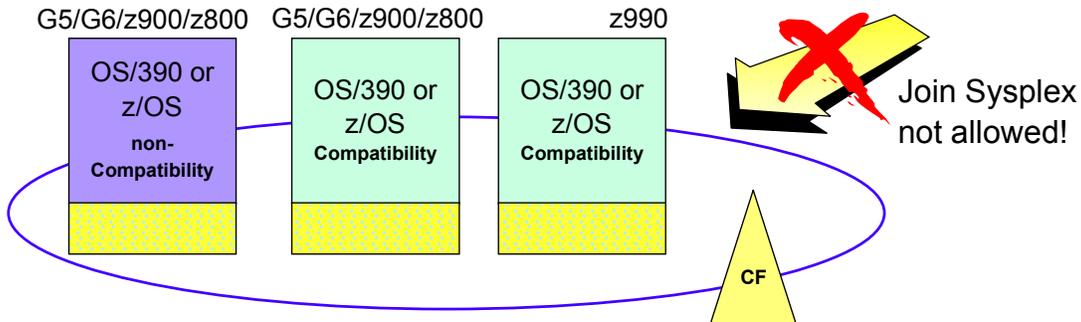
- Compatibility is required on non-z990 images in Parallel Sysplex when:
 - ▶ z990 CF LPAR ID > 15
 - LCSS of CF LPAR irrelevant
 - ▶ z990 z/OS 1.4 with Exploitation feature LPAR ID > 15
 - LCSS of z/OS 1.4 Exploitation image is irrelevant

Restriction: Can not assign LPAR ID > 15 to CF LPARs until Compatibility software rolled out across all images in the Sysplex



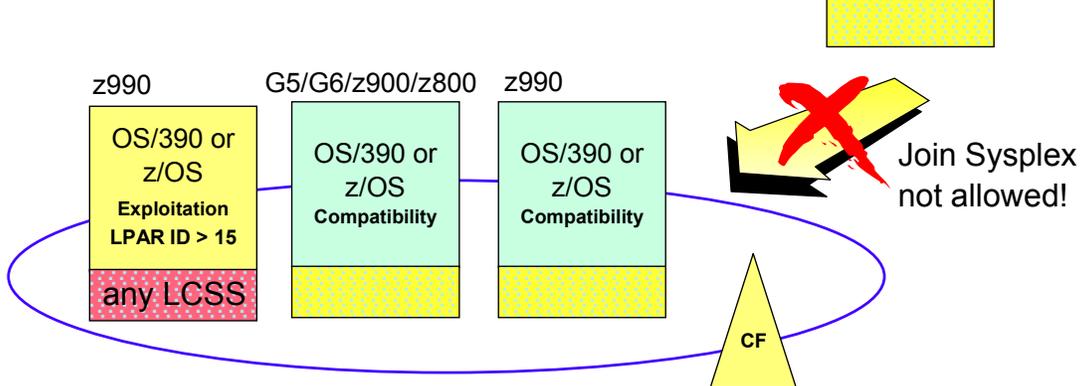
z990 migration - sysplex considerations

- OS/390 and z/OS Compatibility, and z/OS Exploitation, code provide the following protection
 - z/OS with LPAR ID > 15 will not join sysplex unless all images in the sysplex are running on at least Compatibility code

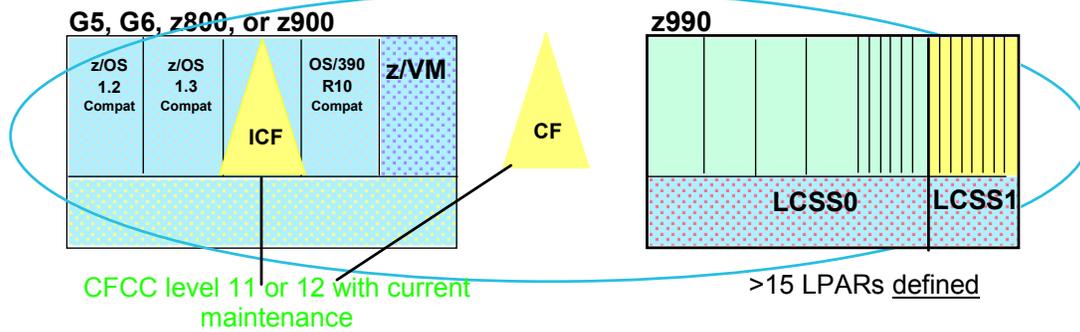


z990 migration - sysplex considerations

- OS/390 and z/OS Compatibility, and z/OS 1.4 Exploitation, code provide the following protection
 - A system not running on at least Compatibility code will not be able to join a sysplex that already has a member with LPAR ID > 15
 - The system will actually join but will immediately be removed (wait state code of 0A2/154)



z990 migration - sysplex considerations



- ▶ Required non-z990 CFCC Levels
 - G5, G6 CFs must be CFCC 11 Service Level 5.04
 - z800, z900 must be CFCC 12 Service Level 10.05
- ▶ For some workloads, using G5 CFs in a Parallel Sysplex with z990 is not recommended
 - Use only as a temporary migration step

Recommendation: Roll out non-z990 CFCC code to support z990 as soon as possible.



z990 sysplex considerations

When defining a CF in the CFRM policy:

- If the Coupling Facility resides on a non-z990 processor, then the partition number specified in the CFRM policy is the same as the partition number defined in HCD.
- If the Coupling Facility resides on a z990 processor, then the partition number specified in the CFRM policy should be the logical partition number specified in the PR/SM image profile.
- When moving a Coupling Facility to a z990 processor, remember to update the CFRM policy.

```

DATA TYPE(CFRM) REPORT(YES)
DEFINE POLICY NAME(POLICY1)
REPLACE(YES)

CF NAME(CF01)
TYPE(123456)
MFG(IBM)
PLANT(02)
SEQUENCE(123456789012)
PARTITION(01)
DUMPSPACE(2000)

CF NAME(CF02)
TYPE(123456)
MFG(IBM)
PLANT(02)
SEQUENCE(123456789012)
PARTITION(11)
DUMPSPACE(2000)
    
```



z990 Sysplex support

z990 CF Link Connectivity to G5/G6 and z900/z800

Connectivity Options	z990 ISC-3	z990 ICB-2	z990 ICB-3	z990 ICB-4
Pre G5 9672	Not Supported	Not Supported	Not Supported	Not Supported
G5/G6 ISC	1 Gbps Compat Mode	n/a	n/a	n/a
z800/z900 ISC-3	2 Gbps* Peer Mode	n/a	n/a	n/a
z990 ISC-3	2 Gbps Peer Mode	n/a	n/a	n/a
G5/G6 ICB	n/a	333 MBps Compat Mode	n/a	n/a
z900 ICB-2	n/a	Not supported	n/a	n/a
z990 ICB-2	n/a	Not Supported	n/a	n/a
z800/z900 ICB-3	n/a	n/a	1 GBps Peer Mode	n/a
z990 ICB-3	n/a	n/a	1 GBps Peer Mode	n/a
z990 ICB-4	n/a	n/a	n/a	2.0 GBps Peer Mode



Message Time Ordering

If CF is on z900 Turbo or faster, *and* connected systems are on z900 Turbo or faster:

- The CF must be connected to the same Sysplex Timer as all connected systems
- APAR OW53831 **MUST** be applied to all connected systems running on z900 Turbo or faster



System Managed CF Duplexing status

System Managed Duplexing experiences:

- System Managed Duplexing is now generally available, back to z/OS 1.2 with enabling APAR OW41617
- At least 10 customers in production, at least 50 more in some stage of roll out
- Very few performance-related problems reported
 - Response times very dependent on distance and technology level, but have seen duplexed response times as low as the mid-50 microseconds when using z990 and ICB4.
 - At least one customer happily exploiting SM Duplexing for DB2 Lock and SCA structures for production data sharing at 3km. - "acceptable" performance depends on your requirements
- At the time of writing, there is only one open SM Duplexing-related APAR



System Managed CF Duplexing status

System Managed Duplexing documentation:

- See PSP Bucket Upgrade CFDUPLEXING for information about APARs for z/OS and all subsystems that support System Managed Duplexing - lists a number of HIPER APARs.
- Use the self-assessment questionnaire available on ResourceLink <http://www.ibm.com/servers/resourcelink>
- See SM Duplexing White Paper at: <http://www.ibm.com/servers/eserver/zseries/library/techpapers/pdf/gm130103.pdf>
- See SM Duplexing Implementation Summary White Paper at: <http://www-1.ibm.com/servers/eserver/zseries/library/techpapers/pdf/gm130540.pdf>
Summary of information in SM Duplexing White Paper above



CFLevel 12

With CFLevel 12, the number of concurrent tasks increased from 16 to 48

CFLevel 12 provided support for Message Time Ordering

CFLevel 12 provided 64-bit support, alleviating constraints for customers with large number of structures or very large lock structures.

- Whether a CF runs in 31-bit or 64-bit mode is COMPLETELY transparent to the connected operating systems

CFLevel 12 provided support for System-Managed CF Structure Duplexing

CF Level 12 provided batch CF request enhancements for DB2 - exploited by DB2 V8



CFLevel 12

For workloads that are very update-intensive (like large DB2 batch jobs), sending requests to a CF one at a time is not very efficient

To address this, CFLevel 12, together with z/OS 1.4 introduces three new options on the IXLCACHE command that allows write, castout, and cross-invalidate requests to be batched

This support is exploited by DB2 V8

Tests within IBM with update-intensive DB2 batch job resulted in:

- 30% reduction in DB2 DBM1 address space CPU time
- Similar elapsed time
- 50%+ decrease in CF Requests, small increase in resp time, small decrease in CF CPU busy



CFLevel 13

Delivered as part of z990 GA3/z890 GA1 - also available on z800/z900

Two significant new functions:

- Enhancements to Read Castout Class (RCC) command.
- Ability to install patches and new CF Levels without requiring a POR ([z890/z990 only](#))

Requires z/OS APAR OA01517

Minimal increase in size of some cache structures, no impact on list or lock structures



CFLevel 13

Enhancements to RCC function:

- Prior to this enhancement, requests to castout pages by data set (rather than castout class) can run for a long time, potentially driving CF CPU utilization to 100%. This is most likely to occur if a GBP contains pages from both data and index components and online utilities are run.
- The enhanced function is available as soon as CFLevel 13 is activated - no APARs are required for z/OS or DB2
- Further exploitation will be provided by a future version of DB2



CFLevel 13

Ability to activate patches and new CF Levels without requiring a POR:

- Prior to this enhancement, some CFCC patches and moving to a new CFLevel required a POR of the CPC to activate the patch
- This enhancement provides the ability to activate these patches, and new CFLevels, one CF LPAR at a time
- Especially attractive for environments where a single CPC contains operating system images, test CFs, and production CFs



CFLevel 14

Delivered as part of z990 GA4/z890 GA2 - not available on z800/z900

Implements a re-write of the CFCC Dispatcher

- The enhanced dispatcher is expected to deliver the most benefits for environments that are XCF-signalling intensive and have a large number of CPs in the CF or those exploiting SM Duplexing

Requires z/OS APAR OA08742

Minimal impact on size of all structures



SETXCF REALLOCATE command

Customers have experienced problems due to structures being in CFs other than the first one in the PREFLIST

- Some structures should be failure isolated, but can end up in the "wrong" CF which might not be failure isolated

Traditional rebuild processes not always effective

- SETXCF REBUILD,CFNM rebuilds all structures in parallel, rebuilds structures that don't need to be moved, won't rebuild duplexed structures, even if the OLD instance is not in the first CF in the PREFLIST, and won't move XCF structures

POPULATECF is an improvement, however:

- POPULATECF won't move a duplexed structure, even if the OLD instance is not in the most preferable CF as specified in the CFRM policy. Also, POPCF doesn't consider all structures, only those that would end up in the named CF



SETXCF REALLOCATE command

New SETXCF START,REALLOCATE command:

- Clears all CFRM "policy change pending" conditions
- Moves all simplex structures into their "most preferred" CF location
- Moves all duplexed structure instances into their two "most preferred" CF locations, in the correct order (automatically corrects "reversal" of primary and secondary structure locations)
- Serial, one-structure-at-a-time processing to minimize any disruption caused by reallocation actions
- Issues a message describing the evaluation process for each allocated structure
- Issues a summary message upon completion of all structures, summarizing actions taken
- Can be used to simplify CF structure movement during disruptive CF maintenance/upgrade scenarios



SETXCF REALLOCATE command

```

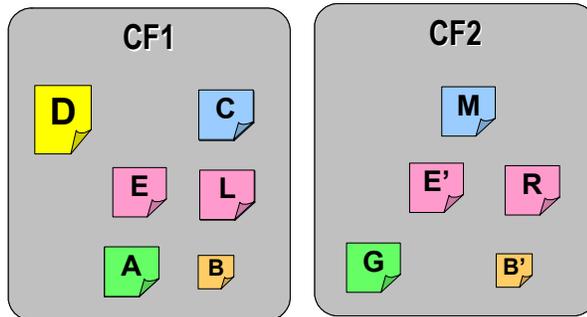
SETXCF START,REALLOCATE
IXC574I EVALUATION INFORMATION FOR REALLOCATE PROCESSING 407
OF STRUCTURE ISGLOCK
SIMPLEX STRUCTURE ALLOCATED IN COUPLING FACILITY: FACIL03
ACTIVE POLICY INFORMATION USED.
CFNAME      STATUS/FAILURE REASON
-----
FACIL04     PREFERRED CF 1
                                     INFO110: 00000064 CC007800 0000000D
FACIL03     PREFERRED CF ALREADY SELECTED
                                     INFO110: 00000064 CC007800 0000000D
.....

IXC545I REALLOCATE PROCESSING RESULTED IN THE FOLLOWING: 168
  1 STRUCTURE(S) REALLOCATED - SIMPLEX
  0 STRUCTURE(S) REALLOCATED - DUPLEXED
  0 STRUCTURE(S) POLICY CHANGE MADE - SIMPLEX
  0 STRUCTURE(S) POLICY CHANGE MADE - DUPLEXED
 12 STRUCTURE(S) ALREADY ALLOCATED IN PREFERRED CF - SIMPLEX
  4 STRUCTURE(S) ALREADY ALLOCATED IN PREFERRED CF - DUPLEXED
  0 STRUCTURE(S) NOT PROCESSED
 41 STRUCTURE(S) NOT ALLOCATED
 142 STRUCTURE(S) NOT DEFINED
-----
 200 TOTAL

  0 ERROR(S) ENCOUNTERED DURING PROCESSING
IXC543I THE REQUESTED START,REALLOCATE WAS COMPLETED. 169
    
```



SETXCF REALLOCATE command



- Operator needs to clear out CF1 for maintenance
- Operator issues setxcf stop,duplex,keep=new for structures E B
- Operator issues setxcf rebuild,cfname=CF1,loc=other

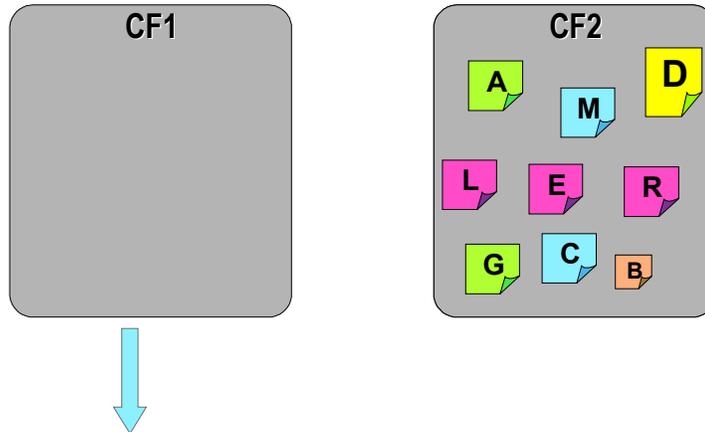
Active CFRM Policy

```

STRUCTURE NAME(C) PREFLIST(CF1,CF2)
STRUCTURE NAME(E) PREFLIST(CF1,CF2) DUPLEX(ALLOWED)
STRUCTURE NAME(B) PREFLIST(CF2,CF1) DUPLEX(ALLOWED)
STRUCTURE NAME(A) PREFLIST(CF1,CF2)
STRUCTURE NAME(G) PREFLIST(CF2,CF1)
STRUCTURE NAME(D) PREFLIST(CF1,CF2)
  0
  0
  0
    
```



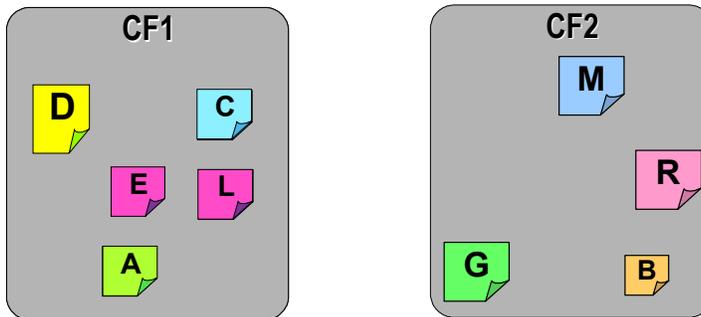
SETXCF REALLOCATE command



Operator configures CF offline and it is then given to CE



SETXCF REALLOCATE command



Active CFRM Policy

```

STRUCTURE NAME(C) PREFLIST(CF1,CF2)
STRUCTURE NAME(E) PREFLIST(CF1,CF2) DUPLEX(ALLOWED)
STRUCTURE NAME(B) PREFLIST(CF2,CF1) DUPLEX(ALLOWED)
STRUCTURE NAME(A) PREFLIST(CF1,CF2)
STRUCTURE NAME(G) PREFLIST(CF2,CF1)
STRUCTURE NAME(D) PREFLIST(CF1,CF2)
    
```



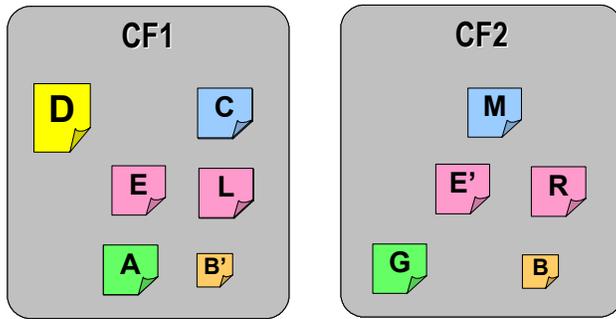
Operator needs to repopulate CF1 according to PREFLIST 0

Operator issues **SETXCF REALLOCATE** 0

All structures get relocated according to PREFLIST



SETXCF REALLOCATE command



Active CFRM Policy



Operator establishes duplexing
 setxcf start,duplex,strname=b
 setxcf start duplex,strname=e
 Etc.

```
STRUCTURE NAME(C) PREFLIST(CF1,CF2)
STRUCTURE NAME(E) PREFLIST(CF1,CF2) DUPLEX(ALLOWED)
STRUCTURE NAME(B) PREFLIST(CF2,CF1) DUPLEX(ALLOWED)
STRUCTURE NAME(A) PREFLIST(CF1,CF2)
STRUCTURE NAME(G) PREFLIST(CF2,CF1)
STRUCTURE NAME(D) PREFLIST(CF1,CF2)
```

```
0
0
0
```



SETXCF REALLOCATE command

REALLOCATE is structure- rather than CF-based

It will move *any* structure that is not in the most preferable available CF, even if the structure is duplexed

Function provided with APAR OA03481 back to z/OS 1.4

APAR needs to be installed prior to IPLing before REALLOCATE cmd is available

APAR needs to be installed on all participating members

- REALLOCATE process will NOT be started when XCF discovers an active system in the sysplex without the prerequisite code.



SETXCF REALLOCATE command

Process to empty a CF:

- SETXCF START,RB,CFNM=xxx,LOC=OTHER
- SETXCF START,RB,STRNM=xxx,LOC=OTHER (for xcf structures)
- SETXCF STOP,RB,DUPLEX,CFNM=xxx

Process to bring CF back to "normal":

- Bring CF online to all systems
- SETXCF START,REALLOCATE



SETXCF REALLOCATE command

There are some exploiters that do not provide robust support for structure rebuild - to stop REALLOCATE touching those structures, a new CFRM keyword will be added to let you specifically exclude certain structures from REALLOCATE processing - see APAR OA08688.



XCF/XES usability enhancements

APAR OA03993 will be shipped back to z/OS 1.4:

- Displays Dedicated/Shared status of CF PUs on D CF
- Batch program to delete failed XCF members
- Structure object utilization information is now contained in D XCF,STR output
- Adds more diagnostic information to logrec records for CF-related problems

APAR OA02620

- Adds structure type to response to D XCF,STR command
- Disallow FORCE of Failed-Persistent Connections
- You CAN now force a Failed-Persistent structure



Accurately sizing z990 processors

Single number metrics - MIPS and MSUs - are based on averages and should not be used for sizing capacity

Actual workload performance can vary significantly around the average

- Greatly affected by workload characteristics particularly as pertains to Nway scalability
- Also affected by LPAR configuration and coupling technology

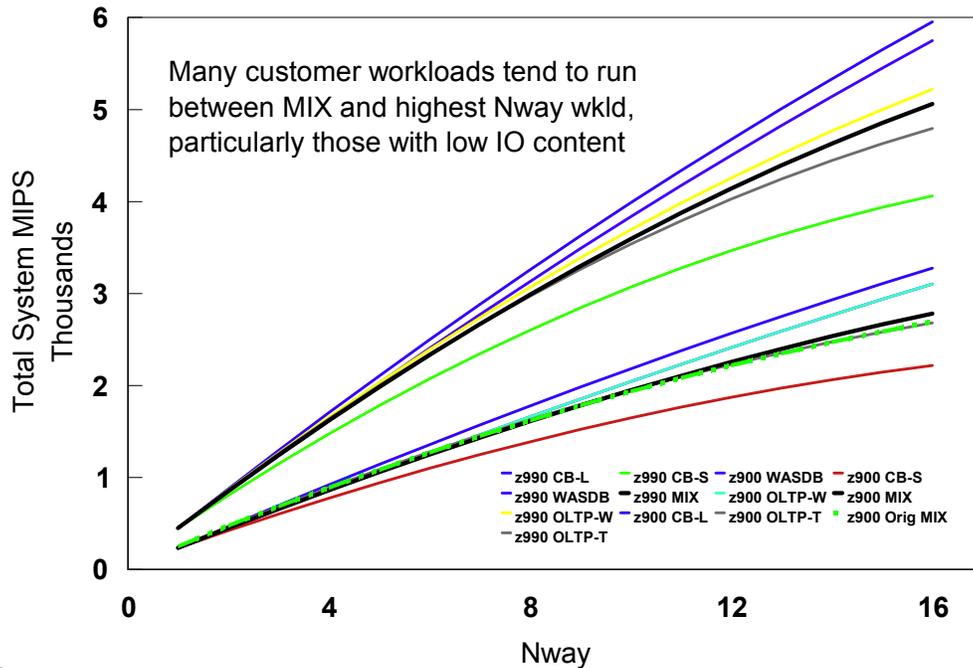
Bigger processors have bigger variability

- Wider variety and changing mix of applications
- Unexpected latent demand can be "let loose" by significantly faster engines



Accurately sizing z990 processors

z900 to z990 upgrades - sensitivity to workload scaling



Accurately sizing z990 processors

Perform a customized capacity plan using sizing tools

- Customize for workload (check for low DASD IO)
- Customize for LPAR configuration

Consider need for tuning changes for LPAR and subsystems

Consider upgrading coupling technology

- Understand capacity impacts of not changing

Monitor need to adjust WLM tuning controls

- Watch out for latent demand



IBM Early Support Programs



What is an Early Support Program (ESP)?

Product Introduction Program providing an opportunity for customers to install, test, use and comment on IBM products/solutions prior to General Availability.

Duration: 2 months to 1 year, depending on program/product

Types of Product Introduction Programs:

- Beta Programs
- Early Support Programs (ESPs)
- Joint Project Development
- Quality Partnership Program

Product Introduction Programs are available for:

- iSeries
- pSeries
- TotalStorage
- zSeries



Early Program Objectives

Evaluate new hardware and software in customer environments

- verify installation, migration, technical validation, documentation, and user experience
- ensure that the product meets/exceeds customer needs and expectations

Improve Product Quality

Improve Customer Satisfaction

Validate IBM's support structure

Provide feedback to development and support organization

Strengthen IBM/Customer partnerships

Solicit and incorporate customer suggested enhancements as appropriate

Obtain References, endorsements and/or testimonials



Customer Benefits

Early availability of products and documentation

- Competitive edge: function, performance and cost saving benefits

Early Program Education

- Learn about new functions

Assigned Early Program Representative

- Liaison to Development/Service personnel
- Facilitate problem resolution

Priority on Customer Problem Resolution

- Expedited channel to Development

ISV assistance

Q&A Support

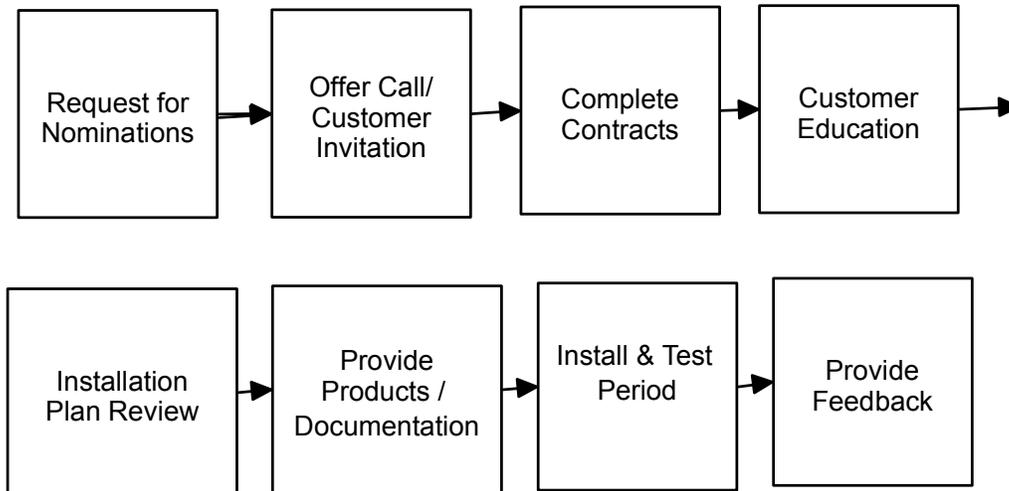
- Forum, Development calls, Information sharing, Early documentation

Disclosure of future products and product directions

Opportunity to influence future product directions



What happens during an Early Program?



by

Early Program Support Services

Early product availability

- Delivery process dependent on program/product

Installation Assistance

- Installation Plan Reviews
- Systems Assurance
- IBM Services, as appropriate

Service

- Identified by Product Introduction Program (PIP) customer number or HW serial number, depending on program/product
- Normal service process as defined by program/product
- Q & A support



Early Program Support Services

Other information:

- You do **not** have to take part in EVERY ESP to be in the program
- z/OS and zSeries processor programs are handled by the same group
- There is no set limit on number of participants, beyond what can be supported by the ESP group
- If you take part in a hardware ESP, you must purchase the hardware - you do NOT get a loan of it!
- Customer reaction generally very positive - once involved, customers usually join again and again



What do you need to do to participate?

Contact your local IBM rep or one of the following Early Program contacts:

zSeries contact

- Kathy Kulchock - kathyk@us.ibm.com

iSeries contact

- Stacy Haugen - shaugen@us.ibm.com

pSeries contact

- Beth Varuola - varuola@us.ibm.com

TotalStorage contact

- Jerry Boyle - jeboyle@us.ibm.com
- Andy Schneider - andyjs@us.ibm.com
- Brad Garland - bgarland@us.ibm.com



zSeries Data Scalability

Addressing the Growth in data in the sysplex



zSeries Data Scalability

What is the issue:

- Once again, our largest customers are moving towards the limits of the number of devices supported per LP

Why?

- Far more use made of data (data mining, data warehouses, etc)
- Server consolidation
- Good system management practices say that all systems in a sysplex should see the same I/O configuration
 - However IBM pricing rules result in diverse system types in the same sysplex ==> more data online to each system
- Function-rich data replication technologies (like FlashCopy) require device numbers
- Need for disaster recovery and continuous availability can double number of used device numbers overnight
- Techniques to deliver high performance on large capacity devices (PAV) require device numbers



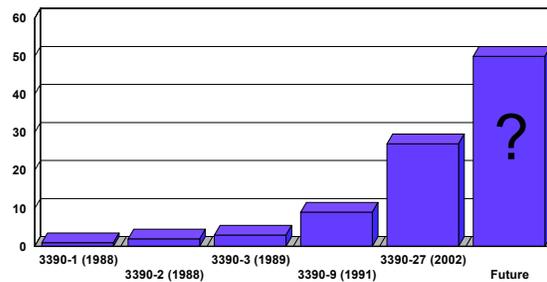
zSeries Data Scalability

This isn't a new challenge.....

Date	Event	# of Cyls
1988	3390-001	1113
1988	3390-002	2226
1989	3390-003	3339
1991	3390-009	10017
2002	3390-027(*)	32760
???	???	>32K

Year	Event	# devs in LP
	3-digit addresses	1970
1994	9672 G1	16384
1995	9672 G2	24576
1999	9672 G5	36864
2000	z900	64512

* "3390-027" reported by QDASD command as 3390-9



zSeries Data Scalability

So, how do we address it:

- Short/medium term:
 - Make full use of existing large devices
 - Maximize use of device numbers
 - Introduce still larger devices
- Longer term:
 - Maybe larger device numbers?
 - Maybe fuller exploitation of existing CCHHR architecture?



zSeries Data Scalability

Using larger devices:

- Not a trivial exercise. Need to consider:
 - Performance/contention of larger volumes
 - ▶ PAV has proven to be very effective at providing excellent performance even with very large volumes
 - Chargeback/ownership
 - ▶ Need to move away from concept of applications "owning" volumes
 - Backups (is tape backup still a viable option?)
 - ▶ This is an issue regardless of device type. How do you backup 10TB to tape with minimal application impact? How do you identify points of consistency to back up databases, application libraries, catalogs, spool, etc?
 - Data movement is disruptive
 - Software support



zSeries Data Scalability

Migrating to larger volumes:

- 3 categories of files:
 1. Constantly being deleted and reallocated
 2. Inactive (not deleted and not in use)
 3. Constantly in use and rarely deleted/reallocated
- Each requires a different strategy...
 1. Use SMS to direct all allocations to larger device types
 2. Use a data mover tool (such as DFSMSdss) to move the data sets to the new devices
 - Review I/O rate skews - can data sets from volumes with very low access rates be moved to tape sooner?
 3. Despite tools such as FlashCopy and HyperSwap, there is no easy, non-disruptive way to merge multiple small volumes onto a larger one
 - Currently, an impact to the data availability *must* be scheduled to move data sets from one volume to another



zSeries Data Scalability

Considerations for moving data sets:

- Data sets that are accessed via VOL=SER (uncataloged)
- Data sets that are cataloged in >1 catalog (JOB CAT/STEP CAT removal should tidy these up)
- ISAM data sets (ISAM will be removed in z/OS 1.7)
- Unmoveable data sets
- Data sets allocated with absolute track
- Data sets that are "cataloged" elsewhere (HSM CDS, spool managers, etc)
- Data sets that have size restrictions or must be placed within a certain range of tracks
- Beware of data sets using non-standard access methods
- Consider VTOC and VTOCIX sizes



zSeries Data Scalability

Considerations for moving to large capacity volumes

- JES2 supports having the spool on a large capacity volume, but the spool data set size is limited to 64K tracks
 - Check RELADDR field in \$D SPOOLDEF
- Page data sets still limited to 4GB, however they can reside anywhere on the volume
- Non-extended format VSAM data sets still limited to 4GB
- Non-extended format SEQ data sets limited to 64K tracks.
 - Stand alone dump data sets - require APAR OA04140 to support data sets larger than 64K tracks (extended format sequential data sets)
- Extended format SEQ can be > 64K tracks
- Extended format VSAM can be > 4GB



zSeries Data Scalability

Maximizing device numbers:

- Review old practices - do you reserve ranges of device numbers for particular device types? This is convenient but no longer necessary
- Are you over-configured for CTCs? Maximize use of CF structures for system-to-system traffic and reduce the number of CTCs to the level you require
- Do not over-configure PAV aliases. It is typical to see 2-3 aliases per base address.....



zSeries Data Scalability

Two formulas for estimating the number of required aliases per LSS:

Method 1:

3-6 Aliases per ESCON ESS connection

6-12 Aliases per FICON ESS connection

Method 2:

Get average device utilization for an LSS using RMF report

Multiply by the number of devices in the LSS

Divide by 100 then multiply by 2

For example, if average utilization is 12% and 168 devices,

$$(12 * 168) / 100 * 2$$

$$= 2016 / 100 * 2$$

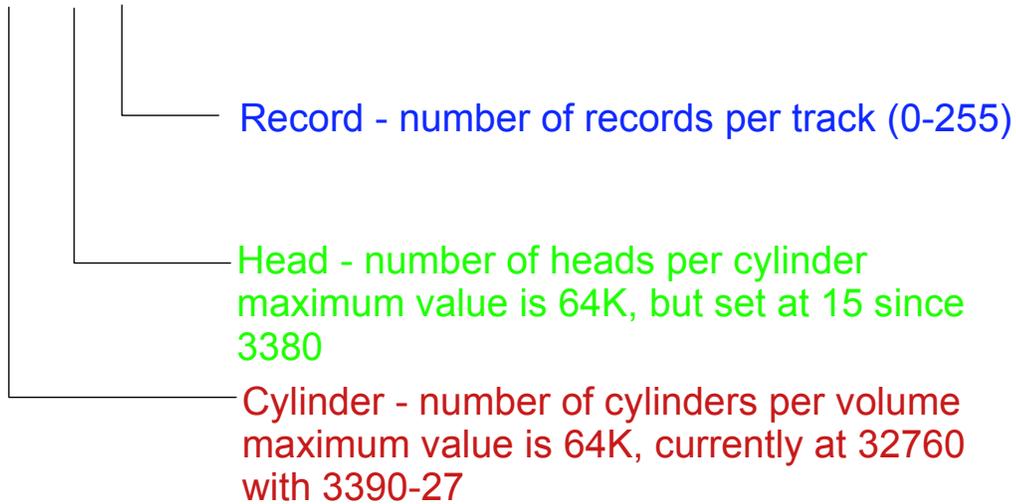
$$= 40 \text{ aliases}$$



zSeries Data Scalability

Creating larger volumes

CCHHR



zSeries Data Scalability

The future....

- Options of more heads per cylinder
- More cylinders per device
- More device numbers



zSeries Data Scalability

Summary:

- Do an inventory of your I/O config and your plans for remote copy and FlashCopy - where will this position you against the 64512 device limit?
- Do an inventory of "unusual" data set types - BDAM, non-standard access methods
 - Is there a way to identify these through the DSN?
 - Identify a move method for each one
- Tidy up unused device numbers - old devices, over-configured CTCs, over-allocated PAV aliases (prior to z/OS 1.6, you must vary off unneeded aliases prior to doing the ACTIVATE)
- Start implementing large volumes (Mod 27) now, initially for temp and work data sets, then database volumes
 - Develop a methodology for moving to larger volumes - this isn't the last time you will have to do this



Miscellaneous



End-of-service for Sysplex Timer Model 1

Support for 9037-001 Timers **ENDED** on December 31, 2003

- Replacement is 9037-002, however 001 is not upgradeable - you must remove the 001 and replace it with a 002
- Replacing 9037-001 with 9037-002 requires a nearly sysplex-wide outage
- See the ITSO Redpaper entitled "Migration Planning for the 9037 Model 2 Sysplex Timer" for a suggested migration procedure

Next generation HMC will NOT provide ability to run the Sysplex Timer Console application:

- Order feature code 4900 to add Sysplex Timer Console (at least 2)



Withdrawal from Marketing of 9032-005

IBM announced withdrawal from marketing of ESCON Director Model 5 *and all associated features* in April 2004 effective Dec 31, 2004

Next generation HMC will NOT provide ability to run the ESCON Director Console application:

- Order console feature code 5900 to get an ESCON Director Console if you will be replacing all HMCs with next generation HMCs before all ESCON Directors are removed



JOBSTAT/STEPSTAT

Starting with z/OS 1.5:

- Use of JOBSTAT/STEPSTAT now optional at the installation level
 - Default is to no longer honor JOBSTAT/STEPSTAT
- New parameter to MODIFY CATALOG command:
 - F CATALOG,ENABLE(JOBSTEPSTAT)
 - F CATALOG,DISABLE(JOBSTEPSTAT)
 - Current setting retained across CAS restart
- If support is disabled and a catalog request is issued with JOBSTAT/STEPSTAT message IDC3009I will be issued with a return code 54 reason code 6 and the request will not be processed



Catalog space monitoring

Prior to z/OS 1.5, there was no warning when a catalog is about to run out of space

New enhancement issue message when a specified percentage of maximum extents is reached

- IEC361I CATALOG catalogname (comptype) HAS REACHED xxx% OF THE MAXIMUM
- Issued once per catalog per extent and reissued for each subsequent catalog extent
- "comptype" is either DATA or INDEX

Invoked by issuing F CATALOG,NOTIFYEXTENT(yyy) where yyy is a value from 0-100. Default is 80. A setting of 0 or 100 will disable the function

F CATALOG,ALLOCATED will display the percentage of allocated extents for each catalog in the list in message IEC348I



Recreating lost SMS SCDS

New function in z/OS 1.5:

- SETSMS SAVESCDS(scds_dsname) command

Relieves the tedious work of recreating the SCDS if the source SCDS and all of its backups are lost

SMS will verify that the 'scds_dsname' is not the currently active ACDS or COMMDS



Decreased IPL times

One aspect of high availability is restoring service as quickly as possible following an outage (planned or otherwise).

APAR OA07335 (integrated in z/OS 1.6 and rolled back to z/OS 1.4 via PTF) can help reduce IPL times:

- Prior to this APAR, dynamic path and PAV initialization was a serial process, so the more DASD you have, the longer it took.
- OA07335 now does this processing in parallel.
- In a number of tests in controlled environments, Master Scheduler Initialization time reduced by up to 70%.



ICF/IFL/zAAP Pooling

It is documented, though not widely known that ICFs, IFLs, and zAAP PUs are treated as a single resource pool by PR/SM on z990/z890

As a result, it is possible that one type of workload (Linux, for example) can use resource that you expect to be dedicated to another workload type (CFs, for example)

It is not immediately obvious when this happens

The weights of different LPAR types (Linux, CF, z/OS) can now have an impact on each other

More information will be provided in the hardware track



Interesting sources of information

"Hints & Tips", RedPapers, RedPieces on Redbooks Web site

- <http://www.redbooks.ibm.com>
- Other ITSO news:
 - Upcoming residencies: IMS APPC, MLS, JES2 Performance Redpaper, Update to Parallel Sysplex Cookbook, RMF Update
 - Recent Redbooks: Systems Programmer's Guide to RRS, Application Programming Considerations for Sysplex

The "MainStream" IBM mainframe software newsletter:

- <http://www-306.ibm.com/software/os/zseries/newsletter/mainstream/selection/>

eServer Magazine, mainframe edition CF:

- <http://www.eservercomputing.com/mainframe/>

z/journal (independent zSeries publication)

- <http://www.zjournal.com/>



Interesting sources of information

Hot Topics Newsletter

- Published twice a year (Feb and August)
- Available from:
 - http://www.ibm.com/servers/s390/os390/bkserv/hot_topics.html

SEARCH390.COM

- <http://www.search390.com>

RMF Newsletters

- <http://www.ibm.com/servers/eserver/zseries/zos/rmf/rmfhtmls/rmflett.htm>

HCD and HCM Newsletter and latest news

- <http://www.ibm.com/servers/eserver/zseries/zos/hcm/hcmhtmls/hcmnltab.html>

PUTDOC tool to FTP dumps to IBM

- <http://techsupport.services.ibm.com/server/nav/zSeries/putdoc/putdoc.html>



Interesting sources of information

It is now (finally!) possible to subscribe to WSC Flashes!

To subscribe, go to <http://www.ibm.com/support/mysupport> then:

- 1) Select "Edit Profile" in the left hand navigation bar. This is where you decide on which products you would like see a FLASH notice on.
- 2) Select product family from the "Select a product family" pull down in the center of the page
- 3) Next, select one or all (but at least one), of the topics found under the product family you selected (make the selection by putting a check in the box beside it).
- 4) Scroll to the bottom of the page and click on "Save & return"
- 5) Go to the bottom of the page that you returned to, under the blue bar header of "Select mail preferences", and under the text "Yes, you may send me this information by e-mail", select Flashes and (if you wish) Downloadable files
- 6) Click on "Submit"



Library Center

New way to access z/OS and OS/390 product information

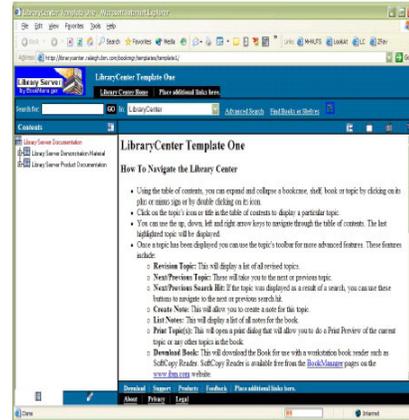
- Formerly BookServer

Will serve as a single repository for all product information

Available in z/OS V1R5

Based on BookServer technology

- Customizable
- Improved end user Interface
- Enhanced Bookshelves and Bookcases
- Enhanced tables
- Enhanced bookshelf (XKS) support
- Handheld device support
- LibraryServer is a multi-platform product (z/OS, Windows, AIX, Linux)



http://publibz.boulder.ibm.com/bookmgr_OS390/libraryserver/zosv1r5/

Will no longer need to know which bookshelf to search



Library Center

New view of IBM's z/Series Internet Library

Uses the new IBM Library Server product

One Library Center for each new release of z/OS (all books on DVD)

- Based on z/OS DVD collections
- Over 2000 manuals under one umbrella

Provides an "Explorer-like" view of BookManager and PDF repositories

Provides advanced searches based on type of information

- Messages
- Commands
- Examples

Integrated PDF download

Ability to launch external searches such as

Google



Updating system symbols

ITSO used to provide a program called SYMUPDTE to dynamically update System Symbols

- SYMUPDTE can be used to update existing symbols without an IPL. It can also add new symbols without an IPL
- Be sure to pull the documentation to understand the limitations of using this program

This program is now delivered in SYS1.SAMPLIB as IEASYMUP, however no JCL or documentation is provided

Note that IEASYMUP REQUIRES a RACF profile

The documentation on the Redbooks Web site has been updated and is still applicable

- Get from Additional Materials section of Redbooks Web site for SG245451



Updating system symbols

JCL to link IEASYMUP:

```
//KYNEFL JOB (0,0), 'LINK SYM', CLASS=A, MSGCLASS=X, NOTIFY=KYNEF
//S2 EXEC PGM=IEWL,
// PARM='XREF, NCAL, LIST, RENT, LET, AC=1'
//SYSPRINT DD SYSOUT=*
//SYSUT1 DD UNIT=SYSDA, SPACE=(TRK, (20, 10))
//SYSLMOD DD DSN=KYNEF.SYMUPDTE.LOADLIB, DISP=(, CATLG),
// SPACE=(CYL, (1, 1, 5)), RECFM=U, LRECL=0, BLKSIZE=6144,
// UNIT=SYSDA
//SYSOBJS DD DSN=SYS1.SAMPLIB, DISP=SHR
//SYSLIN DD *
INCLUDE SYSOBJS (IEASYMUP)
NAME IEASYMUP (R)
```

JCL to run IEASYMUP:

```
//KYNEFR JOB (0,0), 'TEST SYMUPDTE', CLASS=A, MSGCLASS=X
//SYMUPDTE EXEC PGM=IEASYMUP, PARM='TESTFK=TEST1'
//STEPLIB DD DSN=KYNEF.SYMUPDTE.LOADLIB, DISP=SHR <==== APF library
```



Statements of direction

Effective in z/OS 1.6:

- IBM intends to support up to 30 LPARs on a z990
- The default currency symbol in Euro countries is being changed to be the Euro
- ISPF C/C++ Panels are being removed. Various other changes related to C/C++
- Language Environment RTLS is being withdrawn
- DCE Application Services is being removed
- Encina Toolkit Executive is being removed
- Text Search element of z/OS will be available as a Web download rather than being shipped as part of z/OS
- IBM intends to provide a 64-bit SDK 1.4 level Java product concurrently with the general availability of z/OS and z/OS.e V1.6.



Statements of direction

In future releases:

- 1.6 is last release that will support ISAM. The ISAM Compatability Interface that lets you use ISAM programs with VSAM files will continue to be supported
- 1.6 is NOT the last release that will support BIND DNS 4.9.3. - this means that DNS/WLM support will disappear. This support will be withdrawn in a future release
- 1.6 is last release that will support OROUTED. Should be replaced with OMPROUTE
- In a future release, AnyNet will be removed. Enterprise Extender should be used to replace AnyNet
- In a future release of z/OS, IBM intends to certify z/OS to the Common Criteria Controlled Access Protection Profile (CAPP) at EAL2 and Labeled Security Protection Profile (LSPP) at EAL3+.



Statements of direction

In z/OS 1.7:

- IBM intends to withdraw support for ISAM. The only actions that will be possible on ISAM data sets will be to delete them
- Support for JOBCAT and STEPCAT JCL statements will be removed
- Support for JES2 compatibility mode will be removed. Spool MUST be in R4 format. \$ACTIVATE command is being removed
- The z/OS Optional Source Code feature (containing source for some programs in z/OS BCP, BDT base, BDT SNA NJE, BDT File-to-File, DFSMS, MICR/OCR, BCP JPN, and Security Server RACF elements) will no longer be available



Statements of direction

In future releases:

- IBM intends to support SNA for the foreseeable future. To help with the replacement of 3745/3746 hardware, IBM plans to provide in 2005 a version of NCP that will run under Linux on zSeries
- In 2005 IBM Plans to deliver a VSAM JDBC connector
- In the future, IBM plans to remove support for VSAM IMBED, REPLICATE, and KEYRANGE attributes. At that time, data sets with these attributes will not be processed
- In the future, IBM will remove the DFSORT English and Japanese ISPF panels



Statements of direction

For summary of z/OS "Statement of Direction" announcements, see:

- http://www.ibm.com/servers/eserver/zseries/zos/zos_sods.html#040810



Supported releases

CICS TS

- CICS/ESA V4.1, CICS TS 1.1, 1.2 all now out of support.
- CICS TS 1.3 runs out of support on April 30, 2006
- CICS TS 2.3 current release - available 12/19/2003

DB2

- DB2 V6.1 runs out of support June 30, 2005
- DB2 V7.1 no date announced yet
- DB2 V8 avail since Mar 2004, requires zSeries

IMS

- IMS V6.1 ran out of support Sept 30, 2003
- IMS V7.1 runs out of support Nov 8, 2005
- No end of support date for V8.1
- IMS V9.1 QPP running, no GA date yet



Supported releases

MQ

- MQSeries V1.2, V2.1, and V5.2 ran out of support April 2004
- Websphere MQ V5.3.0 runs out of support Oct 31, 2005

OS/390

- All releases prior to V2.10 already out of support.
- V2.10 runs out of support on Sept 30, 2004.

z/OS

- z/OS 1.1 already out of support (Mar 31, 2004).
- z/OS 1.2 runs out of support on Oct 31, 2004.
- z/OS 1.3 runs out of support on Mar 31, 2005.



Supported releases

z/OS Support



		G3-G4	G5/G6 MP3000	z800 z900	z890 z990	End of Service	Coexistence Migration Policy	Ship Date
OS/390	2.8	x	x	x		9/02	1.2	
	2.9	x	x	x		9/03	1.3	
	2.10	x	x	x	x ^C	9/04	1.4	
z/OS	1.1		x	x		3/04	1.4	
	1.2		x	x	x ^C	10/04	1.5	
	1.3		x	x	x ^C	3/05	1.6	
	1.4		x	x	x	3/07	1.7	9/02
	1.5		x	x	x	3/07*	1.8	3/04
	1.6			x	x	9/07*	1.8	9/04
	1.7			x	x	9/08*	1.9	9/05*

* Planned dates

1.4 and 1.5 will remain orderable until Sept. 9, 2004

x^C - Compatibility support – does not exploit new z990 features: 30 LPARs and multiple Logical Channel SubSystems
Bimodal Accommodation Offering is available for z/OS 1.2, 1.3, and 1.4. It will not be provided for z/OS 1.5



Supported releases

Why you should stay up to date:

- CICS TS 2.3 requires z/OS 1.4
- DB2 V8 requires zSeries and z/OS 1.3
- Lotus Domino 6 requires z/OS 1.2
- Java SDK 1.4.1 requires z/OS 1.2

All dates are US dates - dates can vary by country

See the following URLs for more information:

- <http://www.ibm.com/services/sl/products/java.html>
- <http://www-306.ibm.com/software/info/supportlifecycle/>



Interesting APARs

OW56355 - provides benefits for users of VTAM/XCF experiencing NO BUFFER conditions

OA02620 - Disallow force of Failed-Persistent CONNECTION

OA03481 - New REALLOCATE command for CF Management. See WSC FLash 10285.

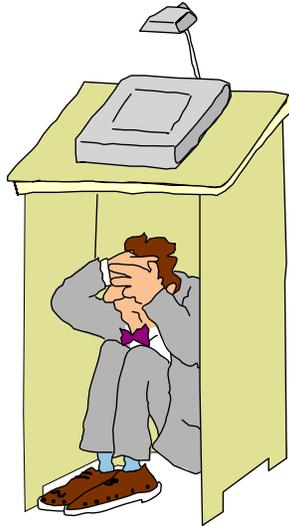
OA03676 - PE APAR that can cause VSAM data loss.

OA07335 - APAR to provide parallelism during IOS initialization. Can significantly decrease IPL times for large I/O configurations

OA03765 - Related to problem removing a PAV alias from a paging volume



Questions?



Thanks!!

Hope I see you again next
year

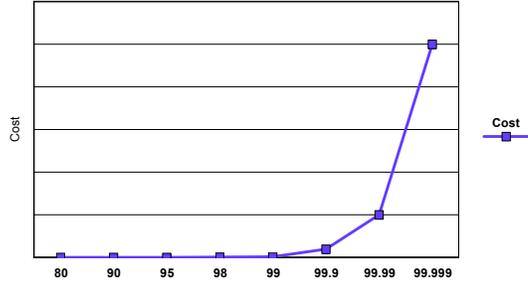


Appendix 1: Availability Checklists



Management considerations

Consideration	Addressed?
Which are your really critical applications - the ones where all the phones ring when they are down? Does your IT spending reflect the criticality of the applications?	
When a new software release goes in, is education about the new features provided for ops and appl developers?	
Is there an executive that has responsibility for availability AND has the budget to improve availability?	
Does the group that is responsible for availability (typically Ops) have the authority to stop an application going into production if that application does not support high availability?	
Does everyone that plays a role in availability have availability as one of their objectives? Ops, sysprogs, and appl developers MUST all work together, but often these groups have different objectives	



Systems Management considerations

Consideration	Addressed?
Do you know your current availability? If you don't know what you are achieving today, how will you know if changes make things better or worse?	
Where is your pain? What is the root cause of your past outages? The pattern varies from installation to installation, so you need to be able to identify where your weaker areas are. If a particular component is causing a large percent of outages, can it be removed or replaced? You also need this information to justify expense of improving availability	
A recent (non-IBM) survey, found that 80% of unplanned outages are as a result of poor Systems Management practices	
Communication (Problem and Change Management) must cover <i>all</i> the people that touch any part of an application, including application developers, database admin, network, distributed servers, ops and the MVS System Programmers	
Software updates - large percent of outages are caused by problems where the fix has been available for more than 6 months. Consider using DR tests as a way of testing the new software levels.	
Subscribe to Red Alerts and apply/address identified fixes as soon as possible	
Access control to critical resources - Do only the people that really should have access actually have access to critical system files?	
Documentation - are procedures for failure scenarios documented, written for the target audience, and <i>regularly tested</i> ?	
Problem determination - how often are you told that you must recreate a problem to get the information required to debug the problem? If this happens a lot, perhaps your procedures should be reviewed? You learn NOTHING from an unexplained outage	
Have you identified all Single Points of Failure? Do you repeat this process regularly? Do you have procedures to handle the recovery should one fail?	
Automation - how do you verify it when moving to new software levels? What do you do if it breaks - do you have the skills to manage manually? Do you know what it should do so you can do that manually until it is fixed?	



CPU considerations

Feature	Using?
Use latest technology you can afford - each new processor generation has availability/reliability and dynamic upgrade/downgrade improvements over previous generations	
Exploit concurrent upgrade capabilities like CUoD (all types of PUs), On/Off Capacity on Demand (all PU types), CBU (CPs only), Concurrent PU conversion, I/O upgrade plan ahead, and Concurrent Memory Upgrade - See GM13-0257-01 for summary of capabilities.	
Ensure every multi-path critical device is configured with <i>at least</i> two channels, through different switches/directors. Use CHPID Mapping Tool to obtain minimum single points of failure, both within the CPC and within the channel subsystem	
Ensure Driver levels are kept up to date, and install Hiper MCLs in a timely manner (nearly all patches are non-disruptive) - establish procedures and communication with hardware service representative	
Implement <i>at least</i> two HMCs and at least two HMC Phone Servers, so that CPCs can report a problem even if one HMC is in the middle of downloading patches	
Consider placing two LPARs from each sysplex on each CPC - ensures that MIPS on that CPC will still be available to the sysplex even if 1 LPAR is shutdown. Use IRD to optimize capacity allocation across LPARs	
Pre-define spare LPARs on z990 (partition name of *) - can then be renamed and used without a POR - requires GA3 and z/OS 1.6	
Every LPAR should have some Reserved CPs in LPAR Profile to allow more CPs to be added non-disruptively. On a z990, INITIAL+RESERVED can = 24 (16 if < z/OS 1.6), even on an A08...	
Enable Automatic I/O Reset Facility in CPC Reset profile	
On z900 or earlier, configure sufficient HSA to allow for Dynamic I/O Reconfigurations. On z890/z990, set the MAXDEV value in HCD high enough to provide I/O configuration definition growth	
Don't configure an LPAR with extremely low weights (<5% of CPC) in a production sysplex	
LPARs should either be successfully IPLed and running, or they should be DEACTIVATED. Don't have LPARs that are ACTIVATED but not IPLed, or ones that have failed IPLs that are not addressed (typically sysprog test LPARs)	
Capacity planning - if consolidating onto a smaller number of larger CPUs (z990 for example), do you have sufficient capacity to failover in case of a planned or unplanned CPU outage?	
If changing the model type of a CPC (2064 to 2084, for example), remember to vary off all CTCs talking to that CPC before the upgrade to avoid boxed devices.	

Switches, directors, and timer considerations

Feature	Exploiting?
Always spread paths to critical devices across more than one director/switch	
Define switch Control Unit Port in HCD, to ensure errors can be reported back to the operating system	
If appropriate, order High Availability configuration when purchasing FICON switch	
Use Model 2 Sysplex Timers - support on Model 1s ran out at the end of 2003. See Redpaper on migration options at: http://www.redbooks.ibm.com/redpapers/pdfs/redp3666.pdf	
Implement the Expanded Availability Configuration with Sysplex Timer Model 2s	
Especially when using recent technology, ensure that microcode levels are kept up to date.	
Next generation HMC will not support console application for Sysplex Timer or ESCON Director, so make sure you have an alternate available	
Ensure that fibre quality and DB loss are within IBM-recommended limits	
If you are experiencing high response times on FICON CTCs, a patch will be available that may address this	
z/OS 1.5 adds ability to re-enable Index VTOCs without taking the volume offline to all sharing systems	
z/OS 1.5 enhances REFORMAT command to support dynamically extending VTOC and Index VTOCs	



CF considerations

Feature	Exploiting?
CFLevel 11 is only supported CF Level on 9672 G5/G6. Ensure you are running this level and remember that upgrading to this level is disruptive	
CFLevel 12 (and 13 soon) is only supported CF Level on 2064 and 2066. Ensure you are running this level and remember that upgrading to this level is disruptive. Also, remember that migrating to CFLevel 12 can significantly increase your structure's storage requirements.	
Most CFLevel and CFCC service level upgrades should no longer require a POR following z990 GA3 (or z890 GA1)	
CFLevel 13 is not expected to significantly impact structure sizes compared to CFLevel 12	
If you are a large DB2 data sharer, especially if you have both data and indexes in the same GBP, install CF Level 13	
MUST have at least 2 CFs, even for Resource Sharing	
CFs must have sufficient capacity (CF Links, storage, MIPS) to take over workload from a failed CF. Should not run higher than 50% busy (max)	
If doing data sharing and not using SM Duplexing, at least one CF must be failure-isolated and all 'category 3' structures should be in that CF. See WSC FLASH 98029	
Regularly make sure that all structures are in the CF that you expect them to be in, especially if one CF is an ICF - msys for Operations helps with this; also SETXCF START,REALLOCATE in APAR OA03481 (z/OS 1.4 and later).	



CF considerations

Feature	Exploiting?
Each CF should have enough storage to hold <i>all</i> structures, especially if you only have 2 CFs - but remember to allow for duplexed structures. Remember Control Storage if CF still CFLevel 11 - check available control space before attempting to empty one CF into the other	
Monitor CF and structure storage utilization as you add structures and upgrade CF Levels	
CF with 1 CP should not run at >50% CPU utilization. If there is more than 1 CP, higher utilization is possible without impacting performance	
Must consider impact of high CF CPU utilization on CF response times, especially during CF or system recovery	
Should be <i>at least</i> 2 sender or peer links from every CPC to every connected CF - with z990 size CPCs, it is no longer safe to assume that 2 links will provide acceptable performance - must monitor using RMF or similar	
If possible, avoid the use of shared engines for production CFs - See ITSO Hint and Tip entitled "Use of Shared Engines for Coupling Facilities"	
Use highest speed links possible, especially for structures with large data transfers - don't permit path busy conditions caused by shared CF links to exceed 20%.	
If considering using System Managed Duplexing, refer to White Papers http://www.ibm.com/servers/eserver/zseries/library/techpapers/gm130103.html and http://www.ibm.com/servers/eserver/zseries/library/techpapers/pdf/gm130540.pdf	



z/OS Health Checker

Feature	Exploiting?
Use z/OS HealthChecker either in SA/390, msys for Ops, or standalone and ensure you have the latest level - download from http://www14.software.ibm.com/webapp/download/search.jsp?go=y&rs=hchk	
Once you have the HealthChecker running cleanly, schedule it to run every day on every system . You should only ever have to take manual action if there is a non-zero RC	
If you have automation that addresses some of the things HealthChecker monitors for, modify it to work with the WTOs that HealthChecker can put out	
HealthChecker future plans....	



z/OS considerations

Feature	Exploiting?
Paging volumes - should be high performance, support DASD fast write, and use <i>dynamic</i> PAV if available (z/OS 1.3 and later only)	
Page data sets limited to 4GB, however PAV support allows you to place multiple page on the same volume without losing performance advantages of SUSPEND/RESUME	
z/OS 1.3 and later protects page data sets from use by multiple systems by placing system name and other information in page data set and checking this at IPL	
Stripe standalone dump data sets across multiple volumes (4 seems to be optimum number) using AMSDADD Rexx exec (delivered in SAMPLIB) to minimize SAD duration	
Extended format sequential data sets larger than 64K tracks will be useable for Stand Alone dumps - supported in z/OS 1.6 - rolled back by APAR OA04140	
Check that RNLs for RACF resources conform to guidance in GRS Planning manual	
GRS Star - higher performance, superior availability and recoverability compared to GRS Ring. Considered a pre-req for large sysplexes	
Enable GRS SYNCHRES option in GRSCNF member prior to z/OS 1.6 - in 1.6 this is default	
Exploit IBM-provided automation in SA/390 or msys for Ops	
EXITs - eliminate where possible (is the exit still needed?), for those that remain, use Dynamic Exit support where possible.	
OPERLOG - used for sysplex-wide problem determination - can be used together with syslog	
LOGREC - sysplex-wide problem recording medium that never fills	
z/OS 1.5 Security Server adds support for dynamic update of RACF Templates (avoids an IPL) and prevents accidentally installing downlevel templates	
z/OS 1.6 Security Server adds support for dynamic updates of RACF Class Descriptor Tables (avoid an IPL)	
z/OS 1.6 Security Server removes the need for a router table entry (ICHRFR0X) for most user-defined classes (avoids an IPL)	



z/OS considerations

Feature	Exploiting?
Place JES2 Checkpoint in CF - equivalent-to-better performance, especially with mixed CPC sizes, and elimination of Reserves	
JES2 Health Monitor started automatically in z/OS 1.4 and later - monitors health of JES2 from separate address space. Use \$J commands to communicate with it.	
z/OS 1.4 added checking to JES2 to avoid corruption if the wrong checkpoint data set is accidentally used (HASP863)	
Sysplex HFS sharing - improved flexibility and availability for end-users and system programmers	
New REMOUNT option lets you unmount and remount a sysplex-shared HFS without having to unmount all filesystems under the one being remounted - APAR OA02584	
Use of system symbols in Parmlib, Proclib, VTAM, TCP, NetView, OPC, AOC, Clists to minimize multi-system maintenance overhead - with z/OS 1.4, you can have up to 800 symbols	
Exploit Console Restructure feature in z/OS 1.5 and available as a feature on top of z/OS 1.4 - designed to minimize impact of WTO buffer shortages	
z/OS 1.6 will add the ability to have multiple GTF traces running concurrently	
Prior to z/OS 1.6, DPS array initialization during MSI was a serial process. Changes to IOS in z/OS 1.6 parallelize this process, resulting in faster IPLs, especially with large I/O configurations - available back to 1.4 via APAR OA07335	
z/OS 1.6 moves EXCP control blocks above the line - will help relieve constraint below the 16MB line	
In z/OS 1.6, size of buffers for SMF data can be increased to 1GB from the prior maximum of 128MB (BUFSIZEMAX in SMFPRM) - in 1.5 use APAR OW56001 and USERMOD.	



XCF considerations

Feature	Exploiting?
There should be 1 XCF structure per Transport Class per CF. Every structure should be defined as both PATHIN and PATHOUT. If you already have CTCs defined, use those in addition to CF structures, especially if there is a large distance between systems	
Ensure every transport class has at least two paths defined	
Place CDSs on dedicated, high performance volumes with no RESERVEs	
Use a backup method that does not cause RESERVE problems - consider Disaster Recovery requirements	
Place primary, alternate, and spare CDSs on <i>separate devices and separate physical control units</i>	
Place the primary sysplex, CFRM, and LOGR CDSs on different volumes	
Ensure there is automation in place to add a spare CDS if primary or alternate is lost	
z/OS 1.6 increases the maximum number of members in an XCF group from 1023 to 2047 - OA 04034 rolls support back to z/OS 1.4	



XCF considerations (cont)

Feature	Exploiting?
If APAR OA03481 is applied, always use the SETXCF START,REALLOCATE command to repopulate a CF - if not, use the POPULATECF command	
Implement automation to check that each structure is in the first CF in its preference list - REALLOCATE will do this for you as well.	
Implement automation to act on potential hang situations (messages IXL040E and IXL041E).	
If XCF message rates above 1000/sec, see Setting Up a Sysplex for tuning recommendations, then see Mark Brooks XCF Tuning presentation on SHARE Web site.	
Use enhanced D XCF,PI,STRNM=nnnn command to get response time information for XCF paths - this information is also available in the RMF SMF records, but not in the RMF PP reports.	



CFRM considerations

Feature	Exploiting?
Ensure that System-Managed Rebuild is enabled - use D XCF,C,TYPE=CFRM command (should say SMREBLD(1)).	
Enable Structure Full Monitoring for appropriate structures and set up automation to monitor for resulting messages (to alert appropriate tech support staff)	
AutoAlter enabled for <i>recommended</i> structures. Plus monitor for messages so that CFRM policy can be adjusted accordingly	
Always update the CFRM policy any time a structure size is changed using SETXCF ALTER or Auto Alter	
All structures should have at least 2 CFs in preference list	
Ensure a CFRM policy naming convention is used so that policy changes can be easily backed out	
Use appropriate structure sizes - validate with CFSizer (for rough figure) or check RMF reports (for more accurate indication if existing structure is large enough)	
Never use REBUILDPERCENT > 1 - recommend to allow REBUILDPERCENT to default	
z/OS 1.4 delivered further enhancements to reduce the elapsed time for structure rebuilds	



Logger considerations

Feature	Exploiting?
Refer to recent ITSO Redbook on System Logger (z/OS Systems Programmers Guide to: System Logger (SG24-6898)) for detailed information on Logger usage and tuning	
Ensure LOGR CDS is at HBB7705 format level (D XCF,C,TYPE=LOGR) - get this by defining LOGR CDS with SMDUPLEX option	
If all systems > z/OS 1.2, and at least one system is at z/OS 1.3, most log stream attributes can now be changed while the log stream is connected	
Ensure that APAR OW51854 is applied, and set up automation to monitor for IXG310I/IXG311I/IXG312E messages - all relate to offload problems (APAR integrated into z/OS 1.4 and rolled back to OS/390 2.10)	
z/OS 1.5 adds Logger Task monitoring and messages IXG271I, IXG272E. See the section entitled "Offload and Service Task Monitoring" in Setting up a Sysplex for guidance on handling these messages	
In z/OS 1.5, System Logger address space will be automatically restarted unless it is terminated using FORCE ARM (rolled back to 2.10 - 1.4 via OW53349)	
Set up automation to monitor for IXG267I message - issued at data set allocation time and indicates that offload or staging data sets do not have the correct share options (3 3)	
z/OS 1.4 increased number of DASDONLY connections on a single system from 512 to 1024. (APAR OW51437)	
z/OS 1.2 added support for use of System Managed Duplexing with CF Log streams	
Never delete offload or staging data sets manually without first checking to ensure Logger has no knowledge of those data sets - always do a DELETE LOGSTREAM to remove log stream definitions and all associated data sets.	



Logger considerations

Feature	Exploiting?
Put logstreams with similar average buffer size and request rates in the same structure. Do <i>not</i> mix idle and busy logstreams in same structure	
Even if you don't intend to use staging data sets, specify STG_DUPLEX(YES), DUPLEXMODE(COND) (or apply the PTF for OA03001) and at least STG_SIZE and STG_DATACLAS for every log stream	
Use data class with 24K CI Size for all offload data sets - staging data sets must still have 4K CI Sizes	
When using CF Log streams and duplexing to DASD, ensure that staging data sets are sufficiently large that they are not hitting the High Threshold before the log streams in the CF	
Have at least 2 different systems connected to each LOGR structure - allows peer recovery in case of a system failure	
For CICS DFHLOG, IGWLOG, and RRS Active log streams, monitor offload data sets - if sequence number is large, adjust HIGHOFFLOAD and LOWOFFLOAD or structure size	



DFSMS features

Feature	Exploiting?
Use DFSMSHsm Multiple Address Space HSM - lets you break HSM tasks across multiple address spaces, running higher priority tasks (like recall) in one address space, and lower priority tasks in another (OS/390 2.10)	
Enable Fast Subsequent Migration, but be aware that this will result in larger MCDS	
Use DFSMSHsm Common Recall Queue - balances HSM recall processing across hosts in a sysplex and provides persistency of recall requests should the HSM address space fail or be stopped - see 2.3.1 in HSM Stor Admin Guide	
z/OS 1.6 adds support for multiple HSM hosts to concurrently run all phases of Secondary Space Management, except for L1 DASD to L2 DASD migration	
OS/390 2.10 increased number of concurrent DFSMSHsm backup tasks from one per system to up to 64 per system. You can also backup data sets directly to tape. See SETSYS DSBACKUP keyword. HOWEVER, we recommend a max of 6 tasks per function per host. To increase throughput, increase the number of hosts running each function, not the number of tasks.	
Use DUMPIO(3 4) not (4 4) in HSM	
If there are a very large number of HSM-managed data sets or volumes, use VSAM/RLS for CDSs to provide improved performance for CDS-intensive activities	
z/OS 1.5 adds the ability to use CDSQ or CDSR serialization with VSAM Extended Attribute HSM CDSs (removes requirement for VSAM/RLS)	
For more tuning suggestions, refer to Chapter 2.5 of DFSMSHsm Implementation and Customization	
If you still have a DFSMSHsm performance bottleneck, IBM offers a fee service to assist with DFSMSHsm Tuning - ask for Stan Kissinger	



SFM considerations

Feature	Exploiting?
USE SFM. Single change that would have eliminated the largest number of multi-system outages	
If the policy is defined and started, SFM can: <ul style="list-style-type: none"> ✎ Partition a system out of the plex if status update is missing ✎ Partition a system out of the plex if XCF signalling to the system is lost ✎ Assist with V XCF,xxx,OFFLINE processing ✎ Control structure rebuild processing based on REBUILDPERCENT 	
Don't set INTERVAL in COUPLExx - let system determine it dynamically	
You must have a CF to be able to use SFM's system isolation function	
Set CLEANUP interval to 15 seconds - should be sufficient for planned shutdowns	
Do <i>not</i> specify PROMPT in the SFM - use ISOLATETIME(0) instead to get sick system immediately removed from sysplex	
Don't specify OPNOTIFY- let it default to the calculated INTERVAL plus 3 seconds	
Ensure that "Enable I/O Interface Reset" option is enabled in CPC Reset profile	



DFSMS features

Feature	Exploiting?
If you have >1 HSMplexes in the same sysplex, implement DFSMSHsm Single GRSpplx support to remove false contention on the CDSs	
Enable DFSMSHsm Secondary Host Promotion (OS/390 2.8)	
Specify RESTART='(a,b)' on HSM startup JCL to automatically restart HSM following a failure	
Exploit DFSMSHsm support for creating duplex tapes	
DFSMSHsm supports Concurrent Copy for all data sets, both SMS-managed and non-SMS managed	
DFSMSHsm in z/OS 1.5 adds copy pool support that is used by DB2 V8 to take point in time copies (using FlashCopy) of an entire DB2 subsystem	



DFSMS features

Feature	Exploiting?
Enhanced Catalog Sharing - performance equivalent to non-shared DASD environment - OS/390 2.10 added support for user managed rebuild for improved recovery from CF problems	
Ensure that SYSIGGV2 and SYSZVDS ENQs are handled "correctly" in GRS RNLs - see GRS Planning manual for details	
Data set separation support feature (z/OS 1.3) allows you to give SMS a list of SMS-managed data sets that should not reside on the same physical control unit	
OAM sysplex support - ability to access OAM data from any system in the sysplex, and remove OAM affinity to just one system	
In z/OS 1.5, GDS reclaim processing default can be changed without an IPL	
Warning messages (IEC361I) issued if catalog space utilization exceeds a user-specifiable threshold - set using F CATALOG,NOTIFYEXTENT (z/OS 1.3)	



DFSMS features

Feature	Exploiting?
"Space Constraint Relief" (OS/390 2.4) decreases chances of a data set allocate or extend failing due to lack of DASD space - see Data Class definitions. Also, see "Reduce Space By %" parameter also in Data Class definition	
Specify Overflow and Extend Storage Groups for production storage groups to reduce out-of-space abends (z/OS 1.3)	
z/OS 1.5 adds ability to automatically kick off space management when a storage group reaches its high threshold. Also, SMS issues WTOs that can be used to kick off automation actions	
Also in 1.5, SMS will place data sets in storage groups in the sequence they are specified in the ACS routines. The first specified storage group will be used till it fills, then the next one till it fills, and so on	
Ability to rename duplicate data sets - support added in OS/390 2.10. Protected by RACF FACILITY profile called STGADMIN.DPDSRN.*	
z/OS 1.6 adds ability to take existence of PAV into account when allocating SMS-managed data sets (new PAV keyword in Storage Class)	
Exploit Enhanced Integrity for sequential data set support in z/OS 1.5 - see member IFGPSEDI in SYS1.PARMLIB	
OS/390 2.10 delivered improvements in trace facility for VSAM data sets	
Peer to Peer VTS support provides ability to keep a duplex copy of volumes in a Virtual Tape Server - delivered in OS/390 2.10	
DFSMSdss support for data set extent reduction while DEFRAging a volume - added in OS/390 2.10 (CONSOLIDATE keyword on DEFRAg)	
Data set level FlashCopy (requires FlashCopy V2) available with DFSMSdss in z/OS 1.3 and rolled back to OS/390 2.10.	



DFSMS features

Feature	Exploiting?
z/OS 1.4 creates a new address space, SMSPDSE, to contain all PDSE-related control blocks, removing them from ESQA and ECSA	
z/OS 1.6 adds the ability to restart one of the PDSE address spaces (SMSPDSE1), potentially avoiding an IPL if a restart is required. Must specify PDS_RESTARTABLE_AS(YES) in IGDSMSxx member	
The SMS SCDS data set can be recreated from the ACDS using the SETSMS SAVESCDS command - new in z/OS 1.5 (See dfp Stor Admin Ref for details)	
Coupled XRC supports lets you XRC more volumes with time consistency across the full set. Also, XRC parameters can now be placed in parmlib members rather than having to issue XSET (z/OS 1.3)	



Network/sysplex considerations

Feature	Exploiting?
See Comm Server session this week for MUCH more information	
VTAM Generic Resources for: CICS DB2 IMS TSO NetView Access Services	
VTAM Multi Node Persistent Sessions for: CICS (CICS TS 2.2 supports signon persistence) IMS Rapid Network Recovery	
TCP/IP can use Dynamic XCF for intra-sysplex IP communication - in 1.7, TCP will select the optimal paths packet routing - will be able to use OSA Express Gigabit Ethernet as well as XCF	
TCP/IP use XCF signalling to provide auto notification when TCP/IP stacks join or leave the sysplex	
Exploit TCP/IP VIPA, Dynamic VIPA, and DVIPA takeover and takeback to allow movement between systems in the sysplex	
Exploit Sysplex Distributor to balance work across systems. Enhancements in z/OS 1.6 and more coming in 1.7.	
TCP/IP uses the Coupling Facility to coordinate Sysplex-wide IPSEC security associations and assignment of client port numbers when a Sysplex-wide source VIPA address is in use	



CICS considerations

Feature	Exploiting?
Make sure all application programs (COBOL, PL/I) are running above the line	
Make sure all application programs (COBOL, PL/I) are running above the line (worth saying it twice!)	
Enable CICS storage protection to protect CICS code and control blocks from errant user programs	
Make sure CICS transaction isolation is used to protect CICS transactions from each other	
If possible, use the TCP/IP sockets support provided with CICS rather than the CICS sockets support provided with Communications Server	
Define CICS TORs as VTAM Generic Resources	
Define two cloned TORs for each CICSplex on each MVS image	
Use a separate Sockets Owning Region (SOR) on each image	
If using CICS Java support, define a separate Java Owning Region (JOR)	
Make sure there is enough capacity (MIPS, below the line storage, MAXTASKS, DB2 and DBCTL threads, etc) in the CICS regions to handle failover should a z/OS image fail	
CICS supports DB2 group attach in CICS TS 2.2 - be sure to specify RESYNCMEMBER(YES) if you restart CICS and DB2 on another MVS image in case of failures	



CICS considerations

Feature	Exploiting?
Use CICS Shared Temporary Storage in CF to eliminate non-recoverable temporary storage as an affinity - also, CICS TS 2.2 supports System Managed Rebuild and Duplexing for this structure	
Use CICS Named Counter Server in CF to provide unique counter values across multiple CICS regions - also, CICS TS 2.2 supports System Managed Rebuild and Duplexing for this structure	
Use CICS Global ENQ/DEQ to propagate application ENQs for CICS resources across multiple CICS regions in a CICSplex, removing the affinity to a particular region	
Use CICS Data Tables in CF to share information currently held in a CICS data table across multiple CICS regions - also, CICS TS 2.2 supports System Managed Rebuild and Duplexing for this structure	
Use CICS VSAM Record Level Sharing to remove CICS FORs as a single point of failure	
Use CICSplex System Manager (part of CICS TS since CICS TS 1.1) for CICSplex management and dynamic transaction routing	
Use CICS Transaction Affinities Utility or CICS Interdependency Analyzer to identify, remove, or manage transaction affinities	



DB2 considerations

Feature	Exploiting?
Read DB2 chapter of SG24-6061, now available as a Redpiece on redbooks Web site	
Use DB2 data sharing	
Always duplex DB2 GBPs. If the DB2 Lock or SCA structure shares a single point of failure with a connected DB2, use System Managed Duplexing for those structures	
If restarting DB2 on another image after a system failure, use the DB2 RESTART LIGHT option added in DB2 V7 - Note that RESTART LIGHT will not connect to other resource managers to resolve IN-DOUBT work. THIS IS CHANGED IN DB2 V8	
Prior to z/OS 1.6, if RRS did NOT fail, DB2 must restart using the same RRS instance. If RRS DID fail, DB2 can restart on another image, and all involved resource managers must start on the same system UNLESS you are z/OS 1.2 or later (removes restriction on grouping). z/OS 1.6 lets DB2 and other RMs start anywhere, even if the original RRS is still up.	
Partition databases so the whole database does not need to be offline for maintenance. DB2 V6 allows you to change partition boundaries and only impact the affected partitions, NOT the whole tablespace (SQL also supports query parallelization across partitions - the more partitions, the more potential parallelization)	
Prior to CFLevel 13, use different buffer pools for data and index portions of databases. CFLevel 13 addresses high CF CPU utilization that can occur during large online reorg if indexes and data share the same GBP	
Online reorg allows you to reorg tablespaces or partitions transparently to the user (slight pause while DB2 switches to the new copy)	



DB2 considerations

Feature	Exploiting?
In every new DB2 release, the performance and function of DB2 utilities is enhanced to provide better data availability - do your Application developers and DBAs exploit these new capabilities?	
As DB2 reliability and availability improve, an increasing percent of outages are caused by application problems. Do you have tested and consistent plans for restoring databases to a point in time?	
Partitioned tablespaces in DB2 V8 allow you to Alter partition boundaries and rebalance data with Online Reorg while access to data is maintained. SQL also supports query parallelization across partitions - the more partitions, the more potential parallel access.	
With DB2 V8 you can add columns to tables concurrently using Alter. The changes take effect after Online reorg, without the need to drop and redefine tables.	
DB2 V8 increases the number of DSNZPARMs values that can be altered without having to restart DB2.	
IRLM 2.2, shipped with DB2 V8, places all locks above the 2GB bar, and can support up to 100,000,000 concurrent locks, 16 times more than IRLM 2.1. This allows more locks AND frees up storage in ECSA that was previously used to hold lock info (if you specified PC=NO).	
Change IRLM timeout interval from default of 300 seconds to 10-30 seconds with F IRLM,SET,TIMEOUT=xx,DBMS.	
Especially for batch jobs, do NOT let them do many many updates without issuing a COMMIT. Monitor and report on any job that holds over xxxx locks - be especially careful with ad-hoc jobs	
Don't forget about user access to DB2. Use TCP/IP and SNA features like Dynamic VIPA and Generic Resources to mask outages. See DB2 chapter in SG24-6061 for more info	



IMS considerations

Feature	Exploiting?
Have multiple cloned IMSs - cloning critical so apps can run in any of the IMSs	
Share all databases, and use Fast Database Recovery (FDBR) region to free up locks in case of a failure - enhancements in IMS V9 for FDBR	
Use shared queues to protect from failure of a control region - If using shared queues with IMS V8 or later, enable CSL and Sysplex Terminal Management	
Use Extended Terminal Option (ETO) to add IMS terminals dynamically	
Use VTAM Generic Resources to mask control region unavailability from logging-on users - if using VTAM GR with cloned regions, do NOT use RNR features	
If using VTAM GR, set up IMS so that it releases the terminal back to VTAM if IMS dies. The method to do this varies with IMS version - look up GRAFFIN to find how it applies to your version of IMS.	
For TCP/IP connected devices, use Sysplex Distributor and/or Websphere Edge Server	
IMS Connect (now part of IMS) supports connections to multiple IMS Control Regions using XCF - IMS Connect does not have to be on the same z/OS system as IMS	
Use IMSGROUP to allow spreading of BMPs across data sharing members	
If IMS V8 or later, implement IMS Common Service Layer and Global Online Change	
Define more than one RACF TCB to avoid serialization of RACF calls - RCFTCB in IMS Parmlib	
Exploit High Availability Large databases (HAL DBs) to let you manage parts of large databases without impacting the whole database	
Exploit concurrent backup and reorg capabilities - IMS tools provides online reorg for ALL IMS database types, but with minimal outage	
IMS V9 removes need to IPL to install Type 4 SVC and Resource Manager Cleanup routine	
If you need to be able to restore IMS and DB2 DBs to same point in time, consider IMS Database Recovery Facility (DRF) - allows restoring IMS DB to any point in time	



Testing (that horrible word)

Feature	Exploiting?
Do you have a test environment? Is it used for destructive testing or just to test things that you know will work?	
How closely does the test environment reflect the production one?	
What percentage of bugs (software or hardware) are found in the test system? What percentage in production?	
When you hit a bug in production, do you analyze why it wasn't found in the test environment? Do you adjust the test environment to rectify this?	
Do you have a Training environment for operators and systems programmers or do they try out their procedures on the production system?	
Do the operators <i>understand</i> their procedures? Do you provide training so they understand what is happening (IBM has an offering called Parallel Sysplex Training Environment that provides this capability)	
Do you test all changes you will make to the production system? Do you test the backout procedures?	



Other MVS-related exploiters

Feature	Exploiting?
Utilize WLM-Managed Batch Initiators where appropriate - can adjust to provide an "appropriate" number of batch initiators if work gets re-routed around a system outage - improved in z/OS 1.4	
Use WLM Resource Affinity Scheduling together with automation to control where jobs can run	
Set up an OPC Hot Standby Controller to take over in case of planned or unplanned outage of the OPC Controller	
Exploit the OPC Interface to WLM to alter WLM Service Class for selected jobs/applications	
Sysplex enhancements in System Automation for OS/390 (also available to non-SA/390 customers as "msys for Operations")	
GRS Enhanced Contention Analysis (enhanced D GRS command)	
SNA Console support - full function MVS console (except NIP support) from any SNA or IP screen	
Use DFSMS Data Set Separation feature to keep software-duplexed data sets on separate control units	

