

z/OS Communications Server – SNA/EE enhancements



© 2011 IBM Corporation

This presentation describes several SNA and Enterprise Extender enhancements in z/OS® V1R13 Communications Server. The SNA and Enterprise Extender theme includes four enhancements: EE firewall-friendly connectivity test, improved APPN routing resilience, an HPR packet trace analyzer for EE, and IDS support for EE. The IDS support for EE is described in a separate presentation.

EE firewall-friendly connectivity test

- `DISPLAY EEDIAG,TEST=YES` command provides information about an EE partner and all routers in between
 - Uses ICMP messages to provide information about all routers in between
 - Firewalls configured to block ICMP traffic can delay command results
- `DISPLAY EEDIAG,TEST=YES,LIST=SUMMARY` command
 - Enhanced to quickly determine the reachability of the remote partner
 - Intermediate hop count is no longer determined
- `DISPLAY EEDIAG,TEST=YES,LIST=DETAIL` command processing is unchanged
 - Intermediate hop count and router information included in output

The EE connectivity test command `DISPLAY EEDIAG,TEST=YES` provides a wealth of information about the reachability of a potential (or actual) EE partner. It includes the list of the routers along the path to that destination and the round-trip time associated with the path to each of the routers. To learn the path of the EE connection, the time-to-live (TTL) is initially set to one in the first test probe packet. The router will return a time-exceeded ICMP message since it cannot forward the packet. Upon receipt of the time-exceeded ICMP message, the TTL is incremented by one each time a new hop is learned. This allows you to learn all the hops in the EE path.

The command can take a long time to determine the availability of the remote partner. Some customers configure their firewalls to block ICMP traffic. In those cases, the host issuing the command never receives ICMP responses from any router on the opposite side of the firewall. The command still satisfies the main goal of verifying EE partner reachability because the actual destination does not return an ICMP response, but instead sends an EE LDLC PROBE reply. The EE LDLC PROBE reply is a UDP packet just like the request packets. However, proving reachability can take some time. Every time the firewall discards an ICMP response, the host has to go through a timeout/retry sequence that totals nine seconds, before attempting to probe the next hop in the route. So, the time necessary to prove destination reachability is approximately equal to the number of hops past the firewall times nine seconds, which can be a fair amount of time. In some cases, if the total test time exceeds the MAXTIME value, the connectivity test might fail.

EE firewall-friendly connectivity test: sample output

```

D NET,EEDIAG,TEST=YES,IPADDR=(9.67.1.1,9.67.1.5),LIST=SUMMARY
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EEDIAG
IST2119I ENTERPRISE EXTENDER DISPLAY CORRELATOR: EE000001
IST2067I EEDIAG DISPLAY ISSUED ON 08/29/05 AT 15:41:22
~~~~~
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 9.67.1.5
IST924I -----
IST2133I INTFNAME: LTRLE1A                INTFTYPE: MPCPTP
IST2134I CONNECTIVITY SUCCESSFUL          PORT: 12000
IST2137I *NA 9.67.1.5                      RTT:    6
IST2134I CONNECTIVITY SUCCESSFUL          PORT: 12001
IST2137I *NA 9.67.1.5                      RTT:    6
IST2134I CONNECTIVITY SUCCESSFUL          PORT: 12002
IST2137I *NA 9.67.1.5                      RTT:    6
IST2134I CONNECTIVITY SUCCESSFUL          PORT: 12003
IST2137I *NA 9.67.1.5                      RTT:    6
IST2134I CONNECTIVITY SUCCESSFUL          PORT: 12004
IST2137I *NA 9.67.1.5                      RTT:    7
IST924I -----
IST2139I CONNECTIVITY TEST RESULTS DISPLAYED FOR 1 OF 1 ROUTES
IST314I END

```

The DISPLAY EEDIAG,TEST=YES,LIST=SUMMARY command output will provide the remote partner reachability information quickly. VTAM® will not attempt to probe each router along the path, but will set the maximum hop limit when sending an LDLC probe to the remote partner using all five EE ports. This determines quickly if the partner is reachable or not on all five ports. Since VTAM sets the hop count to 255, the path hop count is no longer determined for LIST=SUMMARY output.

The EEDIAG,TEST=YES,LIST=DETAIL command output is unchanged and includes the path hop count and intermediate router information.

HPR packet trace analyzer for EE

- Enterprise Extender (EE) is used to transport SNA data over an IP network
 - UDP is used for IP network transport
 - RTP component of HPR is responsible for reliable data delivery
- Basic functions performed by the RTP component are:
 - Handling of sequence numbers of packets
 - Sending data acknowledgments
 - Handling out of order packets
 - Controlling data flow rate (Adaptive Rate Basing)
- Businesses that heavily use EE need to know their network behavior

EE networks might not operate normally if the underlying infrastructure is not properly configured or if there is any hardware or software problem in the network. Statistical information can provide you with an overview of your network health.



HPR packet trace analyzer for EE: sample formatter output

```
COMPONENT TRACE SUMMARY FORMAT
COMP(SYSTCPDA)SUBNAME((TCPCS))
OPTIONS(HPRDIAG(SUMMARY))
z/OS TCP/IP Packet Trace Formatter, Copyright IBM Corp. 2000, 2010; 2009.339
DSNAME('IPCS.P99999A.PKTA')
```

```
Start time: 2010/01/13 16:09:39.983828
End time: 2010/01/13 16:10:22.807176
```

TCID	Pkt	Byte	OfOrder	OfOrder	ReXmit	ReXmit	SlowDown
1B7580F800010034	1	101	0	0	0	0	0
1B7580F900010035	3789	220212	0	0	0	0	0
1B762A7F0001001E	1	156	0	0	0	0	0
1B762A800001001F	3340	218802	0	0	0	0	0
1B762A800001002D	4340	518802	0	0	3	258	0

The packet trace HPRDIAG summary report is intended to provide the operators of the network with any signs of trouble in the network flow. If network problems are observed, then further diagnostic procedures should be taken to analyze the problem. This summary report is used only for checking the overall health of EE data flow. The report is generated using active Transport Connection Identifiers (TCIDs). It is possible that some of the TCIDs, such as those for CP-CP sessions, might not show any statistics on them. It is not possible to correlate the local TCID with the remote TCID, because the packet trace does not have access to the RTP layer of the session. The VTAM DISPLAY ID=rtp_pu command should be used to find the TCIDs that are being used for each SNA session.

Improved APPN routing resilience

- APPN session routes selected by APPN Topology and Routing Services (TRS)
 - Directed search routes used to locate resources
 - Routes for LU-LU sessions
- Optimal route determined by specified criteria
 - Line speed
 - Cost of data transmitted
 - Security
- TRS builds complex routing trees to determine best path
 - Tree records represent nodes along preferred route
 - Transmission groups (TGs) between nodes are associated with tree records

The APPN Topology and Routing Services, or TRS, component is responsible for selecting the optimal routes for sessions. The best path is based on specified criteria such as line speed, cost of data transmitted, security, and user defined values. These routes are used to locate resources that are represented by logical units (LUs) and to send and receive data between these LUs on LU-LU sessions. To determine the best route for a session, TRS builds complex routing trees that consist of tree records representing network nodes, virtual nodes, and border nodes along the preferred path. In addition, transmission groups (TGs) between nodes are associated with the tree records. These tree records are linked to form a routing tree.

Improved APPN routing resilience: route selection ABEND recovery

- Recovery routine for route selection ABENDs is improved
 - Prevents recursive ABENDs caused by compromised tree record pointers
 - Frees all storage that contains routing trees (TREEBLD storage pool)
- Existing sessions are not affected
- Can temporarily impact VTAM performance in large APPN networks
 - New routing trees created for new session route requests, as needed
 - Same tree build processing as at VTAM initialization
- Topology database is unaffected, so no network impact due to TDU flows

On rare occasions, a routing tree record can be corrupted, resulting in recursive ABENDs. For example, an overlay of a pointer in a tree record can cause this corruption. Once the tree record is compromised, every time TRS attempts to select a route using that routing tree additional ABENDs will occur and the session will fail. You cannot recover from these situations without restarting VTAM.

In V1R13, when an ABEND occurs during the TRS route selection process, the recovery routine will automatically free all storage that contains routing trees (TREEBLD storage pool). The recovery routine will not check any pointers in the tree records, so no additional ABENDs will follow the first ABEND. Existing sessions are not affected.

Clearing the APPN routing trees can temporarily impact VTAM performance in APPN networks with a large number of network nodes as routing trees are created for new session route requests. Once the new routing trees are built, performance is no longer impacted by these initial tree rebuild activities. This is the same tree build processing that occurs when VTAM is initialized and sessions are started.

The topology database is unaffected, so there is no additional network impact due to TDU flows.



Improved APPN routing resilience: non-ABEND recovery

- Or session setup failures with sense codes
 - 80140001, 087F0001, 087D0001, or 08400007
- New FUNCTION=CLRTREES operand on MODIFY TOPO command
- Processing is the same as recovery routine for route selection ABENDs
- Do not use the MODIFY *procname*,TOPO,FUNCTION=CLRTREES command unless you are advised by IBM service to do so

```
F NET, TOPO, FUNCTION=CLRTREES
IST097I MODIFY ACCEPTED
IST223I MODIFY TOPO COMMAND COMPLETED
```

It is possible that a compromised tree record will not cause an ABEND, but you will see session setups fail with sense codes such as 80140001, 087F0001, 087D0001, and 08400007. Again, before V1R13, you cannot recover from these situations without restarting VTAM.

In V1R13, a new value, CLRTREES, is added for the FUNCTION operand on the MODIFY TOPO command.

The MODIFY TOPO,FUNCTION=CLRTREES command performs the same processing that was added to the recovery routine for route selection ABENDs.

FUNCTION=CLRTREES will free all routing trees (TREEBLD storage pool) to prevent further problems caused by compromised tree records. Again, existing sessions are not affected.

Typically, you should not need to clear the APPN routing tree information. APPN routing problems are rarely caused by compromised routing trees, and are not resolved with this command. IBM service can determine when it is necessary to use the MODIFY *procname*,TOPO,FUNCTION=CLRTREES command to clear the routing trees.

Improved APPN routing resilience: new message

- DISPLAY NET,TOPO,LIST=SUMMARY output indicates if routing trees have been cleared

```

D NET, TOPO, LIST=SUMMARY
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TOPOLOGY
IST1306I LAST CHECKPOINT  ADJ  NN  EN  SERVED EN  CDSERVR  ICN  BN
IST1307I 07/23/10 12:32:30 2   3   0   0       0       2   2
IST1781I INITDB CHECKPOINT DATASET  LAST GARBAGE COLLECTION
IST1785I 07/21/10 08:00:20          07/23/10 08:00:22
IST2360I ROUTING TREES LAST CLEARED AT 07/22/10 18:43:42 BY VRR
IST314I END
  
```

When the TRS routing trees are cleared by the MODIFY NET,TOPO,FUNCTION=CLRTREES command, message IST2360I displays “MODIFY” on the DISPLAY NET,TOPO,LIST=SUMMARY command output. When the TRS routing trees are cleared by the VTAM recovery routine (VRR) for TRS, as the result of a route selection ABEND, message IST2360I displays “VRR” on the DISPLAY NET,TOPO,LIST=SUMMARY command output. When the TRS routing trees have not been cleared, message IST2360I is not displayed.



Feedback

Your feedback is valuable

You can help improve the quality of IBM Education Assistant content to better meet your needs by providing feedback.

- Did you find this module useful?
- Did it help you solve a problem or answer a question?
- Do you have suggestions for improvements?

Click to send email feedback:

mailto:iea@us.ibm.com?subject=Feedback_about_SNAmisc.ppt

This module is also available in PDF format at: ../SNAmisc.pdf

You can help improve the quality of IBM Education Assistant content by providing feedback.



Trademarks, disclaimer, and copyright information

IBM, the IBM logo, ibm.com, VTAM, and z/OS are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of other IBM trademarks is available on the web at "[Copyright and trademark information](http://www.ibm.com/legal/copytrade.shtml)" at <http://www.ibm.com/legal/copytrade.shtml>

THE INFORMATION CONTAINED IN THIS PRESENTATION IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY. WHILE EFFORTS WERE MADE TO VERIFY THE COMPLETENESS AND ACCURACY OF THE INFORMATION CONTAINED IN THIS PRESENTATION, IT IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. IN ADDITION, THIS INFORMATION IS BASED ON IBM'S CURRENT PRODUCT PLANS AND STRATEGY, WHICH ARE SUBJECT TO CHANGE BY IBM WITHOUT NOTICE. IBM SHALL NOT BE RESPONSIBLE FOR ANY DAMAGES ARISING OUT OF THE USE OF, OR OTHERWISE RELATED TO, THIS PRESENTATION OR ANY OTHER DOCUMENTATION. NOTHING CONTAINED IN THIS PRESENTATION IS INTENDED TO, NOR SHALL HAVE THE EFFECT OF, CREATING ANY WARRANTIES OR REPRESENTATIONS FROM IBM (OR ITS SUPPLIERS OR LICENSORS), OR ALTERING THE TERMS AND CONDITIONS OF ANY AGREEMENT OR LICENSE GOVERNING THE USE OF IBM PRODUCTS OR SOFTWARE.

© Copyright International Business Machines Corporation 2011. All rights reserved.