



# z/VM TCP/IP Stack Configuration

**Miguel Delapaz**  
**z/VM TCP/IP Development**  
**Sunday, 20 May 2007**

**This presentation is an in depth look at configuration of the z/VM TCP/IP server. Basic and advanced configuration topics will be discussed, with an emphasis on practical examples. Topics such as elementary routing, network hardware, and security are discussed in as much depth as necessary to provide an understanding of how to configure them on the z/VM TCP/IP server. Common configuration errors will also be addressed. While prior experience with z/VM TCP/IP is not necessary for attendees, some basic knowledge of z/VM mini disk structure is assumed.**

**References to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe on any of the intellectual property rights of IBM may be used instead. The evaluation and verification of operation in conjunction with other products, except those expressly designed by IBM, are the responsibility of the user.**

**The following terms are trademarks of IBM Corporation in the United States or other countries or both:**

IBM

z/VM

**Linux is a trademark of Linus Torvalds in the United States, other countries, or both. Other company, product, and service names, which may be denoted by double asterisks (\*\*), may be trademarks or service marks of others.**

**© Copyright International Business Machines Corporation, 1998-2007**

# Agenda

- **Defining a TCP/IP Server**
- **Configuring Your TCP/IP Server**
- **General CMS TCP/IP Client Configuration**
- **Commonly Occurring Errors**

# Defining a TCP/IP Server

## The DTCPARMS File

- Defines a server and various startup parameters
- Formatted as a CMS Names file (:<tag>.<value> pairs)
- Search Order
  - <userid> DTCPARMS
  - <nodeid> DTCPARMS
  - SYSTEM DTCPARMS
  - IBM DTCPARMS
- All but IBM DTCPARMS should reside on the TCP/IP server's 198 disk

## Sample SYSTEM DTCPARMS

```
:nick.TCPIP           :type.SERVER      :class.STACK  
:attach.1f08-1f09, 1e00-1e02  
:vnic.e100 SYSTEM LAN1  
:vctc.800 LINUX1 500,  
      801 LINUX1 501
```

# DTCPARMS Tags

## ■ Required Tags

### ➤:NICK.

- The user ID of the server being configured

### ➤:TYPE.

- For IBM servers the type is “server”

### ➤:CLASS.

- Defines what sort of server is being configured (FTP, TCP/IP Stack, MPRoute)

## DTCPARMS Tags (cont.)

### ■ Other Important Tags

#### ➤:ATTACH.

- Attach a device to your TCP/IP Server

#### ➤:VNIC.

- Define and couple a virtual network card to a guest LAN or VSWITCH

#### ➤:VCTC.

- Define and couple a virtual channel-to-channel device to another user

### ■ Also Typical

#### ➤:OWNER.

- Define the owner of the TCP/IP Server (Default: TCPMAINT)

#### ➤:EXIT.

- Run a user defined exit

## Server Profile Exits

- Global exit called for all servers: TCPRUNXT EXEC
- Server-specific exit called via :Exit. tag
- Input: when called, server class
- Output: DTCPARMS-type tags
- Call 'types': SETUP, BEGIN, END, ADMIN and ERROR
- Example:

```
arg calltype class .
```

```
if calltype = "SETUP" & class = "STACK" then
```

```
begin
```

```
  /* ... find current node ID... */
```

```
  if nodeID = DRnodeID then
```

```
    /* Do some DR setup here */
```

```
end
```

# Configuring Your TCP/IP Server

## PROFILE TCPIP

- **Primary TCP/IP server configuration file**
- **Search order:**
  - **<userid> TCPIP**
  - **<nodeid> TCPIP**
  - **PROFILE TCPIP**
- **Only one is used**
- **Should reside on the TCP/IP server's 198 disk**
- **Sample: PROFILE STCPIP**

## Configuration Statements - Pool Sizes

- Pool size configuration statements **MUST** appear first in the TCP/IP configuration file
- These may not be changed while the stack is running
- Values determined by TCP/IP stack's workload
- 16 different pool statements are available in 3 formats:
  - **xxxBUFFERPOOLSIZE** or **xxxENVELOPEPOOLSIZE**
    - **<statement> <number of buffers allocated> <buffer size>**
  - **xxxPOOLSIZE**
    - **<statement> <number of control blocks allocated>**
  - **FIXEDPAGESTORAGEPOOL**
    - **FIXEDPAGESTORAGEPOOL <initial number> <maximum number>**

# Configuration Statements - Pool Sizes (cont.)

```
netstat pool
VM TCP/IP Netstat Level 520
```

```
TCPIP Free pool status:
```

Object	No. alloc	No. free	Lo-water	Permit size
=====	=====	=====	=====	=====
ACB	5006	4958	4794	500
CCB	755	662	654	75
Dat buf	1200	1165	1123	120
Sm dat buf	5000	4804	4787	500
Tiny dat buf	108	98	96	10
Env	1250	1229	1154	125
Lrg env	75	74	75	7
RCB	51	49	48	5
SCB	2014	1949	1823	201
SKCB	256	216	195	25
TCB	5000	4764	4740	500
UCB	512	501	499	51
Add Xlate	1512	1512	1512	5
NCB	1501	1501	1501	5
IP Route	3015	2837	2832	60
IPv6 Route	3018	3018	3018	60
Segment ACK	100000	99967	99877	10000

FPSP total locked pages: 852, Unused locked pages: 270  
 FPSP allocation threshold: 41313, Low-water mark: 0  
 TCPIP machine size: 256M, Pools: 60573K, Avail: 170956K, Max block: 159400K  
 Ready; T=0.01/0.01 11:13:12

## Configuration Statements - LargeEnvelopePoolSize

- **<buffer size> must be at least as large as that of smaller envelope buffers**
- **Acts as an upper bound on the MTU value**
- **CTC connected hosts should have matching buffer sizes**
- **Example:**
  - **LARGEENVELOPEPOOLSIZ 100 32K**

# Configuration Statements - Privileged Users

## ■ INFORM

- Specifies users who should be informed of major stack events
  - Pool expansion
  - Denial of service attacks
  - Dynamic configuration changes

## ■ OBEY

- Specifies users which can issue privileged stack commands
  - OBEYFILE, NETSTAT OBEY
  - Use of raw sockets
  - Use privileged services on IBM servers (via SMSG)
  - Note: As of 5.3.0, OBEY authority is no longer needed to run TRACERTE

## ■ Examples:

- **INFORM tcpmaint ENDINFORM**
- **OBEY tcpmaint maint miguel mproute ENDOBEY**

## Configuration Statements - AssortedParms

- Various miscellaneous stack settings
- Look at what is available and decide what you need
- Of particular interest
  - EqualCostMultipath
  - IgnoreRedirect (if not running a dynamic router)
  - SourceVipa (if using virtual IP addresses)
- Example:
  - ASSORTEDPARMS
  - equalcostmultipath
  - ignoreredirect
  - ENDASSORTEDPARMS

# Configuration Statements - AUTOLOG, PORT

## ■ AUTOLOG

- Defines which servers to start when the stack comes up
- Stack will restart the server if it is logged off

## ■ PORT

- Gives permission for a server to listen on a port
- Low ports (0-1023) are restricted by default
- Listed ports are monitored unless NOAUTOLOG is specified
- Used to start the Telnet server (assign a port to INTCLIEN)

## ■ Examples:

### ➤ AUTOLOG

```
ftpserve 0  
ENDAUTOLOG
```

### ➤ PORT

```
20 tcp ftpserve noautolog  
21 tcp ftpserve
```

# Configuration Statements - INTERNALCLIENTPARMS

- Specifies settings for the Telnet server (internal client)
  - Timeout values
  - Exit routines
  - Linemode options
  - Dynamic TLS Settings (5.3.0)
- Look through the options and decide what works best for you
- Example:
  - INTERNALCLIENTPARMS  
NOTN3270E  
PORT 23  
SECURECONNECTIONS allowed  
TSLABEL mylabel  
ENDINTERNALCLIENTPARMS

# Configuration Statements - Device and Link

- **Configure network interfaces to the stack**

- **Examples:**

- **Real QDIO Ethernet Device**

- **DEVICE qdio0 OSD 1e00 PRIROUTER**  
**LINK eth0 QDIOETHERNET qdio0 MTU 1500**

- **Virtual QDIO Ethernet Device**

- **DEVICE qdio1 OSD e100 NONROUTER**  
**LINK veth0 QDIOETHERNET qdio1 MTU 32768**

- **LCS Ethernet Device**

- **DEVICE lcs0 LCS 1f08**  
**LINK eth1 ETHERNET 0 lcs0 MTU 1500**

- **Virtual Channel to Channel Device**

- **DEVICE ctc0 CTC 800**  
**LINK vctc0 CTC 1 ctc0 MTU 32760**

# Configuration Statements - HOME

- Configures IP addresses and subnet masks for each link
  - If a subnet mask is specified, a subnet route will be generated for that interface (i.e. You won't need a GATEWAY entry for it)

- Determines which VIPA address is associated with each link

- Examples:

- With VIPA

- HOME
 

7.0.0.1	255.255.255.0	lcs1
8.0.0.1	255.255.255.252	vipa1
9.1.0.3	255.255.255.0	qdio1
8.0.0.2	255.255.255.252	vipa2
9.2.0.2	255.255.255.0	qdio2

- Without VIPA

- HOME
 

10.6.3.159	255.255.255.0	eth0
192.8.12.19	255.255.255.240	veth0
192.8.12.12/27		eth1
192.4.0.1	255.255.255.252	vctc0

# Configuration Statements - GATEWAY and START

- **GATEWAY** – Defines static routes
- **START** – Starts devices
- **Examples:**

- **GATEWAY**

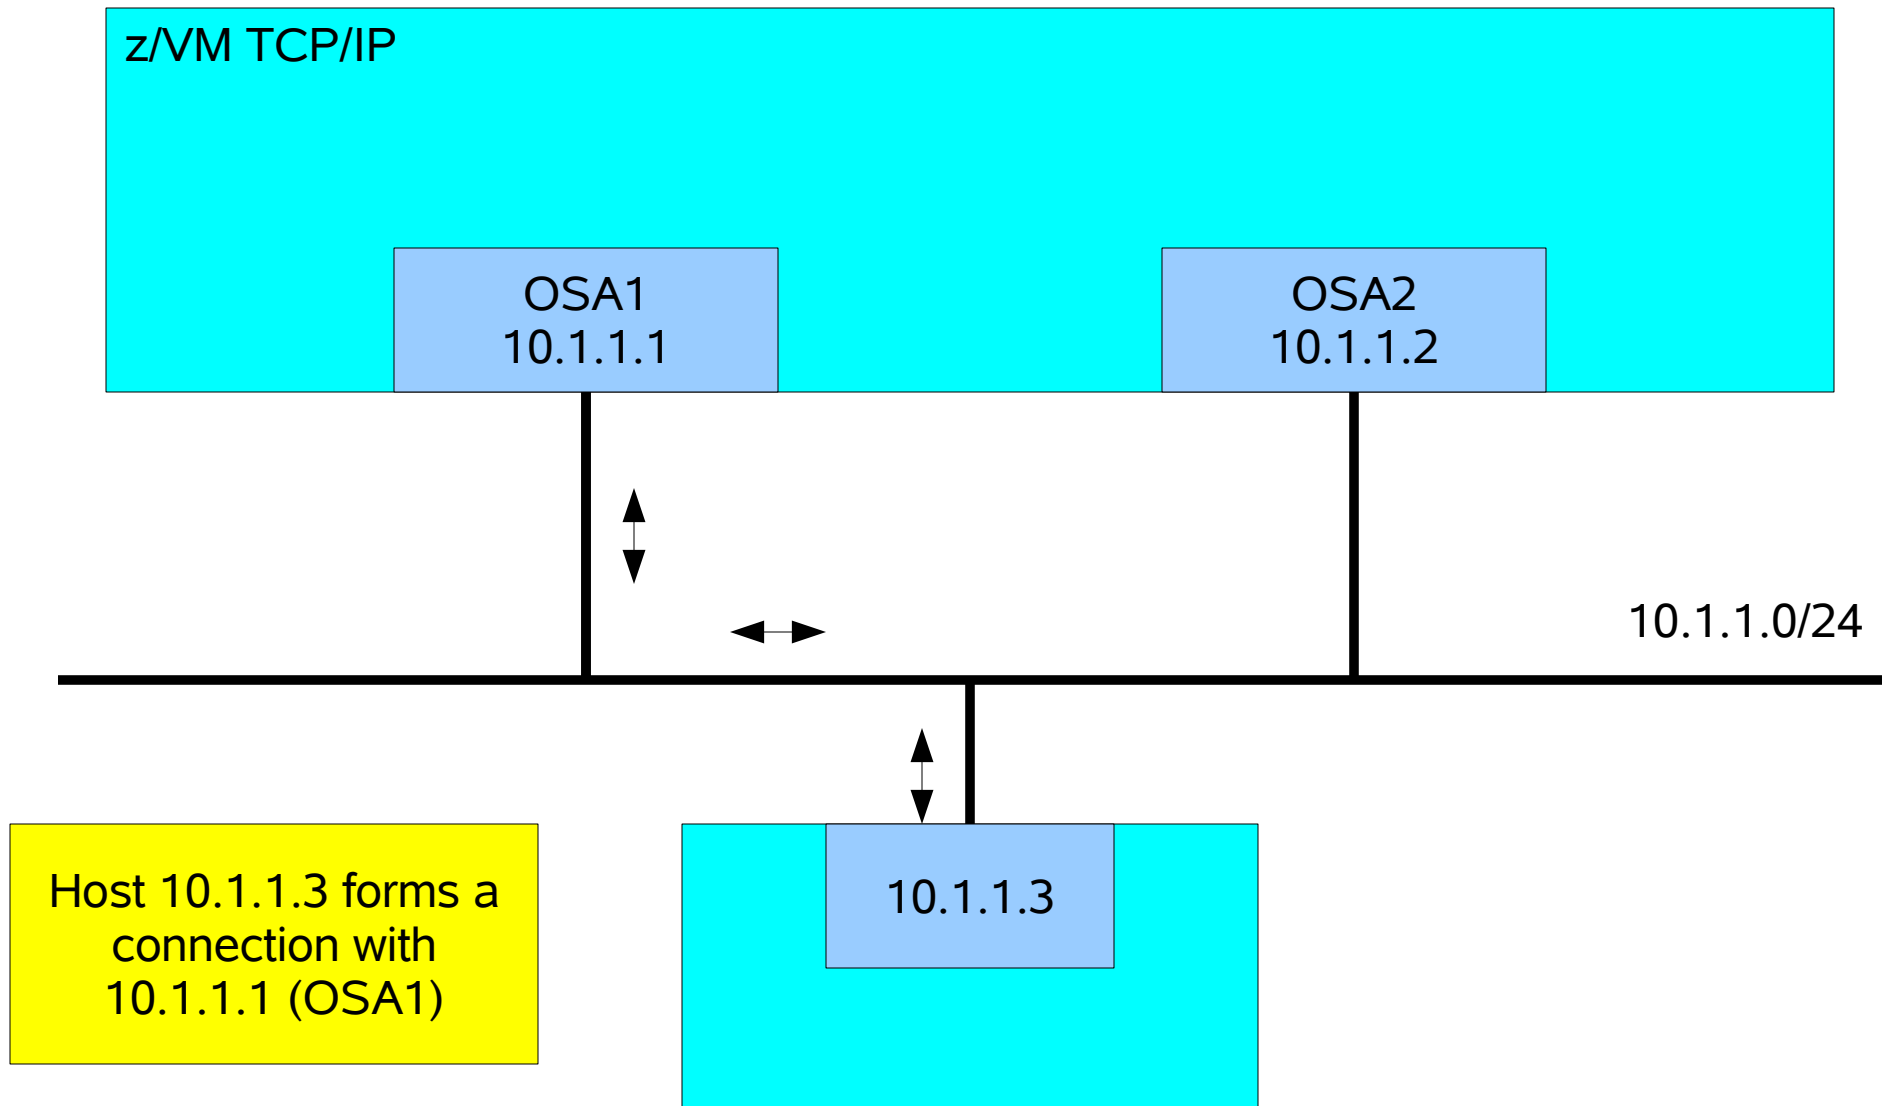
10.6.3.0	255.255.255.0	=	eth0	0
192.8.12.19	255.255.255.240	=	veth0	0
192.8.12.12/27		=	eth1	0
192.4.0.2	HOST	=	ctc0	0
DEFAULTNET		10.6.3.1	eth0	0
DEFAULTNET		192.8.12.1	eth1	0

- **START qdio0**
  - START qdio1**
  - START lcs0**
  - START ctc0**

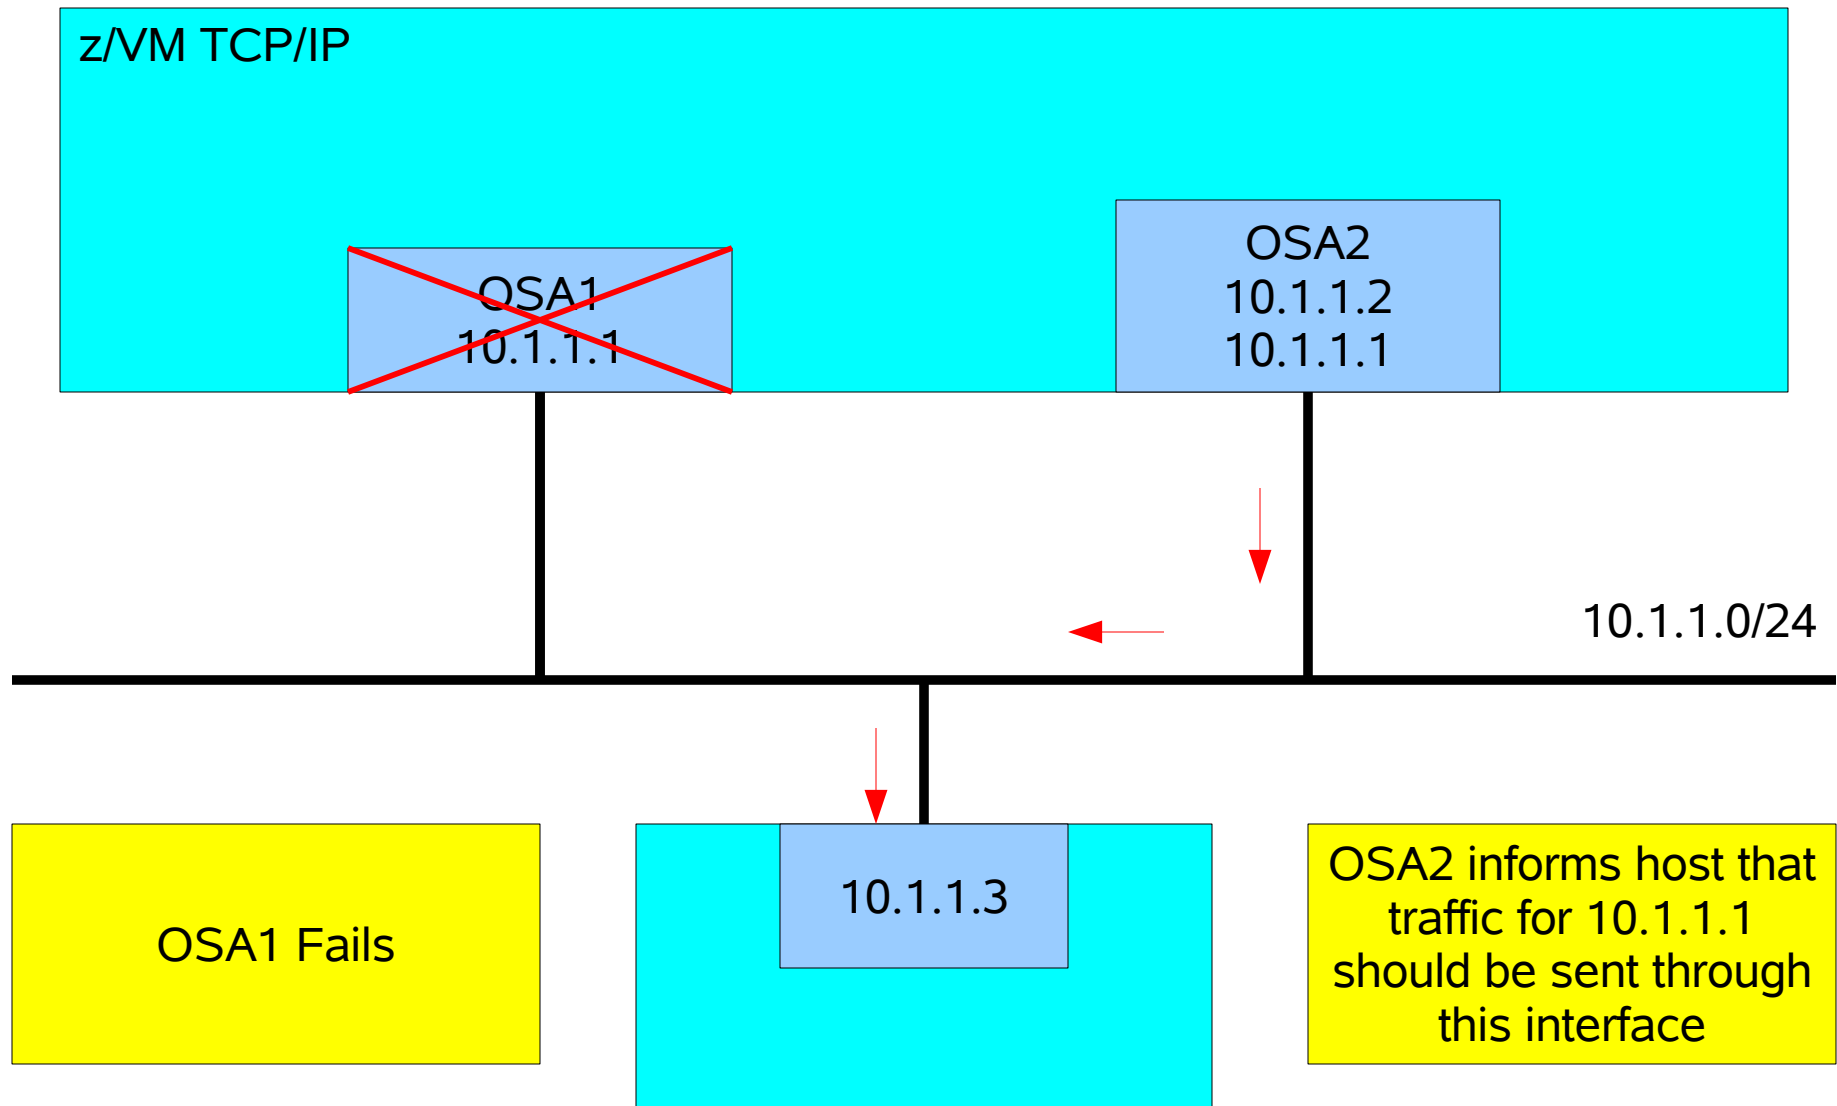
## z/VM TCP/IP 5.3.0 and Interface High Availability

- **IP takeover is supported to minimize the impact of an hardware interface failure**
  - **QDIO ethernet and LCS ethernet devices only**
- **No special parameters or options necessary**
  - **If the TCP/IP stack determines two interfaces are on the same network, IP takeover will be enabled for those interfaces**
  - **For IPv4, determination is based on the IP addresses and subnet masks of the interfaces**
    - **Subnet masks may be defined on the HOME statement, the GATEWAY statement, or in the MPROUTE CONFIG file**

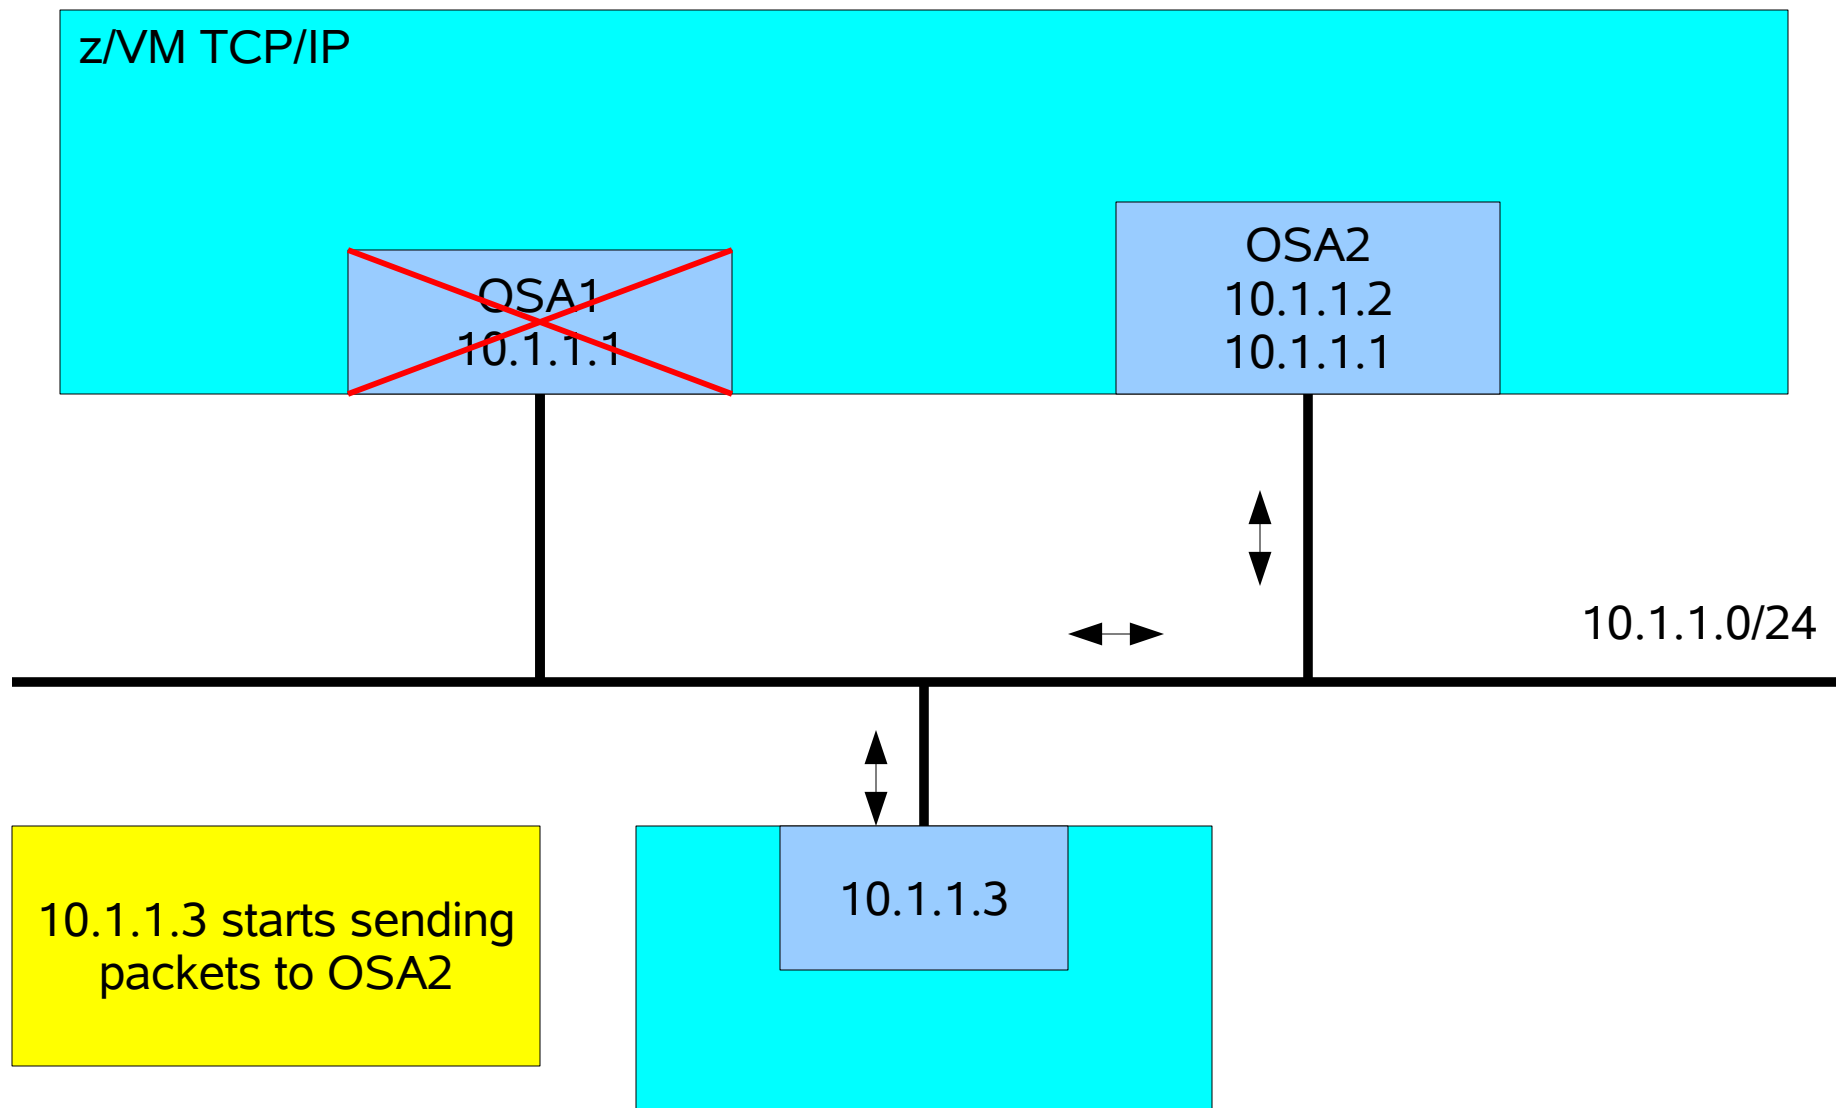
# IP Takeover Details



## IP Takeover Details (cont.)



## IP Takeover Details (cont.)



## Verifying the Configuration

- **Use NETSTAT once the stack is up and running**
  - **NETSTAT GATE**
    - Display the stack's routing table
  - **NETSTAT HOME**
    - Display the stack's HOME list
  - **NETSTAT DEVLINKS**
    - Display the DEVICE and LINK information and status
  - **NETSTAT CONFIG ALL**
    - Display information on ASSORTEDPARMS, INTERNALCLIENTPARMS, OBEY, PERMIT, PORT, RESTRICT, TRACE and TRANSLATE statements.

# Sample PROFILE TCPIP

LARGEENVELOPEPOOLSIZE 100 32K

OBEY tcpmaint maint migueld ENDOBEY

INFORM tcpmaint ENDINFORM

ASSORTEDPARMS

equalcostmultipath

ignoreredirect

ENDASSORTEDPARMS

INTERNALCLIENTPARMS

notn3270E

port 23

ENDINTERNALCLIENTPARMS

PORT

23 tcp INTCLIEN

## Sample PROFILE TCPIP (cont.)

```
DEVICE qdio0 OSD 1e00 PRIROUTER  
LINK eth0 QDIOETHERNET qdio0 MTU 1500
```

```
DEVICE qdio1 OSD e100 NONROUTER  
LINK veth0 QDIOETHERNET qdio1 MTU 32768
```

```
DEVICE lcs0 LCS 1f08  
LINK eth1 ETHERNET 0 lcs0 MTU 1500
```

```
DEVICE ctc0 CTC 800  
LINK vctc0 CTC 1 ctc0 MTU 32760
```

```
HOME  
10.6.3.159      255.255.255.0   eth0  
192.8.12.19    255.255.255.240 veth0  
192.8.12.12    255.255.255.240 eth1  
192.4.0.1      255.255.255.252 vctc0
```

## Sample PROFILE TCPIP (cont.)

```
GATEWAY
192.4.0.1      HOST      192.4.0.2      ctc0      0
DEFAULTNET    10.6.3.1      eth0            0
DEFAULTNET    192.8.12.1    eth1            0
```

```
START qdio0
START qdio1
START lcs0
START ctc0
```

# General CMS Client Configuration

# TCPIP DATA File

- **Information used by both clients and servers**
  - **Resolver Information**
    - host name
    - domain name
    - resolver preferences
    - resolver tracing
    - DNS servers to use
  - **Stack virtual machine name**
  - **E-mail servers**
- **Lives on the TCP/IP Server's 592 disk (client code)**
- **Sample: TCPIP SDATA**

## Sample TCPIP DATA

**TCPIPUSERID**      **TCPIP2**

**NSINTERADDR**    **10.6.3.252**

**NSINTERADDR**    **10.6.3.253**

**HOSTNAME**        **vmhost1**

**DOMAINORIGIN**   **testnet.bigblue.com**

# Commonly Occurring Errors

## Common Hardware Errors

- **Forgot to attach all addresses for a particular adapter**
- **Wrong device type on DEVICE statement**
  - **You have to relate chpid type to device type**
- **Wrong adapter number on LINK statement**
  - **Identifies which port on a multiport device**
  - **'0' and '1' on CTC (or cross-couple instead)**

## Common Routing Errors

- **Incorrect MTU size in routing configuration statements**
  - **Consult hardware documentation**
  - **Talk to your network administrator!**
  - **Everyone on LAN segment should use the same value**
- **Wrong subnet masks or values**
  - **BSD format should help with this**
- **Trying to put guest connected via a point-to-point link in same subnet as VM without PROXYARP**

## The #1 Problem: Incorrect TCPIP DATA file

- **Copy on the incorrect disk (should be on TCPMAINT 592)**
  - **User's A-disk**
  - **Server's A-disk**
  - **TCPMAINT 198**
- **Incorrect NSINTERADDR value**
  - **Ignorance of network changes is no excuse!**
- **TCPIPUSERID pointing to wrong TCP/IP stack virtual machine**
  - **IBM servers will not come up**

## Summary

- **A TCP/IP requires three (3) configuration files**
  - **A DTCPARMS file (<userid>, <nodeid>, or SYSTEM)**
  - **A TCPIP file (<userid>, <nodeid>, or PROFILE)**
  - **A TCPIP DATA file**
- **Communication with networking team is essential to having z/VM happily running on the network**

## Read More About It

- **z/VM Networking Website**

- <http://www.vm.ibm.com/networking/>

- ***TCP/IP Planning and Customization***

- ***TCP/IP Solutions for VM/ESA***

- IBM redbooks at <http://www.redbooks.ibm.com>

- **IETF RFCs**

- <http://www.rfc-editor.org>

- ***Internetworking with TCP/IP*, Comer, Prentice Hall,  
ISBN 0-13-216987-8**

- ***TCP/IP Illustrated, Vol. 1*, Stevens, Addison Wesley,  
ISBN 0-201-63346-9**

## Contact Information

■ By e-mail: **migueld@us.ibm.com**

■ In person: **USA 503.578.3561**

■ Mailing lists:

**A number of listservs relevant to z/VM are available. Information on how to subscribe and view/search archives can be found at the following website:**

**<http://www.vm.ibm.com/techinfo/listserv.html>**

**Of particular interest:**

**[IBMVM@listserv.uark.edu](mailto:IBMVM@listserv.uark.edu)**

**[IBMTCP-L@vm.marist.edu](mailto:IBMTCP-L@vm.marist.edu)**

**[LINUX-390@vm.marist.edu](mailto:LINUX-390@vm.marist.edu)**