

z/VM 6.2 ESP Experiences

Dennis O'Brien

z/VM Chief Engineer

24 Apr 2012

MVMUA

Environment

- z/VM in multiple datacenters
- CMS applications, z/OS guests, Linux guests
- 2-node CSE clusters host production Linux guests.
- Due to time constraints, ESP testing limited to guests.
- Four first-level systems running z/VM 6.2 now. More to follow.

Single System Image

- Some teething pains with mismatched service levels during ESP.
- Directory syntax takes some getting used to.
- Put SFS servers on shared DASD, just in case.

Live Guest Relocation

- Beware of constrained systems (e.g. 512M guest on 128M host)
- Problems using VSWITCH under VSWITCH. Use dedicated OSA's for second-level z/VM systems that have their own virtual switches.
- Dedicated OSA's had EQID mismatch on backup OSA's (SR opened).
- Tested relocation of a disconnected CMS user. It works with shared CMS disks, but is not supported.

Changes Required for LGR

- Add OPTION CHPIDV ONE.
- Detach 190, 19D, and 19E just before Linux IPL in PROFILE EXEC.
- Replace virtual disk MDISK statements with DEFINE commands.
- Make sure VDISK USERLIM is high enough for your virtual disks.
- Define OSA's in SYSTEM CONFIG and specify an EQID.

Migration Tool

- Problem with TSAFVM when first-level and second-level node names are the same. Temporarily change second-level node name or don't use migration tool.
- Use caution when adding second-level system to ISFC. Even if you change the node name, the old name is in the CRR server.
- If you do change the name, don't forget owner information on CP-owned DASD.

Local Tools

- Examine impact of clusters and directory changes on tools and housekeeping jobs.
- Userid starts with either USER or IDENTITY.
- Some jobs run on every member (SPOOL file purge, audit file extract, time zone change).
- Some jobs run on one member (password expiration, password sync, directory backup, tape scratch).
- Location of SYSTEM CONFIG has changed.

VM:Secure

- Commands go to master unless LOCAL prefix is used. This even applies to AUDITEXT, END, and TAKEOVER.
- VMSECURE MAP shows SUBCONFIG name for private minidisks, but you LINK the IDENTITY name.
- VMXGNR puts all comments that occur after the last directory control statement for one directory entry after the USER or IDENTITY statement for the next directory entry. This is documented behavior. Result is that commented-out SUBCONFIG's end up as comments in unrelated directory entries.

VMXGNR Example – part 1

IDENTITY DTCSMAPI AUTOONLY 32M 128M BG

BUILD ON * USING SUBCONFIG DTCSMA-1

* BUILD ON @@member2name USING SUBCONFIG DTCSMA-2

SUBCONFIG DTCSMA-1

LINK TCPMAINT 591 591 RR

LINK TCPMAINT 592 592 RR

LINK TCPMAINT 198 198 RR

MDISK 191 3390 3333 005 M01RES MR READ WRITE MULTIPLE

*SUBCONFIG DTCSMA-2

* LINK TCPMAINT 591 591 RR

* LINK TCPMAINT 592 592 RR

* LINK TCPMAINT 198 198 RR

VMXGNR Example – part 2

IDENTITY DTCSMAPI AUTOONLY 32M 128M BG

BUILD ON * USING SUBCONFIG DTCSMA-1

* BUILD ON @@member2name USING SUBCONFIG DTCSMA-2

SUBCONFIG DTCSMA-1

LINK TCPMAINT 591 591 RR

LINK TCPMAINT 592 592 RR

LINK TCPMAINT 198 198 RR

MDISK 191 3390 3333 005 M01RES MR READ WRITE MULTIPLE

IDENTITY LOHCOST AUTOONLY 768M 2G G

*PW= 12/04/05

*ED=

*SUBCONFIG DTCSMA-2

* LINK TCPMAINT 591 591 RR

* LINK TCPMAINT 592 592 RR

* LINK TCPMAINT 198 198 RR

VM:Tape

- Scratch tapes mounted on one system don't show up in VMTAPE LIST on other systems until VMTAPE REFRESH is run. Working as designed.

VM:Backup

- Run one server per cluster to back up shared DASD, and one per member to back up private DASD. Put all of the servers on shared DASD.

Shared SPOOL

- SPOOL files created on one member are only visible on other members when the owning userid is logged on.

Reserve/Release Changes

- New rules to preserve integrity. Details in CP Planning and Admin.

Finally...

- Goodbye CSE. Clusters automatically provide link protection for all shared DASD. No need to configure link protection.