

# Understanding Enterprise Extender

## Part II. Nuts and Bolts

Sam Reynolds

[samr@us.ibm.com](mailto:samr@us.ibm.com)

IBM z/OS Communications  
Server Design

# Agenda

---

- Understanding EE: Concepts and Considerations
  - ▶ SNA: Dead or Alive?
  - ▶ Enterprise Extender - What is it?
  - ▶ High Performance Routing (HPR) Concepts
  - ▶ Link Characteristics and TGP
  - ▶ Connection Network
  - ▶ Planning for Enterprise Extender

- Understanding EE: Nuts and Bolts
  - ▶ HPR and EE: Layers and Timers
  - ▶ EE Enhancements in V1R5
  - ▶ Coding and Operating EE on z/OS
  - ▶ Testing and Migration for EE
  - ▶ EE Performance

---

# HPR and EE: PUs and Inactivity Timers

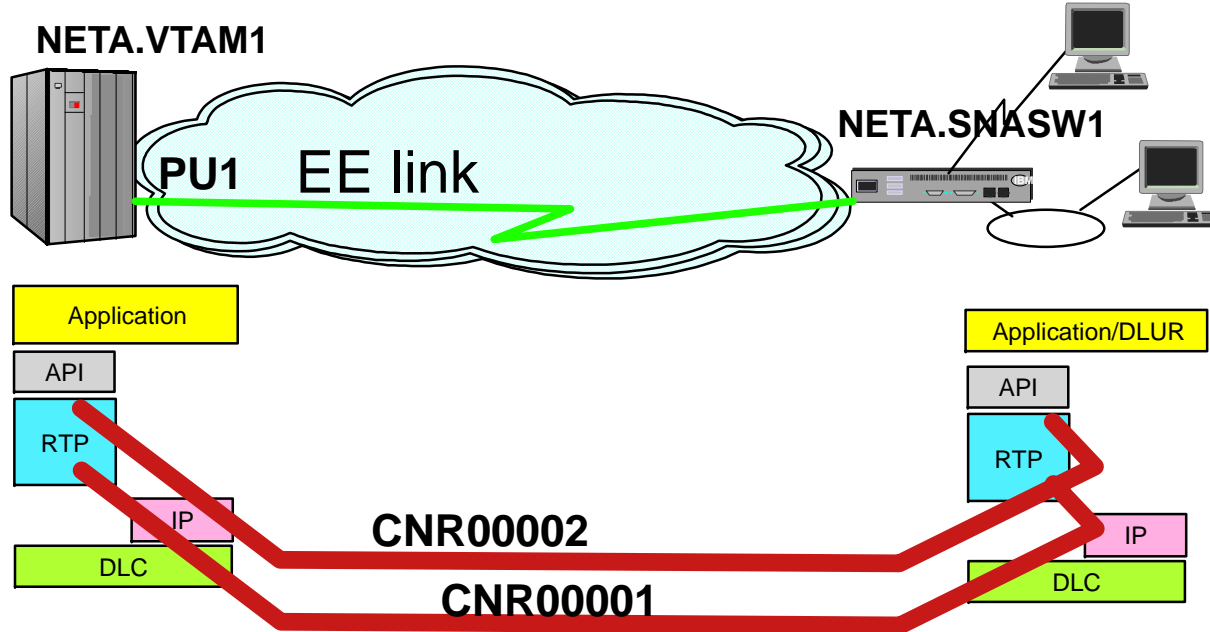
# EE - A Potpourri of PUs

```
VBUILD TYPE=XCA
LINE
PU
LINE
PU
```

```
VBUILD TYPE=SWNET
PU1 PU CPNAME=SNASW1
```

```
ISTRTPMN:
CNR00001
CNR00002
```

```
ISTDSWMN:
CNV00001
CN00001E
```



- XCA Major Node has one LINE/PU for each concurrent EE partner
  - This is a dummy PU and is replaced with an actual PU (linkstation) at connection time
- Switched Major node has PU with CPNAME= matching the remote EE partner's CPname
  - PUs for EE partners may be dynamically defined (ISTEXCCS or DYNPU on AUTOGEN lines) rather than in predefined SWNET
    - DYNPU PUs named CNxxxxxx
    - ISTEXCCS exit supplies PUnames
- If connection network is used, PU may be dynamically defined
  - Named CNVxxxxx
  - Put in Major Node ISTDSWMN
    - Same major node as ISTECCS or DYNPU PUs
- As sessions start, RTP connection PU is dynamically defined (one or more to each EE partner for each APPNCOS)
  - Named CNRxxxxx
  - Put in Major Node ISTRTPMN

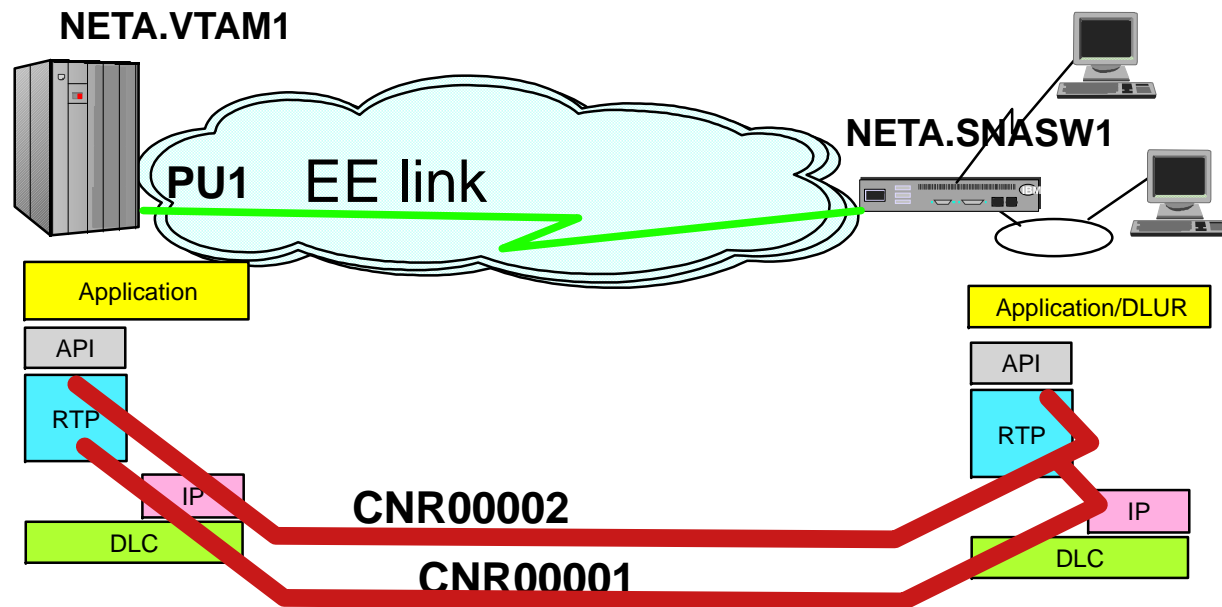
# When Do the PUs Go Up and Down?

```
VBUILD TYPE=XCA
LINE
PU
LINE
PU
```

```
VBUILD TYPE=SWNET
PU1 PU CPNAME=SNASW1
```

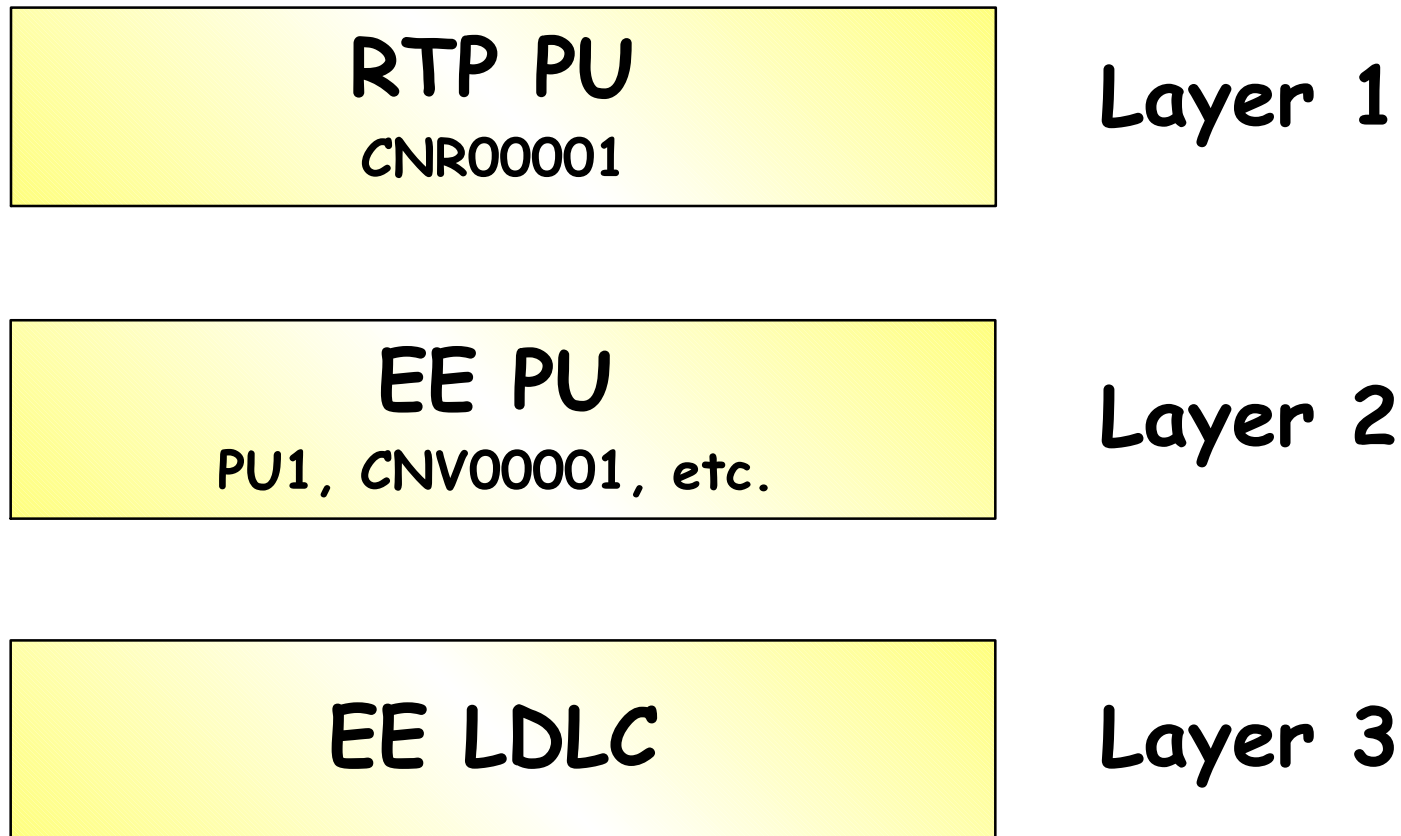
```
ISTRTPMN:
CNR00001
CNR00002
```

```
ISTDSWMN:
CNV00001
CN00001E
```

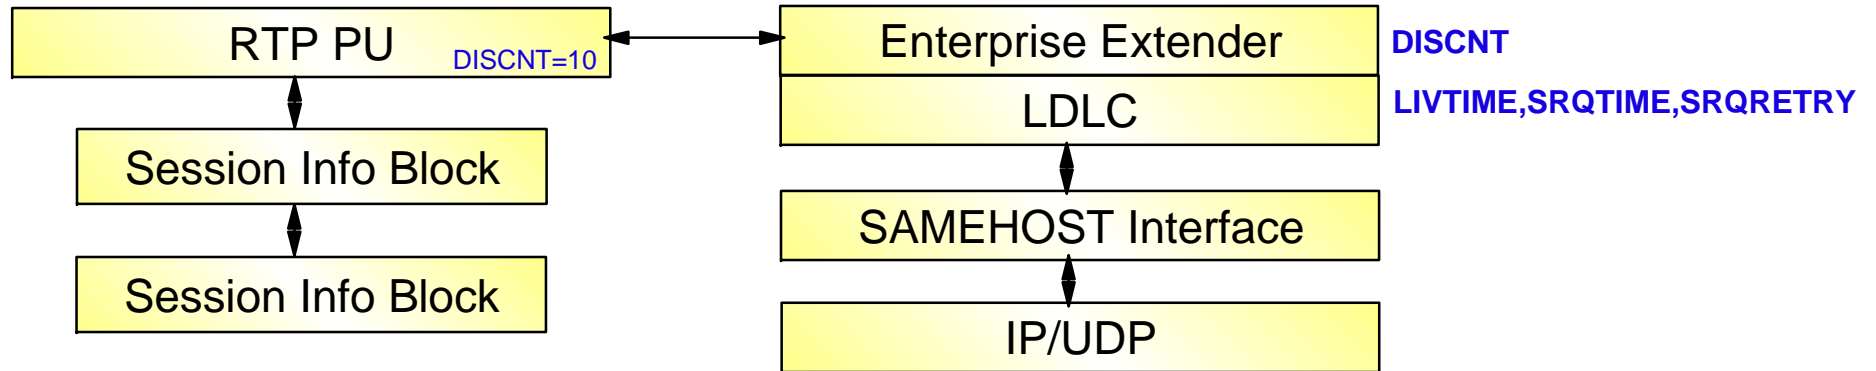


- XCA Major Node has one LINE/PU for each concurrent EE partner
  - Lines display as ACTIVE when VTAM establishes communication with TCP/IP
  - LINE displays as ACTIVE with IST089I showing PUname (i.e. PU1) once link to EE partner comes up
- Dynamic PUs (RTPs, Connection Network Links, DYNPU, etc) are created at connection time and go away when no longer in use (assumes DISCNT=YES)
- SWNET PU is CONCT after activated and ACTIVE once link to EE partner comes up
  - PU is active as long as sessions exist
  - DISCNT=NO specifies link should stay up even with no activity on link
  - For VTAM-VTAM links, DISCNT=NO must be specified on both sides for link to stay up when no activity

# Three Layers of Inactivity Timers




# Layer 1: When Does the HPR Pipe Go Away?



- In z/OS CS, the RTP endpoint is represented by a dynamic PU, and acts like a "delayed disconnect" PU
- Once the last session using the RTP terminates, the RTP will go inactive if no new session is queued to it before the disconnect time expires
- The disconnect time defaults to 10 seconds
- The disconnect time can be altered by defining a Model RTP PU
  - ▶ Note: Need to activate on both sides

```
MODMAJND VBUILD TYPE=MODEL
RTPPU PU DYNTYPE=RTP,DISCNT=(DELAY,,60)
```

# Layer 2: When Does the EE Connection Go Away?



- Since the EE link is itself represented by a switched PU (which may be either predefined or dynamic), DISCNT can be coded to inop the EE link
  - ▶ If DISCNT=YES specified, disconnect time will be specified by the DISCNTIM start option
  - ▶ Disconnect time may be directly specified by coding DISCNT=(DELAY,,time)
  - ▶ DISCNT defaults to NO for predefined links
  - ▶ DISCNT defaults to YES for dynamic links prior to V1R5. In V1R5, the default for non-connection network dynamic PUs is changed to NO.
- Since the EE PU (like any HPR PU) has no sessions queued directly to it, the disconnect decision is made based on a period of inactivity rather than the termination of the last session (as with non-HPR) PUs
- RTP's "ALIVE" timer is set to half of the disconnect time so that RTP status messages ("keep alive" flows) will keep the EE link from dropping as long as at least one RTP is pointing to the EE PU



# 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)**

**NOTE:** The LDLC parameters associated with the EE connection are completely independent from the ARB parameters associated with the RTP pipes!

# Inactivity Example #1

```
d net,rtps
```

```
...
```

IST1695I	PU NAME	CP NAME	COS NAME	SWITCH	CONGEST	SESSIONS
IST1696I	CNR00004	NETA.SSCP2A	#INTER	NO	NO	1
IST1696I	CNR00003	NETA.SSCP2A	RSETUP	NO	NO	0

```
...
```

11:43:32 **Break the IP connectivity.**

```
v tcpip,,stop,trlela
```

```
EZZ0060I PROCESSING COMMAND: VARY TCPIP,,STOP,TRLELA
```

```
EZZ0053I COMMAND VARY STOP COMPLETED SUCCESSFULLY
```

```
EZZ4315I DEACTIVATION COMPLETE FOR DEVICE TRLELA
```

11:44:14 **RTP CNR00004 detects failure and goes into path switch.**

```
IST1494I PATH SWITCH STARTED FOR RTP CNR00004
```

11:44:49 **The LDLC layer detects loss of IP connectivity and inops EE connection**

```
IST1411I INOP GENERATED FOR LNIP1
```

```
IST1430I REASON FOR INOP IS APPN CONNECTION ENDED DUE TO INACTIVITY
```

```
IST314I END
```

```
IST259I INOP RECEIVED FOR SWIP2A1 CODE = 01
```

```
IST619I ID = SWIP2A1 FAILED - RECOVERY IN PROGRESS
```

```
IST1196I APPN CONNECTION FOR NETA.SSCP2A INACTIVE - TGN = 21
```

```
IST590I CONNECTION TERMINATED FOR PU SWIP2A1 ON LINE LNIP1
```

```
IST621I RECOVERY SUCCESSFUL FOR NETWORK RESOURCE SWIP2A1
```

```
IST1488I INACTIVATION FOR RTP CNR00003 AS PASSIVE PARTNER COMPLETED
```

```
IST619I ID = CNR00003 FAILED - RECOVERY IN PROGRESS
```

```
IST129I UNRECOVERABLE OR FORCED ERROR ON NODE CNR00003 - VARY INACT SCHED
```

```
IST105I CNR00003 NODE NOW INACTIVE
```

```
IST871I RESOURCE CNR00003 DELETED
```

11:46:14 **CNR00004's path switch timer (2 min for interactive TP) expires. CNR00004 goes inactive.**

```
IST1494I PATH SWITCH FAILED FOR RTP CNR00004
```

```
IST1495I NO ALTERNATE ROUTE AVAILABLE
```

```
IST314I END
```

```
IST1488I INACTIVATION FOR RTP CNR00004 AS ACTIVE PARTNER COMPLETED
```

```
IST619I ID = CNR00004 FAILED - RECOVERY IN PROGRESS
```

```
IST129I UNRECOVERABLE OR FORCED ERROR ON NODE CNR00004 - VARY INACT SCHED
```

```
IST105I CNR00004 NODE NOW INACTIVE
```

```
IST871I RESOURCE CNR00004 DELETED
```

**Example #1**  
DISCNTIM=(15,0)  
EE PU:  
DISCNT=YES  
PORT:  
LIVTIME=10  
SRQTIME=15  
SRQRETRY=3

# Inactivity Example #2

d net,rtps

...

IST1695I	PU NAME	CP NAME	COS NAME	SWITCH	CONGEST	SESSIONS
IST1696I	CNR00004	NETA.SSCP2A	#INTER	NO	NO	1
IST1696I	CNR00003	NETA.SSCP2A	RSETUP	NO	NO	0

...

11:34:56 **Break the IP connectivity.**

v tcpip,,stop,trlela

EZZ0060I PROCESSING COMMAND: VARY TCPIP,,STOP,TRLELA

EZZ0053I COMMAND VARY STOP COMPLETED SUCCESSFULLY

EZZ4315I DEACTIVATION COMPLETE FOR DEVICE TRLELA

11:36:23 **The LDLC layer detects loss of IP connectivity and inops EE connection**

IST1411I INOP GENERATED FOR LNIP1

IST1430I REASON FOR INOP IS APPN CONNECTION ENDED DUE TO INACTIVITY

IST314I END

IST259I INOP RECEIVED FOR SWIP2A1 CODE = 01

IST619I ID = SWIP2A1 FAILED - RECOVERY IN PROGRESS

IST1196I APPN CONNECTION FOR NETA.SSCP2A INACTIVE - TGN = 21

IST590I CONNECTION TERMINATED FOR PU SWIP2A1 ON LINE LNIP1

IST1494I PATH SWITCH STARTED FOR RTP CNR00004

IST621I RECOVERY SUCCESSFUL FOR NETWORK RESOURCE SWIP2A1

IST1488I INACTIVATION FOR RTP CNR00003 AS PASSIVE PARTNER COMPLETED

IST619I ID = CNR00003 FAILED - RECOVERY IN PROGRESS

IST129I UNRECOVERABLE OR FORCED ERROR ON NODE CNR00003 - VARY INACT SCHED

IST105I CNR00003 NODE NOW INACTIVE

IST871I RESOURCE CNR00003 DELETED

11:38:23 **CNR00004's path switch timer (2 min for interactive TP) expires. CNR00004 goes inactive.**

IST1494I PATH SWITCH FAILED FOR RTP CNR00004

IST1495I NO ALTERNATE ROUTE AVAILABLE

IST314I END

IST1488I INACTIVATION FOR RTP CNR00004 AS ACTIVE PARTNER COMPLETED

IST619I ID = CNR00004 FAILED - RECOVERY IN PROGRESS

IST129I UNRECOVERABLE OR FORCED ERROR ON NODE CNR00004 - VARY INACT SCHED

IST105I CNR00004 NODE NOW INACTIVE

IST871I RESOURCE CNR00004 DELETED

## Example #2

DISCNTIM=(240,0)

EE PU:

DISCNT=YES

PORT:

LIVTIME=20

SRQTIME=15

SRQRETRY=4

# Disconnect/Inactivity Summary

- The LDLC layer monitors the connection, testing it if no activity is detected, and inop'ing the EE link if the tests go unanswered
- The Disconnect timer associated with the EE switched PU can be 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)

# Disconnect Timer Recommendations

---


- Consider lengthening the EE LDLC timer parameters (LIVTIME, SRQTIME, SRQRETRY), especially if RIP is being used
  - ▶ 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 (as calculated using the formula on foil 9).
  - ▶ 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.
- For predefined EE connections, keep DISCNT=NO (default).
- For EE VRN-based dynamic connections, consider coding a DYNTYPE=VN model, with DISCNT=NO, or a delay value of 60+ seconds.
- For DYNPU=YES dynamic connections, DISCNT=YES is the default in V1R4. DISCNT=NO is often preferred, and is the default beginning in V1R5.
- Coding a DYNTYPE=RTP model with DISCNT=NO will keep all RTP pipes active, even with no sessions active.
  - ▶ This may have storage/CPU considerations, but saves the network overhead of RTP pipe setup/takedown activities, and promotes more consistent latency/response times.
  - ▶ APAR OA04393 is recommended if adopting this approach. (See foil 37.)

---

# EE Enhancements in z/OS V1R5

# EE Enhancements in V1R5

- Enhanced HOSTNAME Support removes the incompatibility between EE and network address translation (NAT).
  - ▶ EE architecture has been updated to allow the EE connection network control vectors to carry the hostname corresponding to the EE VIPA.
- V1R5 allows the specification of multiple local and/or multiple global EE connection networks.
- Allows multiple (static) VIPAs, defined on a GROUP basis in the EE XCA major node.
  - ▶ All EE VIPAs must still belong to a single TCP/IP stack
- V1R5 adds EE support for IPv6
- A new model type, DYNTYPE=EE, is added to the PU statement in the Model Major Node, allowing dynamic non-connection network PUs to be customized by:
  - ▶ Coding the DISCNT operand so that a disconnect delay time can be specified.
  - ▶ Specifying the DWINOP, REDIAL, and REDDELAY operands.
  - ▶ Overriding the default TG characteristics that in the past were used for these dynamic PUs.



```
MODEL A1A VBUILD TYPE=MODEL
*
EEMODEL PU DYNTYPE=EE,
CAPACITY=100M,
COSTTIME=0,
CPCP=YES,
DISCNT=NO,
DWINOP=YES,
REDIAL=30,
REDDelay=60
```

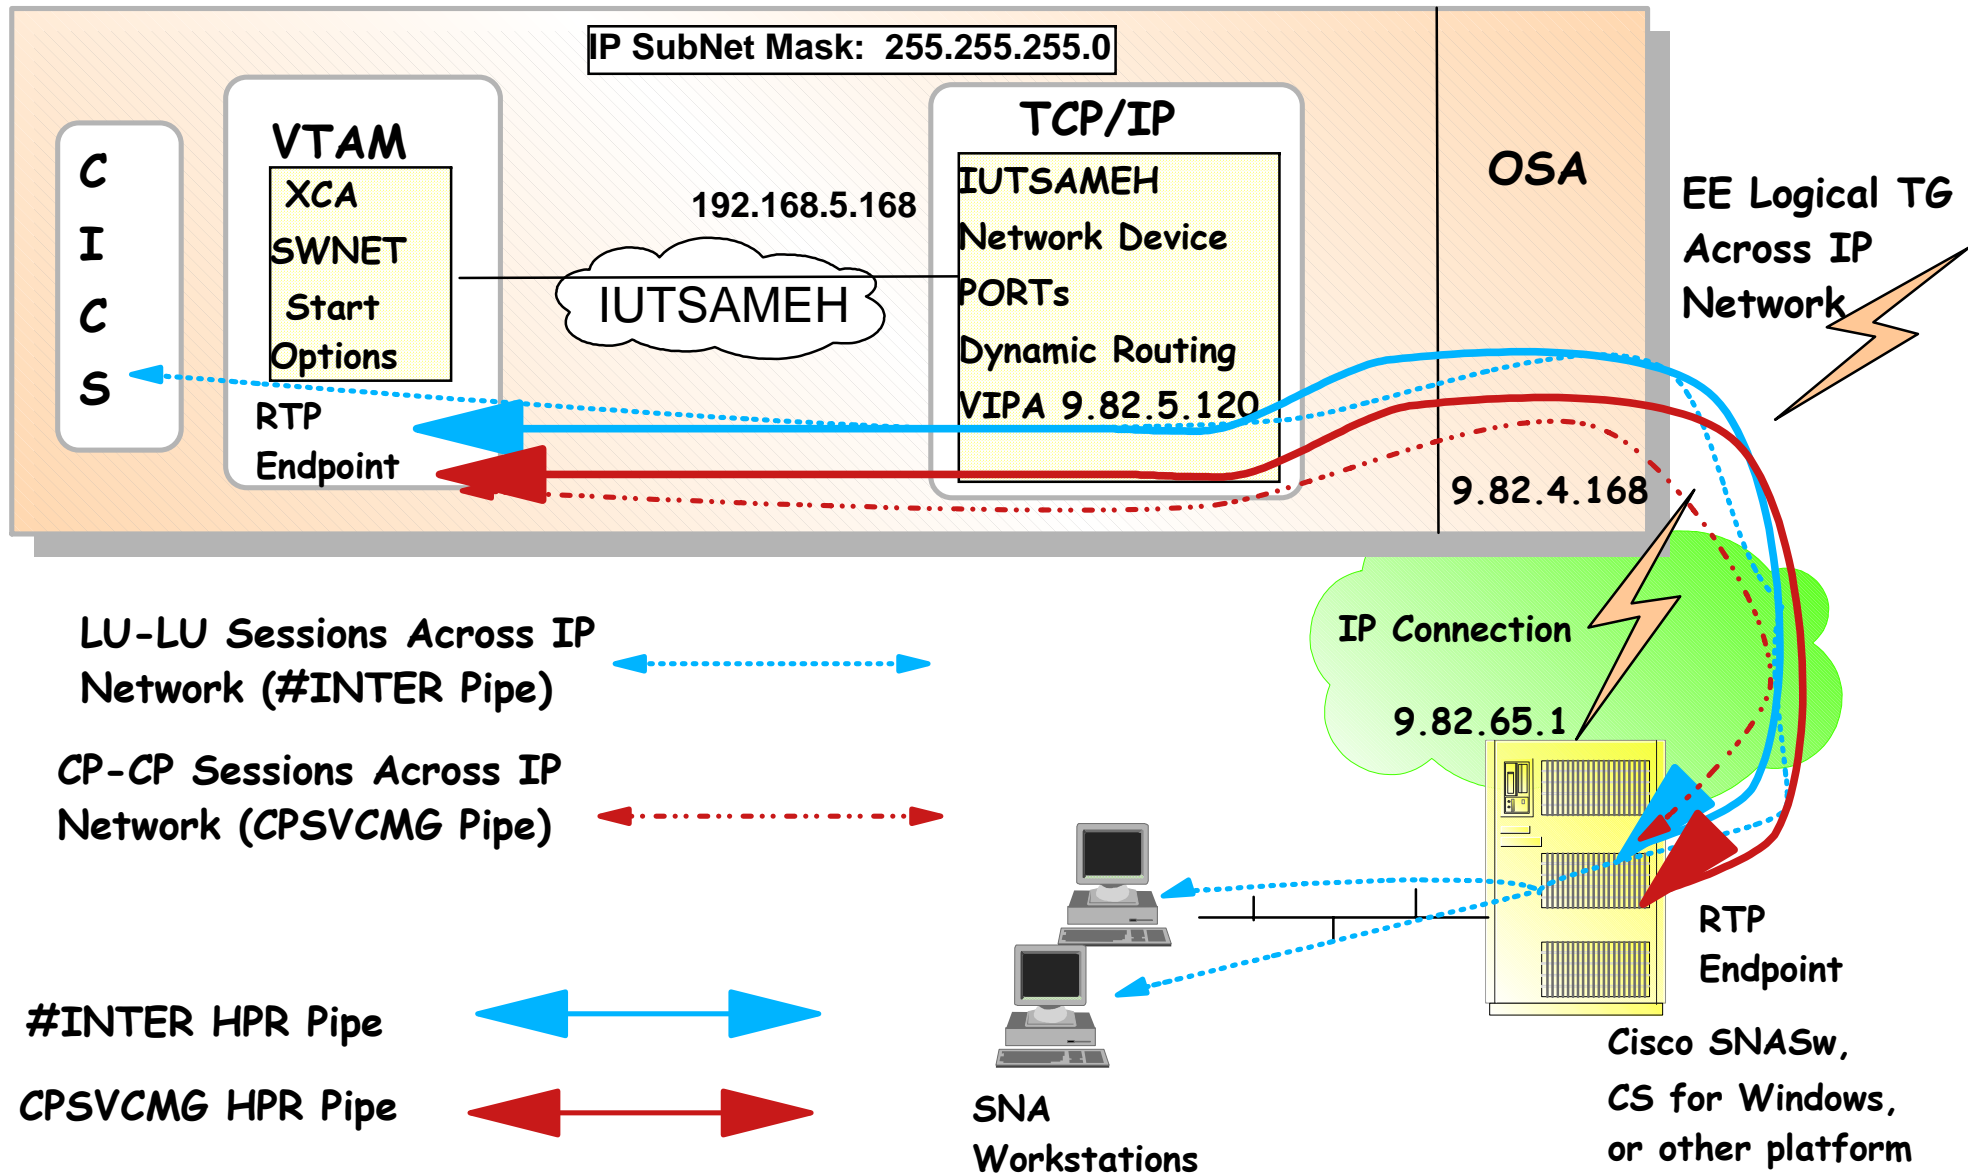
---

# Coding and Operating EE on z/OS



# Enterprise Extender on zServer

z/OS or OS/390



# z/OS CS Enterprise Extender Definitions

---

- **VTAM Definitions:**

- ▶ **Start Options**

- IPADDR, TCPNAME, and HPRARB

- ▶ **XCA Major Node for EE DLC (Medium=HPRIP)**

- ▶ **Switched Major Nodes for Linkstations**

- **TCP/IP Definitions:**

- ▶ **Profile Definitions:**

- Port reservations - by default, EE uses PORTS

- 12000-12004 and TOS C0,C0,80,40,20 (respectively)

- IUTSAMEH device and link (or use DYNAMICXCF)

- Static VIPA address

- ▶ **Other considerations:**

- Dynamic routing is recommended but not required

- EE uses five UDP sockets

# EE XCA Major Node

XCAEEGVN	VBUILD	TYPE=XCA	
PORTEE	PORT	MEDIUM=HPRIP	
GRPEEP	GROUP	DIAL=YES,AUTOGEN=(10,E,X), CALL=INOUT,ISTATUS=ACTIVE	C
GRPEEG	GROUP	DIAL=YES,AUTOGEN=(10,G,V),VNTYPE=GLOBAL, CALL=INOUT,ISTATUS=ACTIVE,TGP=EEXTWAN	C
GRPEEL	GROUP	DIAL=YES,AUTOGEN=(10,L,R),VNNAME=NETA.LVRN, CALL=INOUT,ISTATUS=ACTIVE,TGP=EEXTCAMP	C

- Only one XCA with MEDIUM=HPRIP may be active
  - ▶ AUTOGEN needs to specify maximum number of EE partners expected to be concurrently active
- Coding DYNPU=YES on GROUP allows dynamic definition of APPN PUs (CNxxxxxx)
  - ▶ Coding DYNPU=YES is not needed for connection network links to be dynamically defined
- VNTYPE defaults to LOCAL
  - ▶ if VNTYPE=GLOBAL is coded, the connection network name defaults to IP.IP if VNNAME is not coded
  - ▶ May specify only one LOCAL VRN and one GLOBAL VRN prior to V1R5
- TGPs may be coded on the GROUP Statement (recommended) or PORT
  - ▶ Only applicable to connection networks
  - ▶ Different connection network groups may have different TGPs

# EE Switched Major Node Coding

- Has CPNAME (and NETID, if different) of EE partner
  - ▶ DWACT=YES causes "dial-out" to occur when SWNET activated
  - ▶ PATH Statement contains IP address or HOSTNAME of EE partner for dial-out
    - With different SAPADDR= on PATH, parallel TGs may be coded but this is of limited value, and is not recommended. This capability will be removed after z/OS V1R7.
    - PATH statement not needed if partner always dials-in

```
CSS1SWEE VBUILD TYPE=SWNET
CSS1PUE  PU      ADDR=22,DWACT=YES,TGP=TRING16M,          *
                                     DISCNT=NO,CONNTYPE=APPN,PUTYPE=2,      *
                                     CPNAME=CSS1,NETID=CSSNET,REDIAL=3,     *
                                     REDDELAY=30,DWINOP=NO
CSS1EEPT PATH  IPADDR=9.82.5.120,GRPNM=GRPEEP
```

- ▶ If DWINOP=YES specified, it is recommended that it only be coded on one end of the EE connection to prevent dial conflicts

# Definitional Changes in V1R5

---

- EE support in V1R5 brought a number of changes to z/OS CS EE definition, including a new start option and new XCA GROUP-level operands.
- These definitional changes work together to allow for:
  - ▶ NAT compatibility with EE connection networks
  - ▶ EE IPv6 support
  - ▶ Multiple EE VRNs and VIPAs

**Note:** Existing EE definitions (start options, XCA, and SMN) do not have to be changed unless you want to exploit one (or more) of the above new functions.

# EE XCA in V1R5

```
XCAEE1A  VBUILD TYPE=XCA
PORTEE   PORT   MEDIUM=HPRIP
GPEE1    GROUP  ANSWER=ON,
          AUTOGEN=(1,LNEE1,PUEE1),
          CALL=INOUT,
          DIAL=YES,
          IPADDR=10.1.1.1
GPEE4    GROUP  ANSWER=ON,
          AUTOGEN=(1,LNEE4,PUEE4),
          CALL=INOUT,
          DIAL=YES,
          HOSTNAME=VIP24.SSCP1A
GPEE6    GROUP  ANSWER=ON,
          AUTOGEN=(1,LNEE6,PUEE6),
          CALL=INOUT,
          DIAL=YES,
          HOSTNAME=VIP26.SSCP1A
```


- Enterprise Extender in V1R5 was updated to allow for the exploitation of an IPv6-enabled network, and to allow EE/NAT compatibility.
- Existing IPADDR keywords (start option, path definition in switched major node) are IPv4-only
- IPv6 support requires use of the HOSTNAME keyword (start option, GROUP, path definition)
- EE Connection networks must be IPv4-only or IPv6-only
- To use connection network, nodes supporting both IPv4 and IPv6 will define an IPv4 VRN (local and/or global) and an IPv6 VRN

# VARY Command Enhancements for EE XCA (V1R6)

- 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, prior to V1R6 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.

# Display EE Command (V1R6)

- 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





# D EE: Specific Connection Information (V1R6)

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

IST097I DISPLAY ACCEPTED

IST350I DISPLAY TYPE = EE

IST2001I 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 SWEE2A1

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 )

.

# D EE: Specific Connection Info...


---

```
.
.
.
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
```

---

# Testing and Migration for EE

# Migrating from Native HPR to EE



- For this example, we want to begin introducing an EE link into an existing APPN-enabled network using MPC+ connectivity.
- During our first test window, we will activate an EE link, verify distribution of traffic between the two links, and then restore the network to its original state.
- By default, an MPC+ link will have a weight of 30 for the #INTER APPN COS. An EE link will have a weight of 60.


# Migrating from Native HPR to EE...

- The MPC+ connection is activated, and we start a single session across it.

```
d net,topo,orig=sscp1a,dest=sscp2a,appncos=#inter
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = TOPOLOGY
IST1299I TRANSMISSION GROUPS ORIGINATING AT CP NETA.SSCP1A
IST1357I
IST1300I DESTINATION CP      TGN      STATUS  TGTYPE  VALUE WEIGHT
IST1301I NETA.SSCP2A        21      OPER    INTERM  YES    30
...
```

## <<< Activate a session >>>

```
IST1488I ACTIVATION OF RTP CNR00003 AS ACTIVE TO NETA.SSCP2A
d net,id=cnr00003,e
...
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
...
d net,rtps
...
IST1695I PU NAME          CP NAME          COS NAME SWITCH CONGEST  SESSIONS
IST1696I CNR00003 NETA.SSCP2A #INTER          NO        NO        1
IST1696I CNR00002 NETA.SSCP2A RSETUP          NO        NO        0
IST1696I CNR00001 NETA.SSCP2A CPSVCMG         NO        NO        2
```



# Migrating from Native HPR to EE...

- We activate the EE connection, and display the parallel connections. A new session continues to use the MPC+ connection.


```
v net,act,id=xcaip
IST097I VARY ACCEPTED
IST093I XCAIP ACTIVE
v net,act,id=lnip1
IST097I VARY ACCEPTED
IST1685I TCP/IP JOB NAME = TCPCS
IST1680I LOCAL IP ADDRESS 9.67.1.1
IST093I LNIP1 ACTIVE
EZZ4324I CONNECTION TO 9.67.1.1 ACTIVE FOR DEVICE IUTSAMEH
v net,act,id=toip2a
IST097I VARY ACCEPTED
IST093I SWIP2A1 ACTIVE
IST093I TOIP2A ACTIVE
IST590I CONNECTOUT ESTABLISHED FOR PU SWIP2A1 ON LINE LNIP1
IST1086I APPN CONNECTION FOR NETA.SSCP2A IS ACTIVE - TGN = 22
```

<<< **Activate another session** >>>

```
d net,rtps
```

...

IST	PU NAME	CP NAME	COS NAME	SWITCH	CONGEST	SESSIONS
IST1695I	CNR00003	NETA.SSCP2A	#INTER	NO	NO	2
IST1696I	CNR00002	NETA.SSCP2A	RSETUP	NO	NO	0
IST1696I	CNR00001	NETA.SSCP2A	CPSVCMG	NO	NO	2




# Migrating from Native HPR to EE...

- Check the TG weights for #INTER

```
d net,topo,orig=sscp1a,dest=sscp2a,appncos=#inter
```

```
...
IST1299I TRANSMISSION GROUPS ORIGINATING AT CP NETA.SSCP1A
IST1357I                                     CPCP
IST1300I DESTINATION CP      TGN      STATUS  TGTYPE  VALUE WEIGHT
IST1301I NETA.SSCP2A        21      OPER    INTERM  YES   30
IST1579I -----
IST1163I RSN                  HPR      TIME LEFT
IST1164I 14                  YES      15
IST1579I -----
IST1302I CAPACITY PDELAY  COSTTIME  COSTBYTE
IST1303I 32M          NEGLIGIB 0          0
IST1579I -----
IST1304I SECURITY UPARM1  UPARM2    UPARM3
IST1305I UNSECURE 128      128      128
IST1579I -----
IST1736I PU NAME
IST1737I AHHCPU1
...
IST924I -----
IST1357I                                     CPCP
IST1300I DESTINATION CP      TGN      STATUS  TGTYPE  VALUE WEIGHT
IST1301I NETA.SSCP2A        22      OPER    INTERM  YES   60
IST1579I -----
IST1163I RSN                  HPR      TIME LEFT
IST1164I 12                  YES      15
IST1579I -----
IST1302I CAPACITY PDELAY  COSTTIME  COSTBYTE
IST1303I 4M          TERRESTR 0          0
IST1579I -----
IST1304I SECURITY UPARM1  UPARM2    UPARM3
IST1305I UNSECURE 128      128      128
IST1579I -----
IST1736I PU NAME
IST1737I SWIP2A1
...

```



# Migrating from Native HPR to EE...

- We associate a new TGP with the EE connection, giving it a weight equal to the MPC+ link.

```
f vtam,tgp,tgpname=eefenet,id=swip2a1
IST097I TGP ACCEPTED
IST1090I TGP FOR PU SWIP2A1 IS SET TO EEFENET
```

```
d net,topo,orig=sscp1a,dest=sscp2a,appncos=#inter
```

```
...
IST1357I                                     CPCP
IST1300I DESTINATION CP      TGN      STATUS  TGTYPE  VALUE WEIGHT
IST1301I NETA.SSCP2A        21      OPER    INTERM  YES   30
...
IST924I -----
IST1357I                                     CPCP
IST1300I DESTINATION CP      TGN      STATUS  TGTYPE  VALUE WEIGHT
IST1301I NETA.SSCP2A        22      OPER    INTERM  YES   30
IST1579I -----
IST1163I                                     RSN      HPR      TIME LEFT
IST1164I                                     14      YES     15
IST1579I -----
IST1302I                                     CAPACITY PDELAY  COSTTIME  COSTBYTE
IST1303I                                     99M      NEGLIGIB 0         0
IST1579I -----
IST1304I                                     SECURITY UPARM1  UPARM2    UPARM3
IST1305I                                     UNSECURE 128    128      128
IST1579I -----
IST1736I                                     PU NAME
IST1737I                                     SWIP2A1
...
```

```
EEFENET  TGP  COSTTIME=0,
          COSTBYTE=0,
          SECURITY=UNSECURE,
          PDELAY=NEGLIGIB,
          CAPACITY=100M
```

- OA12948 will provide five new TGPs for IBMTGPs:
  - ▶ FASTENET
  - ▶ GIGENET
  - ▶ HIPERSOC
  - ▶ FICON
  - ▶ FICONEXP
- OA12948 is a V1R7 APAR, but the TGPs can be copied into IBMTGPS on any level of z/OS CS



# Migrating from Native HPR to EE...

- When we start 10 more sessions, half of them go on the existing pipe over the MPC+ link, while half use the EE link.


```
IST1488I ACTIVATION OF RTP CNR00004 AS ACTIVE TO NETA.SSCP2A
IST1488I ACTIVATION OF RTP CNR00005 AS ACTIVE TO NETA.SSCP2A
```

```
...
d net,rtps
```

```
...
IST1695I PU NAME          CP NAME          COS NAME SWITCH CONGEST  SESSIONS
IST1696I CNR00005 NETA.SSCP2A #INTER          NO        NO        6
IST1696I CNR00004 NETA.SSCP2A RSETUP          NO        NO        0
IST1696I CNR00003 NETA.SSCP2A #INTER          NO        NO        6
IST1696I CNR00002 NETA.SSCP2A RSETUP          NO        NO        0
IST1696I CNR00001 NETA.SSCP2A CPSVCMG        NO        NO        2
```

```
...
d net,id=cnr00005,e
```

```
...
IST1855I NUMBER OF SESSIONS USING RTP = 6
IST1697I RTP PACING ALGORITHM = ARB RESPONSIVE MODE
IST1480I RTP END TO END ROUTE - RSCV PATH
IST1460I TGN  CPNAME          TG TYPE          HPR
IST1461I  22  NETA.SSCP2A      APPN              RTP
IST875I ALSNAME TOWARDS RTP = SWIP2A1
```



# Migrating from Native HPR to EE...

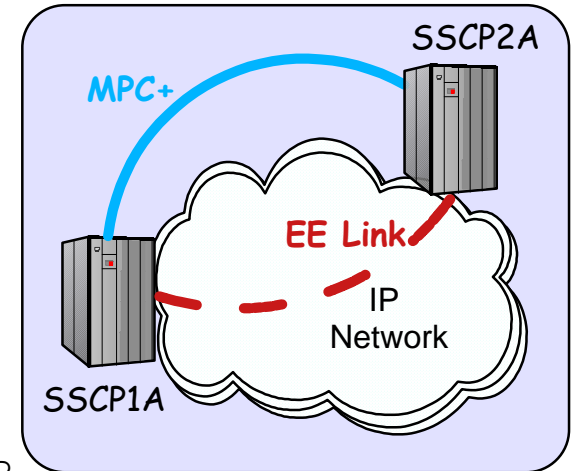
- We now decide to stop using the EE link for new sessions. We use F TOPO to quiesce the EE link.

```
f vtam,topo,orig=sscp1a,dest=sscp2a,tgn=22,function=quiesce
IST097I MODIFY ACCEPTED
IST223I MODIFY TOPO COMMAND COMPLETED
d net,topo,orig=sscp1a,dest=sscp2a,appncos=#inter
```

```
...
IST1357I
IST1300I DESTINATION CP      TGN      STATUS  TGTYPE  VALUE WEIGHT
IST1301I NETA.SSCP2A        21      OPER    INTERM  YES     30
...
IST924I -----
IST1357I
IST1300I DESTINATION CP      TGN      STATUS  TGTYPE  VALUE WEIGHT
IST1301I NETA.SSCP2A        22      QUIES   INTERM  YES     32767
...
```

<<< **Activate 5 more sessions** >>>


```
d net,rtps
...
IST1695I PU NAME          CP NAME          COS NAME  SWITCH  CONGEST  SESSIONS
IST1696I CNR00005 NETA.SSCP2A    #INTER    NO      NO        6
IST1696I CNR00004 NETA.SSCP2A    RSETUP    NO      NO        0
IST1696I CNR00003 NETA.SSCP2A    #INTER    NO      NO       11
IST1696I CNR00002 NETA.SSCP2A    RSETUP    NO      NO        0
IST1696I CNR00001 NETA.SSCP2A    CPSVCMG   NO      NO        2
```



# Migrating from Native HPR to EE...

- We still have a pipe using the EE link. When our test window is over, we inactivate the EE PU, and the pipe switches to the MPC+ link.

```
v net,inact,id=swip2a1
IST097I VARY ACCEPTED
IST1196I APPN CONNECTION FOR NETA.SSCP2A INACTIVE - TGN = 22
IST590I CONNECTION TERMINATED FOR PU SWIP2A1 ON LINE LNIP1
IST1494I PATH SWITCH STARTED FOR RTP CNR00005
IST1819I PATH SWITCH REASON: TG INOP
IST314I END
IST105I SWIP2A1 NODE NOW INACTIVE
IST1488I INACTIVATION OF RTP CNR00004 AS PASSIVE TO NETA.SSCP2A
IST619I ID = CNR00004 FAILED - RECOVERY IN PROGRESS
IST129I UNRECOVERABLE OR FORCED ERROR ON NODE CNR00004 - VARY INACT SCHED
IST105I CNR00004 NODE NOW INACTIVE
IST871I RESOURCE CNR00004 DELETED
IST1494I PATH SWITCH COMPLETED FOR RTP CNR00005
IST1480I RTP END TO END ROUTE - RSCV PATH
IST1460I TGN CPNAME TG TYPE HPR
IST1461I 21 NETA.SSCP2A APPN RTP
IST314I END
d net,rtps
IST1695I PU NAME CP NAME COS NAME SWITCH CONGEST SESSIONS
IST1696I CNR00005 NETA.SSCP2A #INTER NO NO 6
IST1696I CNR00003 NETA.SSCP2A #INTER NO NO 11
IST1696I CNR00002 NETA.SSCP2A RSETUP NO NO 0
IST1696I CNR00001 NETA.SSCP2A CPSVCMG NO NO 2
```



---

# EE Performance

# Inactivity Timer Optimization

---

- Since most HPR deployments are converging on EE, and since most RTP pipes in an EE network are one-hop pipes (or two-hop over an EE connection network), the RTP Alive Timer can be eliminated for those pipes, relying on EE LDLC monitoring to drive path switch and RTP inop processing.
  - ▶ Reduce timer processing overhead
  - ▶ Reduce I/O necessary for RTP status request/reply processing
- This optimization reduces CPU in branch-oriented environments with large numbers of RTP pipes
  - ▶ Reduction of 20% of z/OS CS CPU measured in one large customer network
  - ▶ That customer uses an RTP Model PU with DISCNT=NO to keep pipes up all the time (even during periods of inactivity). Installations that do not employ a similar strategy will probably see smaller benefits.
- APAR OA04393 / PTF UA05982

# EE LDLC Expanding LIVETIME Window

- In a further attempt to reduce CPU and I/O resources associated with keep-alive overhead, we have prototyped an additional optimization whereby the EE LDLC "liveness" window is continually doubled as long as no activity is detected, up until a user-defined maximum (capped at 1 hour).
- If we need to use the EE connection, we immediately collapse the window back to the base value.
- **LIVETIME=(10,3600) Default: LIVETIME=(10,10)**

Base Value      Maximum Value



# EE LDLC Expanding LIVTIME Window...

---

- A prototype of this optimization has been implemented in a large customer environment (2000 EE connections, >10,000 pipes)
- Early measurements indicate that the customer's average z/OS CS CPU utilization has been reduced by over 20% (beyond the reduction achieved with OA04393).
- This customer has a very large number of branch connections with significant periods of inactivity on many of those connections. Installations without similar characteristics will probably see smaller benefits.
- APAR number: OA09535

# EE Considerations & Tips

---

- Consider EE and HPR timer settings. (See foil 13.)
- Consider enabling PSRETRY (off by default) so that HPR pipes will automatically switch to better routes when available
- Configure APPN Link Characteristics
  - New TGP's for EE provided with VTAM - Customization of link speed is recommended
- Do not deploy EE Connection Network over an IP network using NAT until running z/OS V1R5 CS (where the CN/NAT incompatibility is addressed).
- Review EE Info APAR II12223
- Also, see the presentation "z/OS CS Enterprise Extender Hints and Tips"



# For More Information....

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

- Also, see the Redbook:

- ▶ [SG24-5957 Migrating Subarea to an IP Infrastructure](#)

# Summary

---

- Understanding EE: Concepts and Considerations
  - ▶ SNA: Dead or Alive?
  - ▶ Enterprise Extender - What is it?
  - ▶ High Performance Routing (HPR) Concepts
  - ▶ Link Characteristics and TGP
  - ▶ Connection Network
  - ▶ Planning for Enterprise Extender

- Understanding EE: Nuts and Bolts
  - ▶ HPR and EE: Layers and Timers
  - ▶ EE Enhancements in V1R5
  - ▶ Coding and Operating EE on z/OS
  - ▶ Testing and Migration for EE
  - ▶ EE Performance