



IBM Systems and Technology Group

z/VM TCP/IP Stack Configuration

Miguel Diaz
IBM Corporation
Monday, 23 May 2005

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 TCPIP 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

zSeries

HiperSockets

UNIX is a registered trademark of The Open Group in the United States and other countries.

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, 2005

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**
- **Search Order**
 - <userid> DTCPARMS
 - <nodeid> DTCPARMS
 - SYSTEM DTCPARMS
 - IBM DTCPARMS
- **All but IBM DTCPARMS should reside on TCPMAINT 198**

The DTCPARMS file (cont.)

■ 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**
- **Think of it as a dynamic DTCPARMS file**
- **Example:**
arg calltype class .

```
if calltype = "SETUP" then
  begin
    if class = "NFS" | class = "FTP" then
      return "ESM_ENABLE.YES"
    end
```

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  
:owner.TCPMAINT
```


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 TCPMAINT 198**
- **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 - 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)

■ **Examples:**

- OBEY tcpmaint maint migueld mproute ENDOBEY
- INFORM tcpmaint ENDINFORM

Configuration Statements - AssortedParms

- **Various miscellaneous stack settings**
- **Look at what is available and decide what you need**
- **Suggested**
 - RestrictLowPorts (z/VM 4.3.0 and earlier)
 - VarSubnetting (z/VM 4.3.0 and earlier)
 - IgnoreRedirect (if not running a dynamic router)
- **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 (since z/VM 4.4.0)
- 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 – 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 for each link**
- **Determines which VIPA address is associated with each link**
- **Examples:**
 - With VIPA
 - HOME
 - 7.0.0.1 lcs1
 - 8.0.0.1 vipa1
 - 9.1.0.2 qdio1
 - 9.1.0.3 qdio2
 - 8.0.0.2 vipa2
 - 9.2.0.2 qdio3
 - Without VIPA
 - HOME
 - 10.6.3.159 eth0
 - 192.8.12.19 veth0
 - 192.8.12.12 eth1
 - 192.4.0.1 vctc0

Configuration Statements – GATEWAY and START

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

- GATEWAY

10	=	eth0	0	0.255.255.0	0.6.3.0
192.8.12	=	veth0	0	0.0.0.240	0.0.0.16
192.8.12	=	eth1	0	0.0.0.240	0.0.0.0
192.4.0.2	=	ctc0	0	HOST	
DEFAULTNET	10.6.3.1	eth0	0	0	
DEFAULTNET	192.8.12.1	eth1	0	0	

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

Configuration Statements - Miscellaneous

■ **VSWITCH CONTROLLER**

- Specifies the stack is available to control CP-defined virtual switches.
- Requires IUCV *VSWITCH in user directory entry to function

■ **INTERNALCLIENTPARMS**

- Specifies settings for the Telnet server (internal client)
- Look through the options and decide what works best for you

■ **Examples:**

- VSWITCH CONTROLLER ON
- INTERNALCLIENTPARMS
notn3270E
port 23
ENDINTERNALCLIENTPARMS

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 eth0  
192.8.12.19 veth0  
192.8.12.12 eth1  
192.4.0.1 vctc0
```

Sample PROFILE TCPIP (cont.)

GATEWAY

```
10          =          eth0      0      0.255.255.0      0.6.3.0
192.8.12    =          veth0     0      0.0.0.240        0.0.0.16
192.8.12    =          eth1      0      0.0.0.240        0.0.0.0
192.4.0.1   192.4.0.2 ctc0    0      HOST
DEFAULTNET 10.6.3.1   eth0    0      0
DEFAULTNET 192.8.12.1 eth1    0      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 TCPMAINT's 592 disk (client code)**
- **Sample: TCPIP SDATA**

Sample TCPIP DATA

TCPIPUSERID	TCPIP
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
 - Everyone on LAN segment should use the same value
- **Wrong subnet masks or values**
- **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

- ***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 607.429.3504**

■ **Mailing lists:**

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

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

Of particular interest:

VMESA-L@listserv.uark.edu

IBMTCP-L@vm.marist.edu

LINUX-390@vm.marist.edu