In Service to z/VM and Linux on Z

SHARE Atlanta
Session 10894
David Kreuter

Now Showing: In Service to z/VM and Linux on z

Abstract: In this presentation the benefits of building a service zone LPAR for shops with multiple z/VM LPARs are shown. The service LPAR is used to build, service, create and manage z/VM production LPARs and Linux virtual machines. The service zone is the place to remote control other LPARs using standard VM tools in CP, CMS, DIRMAINT and RSCS. Networking using hipersockets in a nicely wrapped CEC box is discussed. The service zone is not your systems programmer playpen sand box!

Presentation Goals

- Describe the design and use of a service zone z/VM LPAR.
- Service zone configuration and definitions.
- Command examples.

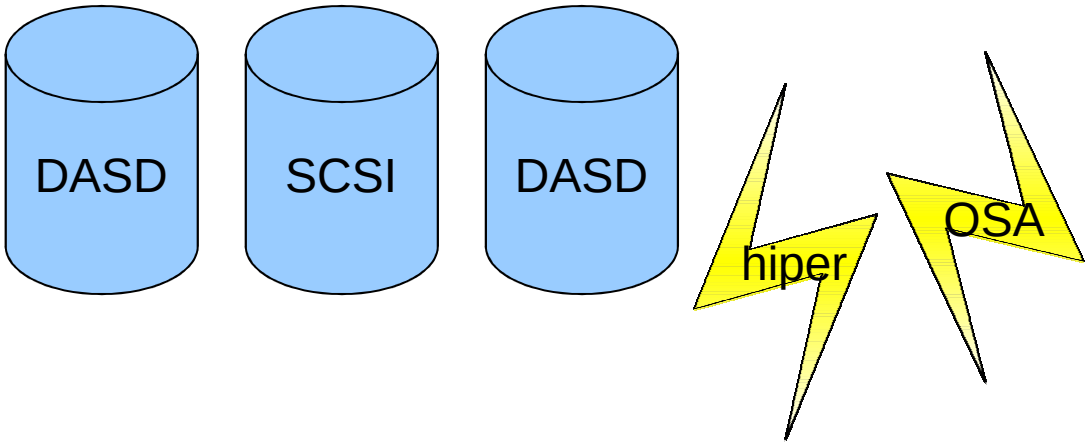

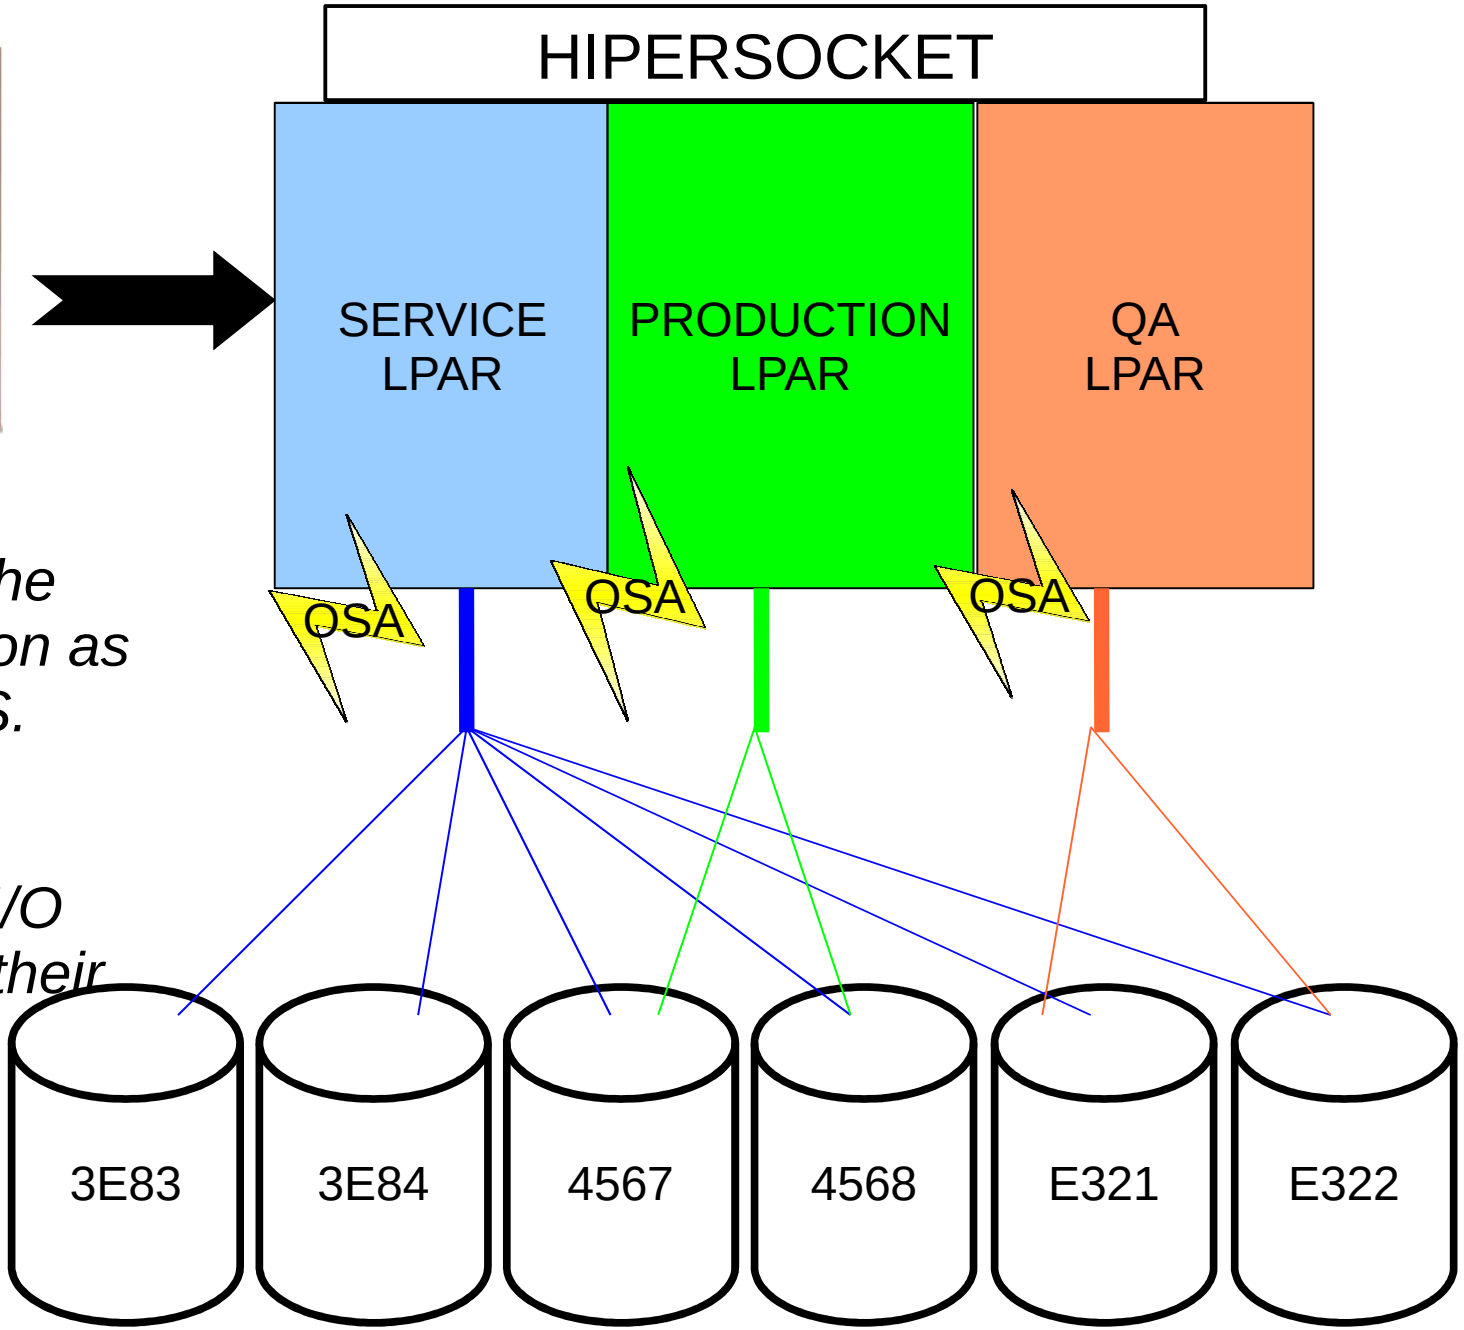

Service LPAR sees the entire I/O configuration as defined in the IOCDS. Great for systems management. Other LPARs only see the I/O devices required for their use.

Backups and Restores

Service initiates backup of 3390 volumes on all LPARs.

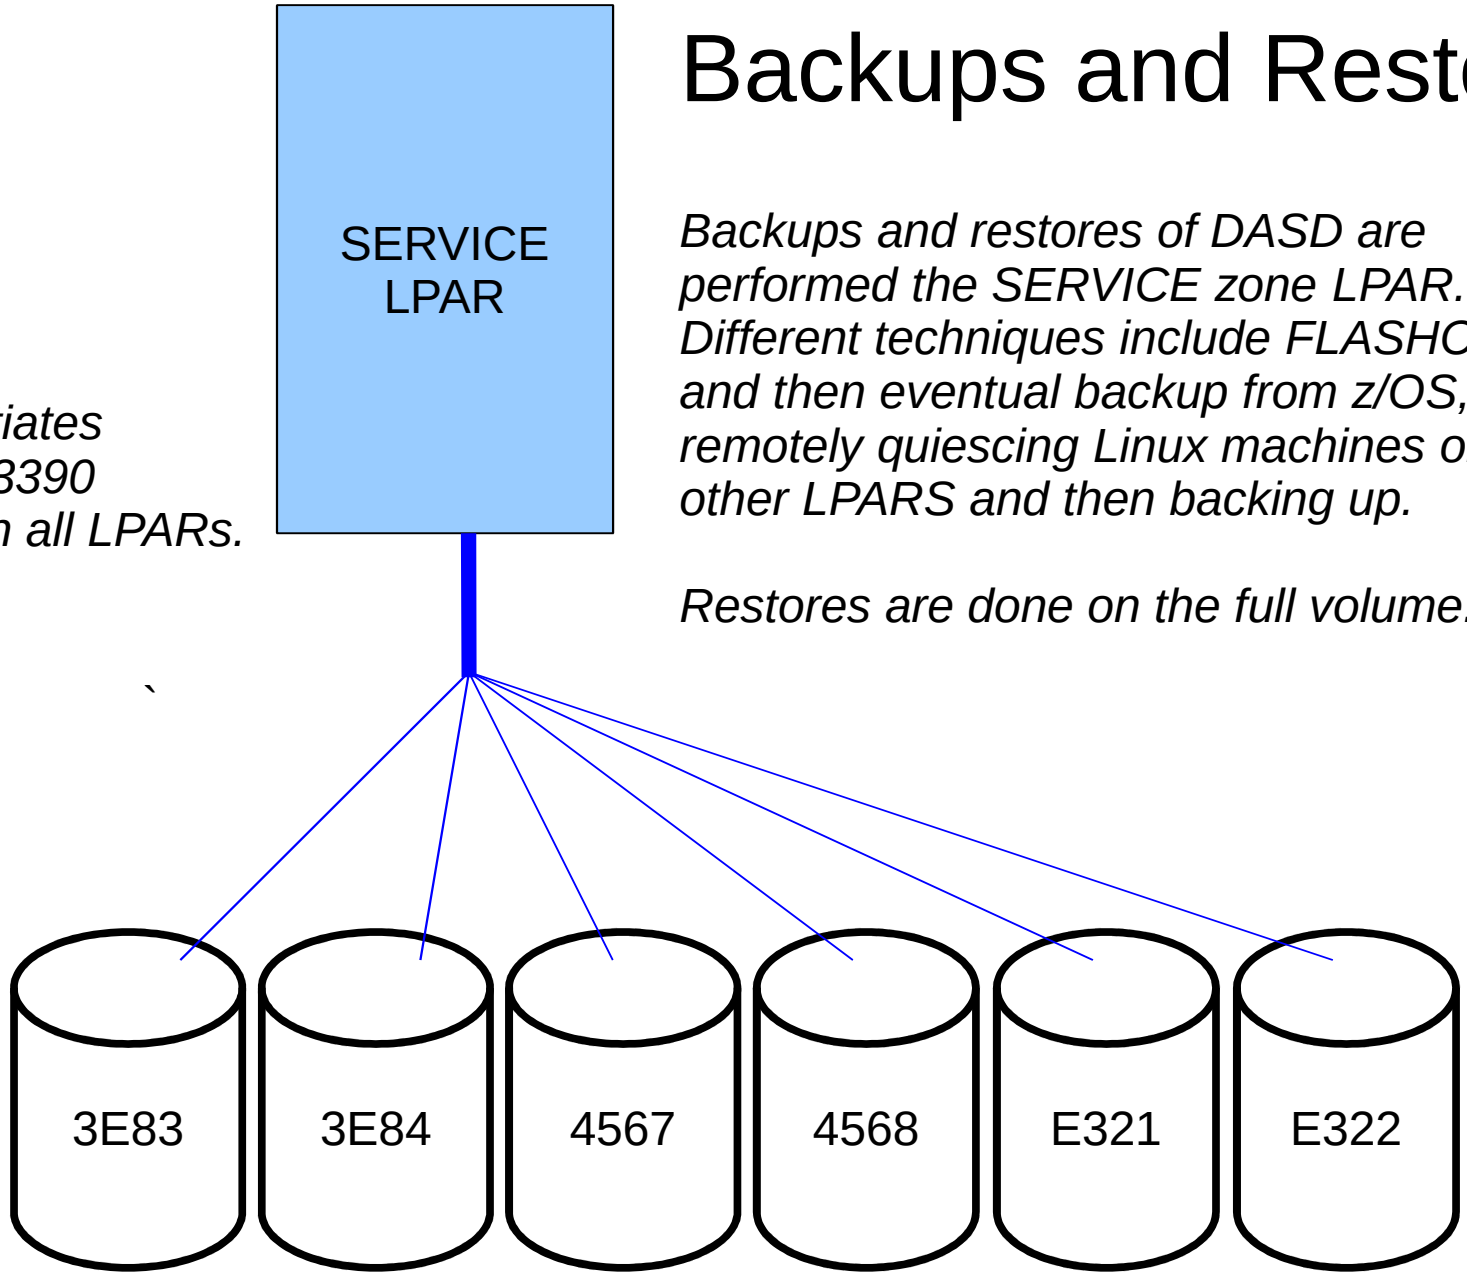

Backups and restores of DASD are performed the SERVICE zone LPAR. Different techniques include FLASHCOPY and then eventual backup from z/OS, and remotely quiescing Linux machines on the other LPARS and then backing up.

Restores are done on the full volume.

HIPERSOCKET

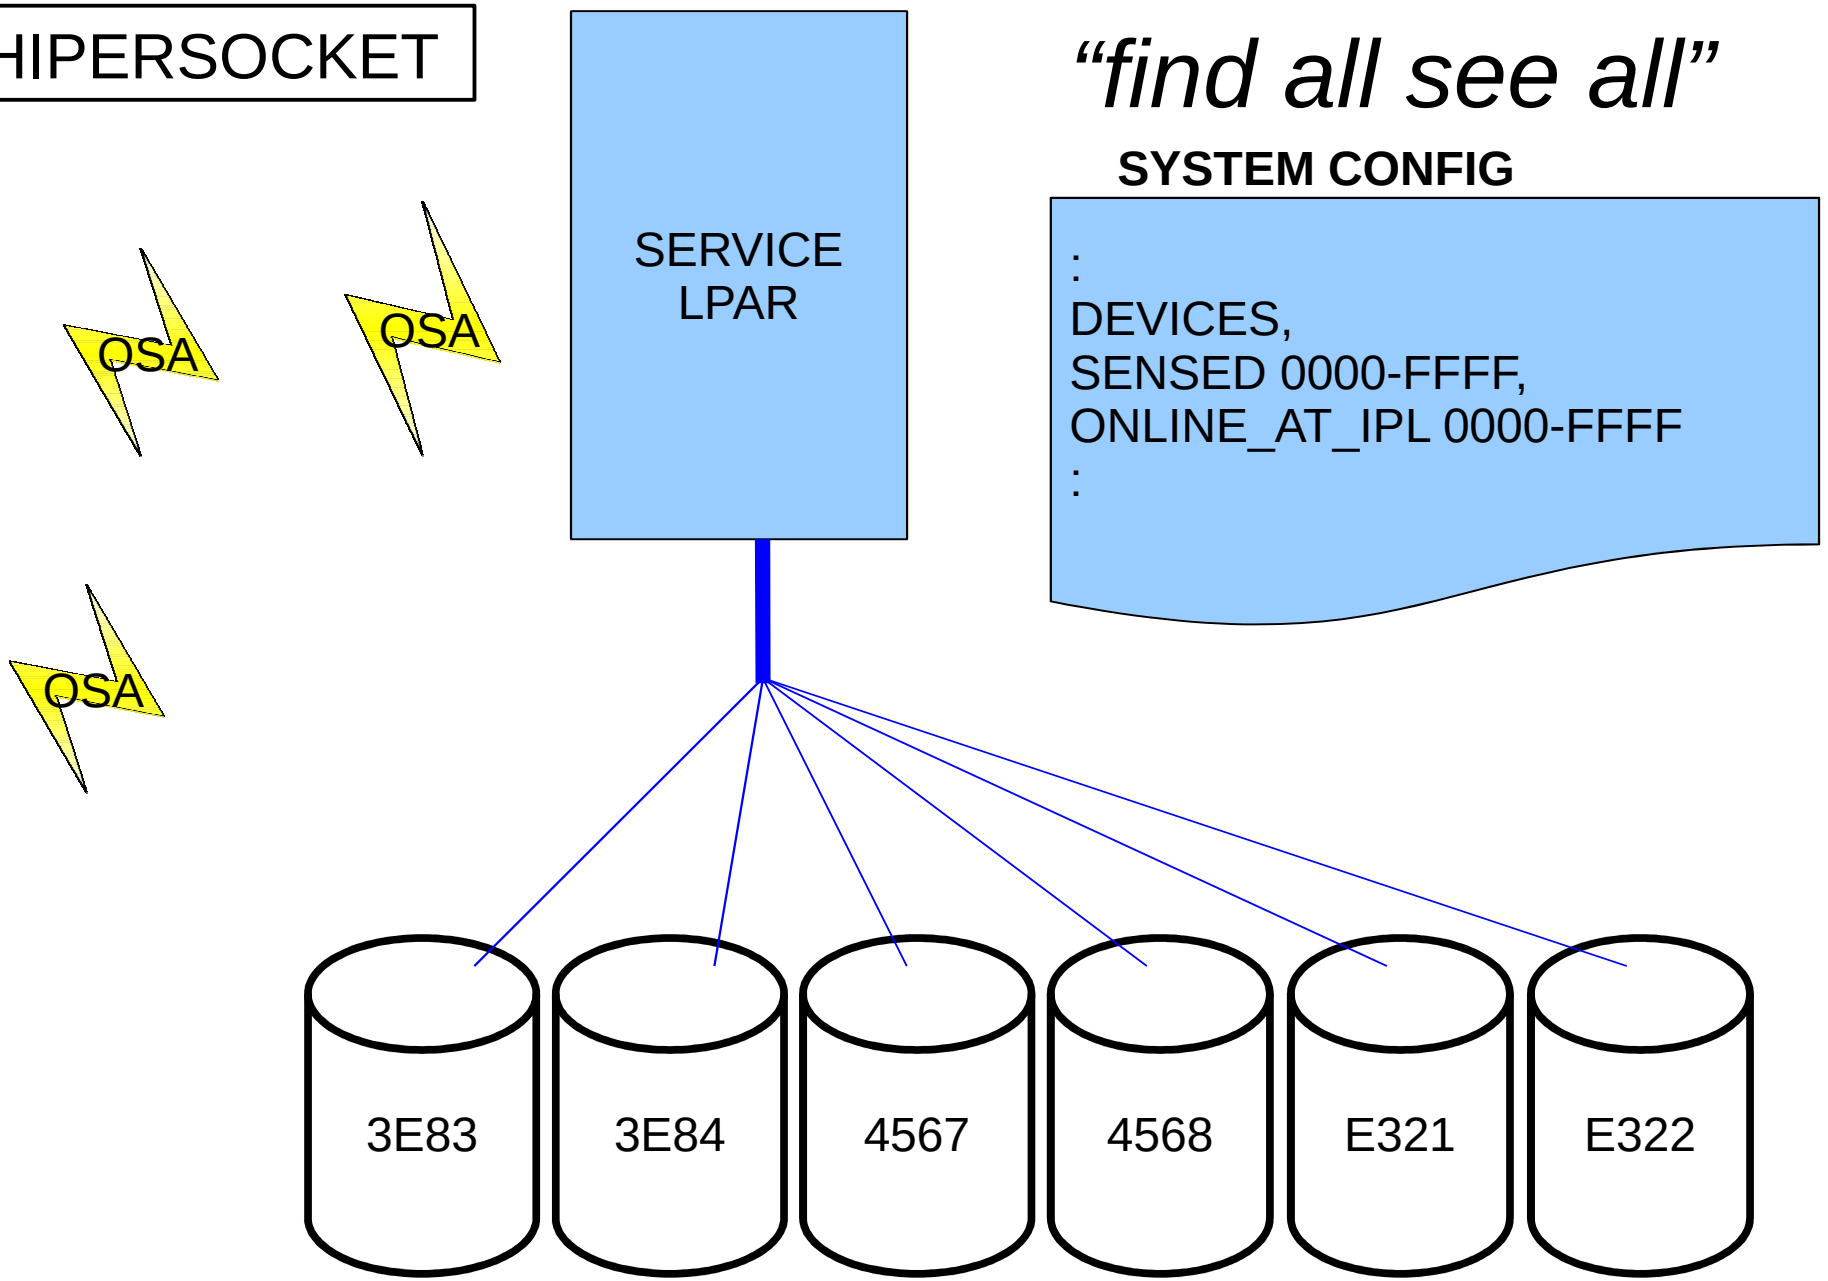

"find all see all"

SYSTEM CONFIG

```
:
DEVICES,
SENSED 0000-FFFF,
ONLINE_AT_IPL 0000-FFFF
:
```

SYSTEM CONFIG

```
.
DEVICES,
SENSED 0000-FFFF,
OFFLINE_AT_IPL 0000-FFFF,
ONLINE_AT_IPL,
4567,
4568,
DF10-DF12,  /* hiper*/
FED0-FED2  /* OSA */
```

PRODUCTION LPAR

The ONLINE_AT_IPL and OFFLINE_AT_IPL statements are OR'ed together. Great way to define only the I/O devices needed by the LPAR. Offline everything and online online what you need!

Devices,
notaccepted …,
accepted …
Is cool too!

HIPERSOCKET

DF10-
DF12

FED0-
FED2

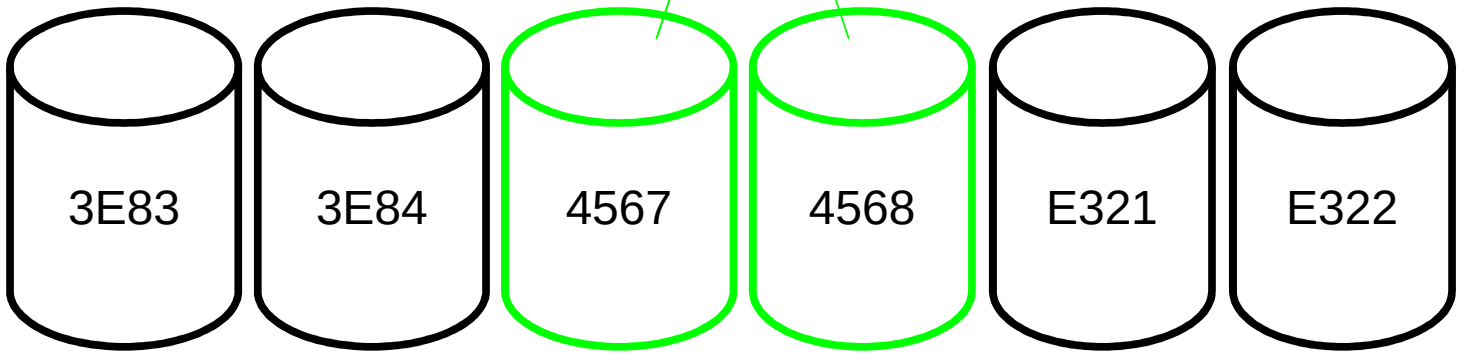

SYSTEM CONFIG

```
:
DEVICES,
SENSED 0000-FFFF,
OFFLINE_AT_IPL 0000-FFFF,
ONLINE_AT_IPL,
E321,
E322
DF13-DF15,    /* hiper */
FED3-FED5   /* OSA */
```

QA
LPAR

HIPERSOCKET

DF13-
DF15

FED3-
FED5

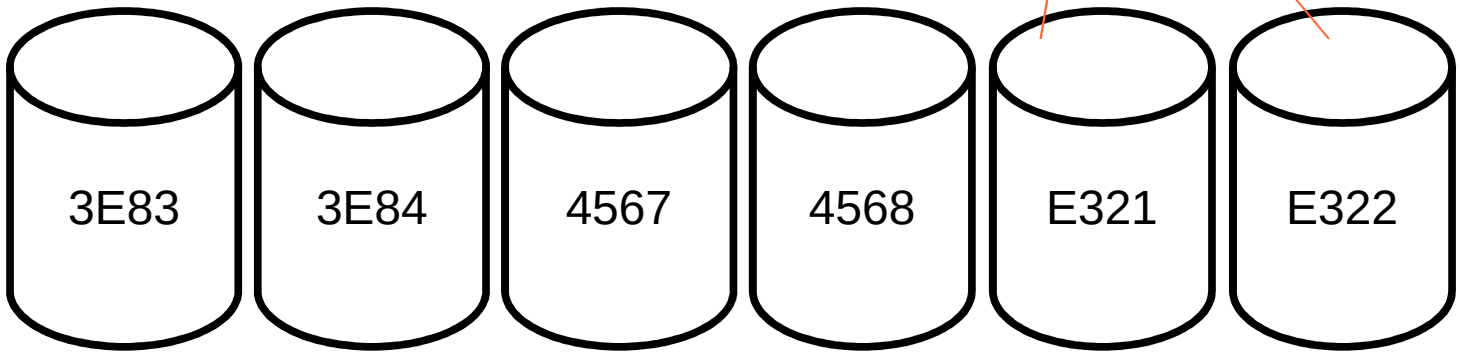

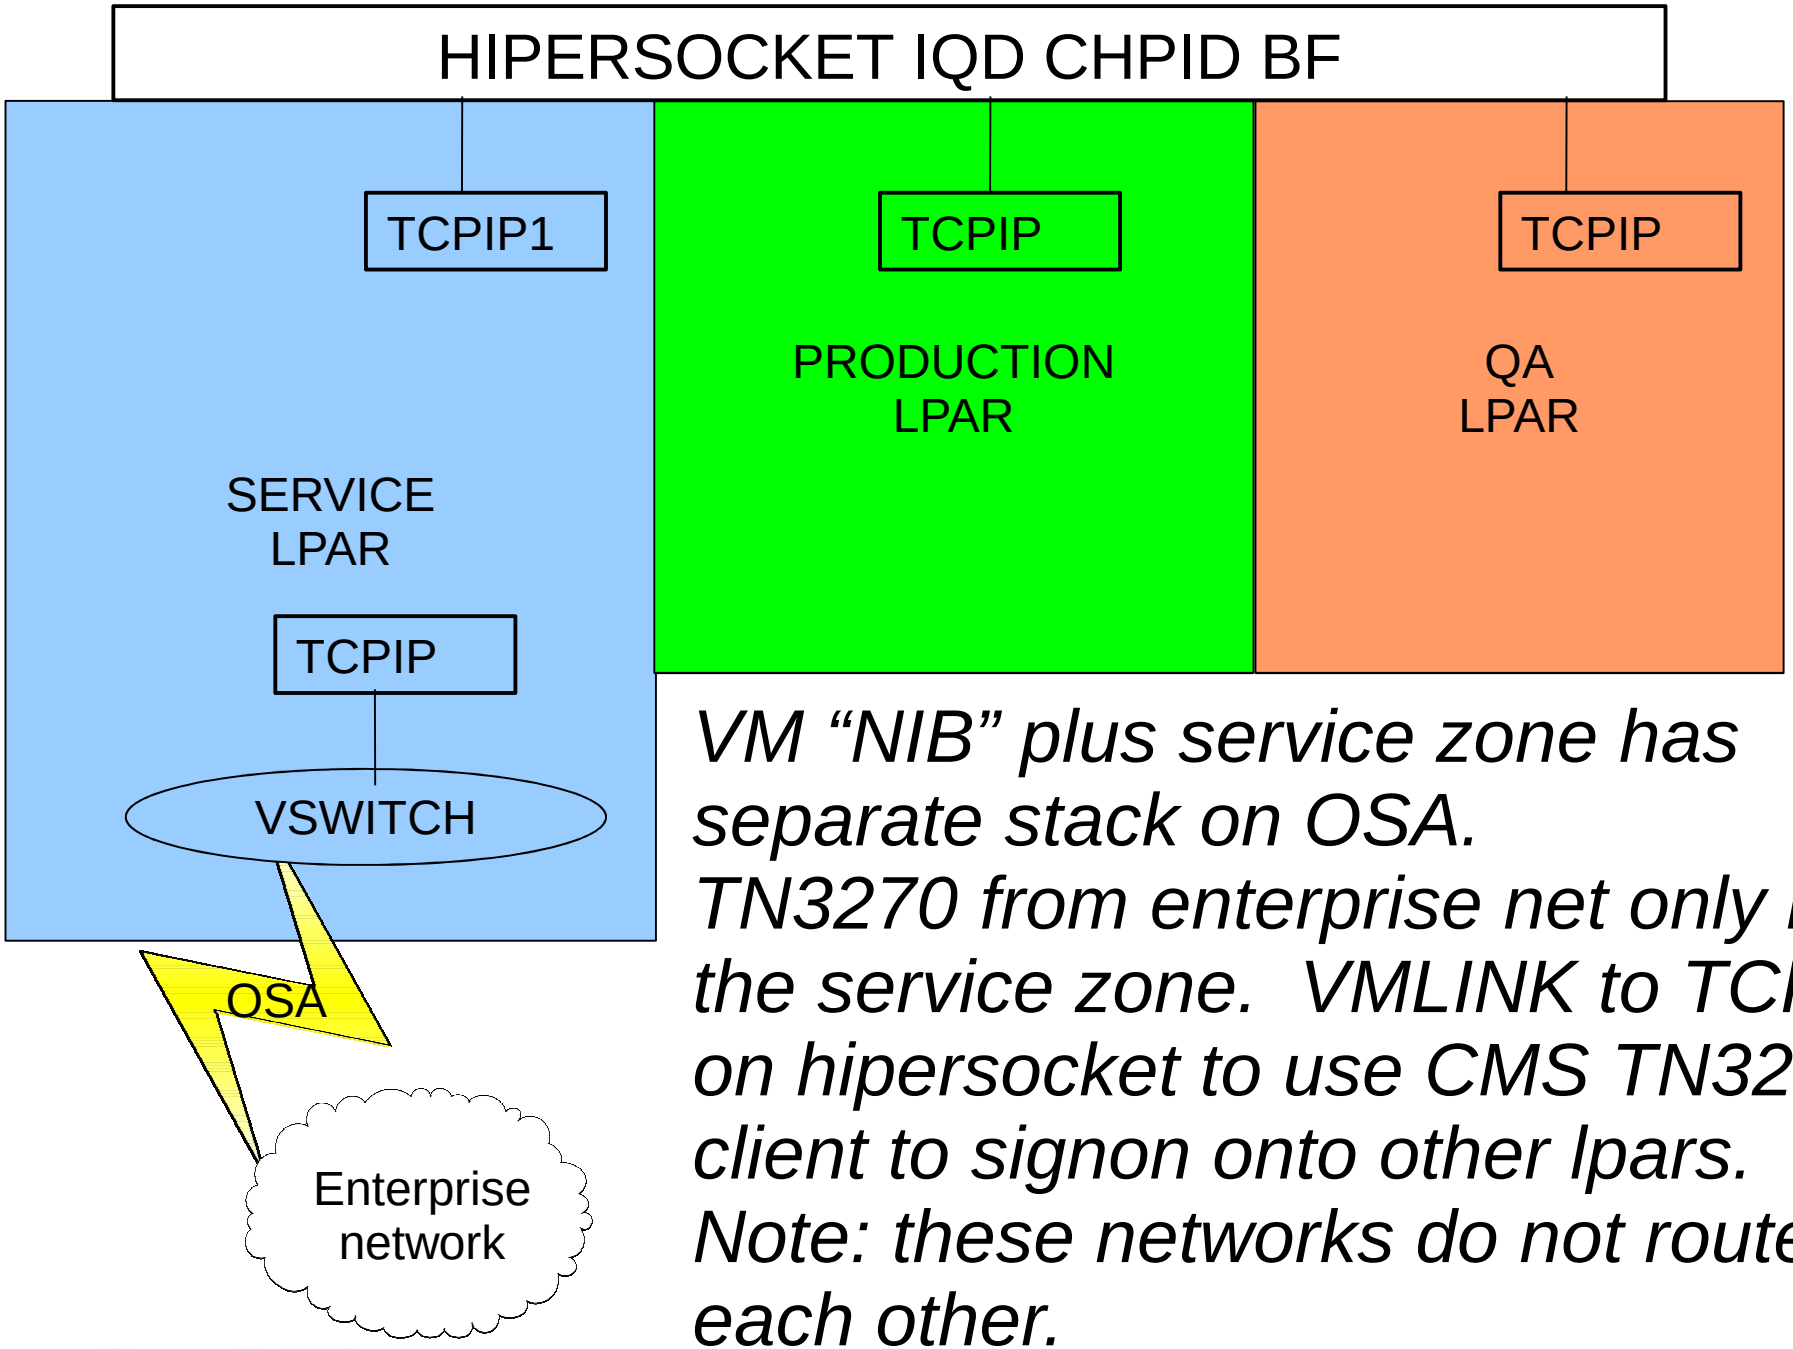

VM "NIB" plus service zone has separate stack on OSA.
TN3270 from enterprise net only into the service zone. VMLINK to TCPIP1 on hipersocket to use CMS TN3270 client to signon onto other lpars.
Note: these networks do not route to each other.

Linux v.m. enterprise networks use OSAs managed by vswitches. Linux IP addresses and ports used. No 3270 access available … or required in the Linux networks.

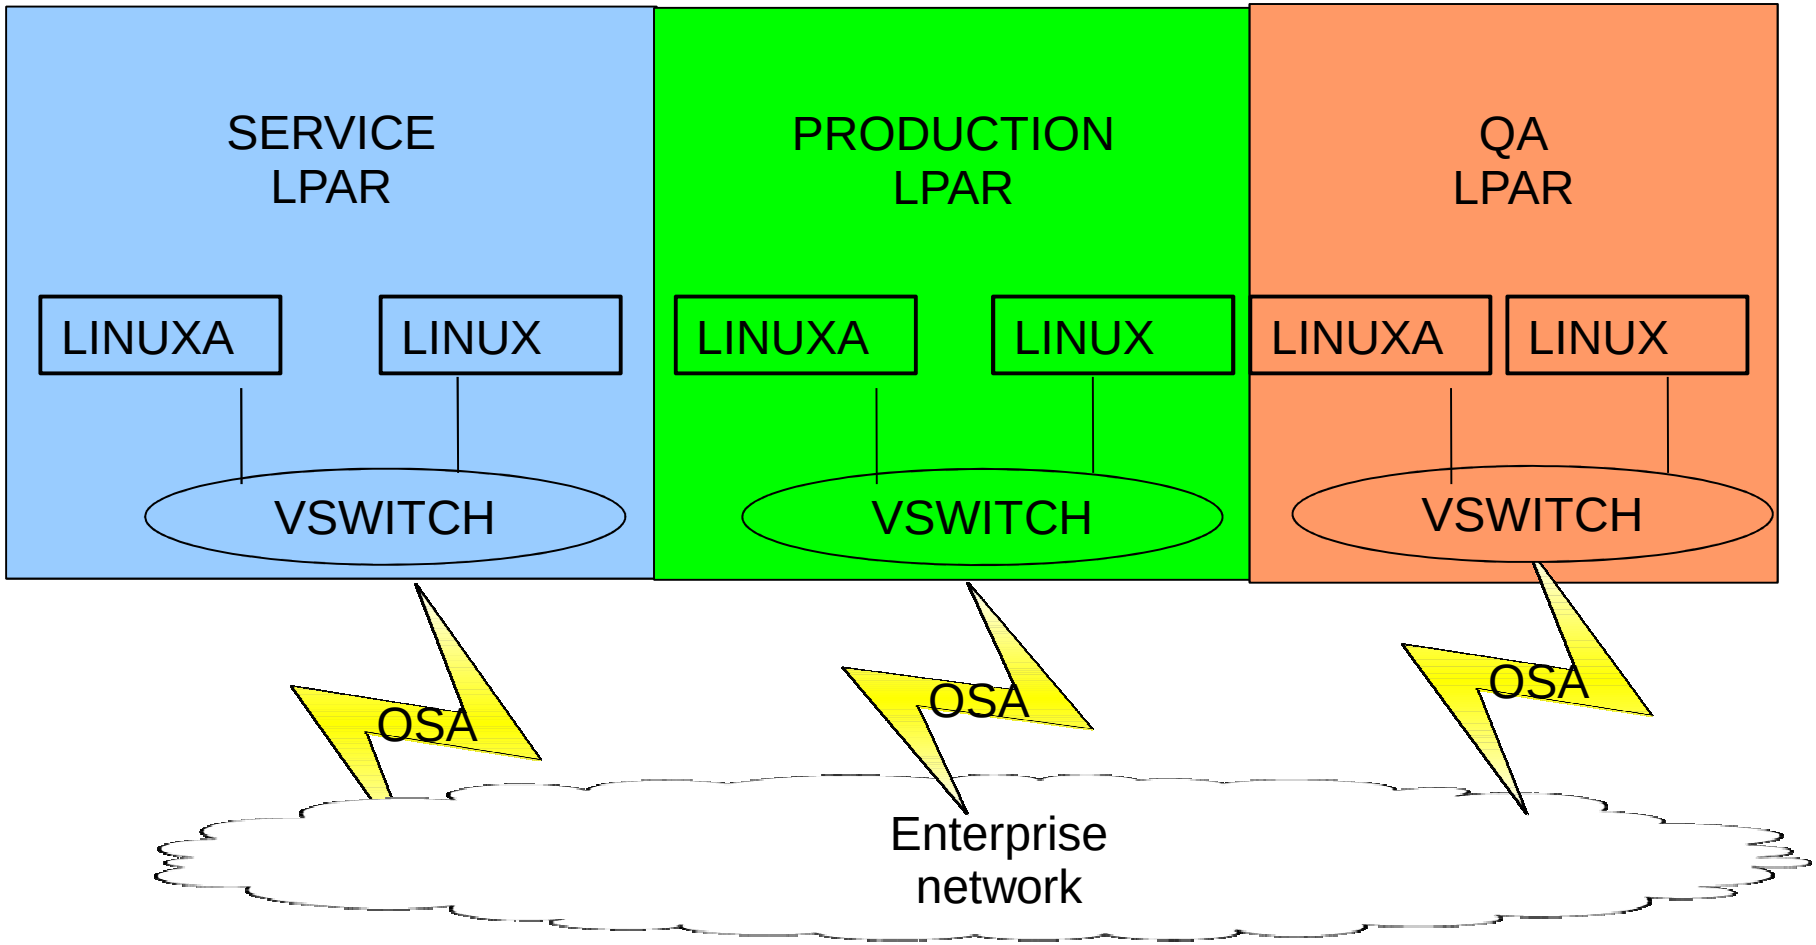

Memory and IFL Definition

Do not need many resources to be great! The two IFLs for the service zone should also be given a low weight in the LPAR profile (HMC). Memory could even be smaller!

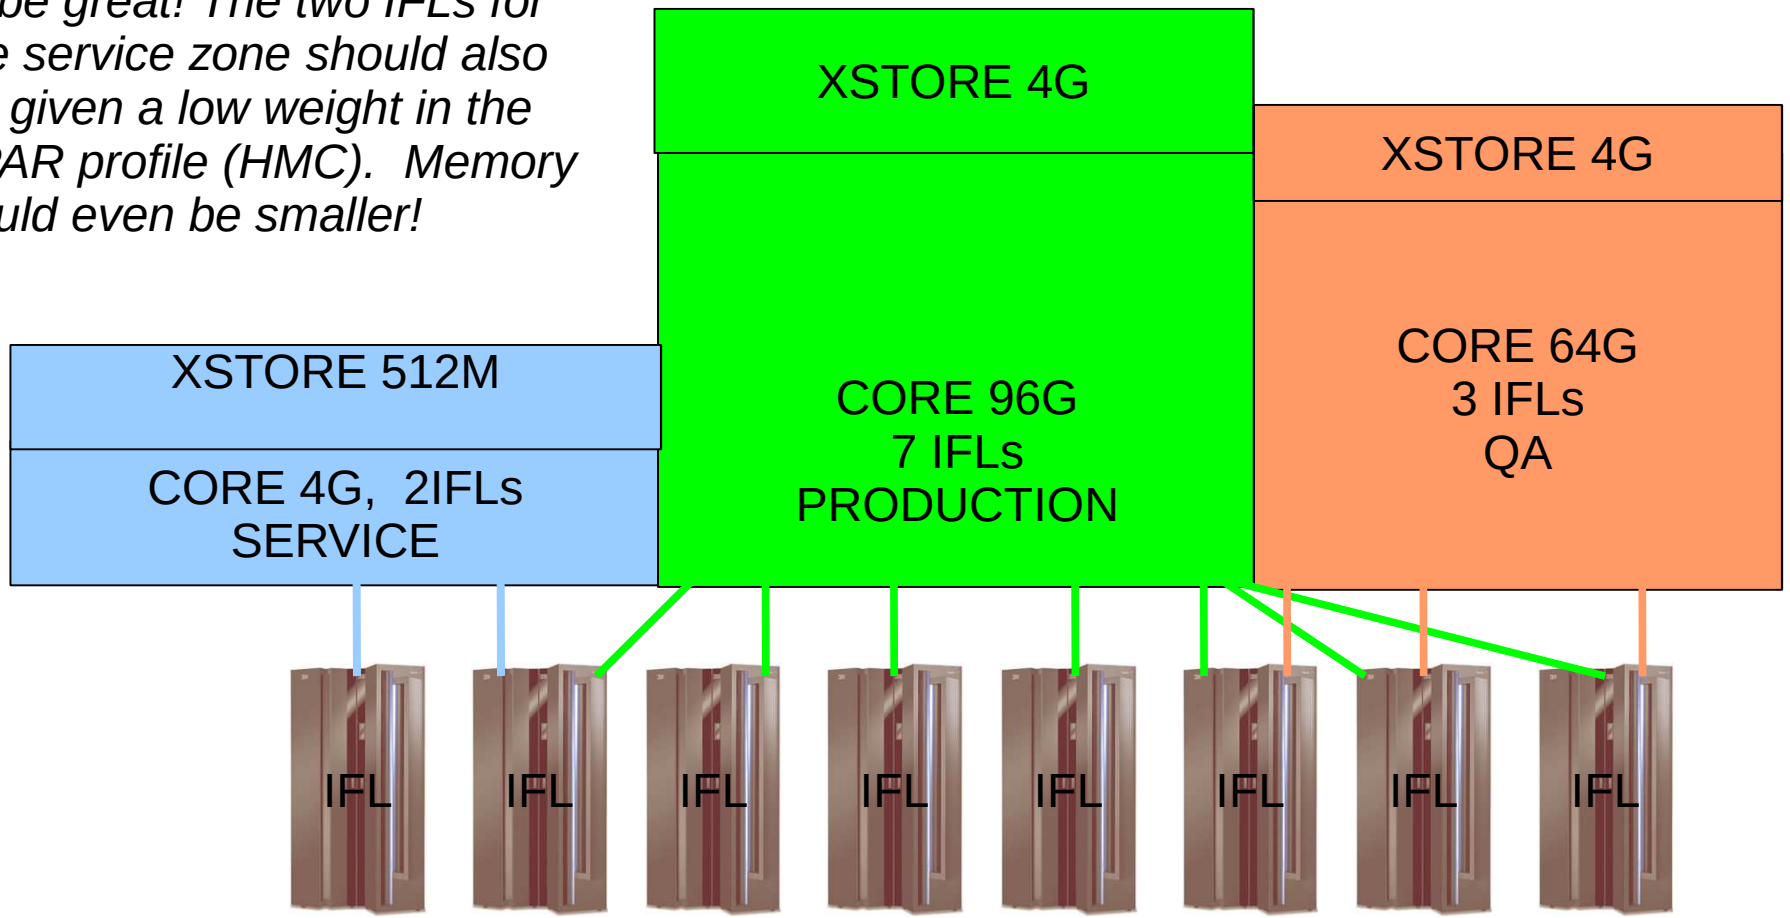

RACF DB Sharing

The RACF database is shared by all LPARs. The DASD are marked as SHAREd in the SYSTEM CONFIG. in LPARs. Most RACF administration tasks performed in the SERVICE zone.

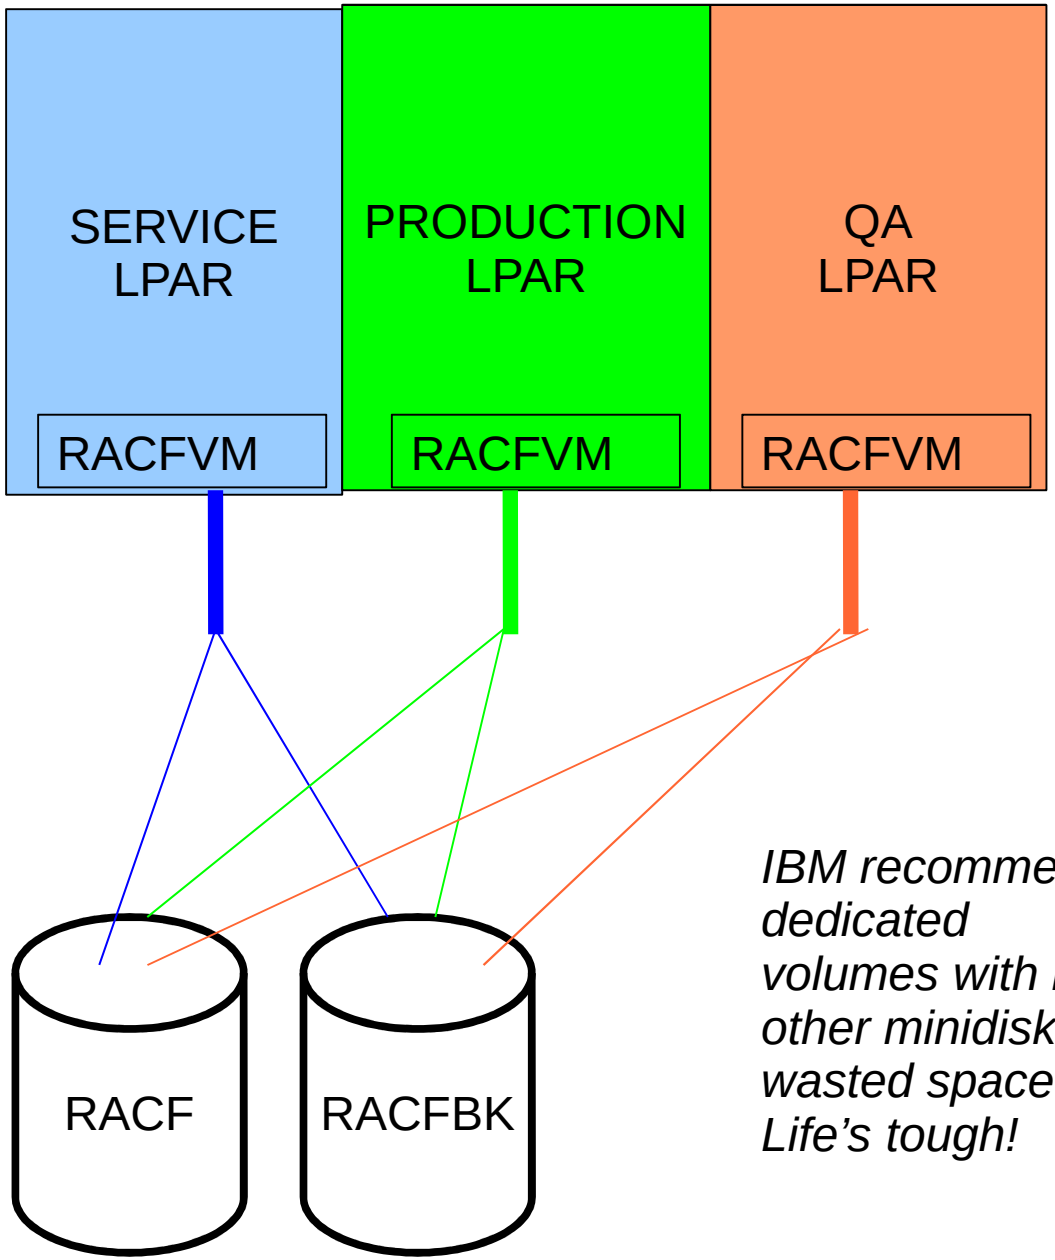

IBM recommends dedicated volumes with no other minidisks … wasted space … Life's tough!

Linux Staging

- Linux virtual machines replicated in the service zone.
- Replicated machine than delivered to the appropriate LPAR.
- New service pack servers built in the service zone.

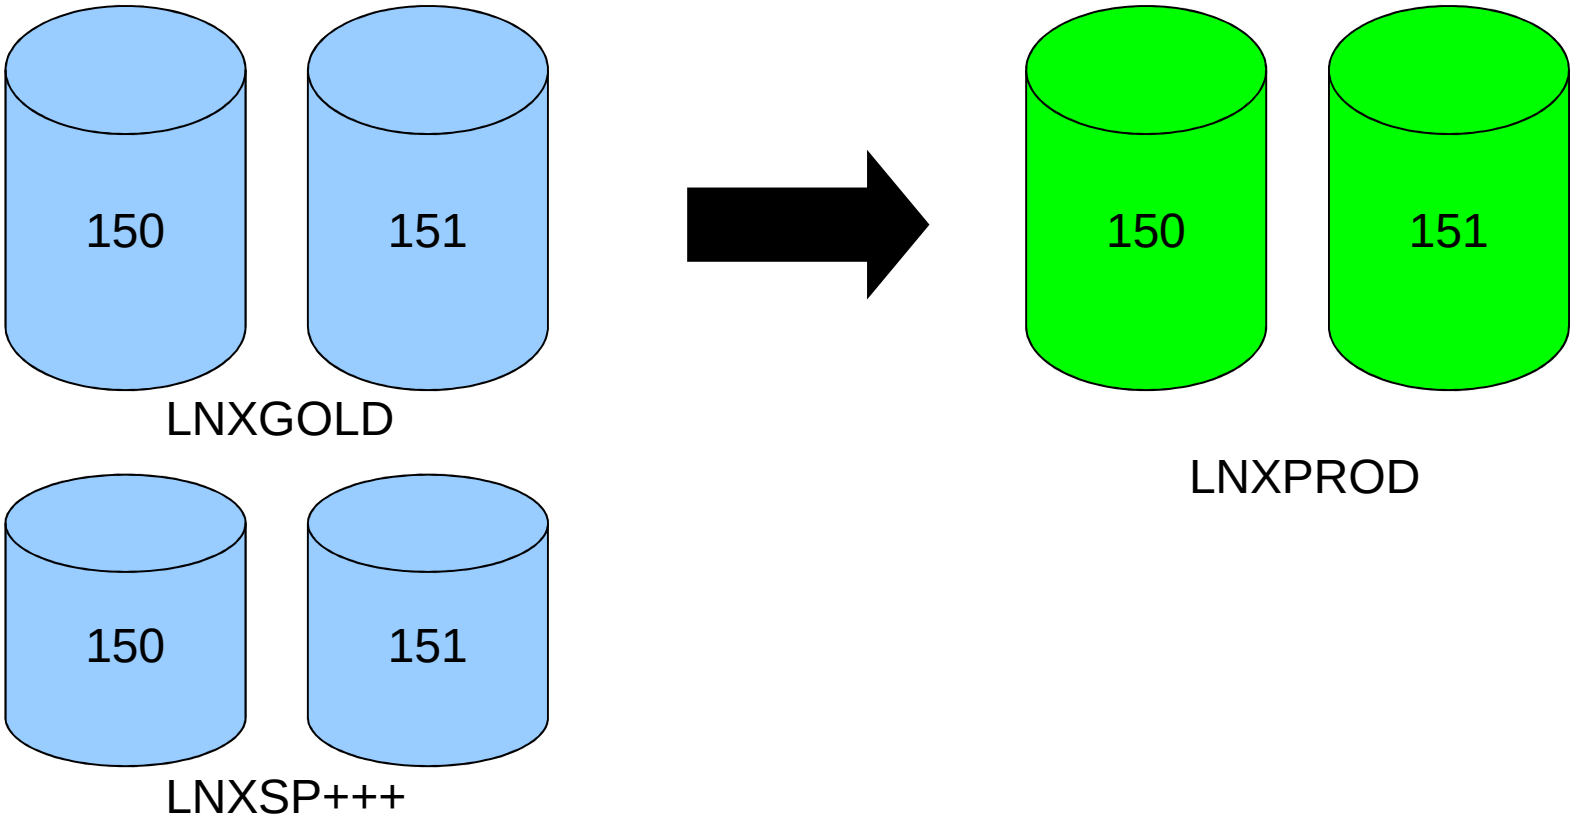

Z/VM 2nd Level Systems

- One or more 2nd level systems should be defined in the service zone – no need for 2nd level z/VM guests in other LPARS

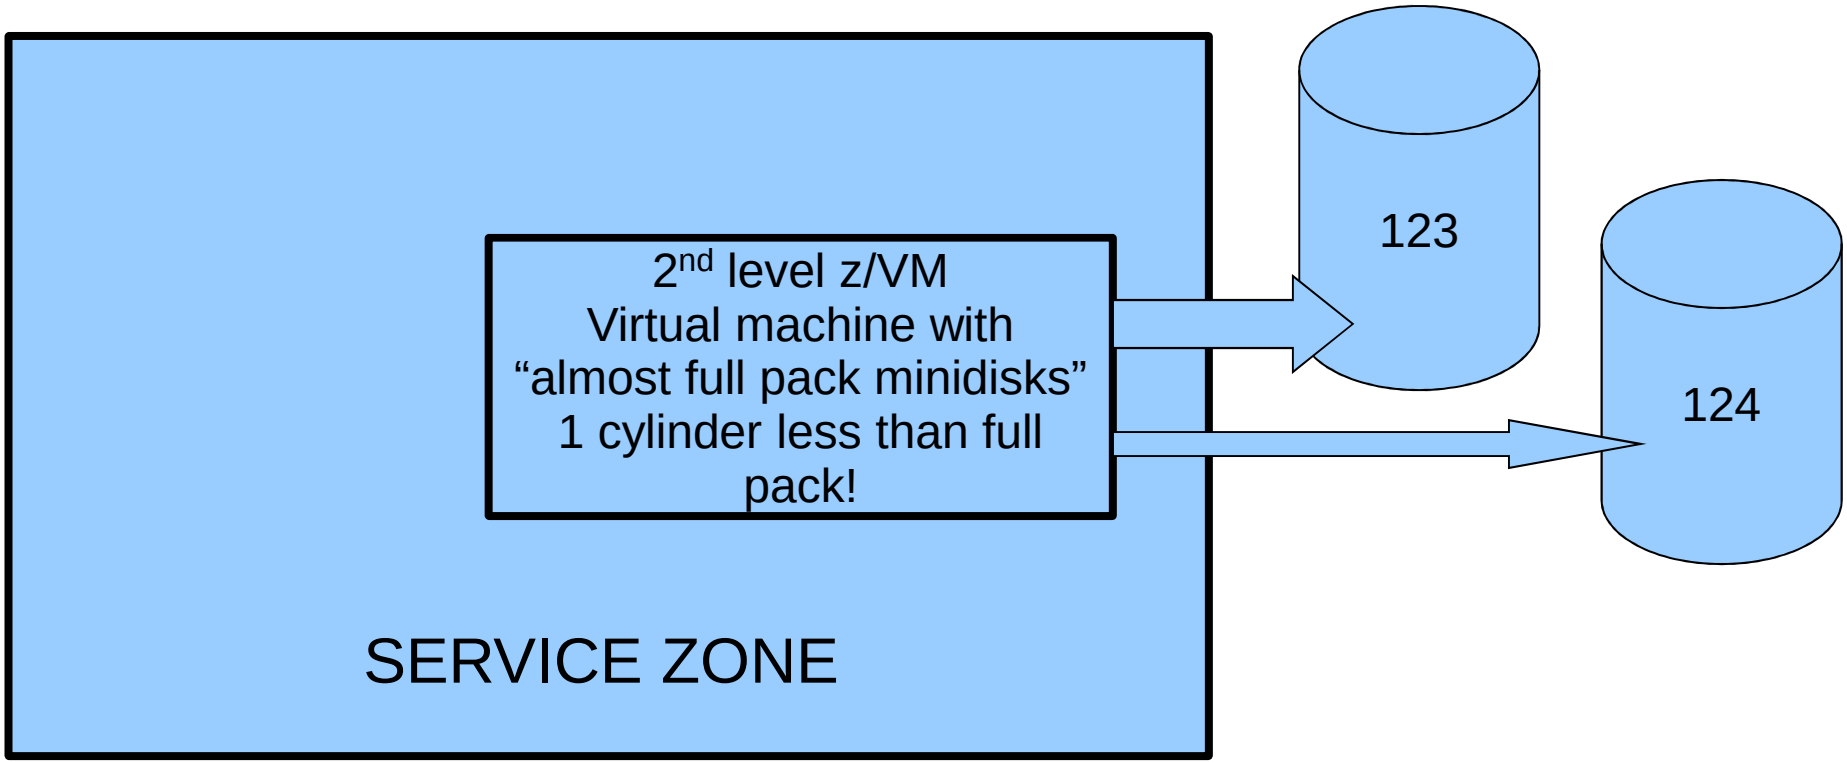

Network Services Stacking

The service zone provides systems management by using a stack of progressive services.

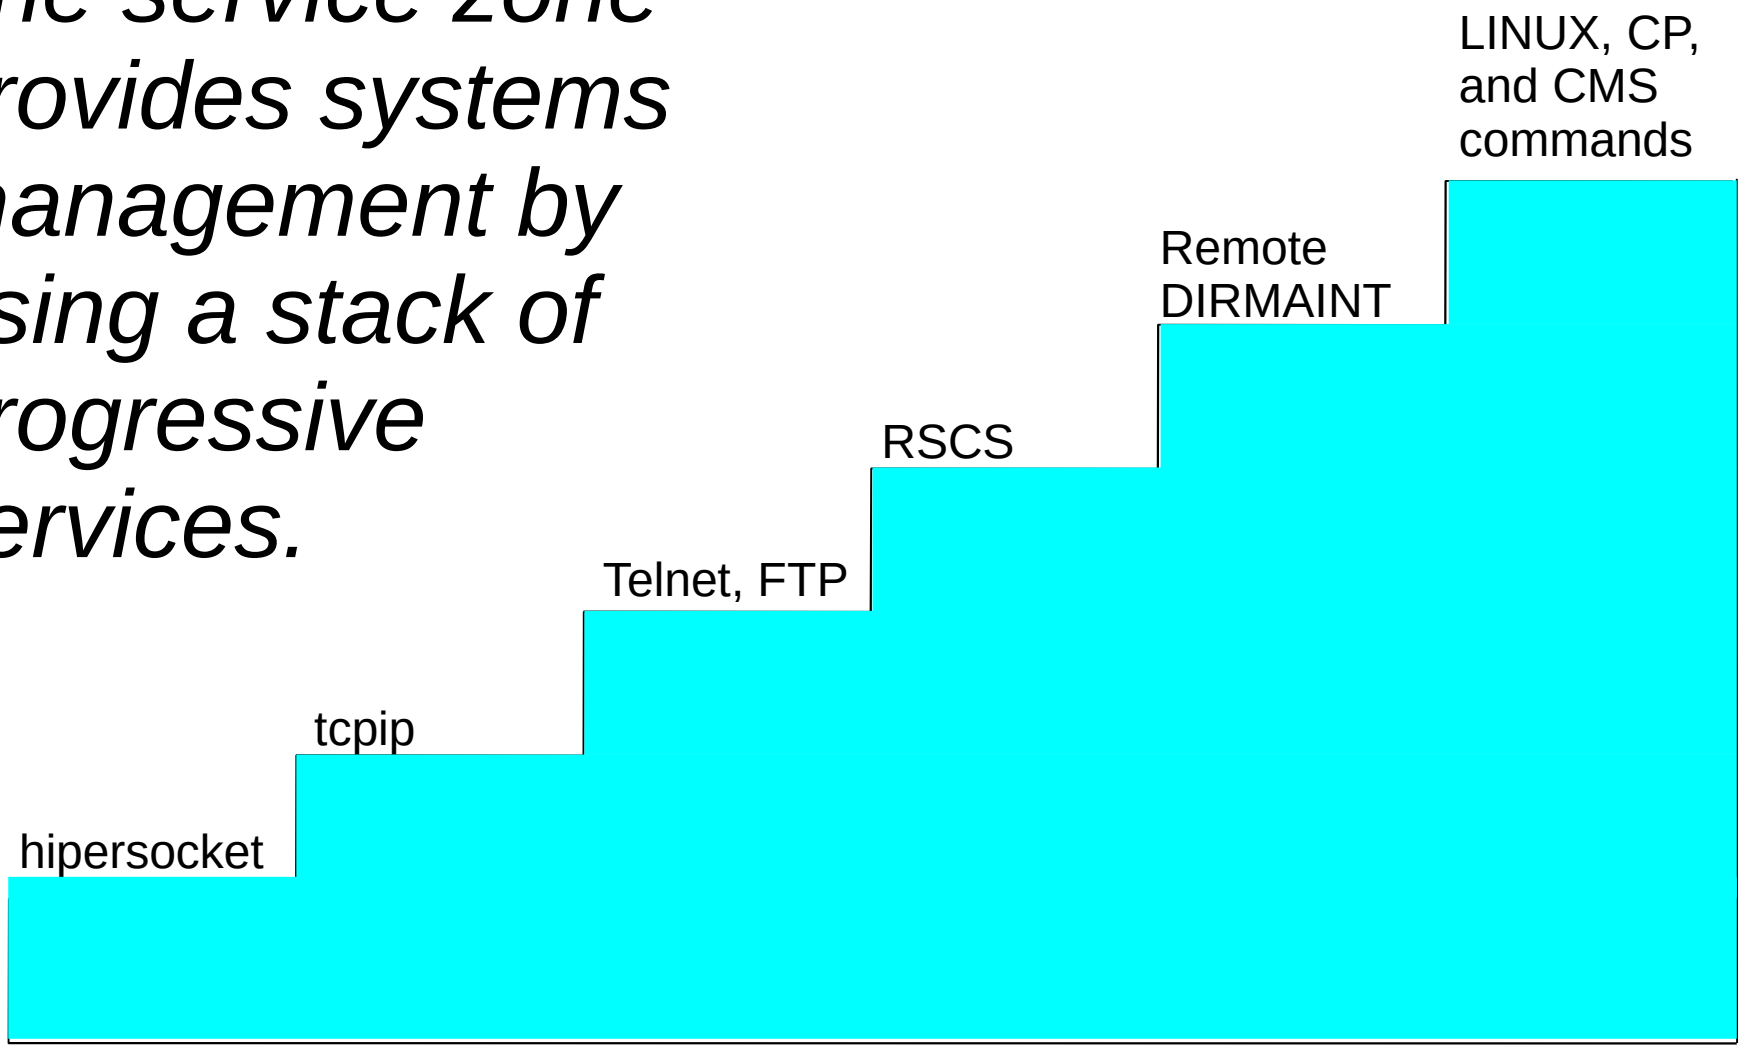

Hipersockets: “Network in a box”

- Network firmware connect between and within LPARs.
- Used to connect service zone to other LPARs.

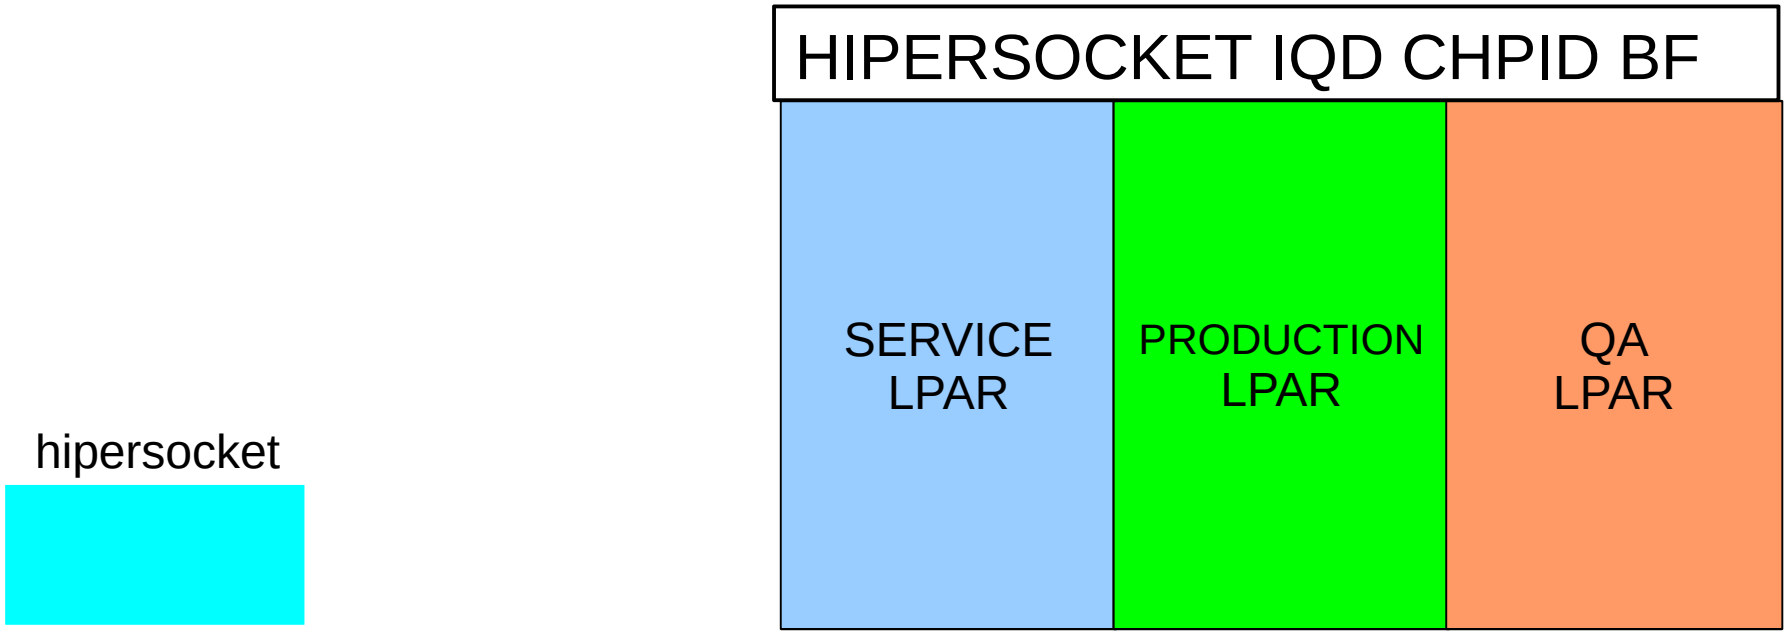

Hipersockets: “Network in a box”

- Configurations: (QUERY OSA Class B command – there is no QUERY HIPER!)

```
SERVICE LPAR:
QUERY OSA
OSA  BF18 ATTACHED TO TCPIP1 BF18 DEVTYPE HIPER  CHPID BF IQD
OSA  BF19 ATTACHED TO TCPIP1 BF19 DEVTYPE HIPER  CHPID BF IQD
OSA  BF1A ATTACHED TO TCPIP1 BF1A DEVTYPE HIPER  CHPID BF IQD
```

HIPERSOCKET IQD CHPID BF

```
    PRODUCTION LPAR:
QUERY OSA
OSA  BF00 ATTACHED TO TCPIP  BF00 DEVTYPE HIPER   CHPID BF IQD
OSA  BF01 ATTACHED TO TCPIP  BF01 DEVTYPE HIPER   CHPID BF IQD
OSA  BF02 ATTACHED TO TCPIP  BF02 DEVTYPE HIPER   CHPID BF IQD
```

hipersocket

VM TCPIP stack machine

VM TCPIP virtual machines in each LPAR connected with hipersockets.

- Configuration sample from the service zone TCPIP1 TCPIP and from the production LPAR.

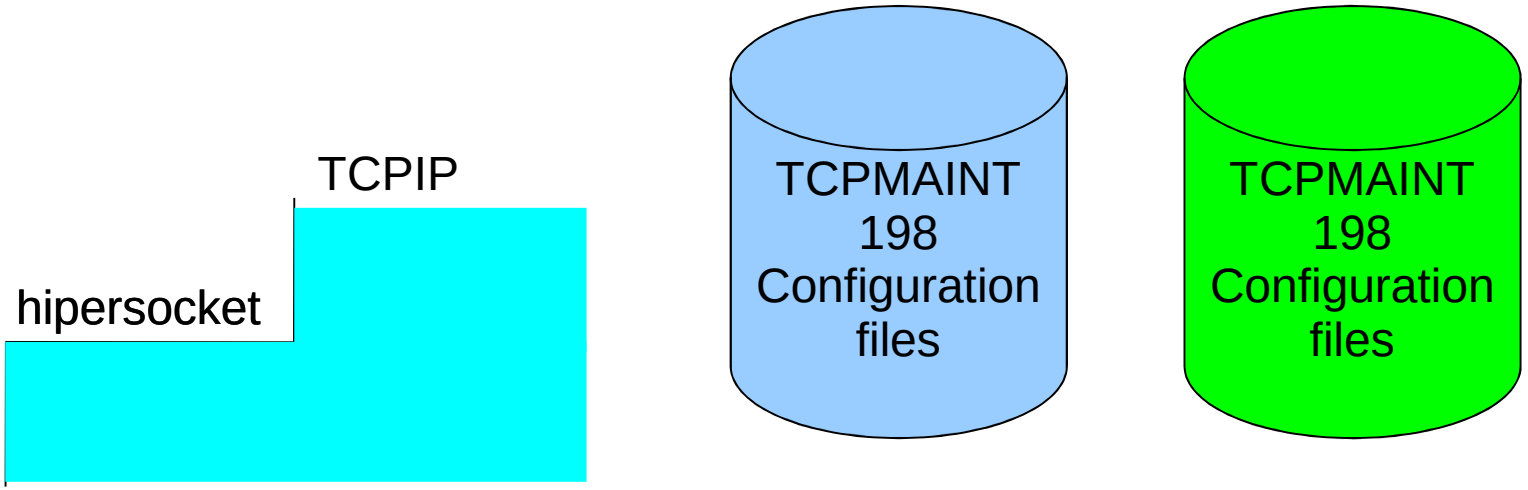

Service: VM TCPIP stack machine

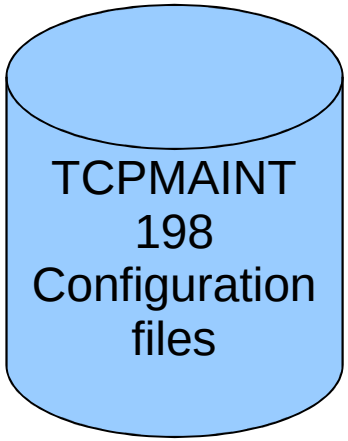


```
system dtcparms:
SYSTEM   DTCPARMS D1
 :NICK.TCPIP1      :TYPE.SERVER   :CLASS.STACK
:ATTACH.BF18-BF1A
```

```
TCPIP1 TCPIP:
DEVICE HIPER HIPERS BF18
LINK HIPER QDIOIP HIPER NOFWD MTU 0 VLAN ANY
HOME
192.168.150.1   255.255.255.0 HIPER
 START HIPER
```

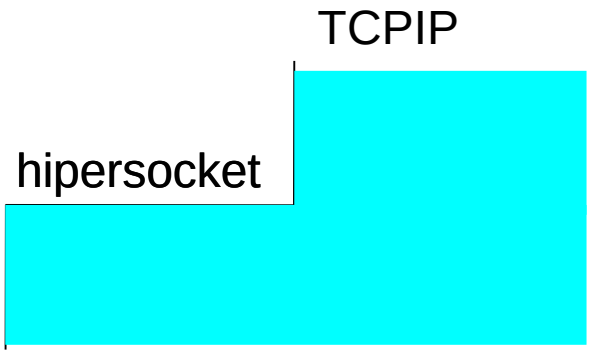


```
SYSTEM NETID:

*CPUID NODEID NETID
0A2DE5 ZGESSEA1 RSCS
111111 TCPIP1 TCPIP1
```

```
TCPIP DATA: (TCPMAINT 592)

TCPIP1: TCPIPUSERID TCPIP1
```

SYSTEM NETID used to point to TCPIP1 machine (TCPIP machine is on the enterprise net). TCPIP DATA points to the TCPIP1 machine.

Production: VM TCPIP stack machine

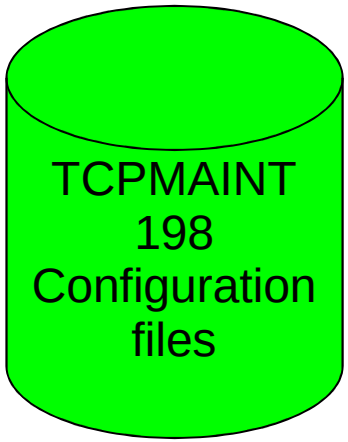


```
SYSTEM DTCPARMS:
:nick.TCPIP     :type.SERVER  :class.STACK
:ATTACH.BF00-BF02
```

```
TCPIP TCPIP:
DEVICE HIPER HIPERS BF00
LINK HIPER QDIOIP HIPER NOFWD MTU 0
HOME
192.168.150.3   255.255.255.0 HIPER

START HIPER
```

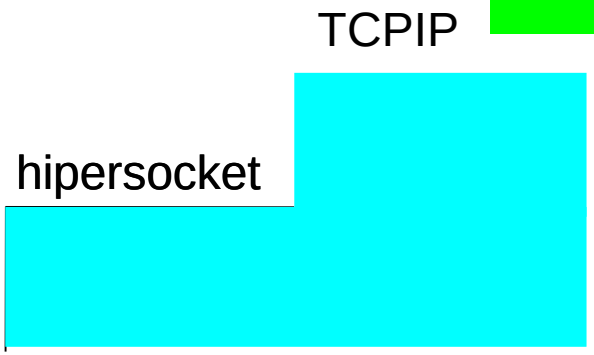

Service: TCPMAINT: Commands

```
netstat home tcp tcpip1
VM TCP/IP Netstat Level 610       TCP/IP Server Name: TCPIP1

IPv4 Home address entries:

Address         Subnet Mask      Link              VSWITCH
-------         -----------      ------            -------
192.168.150.1   255.255.255.0    HIPER             <none>

IPv6 Home address entries: None

set cpuid 111111
Ready; T=0.01/0.01 16:04:35
ping 192.168.150.3
Ping Level 610: Pinging host 192.168.150.3.
                Enter #CP EXT to interrupt.
PING: Ping #1 response took 0.023 seconds. Successes so far 1.
Ready; T=0.01/0.01 16:04:40
```

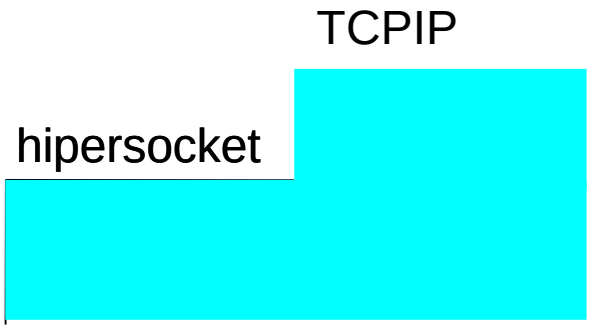

The home address is shown. Changing the CPUID is needed for the ping command to talk with the correct machine (TCPIP1).

Production: VM TCPIP stack machine

```
 netstat home
VM TCP/IP Netstat Level 610       TCP/IP Server Name: TCPIP

IPv4 Home address entries:

Address         Subnet Mask      Link                 VSWITCH
-------         -----------      ------               -------
192.168.150.3   255.255.255.0    HIPER                <none>

IPv6 Home address entries: None

Ready; T=0.01/0.01 16:05:58
ping 192.168.150.3
Ping Level 610: Pinging host 192.168.150.3.
                Enter #CP EXT to interrupt.
PING: Ping #1 response took 0.013 seconds. Successes so far 1.
Ready; T=0.01/0.01 16:06:05
```

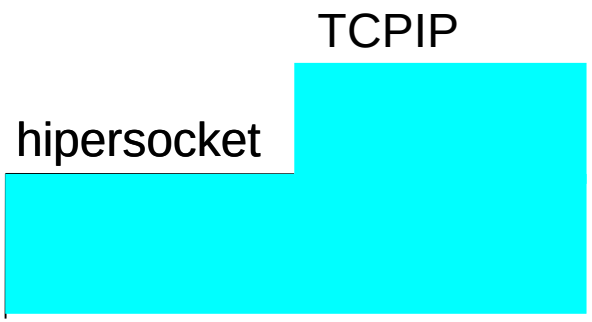

No need to set the CPUID – only 1 VM TCPIP machine on production LPAR!

VM TCPIP Server and Clients

- TN3270, FTP, and SMTP

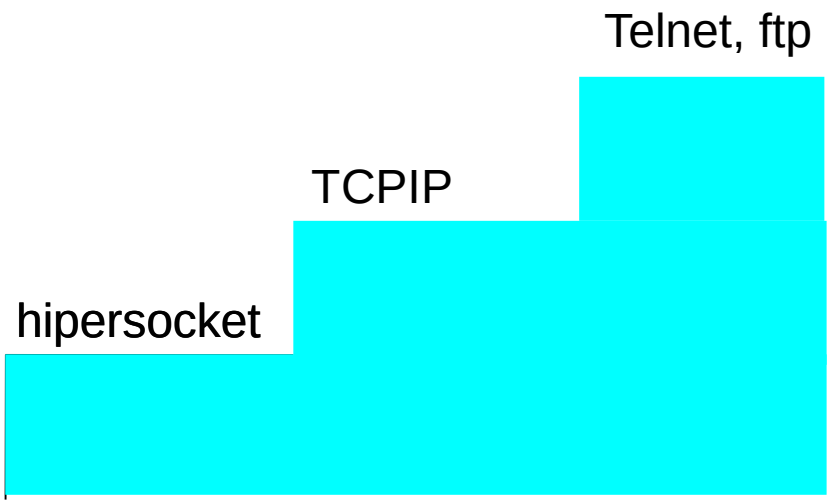

usigtel1:
*** TELNET *** 15:23:13

1) PRODUCTION 1
3) PRODUCTION 2
4) QA 1
5) QA 2
7) SERVICE

Which LPAR do you want to signon to?

This is from the USIGTEL1 virtual machine. It runs an EXEC that asks which LPAR you want to signon to and then will LINK to TCPMAINT user code and run the CMS TELNET EXEC.

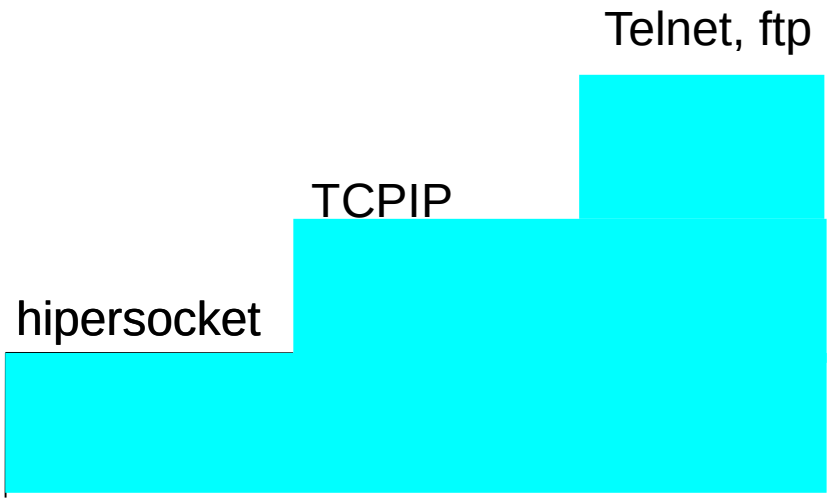

```
SELECT
  WHEN LPAR = "PRODUCTION 1",
     |  LPAR  = '192.168.150.3'   ,
      THEN ip = '192.168.150.3'
  WHEN LPAR = "QA 1",
     |  LPAR  = '192.168.150.4',
      THEN ip = '192.168.150.4'

:
 'VMLINK TCPMAINT 592 (NONAMES'
 'VMLINK TCPMAINT 198 (NONAMES'
 'CP SET CPUID 111111'
 'TELNET ' ip
```

This code extract will assign the IP address, then LINK and run the CMS TELNET command.

The TCPMAINT 592 also has control information in the the TCPIP DATA file.

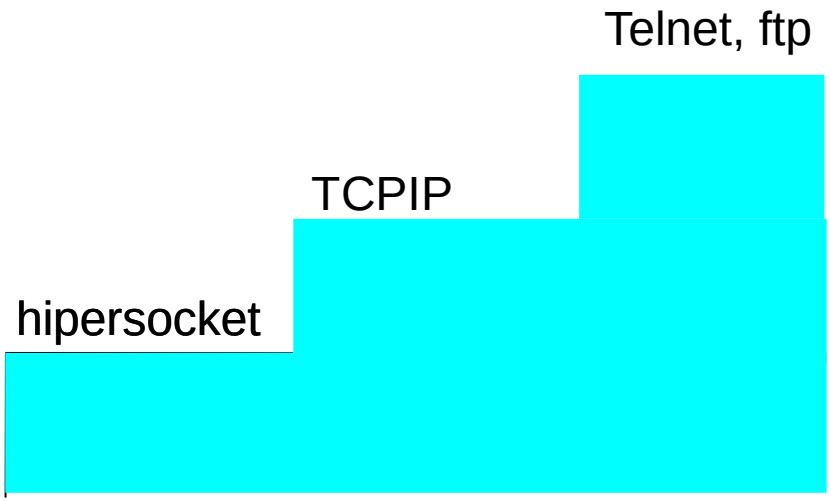

The other session is on the production side

```
 netstat home
VM TCP/IP Netstat Level 610       TCP/IP Server Name:
TCPIP

IPv4 Home address entries:
```

For example this screen shown earlier is the result of the code from USIGTEL1 – production LPAR chosen and TCPMAINT on production LPAR logged onto.

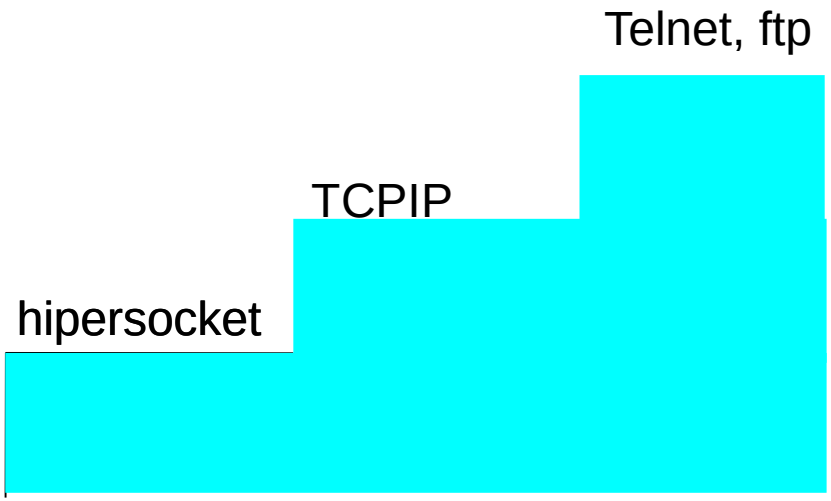

The LOGOFF

```
log
CONNECT= 00:08:01 VIRTCPU= 000:00.02 TOTCPU= 000:00.06
LOGOFF AT 15:43:55 EDT SUNDAY 07/24/11
```

```
Press enter or clear key to continue


Session ended. <ENTER> to return to CMS.
Telnet terminated -- Connection closed
CONNECT= 00:20:54 VIRTCPU= 000:00.03 TOTCPU= 000:00.06
LOGOFF AT 15:44:05 EDT SUNDAY 07/24/11

Press enter or clear key to continue
```

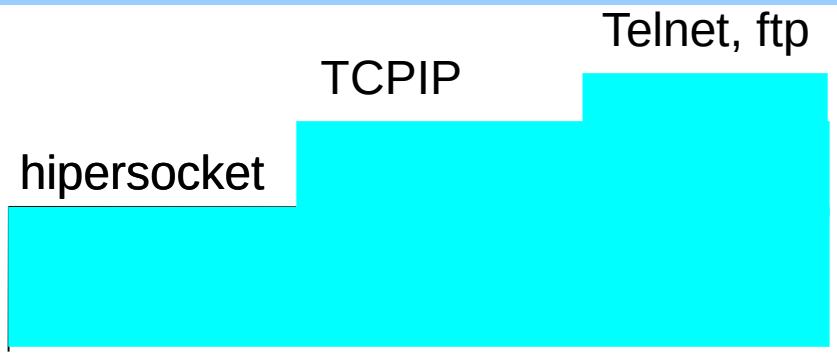

Logoff from TCPMAINT in production will return to USIGTEL1 code. The code will do a LOGOFF of the service zone session.

RSCS

Using RSCS (Remote Spooling Communication Subsystem) is highly recommended. It is great for sending files from the service zone to the other lpars, z/os, and CECs.

It is also the carrier pigeon for delivering and receiving remote DIRMAINT commands, and the issuance and delivery of CP commands.

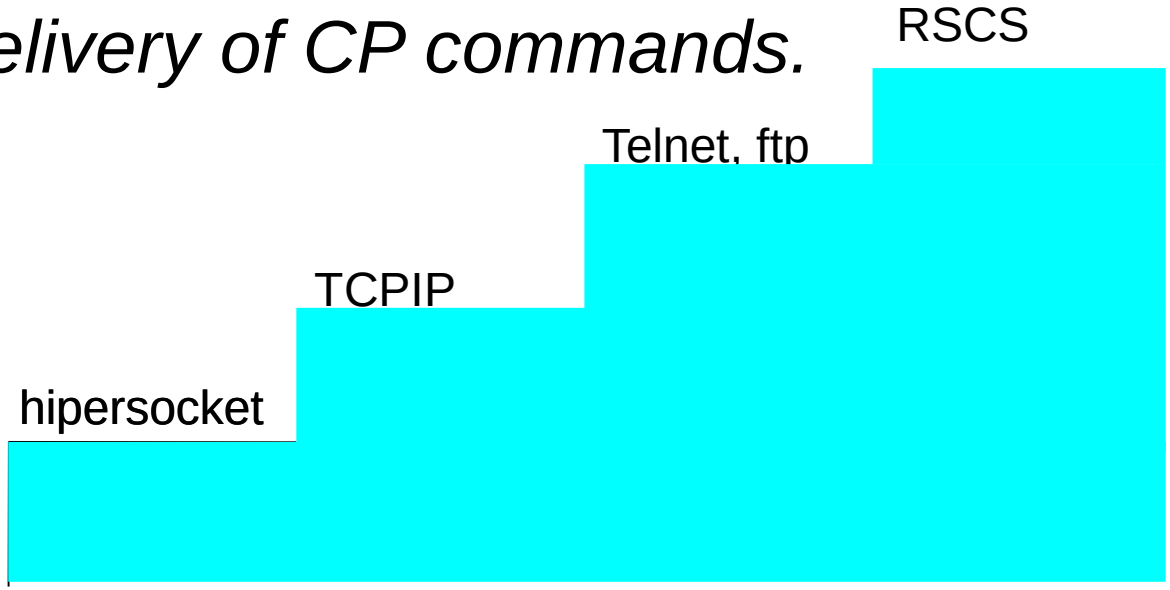

RSCS Configuration

```
RSCS CONFIG:
LOCAL     SERVICE     *  RSCS
:
 LINKDEFINE PRODTION TYPE TCPNJE NODE PRODTION ASTART RETRY
 PARM PRODTION TCPID=TCPIP1 HOST=192.168.150.3
```

SERVICE

6VMRSC10
401
Configuration
Files

PROD

6VMRSC10
401
Configuration
Files

```
RSCS CONFIG:

LOCAL PRODTION * RSCS
LINKDEFINE SERVICE TYPE TCPNJE NODE SERVICE ASTART RETRY
PARM SERVICE TCPID=TCPIP  HOST=192.168.150.1
```

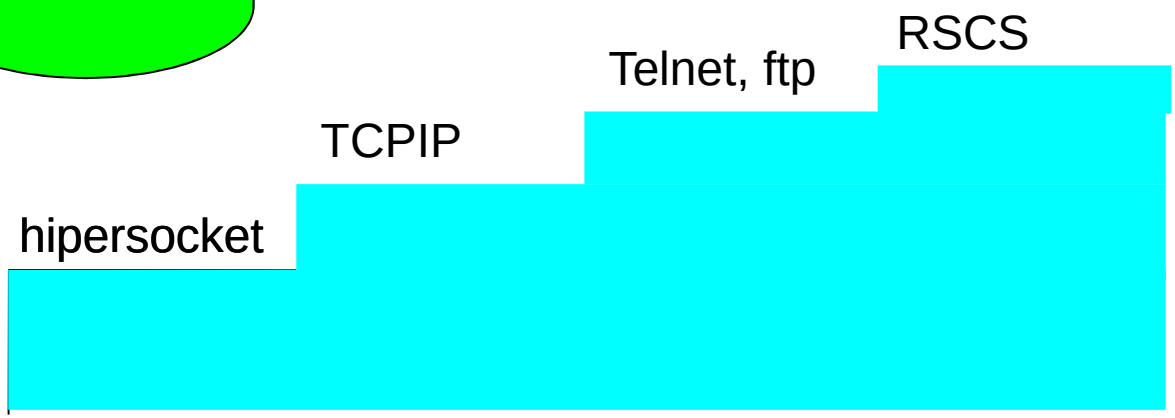

Service zone RSCS Commands

```
sm rscs q links name prodtion show parm
Ready; T=0.01/0.01 15:33:14
 Parm Text
 TCPID=TCPIP1 HOST=192.168.150.3
 1 link found
sm rscs q links name prodtion
Ready; T=0.01/0.01 15:33:19
 Link                          Line
 Name     Status     Type      Addr LU Name  Logmode  Queueing
 PRODTION connect    TCPNJE    0000 ...      ...      priority
 1 link found

:
```

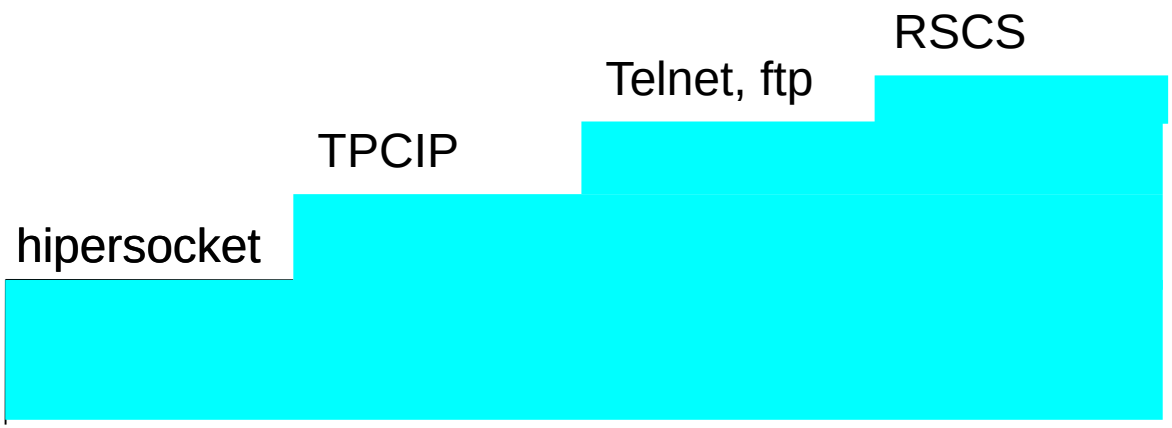


```
sm rscs q sy local
Ready; T=0.01/0.01 16:53:55
 RSCS local ID SERVICE ..
application ID RSCS
```

Service zone send file to production

```
sendfile profile exec to maint at prodtion
File PROFILE EXEC A1 sent to MAINT at PRODTION on 07/24/11 17:01:10
Ready; T=0.01/0.01 17:01:10
 From PRODITON: DMTAXM104I File (5511) spooled to MAINT
 -- origin SERVICE(MAINT) 07/24/11 17:01:09 EDT
:
```

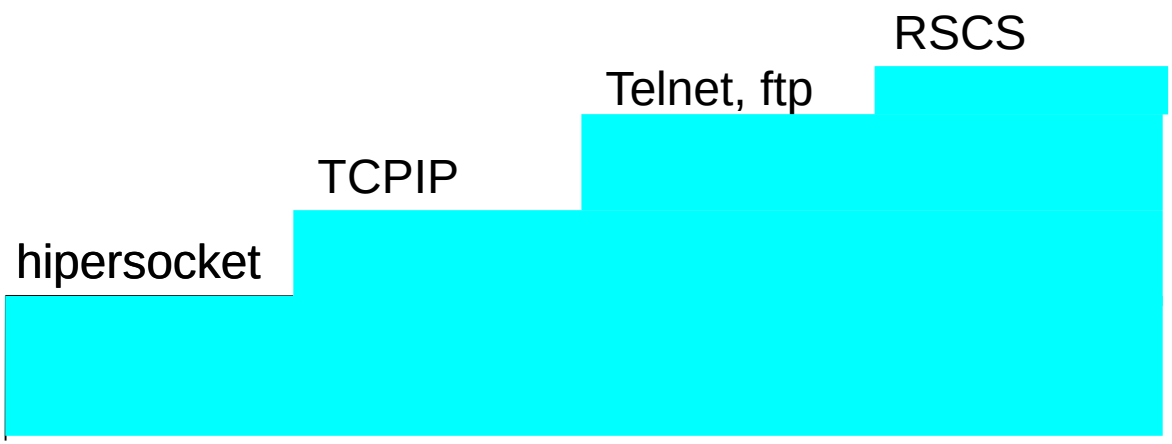


```
sm rscs q sy local
Ready; T=0.01/0.01 16:53:55
 RSCS local ID SERVICE ..
application ID RSCS
```

Production zone RSCS Commands

```
sm rscs q links name service show parm
Ready; T=0.01/0.01 15:43:17
 Parm Text
 TCPID=TCPIP  HOST=192.168.150.1
 1 link found

sm rscs q links name service
Ready; T=0.01/0.01 15:43:23
 Link                              Line
 Name     Status     Type          Addr LU Name  Logmode  Queueing
 service connect    TCPNJE   0000 ...      ...      priority
 1 link found
```

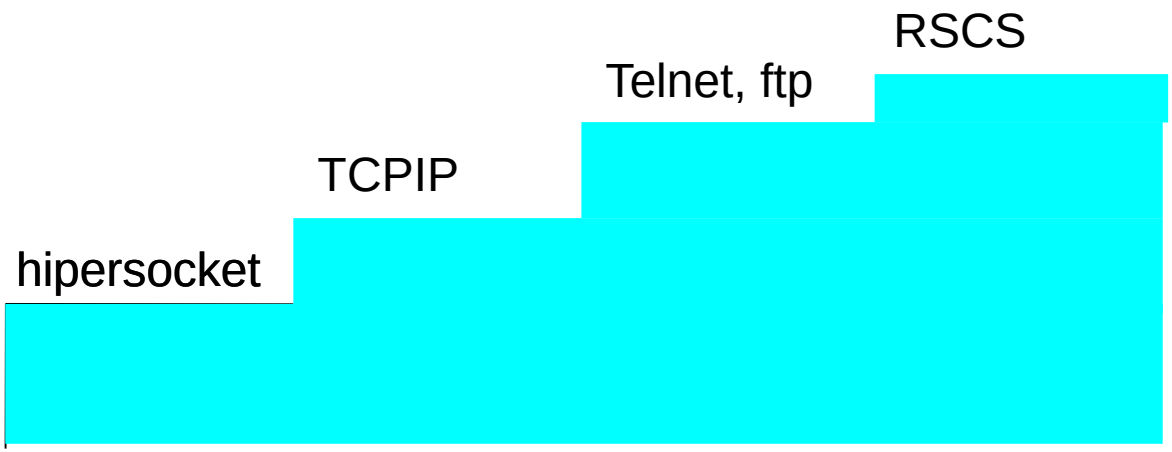


```
sm rscs q sy local
Ready; T=0.01/0.01 16:53:55
 RSCS local ID PRODTION ..
application ID RSCS
```

Remote DIRMAINT

Now that all this wonderful infrastructure is in place greatness is yours o masterful DIRMAINT!! The service can manage DIRMAINT on the other LPARs remotely based on the stacked services plus the information in the SYSTEM NETID!

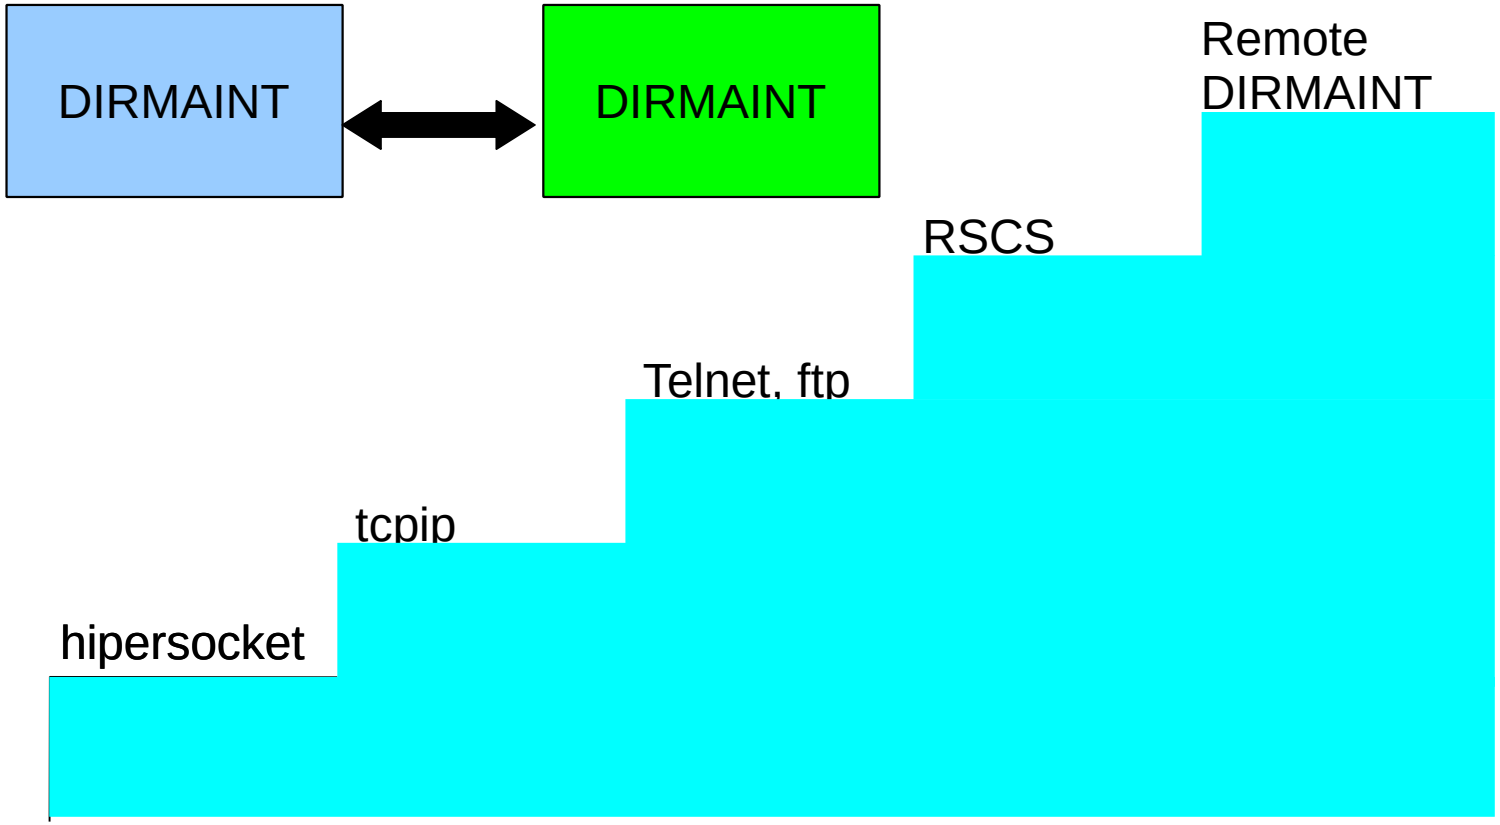

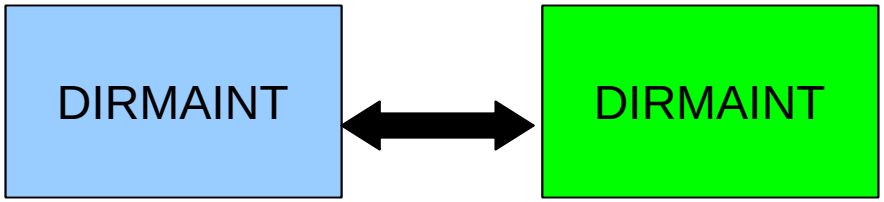


```
dirm to prodtion user withpass
DVHXMT1191I Your USER request has been sent for processing.
Ready; T=0.01/0.02 15:35:00
 From PRODTION(DIRMAINT): DVHREQ2288I Your USER request for MAINT at * has
 From PRODTION(DIRMAINT): DVHREQ2288I been accepted.
 From PRODTION(DIRMAINT): DVHREQ2289I Your USER request for MAINT at * has
 From PRODTION(DIRMAINT): DVHREQ2289I completed; with RC = 0.
RDR FILE 0429 SENT FROM RSCS     PUN WAS 5510 RECS 4949 CPY  001 A NOHOLD NOKEEP
 DMTAXM104I File (1307) spooled to MAINT -- origin PRODTION(DIRMAINT) 07/24/11
 15:35:01 EDT
```

USER WITHPASS from the production LPAR sent to the service zone where varied processing is done asked for by the service zone

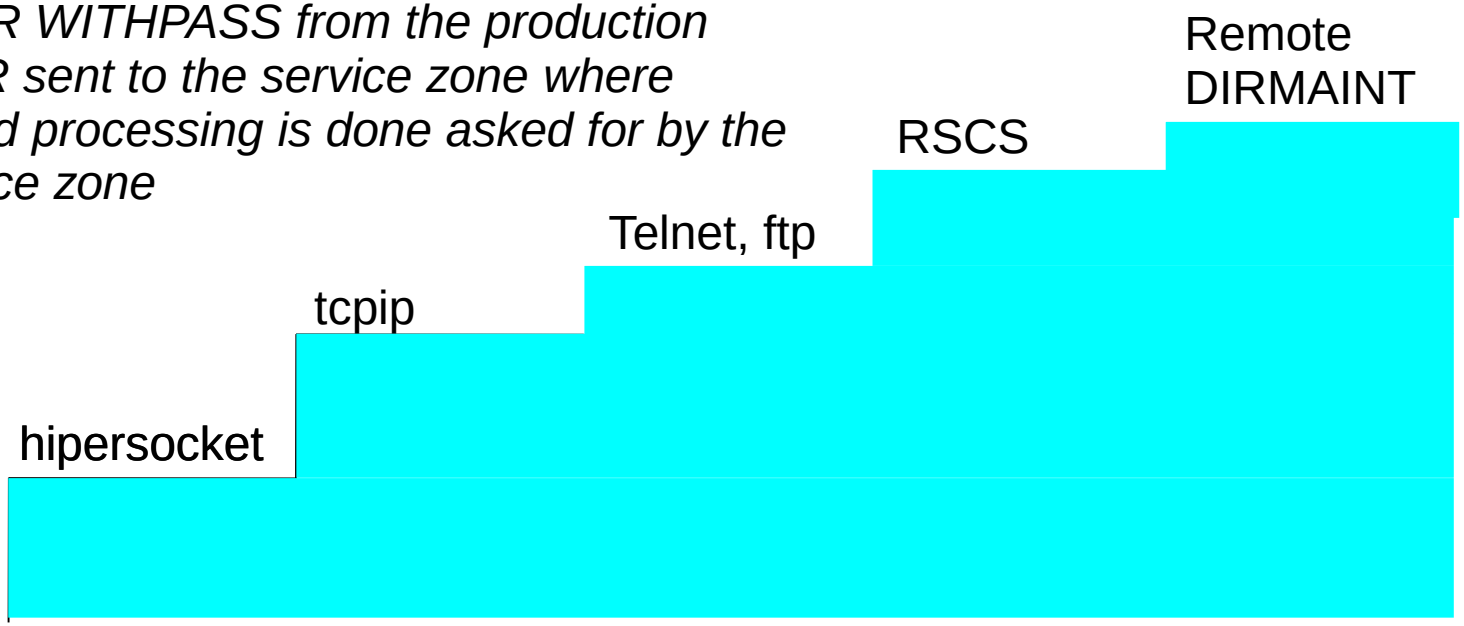

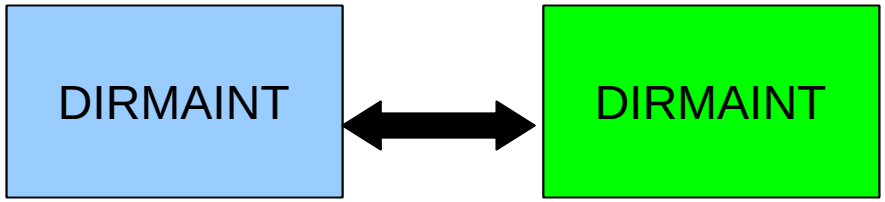

Since DIRMAINT commands for adding, changing and deleting minidisks can be done from the service zone it is important for the directories from the other LPARs to be synchronized on the service zone. Otherwise it would be possible for the service zone to clobber space definitions on the other LPARs.

Directory synchronization code is run nightly.

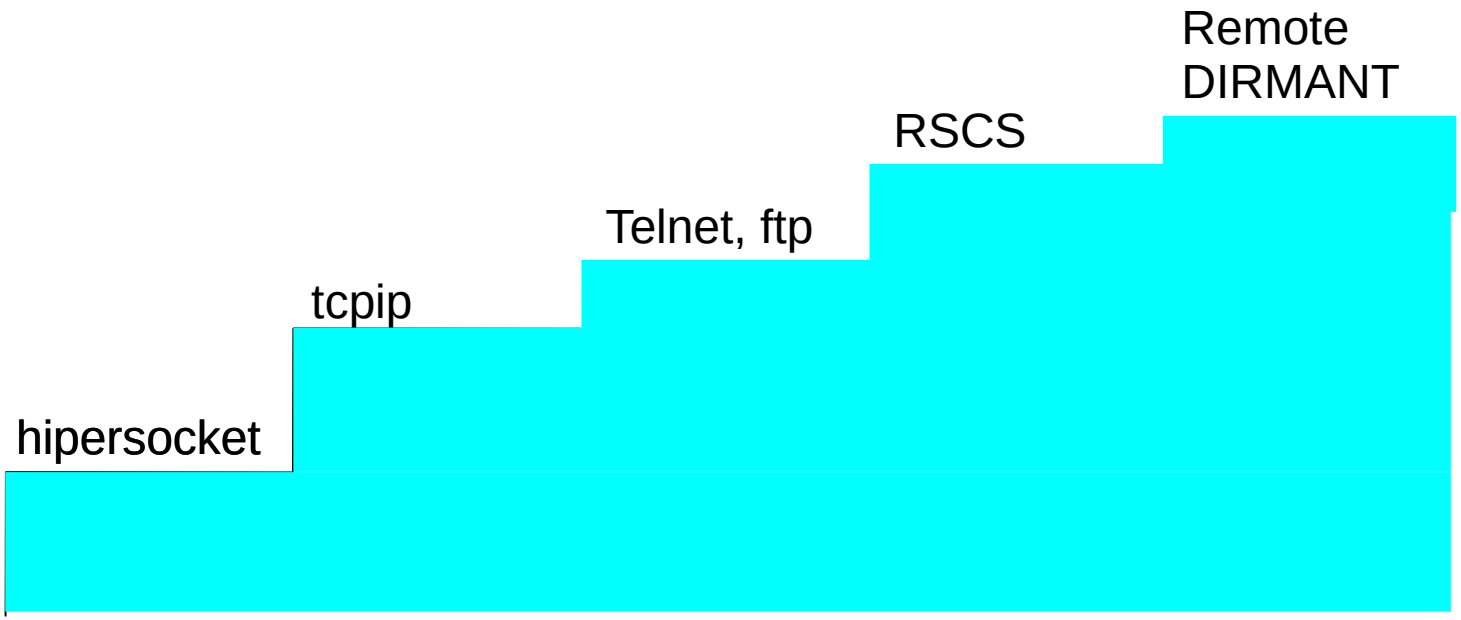

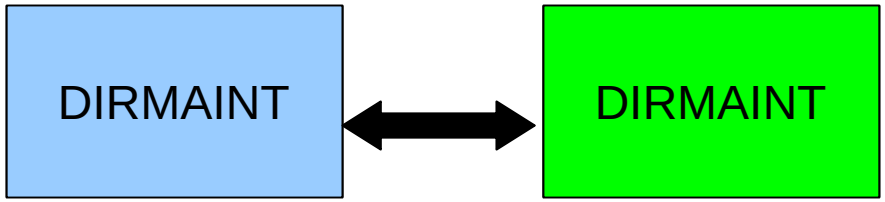

1. 3390 minidisk change made from service zone or directly on the LPAR.
2. If the service zone remains unaware of the change, and adds a minidisk it can result in a destructive overlap – VERY VERY BAD!
3. Directory synchronization code runs nightly in the service zone:
   a. get the full direct from each LPAR (DIRM TO <lpar> USER WITHPASS
   b. process the <lpar> directory forming a list of all minidisks in a userid:

```
PRODTION NOLOG
    MDISK 0001 3390 39 120 PRDVL1 MR MAINT CF1
```

By using the read password as the name of userid and the write password as the minidisk address a handy reference is available. Use DIRMAINT FOR <lpar name> REPLACE. Warning: REPLACE does not consult EXTENT CONTROL – but it's fast! ….
and the USER WITHPASS gives reliable information so it's safe.

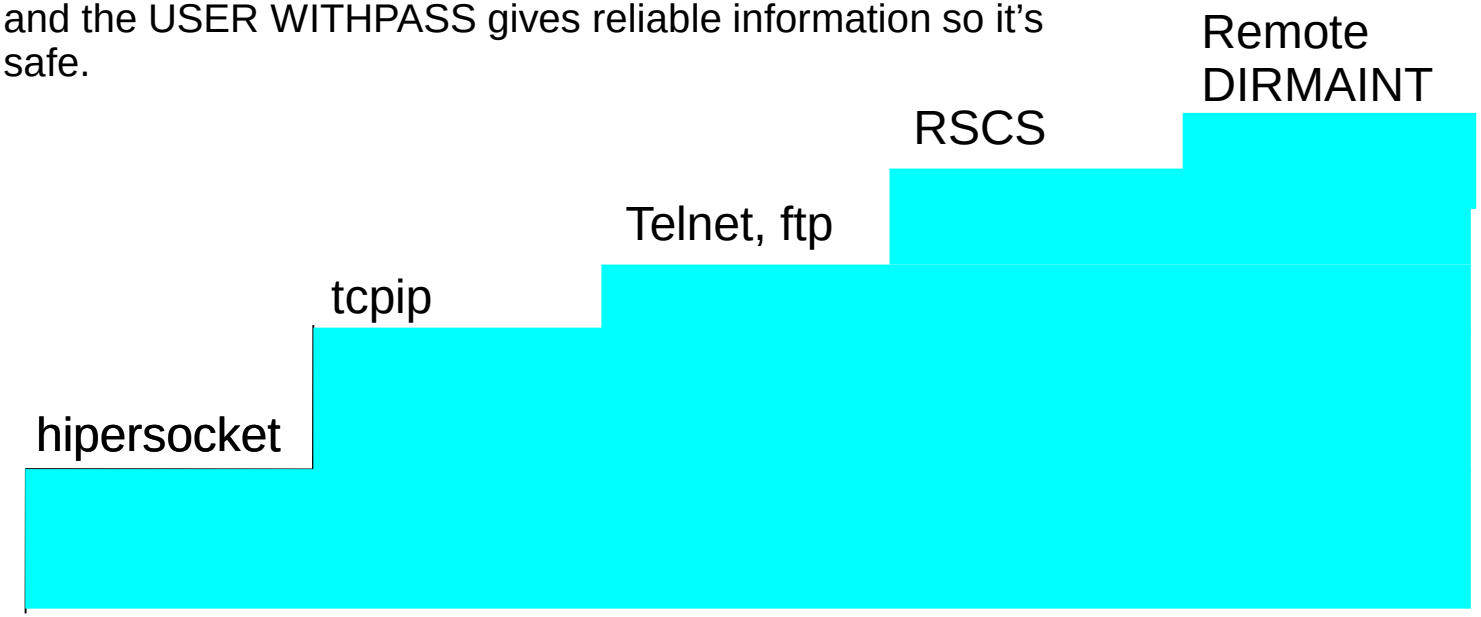

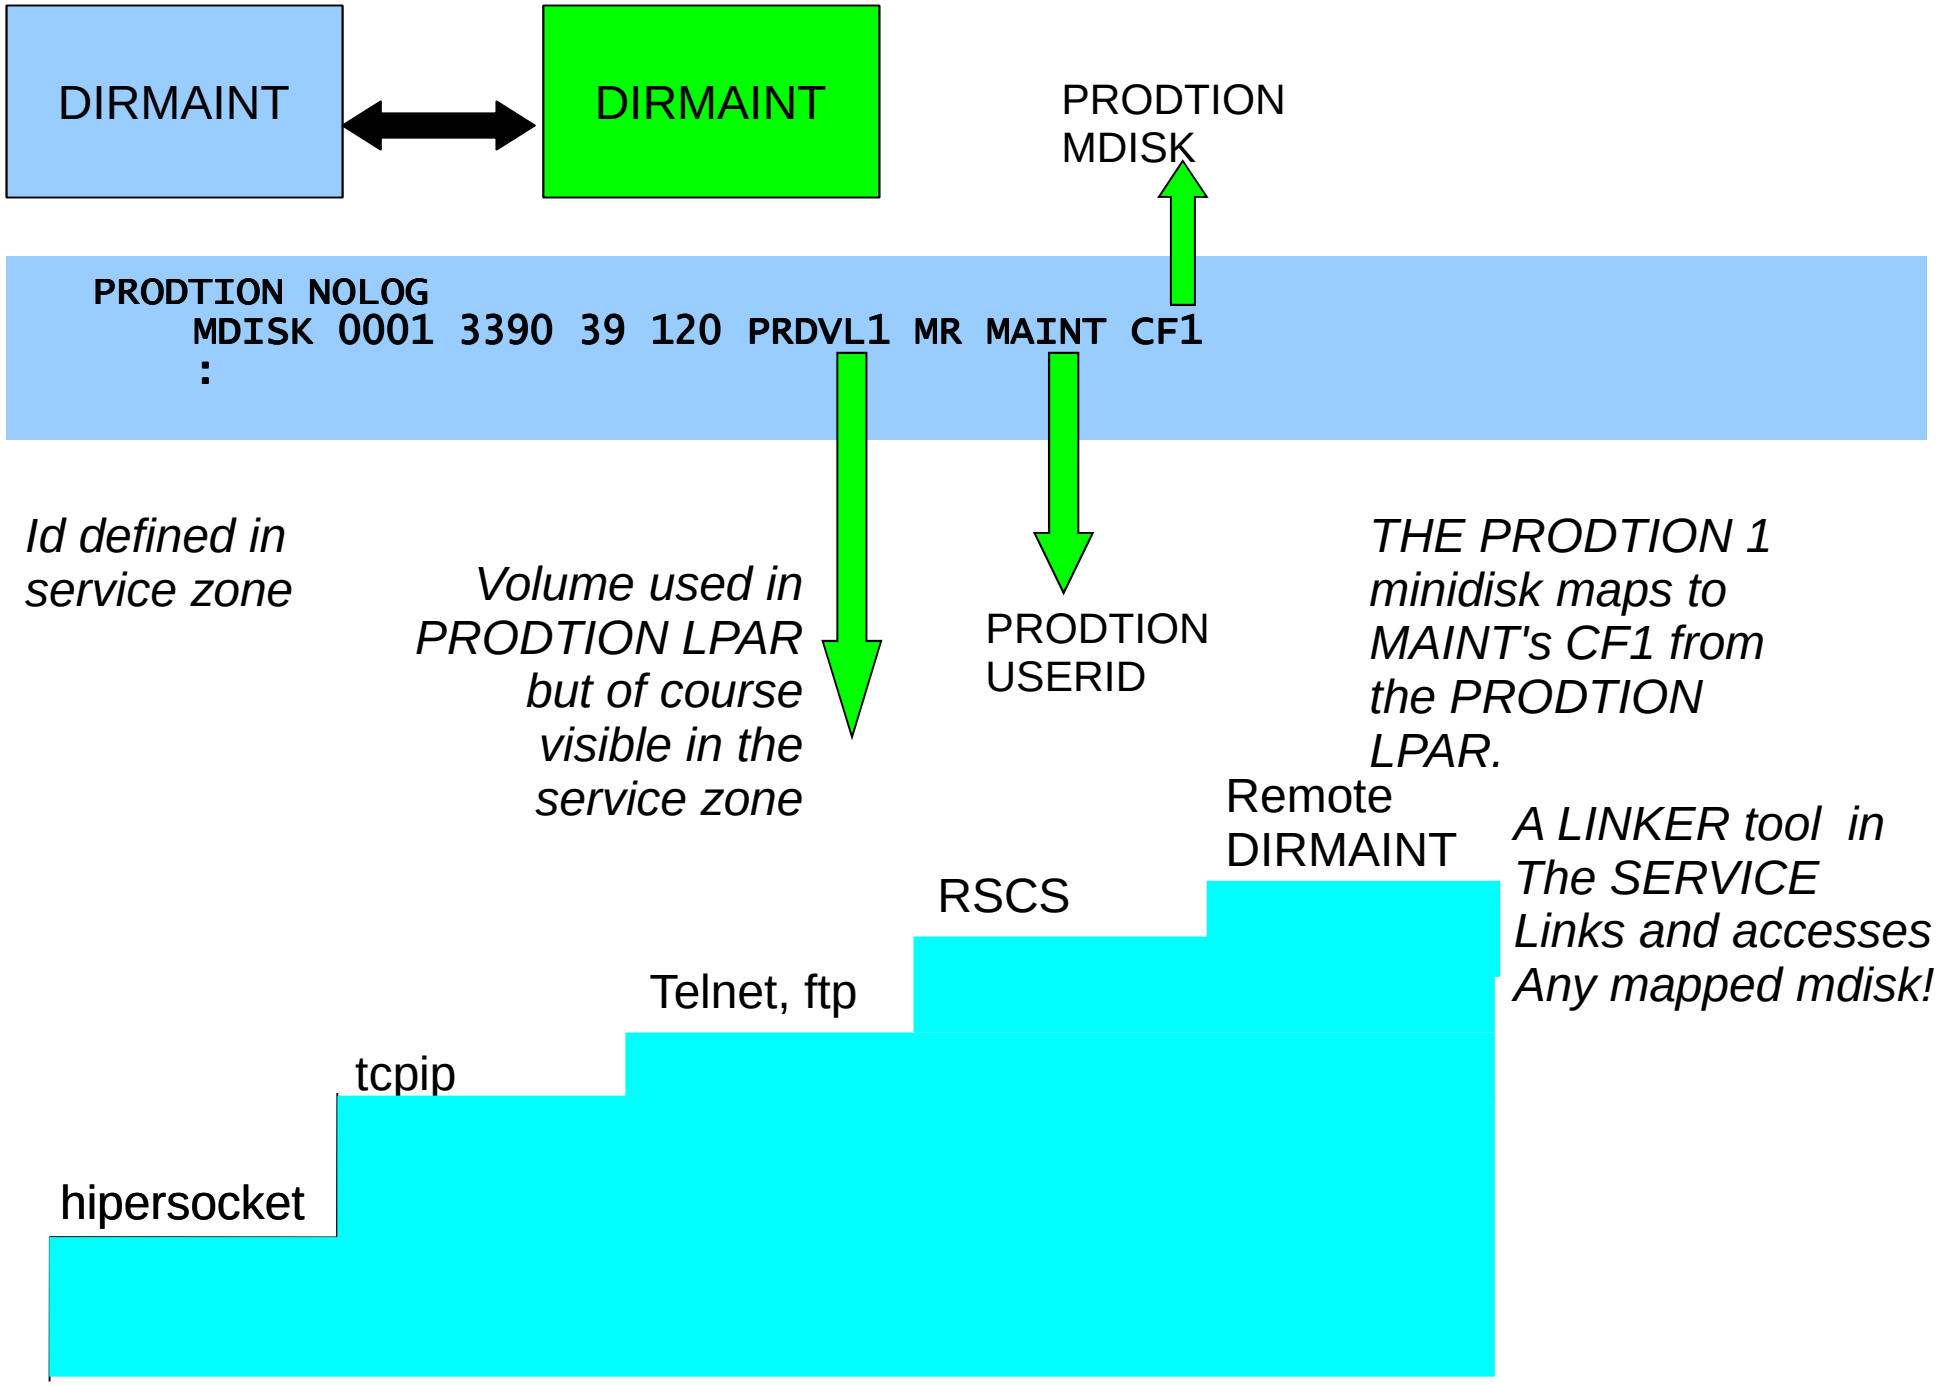

Remote commands

Using RSCS as a carrier pigeon commands may be delivered to remote system including cp, cms and Linux commands.

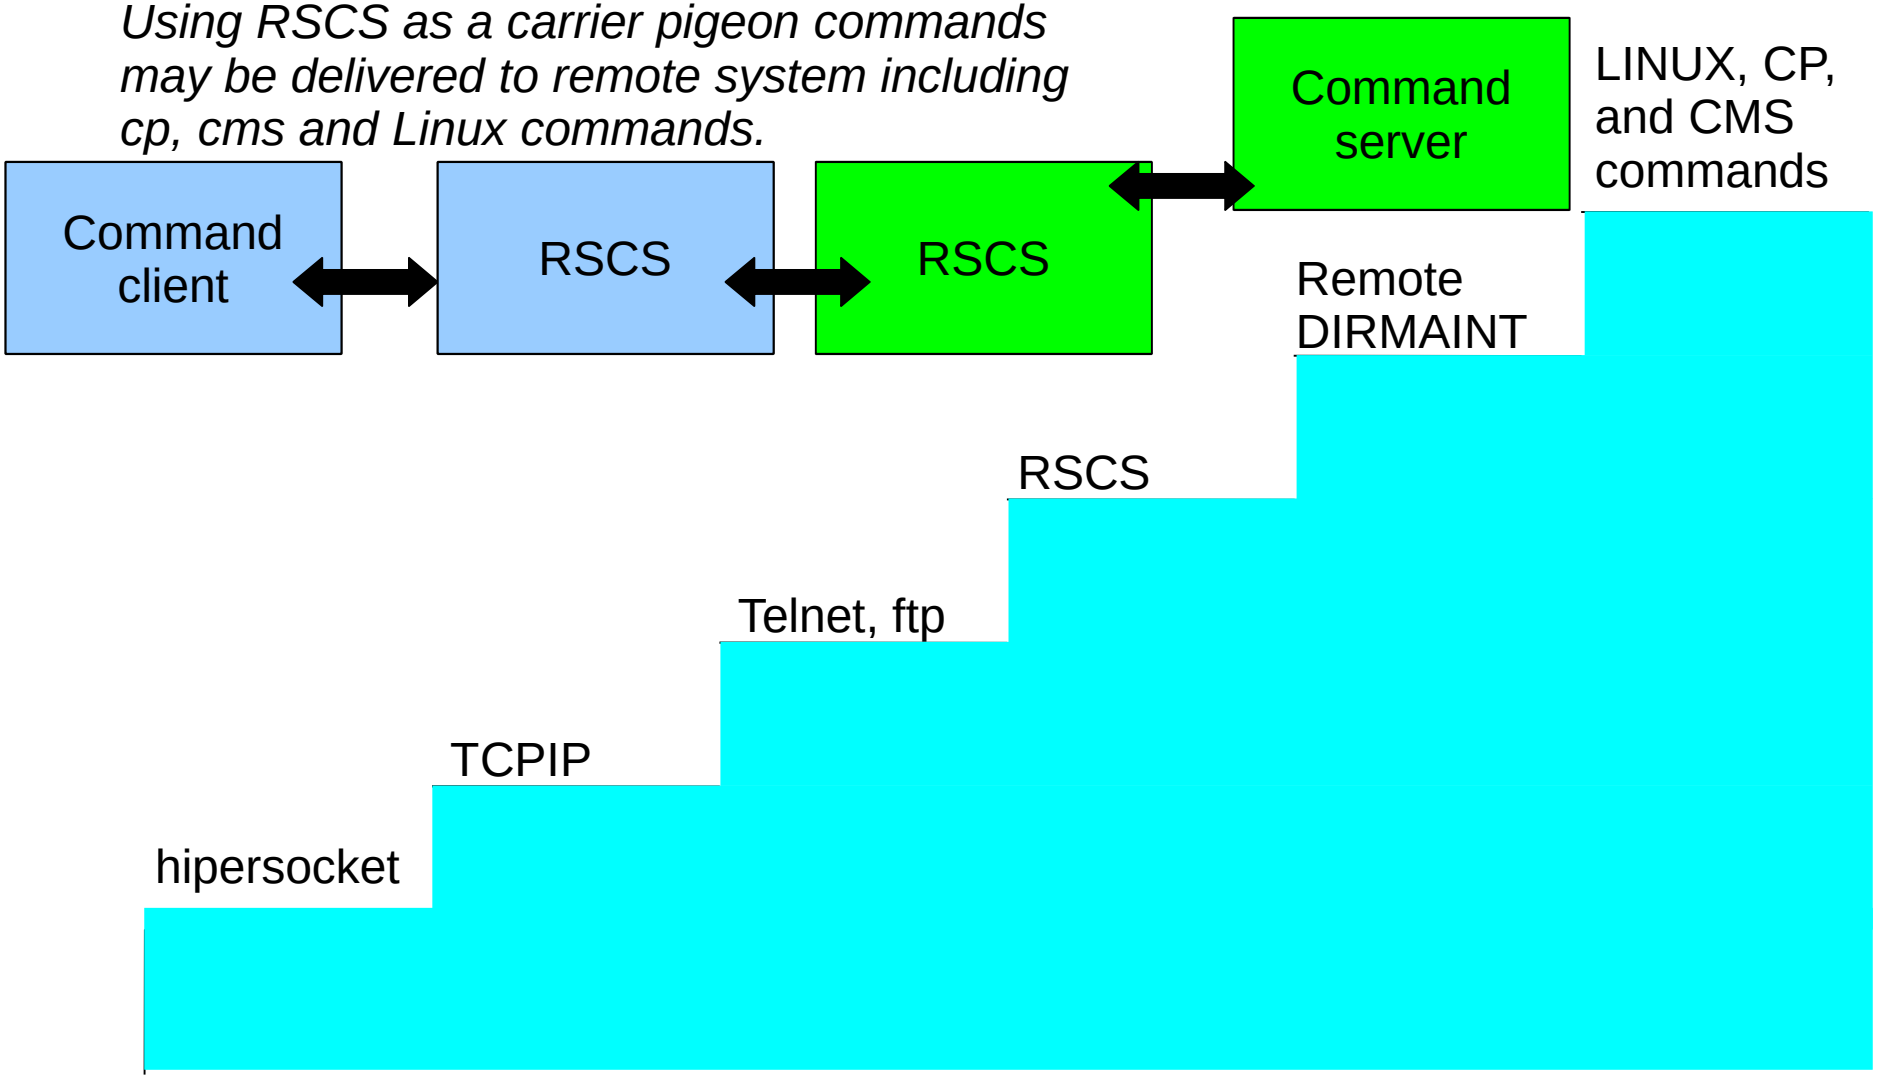

Remote commands

MAINT on service send a Linux command to the production LPAR. The command server on production delivers it to the Linux server.

```
MAINT:
CP MSG DELSERV TO LNXA AT PROD ls
DELSERV: From LNXA at PROD:
DELSERV: /etc /boot /usr /bin /sbin ...
```

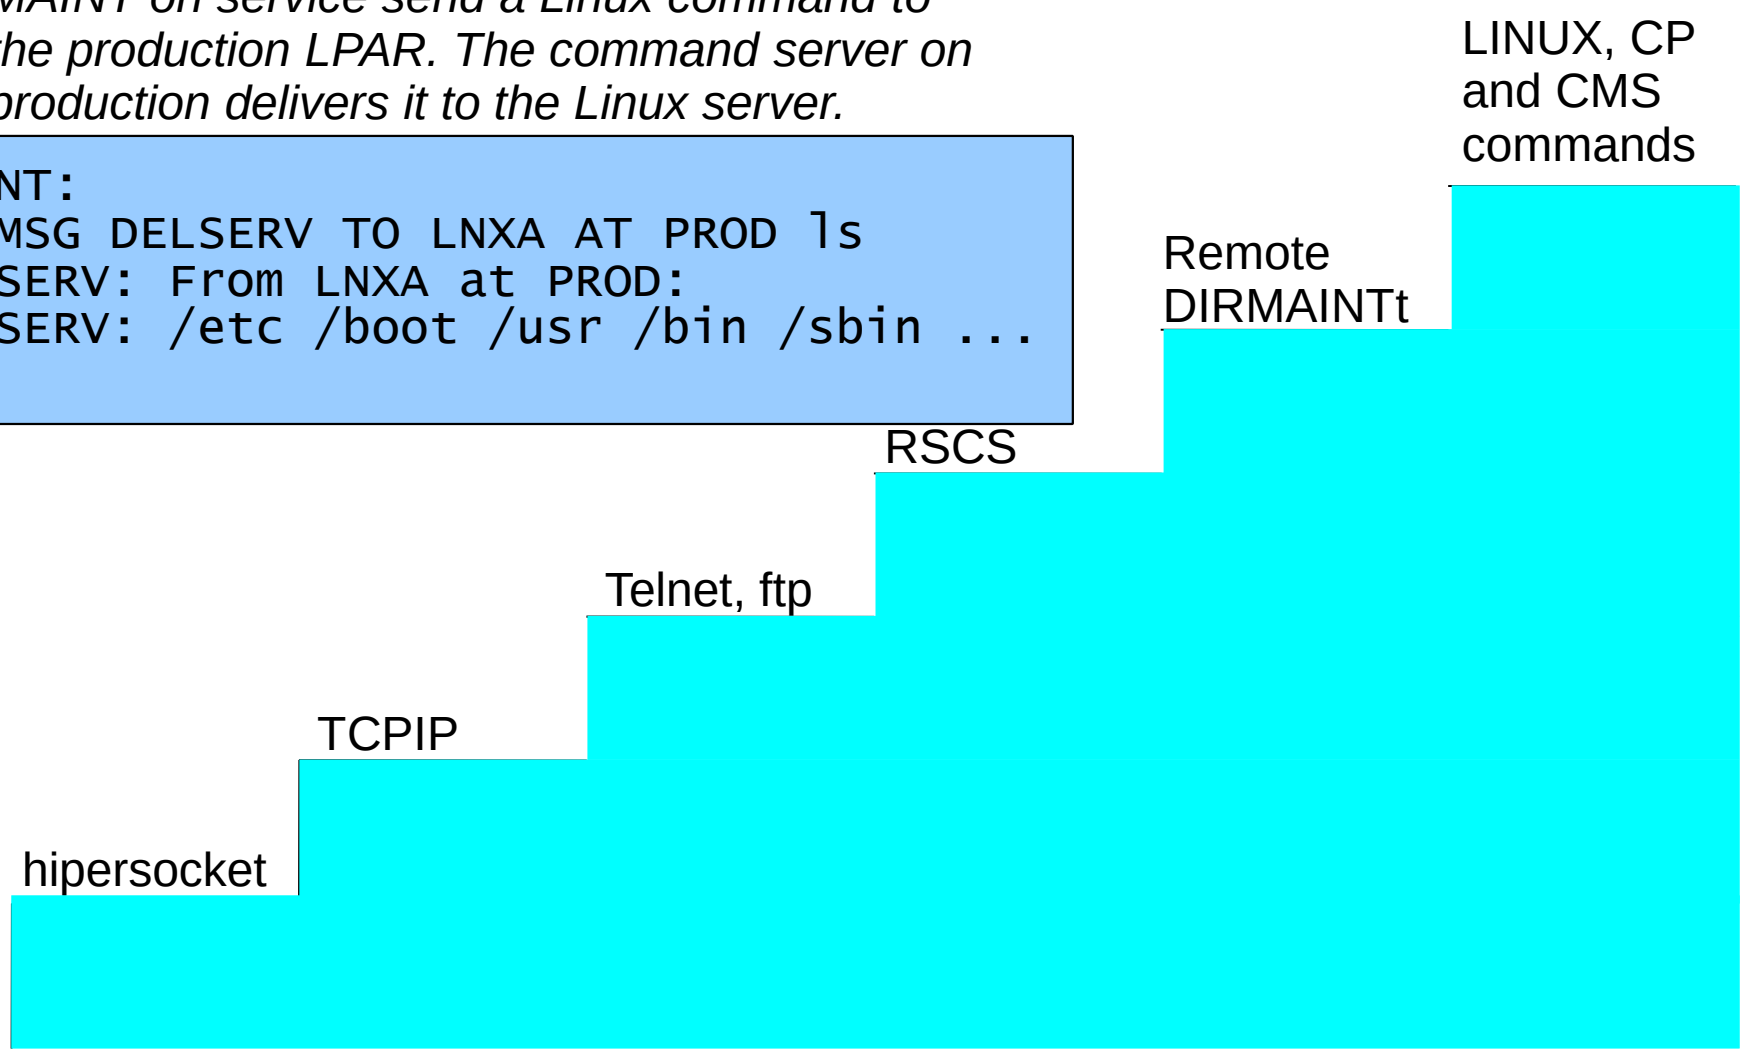

Thoughts on the Service Zone LPAR: ZVM620

- Interesting!
- With SSI Enabled I’m looking forward to PUT2PRODing to other LPARs!
- SERVICE zone will be very valuable for complex implementations.
- Remote DIRMAINT mapping and synchronizing will still be needed.
- Interesting! Can’t wait!

Thoughts on the Service Zone LPAR

- Service zone LPAR is a must have for shops with multiple LPARs.
- Vital for effective systems management.
- Remote control of other LPARs.
- It is *not* a sandbox LPAR!
- Define the 2nd level vm systems in the service zone for staging of RSUs and PTFs, virtual sandbox, etc.
- Build new versions of Linux