#SHAREorg

SHARE
Technology · Connections · Results

Using z/VM DirMaint in an SSI Cluster

Pam Bryant
IBM
Advanced Technical Skills

Monday, March 12, 2012
Session # 10567

Trademarks

U.S. Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

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

- Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.
- For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml

Agenda

- How DirMaint works in an SSI cluster
- DirMaint Command Changes
- Installation and Configuration Tips

What is DirMaint?

- DirMaint, the Directory Maintenance Facility of z/VM, provides:
  - Automated control of the z/VM user directory through a command interface
  - Auditing and controls for access to the directory
  - Error checking
  - Automated facilities for minidisk allocation, deallocation, and copying.
  - The ability to work in conjunction with external security managers, such as RACF.

How DirMaint Works in a Cluster

- In an SSI Cluster, the source user directory file is shared by all members of the cluster, but each member has its own object directory.
- The DIRMAINT server, running on a single member of the cluster, controls the source directory files for the cluster.
- Satellite DirMaint servers, running on the other members of the cluster, provide an interface to users on their local system to the DIRMAINT server.
- DIRMAINT and Satellite servers control the object directory for their local member systems.

How DirMaint Works in a Cluster

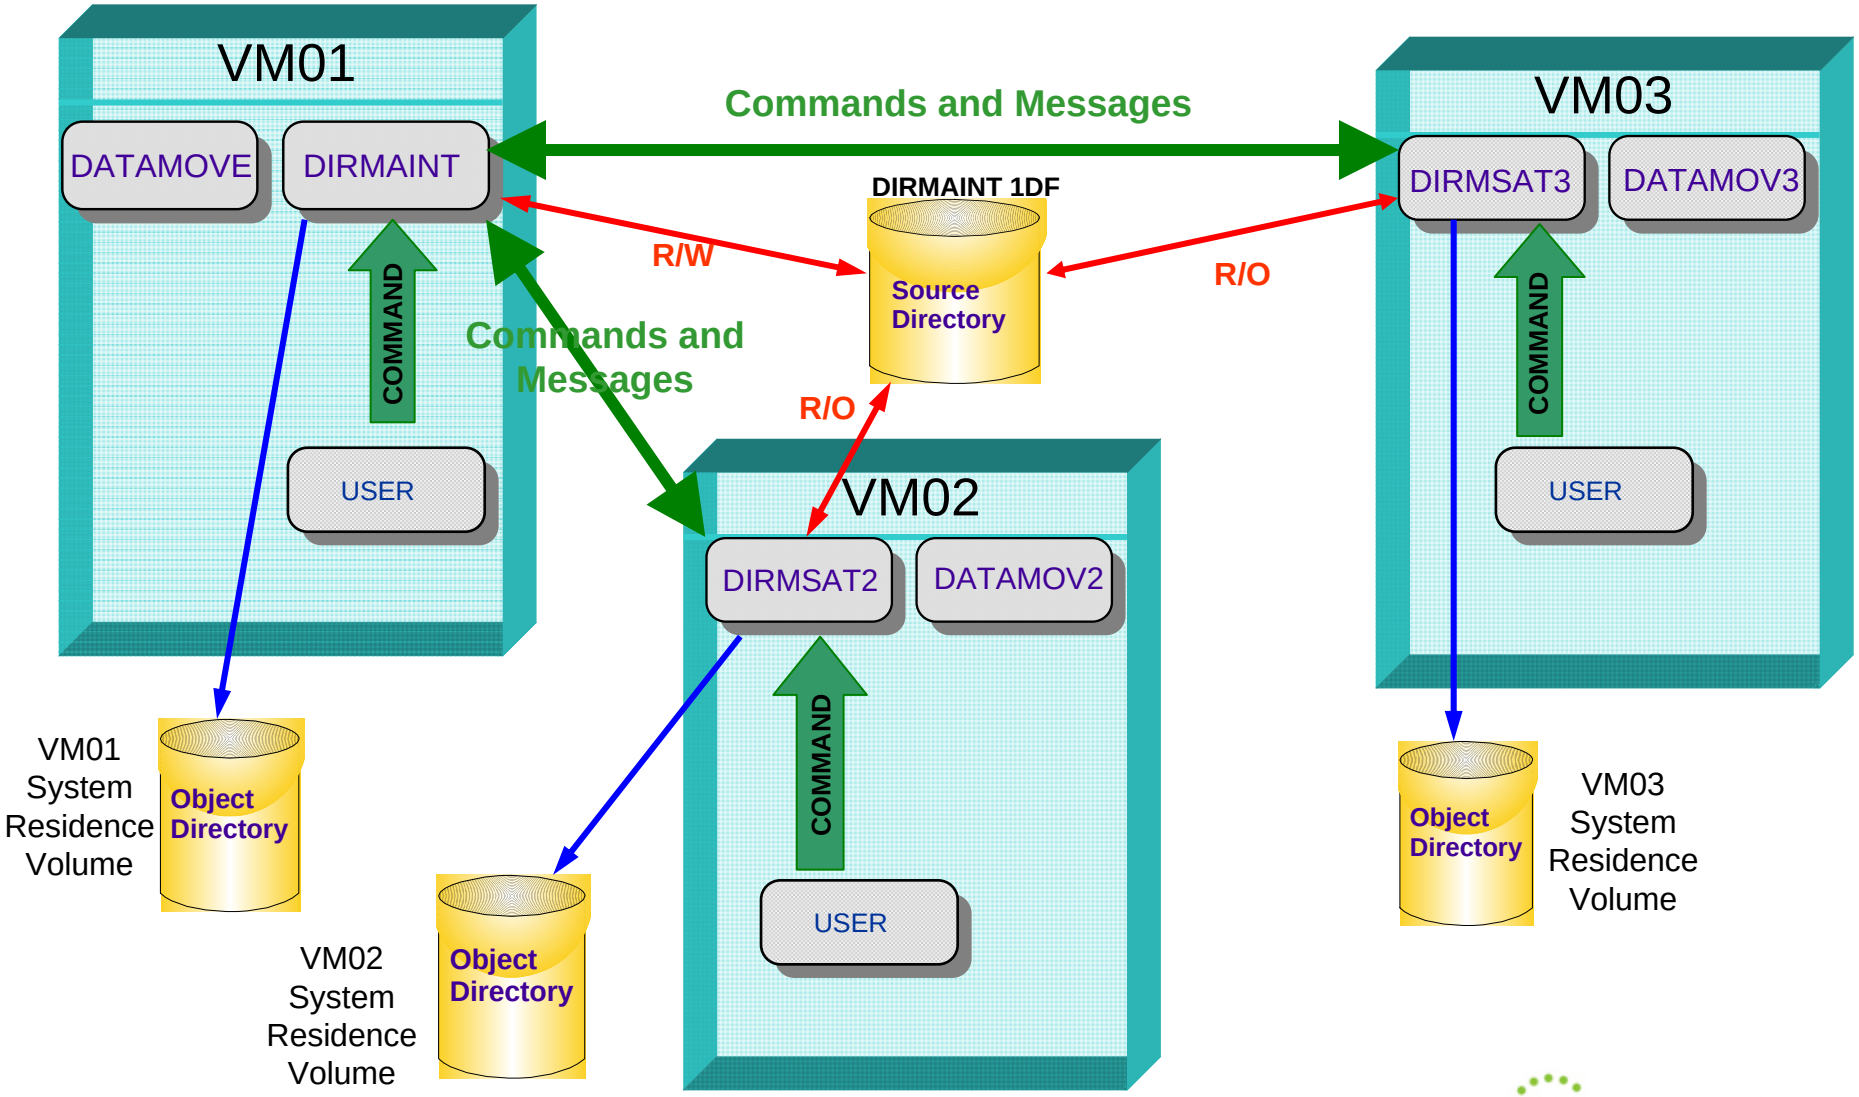

DirMaint Server Communications

- In an SSI cluster, DIRMAINT and the satellite servers communicate through the shared spool.

- At startup, the DIRMAINT server creates a control file SATRELAY DATADVH to identify the satellite servers for each member.

```
SSI
DIRMSAT  VM01
DIRMSAT2 VM02
```

- SATRELAY DATADVH is used by :
  - the DIRMAINT command to determine which server to route commands through
  - DIRMAINT to determine which satellite to send output to.

Where servers can run in the cluster

- The DIRMAINT server can run on any system in the cluster.
  - DIRMAINT **must** run on one member to process commands
  - DIRMAINT is a single-configuration virtual machine (USER)
  - All minidisks are located on common disks
- The Satellite servers (DIRMSATn) and Datamove servers (DATAMOVn) should only run on one member system.
  - DIRMSAT servers are single-configuration virtual machines (USER), but...
  - By default, minidisks are defined on non-shared volumes
  - A Satellite and Datamove server should be defined on every member, even the member where DIRMAINT will run.

Where Servers Run

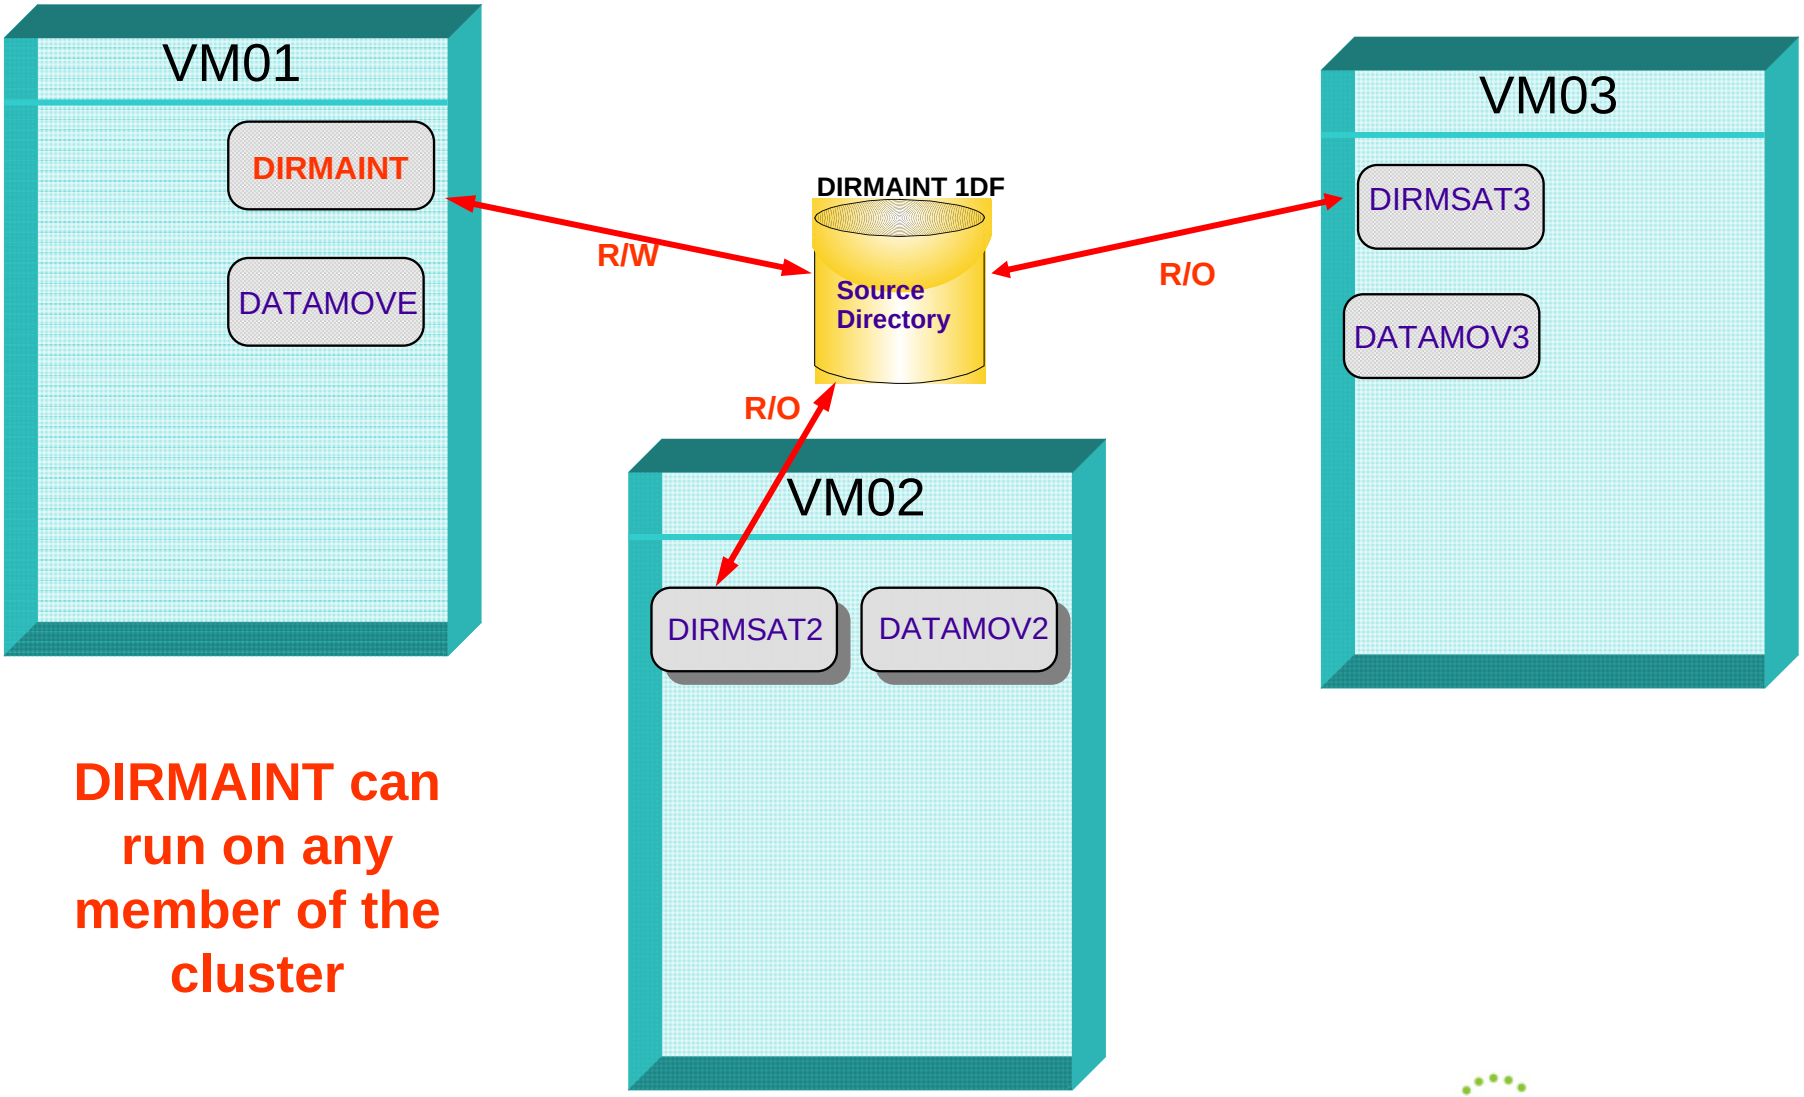

DIRMAINT can run on any member of the cluster

Where Servers Run (continued)

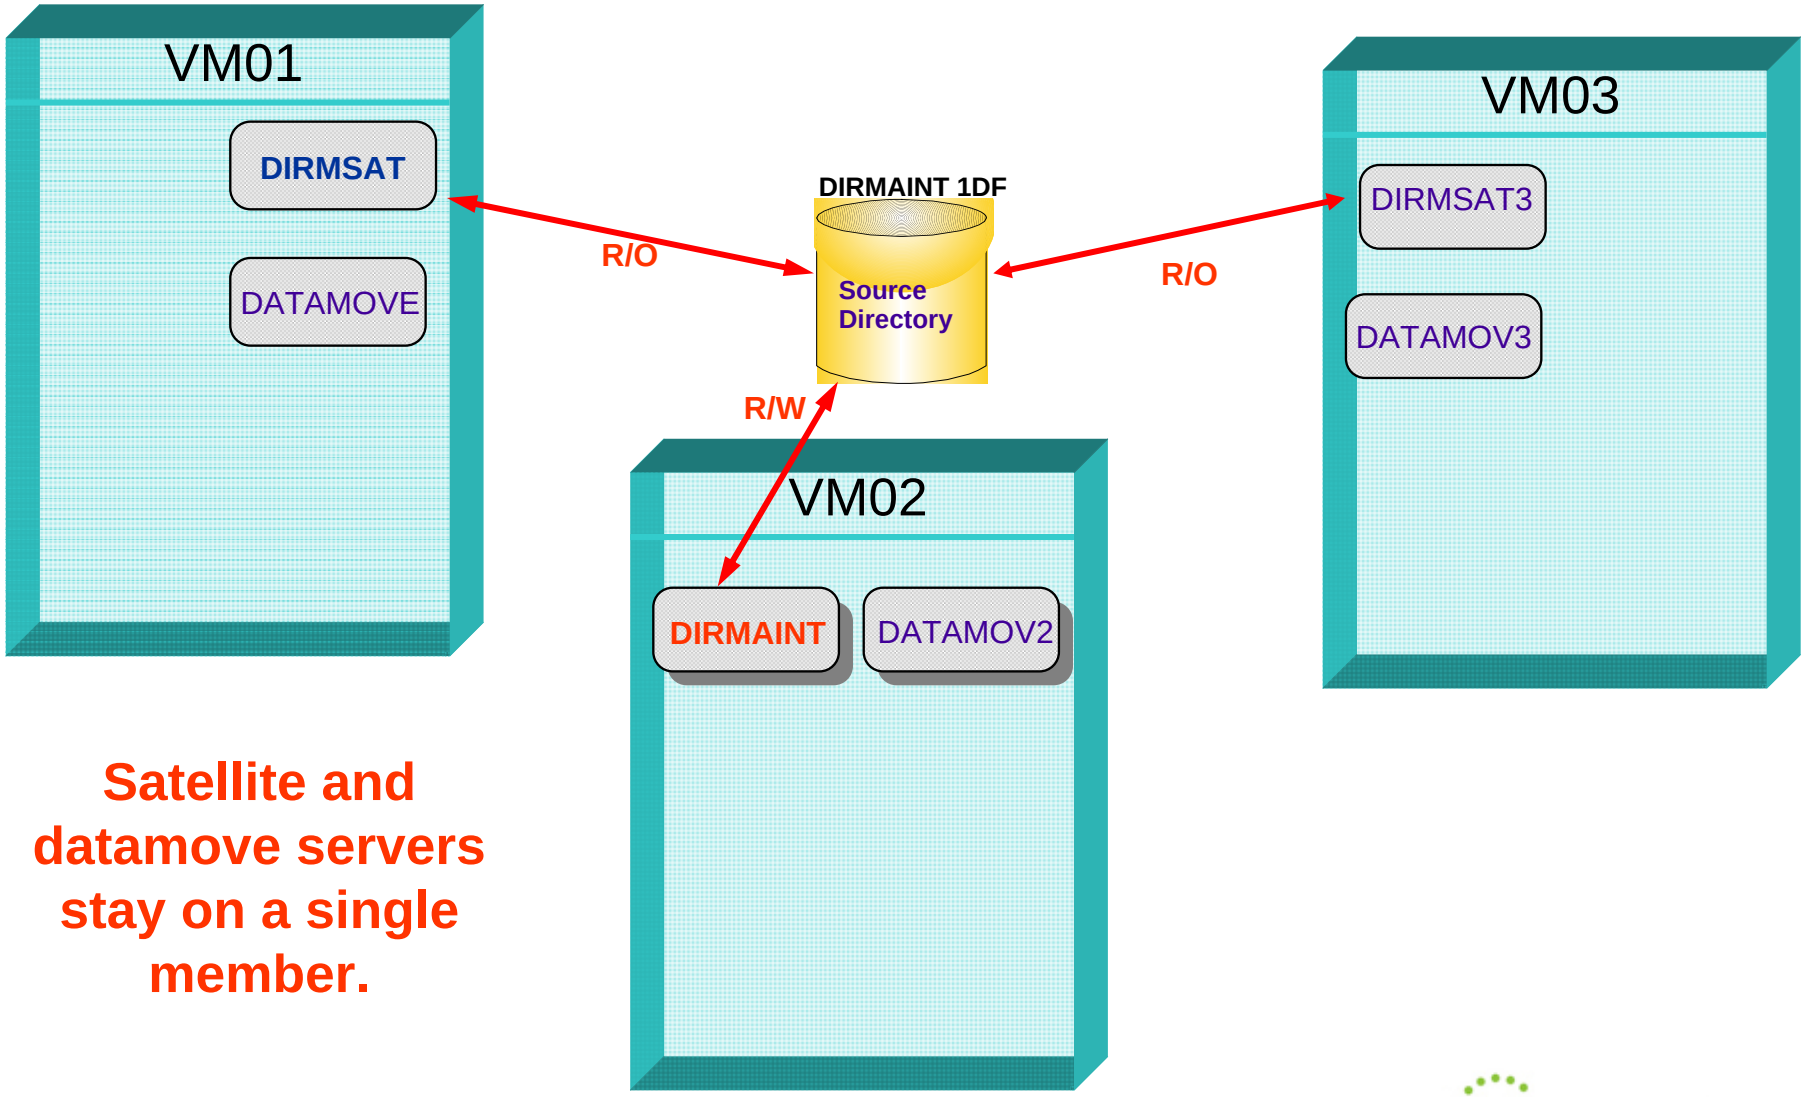

Satellite and datamove servers stay on a single member.

Where Servers Run (continued)

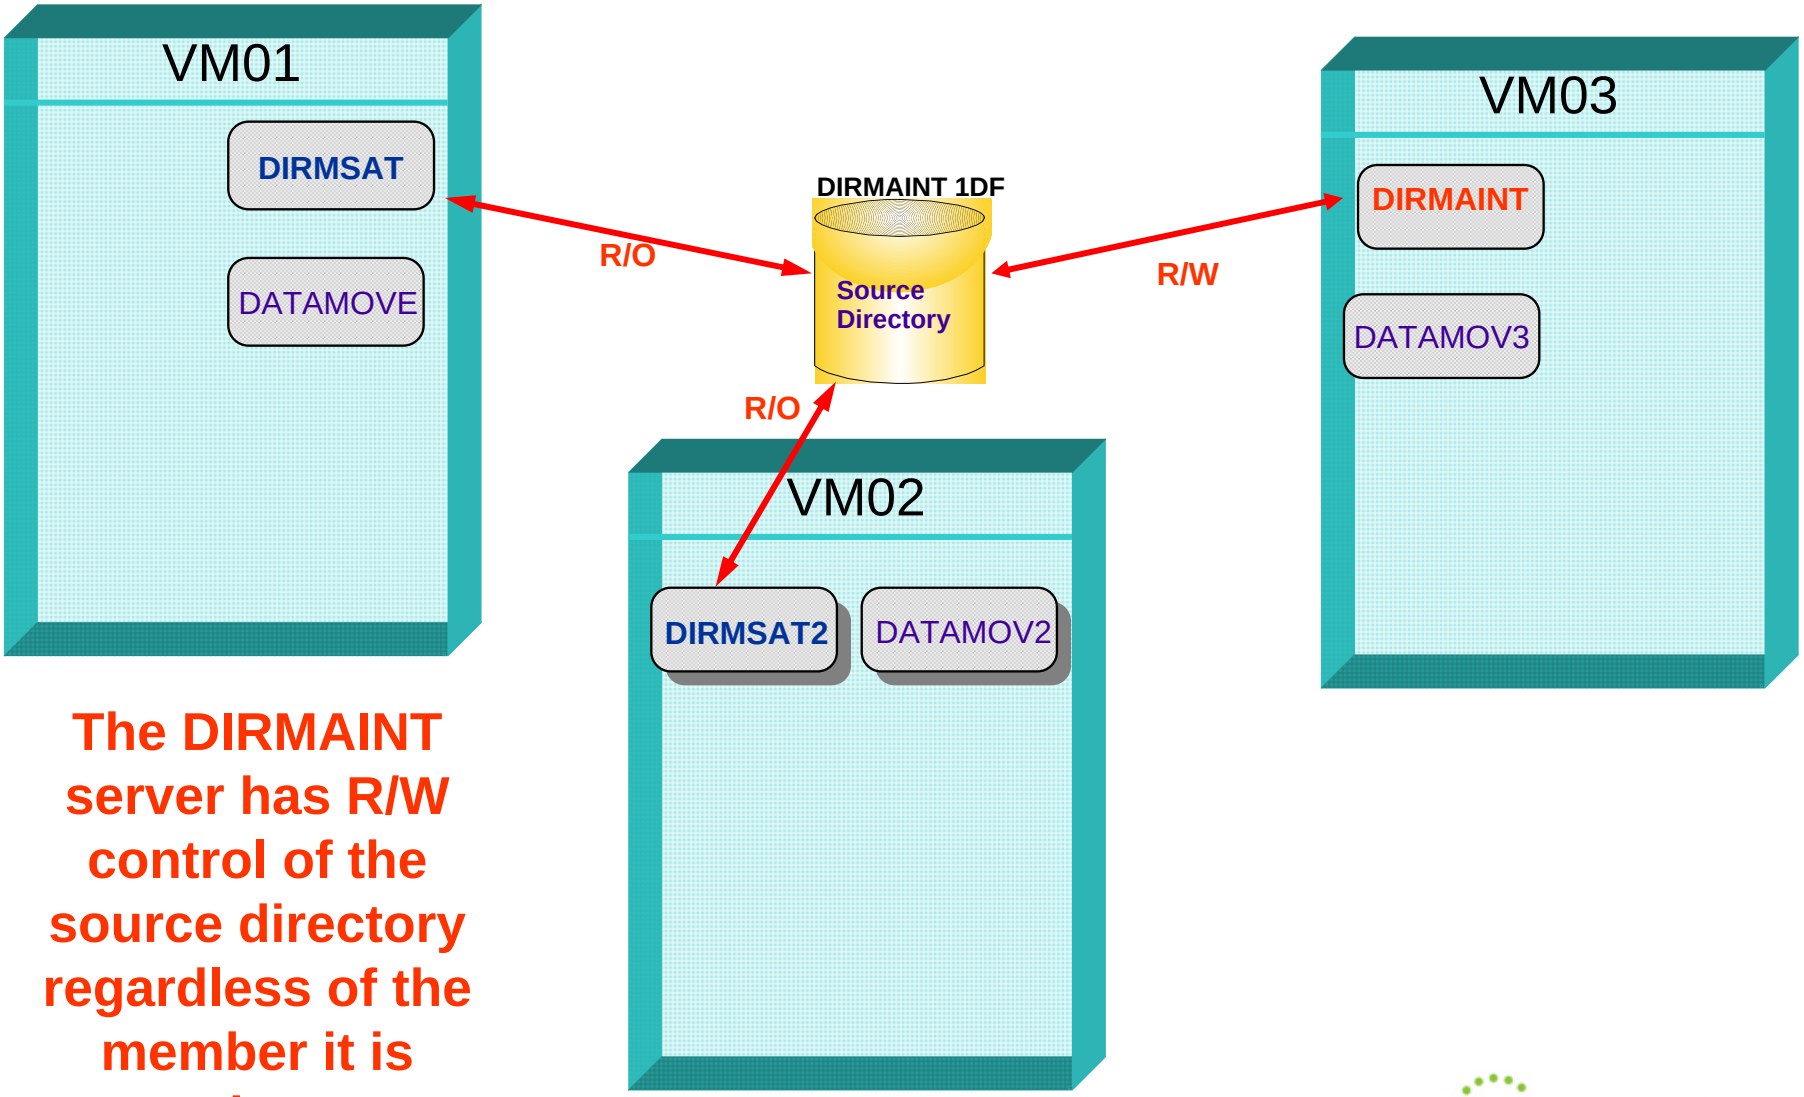

The DIRMAINT server has R/W control of the source directory regardless of the member it is running on.

If DIRMAINT Stops

- If DIRMAINT is stopped for any reason on one member, it can be started on another member.
  - No DIRMAINT commands can be processed until DIRMAINT is started somewhere in the cluster.
- When the member is ready to start using DirMaint again:
  - It must run a satellite server or
  - DIRMAINT on the secondary system is shut down and the server is re-started on the original system. The secondary system then starts its Satellite server.
- DirMaint does not automate the process of starting or stopping servers when the DIRMAINT server stops or re-starts.

If a Satellite server goes down

- If a Satellite server running on a cluster member goes down while the member is still joined:
  - Users on that system cannot issue DIRMAINT commands
  - Changes made to the directory by DIRMAINT will not be reflected in the Satellite system's object directory.
- When the Satellite server is restarted, it will process updates made while it was out of service.

DirMaint Commands

- The DIRMAINT command has been updated to handle the new IDENTITY and SUBCONFIG entries in the directory.
  - IDENTITYs and SUBCONFIGs are treated as separate entities by DirMaint.
  - IDENTITY and SUBCONFIG entries use the same commands as PROFILE and USER entries – ADD, PURGE, GET, REVIEW, LOCK, and UNLOCK.
- Other new and updated directory statements are also supported.
  - For example, VMRELOCATE and CHPID virtualization options of OPTION and GLOBALOPTS statements.

DirMaint Commands – Adding a Multi-Configuration User

- A multi-configuration virtual machine consists of a single IDENTITY entry, with one or more SUBCONFIG entries.
  - Each IDENTITY and SUBCONFIG entry are created with separate ADD commands.

- Example: To create a new multi-configuration user on members VM01 and VM02, you would need 3 DirMaint commands and 3 DIRECT files.

DirMaint Commands – Adding a Multi-Configuration User (*continued)*

DIRM ADD TUSER

TUSER DIRECT A

```
IDENTITY TUSER APASSWD 128M 1000M ABCG
MACHINE ESA
IPL 190
CONSOLE 009 3215
SPOOL 00C 2540 READER *
SPOOL 00D 2540 PUNCH A
SPOOL 00E 1403 A
LINK MAINT 0190 0190 RR
LINK MAINT 019D 019D RR
LINK MAINT 019E 019E RR
```

DirMaint Commands – Adding a Multi-Configuration User (*continued)*

- DIRM ADD TUSER-1 BUILD ON VM01 IN TUSER

TUSER-1 DIRECT A

SUBCONFIG TUSER-1
AMDISK 191 3390 AUTOV 005 M01W01

DirMaint Commands – Adding a Multi-Configuration User (*continued)*

- DIRM ADD TUSER-2 BUILD ON VM02 IN TUSER

TUSER-2 DIRECT A

SUBCONFIG TUSER-2
AMDISK 191 3390 AUTOV 005 M02W01

DirMaint Commands – Adding a Multi-Configuration User (*continued*)

- Results in a directory entry of:

```
IDENTITY TUSER APASSWD 128M 1000M ABCG
   BUILD ON VM01 USING SUBCONFIG TUSER-1
   BUILD ON VM02 USING SUBCONFIG TUSER-2
   IPL 190
   MACHINE ESA
   CONSOLE 0009 3215
   SPOOL 000C 2540 READER *
   SPOOL 000D 2540 PUNCH A
   SPOOL 000E 1403 A
   LINK MAINT 0190 0190 RR
   LINK MAINT 019D 019D RR
   LINK MAINT 019E 019E RR
SUBCONFIG TUSER-1
   MDISK 0191 3390 2733 5 M01W01
SUBCONFIG TUSER-2
   MDISK 0191 3390 2728 10 M02W01
```

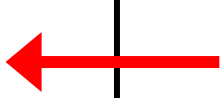

Added by DIRMAINT

DirMaint Commands – Adding a Multi-Configuration User (*continued)*

- IDENTITY must be added before SUBCONFIGs.
- Using prototype directories works the same way.
  - ADD TUSER LIKE MULTISRV
  - ADD TUSER-1 LIKE MULTIS-1 BUILD ON VM01 IN TUSER
  - ADD TUSER-2 LIKE MULTIS-2 BUILD ON VM02 IN TUSER
  - Where MULTISRV PROTODIR, MULTIS-1 PROTODIR, and MULTIS-2 PROTODIR reside on DIRMAINT's A-disk.
- Creating a single-configuration virtual machine (USER) has not changed.

DirMaint Commands – GET and REVIEW, Authorization

- DIRM GET
  - For an IDENTITY, DirMaint responds with the IDENTITY section of the directory entry, including BUILD statements.
  - If you specify AT member, DirMaint responds with the corresponding SUBCONGFIG and not the IDENTITY section.

- DIRM REVIEW
  - For an IDENTITY, DirMaint responds with the IDENTITY and associated SUBCONFIG sections of the directory entry.

- SUBCONFIGs can be the target of a GET or REVIEW. Only the SUBCONFIG section of the directory entry will be retrieved.

- DirMaint commands are authorized by USER id or IDENTITY id, not SUBCONFIG.
  - The IDENTITY is automatically given authority over its associated SUBCONFIGs.
  - Command authorization is defined in AUTHFOR CONTROL.

DirMaint Commands – REPLACE and PURGE

- DIRM REPLACE
  - Cannot change entry type using DIRM REPLACE
  - Cannot have multiple entry types in one entry
  - Cannot remove BUILD statement

- DIRM PURGE
  - When deleting a SUBCONFIG entry, DirMaint will remove associated BUILD statement from IDENTITY entry
  - When deleting an IDENTITY entry, DirMaint will remove all related SUBCONFIG entries

DirMaint Commands – AMDISK and DMDISK

- DIRM AMDISK and DMDISK
  - For single-configuration virtual machines, the prefix keywords refer to the USERid
  - For multi-configuration virtual machines, the prefix keywords refer to the SUBCONFIG id.
    - You can specify the IDENTITY and not the SUBCONFIG id on the AMDISK command, but the result will probably not be what you intended.
  - For operations that require a Datamove machine, the Datamove machine will be selected based on the system node associated with SUBCONFIG on BUILD statement.

    - **DIRM FOR TUSER-1 DMDISK 191 CLEAN**
      - DATAMOVE is assigned the CLEAN task

    - **DIRM FOR TUSER-2 DMDISK 191 CLEAN**
      - DATAMOV2 is assigned the CLEAN task

DirMaint Commands - SSI

```
>>--DIRMaint--.------------------.--SSI--nodeid-----><
              '-Prefix keywords-'
```

- Changes an SSI-Ready source directory to SSI-Enabled:
  - Updates the DIRECTORY statement with the SSI keyword
  - Changes all BUILD ON * statements to BUILD ON *nodeid*

DirMaint Commands - UNDOSSI

```
>>--DIRMaint--.------------------.--UNDOSSI--nodeid----------><
              '-Prefix keywords-'
```

- Rolls back changes made by the DIRM SSI command
  - Changes all BUILD ON n*odeid* statements to BUILD ON *
  - Removes SSI option from DIRECTORY statement
  - Directory must have only one BUILD statement per IDENTITY and the system node on the BUILD statement must equal *nodeid*

Using DIRMAP

The subconfig and the member associated with the minidisk has been added.

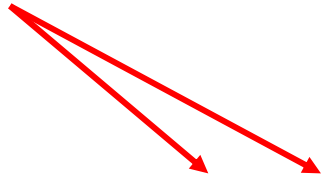

```
USER      DIRECT      Map of Minidisks    16:25:12   20120215
Volser    Type Ownerid   Addr Start       End         Length     Flags Subconfig Member
M01RES    3390 $ALLOC$   0A04           0           0           1 Overlap
               MAINT     0123           0        1112        1113       MAINT-1   VM01
               SYSDUMP1  0123           0        1112        1113       SYSD
               .DRCT.    0300           1          20          20 Overlap
               $DIRECT$  0A01           1          20          20
               $SYSCKP$  0A01          21          29           9
               $SYSWRM$  0A01          30          38           9
               MAINT     0CF1          39         158         120       MAINT-1   VM01
               MAINT     0CFD         159         159           1       MAINT-1   VM01
               MAINT     0CF3         160         279         120       MAINT-1   VM01
               MAINT     0190         280         493         214       MAINT-1   VM01
               MAINT     0191         494         668         175       MAINT-1   VM01
               MAINT     0193         669        1168         500       MAINT-1   VM01
```

Configuring DirMaint

- Installation Tips
- Configuration Files
  - CONFIGSS DATADVH
  - EXTENT CONTROL

- Automating startup with AUTOLOG1

Installation Tips

- An SSI installation will create the service machines, the CONFIGSS DATADVH and EXTENT CONTROL configuration file statements necessary to run DirMaint in the cluster.
- The SERVICE DIRM ENABLE command (to enable the product to VMSES/E and to CP in SYSTEM CONFIG), only has to be run on one member.
- PUT2PROD needs to be run on every member.
- Configuration files are shared. They can be created once from any member of the cluster.
- Change DIRMAINT's default password from AUTOONLY to some other password prior to installation. You can change it back after you've successfully tested DirMaint.

CONFIGSS DATADVH Configuration File

CONFIGSS DATADVH file contains DirMaint configuration statements which override default statements contained in the primary configuration file CONFIG DATADVH. These statements will define the Satellite and Datamove servers for the cluster.

- This override file is created for you during an SSI installation.

```
SATELLITE_SERVER= DIRMSAT VM01
SATELLITE_SERVER= DIRMSAT2 VM02
SATELLITE_SERVER= DIRMSAT3 VM03
DATAMOVE_MACHINE= DATAMOVE VM01   *
DATAMOVE_MACHINE= DATAMOV2 VM02   *
DATAMOVE_MACHINE= DATAMOV3 VM03   *
```

EXTENT CONTROL

- EXTENT CONTROL identifies volumes that will be used for automated minidisk allocation and provides a template of how the space should be used. The four sections of interest are:
  - Regions section – where contiguous spaces on a volume are identified and assigned a name (RegionId).
  - Groups section – where regions are grouped under a single name (GroupName) so that DirMaint can look for available space in any of the regions in the group.
  - Exclude section – where minidisks are identified that should be excluded from consideration when looking for unused space on a volume for a new minidisk.
  - SSI_Volumes section – used when cloning subconfig entries from existing ones.

EXTENT CONTROL – REGIONS Section

- A Region is the basic unit of DASD segmentation used by DirMaint. It defines a single, continguous area on a single DASD volume.

- Every volume, shared or non-shared, that will be used for minidisk allocation by DirMaint must have a region or regions defined for it.

EXTENT CONTROL – REGIONS Section (*continued*)

- Regions are defined as follows:

Regionid	Volser	RegStart	RegEnd	Dev-Type
RegionA	LNXV04	500	2000	3390-03
RegionC	USRPK1	100	END	3390-03
RegionD	USRPK2	25	1000	3390-03
RegionE	USRPK2	1500	END	3390-03

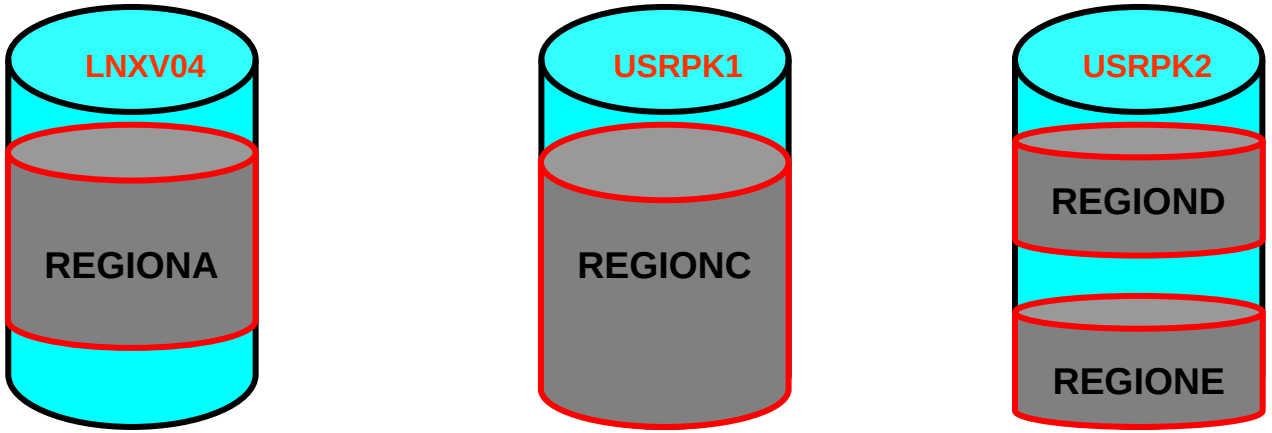

EXTENT CONTROL – GROUPS Section

A Group is a collection of one or more regions. Groups are defined as follows:

groupid	**region1**	**region2**	***.....***	**regionn**
GrpSyS	SystemA	SystemB		
GrpLnx	Lnxpk1	Lnxpk2		Lnxpk3

Regions specified in a group entry must be defined in the Regions section.

Do not mix regions located on shared and non-shared volumes within a group.

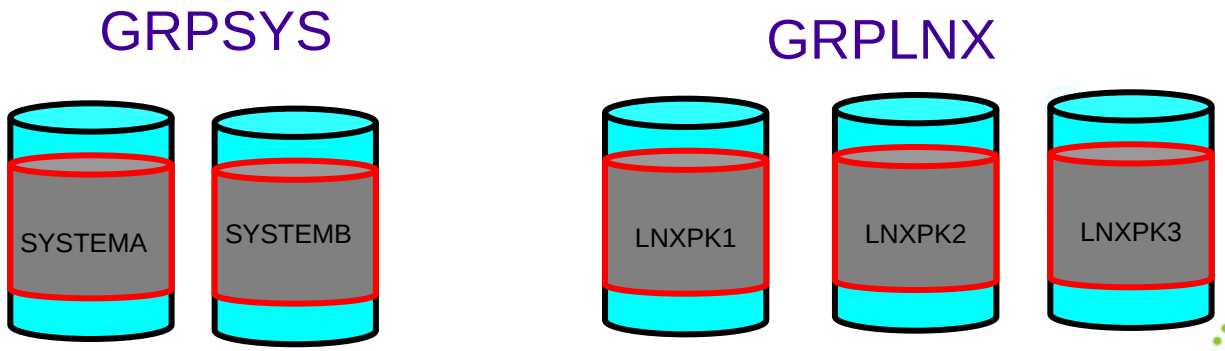

EXTENT CONTROL – EXCLUDE Section

- Overlapping minidisks are commonly defined in the User Directory
  - **MAINT** has full-pack minidisks covering many system volumes for writing to the directory area of the system residence pack, volume backup and restore, etc.
  - **MAINT*vrm*** has full-pack minidisks covering the release volumes.
  - **PMAINT** has full-pack minidisks covering the common volumes.
  - Other virtual machines, such as those that perform system backups, may also have overlapping extents.

- The EXCLUDE section of EXTENT CONTROL allows you to tell DIRMAINT to ignore these overlaps when looking for free extents during minidisk allocation.

EXTENT CONTROL – EXCLUDE Section (continued)

- Tip: You must specify subconfig ids for identity users in the EXCLUDE SECTION. Wildcards can be used to identify multiple minidisks in the same statement.

```
:EXCLUDE.
* entry_name Address
MAINT*      012*
MAINT620    013*
SYSDMP*     012*
PMAINT      014*
:END.
```

MAINT SUBCONFIGs MAINT-1 and MAINT-2 each have overlapping minidisks at addresses 0122, 0123, and 0124.

SYSDUMP1 has SUBCONFIGs SYSDMP-1 and SYSDMP-2 with fullpack 0123 minidisks overlapping the system residence pack.

SSI_Volume Section

- The SSI_Volume section is used for cloning an SSI member to a new member and DirMaint is used to create the new subconfig entries.

```
:SSI_VOLUMES.
*VolumeFamily      Member    VolSer
IBM_RES            VM01      M01RES
IBM_WORK1          VM01      M01W01
IBM_RES            VM02      M02RES
IBM_WORK1          VM02      M02W01
:END.
```

Inserted during a two member installation by the installation tool.

SSI_Volume Section example

- Example – a two-member SSI cluster was installed. Now, a third member is being added. Identity users like TCPIP may need a subconfig for the new member added to their directory entries based on the existing subconfigs for other members.

SSI_Volume Section example (continued)

```
IDENTITY TCPIP    TCPIP     128M  256M ABG
 INCLUDE TCPCMSU
 BUILD ON VM01 USING SUBCONFIG TCPIP-1
 BUILD ON VM02 USING SUBCONFIG TCPIP-2
 OPTION QUICKDSP SVMSTAT MAXCONN 1024 DIAG98 APPLMON
 SHARE RELATIVE 3000
 IUCV ALLOW
 IUCV ANY PRIORITY
 IUCV *CCS PRIORITY MSGLIMIT 255
 IUCV *VSWITCH MSGLIMIT 65535

SUBCONFIG TCPIP-1
 LINK TCPMAINT 491 491 RR
 LINK TCPMAINT 492 492 RR
 LINK TCPMAINT 591 591 RR
 LINK TCPMAINT 592 592 RR
 LINK TCPMAINT 198 198 RR
 MDISK 191 3390 2627 005 M01W01 MR RTCPIP   WTCPIP   MTCPIP

SUBCONFIG TCPIP-2
 LINK TCPMAINT 491 491 RR
 LINK TCPMAINT 492 492 RR
 LINK TCPMAINT 591 591 RR
 LINK TCPMAINT 592 592 RR
 LINK TCPMAINT 198 198 RR
 MDISK 191 3390 2627 005 M02W01 MR RTCPIP   WTCPIP   MTCPIP
```

Existing directory entry for TCPIP for 2-member SSI cluster.

TCPIP on member VM03 needs a non-shared 191 minidisk.

SSI_Volume Section example (*continued*)

The SSI_Volume section of EXTENT CONTROL has been updated to include these statements:

```
:SSI_VOLUMES.
*VolumeFamily       Member    VolSer
SYSRES              VM01      M01RES
SYSRES              VM02      M02RES
SYSRES              VM03      M03RES
SYS_LOCAL           VM01      M01W01
SYS_LOCAL           VM02      M02W01
SYS_LOCAL           VM03      M03W01
:END.
```

The system residence volumes for each member and the local W01 volumes for each member have been associated with a volume family.

When a subconfiguration is cloned from an existing subconfiguration, DirMaint refers to the volume family of the original system to determine the VolSer to be used to create minidisks on the target system.

```
DIRM ADD TCPIP-3 LIKE TCPIP-1 BUILD ON VM03 IN TCPIP
```

SSI_Volume Section example (*continued*)

DIRM ADD TCPIP-3 LIKE TCPIP-1 BUILD ON VM03 IN TCPIP

SUBCONFIG TCPIP-1

```
LINK TCPMAINT 491 491 RR
LINK TCPMAINT 492 492 RR
LINK TCPMAINT 591 591 RR
LINK TCPMAINT 592 592 RR
LINK TCPMAINT 198 198 RR
MDISK 191 3390 2627 005 M01W01 MR RTCPIP   WTCPIP   MTCPIP
```

M01W01 on VM01 is associated with volume family SYS_LOCAL.

SUBCONFIG TCPIP-3

```
LINK TCPMAINT 491 491 RR
LINK TCPMAINT 492 492 RR
LINK TCPMAINT 591 591 RR
LINK TCPMAINT 592 592 RR
LINK TCPMAINT 198 198 RR
MDISK 0191 3390 2627 5 M03W01 MR RTCPIP WTCPIP MTCPIP
```

The SYS_LOCAL volume for member VM03 is MO3W01.

Note: This will not copy the contents of the source minidisk to the target minidisk.

Where to find things

- DirMaint code –
  - on minidisks owned by installation user id 6VMDIR20 on a shared volume.
  - 491/492 - Production/Test server code
  - 11F/41F – Production/Test DirMaint interface code
- Directory files
  - on minidisks owned by primary DIRMAINT server on a shared volume.
  - 1DF, 1DB – Primary source directory files, Primary location of USER BACKUP file

Where to find things (continued)

- EXTENT CONTROL
  - on 1DF minidisk owned by DIRMAINT server on a shared volume.
- CONFIG* DATADVH
  - on 11F minidisk owned by DIRMAINT server on a shared volume.
- Object Directory
  - Each member system has a non-shared object directory located on the local non-shared system residence volume, owned by identity user MAINT.
- Common DIRECTXA, DIRMAP, and DISKMAP utilities
  - on PMAINT 551, these are the versions of the utilities with the highest level of maintenance in the cluster.

Automating startup with AUTOLOG1

- The PROFILE EXEC for AUTOLOG1 that comes with z/VM 6.2 contains logic to determine if DirMaint is installed and running in an SSI cluster.
  - If installed, the DIRMAINT service machine will be AUTOLOGed at system startup.
  - If running in an SSI cluster, satellite servers will also be started.
    - Servers defined in CONFIGSS DATADVH
    - The first system in the cluster to come up will get DIRMAINT, everyone else will get satellites.
  - DATAMOVE servers must be manually added to the AUTOLOG process for each system.