

IBM  
IBM Washington Systems Center

# z/OS Console Restructure

Riaz Ahmad  
IBM Washington Systems Center  
Gaithersburg, Maryland

© 2005 IBM Corporation

IBM Washington Systems Center

For the development effort, Special Thanks to:

Scott Fagen  
David Hom  
z/OS Core Technology Design and Development  
IBM Corporation

© 2005 IBM Corporation

## Topics

- Problems to be solved
- Staging Plan
- Un-enhanced message delivery
- Enhanced message delivery
- Getting ready for stage 1a
- Experiences
- Getting ready for stage 1b
- Console State Data Solution
- Getting ready for stage 2
- 1-byte Migration/Mitigation

## Problems to be solved

- Message delivery:
  - ▶ Buffer shortages
  - ▶ Varying speeds and feeds
- Synchronizing console state information:
  - ▶ Data replicated on every system in the sysplex
  - ▶ State changes, system join/leave cause flurries of activity
  - ▶ More systems => elongated startup, shutdown and recovery time
- Limit of 99 MCS/SMCS/Subsystem consoles in a sysplex:
  - ▶ Installation constraint

## Staging Plan

- Enhancements to be staged:
- Stage 1a solves problem 1
  - ▶ Delivery in z/OS: 1.4.2 feature, 1.5. base
- Stage 1b delivers improved RAS, IPL/recovery time improvements
  - ▶ Scheduled for z/OS 1.7
- Stage 2 solves problems 2 & 3
  - ▶ Future release of z/OS

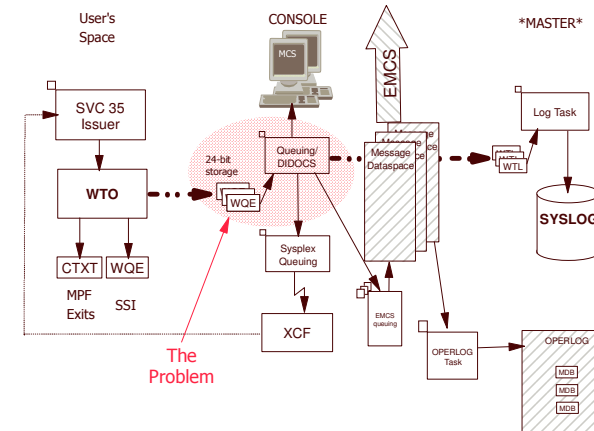
## Console Availability - Staging

- Stage 1a: "Message delivery enhancements"
- Stage 1b: "RAS and other enhancements"
- Stage 2: "Console State Data Solution"

## Console Availability - Staging

- Stage 1a: "Message delivery enhancements:"
  - ▶ z/OS 1.4 feature (JBB7727) - optional
  - ▶ Delivered with base z/OS 1.5 - not optional
  - ▶ Targets WTO buffer shortage problems
    - 40-45% of outages that Consoles Restructure intended to address
  - ▶ Mixed levels of z/OS supported:
    - Compatibility PTFs required
    - Can re-IPL single image to fall out of Stage 1

## Un-enhanced message delivery



## Un-enhanced message delivery...

- Classic producer/consumer problems:
  - ▶ Runaway application can wipe out system
  - ▶ Large systems can overcome small systems
  - ▶ Single task manages all delivery decisions
- Prone to backup from:
  - ▶ A particular console
  - ▶ SYSLOG
  - ▶ Un-ended MLWTO
- Importance of messages 'inflated':
  - ▶ Attempts to deliver message to all destinations with no regard to system impact
- Importance of MCS consoles 'inflated':
  - ▶ Real action occurs in Ops packages, Log browsers
    - Subsystem Interface (SSI)
    - EMCS consoles
    - SYSLOG/OPERLOG browsing

## Enhanced Message Delivery

- Design Points:
  - ▶ No single points of failure!
    - Do not deliver messages to consoles if the system is at stake
    - Messages will be logged
  - ▶ Fault tolerance:
    - Failure of a single component does not take down the system
  - ▶ Provide 'pressure relief valves' to mitigate effects of fast producers/slow consumers
  - ▶ Isolate MCS console processing from the overall mechanism
  - ▶ Isolate logging from other queueing tasks

## Enhanced Message Delivery...

### ■ Design Points:

- ▶ Process as much as possible under the caller's unit of work:
  - Up to and including calling XCF to send message
- ▶ Utilize 'modern' XCF buffer management:
  - Makes local and foreign message processing much more common
- ▶ Mitigate impact of poorly behaved MLWTO issuers on the system

## Enhanced Message Delivery...

### ■ Pressure relief valves:

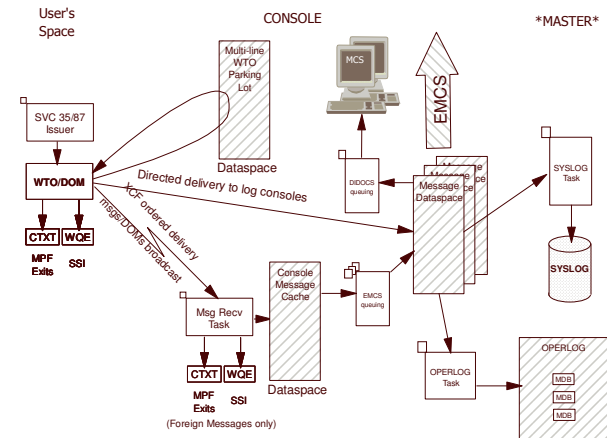
- ▶ All messages fed into a common message cache:
  - Circular
  - Dataspace resident (>>16MB messages)
  - Slow consumers 'fall off the end'
- ▶ (S)MCS consoles now receive messages via an EMCS console:
  - Still have same triggers and messages,
  - MCS WQE buffer shortages can still occur, but impact is isolated to MCS consoles

## Enhanced Message Delivery...

### Pressure Relief Valves:

- ▶ MLWTO "Parking Lot"
  - MLWTOs held in a dataspace until:
    - Last line received
    - Message timeout exceeded
    - EOT/EOM of issuer

## Enhanced Message Delivery...



## System Console Availability

- Additional availability enhancement for the System Console:
  - ▶ Brings System Console into 'PD mode' when no consoles in a particular group are active
  - ▶ Use the AUTOACT keyword on the CONSOLE statement for the System Console (SYSCONS)
  - ▶ Use the AUTOACT keyword on the VARY CN command
  - ▶ When no consoles in the specified group are active, the System Console is automatically placed in PD mode
  - ▶ Especially useful in environments where:
    - There are no MCS consoles
    - The base operating system is up, but the supporting subsystems (e.g. VTAM, TSO, SDSF) are not yet available

```
CONSOLE DEVNUM(SYSCONS) NAME(SYSCON&SYSCLONE)
        AUTOACT(SCONGRP)
        ROUTCODE(1-2)
```

## Getting Ready for Stage 1a

- Operational changes:
  - ▶ All consoles must be explicitly named:
    - CONSOLE NAME parameter now mandatory
  - ▶ ALTERNATE keyword no longer supported:
    - ALTGRP is sole mechanism for specifying back-up consoles
    - ALTCONS parameter removed from VARY command
  - ▶ Undelivered messages no longer detected:
    - UD keyword no longer supported on CONSOLE and HARDCOPY CONSOLxx PARMLIB statements
    - UD keyword no longer supported on VARY CONSOLE and VARY HARDCPY commands



## Getting Ready for Stage 1a...

### Operational changes:

- ▶ Hardcopy can now only be directed to SYSLOG and/or OPERLOG
  - DEVNUM parameter removed from HARDCOPY statement
- ▶ Hardcopy group (HCPYGRP) no longer supported
  - Can no longer specify on HARDCOPY statement
- ▶ "Re-route" parameter (R=) removed from the CONTROL Q command
  - Command will only remove message backlog from target console
  - Messages cannot be redirected to another console
- ▶ System Console MSCOPE now defaults to '\*\*' (this image)
  - (Used to be '\*ALL')
- ▶ 3 new CADs introduced by support:
  - Verify MAXCAD value is sufficient

## Getting Ready for Stage 1a...

### System behavioral changes:

- ▶ A request to delete a message by an exit (SSI/MPF) results in the message being physically deleted:
  - Message will not be logged
  - Visible on the subsystem interface (SSI) on the issuing system (only)
- ▶ Multi-line messages that are built dynamically using CONNECT processing will be "parked" in a dataspace until complete:
  - Not transmitted via XCF or queued to consoles or logged (SYSLOG/OPERLOG) until completed
  - Still sent to MPF, SSI and MTRACE as it is constructed
  - Un-ended messages will be ended at EOT of original issuing task or on timeout (dormant for more than one second)
  - Arrival at console destinations may be long after seen by exits or on SSI

## Getting Ready for Stage 1a...

- System behavioral changes:
  - ▶ Increased number of XCF signals, as well as increased use of local signaling).
- Recommendation of 3 Transport Classes to segregate message traffic according to message length:
  - ▶ CLASSLEN of 1024/8192/20480
  - ▶ CLASSLEN of 1024/8192/32768 with IMS OTMA
- Recommendation to override the LOCALMSG default:
  - ▶ MAXMSG of 1000 or 1500

## Getting Ready for Stage 1a...

- Compatibility/Migration/Coexistence
  - ▶ Before installing the Console Availability Feature (z/OS 1.4.2):
    - For OS/390 R10, z/OS 1.2 and z/OS 1.3:
      - Install APAR OW56244 (compatibility APAR)
    - For z/OS 1.4:
      - Install APAR OW56244 and SDSF APAR PQ73805
  - ▶ Check for other recommended service for JBB7727 and/or z/OS 1.5

## Experiences

- Internal changes affected OEM products:
  - ▶ STK, CA, PSI, BMC
- And IBM products:
  - ▶ SDSF, RMF, NetView
- Old hand-built WTO parameter lists fail new WTO processing
  - ▶ OA06083 allows IBM to provide an exit to temporarily correct parameter list
- Some performance anomalies:
  - ▶ Under extreme message loads, systems with feature can overrun down level systems
  - ▶ Messages can appear on foreign systems before issuing system.
    - Check LOCALMSG MAXMSG and increase value as appropriate
  - ▶ MCS WQE buffer shortages no longer stop system activity
    - MCS consoles can get hours behind
    - Log is up to date
  - ▶ Extended SSI processing times for foreign messages
    - IBM APAR OA08482 to address.
      - Provides ability to limit which subsystems will see foreign messages.

## Experiences ...

- More messages per unit time:
  - ▶ Sample system test run:
    - Sysplex of two 14-CPU machines
    - Ran 90 minute WTO flood (jobs that continuously issue messages) of single line WTOs
    - More messages per unit time (16X), but equivalent CPU utilization
    - Several internal locking constraints relieved, enabling greater throughput

	Lines of Syslog (per minute)	CONSOLE CPU utilization for 90 minutes (SU)	CPU/WTO
<u>z/OS 1.4</u>			
System D0	2248	32786K	161
System D2	2211	34768K	179
<u>z/OS 1.5</u>			
System D0	39840	511213K	143
System D2	32128	482030K	166

## Experiences ...

- Increased CPU consumption:
  - ▶ Several installations have seen increased CPU consumption in Console address space.
    - Especially noticeable when machine capacity is being stressed
  - ▶ Why? Two reasons:
    1. More messages are being processed per unit time
    2. More messages being 'multi-cast' around the sysplex

## Experiences ...

- More messages being "multi-cast" around the sysplex:
  - ▶ Console Restructure design changed to send all messages to all systems, with filtering decisions made at the target
    - Improved diagnostics, as the message cache contains
  - ▶ Turned out to be a poor design choice for some installations:
    - "Sham-plex" – not taking advantage of ability to route messages around the sysplex
  - ▶ Two APARs taken to address CPU consumption:
    - OA08482: New Function
      - Enables installations to prevent foreign WTOs from being presented on the Subsystem Interface (SSI)
    - OA09229: High CPU utilization in the CONSOLE address space after Console Restructure

## Experiences ...

- OA09229 improvements:
  - ▶ With all consoles set to MSCOPE=\*, CPU and message counts for foreign messages significantly reduced
    - Messages are still received via the XCF group, as the new WTO function still uses XCF to pass the message from issuing address space to the CONSOLE address space

	z/OS 1.4	z/OS 1.6	z/OS 1.6 +OA09229
Message Rate	x	9-11x	9-11x
CONSOLE CPU (per message)	y	y or less	y or less
Total CONSOLE CPU for the test (not per message)	0	+1%	+0.01%
Number of messages received on SYSMCS XCF group	150	900000	4000

## Console Availability - Staging

- ✓ Stage 1a: “Message delivery enhancements”
- Stage 1b: “RAS and other enhancements”
- Stage 2: “Console State Data Solution”

## Console Availability - Staging

- Stage 1b: RAS and other enhancements:
  - ▶ Enable deletion of EMCS consoles:
    - Installations can use a new program interface (similar to IEARELCN for MCS consoles) to delete inactive EMCS consoles, without reIPLing the sysplex
  - ▶ Internal task tracking RAS enhancements:
    - Improved PD for 'hang' situations
  - ▶ Program interface to obtain currently retained action messages and WTORs
  - ▶ Activate MONITOR independent of a console
    - SETCON MN command
  - ▶ Eliminated functions:
    - Use of 1-byte console ids on macros interfaces, and commands
    - TRACK command
- Scheduled for z/OS 1.7

## Getting Ready for Stage 1b

- One-byte console IDs no longer supported:
  - ▶ Can no longer be specified on:
    - WTO/WTOR
  - ▶ Subsystem console IDs are now 4 bytes in length
  - ▶ Migration console IDs no longer supported
    - MCSOPER
  - ▶ Recommend use of four-byte ids for INTERNAL and INSTREAM
  - ▶ Removal of 1-byte specification on operator commands
    - D C,CN= D PFK,CN= D R,CN= MSGRT RESET CN= SWITCH CN= VARY CN()
    - L=cc and L=cca no longer supported on commands
      - L=name and L=name-a are still supported
- TRACK command eliminated:
  - ▶ Following commands are also eliminated:
    - TRACK, STOPTR, CONTROL T, CONTROL D,U, CONTROL D,H, MSGRT TR=A
  - ▶ UTME keyword no longer recognized in CONSOLxx

## EMCS Console Removal...

- You can use `DISPLAY EMCS, ST=L` to obtain all the defined EMCS consoles in the sysplex
- Determine which consoles can/should be removed
- Modify 'SYS1.SAMPLIB(IEARELEC)' to remove the console definitions
  
- There is also a HealthCheck that can detect when you have exceeded a certain threshold:
  - ▶ CNZ\_EMCS\_Inactive\_Consoles
  - ▶ Default threshold = 10000 consoles

## Console Availability - Staging

- ✓ Stage 1a: "Message delivery enhancements"
- ✓ Stage 1b: "RAS and other enhancements"
- Stage 2: "Console State Data Solution"

## Console Availability – Staging

- Stage 2: "Console State Data Solution"
  - ▶ New "Distributed Mode" processing
    - Serialization granularity reduced to 1 ENQ per console, not 1 ENQ per "console class"
    - Enables more parallel activity without serialization bottlenecks
    - Data are correct on owning system, "lazy update" to other participants
  - ▶ Solves 50% of Console related MSOs
    - Console state data only maintained for active consoles
  - ▶ Will allow 99 active (250 defined) MCS/SMCS/Subsystem consoles per z/OS image
    - (n\*99 across the sysplex)
- Delivered in a future z/OS release (after z/OS 1.7):
  - ▶ No rollback to prior releases
  - ▶ Requires migration (operator command) to activate
    - All systems must be at required z/OS level
  - ▶ Limited 'reverse-migration' capability

## Console Availability - Staging ...

- Migration:
  - ▶ Preserves console state across the transition
  - ▶ Transition from old 'shared' mode' to 'distributed' mode
- Reverse migration (fallback):
  - ▶ In case of an unacceptable problem  
*not*  
A switch to flip back and forth
  - ▶ Needs to be fast to return to functional state
  - ▶ Will not restore entire console state:
    - Will restore (S)MCS consoles with console ids less than 100
      - Likely to be all the consoles defined prior to the original migration



## Getting Ready for Stage 2

- UCME index not equivalent to console ID
  - ▶ A UCME can contain any console ID
  - ▶ UCMEs not duplicated across sysplex
  - ▶ Access to UCME by console ID must be through IBM supplied service routine
- 1-byte console ids no longer provided in control blocks:
  - ▶ CIB, CSCB, ORE, WQE, XSA
- MCS, SMCS, and subsystem console ids are 4 bytes in length, in console class x'00'
  - ▶ INTERNAL = x'00000000'
  - ▶ INSTREAM = x'00000080'
- The limit of 99 consoles/sysplex is removed:
  - ▶ Up to 99 active MCS, SMCS and subsystem consoles per system are supported
  - ▶ Up to 250 MCS, SMCS and subsystem consoles can be defined per system.

## Getting Ready for Stage 2 ...

- The 'Console Switch' function is being removed
  - ▶ Value of the function has eroded over the years with automation managing the majority of messages
  - ▶ Messages more readily retrieved from SYSLOG or OPERLOG
  - ▶ Installation can just reactivate or activate another console, if needed

## Getting Ready for Stage 2 ...

- The Master Console no longer exists:
  - ▶ MSTCONS parameter is removed from the VARY command
  - ▶ DISPLAY CONSOLES,MASTER command will now report on the status of all consoles with MASTER authority
  - ▶ Route code 1 & 2 messages are not guaranteed to be displayed on a console. These codes were formerly 'forced' on the Master Console.
- Changes to console attributes no longer persist after the console is deactivated:
  - ▶ (S)MCS consoles return to Parmlib values
  - ▶ EMCS consoles return to values returned from OPERPARM segment

## Getting Ready for Stage 2 ...

- Operational Changes:
  - ▶ Console support mode specification added to the CON parameter of IEASYSxx:
  - ▶ ALTGRP parameter no longer supported on the CONSOLE statement.
  - ▶ NOCCGRP parameter no longer supported on the INIT statement.
  - ▶ INTIDS and UNKNIDS parameters supported on:
    - CONSOLE statement
    - DISPLAY EMCS command
    - VARY CN command
  - ▶ MSTCONS keyword no longer supported on the VARY *devnum* command.

## 1-byte Id Migration/Mitigation

### 1-byte Id Detection Tool:

- ▶ Available on any release with console restructure
- ▶ Gives IBM development, OEM Vendors, customers a way to find 1-byte interface usage
  - Includes migration IDs
- ▶ Identifies interface and using programs
- ▶ WTO, MPF, SSI, subsystem L=cca processor, MGCR/MGCRE, CONVCON, MCSOPER, etc.
- ▶ DISPLAY OPDATA,TR command displays usage to operator
- ▶ Parmlib interface provided to suppress known failures

## 1-byte Id Migration/Mitigation ...

```

CNZ1001I 14.19.39 TRACKING DISPLAY
STATUS=ON      INSTANCES=14    MAX=1000
---TRACKING INFORMATION---VALUE--JOBNAME  PROGNAME  -NUM-
CONVCON                65  IBMUSER  ISFMAIN   1
CONVCON                65  IBMUSER  ISFSTOP   2
MCSOPER: Obtain RICK   65  IBMUSER  ISFMAIN   1
MCSOPER: Release      65  IBMUSER  ISFSTOP   1
MGCRE: D T             65  IBMUSER  ISFMAIN   1
Parmlib Reader: ADYSET00  00  *MASTER*  ADYSETP   1
Parmlib Reader: COFVLF04  00  VLF       COFINIT   1
Parmlib Reader: IEFSSN00  00  *MASTER*  IEEMB860  1
Parmlib Reader: SMFPRM00  00  SMF       IFASTART  1
WTO: $HASP003          00  JES2     $HASPWTO  2
WTO: $HASP003 RC=(52   00  JES2     $HASPWTO  2
WTO: $HASP893          00  JES2     $HASPWTO  1
WTO: $HASP893 VOLUME  00  JES2     $HASPWTO  1
WTO: IEF677I WARNING  00  JES2     IEFNB903  1

```

### 1-byte Id Migration/Mitigation ...

- IBM has documented the use of the tracker in RETAIN:
  - ▶ See APAR II13752, also describes process for reporting 1-byte usage to IBM
- Current 1-byte usage list is maintained on the web:
  - ▶ <http://www-1.ibm.com/servers/eserver/zseries/zos/downloads/>
  - ▶ List will be periodically updated
- Install OA05596 which reduces instances to be tracked.

### Console Restructure - Summary

- Stage 1a solves message delivery problems
- Stage 1b to provide
  - ▶ Internal RAS / Problem determination enhancements
  - ▶ Ability to delete EMCS consoles
  - ▶ Operational changes in preparation for stage 2
- Stage 2 to solve the CONSOLE state data problem
  - ▶ Remove 99 console per sysplex constraint
  - ▶ Reduce serialization bottleneck on SYSZMCS global resources

## References

- **Announce material:**  
[http://www-306.ibm.com/common/ssi/rep\\_ca/2/897/ENUS203-132/ENUS203-132.PDF](http://www-306.ibm.com/common/ssi/rep_ca/2/897/ENUS203-132/ENUS203-132.PDF)
- **z/OS 1.5 Publications:**
  - ▶ **MVS Planning: Operations (SA22-7601-04)**  
<http://publibz.boulder.ibm.com/epubs/pdf/iea2g340.pdf>
  - ▶ **MVS Initialization and Tuning Reference (SA22-7592-07)**  
<http://publibz.boulder.ibm.com/epubs/pdf/iea2e240.pdf>
- **Important APARs:**
  - ▶ **II13752:** <http://www-1.ibm.com/support/docview.wss?uid=isgl1113752>
  - ▶ **OA05596:** <http://www-1.ibm.com/support/docview.wss?uid=isgl10A05596>
  - ▶ **OA10016:** <http://www-1.ibm.com/support/docview.wss?uid=isgl10A10016>
  - ▶ **OA08482:** <http://www-1.ibm.com/support/docview.wss?uid=isgl10A08482>
  - ▶ **OA09229:** <http://www-1.ibm.com/support/docview.wss?uid=isgl10A09229>
- **One-byte Id Tracker:**
  - ▶ <http://www-1.ibm.com/servers/eserver/zseries/zos/downloads/>