



---

**Communications Server z/OS® V1R5 and V1R6 Technical Update**

---

# **Enterprise Extender and SNA**

in z/OS V1R6 CS

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

---



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

## Topics

### ➤ z/OS V1R6

- ┆ Vary update for EE XCA major node
- ┆ D EE command
- ┆ EE connection network reachability awareness
- ┆ D RTP command enhancements
- ┆ EBN session awareness
- ┆ Cleaning up "hung" APPN search
- ┆ Display CSM usage
- ┆ EE packet trace formatting
  
- EE task force initiatives
  - ┆ LDLC timer optimization
  - ┆ RTP queueing optimization
  - ┆ EE security



# Enterprise Extender and SNA in z/OS V1R6 CS

Copyright International Business Machines Corporation 2004. All rights reserved.



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

## VARY Command Enhancements for EE XCA



- V1R5 increases EE flexibility with support for multiple VRNs, multiple VIPAs, IPv6, and NAT compatibility. With the flexibility comes additional complexity, and this often leads to the need to add or change definitions.
- However, currently those definitions cannot be changed or augmented once the XCA major node is active, without inactivating the major node, thereby **disrupting all existing Enterprise Extender connections**.
- V1R6 improves usability by allowing the UPDATE operand on the VARY command for the EE XCA major node, thereby allowing adding of GROUPs and changing of operand values, without bouncing the major node.
  - ⌋ However, note that a GROUP must itself be inactive before its operands can be changed.
  - ⌋ To simplify the inactivation of a GROUP prior to a change (or activation after a change), V1R6 also allows a VARY ACT (or VARY INACT) command to be issued against an EE XCA GROUP, thereby activating (or inactivating) the GROUP and all subordinate LINES/PUs.

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

# VARY Command Enhancements...



```

d net, topo, orig=sscpla, dest=sscpvn
...
IST1302I          CAPACITY PDELAY  COSTTIME  COSTBYTE
IST1303I          10M      NEGLIGIB 0          0
...
<<< Edit XCA major node to change CAPACITY of SSCPVN to 1000M >>>
v net, act, id=xcaip, update=all
IST097I VARY ACCEPTED
IST886I VARY ACT XCAIP CHANGE GPIP FAILED
IST523I REASON = INVALID RESOURCE CURRENT STATE
IST314I END
IST886I VARY ACT XCAIP CHANGE GPIPVN FAILED
IST523I REASON = INVALID RESOURCE CURRENT STATE
IST314I END
IST093I XCAIP ACTIVE
v net, inact, id=gpipvn
IST097I VARY ACCEPTED
IST1225I VIRTUAL NODE NETA.SSCPVN CONNECTION INACTIVE
IST105I GPIPVN NODE NOW INACTIVE
v net, act, id=xcaip, update=all
IST097I VARY ACCEPTED
IST886I VARY ACT XCAIP CHANGE GPIP FAILED
IST523I REASON = INVALID RESOURCE CURRENT STATE
IST314I END
IST093I XCAIP ACTIVE
v net, act, id=lnipvn
IST097I VARY ACCEPTED
IST1168I VIRTUAL NODE NETA.SSCPVN CONNECTION ACTIVE
IST093I LNIPVN ACTIVE
d net, topo, orig=sscpla, dest=sscpvn
...
IST1579I
IST1302I          CAPACITY PDELAY  COSTTIME  COSTBYTE
IST1303I          944M      TERRESTR 0          0
    
```

XCAIP	VBUILD	TYPE=XCA
PORTIP	PORT	MEDIUM=HPRIP, ...
...		
GPIPVN	GROUP	VNNAME=NETA.SSCPVN, ..., TCP-ETHERNET
		<b>CAPACITY=1000M</b>
LNIPVN	LINE	
GPIPVN	PU	

SSCPVN has a CAPACITY of 10M

The update fails because the changed GROUP (GPIPVN) is active

After inactivating the GROUP, the update is successful

The CAPACITY has been updated.

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

## VARY Command Enhancements Notes



➤ The following operands in the EE XCA major node can be changed via UPDATE:

ANSWER	AUTOGEN
CALL	CAPACITY
COSTBYTE	COSTTIME
DYNPU	DYNPUPFX
DYNVNPFX	HOSTNAME
IPADDR	ISTATUS
KEEPACT	PDELAY
SECURITY	SWNORDER
TGP	UNRCHTIM
UPARM2	UPARM2
UPARM3	VNNAME
VNGROUP	VNTYPE
VTAMTOPO	

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

## Display EE Command



➤ V1R6 provides a new operator command to provide additional details about Enterprise Extender connectivity

➤ Three basic forms:

/ General information


- Basic XCA settings
- Local IP addresses and/or hostnames
- RTP pipe and LU-LU session counts
- Connection counts

/ Specific connection information

- Local IP address and/or hostname
- PU information
- LDLC information
- Data transfer statistics

/ Aggregate connection information

- Local IP address and/or hostname
- Connection counts
- Aggregate data transfer statistics



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

## D EE: General Information



D NET, EE, LIST-DETAIL

```
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 = XCAIP1A
IST2004I LIVTIME = 10 SRQTIME = 15 SRQRETRY = 3
IST2005I IPRESOLV = 0
IST924I -----
IST2006I PORT PRIORITY = SIGNAL NETWORK HIGH MEDIUM LOW
IST2007I IPPORT NUMBER = 12000 12001 12002 12003 12004
IST2008I IPTOS VALUE = C0 C0 80 40 20
IST924I -----
IST1680I LOCAL IP ADDRESS 9.67.1.5
IST2009I RTP PIPES = 2 LU-LU SESSIONS = 1
IST2010I INOPS DUE TO SRQRETRY EXPIRATION = 0
IST1324I VNNAME = IP.GVRN5 VNGROUP = GPIP5 (GLOBAL)
IST2011I AVAILABLE LINES FOR THIS EE VRN = 0
IST2012I ACTIVE CONNECTIONS USING THIS EE VRN = 1
IST2013I AVAILABLE LINES FOR PREDEFINED EE CONNECTIONS = 0
IST2014I ACTIVE PREDEFINED EE CONNECTIONS = 0
IST2015I ACTIVE LOCAL VRN EE CONNECTIONS = 0
IST2016I ACTIVE GLOBAL VRN EE CONNECTIONS = 1
IST924I -----
IST2017I TOTAL RTP PIPES = 6 LU-LU SESSIONS = 3
IST2018I TOTAL ACTIVE PREDEFINED EE CONNECTIONS = 0
IST2019I TOTAL ACTIVE LOCAL VRN EE CONNECTIONS = 0
IST2020I TOTAL ACTIVE GLOBAL VRN EE CONNECTIONS = 3
IST2021I TOTAL ACTIVE EE CONNECTIONS = 3
IST314I END
```

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



## D EE: Specific Connection Information



D NET, EE, ID=LNEE1000, LIST=DETAIL

```
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EE
IST200I ENTERPRISE EXTENDER CONNECTION INFORMATION
IST075I NAME = LNEE1000, TYPE = LINE
IST1680I LOCAL IP ADDRESS 9::67:1:1
IST1910I LOCAL HOSTNAME VIPA16.SSCP1A.TCP.RALEIGH.IBM.COM
IST1680I REMOTE IP ADDRESS 9::67:1:2
IST1909I REMOTE HOSTNAME VIPA16.SSCP2A.TCP.RALEIGH.IBM.COM
IST2022I EE CONNECTION ACTIVATED ON 08/31/03 AT TIME 14:14:26
IST2024I CONNECTED TO SWITCHED PU SNEE2A1
IST2025I LDLC SIGNALS RETRANSMITTED AT LEAST ONE TIME = 0
IST2026I LDLC SIGNALS RETRANSMITTED SRQRETRY TIMES = 0
IST2009I RTP PIPES = 4 LU-LU SESSIONS = 3
IST2027I DWINOP = NO REDIAL = *NA* REDDELAY = *NA*
IST2028I KEEPACT = NO
IST2029I MTU SIZE = 1232
IST924I -----
IST2030I PORT PRIORITY = SIGNAL
IST2036I NLPS SENT = 4 ( 000K )
IST2037I BYTES SENT = 526 ( 000K )
IST2038I NLPS RETRANSMITTED = 0 ( 000K )
IST2039I BYTES RETRANSMITTED = 0 ( 000K )
IST2040I NLPS RECEIVED = 4 ( 000K )
IST2041I BYTES RECEIVED = 577 ( 000K )
IST924I -----
IST2031I PORT PRIORITY = NETWORK
IST2036I NLPS SENT = 91 ( 000K )
IST2037I BYTES SENT = 7036 ( 007K )
IST2038I NLPS RETRANSMITTED = 0 ( 000K )
IST2039I BYTES RETRANSMITTED = 0 ( 000K )
IST2040I NLPS RECEIVED = 90 ( 000K )
IST2041I BYTES RECEIVED = 6693 ( 006K )
```

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

## D EE: Specific Connection Info (Continued)



```

IST924I -----
IST2032I PORT PRIORITY = HIGH
IST2036I NLPS SENT = 56390 ( 056K )
IST2037I BYTES SENT = 5184997 ( 005M )
IST2038I NLPS RETRANSMITTED = 0 ( 000K )
IST2039I BYTES RETRANSMITTED = 0 ( 000K )
IST2040I NLPS RECEIVED = 56394 ( 056K )
IST2041I BYTES RECEIVED = 5185656 ( 005M )
IST924I -----
IST2033I PORT PRIORITY = MEDIUM
IST2036I NLPS SENT = 0 ( 000K )
IST2037I BYTES SENT = 0 ( 000K )
IST2038I NLPS RETRANSMITTED = 0 ( 000K )
IST2039I BYTES RETRANSMITTED = 0 ( 000K )
IST2040I NLPS RECEIVED = 0 ( 000K )
IST2041I BYTES RECEIVED = 0 ( 000K )
IST924I -----
IST2034I PORT PRIORITY = LOW
IST2036I NLPS SENT = 0 ( 000K )
IST2037I BYTES SENT = 0 ( 000K )
IST2038I NLPS RETRANSMITTED = 0 ( 000K )
IST2039I BYTES RETRANSMITTED = 0 ( 000K )
IST2040I NLPS RECEIVED = 0 ( 000K )
IST2041I BYTES RECEIVED = 0 ( 000K )
IST924I -----
IST2035I TOTALS FOR ALL PORT PRIORITIES
IST2036I NLPS SENT = 56485 ( 056K )
IST2037I BYTES SENT = 5192559 ( 005M )
IST2038I NLPS RETRANSMITTED = 0 ( 000K )
IST2039I BYTES RETRANSMITTED = 0 ( 000K )
IST2040I NLPS RECEIVED = 56488 ( 056K )
IST2041I BYTES RECEIVED = 5192926 ( 005M )
IST314I END
    
```

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

## D EE: Aggregate Connection Information



D NET,EE,HOSTNAME=VIP16.SSCP1A.TCP.RALEIGH.IBM.COM

```
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EE
IST2002I ENTERPRISE EXTENDER AGGREGATE CONNECTION INFORMATION
IST924I -----
IST1680I LOCAL IP ADDRESS 9::67:1:1
IST1910I LOCAL HOSTNAME VIP16.SSCP1A.TCP.RALEIGH.IBM.COM
IST2009I RTP PIPES =          7          LU-LU SESSIONS =          6
IST2010I INOPS DUE TO SQRETRY EXPIRATION =          0
IST2013I AVAILABLE LINES FOR PREDEFINED EE CONNECTIONS =          9
IST2014I ACTIVE PREDEFINED EE CONNECTIONS =          1
IST2015I ACTIVE LOCAL VRN EE CONNECTIONS =          0
IST2016I ACTIVE GLOBAL VRN EE CONNECTIONS =          0
IST2044I TOTAL ACTIVE EE CONNECTIONS FOR LOCAL IPADDR =          1
IST924I -----
IST2035I TOTALS FOR ALL PORT PRIORITIES
IST2036I NLPS SENT =          78877 ( 078K )
IST2037I BYTES SENT =          8671229 ( 008M )
IST2038I NLPS RETRANSMITTED =          0 ( 000K )
IST2039I BYTES RETRANSMITTED =          0 ( 000K )
IST2040I NLPS RECEIVED =          78923 ( 078K )
IST2041I BYTES RECEIVED =          7951928 ( 007M )
IST314I END
```

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

## EE Connection network reachability awareness background information




- A connection network is a representation of a shared access transport facility (SATF) that enables nodes identifying their connectivity to the SATF by a common virtual routing node (VRN), which allows communication without having individually defined connections to one another.
  - ⌘ Connection network technology was originally introduced when LANs became popular with many SNA nodes attached to the same LAN
  - ⌘ Instead of pre-defining links to all other SNA nodes on the LAN, each node predefined a link to a single virtual node
  
- An Enterprise Extender connection is defined in an External Communication Adapter (XCA) major node
  - ⌘ PORT definition statement defines the VTAM connection to the Enterprise Extender SATF
  - ⌘ MEDIUM=HPRIP is required on the PORT definition statement for Enterprise Extender
  
- An Enterprise Extender connection network is defined on the PORT definition statement or the GROUP definition statement
  - ⌘ VNNAME is the CPNAME for the connection network. This name is reported to network topology as a virtual node (VRN) and is treated as an adjacent CP to this node. TGs to and from the VRN are also reported to network topology.
  - ⌘ VNTYPE specifies whether an Enterprise Extender Connection Network is permitted to span network boundaries.
  - ⌘ VNGROUP (on a PORT definition statement only) specifies the name of the GROUP containing dial-out links available for use on the connection network.

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

## EE Connection network reachability awareness background information



- When a session is established or when a path switch occurs, Topology and Routing Services (TRS) selects the session path.
- Normally the network node on which the primary LU (PLU) resides, or the NNS of the end node on which the PLU resides, determines the optimal route from the PLU node to the secondary LU (SLU) node.
- The optimal route is calculated by TRS based on weight values assigned to nodes and TGs. The total weight of a route includes the weights of all TGs and intermediate nodes between the two endpoints of a route.
- Weights are determined by characteristics assigned to each node and each TG, and the APPN Class of Service (CoS) associated with the session.
  - ⌋ The route resistance node characteristic is defined in VTAM by the ROUTERES start option.
    - A virtual routing node has a weight of zero.
  - ⌋ TG characteristics are defined on the PORT definition statement or the GROUP definition statement in an Enterprise Extender XCA major node for a TG to a connection network, or on the switched PU definition statement for a connection that is not connection network.
    - Examples of TG characteristics are CAPACITY, PDELAY, SECURITY, and TGP.



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

## EE Connection network reachability awareness background information



- When a dial failure or a connection INOP occurs for a connection over an Enterprise Extender connection network, the partner node cannot be reached using the connection network path across the virtual node (VRN). This connection network path might have been chosen for this connection because it had a lower weight than any alternate path available at one of the following times:
  - the time of this failing dial
  - the time of the dial that set up the existing connection
  - the time of a path switch to this connection network path for an existing connection
- If this path still has the lowest weight of any available path to the partner node, any attempt to redial the partner node will continue to try the path over this particular VRN, which is likely to result in failure until the underlying problem with the path is corrected.
- This occurs because, with connection network, Topology and Routing Services (TRS) has knowledge of the virtual path, but no knowledge of the underlying physical connection.
- APPN Architecture does not provide a way to inform TRS when a dial failure occurs for session setup over the physical path, or when a station INOP occurs over the physical path, breaking an existing connection. The virtual path through the connection network is still considered by TRS to be available.
- z/OS V1R5 Communications Server included an enhancement to issue a more focused failure message for these cases, so that it is obvious that the dial failure or INOP has occurred on the physical connection underlying a connection network, and it is clear what VRN and partner node are involved. The customer can then take action based on that message. This improved the situation, but still put the onus on the customer to design automation to deal with the situation.

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

## EE Connection network reachability awareness Connection Network Failures



### ➤ Typical VRN Configuration

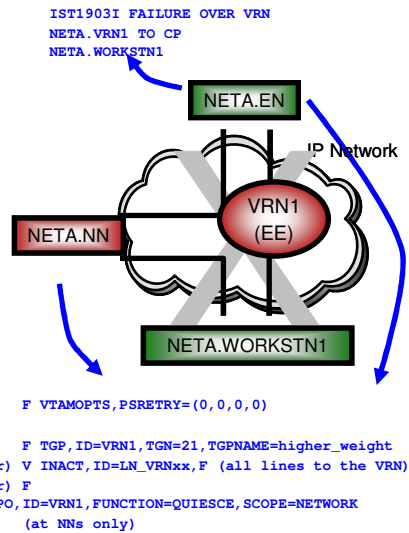
- ENs predefine links to NNS for CP-CP
- Dynamic CN links used EN-to-EN

### ➤ What happens if IP network fails?

- Affected RTPs begin path switching
- VRN path is chosen again
  - Topology component is \*not\* notified of the failure
  - VRN path still has lowest weight
- Path switch fails even though a functional alternate path exists (through NN)

### ➤ V1R5 Will Issue New Message

- For VRN dial failures and VRN link INOPs
  - IST1903I - Identifies the VRN and partner node
- Allows network operator to take action
  - Disable PSRETRY
  - Prevent VRN from being used for new sessions
  - Re-enable both after IP network is restored
- Similar actions may be required at the partner node and/or at an EN's NNS



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

## EE Connection network reachability awareness



- 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.
- The period of time that a path through the EE VRN to the unreachable partner will remain unavailable is configurable.
- 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).

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



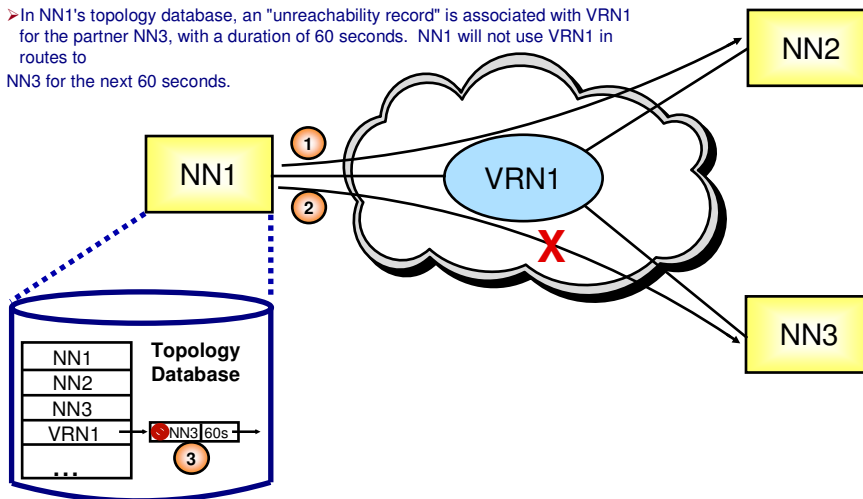


- Once the unreachability period expires, a subsequent attempt to use the VRN will result in failure if the underlying problem with the connection has not been corrected. The new dial failure will again prevent selection of the path through this VRN to the unreachable partner node for the same period of time. This will continue until the problem with the connection path is corrected.
- It is possible for a path between two nodes, through the VRN, to be usable for route selection in one direction but not the other. This function can detect that distinction and will allow routing in the direction that is usable while preventing the path in the direction that is not usable from being selected when new sessions are established and HPR path switches occur.
- A new control vector has been architected to carry unreachable partner information on TDUs.

## EE Connection network reachability awareness



- NN1 successfully contacts NN2 across VRN1.
- NN1's attempt to contact NN3 across VRN1 fails.
- In NN1's topology database, an "unreachability record" is associated with VRN1 for the partner NN3, with a duration of 60 seconds. NN1 will not use VRN1 in routes to NN3 for the next 60 seconds.



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

## EE Connection network reachability awareness Controls



- An 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 unreachability durations for connection networks of different characteristics, or to different business partners, for example.
- Current unreachability information can be displayed by using a new flavor of the DISPLAY TOPO command.
- Unreachability records may be cleared by using a new flavor of the MODIFY TOPO command.
- A new message is now issued with IST1903I (introduced in V1R5) to convey unreachability status at time of failure:

```
IST1903I FAILURE OVER VRN NETA.VRNI TO CP NETA.SSCPAA  
IST2050I THIS PATH WILL NOT BE SELECTED FOR UNRCHTIM = 300 SECONDS  
IST314I END
```

## EE Connection Network Reachability Awareness Controls



- Range for UNRCHTIM is 0, or 10-65535 seconds. The default value is 0, indicating that paths through EE connection networks will always be considered for routing.
- A minimum value of 10 seconds is set for UNRCHTIM to reduce the risk of performance concerns due to excess TDU flows.
- You can change the value of the UNRCHTIM start option with the MODIFY VTAMOPTS command while VTAM is running. The UNRCHTIM for existing paths through connection networks to partner nodes that were determined to be unreachable before the start option value was modified will not be changed.
- If UNRCHTIM is specified on the PORT or GROUP, then VNNAME or VNTYPE must also be specified on that PORT or GROUP.
- You can dynamically change the UNRCHTIM parameter on the PORT or GROUP by editing the VTAMLST member and then issuing the VARY ACT,UPDATE=ALL command.

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

## EE Connection Network Reachability Awareness D TOPO,LIST=UNRCHTIM



➤ To display unreachable partner information for a specific EE VRN:

```
d net,topo,id=vrn1,list=unrchtim
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TOPOLOGY
IST2057I UNREACHABLE PARTNER INFORMATION:
IST924I -----
IST2051I VIRTUAL NODE BETWEEN ORIGIN AND PARTNER - NETA.VRN1
IST2052I   ORIGIN NODE           PARTNER NODE      UNRCHTIM  EXPIRES
IST2055I NETA.SSCP1A             NETA.SSCPAA       300S      24:15:42
IST2055I NETA.TEST1             NETWORK1.TEST4444 14500S    01:15:24
IST2055I NETA.SSCP1A             NETWORKB.SSCP7B   45S       23:18:19
IST2055I NETA.SSCP2A             NETA.SSCPAA       780S      24:24:10
IST2055I NETA.TEST1234          NETA.SSCPAA       300S      24:16:59
IST924I -----
IST314I END
```

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

## EE Connection Network Reachability Awareness D TOPO,LIST=UNRCHTIM...



➤ To display all unreachable partner information known in this node:

```
d net,topo,list=unrchtim
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TOPOLOGY
IST2057I UNREACHABLE PARTNER INFORMATION:
IST924I -----
IST2051I VIRTUAL NODE BETWEEN ORIGIN AND PARTNER - NETA.VRN1
IST2052I  ORIGIN NODE          PARTNER NODE      UNRCHTIM  EXPIRES
IST2055I NETA.SSCP1A          NETA.SSCPAA        300S     24:15:42
IST2055I NETA.TEST1          NETWORK1.TEST4444  14500S   01:15:24
IST2055I NETA.SSCP1A          NETWORKB.SSCP7B    45S      23:18:19
IST2055I NETA.SSCP2A          NETA.SSCPAA        780S     24:24:10
IST2055I NETA.TEST1234       NETA.SSCPAA        300S     24:16:59
IST924I -----
IST2053I ORIGIN END NODE - NETA.ENDNODE1
IST2054I  VIRTUAL NODE          PARTNER NODE      UNRCHTIM  EXPIRES
IST2055I NETA.VIRTUAL1        NETA.SSCPAA       3000S    24:04:15
IST2055I NETB.VIRTUAL2        NETA.SSCPAA       65535S   02:54:02
IST2055I NETA.VIRTUAL3        NETWORK1.TEST4444  65535S   02:32:34
IST924I -----
IST2051I VIRTUAL NODE BETWEEN ORIGIN AND PARTNER - NETB.GVRN2
IST2052I  ORIGIN NODE          PARTNER NODE      UNRCHTIM  EXPIRES
IST2055I NETA.TEST1234       NETA.SSCPAA        300S     24:16:59
IST924I -----
IST314I END
```

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

## Enhanced RTP Display Output



➤ V1R6 adds yet more information to the detailed RTP PU output, and also improves the organization of the output:

```
d net,id=cnr00004,hprdiag=yes
IST097I  DISPLAY ACCEPTED
IST075I  NAME = CNR00001, TYPE = PU_T2.1
IST1392I DISCNTIM = 00010 DEFINED AT PU FOR DISCONNECT
IST486I  STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I CP NAME = SSCP2A, CP NETID = NETA, DYNAMIC LU = YES
IST1589I XNETALS = YES
IST231I  RTP MAJOR NODE = ISTRTPMN
IST654I  I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST1964I APPNCOS = #INTER - PRIORITY = MEDIUM
IST1476I TCID X'3A06F08600000011' - REMOTE TCID X'3A0724A100000013'
IST1481I DESTINATION CP NETA.SSCP2A - NCE X'D400000000000000'
IST1587I ORIGIN NCE X'D400000000000000'
IST1967I ACTIVATED AS PASSIVE ON 05/26/03 AT 08:40:21
IST1479I RTP CONNECTION STATE = CONNECTED - MNPS = NO
IST1959I DATA FLOW STATE: NORMAL
IST1855I NUMBER OF SESSIONS USING RTP = 1
IST1480I RTP END TO END ROUTE - RSCV PATH
IST1460I TGN  CPNAME          TG TYPE      HPR
IST1461I  21  NETA.SSCP2A      APPN        RTP
IST875I  ALSNAME TOWARDS RTP = AHHCPU1
IST1738I ANR LABEL          TP          ER NUMBER
IST1739I 8001000A00000000  *NA*      *NA*
IST924I -----
```

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

## Enhanced RTP Display Output...



```
IST1968I ARB INFORMATION:
IST1844I ARB MODE = GREEN
IST1697I RTP PACING ALGORITHM = ARB RESPONSIVE MODE
IST1477I ALLOWED DATA FLOW RATE = 3200 KBITS/SEC
IST1516I INITIAL DATA FLOW RATE = 1600 KBITS/SEC
IST1841I ACTUAL DATA FLOW RATE = 10 KBITS/SEC
IST1969I MAXIMUM ACTUAL DATA FLOW RATE = 3200 KBITS/SEC
IST1862I ARB MAXIMUM SEND RATE = 32000 KBITS/SEC
IST1846I CURRENT RECEIVER THRESHOLD = 37000 MICROSECONDS
IST1846I MAXIMUM RECEIVER THRESHOLD = 37000 MICROSECONDS
IST1846I MINIMUM RECEIVER THRESHOLD = 17000 MICROSECONDS
IST1970I RATE REDUCTIONS DUE TO RETRANSMISSIONS = 0
IST924I -----
IST1971I TIMER INFORMATION:
IST1852I LIVENESS TIMER = 10 SECONDS
IST1851I SMOOTHED ROUND TRIP TIME = 519 MILLISECONDS
IST1972I SHORT REQUEST TIMER = 500 MILLISECONDS
IST924I -----
IST1973I OUTBOUND TRANSMISSION INFORMATION:
IST1974I NUMBER OF NLPs SENT = 12 (OK)
IST1975I TOTAL BYTES SENT = 1823 (1K)
IST1849I LARGEST NLP SENT = 161 BYTES
IST1980I SEQUENCE NUMBER = 372 (X'00000174')
IST1842I NUMBER OF NLPs RETRANSMITTED = 0
IST1976I BYTES RETRANSMITTED = 0 (OK)
IST1478I NUMBER OF UNACKNOWLEDGED BUFFERS = 0
IST1958I NUMBER OF ORPHANED BUFFERS = 0
```

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



## Enhanced RTP Display Output...



```
IST1843I NUMBER OF NLPs ON WAITING-TO-SEND QUEUE = 0
IST1847I NUMBER OF NLPs ON WAITING-FOR-ACKNOWLEDGEMENT QUEUE = 0
IST1977I MAXIMUM NUMBER OF NLPs ON WAITING-FOR-ACK QUEUE = 6
IST1978I WAITING-FOR-ACK QUEUE MAX REACHED ON 05/26/03 AT 13:44:03
IST1511I MAXIMUM NETWORK LAYER PACKET SIZE = 16410 BYTES
IST924I -----
IST1979I INBOUND TRANSMISSION INFORMATION:
IST2059I NUMBER OF NLPs RECEIVED = 13 (OK)
IST1981I TOTAL BYTES RECEIVED = 1488 (1K)
IST1850I LARGEST NLP RECEIVED = 137 BYTES
IST1980I SEQUENCE NUMBER = 204 (X'000000CC')
IST1853I NUMBER OF NLPs ON OUT-OF-SEQUENCE QUEUE = 0
IST1854I NUMBER OF NLPs ON INBOUND SEGMENTS QUEUE = 0
IST1982I NUMBER OF NLPs ON INBOUND WORK QUEUE = 5
IST1983I MAXIMUM NUMBER OF NLPs ON INBOUND WORK QUEUE = 8
IST924I -----
IST1984I PATH SWITCH INFORMATION:
IST1856I LAST PATH SWITCH OCCURRENCE WAS ON 05/26/03 AT 12:53:33
IST1937I PATH SWITCH REASON: INITIATED BY REMOTE PARTNER
IST1985I PATH SWITCHES INITIATED FROM REMOTE RTP = 4
IST1986I PATH SWITCHES INITIATED FROM LOCAL RTP = 2
IST1987I PATH SWITCHES DUE TO LOCAL FAILURE = 0
IST1988I PATH SWITCHES DUE TO LOCAL PSRETRY = 2
IST924I -----
IST1857I BACKPRESSURE REASON COUNTS:
IST1858I PATHSWITCH SEND QUEUE MAX STORAGE FAILURE STALLED PIPE
IST1859I      2          0          0          0
IST924I -----
IST314I END
```

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

## DISPLAY RTPS by TCID



- When attempting to diagnose a possible problem with a specific RTP connection, customers often find it necessary to monitor the status and/or performance of a given RTP connection from "both ends".
- There is currently no easy way for customers to correlate the RTP (CNRxxxx) PU name on one side of an RTP connection to the corresponding RTP PU name on the other side of that same RTP connection.
  - ∫ The only information associated with an RTP connection that can be used to correlate the local RTP PU name to the corresponding remote RTP PU name on the other side of the RTP connection is the "Remote TCID" (Transport Connection Identifier).
  - ∫ The Local and Remote TCIDs for an RTP connection are provided in the output of the DISPLAY ID=CNRxxxx command (on message IST1476I), but this requires that the RTP PU name already be known.
- The DISPLAY RTPS command is expanded to include the new TCID= operand, which allows an RTP PU to be found and displayed by its Local TCID.

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

## DISPLAY RTPS by TCID...



- 1 From the local host (NETA.SSCPAA), issue the DISPLAY ID=CNRxxxxx command and remember the value of the remote partner CP (from message IST1481I) and the Remote TCID (from message IST1476I).

```
d net,id=cnr00006
IST097I DISPLAY ACCEPTED
IST075I NAME = CNR00006, TYPE = PU_T2.1
:
IST1963I APPNCOS = #INTER - PRIORITY = HIGH
IST1476I TCID X'14AB34050001001F' - REMOTE TCID X'14AB300100010020'
IST1481I DESTINATION CP NETB.SSCPBA - NCE X'D00000000000000000'
IST1587I ORIGIN NCE X'D000000000000000000'
:
IST1480I RTP END TO END ROUTE - RSCV PATH
IST1460I TGN CPNAME TG TYPE HPR
IST1461I 21 NETA.VRNA APPN RTP
IST1461I 21 NETB.SSCPBA APPN RTP
:
IST314I END
```


- 2 From the remote host (NETB.SSCPBA), use DISPLAY RTPS,TCID= to display the corresponding RTP PU. (If desired, the TEST=YES operand can also be included on this command.)

```
d net,rtps,tcid=14AB300100010020
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = RTPS
IST1695I PU NAME CP NAME COSNAME SWITCH CONGEST STALL SESS
IST1960I CNR00007 NETA.SSCPAA #INTER NO NO NO 1
IST1454I 1 RTP(S) DISPLAYED
IST314I END
```

## EBN Session Awareness



- When HPR is used across subnet boundaries, the border nodes (EBNs) lose session awareness. This complicates authorization (DSME required instead of SME) and prevents accounting (at the EBN).
- Some EE/HPR exploiters (primarily service providers) require a method to:
  - ⌘ Control and monitor which sessions are using their EBNs as a transport
  - ⌘ Control which business partners can exploit Global VRNs (into or around their network)
- V1R6 provides controls to force back-to-back RTP pipes at the EBN, thereby restoring session awareness on the border node




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

## EBN Session Awareness...



➤ New RTPONLY operand on ADJCP definitions allows control of session awareness at APPN subnetwork boundaries

↳ Controls where EBNs are allowed to perform ANR routing and/or exploit GVRNs



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

## EBN Session Awareness...



➤ EBN Session Awareness is enforced during RTP pipe establishment

- ⌈ New function has no affect on how EBNs perform APPN searching:
  - EBN searching performed based on DS DB, BNORD, BNDYN and ADJCLUST tables
  - Search path dictates which EBNs are on the session path (except when GVRNs are used)
- ⌈ End-to-end session path (RSCV) is computed subnet-by-subnet as usual
  - Global VRN TG Vectors are discarded by Entry and Exit EBNs when ANR is not allowed
  - Forces session path to pass through the EBNs that processed the search request
- ⌈ Route\_Setup processing (after the RSCV is computed) enforces the RTPONLY value
- ⌈ Coding RTPONLY=YES does have drawbacks
  - Increased CPU and storage consumption at the EBN due to more RTP endpoints
  - May increase route setup flows
  - Back-to-back pipes make the EBN a single point of failure
- ⌈ Sample ADJCP major node with RTPONLY coded:

```
ADJBNRTP  VBUILD  TYPE=ADJCP
SSCP2A    ADJCP    NETID=NETA
SSCPAA    ADJCP    NETID=NETA
SSCPBA    ADJCP    NETID=NETA
SSCP7B    ADJCP    NETID=NETB, RTPONLY=YES
SSCP9C    ADJCP    NETID=NETC, RTPONLY=NO
```

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

## Cleaning up "hung" APPN searches



- On some occasions, an APPN search may hang. Resource Discovery search (RDS) and Management Services searches are examples of APPN searches that can hang and do not have a SIB, and thus can not be terminated by the customer.
- In some cases, setting the IOPURGE start option can clear up these hung sessions. However, there are some cases where IOPURGE does not help.
- If the hung search happens to be a Resource Discovery Search (RDS), the problem can be much more severe because any subsequent search for the same target resource will be "concentrated" behind this RDS until the RDS completes (which may never happen).
- Although hung APPN searches are not intended and (for most customers) are unlikely to occur, when it does occur, these hung searches can only be cleaned up recycling VTAM.
- The solution to this problem is an enhancement to the existing VARY TERM command.
  - ⌈ A new SCOPE= operand value (APPN) is now available on the VARY TERM command.
  - ⌈ SCOPE=APPN is only valid with the SID= flavor of the command.
  - ⌈ The D NET,SRCHINFO command may be used to display potentially hung searches, and to obtain the SID for a specific (hung) search request.
    - V NET,TERM,SCOPE=APPN,SID=F1EACB23B5FD9823
- Some searches may not be terminated by this command. But normal processing and/or IOPURGE processing should clean up those search requests.

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

## Persistent Sessions Forced Takeover Background Information - SNPS and MNPS



- VTAM Persistent Sessions function has had a long history:
- "Single Node" Persistent Sessions (SNPS) support was introduced in VTAM V3R4
  - Allowed recovery of application sessions on the same VTAM node after application failure or temporary shutdown
  - Allowed "forced" takeover of application sessions by a second image of the application
    - Any persistent application, by default, supported being taken over in this manner
- Multi Node Persistent Sessions (MNPS) was added to the product with VTAM V4R4
  - Extended the SNPS function to allow recovery of application sessions on another VTAM node within the same Sysplex
    - Application must declare its intent to perform MNPS function via the **PERSIST=MULTI** operand on the APPL definition
    - Nodes involved in MNPS must be connected to the same coupling facility data structure (commonly called **ISTMNPS**)
  - Recovery at this point was permitted only in cases of node or system failure involving the VTAM that owned the application
- MNPS was enhanced in VTAM V4R4.1 to provide Planned Takeover capability
  - Recovery was now permitted for applications that had failed or been shutdown intentionally without any corresponding node or system failure occurring
  - Still no MNPS "forced" takeover capability that was similar to SNPS processing

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



## Persistent Sessions Forced Takeover Background Information - XRF



➤ XRF (Extended Recovery Facility) has an even longer history

• Provides backup capability for an application

- All sessions are owned by the primary application image, but a backup application image has a secondary "hot standby" session
- Backup image can request that the session ownership be changed ("switched") from the primary image (primary session) to the backup image (secondary session)
- Primary image can later regain ownership of the session by "switching" them again

• Function requires 37xx/NCP to assist in establishing the primary and secondary XRF sessions, and to support session switching between these two sessions

➤ Applications that use XRF, such as IMS, have requested an alternative to XRF. XRF depends on an NCP which so far has depended on an IBM 3745/46

➤ MNPS provides similar, but not equivalent, support as XRF

• Session recovery capability allows for session movement from one application image to another, similar to what can be accomplished via XRF "switching"

• MNPS requires, however, that the "primary" (or active) application image be shutdown gracefully prior to allowing another image (the "backup") to initiate recovery of the sessions

➤ Unacceptable limitation for XRF capable applications, which are used to "switching" the sessions in one step without first shutting down the primary application image

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

## Persistent Sessions Forced Takeover MNPS enhancement for forced takeover



➤ Enhance MNPS to provide "forced" takeover capability similar to SNPS

/ Allow another image of the application, on a different VTAM image, to acquire (or recover, in MNPS terminology) sessions from an existing application image without requiring the existing image to fail or to be shutdown

- Similar to "backup" application "switching" XRF sessions from the primary application

/ For migration reasons, MNPS "forced" takeover requires indicators that both the existing and the "recovering" applications images want or will allow forced takeover to happen:

- Existing application image must indicate forced takeovers are permitted
- Recovering application image must indicate this OPEN ACB should be allowed to trigger forced takeover processing

/ MNPS forced takeovers will still be limited to application states that imply some level of session recovery is possible

- Cleanup or termination states not acceptable, as the application is in the process of shutting down and terminating the sessions anyway
- Persistence disabled state not acceptable either, as the application has not indicated that it wants its sessions to persist in the case of application or node failure
- Consistent with existing SNPS forced takeover limitations

➤ Permit SNPS applications to reject SNPS forced takeovers

/ Application can indicate support for SNPS forced takeovers using the same mechanism used to indicated support of MNPS forced takeovers

/ Default, for migration reasons, remains that SNPS forced takeovers are supported

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

## Persistent Sessions Forced Takeover (ACB macroinstruction)



➤ New PARM operand: PARM=(FORCETKO=YES|NO)

┆ **FORCETKO** only valid when PERSIST=YES is also specified

┆ **FORCETKO** indicates willingness of application to initiate forced takeover processing

• **FORCETKO=YES** indicates any OPEN ACB associated with this ACB can initiate MNPS forced takeover processing if necessary

• **FORCETKO=NO** indicates any OPEN ACB associated with this ACB cannot initiate MNPS forced takeover processing

┆ Default is FORCETKO=NO

➤ Regardless of setting, other forms of MNPS processing (planned takeover, standard recovery, etc.) are permitted for this OPEN ACB

┆ Determination of type of MNPS processing required, and whether forced takeover is applicable or not, is made dynamically by VTAM based on the application status at the time of the OPEN ACB

┆ Planned takeover will be performed instead of forced takeover if the application state allows for planned takeover

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

## Persistent Sessions Forced Takeover (SETLOGON macroinstruction)



- New PARM operand: PARM=(FORCETKO=ALL|MULTI|NONE|SINGLE)
  - ⌈ **FORCETKO** only valid when OPTCD=PERSIST is also specified on the SETLOGON
  - ⌈ **FORCETKO** indicates which level of forced takeover processing is permitted
    - **FORCETKO=ALL** indicates that this application supports receipt of both SNPS and MNPS forced takeover attempts
    - **FORCETKO=MULTI** indicates that this application will accept MNPS forced takeover attempts, but not SNPS forced takeovers
    - **FORCETKO=NONE** indicates that this application will not accept any forced takeover attempts
    - **FORCETKO=SINGLE** indicates that this application will accept SNPS forced takeover attempts, but not MNPS forced takeovers
- When an application issues OPEN ACB, FORCETKO=SINGLE is the initial setting
  - ⌈ This setting was chosen for migration purposes for existing applications
- Current setting of FORCETKO (either the initial default setting, or the most recently specified SETLOGON value) is unaffected by subsequent SETLOGON OPTCD=PERSIST invocations that do not specify FORCETKO
  - ⌈ Similar to the existing PARM=(PSTIMER) operand on SETLOGON
- Current setting of FORCETKO is also unaffected by SETLOGON OPTCD=NPERSIST invocations
  - ⌈ However, applications that are not enabled for persistence cannot be taken over via either MNPS or SNPS forced takeover requests

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

## Persistent Sessions Forced Takeover (VTAM Displays)



➤ Added new message to the DISPLAY ID=*applname* displays

⌈ One of three possible messages can be displayed:

- **IST2061I NO FORCED TAKEOVER REQUESTS ARE ACCEPTABLE**
- **IST2062I {SNPS|MNPS} FORCED TAKEOVER REQUESTS ARE ACCEPTABLE**
- **IST2063I ALL FORCED TAKEOVER REQUESTS ARE ACCEPTABLE**

⌈ Appropriate message is included in the following displays:

- *applname* is an MNPS application located within the Sysplex
- *applname* is a local persistent-capable application
- *applname* is locally defined as a CDRSC, but is also an MNPS application located within the Sysplex.

➤ Example when *applname* is an MNPS application located within the Sysplex:

```
d net,id=mappl1,e
<<TIME>> 9.19.13.23-2003.329
JOB  2  IST097I DISPLAY ACCEPTED
JOB  2  IST075I NAME = NETA.MAPPL1, TYPE = APPL
IST1549I OWNER = NETA.SSCP1A MNPS STATE = ENABLED
IST2061I NO FORCED TAKEOVER REQUESTS ARE ACCEPTABLE
IST314I END
```

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



## Display CSM usage Sample report



```
d net, csmuse
IVT5508I DISPLAY ACCEPTED
IVT5572I PROCESSING DISPLAY CSMUSE COMMAND - OWNERID NOT SPECIFIED
IVT5532I -----
IVT5575I USAGE SUMMARY - 4KECSA POOL TOTAL (ALL USERS) = 256K
IVT5576I AMOUNT  MONITOR ID  OWNERID  JOBNAME
IVT5577I      64K      B2           01F6    TCPCS1
IVT5577I      56K      B2           01F7    TCPCS2
IVT5577I      48K      B1           01F6    TCPCS1
IVT5577I      32K      22           0112    NET
IVT5578I DISPLAY TOTAL FOR 4KECSA POOL (4 USERS) =      202K
IVT5532I -----
IVT5575I USAGE SUMMARY - 16KECSA POOL TOTAL (ALL USERS) = 128K
IVT5576I AMOUNT  MONITOR ID  OWNERID  JOBNAME
IVT5577I      102K     98           01F6    TCPCS1
IVT5578I DISPLAY TOTAL FOR 16KECSA POOL (1 USERS) =      102K
IVT5532I -----
IVT5575I USAGE SUMMARY - 32KECSA POOL TOTAL (ALL USERS) = 1024K
IVT5576I AMOUNT  MONITOR ID  OWNERID  JOBNAME
IVT5577I      512K     B2           01F6    TCPCS1
IVT5577I      256K     42           0112    NET
IVT5577I      64K      B1           01F7    TCPCS2
IVT5577I      32K      22           0112    NET
IVT5578I DISPLAY TOTAL FOR 32KECSA POOL (4 USERS) =      864K
IVT5532I -----
IVT5599I END
```

## EE IP packet trace formatter



- A part of the TCP/IP SYSTCPDA packet trace formatter.
- Use `OPTIONS((FORMAT EE))` to select and format EE packets.
- The keyword **EE(12000)** can be used to select and format packets for port 12000 and the next four ports.
- There are no migration issues.

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



## EE IP packet trace formatter Sample trace - part 1 of 2



```

EE: 140
LDLC:
Remote Sap      : 04          Source Sap: 04 Request Control: AF (XID3 without poll)
STID... 34     LENTH... 89     BLKNM.. FFF.B9900
NCHR... 10CB   BIND... 41     NGFL... 00      TGSU... B0
TGNO... 00     DLCT... 01     LEN.... 0B
SDLC:
LSCP... 7100   ABCN... 00     MBTU... 7FFF    PROF... 00
MODE... 00     MAXF... 00
Control Vectors:
Cv..... 0E     Len.... 09     Type... F1      PU..... INDIAPU
Cv..... 0E     Len.... 0e     Type... F4      CPNAME NEEAST.INDIA
Cv..... 0E     Len.... 08     Type... F7      LINK... EIN2SW1
Cv..... 46     Len.... 09     Transmission group
SV..... 80     Len.... 09     Transmission Group
00 09800100 02000000 00     *.....      *
Cv..... 10     Len.... 3a     Product Set ID
SV..... 11     Len.... 23     IBM Software
SF... 02     Len.... 0E     Product component 569511701150
SF... 04     Len.... 08     Product level 060105
SF... 06     Len.... 0A     Product name ACF/VTAM
SV... 11     Len.... 16     IBM Hardware
SF... 00     Len.... 13     Hardware product
00 130011F9 F6F7F200 00000000 F0F3F7F3 *...9672.....0373*
10 F3F4F9 *349 *
5 control vectors found
    
```

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

## EE IP packet trace formatter Sample trace - part 2 of 2



```

EE: 216
LDLC:
Remote Sap      : 04          Source Sap: 08 Request Control: 03 (UI Control)
NLH Anr Route
TpF             : NET          Flags: No_Delay
Link           : D4000000 00000000
Thdr
TCID           : 32884A34 00000093
Reuse_Ct       : 32884A34      Index: 0000          Element: 0093
Flag1         : 38            Flag2: 04
Offset        : 000A         Length: 000000A1    Sequence: 00000000
Segment       : 0E           Size: 05            Status
                00 050E0000 00010002 00000000 00000000 *.....*
                10 00000000 *.....*

TH5
Flags..5D00      SNF....000F      SA.....80000000 01000001
Rh - Session Control Request - Bind Session
RH.....6B8000 - FI BCI ECI DR1
Ru: 146
000000 31001307 B0B050B3 3F879797 873F0602 |.....&..gppg... 1.....P.?....?..|
000010 00000000 00000000 2300000B D5C5E3C1 |.....NETA.....#. ....|
000020 4BE2E2C3 D7F2C121 000802C3 D7E2E5C3 |.SSCP2A....CPSVC K.....!.....|
000030 D4C70903 01ABEEC3 4D6890F9 0C04D5C5 |MG.....C(...9..NE.....Mh.....|
000040 E3C14BE2 E2C3D7F2 C1000BD5 C5E3C14B |TA.SSCP2A..NETA...K.....K|
000050 E2E2C3D7 F1C16014 F6ABEEC3 4D6890F9 |SSCP1A-.6..C(...9.....Mh...|
000060 0BD5C5E3 C14BE2E2 C3D7F2C1 2B180101 |.NETA.SSCP2A.....K.....+....|
000070 16461480 150BD5C5 E3C14BE2 E2C3D7F1 |.....NETA.SSCP1 .F.....K.....|
000080 C1218000 00022C0A 0708C3D7 E2E5C3D4 |A.....CPSVCM !.....|
000090 C740          |G.....g
    
```

## EE IP packet trace formatter



- The latest packet trace formatter may be downloaded from [www.ibm.com](http://www.ibm.com)
- <ftp://ftp.software.ibm.com/software/network/commserver/downloads/systcpda.zip>
- Contains three files
  - , SYSTCPDA.XMI - a file in TSO RECEIVE format
  - , SYSTCPDA.HLP - a TSO HELP file
  - , README - instructions
- Updates for the following packet formatters:
  - , RipNg - RIP for IPv6 packets
  - , EE - Enterprise Extended packets
  - , OSPF - OSPF version 3 packets
  - , TELNET - Format 3270 data streams
  - , LPR - Line printer request
  - , SYSLOG - syslog data
  - , POP3 - POP3 data
  - , SMTP - Simple Mail Transfer Protocol
- SESSION(PIPE) report
  - , This alternative session report shows the amount of data waiting for acknowledgement.
- CHECKSUM(DETAIL|SUMMARY) | NOCHECKSUM
  - , Check sum all selected packet and report any errors
- FLAG(DATA) and DATASIZE(nnnnn) to select packets that are larger than DATASIZE
- SNIFFER(TCPDUMP) to write packet trace records in TCPDUMP format.

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

## Miscellaneous SNA items



### ➤ Enhanced Addressing for DLUR and RTP PUS

- DLUR and RTP PUs (CNRxxxx) are now assigned high order element addresses.
- There are no modifications needed to enable this feature.
- There are no migration issues.

### ➤ TelNet Visibility

- The TelNet server now informs the local VTAM of a change in the client IP address following a session takeover (e.g. TKOSPECLURECON).
- A new bit is passed in the CV64 indicating a particular TelNet LU could have its address changed following a session takeover; this new bit flows to all exits to which the CV64 currently flows (e.g. LOGON exit)
- There are no modifications needed to enable this feature.
- There are no migration issues.

### ➤ Stalled HPR pipe detection

- Rare error conditions can result in an HPR pipe stalling
- This condition has so far gone un-detected
- In z/OS V1R6 new detection and recovery logic has been added - and new messages will be displayed when an HPR pipe stall is detected
  - IST1955I STALL DETECTED FOR RTP puname TO cpname
  - IST1956I STALL CONTINUES FOR RTP puname TO cpname
  - IST1957I STALL ALLEVIATED FOR RTP puname TO cpname

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

## EE Task Force Activity

Copyright International Business Machines Corporation 2004. All rights reserved.



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

## Task Force Focus Areas



➤ While EE has been functioning well, and has in fact been surprisingly easy to deploy for some of our customers, a number of areas may need additional focus to ensure that EE meets the needs of large enterprise networks:

✓ Initial focus areas:

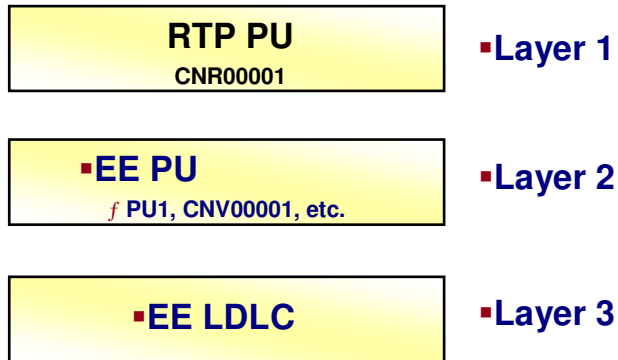
- ✓ Improved throughput
- ✓ Reduced CPU utilization

✓ Follow-on focus areas:

- Architectural evolution
- Security
- Usability
- Network management

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

## Three Layers of Inactivity Timers



### Layer 3: When Does the EE Connection Go Away?



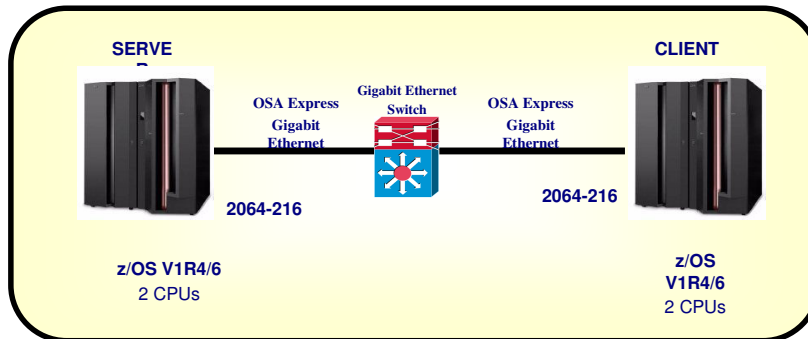
- The LDLC layer monitors the EE connection, and will terminate the EE connection if contact is lost with the partner
- The LDLC inactivity trigger is controlled by three parameters on the PORT statement:
  - ,LIVTIME: The amount of time of inactivity before LDLC tests the connection
  - ,SRQTIME: The amount of time LDLC waits for a response to its test
  - ,SRQRETRY: The number of times the test is retried
- The connection will be terminated if no activity/response for a duration of approximately:
  - ,LIVTIME + ((SRQRETRY+1) \* SRQTIME)

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



- Expand the LDLC "liveness window" during periods of inactivity on the EE connection.
- Start out with the LIVTIME value that is coded (or defaulted) on the XCA port, but if no activity is detected, double the LIVTIME value for the next interval.
- Continue doubling the LIVTIME interval until it reaches a user-defined maximum (up to 1 hour), or until data again starts flowing on the connection (at which time it immediately reverts back to the original LIVTIME value).
- XCA LIVTIME operand now has an initial and maximum value. For example, LIVTIME=(10,3600)
- Goal: Reduce the keep-alive chatter to as little as reasonably possible. Tests of the prototype have been encouraging, and indicate a modest CPU savings may be possible even in the asymmetric case (where the change is only on VTAM).

## OSA Express Gigabit Ethernet Configuration for EE:



➤ Simulated 99 EE Links for LDLC Timer measurements, using LIVTIME=(10,3600)

The numbers in this test are based on a single RTP pipe between the two z/OS images, with 20 sessions using the pipe.

➤ **Performance Comparison of 99 EE link with  
disconnet=no, with and without bootleg LIVTIME**

## RTP Queueing Optimization (V1R6)




- This code change optimizes the management of two RTP queues, the RPN\_Pending\_Sends\_Q and the RPN\_Wait\_For\_Ack\_Q. Both of these queues are singly-linked, ordered lists of control blocks representing NLPs waiting to be sent, and NLPs already sent but waiting for partner acknowledgement, respectively.
- Under heavy loads, it is possible for there to be thousands of control blocks queued to these lists, and the current queue handling code does full traversal of the entire list(s) for every operation.

**Performance Comparison with and without  
Queue optimization modification**

**Performance Comparison with and without  
Queue optimization modification**

- Current options include IPSec (host-based, and/or offloaded to router), and SNA Session-Level Encryption.
- SSL/TLS has not been applicable to EE traffic since EE is UDP-based, and SSL is a TCP-based protocol.
- In addition to authentication and encryption requirements, a remaining inhibitor for many enterprises considering EE is the necessity of opening up firewall ports to UDP traffic.
- Are there any alternatives?

## EE SSL Proxy



➤ A proxy approach could allow the transport of the EE traffic over a TCP connection, thereby allowing several options:

- TCP transport without SSL. This avoids opening the firewall ports to UDP.
- TCP transport with SSL authentication only. This addresses the firewall concern and provides partner authentication without incurring encryption costs.
- TCP transport with SSL encryption.



➤ Issues/Concerns:

• Development Cost / Time to Delivery

• RTP/TCP interaction - Dueling Transport Layers

- ARB flow control vs. TCP flow control
- RTP packet retransmission vs. TCP packet retransmission

• Scalability

- Throughput and latency impact with and without encryption
- CPU cost to enable encryption

➤ ARB Desensitization

┆ An RTP endpoint makes decisions on initiating path switch based on ARB computations of smoothed round trip times (SRTT) and smoothing deviations. If the RTP partner endpoint cannot provide predictable response times (underpowered or maxed-out CPU, etc.), then many extraneous path switch attempts may occur, wasting CPU and network bandwidth, and producing large amounts of message output. The task force is exploring mechanisms (likely additional controls) to tune the sensitivity of the ARB algorithm to reduce the likelihood of unnecessary path switch attempts.

➤ Intelligent re-FIFO timer management

┆ If the RTP path contains a link type (such as EE) that does not guarantee in-order delivery of packets, then the RTP endpoint must set a re-FIFO timer whenever a gap is detected, and that timer must pop before the gap is reported to the partner (requesting retransmission). The current algorithms are formulas based on fixed multipliers of SRTT and the smoothing deviation. However, we have discovered that in some cases this results in non-optimal performance (waiting too long for packet retransmission, or worse, forcing unnecessary retransmissions). The task force plans to explore enhancements to improve the re-FIFO timer choice by incorporating other statistics such as input queue depth and duplicate packet arrivals.

# Trademarks, Copyrights, and Disclaimers



The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both:

IBM	CICS	IMS	MQSeries	Tivoli
IBM (logo)	Cloudscape	Informix	OS/390	WebSphere
e(Business)	DB2	Series	OS/400	xSeries
ALX	DB2 Universal Database	Lotus	pSeries	zSeries

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, ActionMedia, LANdesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

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

Linux is a registered trademark of Linux Torvalds.

Other company, product and service names may be trademarks or service marks of others.

Product data has been reviewed for accuracy as of the date of initial publication. Product data is subject to change without notice. This document could include technical inaccuracies or typographical errors. IBM may make improvements and/or changes in the product(s) and/or program(s) described herein at any time without notice. Any statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only. References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business. Any reference to an IBM Program Product in this document is not intended to state or imply that only that program product may be used. Any functionally equivalent program, that does not infringe IBM's intellectual property rights, may be used instead.

Information is provided "AS IS" without warranty of any kind. THE INFORMATION PROVIDED IN THIS DOCUMENT IS DISTRIBUTED "AS IS" WITHOUT ANY WARRANTY, EITHER EXPRESS OR IMPLIED. IBM EXPRESSLY DISCLAIMS ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NONINFRINGEMENT. IBM shall have no responsibility to update this information. IBM products are warranted, if at all, according to the terms and conditions of the agreements (e.g. IBM Customer Agreement, Statement of Limited Warranty, International Program License Agreement, etc.) under which they are provided. Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products in connection with this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. IBM makes no representations or warranties, express or implied, regarding non-IBM products and services.

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents or copyrights. Inquiries regarding patent or copyright licenses should be made, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. The actual throughput or performance that any user will experience will vary depending upon considerations such as the amount of multiprocessing in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput or performance improvements equivalent to the ratios stated here.

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

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.

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