



---

# **z/OS Communications Server TCP/IP: Hints & Tips**

© Copyright International Business Machines Corporation 2004,2005. All rights reserved.

---

# SOURCEVIPA

Copyright International Business Machines Corporation 2004,2005. All rights reserved.

---



# When is a source VIPA address being used?

A source VIPA address is used when the following conditions are met:

- The SOURCEVIPA option has been enabled in the IPCONFIG statement

**AND**

- An outbound connection or UDP association is being established from z/OS

**AND**

- The application has not bound the local socket to a specific interface IP address before establishing an outbound connection or UDP association

**AND**

- The application has not disabled the use of SOURCEVIPA through a setsockopt call

SOURCEVIPA is not being used when outbound data is sent on a connection that was established inbound to z/OS (such as data sent as a response on a TN3270(E) connection that was established from a remote TN3270(E) client to the TN3270(E) server on z/OS).

An FTP outbound data connection is always established from a socket that was bound to the same server IP address as to where the control connection was directed.

# Which source VIPA address is being used?

➤ Two basic rules:

1. If TCPSTACKSOURCEVIPA *is not enabled* on the IPCONFIG statement, then the SOURCEVIPA address to use is selected based on the order of the HOME list

```
10.0.0.1    VIPALINK1
192.168.1.1 REALLINKA ; Will use 10.0.0.1 as SOURCEVIPA
10.0.0.2    VIPALINK2
192.168.2.1 REALLINKB ; Will use 10.0.0.2 as SOURCEVIPA
```

If the connection setup request is sent over REALLINKA, then 10.0.0.1 will be used as source VIPA. If the connection setup request is sent over REALLINKB, then 10.0.0.2 will be used as source VIPA.

2. If TCPSTACKSOURCEVIPA *is enabled* on the IPCONFIG statement, then the IP address specified as TCPSTACKSOURCEVIPA will be used as source VIPA address for all outbound TCP connections, while UDP and RAW associations will continue to use a source VIPA address based on the order of the HOME list

- If the IP address that is specified on the TCPSTACKSOURCEVIPA option is used on multiple LPARs in the Sysplex, then the address must be defined as a **VIPADISTRIBUTE** IP address and the **SYSPLEXPOR**T option must be specified on the VIPADISTRIBUTE statement.

```
VIPADYNAMIC
VIPADefINE 255.255.255.192 201.2.10.11 201.2.10.12
VIPADISTRIBUTE DEFINE SYSPLEXPORtS 201.2.10.11 PORT 9999
                DESTIP 201.3.10.10 201.3.10.11
VIPABACKUP 100 201.2.10.13
ENDVIPADYNAMIC
IPCONFIG .... TCPSTACKSOURCEVIPA 201.2.10.11 ....
```

---

# ECSALIMIT

Copyright International Business Machines Corporation 2004,2005. All rights reserved.

---



# ECSALIMIT usage

- ▶ ECSALIMIT is part of the GLOBALCONF statement in the TCPIP PROFILE data set
  - GLOBALCONF  
ECSALIMIT 200M
    - Specifies the amount of ECSA that TCPIP can use
      - ✓ 10240K and 2 096 128K inclusive
      - ✓ 10M and 2047M inclusive
      - ✓ 0 (zero)...unlimited ECSA usage
  - Cannot be set to a value that is 80%, or greater, of the current system usage of ECSA
  - Can be adjusted with OBEYFILE
  - Does not include ECSA used by CSM (Communications Storage Manager)

# ECSALIMIT Messages

---

- ▶ EZZ4360I jobname ECSA CONSTRAINED
  - Issued when ECSA usage reaches 80% or higher of ECSALIMIT
  - Relieved when ECSA usages reaches 75% or lower of ECSALIMIT
  
- ▶ EZZ4361I jobname ECSA CRITICAL
  - Issued when ECSA usage reaches 90% or higher of ECSALIMIT
  - Relieved when ECSA usages reaches 85% or lower of ECSALIMIT
  
- ▶ EZZ4362I jobname ECSA EXHAUSTED
  - Issued when ECSA usage reaches 98% or higher of ECSALIMIT
  
- ✓ Issue DISPLAY TCPIP,,STOR command to determine current storage usage and limits
  
- ✓ Save system log and request a dump of TCPIP for problem determination

# ECSA Storage Notes

- ✓ Use a system monitoring tool (such as RMF) to monitor ECSA usage
- ✓ Collect regular reports of TCPIP usage of ECSA using the DISPLAY TCPIP,[stackname],STOR command. If using multiple stacks, this needs to be issued against each one
- ✓ Note the allocations used during peak periods. Calculate the amount of ECSA used for other components by subtracting the CSM and TCPIP amounts from the recorded total ECSA for the same time frame.
- ✓ Update the IEASYSxx specification of ECSA size to accommodate the allocations if all of these uses were to have a peak at the same time (plus a growth factor).
- ✓ Specify reasonable limits for CSM and TCPIP ECSA usage based on the above analysis. If some unanticipated event were to cause either one of these to limit communication processing, that would be preferable to having them exhaust ECSA (potentially causing a full system outage).
- ✓ Regularly evaluate these figures, especially after system upgrades or workload increases.



---

# Replaceable Static Routes

Copyright International Business Machines Corporation 2004,2005. All rights reserved.

---



# Static routes prior to z/OS V1R2

---

- ▶ Static routes are defined in GATEWAY or BEGINROUTES/ENDROUTES blocks in the PROFILE
- ▶ A static route always take priority
- ▶ Static routes can never be replaced with dynamic routes
- ▶ Many customers wanted flexibility in defining static routes in z/OS

# Replaceable Static Routes

---

- ▶ Used only if there is no dynamic route is known to destination
  - Can be replaced by dynamic route
  - Used as "last resort", backup route
- ▶ More suitable for use along side OMPROUTE
- ▶ TCPIP "remembers" replaceable static routes
  - Useful when a dynamic route is down or lost
  - Replaceable static route can "step in" to take over when dynamic route has failed
  - Useful if you have older technology that you only want to use in emergency situations
- ▶ You can define multiple replaceable static routes
  - Multipath is supported

# Replaceable Static Routes - example

- Replaceable Static Routes can only be coded with BEGINROUTES/ENDROUTES block

## BEGINROUTES

```
;ROUTE DEST/MASK    FIRST HOP  LINK           MTU           OPTS
;-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
ROUTE 130.200.1.1/32    =          CTC3TO1       MTU 4096
ROUTE 130.200.1.5/32  130.200.1.1 CTC3TO1       MTU 4096
ROUTE 130.200.0.0/16  130.200.1.1 CTC3TO1       MTU 2000      REPL
ROUTE 130.201.0.0/16  130.200.1.1 CTC3TO1       MTU 2000      NOREPL
ROUTE 130.202.0.0/16  130.200.1.1 CTC3TO1       MTU 2000      NOREPL
ROUTE 9.67.111.0/24   =          EZASAMEMVS    MTU 5096
ROUTE 9.67.101.4/32   =          CTC3TO4       MTU 4096      REPL
ROUTE 9.67.102.7/32   =          CTC3TO7       MTU 4096
ROUTE 5.5.5.5/32      9.67.111.3  EZASAMEMVS    MTU 4096
ROUTE 197.1.1.99/32   9.67.111.3  EZASAMEMVS    MTU 4096      REPL
ROUTE 197.1.1.99/32   9.67.101.4  CTC3TO4       MTU 4096
ROUTE 197.1.1.99/32   9.67.102.7  CTC3TO7       MTU 4096      REPL
ROUTE 197.11.3.99/32  130.200.1.1 CTC3TO1       MTU 4096
ROUTE 197.11.3.99/32  9.67.102.7  CTC3TO7       MTU 4096      REPL
ROUTE 197.11.3.99/32  9.67.111.3  EZASAMEMVS    MTU 4096
ROUTE 4.4.4.4/32      9.67.101.4  CTC3TO4       MTU 4096      REPL
ENDROUTES
```

# Replaceable Static Routes - NETSTAT flags

---

- ▶ To assist the user in determining where routes in the IP routing table come from, new flags have been added to the netstat route output.
  - S -- a nonreplaceable static route
  - Z -- a replaceable static route
  - O -- a route learned via OSPF
  - R -- a route learned via RIP
  - C -- a route learned by other means
    - direct connection, OMPROUTE defined default

# Replaceable Static Routes - NETSTAT RSTAT

---

- ▶ To see the Replaceable Static Routing table NETSTAT has been enhanced with the RSTAT parameter
  - *omvs: netstat -r RSTAT*
  - *mvs console: d tcpip,,n,route,rstat*
  - *tso command line: NETSTAT ROUTE RSTAT*

# Replaceable Static Routes - NETSTAT RSTAT

Using the previous example and issuing NETSTAT from the MVS console... `D,TCPIP,,N,ROUTE,RSTAT`

15.58.57 EZZ2500I NETSTAT CS V1R4 TCPCS3

DESTINATION	GATEWAY	INTERFACE
4.4.4.4	9.67.101.4	CTC3TO4
9.67.101.4	0.0.0.0	CTC3TO4
130.200.0.0	130.200.1.1	CTC3TO1
197.1.1.99	9.67.111.3	EZASAMEMVS
197.1.1.99	9.67.102.7	CTC3TO7

5 OF 5 RECORDS DISPLAYED

# Replaceable Static Routes - summary

---

- ▶ Replaceable Static routes can be thought of as "last resort" or "worst case" routes
  - Any dynamic route learned by OMPROUTE is considered better than Replaceable Static Route
- ▶ Nonreplaceable static routes can be thought of as "must use" routes
  - No dynamic routes will be learned to a destination
  - Prior release(s)...all static routes were Nonreplaceable



---

# Resolver Trace Output

Copyright International Business Machines Corporation 2004,2005. All rights reserved.

---



# Resolver Trace Output

---

- ▶ Ways to activate Resolver trace
  - z/OS UNIX **RESOLVER\_TRACE** environment variable
  - SYSTCPT DD allocation
    - Both of these options will dynamically activate Resolver tracing
  - Specifying the TCPIP.DATA statement TRACE RESOLVER or OPTIONS DEBUG
  - Setting the debug option (resDebug) in an application's \_res structure

# Resolver Trace Output

- ▶ The search order used by the resolver to determine if Trace Resolver output is desired
  - Environment variable `RESOLVER_TRACE` (z/OS UNIX only)
  - `SYSTCPT DD` allocation
  - `TRACE RESOLVER` or `OPTIONS DEBUG` statement
    - `stdout` or `SYSPRINT` must be allocated or no trace data is generated.
    - should be the *\*first\** `TCPIP.DATA` statement so that all trace information can be collected.
  - `resDebug` bit set on in the options field of the `_res` structure
    - `stdout` or `SYSPRINT` must be allocated or no trace data is generated.

# Resolver Trace Output

- ▶ Where does the trace output go?
  - A TSO user's terminal screen
  - z/OS UNIX stdout (standard out)
  - JES sysout
  - An MVS Sequential data set
    - The data set must already exist or be allocated as NEW with the following DCB characteristics. An LRECL between 80 and 256 with a RECFM of Fixed Block
    - PDS or PDSE members are not supported
  - An HFS file
    - The file can either be an existing file or it will be dynamically allocated by the resolver when needed
  - NOTE: If the Trace Resolver output is using an MVS data set or HFS file the output will be for the resolver services invoked by the last command or UNIX process. If possible, it is recommended to use SYSOUT=\* or z/OS UNIX standard out so that multiple resolver service invocations (for example, a multitask environment) can be traced.

# Resolver Trace Output

- ▶ Okay, how do I control where it goes?
  - For TSO environment
    - To use the user's terminal:
      - ✓ `alloc dd(systcpt) da(*)`
      - ✓ When directing Trace Resolver output to a TSO terminal the screen size must be defined to be only 80 columns wide. Wider screen sizes will cause trace output to be wrapped and hard to read.
    - To use an existing MVS data set:
      - ✓ `alloc dd(systcpt) da(appl.restrace)`. The userid will be used as the first qualifier for the data set. `alloc dd(systcpt) da('user3.resolver.trace')`
    - To disable the Trace Resolver output:
      - ✓ `free dd(systcpt)`

Continue to next page...

# Resolver Trace Output

- For z/OS UNIX shell environment
  - To use standard out:
    - ✓ `export RESOLVER_TRACE=stdout`. If desired, stdout can be redirected when the z/OS UNIX command is issued.
  - To use a new HFS file or existing MVS data set:
    - ✓ `export RESOLVER_TRACE=/tmp/myjob.resolv.trace` or  
`export RESOLVER_TRACE="//appl.restrace"` or  
`export RESOLVER_TRACE="//user3.resolver.trace"`
    - ✓ The userid will be used as the first qualifier for the data set
  - To disable the Trace Resolver output:
    - ✓ `set -A RESOLVER_TRACE`

Continue to next page...

# Resolver Trace Output

---

- For MVS batch job environment
  - To use the recommended JES SYSOUT
    - ✓ //SYSTCPT DD SYSOUT=\*
    - ✓ //SYSPRINT DD SYSOUT=\*
    - ✓ The SYSPRINT DD must be allocated if the TCPIP.DATA statements TRACE RESOLVER or OPTIONS DEBUG are specified. If the DD does not exist no trace output will be written.

---

# Common Reported Problems

Copyright International Business Machines Corporation 2004,2005. All rights reserved.

---





# SSL/TLS

## ■ Setting FTP with SSL/TLS (Socket Secure Layer/Transport Layer Security)

- ▶ Implicit TLS connections over port 990
  - We recommend that you NOT use port 990
- ▶ AUTHSSL support
  - Replaced with AUTH TLS
  - FTP client does not support AUTH SSL
- ▶ TLS through firewall(s)
  - Only for the brave!

## ■ Please see APAR II13516 for these and other solutions for FTP TLS problems

Other sources of valuable information:

- ▶ Redbook, Communications Server for z/OS V1R2 TCP/IP Implementation Guide Volume 7: Security, SG24-6840-00
- ▶ z/OS V1R2.0 CS: IP Configuration Guide, SC31-8775-01
- ▶ System SSL Programming, SC24-5901-03

# EDC5140I Broken pipe

---

## ■ Any TCP connection has been reset

- ▶ Usually caused when a write() is performed on a reset connection
- ▶ Additional messages may be seen from the application

## ■ To prove:

- ▶ Review some type of network trace
  - ▶ TCPIP CTRACE SYSTCPDA (PKTTRACE)
  - ▶ Sniffer (or similar) trace

# FTP Server Traces

- **FTP server traces are routed to SYSLOGD by default.**
  - ▶ If SYSLOGD is not running FTP server traces go to the MVS system log
  - ▶ If SYSLOGD is misconfigured FTP server traces can end up in MVS system log
  
- **Ensure that SYSLOGD is configured correctly**
  - ▶ Direct debug tracing to a file  
daemon.debug            /tmp/syslogd/server.debug.a
  - ▶ Direct debug tracing to MVS system log  
daemon.debug            /dev/console

**Warning: you don't want to do this!**

# CICS Region hangs, ABENDs0C4, ABENDASRA, etc.

---

- Poor response or hanging CICS region
- ABENDASRA, ABENDAKEB
  - Underlying ABEND0C4, ABENDEC6
- This is an old friend that has taken on a new life with CICS Transaction Server 2.3
  - ▶ Ensure that EZACICAL is linked into the CICS application that is using the TCPIP CICS API.
    - ▶ This includes all transactions that issue
  - ▶ For more information see Technote 1104510 at [www.ibm.com/software/network/commserver/zos/support/](http://www.ibm.com/software/network/commserver/zos/support/)

# Resolver

## ■ Only 1st of multiple defined DNSs queried by Resolver

→ NSINTERADDR

→ NAMESERVER

- ▶ After the Resolver has successfully contacted a name server, it stops without contacting the remaining name servers for that query.
- ▶ A successful contact includes NEGATIVE responses, such as NXDOMAIN (nonexistent domain).
- ▶ Name servers beyond the first in the list are used only if the name server currently being contacted is down or unreachable through the network.
  - \* RESOLVERUDPRETRIES indicates the maximum number of times an attempt is made to reach a given name server if a response is not received within the current time-out interval.
  - \* RESOLVERTIMEOUT is the number of seconds the Resolver will wait for a response from the DNS

# Telnet (TN3270)

- **Telnet connection/session hangs (i.e. Client cannot enter data because keyboard is locked)**
  - ▶ ENSURE That all PTFs mentioned in APAR II13135 are applied.
  - ▶ If problem still exists, gather doc fould in above APAR
    - TCPIP CTRACE with OPTIONS=(TELNET)
    - PKTTRACE filtered on client IP address
    - Rarely - VTAM VIT with OPTIONS=(API,PIU) and VTAM Buffer trace

# TCPIP Hang

## ■ TCPIP hangs. At startup, shutdown or normal processing

- ▶ ENSURE THAT all Hiper PTFs are applied.

5655hal00 R\*\*\* (R140, R150, R160)

- ▶ If problem still exists, gather a console dump of TCPIP and OMVS. Include all dataspace

- DUMP COMM=(Descriptive name)

```
r nn,JOBNAME=(OMVS,tcpip_proc),dspname=('OMVS'.*),  
sdata=(allnuc,psa,csa,lpa,trt,sqa,rng,grsq,sum),end
```

# FTP Client Data Connection

## ■ Problem

Data connection cannot be completed. The server is GuildFTPd

## ■ Cause

The FTP client sends out a SITE command before it sends the PORT command. The GuildFTPd FTP server responds with a 214 'SITE' command not understood message which ends with 0D0A 0D0A. Per RFC 959, there should only be one 0D0A. When the client reads the 214 message, it reads the first 0D0A and the second 0D0A stays in the buffer. When the client sends out the PORT command and checks for a response, it reads in this 0D0A. The client interprets this as a failure, and brings down its listen on the port for the data connection. Any SYNs sent from the server during this time will be reset. The client then tries to recover from this and puts up a listen using the control connection port and sends out the STOR command. The remote server is not able to open the data connection, so it may send back message 425 Can't build data connection.

## ■ Solution

Issue the SENDSite command in the session before the put or get is done. This will tell the client not to send out the SITE command.



# MultiNode Load Balancing Packets

---

## ■ Problem

MNLB packets are dropped by Forwarding Agents

## ■ Cause

PQ97846 (R140 UK00038/R150 UK00039/R160 UK00040) changed the z/OS Service Manager to send MNLB packets with a TTL of 4 instead of 1. The MNLB packets are multicast to the forwarding agents. Since the TTL is greater than 1, the forwarding agents are first checking if they are configured to route multicast traffic. If they are not configured this way, the packets will be dropped.

## ■ Solution

After applying PQ97846, **ip pim dense-mode** must be configured on the Forwarding Agents to allow the MNLB packets to be processed.

# Bind() Failure

---

## ■ Problem

TCPIP application cannot start due to bind() failure

Errno=1115 (EADDRINUSE)

Reason Code=12BA00C4

## ■ Cause

Notice that the first four digits of the reason code are in the range 0000-1999. These are issued from the z/OS UNIX component. The BIND is being rejected because z/OS UNIX is configured to support multiple stacks (CINET) and the requested port number is in the range of ports defined by the INADDRANYPORT and INADDRANYCOUNT values specified on the NETWORK statement in the BPXPRMxx PARMLIB member.

## ■ Solution

Please see next page

# Bind() Failure

## ■ Solution

> Select the range of port numbers specified in the z/OS UNIX configuration so that there is no application (server) that will attempt to explicitly BIND to one of the ports for any of the configured stacks.

> Ensure that the INADDRANYCOUNT value is large enough to handle, at the same time, all of the applications that do not have a stack affinity and will BIND (explicitly or implicitly) without requesting a specific port.

> Configure this range in the TCPIP PROFILE for every stack on the system. If, for example, INADDRANYPORT(50000) INADDRANYCOUNT(999) is specified in BPXPRMxx, then add the following two statements to the profile:

```
PORTRANGE 50000 999 TCP *  
PORTRANGE 50000 999 UDP *
```

# Telnet connection drops

## ■ Problem

Telnet connection drops with the following message:

```
EZZ6034I TELNET CONN connid LU luname CONN DROP ERR 1027
```

## ■ Cause

Telnet (or any application) connections are terminated with a error code of 1027 (ECONNABORTED) typically due to a connection RST packet being received from the network.

## ■ Solution

A RST packet can be sent as a result of one of the following conditions:

- A remote TN3270 client closed its socket which had specified the `SO_LINGER` socket option with a linger time of 0 and there was unread data sitting on the TCP receive queue.
- An intermediate firewall may contain a rule which will drop connections which have not been used for a specified period of time.

# Where can I get helpful hints and tips?

■ [www.ibm.com/software/network/commserver/zos/support/](http://www.ibm.com/software/network/commserver/zos/support/)

- ▶ IBM's new URL for support regarding z/OS  
Comminucations Server
- ▶ Search for Hints and Tips
- ▶ Download fixes
- ▶ Find online manuals
- ▶ And more....

# Demo z/OS CS Support

■ [www.ibm.com/software/network/commserver/zos/support/](http://www.ibm.com/software/network/commserver/zos/support/)

- ▶ I'm having a hard time specifying **TCP/IP server parameters** in **JCL**
- ▶ How do I **modify FTP message catalogs** for local use
- ▶ I get message **iec143i 213-07** for **SEZALOAD** when I try to dynamically alter **USS table**

# HIPER APARs

- PK06135 - UK04942/UK04943/UK04944
  - ▶ STORAGE leak in SP227 K6 when using EE
- PK06394 - UK04414/UK04415/UK04416/UK04417
  - ▶ Telnet server causes CSA overlay
- PK07601 - UK04958/UK04959
  - ▶ TCPIP Storage leak for SCBs
- PQ99420 - UK00809/UK00810/UK00811
  - ▶ Dropped packets when using Sysplex Distributor and Multipath Perpacket

# For More Information....

URL	Content
<a href="http://www.ibm.com/servers/eserver/zseries">http://www.ibm.com/servers/eserver/zseries</a>	IBM eServer zSeries Mainframe Servers
<a href="http://www.ibm.com/servers/eserver/zseries/networking">http://www.ibm.com/servers/eserver/zseries/networking</a>	Networking: IBM zSeries Servers
<a href="http://www.ibm.com/servers/eserver/zseries/networking/technology.html">http://www.ibm.com/servers/eserver/zseries/networking/technology.html</a>	IBM Enterprise Servers: Networking Technologies
<a href="http://www.ibm.com/software/network/commserver">http://www.ibm.com/software/network/commserver</a>	Communications Server product overview
<a href="http://www.ibm.com/software/network/commserver/zos/">http://www.ibm.com/software/network/commserver/zos/</a>	z/OS Communications Server
<a href="http://www.ibm.com/software/network/commserver/z_lin/">http://www.ibm.com/software/network/commserver/z_lin/</a>	Communications Server for Linux on zSeries
<a href="http://www.ibm.com/software/network/ccl">http://www.ibm.com/software/network/ccl</a>	Communication Controller for Linux on zSeries
<a href="http://www.ibm.com/software/network/commserver/library">http://www.ibm.com/software/network/commserver/library</a>	Communications Server products - white papers, product documentation, etc.
<a href="http://www.redbooks.ibm.com">http://www.redbooks.ibm.com</a>	ITSO redbooks
<a href="http://www.ibm.com/software/network/commserver/support">http://www.ibm.com/software/network/commserver/support</a>	Communications Server technical Support
<a href="http://www.ibm.com/support/techdocs/">http://www.ibm.com/support/techdocs/</a>	Technical support documentation (techdocs, flashes, presentations, white papers, etc.)
<a href="http://www.rfc-editor.org/rfcsearch.html">http://www.rfc-editor.org/rfcsearch.html</a>	Request For Comments (RFC)