



Enterprise Networking Solutions

Enterprise Extender on z/OS CS: Hints and Tips


Sam Reynolds
samr@us.ibm.com
March 10, 2014



Agenda



- ☐ SNA: Dead or Alive?
- ☐ Modernizing SNA
- ☐ Enterprise Extender Recommendations
- ☐ Diagnostic Tools
- ☐ State of SNA
- ☐ References



SNA: Dead or Alive?

"This report of my death was an exaggeration."

Mark Twain, 1897

SNA, 2014

- SNA Applications: Over a trillion lines of customer written application code based on CICS, IMS, and DB2
- A large percentage (a majority?) of all business data is still accessed via SNA applications
- Numerous market factors including the continued convergence of enterprise networks onto IP technologies, and the withdrawal of the venerable 3745 from marketing, have led to a very rapid adoption of Enterprise Extender as a key component of SNA application access strategy amongst the IBM customer set.
- There have been multiple EE user experience presentations at SHARE conferences, by organizations such as Bank of America, Finanz Informatik, iT-Austria, State Farm, and the Social Security Administration



SNA: A Little History

- SNA developed and distributed by IBM
 - Specifications have been provided so that a large number of other vendors also provide products that implement SNA
- 1974 - SNA Announced
 - Systems Network Architecture
 - SNA is an architecture defining protocols such as:
 - Link Protocols
 - Node Intercommunication Protocols
 - Application Protocols
- VTAM is the mainframe and NCP is the front-end SNA product
 - VTAM runs on MVS, VM, and VSE
- VTAM and TCP/IP for MVS were combined into a single product (Communications Server) in 1996
- NCP (Network Control Program) originally ran on a hardware front-end controller, such as the 3745. It can also run on a 37xx-emulator known as the Communications Controller for Linux (CCL).



SNA History: The Road to Enterprise Extender

- In the beginning, there was hierarchical SNA: VTAM, NCP, subareas, SSCPs, VRs, ERs, SNI...
- In the late 1980's, Advanced Peer-to-Peer Networking (APPN) began to emerge, first appearing on the System 36 and AS/400, and by the early 90's on PCs via NS/2 (later CM/2).
- APPN first appeared on the mainframe with VTAM V4R1 in 1993, although LEN ("APPN without a brain") was supported in V3R4 of VTAM.
- The first support for High Performance Routing (HPR) was introduced in VTAM V4R3 in 1995. Support followed for the 22xx router, the Nways 950, PCs via CS/2, and Cisco routers via the SNASw feature. Other platforms would follow over time.
- Enterprise Extender was first shipped on the mainframe with OS/390 V2R7 in early 1999 (and was simultaneously made available on V2R6 via PTF).
- Enterprise Extender has now been implemented by numerous customers, and continues to evolve, with enhancements in each z/OS release, including V2R1, which shipped last year.



Modernizing SNA



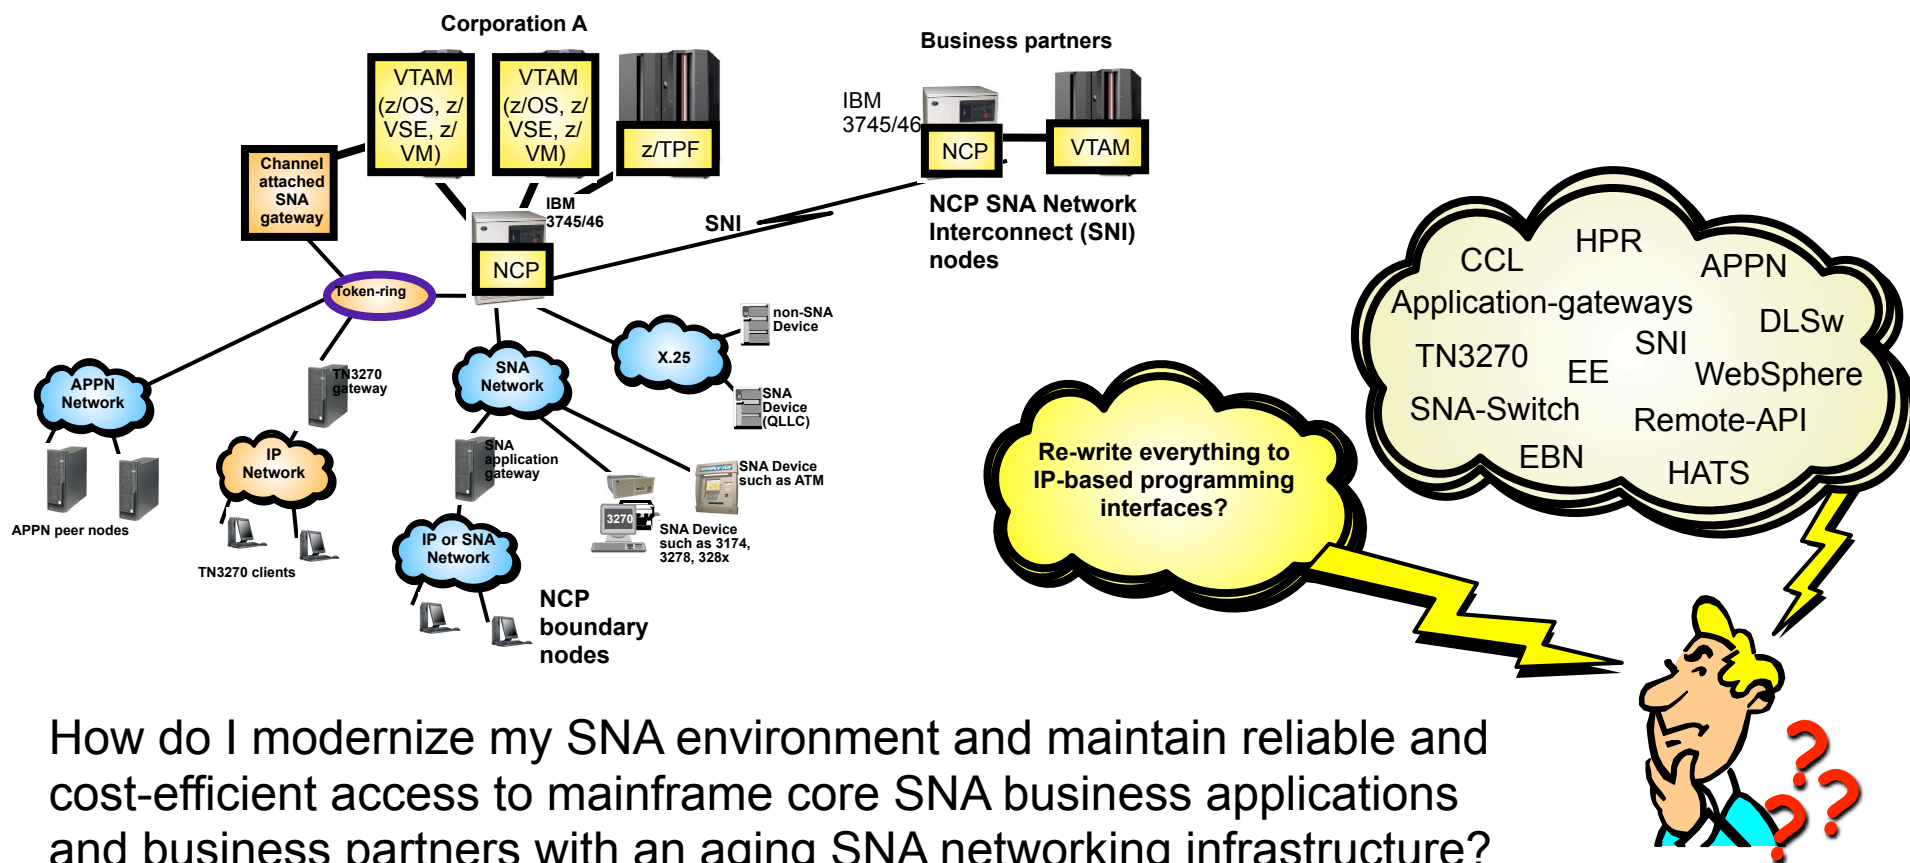
SNA modernization is about preserving SNA applications, not replacing them

- Analysts estimate that 200 billion lines of COBOL code exist today
 - 5 billion lines are added each year
 - Similar inventory of PL/I code
- The typical mainframe customer has:
 - 30M lines of COBOL code
 - Worth \$600M
 - Automating 100,000 business processes
- Any mainframe customer
 - Banking, Insurance, Government, Manufacturing, Travel and Transportation, Distribution and Retail, Media and Utilities, Healthcare Industries
- A majority (70-80% according to some studies) of these existing applications are terminal-access based

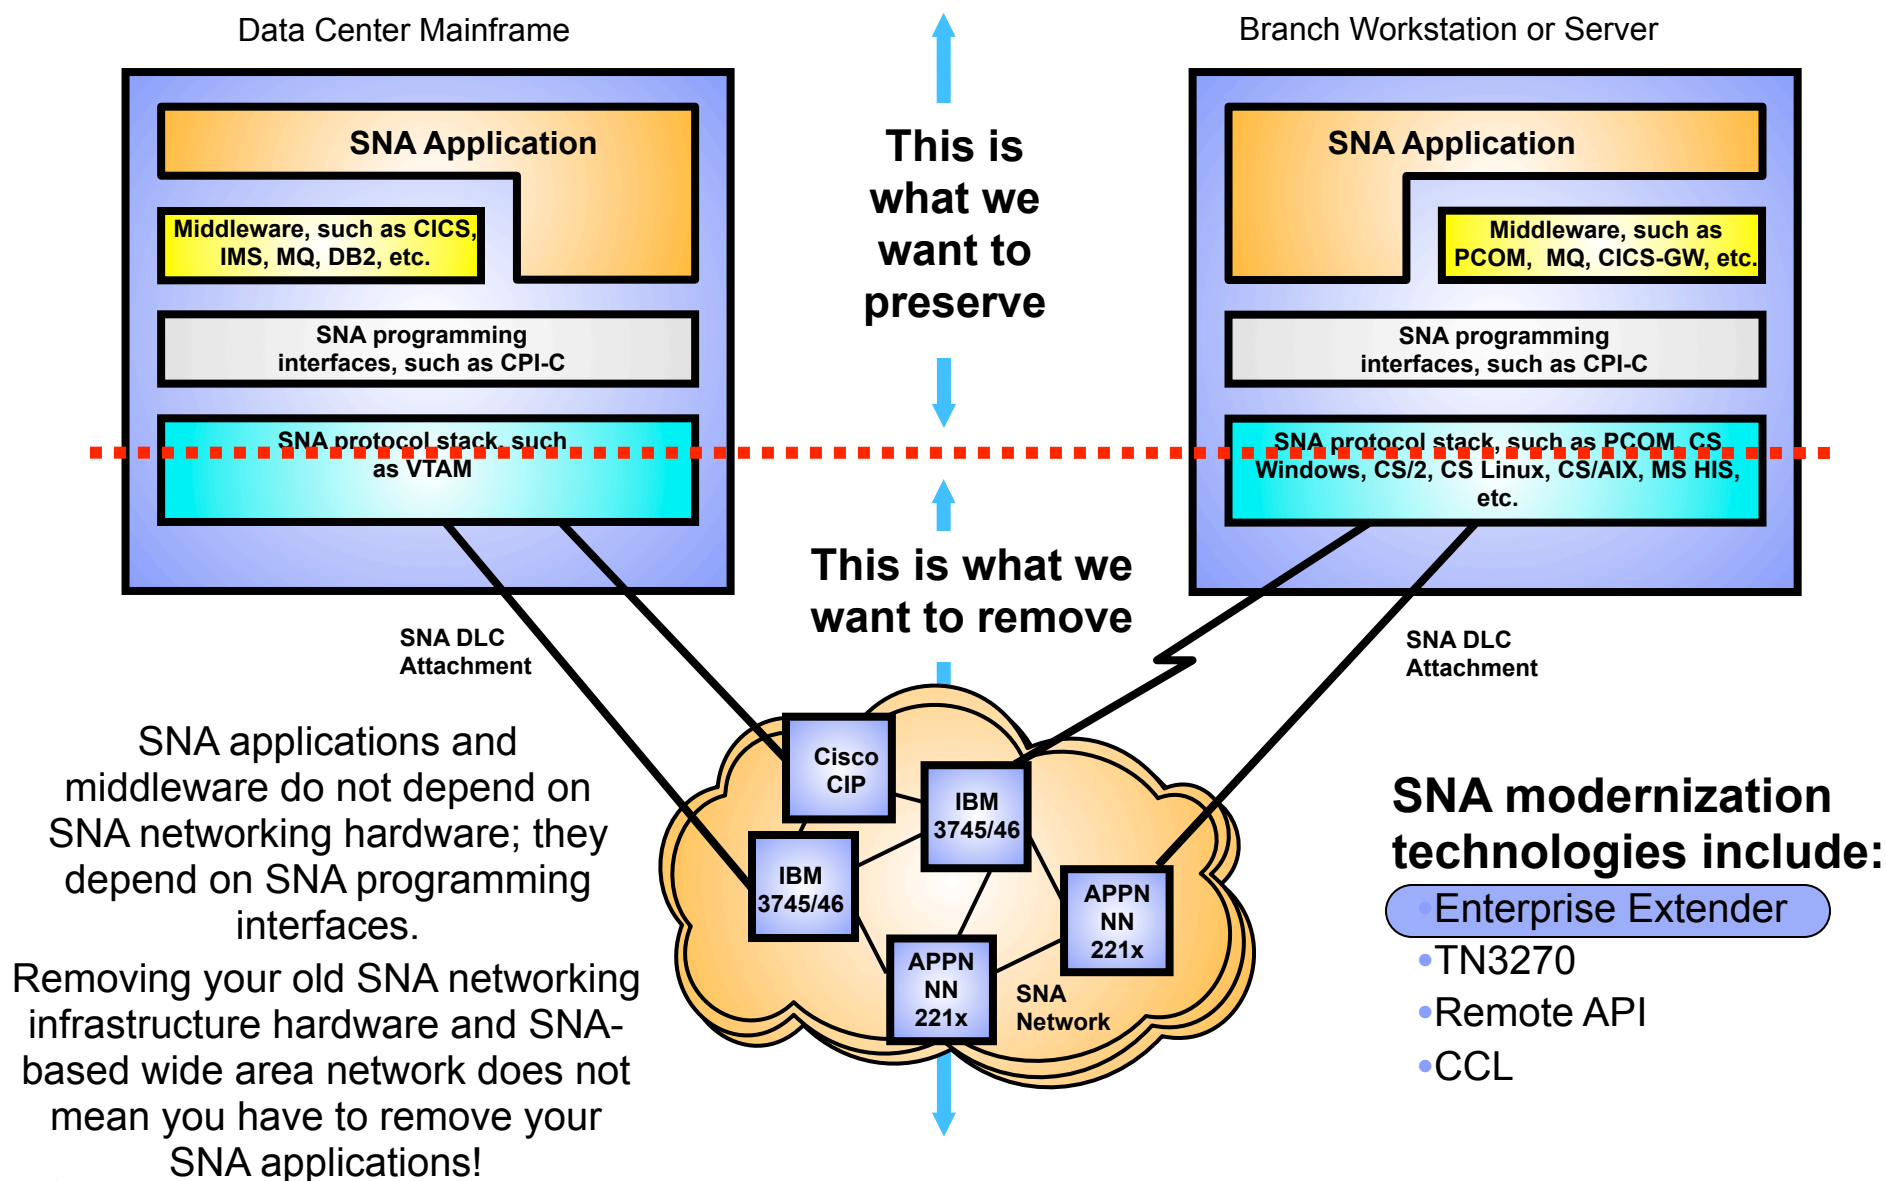
Modernizing SNA is not about re-writing or throwing away SNA applications. It is about preserving core SNA business applications in an IP-based network infrastructure and it is about enabling re-use of those applications in new end-user environments in an application-transparent manner.



SNA networks and SNA applications in 2014 and beyond - what are the questions that need to be asked?




What do we want to remove and what do we want to preserve?



What is Enterprise Extender?

- Allows use of IP network for SNA sessions
- **EE allows enablement of IP applications and convergence on a single network transport while preserving SNA application and endpoint investment.**
- EE is APPN HPR routing over an IP network
 - To the IP network, EE looks like a UDP application
 - To the APPN network, EE looks like an HPR link

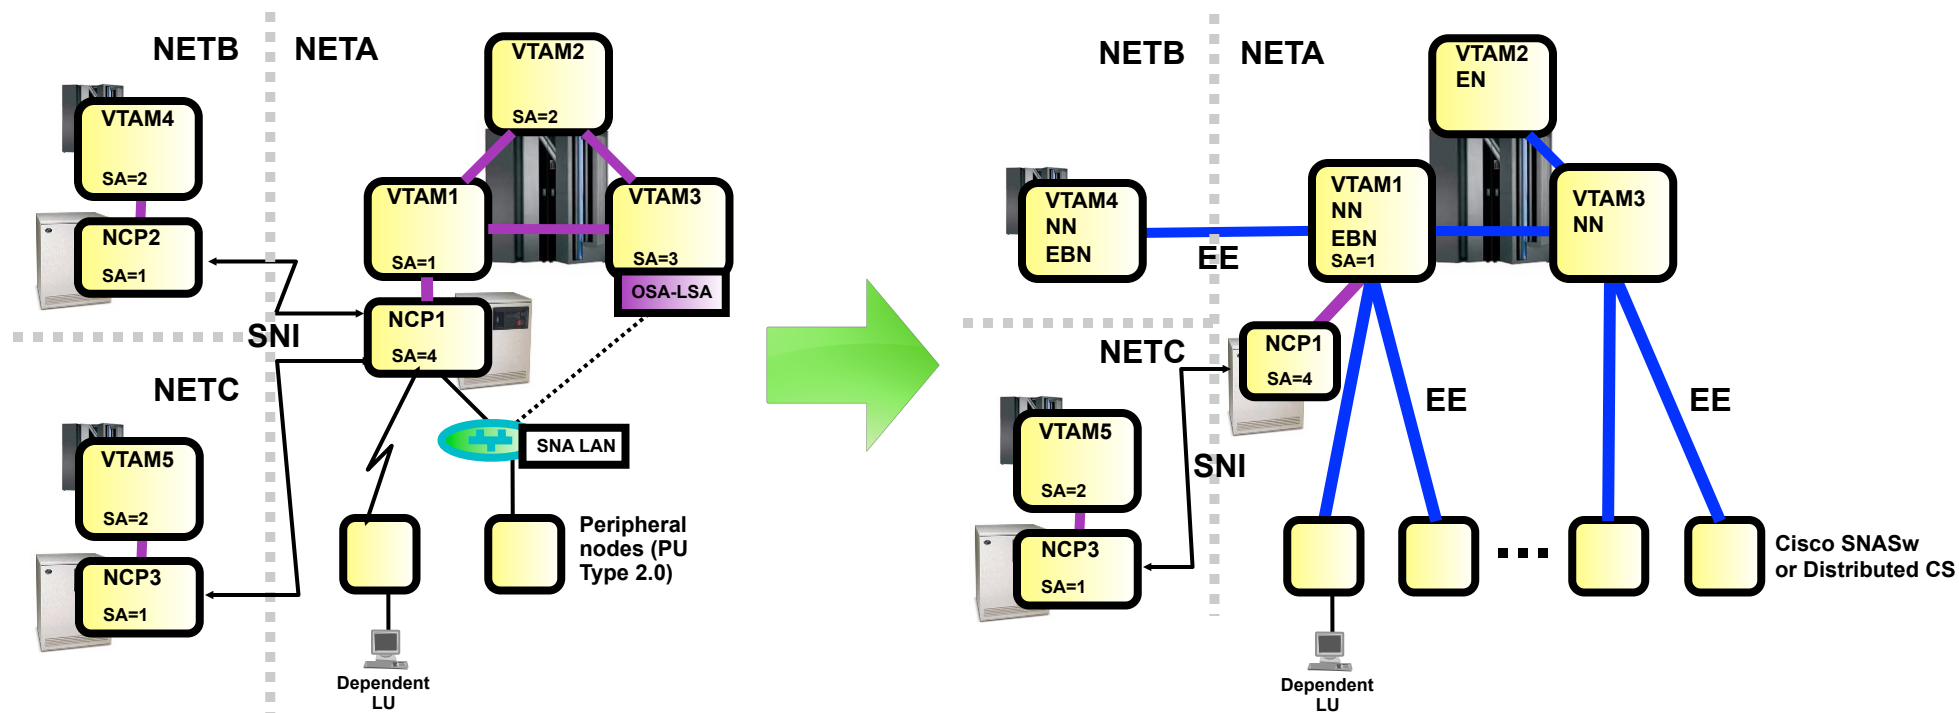


Enterprise Extender characteristics

- The SNA traffic is sent as UDP datagrams over the IP network, each EE endpoint using 5 UDP port numbers
 - Firewalls can be an issue, especially between business partners
- EE can be implemented on the SNA application hosts, or on APPN nodes that act as EE gateways
- Complete APPN/EE nodes include z/OS, CS/Linux, CS/AIX, and CS/Windows
 - Some EE nodes implement an EE-DLC connectivity function without being APPN network node capable - examples include Microsoft's Host Integration Server which can only be an end node, and Cisco SNA Switch which can only be a Branch Extender node
 - i5/OS (iSeries) added EE support to i5/OS V5R4 in 2006
- Since EE is HPR over IP, EE traffic inherits all the APPN/HPR characteristics including non-disruptive path switch
- EE traffic can be secured using IPSec
 - But not with SSL/TLS - SSL/TLS is TCP only
- Business partner connectivity through EE/EBN (z/OS only)



Subarea SNA to Enterprise Extender



Subarea SNA

Enterprise Extender

Subarea (FID4) connections ———
Enterprise Extender logical links ———


Enterprise Extender Recommendations



EE Recommendations

- #1 problem in bringing up EE connections: Firewall issues.

- Any underlying firewall must allow the transport of UDP ports 12000-12004 (possibly limited to known EE endpoint addresses)



- Configure APPN Link Characteristics

- TGP's for EE provided with VTAM - Customization of link speed is recommended
- Consider lengthening the EE LDLC timer parameters (LIVTIME, SRQTIME, SRQRETRY).
 - It is recommended that the LDLC timer parameters be adjusted on both ends of the connection.
- Lengthen HPR path switch timers (HPRPST) as necessary to ensure that all four timers are longer than the LDLC timeout interval.
 - This will ensure that RTP pipes stay in path switch long enough during IP network instability to allow the EE link to inop, and thereby allow another path to be selected.
 - If multiple LDLC parameter sets are in use (by coding different LIVTIME, SRQTIME, and SRQTIME values on different XCA GROUPs), then the HPRPST values should be adjusted relative to the longest of the LDLC timeout intervals.

EE/HPR Timer Notes

NOTES

- The EE LDLC layer monitors the connection, testing it if no activity is detected, and inop'ing the EE link if the tests go unanswered
 - Total LDLC timeout interval = $LIVTIME + ((SRQRETRY+1) * SRQTIME)$
- The Disconnect timer associated with the EE switched PU is used to trigger an inop if no activity is detected for a specified amount of time
- The RTP layer is responsible for driving status requests frequently enough (in the absence of data traffic) to keep the disconnect timer from tripping. The RTP endpoint will inop itself if its last session goes away, and no new session is queued to it for a period of 10 seconds (or the period specified in the RTP model, if used)
- The path switch timer (set by HPRPST) controls the amount of time that an RTP pipe will stay in path switch state, trying to find an alternate route, before abandoning the attempt and inactivating the pipe.

EE Recommendations ...

- Define all of your EE connection networks on GROUP statements, and not on the XCA PORT. This will provide more consistent definitions as your connection networks expand in the future, and allows you to utilize GROUP-based configuration options.
- Having GROUP-based definitions also enables more flexible and less disruptive updating of the EE XCA definitions via the VARY ACT,UPDATE=ALL command
- Leave the IPPORT operand at its default of 12000. If you change it, you must change it on all EE platforms that you connect to, and some do not support such a change.
- If defining an EE Connection Network over an IP network which employs Network Address Translation (NAT), you must define the virtual routing node's addressability using the HOSTNAME operand (not the IPADDR operand).

```

DEMOXCA  VBUILD  TYPE=XCA
DEMOPORT PORT    MEDIUM=HPRIP
*
DEMOGRP  GROUP   DIAL=YES,CALL=INOUT,                                X
                    AUTOGEN=(5,EV4,P),DYNPU=YES,ISTATUS=ACTIVE
*****
* EXAMPLE VRN                                                    *
*****
DEMOVRN  GROUP   DIAL=YES,CALL=INOUT,VNNAME=NETA.LVRN,                X
                    AUTOGEN=(5,LV01,P),DYNPU=YES,VNTYPE=LOCAL,        X
                    HOSTNAME=TUVIPA2.AREA51.SVT390.COM,                X
                    ISTATUS=INACTIVE,TGP=GIGENET

```

Effect of IP Multipath on EE

- The IPCONFIG MULTIPATH parameter in the TCP/IP profile enables the multipath routing selection algorithm for outbound IP traffic. (The default is NOMULTIPATH.)
 - When MULTIPATH is enabled, there are two options to choose between: PERCONNECTION and PERPACKET, with PERCONNECTION being the default choice.
 - If MULTIPATH is enabled, the choice of PERCONNECTION vs. PERPACKET makes no difference for EE.
 - In either case, IP will use a "per-batch-of-packets-from-VTAM" approach.
 - Since there is no "UDP connection", true PerConnection multipath is really not possible
 - PerPacket is too granular as it leads to too much resequencing overhead at the RTP receiving endpoint.
 - **Recommendation:**
 - z/OS V1R11 CS and earlier: If you already have MULTIPATH enabled for your TCP applications, there is no requirement to change it. But in general, MULTIPATH is probably not a particularly good idea for EE as it often leads to extra resequencing overhead at the receiver.
- z/OS V1R12 CS: Leave the VTAM start option MULTIPATH set to the default value of NO



SNA Intra-Sysplex Connectivity: MPC+, XCF, or EE using QDIO?

- MPC+ vs. XCF:

- All things being equal, MPC+ throughput should exceed XCF throughput, although using multiple XCF links can increase throughput.
- XCF links can bypass the coupling facility, increasing throughput if the CF is being used for other functions (GR, MNPS)
- Many customers define both MPC+ and XCF links, but want to prefer MPC+, with XCF available for backup.
- This can be accomplished by adding COSTBYTE=1 to the XCF TGP (in IBMTGPS) which is automatically associated with XCF TGs (assuming IBMTGPS has been activated). This makes the XCF link have a higher weight (and therefore be less desirable) than the MPC+ link for the IBM-supplied APPN Classes of Service.

- EE using QDIO and/or Hipersockets:

- Cross-CEC traffic should realize a significant performance advantage over XCF or MPC+
 - The magnitude of the improvement will vary based on the factors above
- For SNA workloads within the same CEC, EE over HiperSockets will provide superior performance unless CPU availability is limited
- Note that EE cannot use Shared Memory Communications over RDMA links (SMC-R) since SMC-R is TCP-only, and EE is UDP-based.

- Recommendation:


- Configure EE links as primary transport, with XCF and/or MPC+ available as backup links



EE Connection Network


- Connection network is an APPN technology that reduces the need for predefining APPN links between nodes that are connected to a shared transport fabric
 - With EE, the IP network itself is the shared transport.
- The shared transport (IP network in the EE case) is represented as an APPN Virtual Routing Node (VRN).
- All EE nodes participating in the connection network can send EE packets directly to each other without defining links to all the other participating nodes.
 - Connections are dynamically defined between VRN partners as needed.
- For much more detail on EE connection network, download the following presentation:

http://proceedings.share.org/client_files/SHARE_in_Denver/S3206SR100305.pdf

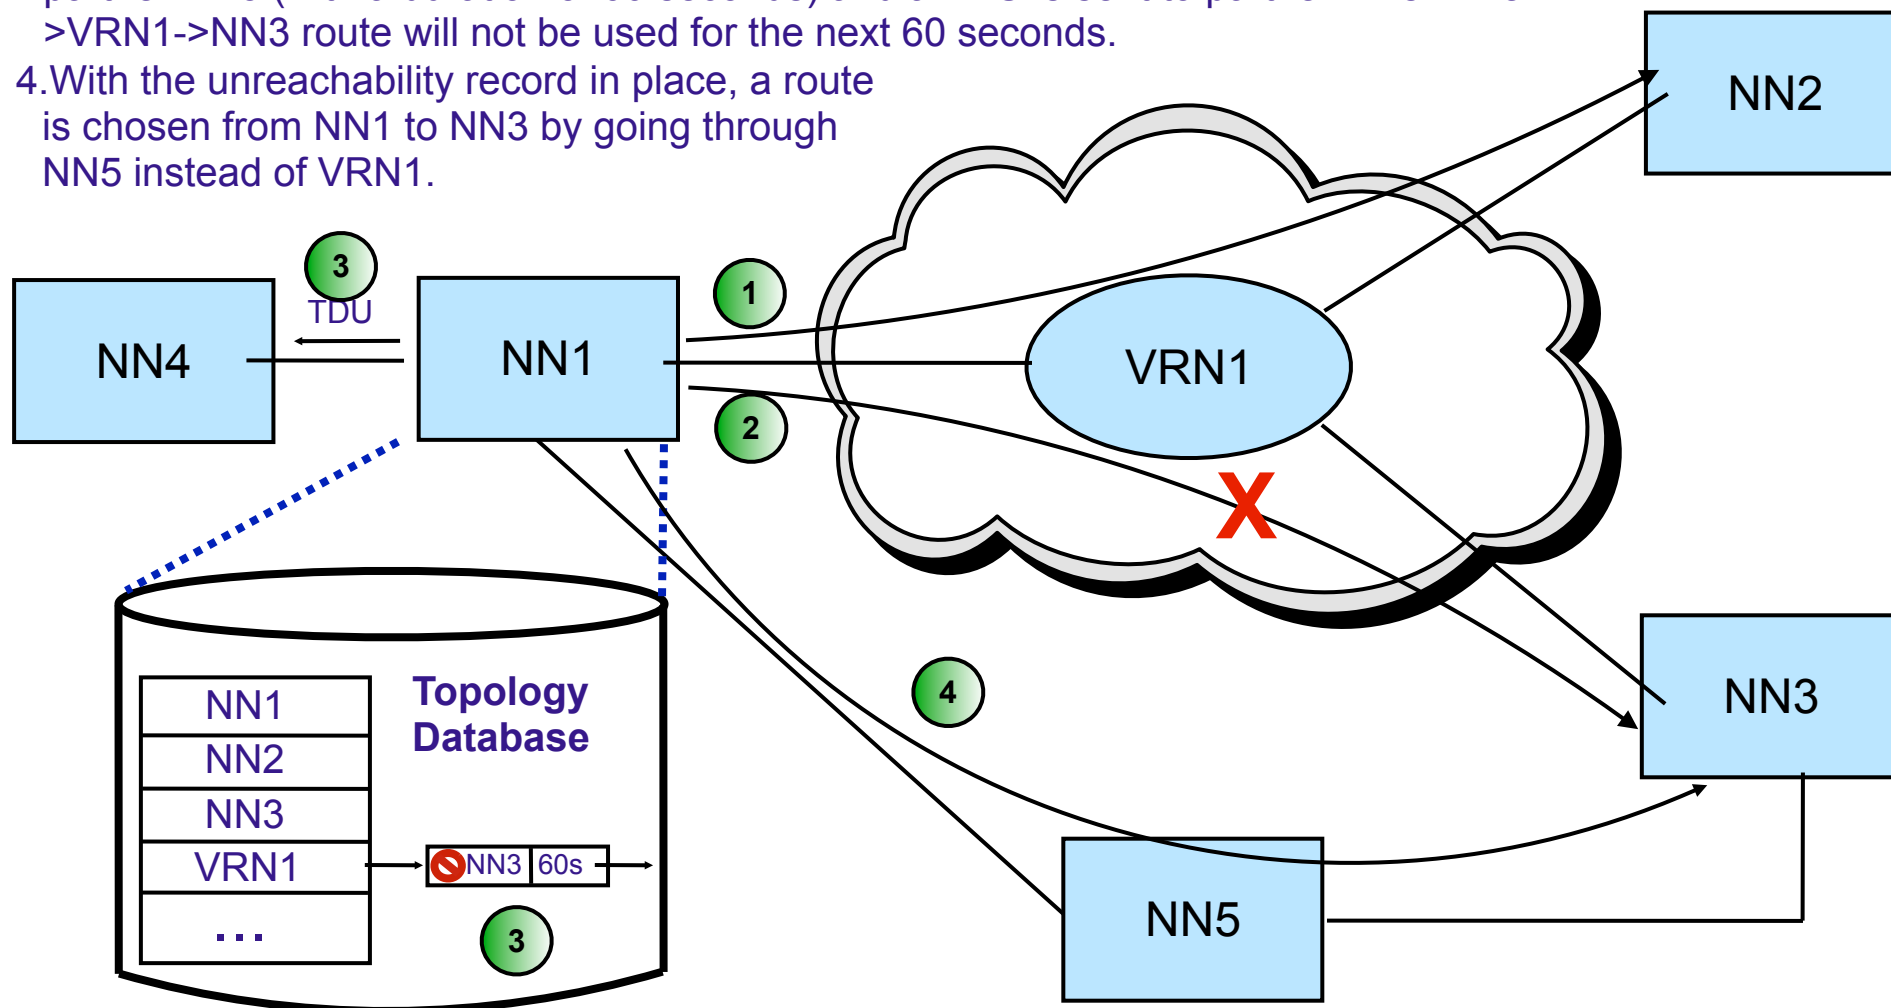


Example: Since EN2 and NN2 both define the VRN, a dynamic connection can be activated between them over the VRN without predefining a static EN2-to-NN2 connection

EE Connection Network Reachability Awareness



1. NN1 successfully contacts NN2 across VRN1.
2. NN1's attempt to contact NN3 across VRN1 fails.
3. In NN1's topology database, an "unreachability record" is associated with VRN1 for the partner NN3 (with a duration of 60 seconds) and a TDU is sent to partner NNs. The NN1->VRN1->NN3 route will not be used for the next 60 seconds.
4. With the unreachability record in place, a route is chosen from NN1 to NN3 by going through NN5 instead of VRN1.



EE Connection Network Reachability Awareness

NOTES

- EE Connection Network Reachability Awareness detects a dial failure or connection INOP for a connection over an Enterprise Extender connection network and prevents that specific path to the partner node from being used for a period of time. If alternate paths are available, APPN Topology and Routing services will select the optimal alternate session path for session establishment or an HPR path switch.
- When the time expires, if the path through the EE virtual routing node (VRN) still has the lowest weight of any available path to the partner node, the path over this particular VRN will be selected on the next attempt to redial the partner node.
- Unreachable partner information is maintained in the Topology Database and is associated with an EE VRN or with an end node that is on the origin side of the VRN.
- Unreachable partner information is sent to an end node's NNS or broadcast to a network node's adjacent network nodes in Topology Database Updates (TDUs).
- The period of time that a path through the EE VRN to the unreachable partner will remain unavailable is configurable. The UNRCHTIM start option allows the specification of the default number of seconds that a partner node for a session path through an EE connection network is considered unreachable after connection network failures.
- During this UNRCHTIM period, the path through the EE VRN to this partner node will not be considered for new sessions or HPR path switches.
- UNRCHTIM can also be specified on the PORT and/or GROUP statements in the EE XCA major node. This provides the capability of specifying different unreachable durations for connection networks of different characteristics, or to different business partners, etc.
- The range for UNRCHTIM is 0, or 10-65535 seconds.
- UNRCHTIM=0 indicates that paths through EE connection networks will always be considered for routing. This is the default value.
- Current unreachable information can be displayed using the DISPLAY TOPO command.
- Unreachability records may be cleared using the MODIFY TOPO command.
- To prevent the performance impact of unreachability lists growing without bound, once a certain number of records is associated with a connection network, that VRN will be treated as quiesced. No further attempts will be made to use it (and no more records will be added) until the record count drops below a threshold.
- A second parameter specifiable with UNRCHTIM allows control of the number of records at which the VRN will be quiesced. The VRN will be considered usable again when the record count drops below 80% of that value. So, for example, UNRCHTIM=(60,20) would set the unreachable time to 60 seconds, with the record count allowed to grow until 20 before the VRN is quiesced. The quiesced status would be cleared when the record count dropped below 16 (80% of 20).
- If UNRCHTIM is enabled, the default value for the record limit is 10.
- Consider enabling PSRETRY (off by default) so that HPR pipes will automatically switch to better routes when available.
 - This will allow HPR pipes to switch back onto the connection network path once the unreachability condition is resolved.

HPR Path Switch Summarization

- HPR Path Switch Summarization reduces the number of path switch message groups VTAM issues across a 60-second interval
- HPRPSMSG=ALL|count, where count is in the range 10-100
- The HPRPSMSG value specifies the number of IST1494I ("Path Switch Started") messages that will be issued, before suppressing the remaining ones across the 60-second interval.
 - If the STARTED message is issued for a pipe, the associated COMPLETED or FAILED message will always be issued as well
- At the end of the 60-second interval, a summarization report is issued on total path switch activity during the interval
 - The report output is limited to 10 Net IDs and 50 CPs, but the "started", "completed", and "failed" counts will accurately reflect all path switch activity
- If HPRPSMSG=ALL is specified, then all path switch message groups will be issued and no summarization provided
- If you have previously specified IST1494I in the message flooding prevention table, you will probably want to remove it when enabling path switch summarization



HPR Path Switch Summarization ...


IST2191I HPR PATH SWITCH SUMMARY FROM 04/05/06 AT 09:45:14

```

IST924I -----
IST2192I  STARTED      =      12
IST2193I   TGINOP     =      12      SRQTIMER =      0      PSRETRY      =      0
IST2194I   PARTNER    =      0      MNPS      =      0      UNAVAILABLE =      0
IST2195I   NETWORK    =      3      HIGH     =      3      MEDIUM    =      3      LOW     =      3
IST924I -----
IST2196I  COMPLETED   =      8
IST2195I   NETWORK    =      2      HIGH     =      2      MEDIUM    =      2      LOW     =      2
IST924I -----
IST2197I  FAILED       =      4
IST2195I   NETWORK    =      1      HIGH     =      1      MEDIUM    =      1      LOW     =      1
IST924I -----
IST2198I  NETID        STARTED          COMPLETED          FAILED
IST2199I   CPNAME      NET  HI MED LOW  NET  HI MED LOW  NET  HI MED LOW
IST2205I -----
IST2200I  NETA         2    2    2    2    1    1    1    1    1    1    1    1
IST2201I   SSCP2A      1    1    1    1    1    1    1    1    0    0    0    0
IST2201I   SSCP7A      1    1    1    1    0    0    0    0    1    1    1    1
IST2205I -----
IST2200I  NETB         1    1    1    1    1    1    1    1    0    0    0    0
IST2201I   SSCP99      1    1    1    1    1    1    1    1    0    0    0    0
IST924I -----
IST2206I  24 PATH SWITCH EVENTS FOR 3 CPS IN 2 NETIDS
IST314I  END
  
```

Progressive Mode ARB

- HPR's Responsive Mode ARB flow control is very sensitive to minor variations in packet round-trip time or unpredictability in response time from the RTP partner node. If the partner suddenly becomes CPU constrained, even for a short period, throughput and response time can be degraded.
- Typical causes:
 - Partner node has a shortage of CPU availability, memory, or network bandwidth
 - Partners in a virtual server environment on a single hardware platform cannot guarantee consistent response time
- V1R11 introduced a new level of the ARB flow control algorithm: Progressive Mode ARB
 - Implements several small changes to the flow control rules to improve responsiveness in a CPU-constrained environment
 - Both partners must agree to use progressive mode ARB
 - Limited to single-hop pipes over an EE connection (including two-virtual-hop connection network paths)



Progressive Mode ARB ...

- HPREEARB = PROGRESS can be specified:
 - On an EE PU in a switched major node
 - **In V2R1: On the GROUP statement in a switched major node**
 - On a connection network GROUP in the EE XCA major node
 - On the EE model PU in a model major node
- The desire/capability to use progressive mode ARB is conveyed between the RTP partners on the XID, route setup reply, and the RTP ARB setup segment. If both partners do not agree to use progressive mode ARB, responsive mode ARB will be used as the flow control protocol.
- A display of the RTP pipe will show that progressive mode ARB is being used:

```
D NET,ID=CNR00004
IST097I DISPLAY ACCEPTED
IST075I NAME = CNR00004, TYPE = PU_T2.1
.
.
IST2267I RTP PACING ALGORITHM = ARB PROGRESSIVE MODE
.
.
```

- Progressive Mode ARB is also available in Distributed CS V6.4.0 (for Windows,AIX, and Linux (both Intel and System z))

HPR Path Switch Delay

- If the RTP endpoint suspects a problem with partner communications, it will make several attempts to contact the partner, with a delay between attempts based on a “short request” (SRQ) timer value based on the round-trip time.
- At times this logic is too sensitive:
 - Transient network or partner conditions can cause temporary swings in round-trip time that can cause unnecessary entry into path switch state
 - In this case, the pipe usually path switches right back onto the same route
 - This wastes cycles and clutters the console with path switch messages
- V1R11 introduced new controls to specify a minimum time period that will be required before entering path switch state
 - Specifies the minimum amount of time a z/OS CS RTP endpoint must wait before initiating a path switch attempt due to an unresponsive partner
 - Does not control path switches initiated due to the PSRETRY function, the MODIFY RTP command, or local TG inops
 - This only affects the path switch logic on the local end of the RTP pipe. The path switch delay value is not negotiated with the RTP partner



```
...  
IST1818 PATH SWITCH REASON: SHORT REQUEST RETRY LIMIT EXHAUSTED  
...
```

HPR Path Switch Delay ...

- New start option: `HPRPSDLY = 0 | ps-delay`
 - Specifying a non-zero value for this start option sets the minimum path switch delay value
 - Range: 0 - 240 seconds
 - Default value of zero indicates that the RTP pipe should enter path switch as soon as a predetermined number of retry attempts have been unsuccessful (prior behavior)
 - Can be modified via MODIFY VTAMOPTS
- `HPRPSDLY = value-of-HPRPSDLY-start-option | ps-delay | EEDELAY` can also be specified:
 - On an EE PU in a switched major node
 - On a connection network GROUP in the EE XCA major node
 - On the EE model PU in a model major node
 - If you set HPRPSDLY to EEDELAY, VTAM will calculate a delay value long enough to allow the EE LDLC mechanism to inop the EE connection in the event the EE partner becomes unreachable
 - Local EE inop will trigger path switch processing

```
D NET,ID=CNR00004,HPRDIAG=YES
IST097I DISPLAY ACCEPTED
IST075I NAME = CNR00004, TYPE = PU_T2.1
.
.
IST924I -----
IST1984I PATH SWITCH INFORMATION:
IST2271I PATH SWITCH DELAY = 90
IST2272I PATH SWITCH DELAYED UNTIL 11/17/08 AT 12:07:34
.
```

Diagnostic Tools



EE Connectivity Test Command



- The Enterprise Extender connectivity test command is useful in debugging various network problems. This command can be used to test an existing Enterprise Extender connection, or it can be used to assist in diagnosing why an EE connection cannot be established.
- The EE connectivity test will verify:
 - EE line availability
 - Address resolution capability
 - EE partner reachability
 - The output generated from this request will show the reachability to the remote EE endpoint over all five UDP ports reserved for EE.
 - When MULTIPATH function is enabled in the Enterprise Extender capable TCP/IP stack, the EE connectivity test is repeated for each valid TCP/IP interface which routes EE traffic.

```


D NET,EEDIAG,TEST=YES,LIST=DETAIL,ID=ETU2HO
...
IST2067I EEDIAG DISPLAY ISSUED ON 07/11/07 AT 10:41:12
IST1680I LOCAL IP ADDRESS 197.51.125.1
IST1680I REMOTE IP ADDRESS 197.51.153.1
...
IST924I -----
IST2133I INTFNAME: LMTU2BR55                      INTFTYPE: MPCPTP
IST2134I CONNECTIVITY SUCCESSFUL                      PORT: 12000
IST2137I 1 197.51.155.14                      RTT: 1
IST2137I 2 197.51.153.1                      RTT: 4
IST2134I CONNECTIVITY SUCCESSFUL                      PORT: 12001
IST2137I 1 197.51.155.14                      RTT: 2
IST2137I 2 197.51.153.1                      RTT: 4
IST2134I CONNECTIVITY SUCCESSFUL                      PORT: 12002
IST2137I 1 197.51.155.14                      RTT: 2
IST2137I 2 197.51.153.1                      RTT: 5
IST2134I CONNECTIVITY SUCCESSFUL                      PORT: 12003
IST2137I 1 197.51.155.14                      RTT: 2
IST2137I 2 197.51.153.1                      RTT: 6
IST2134I CONNECTIVITY SUCCESSFUL                      PORT: 12004
IST2137I 1 197.51.155.14                      RTT: 3
IST2137I 2 197.51.153.1                      RTT: 5
...

```

EE Connectivity Test Example



IPv4 EE connection is established from HostA to HostB, but HPR performance is poor



When MULTIPATH is configured, the connectivity test will verify all computed routes from HostA to HostB

EE Connectivity Test Example ...



D NET,EEDIAG,TEST=YES,IPADDR=(9.67.1.1,9.67.1.6),LIST=DETAIL

```
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EEDIAG
IST2119I ENTERPRISE EXTENDER DISPLAY CORRELATOR: EE00000E
IST2067I EEDIAG DISPLAY ISSUED ON 10/04/05 AT 11:05:50
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 9.67.1.6
IST2023I CONNECTED TO LINE LN11
IST2126I CONNECTIVITY TEST IN PROGRESS
IST314I END
```

.
.

EE Connectivity Test Example ...



```

IST350I DISPLAY TYPE = EEDIAG
IST2130I ENTERPRISE EXTENDER CONNECTIVITY TEST INFORMATION
IST2119I ENTERPRISE EXTENDER DISPLAY CORRELATOR: EE00000E
IST2131I EEDIAG DISPLAY COMPLETED ON 10/04/05 AT 11:05:52
IST2132I LDLC PROBE VERSIONS: VTAM = V1  PARTNER = V1
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 9.67.1.6
IST924I -----
IST2133I INTFNAME: OSA1                      INTFTYPE: OSAFDDI
IST2135I  CONNECTIVITY UNSUCCESSFUL          SENSE: ***NA***    PORT: 12000
IST2137I    1  9.67.1.2                      RTT:          2
IST2137I    2  9.67.1.21                      D-1          RTT:          3
IST2135I  CONNECTIVITY UNSUCCESSFUL          SENSE: ***NA***    PORT: 12001
IST2137I    1  9.67.1.2                      RTT:          2
IST2137I    2  9.67.1.21                      D-1          RTT:          3
IST2135I  CONNECTIVITY UNSUCCESSFUL          SENSE: ***NA***    PORT: 12002
IST2137I    1  9.67.1.2                      RTT:          2
IST2137I    2  9.67.1.21                      D-1          RTT:          4
IST2135I  CONNECTIVITY UNSUCCESSFUL          SENSE: ***NA***    PORT: 12003
IST2137I    1  9.67.1.2                      RTT:          2
IST2137I    2  9.67.1.21                      D-1          RTT:          4
IST2135I  CONNECTIVITY UNSUCCESSFUL          SENSE: ***NA***    PORT: 12004
IST2137I    1  9.67.1.2                      RTT:          2
IST2137I    2  9.67.1.21                      D-1          RTT:          3
.
.

```

EE Connectivity Test Example ...

NOTES

```
IST2137I  hop  ipv4address      flags  RTT:  time
```

This message displays information gathered during the EE connectivity test over an IPv4 route using the OSA1 adaptor. Take the following IST2137I message from the previous display:

```
IST2137I      2  9.67.1.21          D-1      RTT:      3
```

2 is the TTL *hop* count used in the LDLC probe command.

9.67.1.21 is the source IPv4 address (*ipv4addr*) from the ICMP response.

In this case, the *flags* field (D-1 in this example) has a format of *t-ccc*.

Where *t* is a representative character of the ICMP message type returned in response to the LDLC probe and *ccc* represents the specific code associated with the ICMP message type.

The ICMP message type is displayed as one of the following:

D - "Destination Unreachable" ICMP Type 3

P - "Parameter Problem" ICMP Type 12

Q - "Source Quench" ICMP Type 4

Hop 2 returned an ICMP message type 3 with a specific code of 1 in response to the EE probe. Message IST2137I indicates this by displaying the *flags* field as D-1. For a list of the ICMP types and codes, see Appendix E in the z/OS Communications Server: IP System Administrator's Commands, "ICMP/ICMPv6 types and codes".

time is the round trip time for the LDLC probe to be sent to the TTL hop, and for an ICMP response to be received. In this example, the round trip time is 3 milliseconds.

EE Connectivity Test Example ...



```
IST924I -----
IST2133I INTFNAME: OSA2                      INTFTYPE: OSAFDDI
IST2134I CONNECTIVITY SUCCESSFUL              PORT: 12000
IST2137I   1 9.67.1.3                        RTT:      9
IST2137I   2 9.67.1.11                      RTT:     14
IST2137I   3 9.67.1.12                      RTT:     19
IST2137I   4 9.67.1.5                       RTT:     23
IST2137I   5 9.67.1.6                       RTT:     27
IST2134I CONNECTIVITY SUCCESSFUL              PORT: 12001
IST2137I   1 9.67.1.3                        RTT:      8
IST2137I   2 9.67.1.11                      RTT:     14
IST2137I   3 9.67.1.12                      RTT:     17
IST2137I   4 9.67.1.5                       RTT:     21
IST2137I   5 9.67.1.6                       RTT:     25
.
.
IST2134I CONNECTIVITY SUCCESSFUL              PORT: 12004
IST2137I   1 9.67.1.3                        RTT:      7
IST2137I   2 9.67.1.11                      RTT:     11
IST2137I   3 9.67.1.12                      RTT:     12
IST2137I   4 9.67.1.5                       RTT:     17
IST2137I   5 9.67.1.6                       RTT:     23
IST924I -----
IST2039I CONNECTIVITY TEST INFORMATION DISPLAYED FOR 2 INTERFACES
IST314I END
```

EE Connectivity Test Example ...

NOTES

```
IST2137I  hop  ipv4address      flags  RTT:  time
```

This message displays information gathered during the EE connectivity test over an IPv4 route using the OSA2 adaptor. Take the following IST2137I message from the previous display:

```
IST2137I      5  9.67.1.6                RTT:      27
```

5 is the TTL *hop* count used in the LDLC probe command. In this case, the EE connectivity test was successful over each of the 5 EE ports. Each IST2137I message indicating it was a 5 hop route to reach the EE partner with an IPv4 address of 9.67.1.6.

In this case, the *flags* field is blank as there were not any probe retries or any unexpected ICMP messages returned.

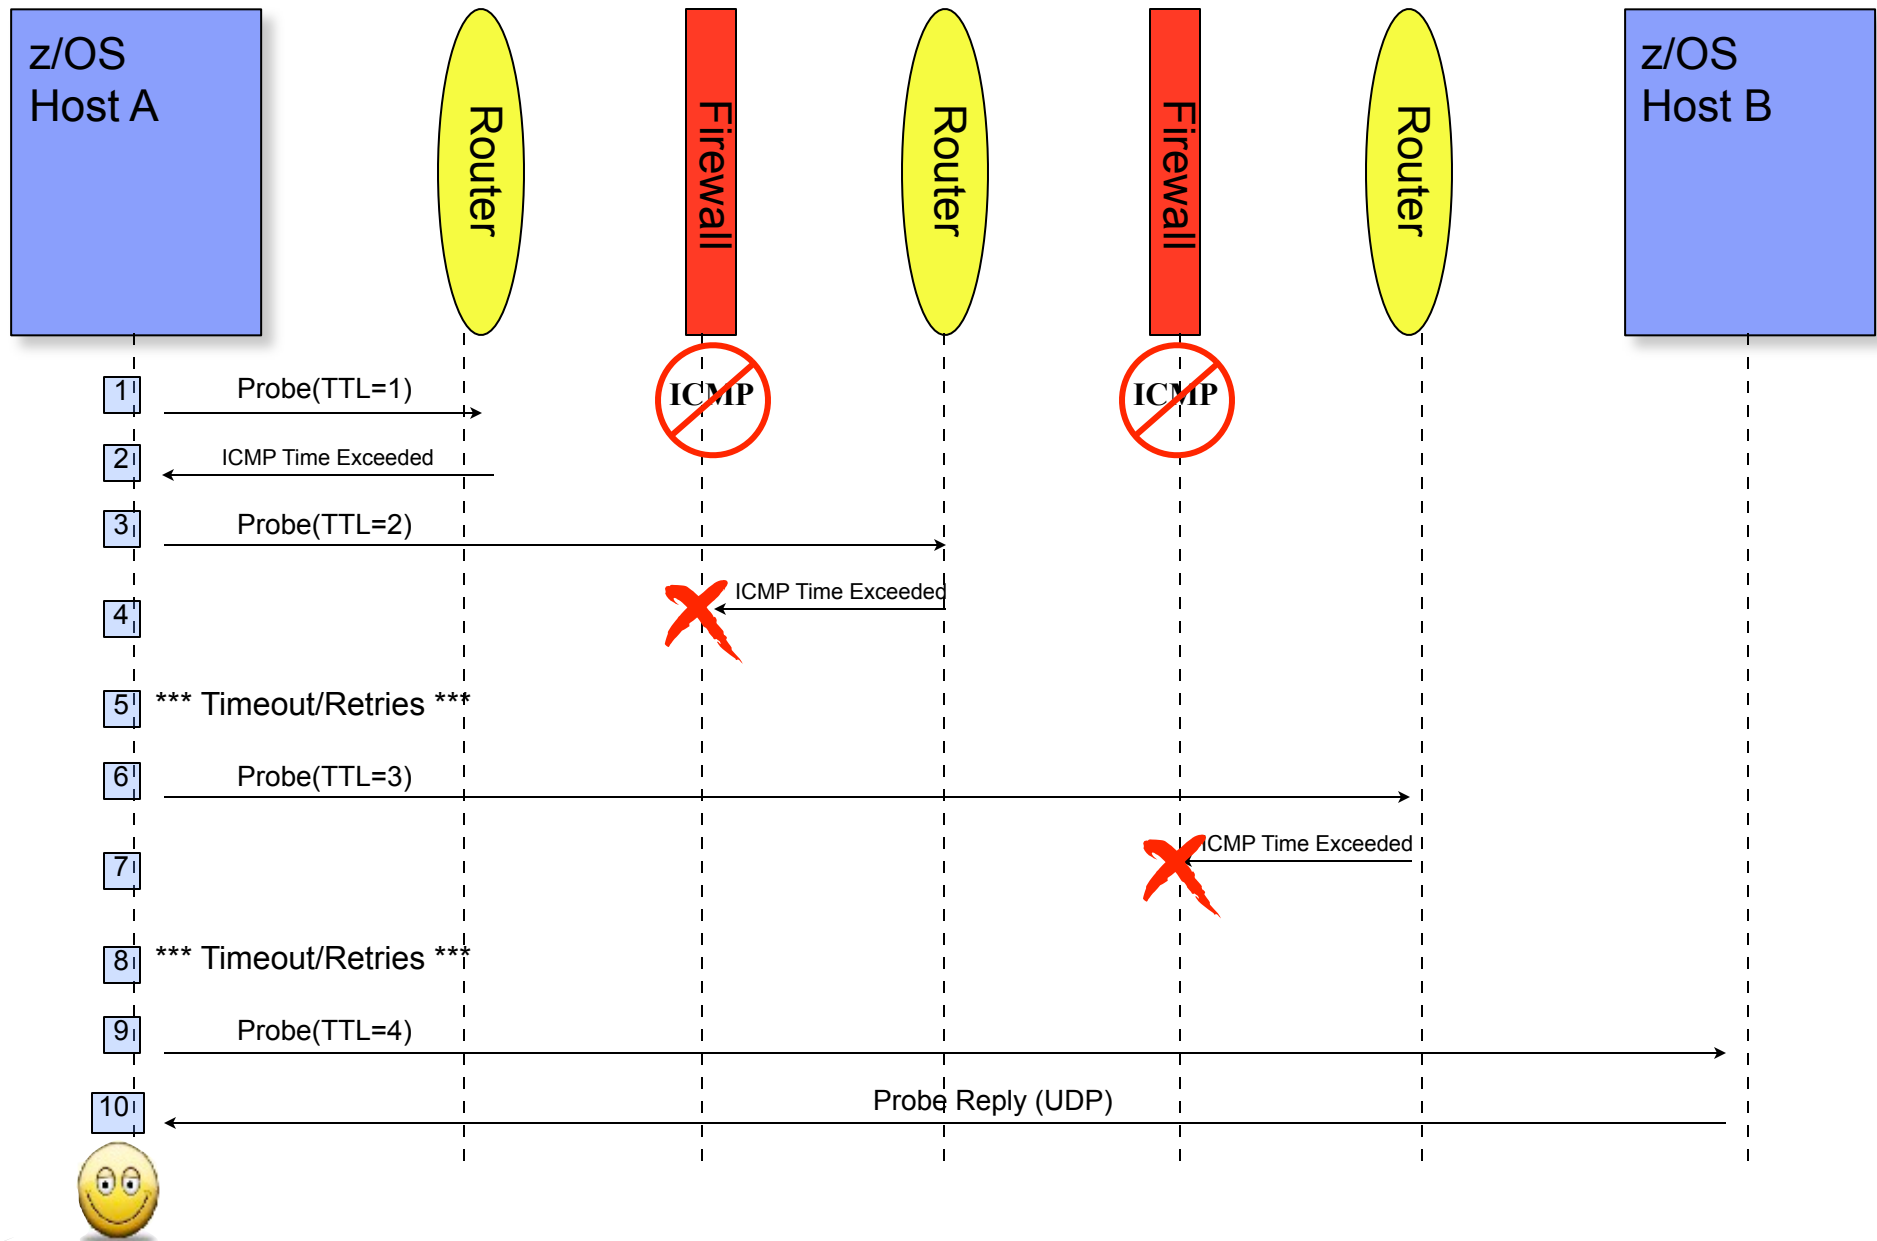
time is the round trip time for the LDLC probe to be sent to the TTL hop, and for the LDLC probe response (UDP datagram) to be received from the EE partner. In this example, the round trip time is 27 milliseconds.

EE Connectivity Test Considerations

- EE connectivity test expects ICMP messages to be returned from intermediate hops
- Firewalls may be configured to block ICMP messages
 - If possible, configure firewall to allow ICMP "Time Exceeded" messages
 - IPv4: ICMP Message Type 11
 - IPv6: ICMPv6 Message Type 3
 - If allowing ICMP responses through the firewall is not possible, intermediate hops past firewall will appear as unresponsive, but final destination reachability can still be determined (see next chart) assuming destination has EE Test probe "responder" support.
- Firewalls may be configured to block UDP traffic
 - Firewalls must allow UDP traffic on EE UDP ports 12000-12004 (both directions)
 - Test probe response may route through firewall
 - Probe response is not an ICMP message. It is a UDP datagram.
- EE Test probe "responder" support
 - Support available for CS/Windows, CS/AIX, CS/Linux, PComm, and Cisco SNASw



EE Connectivity Test w/Firewalls Blocking ICMP



EE Connectivity Test w/Firewalls Blocking ICMP ...

NOTES

The previous chart illustrates what happens when Host A issues an EE connectivity test command with Host B as the partner, but there are intervening firewalls that do not allow the passage of ICMP messages.

- (1) Host A begins trying to verify partner reachability by sending an LDLC Probe command inside of a UDP packet. (The same UDP packet would be sent on all 5 EE ports.) For the first flow, the Time-To-Live (TTL) value is set to one.
- (2) When the UDP packet reaches the first router in the path, the router decrements the TTL value to zero. Since the TTL is zero, the packet is discarded, and an ICMP "Time Exceeded" message is returned.
- (3) Upon receiving the the ICMP response, Host A sends the probe again, but with the TTL value incremented to two.
- (4) When the probe reaches the second router on the path, that router decrements TTL to zero, and again the packet is discarded and an ICMP "Time Exceeded" response generated. However, in this case a firewall between that router and Host A is blocking ICMP, and discards the message.
- (5) After three seconds, Host A will timeout and resend the probe, still with a TTL of two. Again, the ICMP response will be discarded. After a second three-second timeout, Host A will retry one last time, and will again timeout after three more seconds.
- (6) After the three attempts to send the probe with a TTL of two (and three timeouts spanning a total of nine seconds), Host A will resend the probe with a TTL value of three.
- (7) When the probe reaches the third router on the path, that router decrements TTL to zero, and again the packet is discarded and an ICMP "Time Exceeded" response generated. Again a firewall between that router and Host A is blocking ICMP, and discards the message.
- (8) Once again, Host A goes through a sequence of three-second timeouts and probe retries.
- (9) After exhausting its retries with a TTL value of three, Host A sends the probe with a TTL value of four. In this case, the probe makes it all the way to Host B.
- (10) Host B does not respond with an ICMP message, but instead sends an LDLC Probe Response message as a UDP datagram. This packet makes it back to Host A, verifying that Host B is a reachable EE partner.

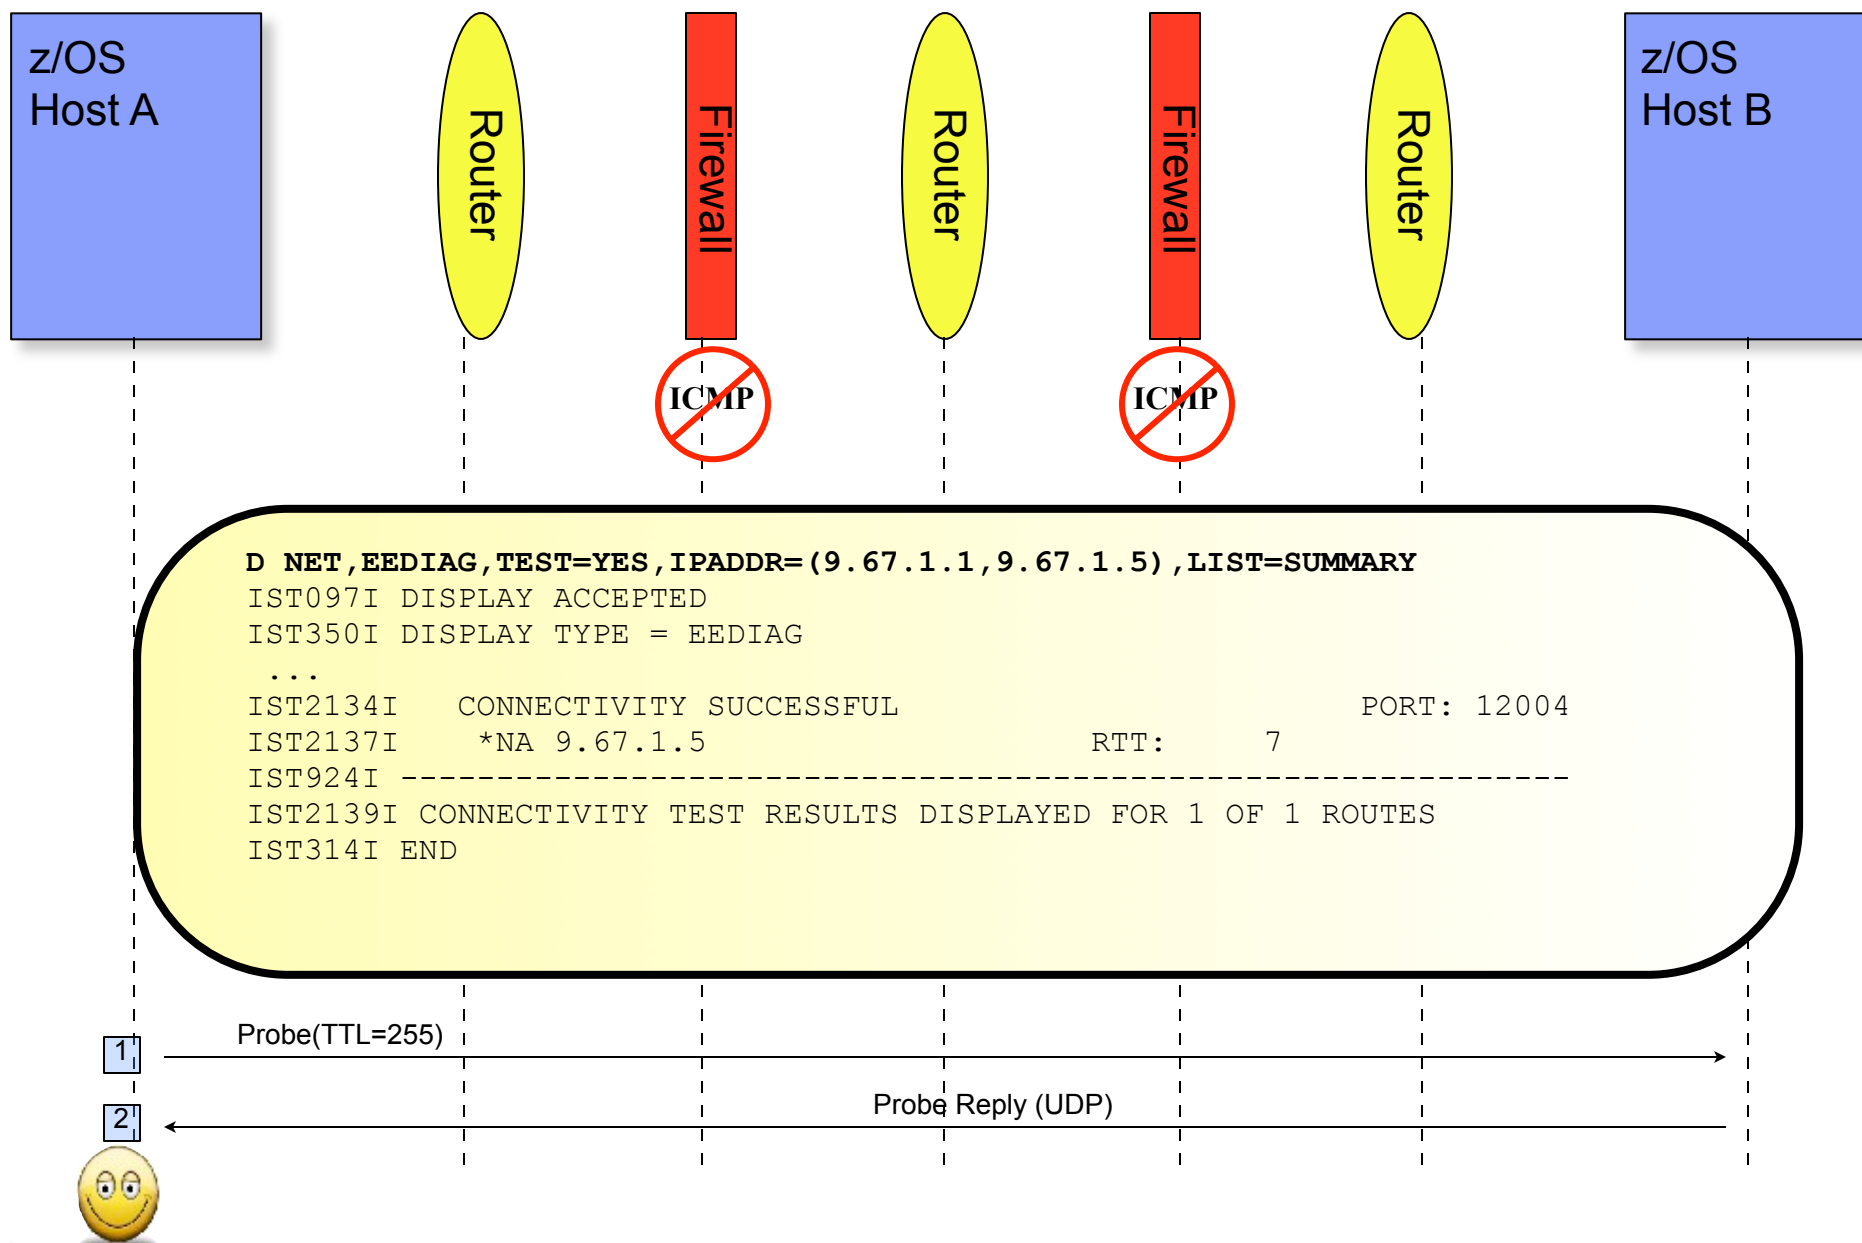
The MAXTIME operand on the EE connectivity test command is used to specify how long VTAM will spend performing the connectivity test before terminating the test, and displaying the available information. The default value of MAXTIME is 60 seconds.

Firewall-Friendly EE Connectivity Test (V1R13)

- V1R13 introduced a new “firewall-friendly” form of the EE Connectivity Test
 - The “DISPLAY EEDIAG,TEST=YES,LIST=SUMMARY” command will quickly verify partner reachability
 - The TTL is set to the maximum hop limit so that the partner can quickly receive the Probe command and generate a UDP Probe reply packet.
 - Intermediate hop analysis is not possible

```
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: 12004
IST2137I   *NA 9.67.1.5                      RTT:      7
IST924I -----
IST2139I CONNECTIVITY TEST RESULTS DISPLAYED FOR 1 OF 1 ROUTES
IST314I END
```

Firewall-Friendly EE Connectivity Test (V1R13) ...



EE Health Verification (V1R12)



- The EE Health Verification function will verify the health of a potential EE connection by sending a probe to the remote partner using all five ports during the connection activation
 - VTAM does not activate the EE connection if the remote partner is not reachable on all ports
 - However, if the remote partner does not support the probe, VTAM will still bring up the EE connection
- EE Health Verification will also optionally verify the health of an active EE connection by sending a probe to the remote partner on all five ports at a user-specified interval
 - VTAM issues a warning message if the remote partner is not reachable on all ports, but will keep the connection active
- EE Health Verification is enabled by the EEVERIFY start option:

```

__EEVERIFY= _ACTIVATE__
>>_|_____|____<<
    |_EEVERIFY= __ _NEVER_|
    | _ACTIVATE_|
    |_time_interval_value_|

```

- Or the EEVERIFY GROUP/PU parameter:

```

>>_____|____<<
    |_EEVERIFY= __ _NEVER_|
    | _ACTIVATE_|
    |_time_interval_value_|

```

- The time_interval_value can range from 1-1440 minutes

EE Health Verification ...



- When EE Health verification fails for an active EE connection, VTAM issues highlighted warning message IST2323E if it is not already present:

IST2323E EE HEALTH VERIFICATION FAILED ON ONE OR MORE CONNECTIONS

- This message stays on the console until the condition is cleared or the message is erased by the operator
- DISPLAY EE,LIST=VERIFY is used to determine which connection(s) have failed verification:

```
d net,ee,list=eeverify
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EE
IST2000I ENTERPRISE EXTENDER GENERAL INFORMATION
IST1685I TCP/IP JOB NAME = TCPCS
IST2003I ENTERPRISE EXTENDER XCA MAJOR NODE NAME = XCAIP
...
IST924I -----
IST2324I EE HEALTH VERIFICATION: FAILED CONNECTION INFORMATION
IST2325I LINE LNIP1 PU SWIP2A1 ON 12/21/09 AT 15:56:39
IST2326I EE HEALTH VERIFICATION TOTAL CONNECTION FAILURES = 1
IST2017I TOTAL RTP PIPES =          1          LU-LU SESSIONS =          2
IST2018I TOTAL ACTIVE PREDEFINED EE CONNECTIONS              =          1
IST2019I TOTAL ACTIVE LOCAL VRN EE CONNECTIONS               =          0
IST2020I TOTAL ACTIVE GLOBAL VRN EE CONNECTIONS              =          0
IST2021I TOTAL ACTIVE EE CONNECTIONS                         =          1
IST314I END
```

EE Health Verification ...



- Using DISPLAY EE to display an individual EE connection will also provide information on the success or failure of EE Health Verification:

```
d net,ee,id=SWIP2A1
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EE
IST2001I ENTERPRISE EXTENDER CONNECTION INFORMATION
IST075I NAME = SWIP2A1, TYPE = PU_T2.1
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST1680I REMOTE IP ADDRESS 9.67.1.2
IST2022I EE CONNECTION ACTIVATED ON 12/21/09 AT 16:21:57
IST2114I LIVTIME:      INITIAL =   10    MAXIMUM =   0    CURRENT =   10
IST2023I CONNECTED TO LINE LNIP1
IST2327I EE HEALTH VERIFICATION OPTION - EEVERIFY = 2 MINUTES
IST2329I EE HEALTH VERIFICATION SUCCESSFUL ON 12/21/09 AT 16:37:21
IST2341I EE HEALTH VERIFICATION HAS NEVER FAILED FOR THIS CONNECTION
IST2025I LDLC SIGNALS RETRANSMITTED AT LEAST ONE TIME      =      0
IST2026I LDLC SIGNALS RETRANSMITTED SRQRETRY TIMES        =      0
...
IST314I END
```

- If a connection is failing EE Health Verification, the DISPLAY EEDIAG,TEST=YES command can be used to further diagnose the cause of the failure.

Subarea SNA



Enterprise Extender

The State of SNA



The State of SNA - Challenges in Converting to IP

Subarea SNA




Enterprise Extender

- Economic challenges
 - Questionable cost/benefit ratio
 - Other projects more important
- Risk in altering long-established business processes
- Challenges when converting existing SNA applications to sockets applications.
 - Declining skills to participate in the conversion project
 - Typically not available to be purchased off-the-shelf
 - For applications that can be purchased, additional customization likely needed
 - Need to find good sockets programmers, which can be a challenge
 - Most known large-scale conversions have taken years
- Once the applications are converted to sockets, there are additional challenges:
 - Change in monitoring/management when moving from an LU-basis to an IP address/port basis
 - Affects tooling and processes
 - Affects overall process




2011 Customer Survey - SNA Questions/Responses

SNA Connectivity Used



Five customers still had 37x5s, with one customer having 16, two having two each, and the other two having one each.

Date SNA No Longer Needed



Subarea SNA



Enterprise Extender



IBM SNA Statements of Direction

2002:

It is IBM's intent to support VTAM in z/OS Communications Server for the foreseeable future. Customers have a substantial investment in 3270 and SNA applications. We continue to support and enhance VTAM's capabilities while integrating it with new technologies. IBM has no plans at this time to discontinue SNA support in z/OS Communications Server.

2004:

IBM's plans to support SNA workloads have not changed since the Statement of Direction made in 2002. As of June 2004, customers can -- for selected SNA workloads -- use Communications Server products for Linux®, Linux on IBM® zSeries®, Microsoft® Windows®, and AIX® to replace some of the old SNA infrastructure components, such as the IBM 3745/46 or other channel-attached SNA controllers. z/OS® Communications Server can replace some (SNA Network Interconnect) SNI workloads using Enterprise Extender and Extended Border Node functions.

It is IBM's intent to introduce an additional solution in 2005 that uses NCP software running within Linux on zSeries. The intent is to provide a migration path for customers who use traditional SNA (including SNA Network Interconnect (SNI)) to communicate with their Business Partners. This solution can allow them to continue using traditional SNA without a dependency on IBM 3745 and 3746 Communications Controller hardware.

This statement represent current intentions of IBM. Any reliance on this statement of direction is at the relying party's sole risk and will not create any liability or obligation for IBM. All statements regarding IBM's plans, directions, and intent are subject to change or withdrawal without notice.

<http://www-01.ibm.com/software/network/commserver/snasupport.html>

These statements of direction are old, but still relevant!

Subarea SNA



Enterprise Extender

References



References

- Redbooks

- SG24-5957-00 Migrating Subarea to an IP Infrastructure
- SG24-7359-00 Enterprise Extender Implementation Guide
- SG24-7334-00 A Structured Approach to Modernizing the SNA Environment

- Screencasts

- APPN Configurations: Recommendations & Limitations:
 - ▶ <http://www.youtube.com/zoscommserver#p/a/u/0/TC1gaiARPgM>
- APPN Logmodes and Class of Service:
 - ▶ <http://www.youtube.com/zoscommserver#p/u/14/-rPxj2ImP-Y>
- Practical Guide to Optimizing APPN and Extended Border Node Searches:
 - ▶ <http://ibm.co/mxWyE3>
- IBM Remote API Client:
 - ▶ <http://www.youtube.com/watch?v=h9K057ujZBs>

Subarea SNA



Enterprise Extender

References ...

- Prior SHARE sessions

- Searching in Mixed APPN/Subarea Networks
 - ▶ http://proceedings.share.org/client_files/SHARE_in_Austin/S3618JH145212.pdf
- APPN Logmodes and Class of Service
 - ▶ Presentation: http://proceedings.share.org/client_files/SHARE_in_Austin/S3620JH145658.pdf
 - ▶ Script: http://proceedings.share.org/client_files/SHARE_in_Austin/S3620JH145712.pdf
- APPN Configurations: Recommendations & Limitations
 - ▶ http://proceedings.share.org/client_files/SHARE_in_Austin/S3608JH144707.pdf
- Searching and Security in APPN/HPR Border Node Networks (Parts 1 and 2)
 - ▶ http://proceedings.share.org/client_files/SHARE_in_Austin/S3615JH145455.pdf
- Enterprise Extender: Implementing Connection Network
 - ▶ http://proceedings.share.org/client_files/SHARE_in_Austin/S3602SR224007.pdf
- SNA Security Considerations
 - ▶ http://proceedings.share.org/client_files/SHARE_in_Austin/S3612RW083850.pdf
- SNA 101: Basic VTAM, APPN, and EE Concepts:
 - ▶ http://proceedings.share.org/client_files/SHARE_in_San_Jose/S3431SR132942.pdf
- Diagnosing Enterprise Extender Problems
 - ▶ http://proceedings.share.org/client_files/SHARE_in_San_Jose/S3611MB092402.pdf
- It's Gr-EE-k to Me! What Do All those Enterprise Extender Messages Mean?
 - ▶ http://proceedings.share.org/client_files/SHARE_in_Orlando/S3618GD171929.pdf



Please Complete Session Evaluation

- Enterprise Extender on z/OS CS: Hints and Tips
- Session # 14679
- QR Code:



Find us on Facebook at
<http://www.facebook.com/IBMCommserver>



Follow us on Twitter at
http://www.twitter.com/IBM_Commserver



Read the z/OS Communications Server blog at
<http://tinyurl.com/zoscsblog>



Visit the z/OS CS YouTube channel at
<http://www.youtube.com/user/zOSCommServer>

Appendix




EE Connection Network Tips

- A limited set of characteristics of the dynamic connection can be customized by coding a DYNTYPE=VN PU entry in the model major node:


```
MODEL A1A VBUILD TYPE=MODEL  
*  
VNMODEL PU DYNTYPE=VN,  
DISCNT=NO
```

- The primary use of the DYNTYPE=VN model is to customize the DISCNT value.
- TG characteristics cannot be specified on a DYNTYPE=VN model PU
 - The DYNTYPE=VN model defines the dynamic PU used for the actual VN connection, not the virtual routing node (VRN) and its associated TGs. The TG created with the dynamic PU is not reported to Topology and Routing Services and is not involved in route calculations. To configure the weight of the TGs associated with the VRN, specify the TG characteristics on the GROUP statements in the XCA major node.
- Consider coding a DYNTYPE=VN model, with DISCNT=NO, or a delay value of 60+ seconds.
 - **Important note for CICS LU6.2 Users:** Specifying DISCNT=NO will prevent CICS from terminating its sessions at the end of every transaction.
- For any APPN TG to be activated, whether predefined or dynamically-activated (as is usually the case with connections associated with connection networks), an ADJCP entry must exist. If no ADJCP entry is coded for the partner CP, and DYNADJCP=NO, then the APPN connection activation will fail. So, to allow connection network connectivity you must set DYNADJCP=YES, or code ADJCP definitions for the partner CPs.
 - This provides a mechanism to control which partners are allowed to connect to your CP




Path MTU Discovery for EE

- Path MTU Discovery (PMTU) for EE:
 - IPv4: Stacks set "don't fragment" (DF) bit in all outbound IPv4 EE packets
 - IPv6 does not support fragmentation in the network
 - Stack monitors for ICMP/ICMPv6 "Packet too big" messages
 - Stack updates VTAM with learned path MTU
 - Local RTP pipes segment to new size
- PMTU originally architected as TCP-based solution and did not apply to UDP
 - PMTU for EE is an EE-specific adaptation of PMTU, and does not apply to UDP in general



Path MTU Discovery for EE ...

- Path MTU discovery requires ICMP messages to be returned
 - Some firewalls are configured to block ICMP
 - Minimally configure firewalls to permit the following ICMP messages for ports 12000-12004:
 - IPv4: ICMP Message Type 3 - Destination Unreachable
 - IPv6: ICMPv6 Message Type 2 - Packet Too Big
 - Do NOT enable PMTU discovery if ICMP messages are not permitted through the firewalls



- For situations where enabling PMTU discovery for EE is not viable, or is not an adequate solution, the MTU operand allows static configuration of the MTU size to be used for the EE connection:
 - $MTU = mtu_size$
 - Specifiable on EE switched PU, EE XCA GROUP (connection network), and EE model PU