



IN THE BUSINESS OF YOUR SUCCESSSM

ADP – zFCP and Oracle on System z

Joe Comitini
Principal Architect System Z

HR. Payroll. Benefits.

Agenda

- **Why Oracle & FCP on Z**
- **Benefits**
- **Team Effort**
- **Standards**
- **Engineering**
- **How it all comes together**

Why Oracle & FCP on Z

- It does work ...
- Make sure to patch (Linux & Oracle)
- ASM is your friend
- FCP over DASD for ASM disks
- To NPIV or not to NPIV ...

FCP vs DASD

Communication is key ...

Storage Admin

DBA

System Admin

- FCP SAN LUN layout
- Multipath Binding Alias Naming convention

Standards – storage allocation spreadsheet

| Hostname (Guest) | File System | VMAX Device ID | zLinux LUN ID | Size | Thin Pool |
|---------------------------|---------------------------|----------------|---------------|--------------|--------------------------------|
| All ClusterNodes | ASM Voting (0V01) | 1acd | 00a | 2 Gb | Thin Pool = RAID1_600GB |
| All ClusterNodes | ASM Voting (0V01) | 1ace | 00b | 2 Gb | Thin Pool = RAID1_600GB |
| All ClusterNodes | ASM Voting (0V01) | 1acf | 00c | 2 Gb | Thin Pool = RAID1_600GB |
| All ClusterNodes | ASM Voting (0V01) | 1ad0 | 00d | 2 Gb | Thin Pool = RAID1_600GB |
| All ClusterNodes | ASM Voting (0V01) | 1ad1 | 00e | 2 Gb | Thin Pool = RAID1_600GB |
| | | | | | |
| All ClusterNodes | ASM (DG01) DATA | 234a | 101 | 128 Gb | Thin Pool = RAID1_600GB |
| All ClusterNodes | ASM (DG01) DATA | 236a | 102 | 128 Gb | Thin Pool = RAID1_600GB |
| All ClusterNodes | ASM (DG01) DATA | 238a | 103 | 128 Gb | Thin Pool = RAID1_600GB |
| All ClusterNodes | ASM (DG01) DATA | 23aa | 104 | 128 Gb | Thin Pool = RAID1_600GB |
| All ClusterNodes | ASM (RA01) DATA | 242a | 108 | 8 Gb | Thin Pool = RAID1_600GB |
| All ClusterNodes | ASM (RA01) DATA | 244a | 109 | 8 Gb | Thin Pool = RAID1_600GB |
| All ClusterNodes | ASM (RA01) DATA | 246a | 10A | 8 Gb | Thin Pool = RAID1_600GB |
| | | | | | |
| Cluster Node 01 | /oracle (Oracle) | 248a | 011 | 32 Gb | Thin Pool = RAID1_600GB |
| Cluster Node 01 | /oraadmin (Oracle) | 248c | 012 | 32 Gb | Thin Pool = RAID1_600GB |
| Cluster Node 01 | /oracrs (Oracle) | 248e | 013 | 32 Gb | Thin Pool = RAID1_600GB |
| | | | | | |
| Cluster Node 02 | /oracle (Oracle) | 2490 | 021 | 32 Gb | Thin Pool = RAID1_600GB |
| Cluster Node 02 | /oraadmin (Oracle) | 2492 | 022 | 32 Gb | Thin Pool = RAID1_600GB |
| Cluster Node 02 | /oracrs (Oracle) | 2494 | 023 | 32 Gb | Thin Pool = RAID1_600GB |
| | | | | | |
| Cluster Spare Node | /oracle (Oracle) | 24a8 | 0b1 | 32 Gb | Thin Pool = RAID1_600GB |
| Cluster Spare Node | /oraadmin (Oracle) | 24aa | 0b2 | 32 Gb | Thin Pool = RAID1_600GB |
| Cluster Spare Node | /oracrs (Oracle) | 24ac | 0b3 | 32 Gb | Thin Pool = RAID1_600GB |

Standards – CMS FCP file

| <u>SAN WWPN</u> | <u>LUN Id</u> | <u>Guest Name</u> | <u>Role</u> |
|---------------------------------------|---------------|--------------------|-------------|
| 0x5000097300140530:0x0011000000000000 | | E521601A:/oracle | |
| 0x5000097300140530:0x0012000000000000 | | E521601A:/oraadmin | |
| 0x5000097300140530:0x0013000000000000 | | E521601A:/oracrs | |
| 0x5000097300140530:0x000a000000000000 | | E521601A:voting1 | |
| 0x5000097300140530:0x000b000000000000 | | E521601A:voting2 | |
| 0x5000097300140530:0x000c000000000000 | | E521601A:voting3 | |
| 0x5000097300140530:0x000d000000000000 | | E521601A:voting4 | |
| 0x5000097300140530:0x000e000000000000 | | E521601A:voting5 | |
| 0x5000097300140530:0x0101000000000000 | | E521601A:diskdg001 | |
| 0x5000097300140530:0x0102000000000000 | | E521601A:diskdg002 | |
| 0x5000097300140530:0x0103000000000000 | | E521601A:diskdg003 | |
| 0x5000097300140530:0x0104000000000000 | | E521601A:diskdg004 | |
| 0x5000097300140530:0x0108000000000000 | | E521601A:diskra001 | |
| 0x5000097300140530:0x0109000000000000 | | E521601A:diskra002 | |
| 0x5000097300140530:0x010a000000000000 | | E521601A:diskra003 | |

Engineering – Script functions for FCP LUNs

```
setFCP ora                # configure FCP/SCSI disks
mkBindingsFile ora       # make the multipath bindings file
mkFileSystems ora        # make partitions for the FCP cooked file systems
append2fstab ora         # add entries to the /etc/fstab file
(possibly some more ...)
```

Engineering – Script function mkBindingsFile

```
lines=`cmsfscat -a -d $dev191 $userID.FCP1000`
for nextLine in $lines; do
  tgtID=`echo $nextLine | awk -F: '{print $3}'`
  if [ "$userID" != "$tgtID" ]; then # this record is not for this ID => skip
    continue
  fi
  lun=`echo $nextLine | awk -F: '{print $2}'`
  let lunInDecimal=${lun:0:6} # convert LUN from hex to decimal
  fileSys=`echo $nextLine | awk -F: '{print $4}'`
  if [ ${fileSys:0:1} = '/' ]; then # "cooked" file system
    if [ ${SLESversion} = 10 ]; then # for SLES 10 mpath<x> is friendly name
      letter=${alphabet:$mpathIndex:1}
      friendlyName="mpath$letter"
      let mpathIndex=$mpathIndex+1
    else # for SLES 11 file system name is friendly name
      friendlyName=`echo $fileSys | sed -e 's:/::g'` # strip leading /
    fi
  fi
done
```

Engineering – Script function mkBindingsFile

```
# create a list of cooked file systems
  if [ $#cookedFileSystems = 0 ]; then # list is empty
    cookedFileSystems="$fileSys"
  else # add new file system to list
    cookedFileSystems="$cookedFileSystems $fileSys"
  fi
else # this must be an ASM disk
  if [ ${SLESversion} = 10 ]; then # for SLES 10 mpath<x> is friendly name
    letter=${alphabet:$mpathIndex:1}
    friendlyName="mpath$letter"
    let mpathIndex=$mpathIndex+1
  else # for SLES 11 friendly name is in FCPx000 file
    friendlyName=$fileSys
    if [ "$friendlyName" = "ASM" ]; then # ASM is a SLES 10 keyword => fail
      exitNow 26 "Friendly name of ASM in $userID.FCP1000 not correct for
SLES 11"
    fi
  fi
fi
```

Engineering – Linux bindings file

```
# Multipath bindings, Version : 1.0
# NOTE: this file is automatically maintained by the multipath program.
# You should not need to edit this file in normal circumstances.
#
# Format:
# alias (friendly name) wwid (ugly name)
#
oracle 360000970000195701281533032344138
oraadmin 360000970000195701281533032344141
oracrs 360000970000195701281533032344143
voting1 360000970000195701281533031414344
voting2 360000970000195701281533031414345
voting3 360000970000195701281533031414346
voting4 360000970000195701281533031414430
voting5 360000970000195701281533031414431
diskdg001 360000970000195701281533032333441
diskdg002 360000970000195701281533032333641
diskdg003 360000970000195701281533032333841
diskdg004 360000970000195701281533032334141
diskra001 360000970000195701281533032343241
diskra002 360000970000195701281533032343441
diskra003 360000970000195701281533032343641
```

References

boot.oracle script

Experiences with Oracle 11gR2 on Linux on System z

<http://www.redbooks.ibm.com/abstracts/sg248104.html?Open>

Fibre Channel Protocol for Linux and z/VM on IBM System z

<http://www.redbooks.ibm.com/abstracts/sg247266.html>

Introduction to SCSI over FCP for Linux on System z

<http://www.vm.ibm.com/education/lvc/lvc1020c.pdf>

Linux on System z - Disk I/O Performance – Part 1

<http://www.vm.ibm.com/education/lvc/LVC0918.pdf>

