



## **z/OS Communications Server**

---

# **Enterprise Extender: Recent Enhancements and Coming Attractions**

z/OS CS Development, Raleigh  
Sam Reynolds - [samr@us.ibm.com](mailto:samr@us.ibm.com)

© Copyright International Business Machines Corporation 2009. All rights reserved.

---



# Agenda

---

- Enterprise Extender with z/OS V1R9 CS
  - HPR Path Switch Summarization
  - EE LDLC Granularity
  - Local MTU Discovery for EE
- Enterprise Extender with z/OS V1R10 CS
  - Path MTU Discovery for EE
  - RTP Pipe Session Limit Control
  - TGN Parameter for EE Model PUs
  - RTP Display Enhancements
- Enterprise Extender with z/OS V1R11 CS
  - HPR Performance Enhancements
  - Reduction in ECSA Requirements for RTP Pipes
  - EE/IPSec Performance Enhancements
  - Additional EE-Related items in V1R11

---

# **Enterprise Extender with z/OS V1R9 CS**

Copyright International Business Machines Corporation 2009. All rights reserved.

---



# HPR Path Switch Enhancements

---

- When a network failure drives a large number of RTP pipes into path switch state:
  - A significant spike in CPU usage may occur
  - An excessive number of path switch message groups are issued
    - Can be a usability/manageability concern
    - Also can lead to WTO buffer shortages
- V1R9 implemented the following path switch enhancements:
  - Reduced storage allocation/deallocation
  - Optimized handling of path switch timers
  - HPR Path Switch Summarization

# HPR Path Switch Summarization

---

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

# HPR Path Switch Summarization...

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

IST924I -----

IST2192I STARTED = 12

IST2193I TGINOP = 12 SRQTIMER = 0 PSRETRY = 0

IST2194I PARTNER = 0 MNPS = 0 UNAVAILABLE = 0

IST2195I NETWORK = 3 HIGH = 3 MEDIUM = 3 LOW = 3

IST924I -----

IST2196I COMPLETED = 8

IST2195I NETWORK = 2 HIGH = 2 MEDIUM = 2 LOW = 2

IST924I -----

IST2197I FAILED = 4

IST2195I NETWORK = 1 HIGH = 1 MEDIUM = 1 LOW = 1

IST924I -----

IST2198I NETID	STARTED				COMPLETED				FAILED			
IST2199I CPNAME	NET	HI	MED	LOW	NET	HI	MED	LOW	NET	HI	MED	LOW

IST2205I -----

IST2200I NETA	2	2	2	2	1	1	1	1	1	1	1	1
---------------	---	---	---	---	---	---	---	---	---	---	---	---

IST2201I SSCP2A	1	1	1	1	1	1	1	1	0	0	0	0
-----------------	---	---	---	---	---	---	---	---	---	---	---	---

IST2201I SSCP7A	1	1	1	1	0	0	0	0	1	1	1	1
-----------------	---	---	---	---	---	---	---	---	---	---	---	---

IST2205I -----

IST2200I NETB	1	1	1	1	1	1	1	1	0	0	0	0
---------------	---	---	---	---	---	---	---	---	---	---	---	---

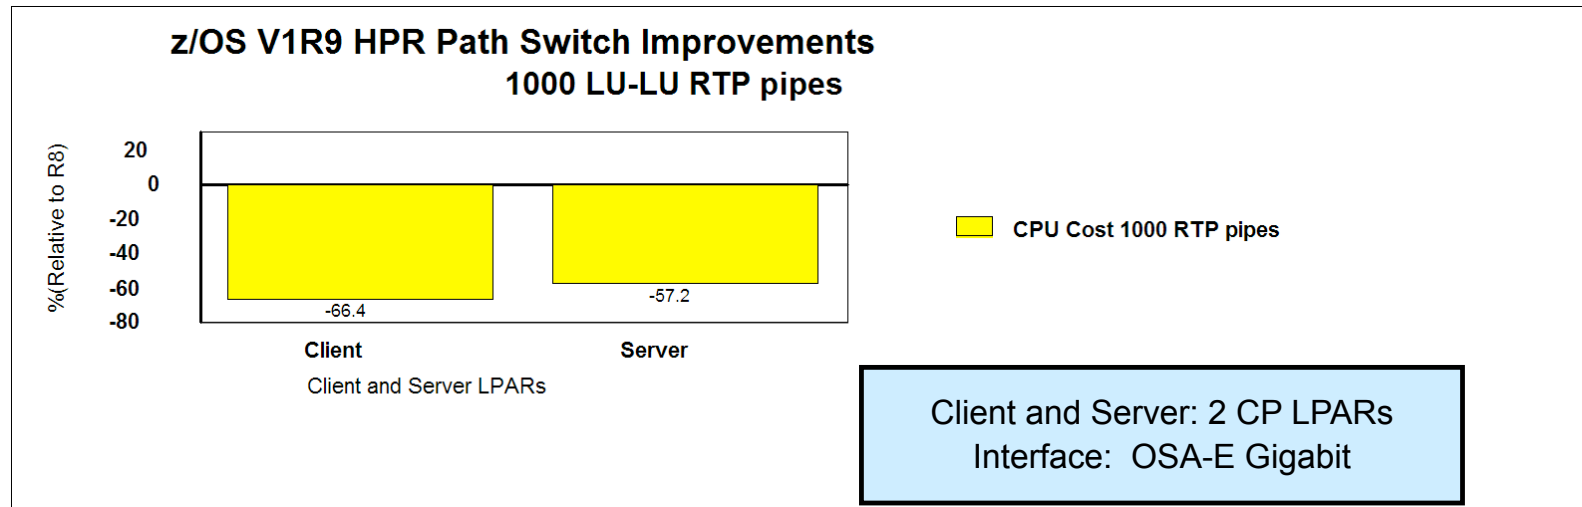
IST2201I SSCP99	1	1	1	1	1	1	1	1	0	0	0	0
-----------------	---	---	---	---	---	---	---	---	---	---	---	---

IST924I -----

IST2206I 24 PATH SWITCH EVENTS FOR 3 CPS IN 2 NETIDS

IST314I END

# HPR Path Switch Summarization...



- V1R8 versus V1R9 path switch performance comparison
  - NN – NN configuration over one-hop EE (1Gb OSA)
    - Two EE TGs defined between the network nodes
  - 1000 LU-LU RTP pipes path switched from one TG to the other
  - V1R8
    - Message IST1494I added to the VTAM message-flooding prevention table
  - V1R9
    - Message IST1494I removed from the VTAM message-flooding prevention table
    - HPRPSMSG = 10
- 57% - 66% reduction in CPU consumption for path switch processing
  - Results will vary depending on configuration

# EE LDLC Review

---

- The EE 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:
    - $LIVTIME + ((SRQRETRY+1) * SRQTIME)$
- Prior to V1R9, only one set of LDLC values can be specified, and that set will govern all EE connections on this node
  - This prevents using different LDLC values for branch vs. data center connectivity, or using different LDLC values for connections to EBN-connected business partners.



# EE LDLC Granularity

- V1R9 provides the ability to specify a different set of LDLC parameters for each local EE VIPA.
- The LDLC parameters are specified on the GROUP statement where the EE VIPA is specified (either via the IPADDR or HOSTNAME operand).
  - If the parameters are coded on multiple GROUPs associated with the same VIPA, the parameters from the first activated GROUP are used.
- The LDLC parameters can still be coded on the PORT
  - GROUP parameters override those on the PORT
  - PORT parameters sift to the GROUP level if not specified there
- LDLC parameters specified on the GROUP can be changed via the VARY ACT,UPDATE command

```
XCAEE      VBUILD TYPE=XCA
PORTEE     PORT   MEDIUM=HPRI, LIVTIME=(10,20), SRQTIME=15, SRQRETRY=3
GRPEE1     GROUP  DIAL=YES, AUTOGEN=(10,E,X), IPADDR=10.1.1.1,          *
              CALL=INOUT, ISTATUS=ACTIVE
GRPEE2     GROUP  DIAL=YES, AUTOGEN=(10,G,V), IPADDR=10.1.1.2,          *
              CALL=INOUT, ISTATUS=ACTIVE, LIVTIME=(10,60), SRQRETRY=5
```

# EE LDLC Granularity...

```
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 = XCAEE
IST2004I LIVTIME = (10,20) 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 10.1.1.1
IST2004I LIVTIME = (10,20) SRQTIME = 15 SRQRETRY = 3
IST2009I RTP PIPES = 2 LU-LU SESSIONS = 1
...
IST924I -----
IST1680I LOCAL IP ADDRESS 10.1.1.2
IST2004I LIVTIME = (10,60) SRQTIME = 15 SRQRETRY = 5
IST2009I RTP PIPES = 2 LU-LU SESSIONS = 1
...
IST924I -----
IST2017I TOTAL RTP PIPES = 4 LU-LU SESSIONS = 2
IST2018I TOTAL ACTIVE PREDEFINED EE CONNECTIONS = 2
IST2019I TOTAL ACTIVE LOCAL VRN EE CONNECTIONS = 0
IST2020I TOTAL ACTIVE GLOBAL VRN EE CONNECTIONS = 0
IST2021I TOTAL ACTIVE EE CONNECTIONS = 2
IST314I END
```

# Local MTU Discovery for EE

---

- Prior to z/OS V1R9, the following rules governed z/OS CS EE connections:
  - MTU size is determined during connection establishment and will not be altered for the duration of the connection.
  - Therefore, the MTU size being utilized for an EE connection may not represent the current value.
    - This can cause EE packet fragmentation which will result in reduced performance
    - May under utilize by transmitting EE packets that are smaller than what currently is permitted by the MTU
  - Associated IP route(s) will be computed during the EE connection establishment. New IP route(s) to the remote EE endpoint will not be utilized unless the existing route is deleted or is no longer active.
    - Therefore, an EE connection may not be utilizing an optimal route between the two endpoints.
    - More optimal routes may become available after the EE connection is established.

# Local MTU Discovery for EE

---

- In z/OS V1R9:
  - When RTP data is being transmitted over an EE connection, changes in the local MTU size will now be learned.
    - Avoids fragmentation of packets being transmitted over an EE connection
    - Better utilizes the EE connection's overall capacity.
  - As more optimal routes are made available for an existing EE connection, they will now be utilized.
    - Avoids using less optimal routes for the life of the EE connection.
- **Note:** RTP pipe(s) routed over an EE connection will only learn of changes in the MTU size when the RTP endpoint and the EE connection reside in the same node, and will only be aware of the MTU size of the local interface. EE is not aware of MTU sizes within the network.

APARs OA25578 and OA25881  
are strongly recommended

---

# **Enterprise Extender with z/OS V1R10 CS**


Copyright International Business Machines Corporation 2009. All rights reserved.

---



# Path MTU Discovery for EE

- z/OS V1R8 CS (and prior):
  - MTU size determined during connection establishment
    - not altered for the duration of the connection.
    - Fragmentation may occur at host and in network
- z/OS V1R9 CS: Local MTU Discovery for EE
  - Local MTU changes are detected
  - Avoids fragmentation at the host, but not in the network
- z/OS V1R10 CS: Path MTU Discovery (PMTU) for EE
  - IPv4: Stacks set "don't fragment" (DF) bit in all outbound IPv4 EE packets
    - IPv6 does not support fragmentation in the network
  - Stack monitors for ICMP/ICMPv6 "Packet too big" messages
  - Stack updates VTAM with learned path MTU
  - Local RTP pipes segment to new size
- PMTU originally architected as TCP-based solution and did not apply to UDP
  - PMTU for EE is an EE-specific adaptation of PMTU, and does not apply to UDP in general



# Path MTU Discovery for EE...

- New modifiable VTAM start option to enable PMTU for EE:
  - PMTUD = TCPVALUE | NO
  - Setting (of defaulting) PMTUD to TCPVALUE indicates that VTAM should accept the TCP/IP setting for PMTU for both IPv4 and IPv6
- MTU size can be viewed on a per-EE-port basis using the detailed view of the DISPLAY EE command:

```
D NET,EE,ID=SWEE42AI,DETAIL
```

```
.  
IST2030I PORT PRIORITY = SIGNAL  
IST2029I MTU SIZE = 1472
```

```
.  
IST924I -----  
IST2031I PORT PRIORITY = NETWORK  
IST2029I MTU SIZE = 1472
```

```
.  
IST924I -----  
IST2032I PORT PRIORITY = HIGH  
IST2029I MTU SIZE = 972
```


```
.  
IST924I -----  
IST2033I PORT PRIORITY = MEDIUM  
IST2029I MTU SIZE = 972
```

```
.  
IST924I -----  
IST2034I PORT PRIORITY = LOW  
IST2029I MTU SIZE = 1472
```

```
.  
.
```

# Path MTU Discovery for EE...

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






# EE MTU Operand

---

- For situations where enabling PMTU for EE is not viable, or is not an adequate solution, V1R10 provides a new operand to allow static configuration of the MTU size to be used for the EE connection:
  - $MTU = mtu\_size$
  - Specifiable on the following definitions:
    - PU in the switched major node
    - GROUP statement in the EE XCA major node (but only for connection network GROUPs)
    - DYNTYPE=EE PU in the model major node
  - The specified MTU value will be reduced by VTAM:
    - Account for the IP and UDP headers
    - Use the lesser of the new MTU value (after subtracting the header sizes) and the learned MTU value provided by the TCP/IP stack (local or path MTU discovery)

# HPR Transmission Stall Support

- An HPR transmission stall occurs when a packet (or packets) is repeatedly dropped in the network, but some communication with the RTP partner is still possible (e.g., for small packets). Possible causes include:
  - PMTU for EE is enabled, but the firewalls are not passing the ICMP "packet too big" messages, so large packets are dropped
  - A multi-hop RTP pipe is routed over EE with PMTU, but the RTP endpoint is not in the same host as the EE connection:



# HPR Transmission Stall Support...

- Detection and notification:

- Notifies operator when RTP pipe suffers a transmission stall
  - Packet resent 6 times, with a minimum of 4 seconds since first transmission

IST2245I XMIT STALL DETECTED FOR RTP puname TO cpname

- Notifies operator at 30-second intervals if stall persists

IST2246I XMIT STALL CONTINUES FOR RTP puname TO cpname


- Notifies operator when stall is alleviated

IST2247I XMIT STALL ALLEVIATED FOR RTP puname TO cpname

- Recovery:

- If HPR packet length exceeds HPR minimum size (768 bytes), reduce HPR maximum NLP size to 768
- Re-segment all outbound data to the new size
- If packets arrive at partner, stall is alleviated. If not, the stall continues.
- After 20 minutes, HPR maximum NLP size is restored to previous value

- Function is automatically enabled



# HPR Transmission Stall Support...

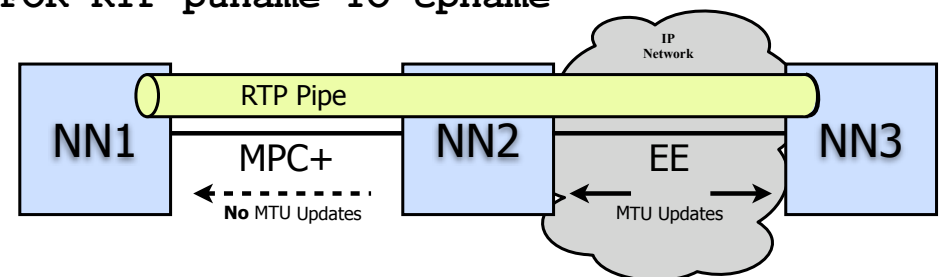
- DISPLAY RTPs can be used to look for pipes in a stalled state:

```
D NET,RTPS,STALL=YES
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = RTPS
IST1695I PU NAME          CP NAME          COSNAME SWITCH CONGEST STALL SESS
IST1960I CNR00004 NETA.SSCP2A          #INTER      NO      NO      YES      5
IST2084I 1 OF 1 MATCHING RTP PIPES DISPLAYED
IST314I END
```

- New modifiable HPRSTALL start option allows automatic termination of continuously-stalled pipes:

- HPRSTALL = 0 | *time\_limit* where *time\_limit* is 30 to 3600 seconds
- Default value of 0 indicates that no automatic termination will occur
- If HPRSTALL is enabled and an RTP pipe remains in a stalled state for *time\_limit* seconds, the pipe will be terminated and the following message issued:

```
IST2253I HPRSTALL TIME EXCEEDED FOR RTP puname TO cpname
```



# RTP Pipe Session Limit Control

---

- In HPR networks, RTP pipes are created (as needed) when new sessions are established
- To minimize the number of RTP pipes created, new sessions are assigned to existing RTP pipes (with matching characteristics: route, class of service, and NCE pair) whenever possible
- Recent customer experience shows the number of sessions assigned to a single RTP pipe is growing
- All data flowing (in or out) over a given RTP pipe is handled by a single process regardless of the number of sessions assigned to that RTP pipe
- Each process uses only one CPU at a time
- If the number of sessions assigned to a single RTP pipe grows large enough, then the CPU becomes a potential performance bottleneck

# RTP Pipe Session Limit Control...

---

- V1R10 provides a new modifiable start option to limit the number of sessions that are assigned to an RTP pipe
  - `HPRSESLM = NOLIMIT | sess_lim` where `sess_lim` is in the range 1000 - 2147483647
- When choosing an RTP pipe for a new session, VTAM will skip over matching RTP pipes that have already reached or exceeded the RTP session limit
- If all matching RTP pipes have already reached or exceeded the RTP session limit, VTAM will create a new RTP pipe for the new session
- The HPRSESLM value used by one APPN node has no affect on the partner RTP endpoint node
  - The partner RTP endpoint can continue to add new sessions to an RTP pipe that has already exceeded this node's HPRSESLM value
    - This is one of several cases where the session count can be exceeded. It should only be regarded as a guideline to VTAM.
  - HPRSESLM works best if the same value is used on all APPN nodes
- APAR OA22854 (V1R7) needs to be applied to any VTAM that will coexist with a V1R10 VTAM

# TGN Parameter for EE Model PUs

- V1R10 adds the capability to specify preferred TG numbers for EE connections defined using DYNPU=YES and a DYNTYPE=EE model PU
- Supports parallel EE TGs (which can exist when multiple EE VIPAs are defined)
- Up to four values can be specified
- ANY can be one of the four values
  - Specifying ANY indicates that if none of the predefined values are available, VTAM will accept other TG numbers for the connection. If ANY is not specified and none of the predefined values are available, the connection establishment will fail.
  - Example: For the EE model PU shown in the sample definition, when we create the DYNPU=YES dynamic PU, we will attempt to assign a TGN of 5. If TGN 5 is not available, we will attempt to assign TGN 6.
  - If TGN 6 is not available, we will allow the connection to establish using any available TG number (which will come from the negotiated TG number space of TGN >= 21)
    - If ANY had not been included (that is, if TGN=(5,6) had been specified), and neither TGN 5 nor TGN 6 were available, the connection establishment would fail

```
MODEL1A  VBUILD  TYPE=MODEL
*
EEMODEL  PU      DYNTYPE=EE,
              CAPACITY=100M,
              COSTTIME=0,
              TGN=(5,6,ANY),
              DISCNT=NO
```

# RTP Display Enhancements

- V1R10 enhances the DISPLAY RTPS command to display all RTPs that meet or exceed a specified retransmission rate
- The DISPLAY RTPS command is also enhanced to allow clearing of the diagnostic counters
- The DISPLAY ID command of an RTP PU (with HPRDIAG=YES) is enhanced to allow clearing of the diagnostic counters
- Example: Display all pipes with retransmission rates meeting or exceeding 0.2%

```
D NET,RTPS,REXMIT=0.2
```

```
IST097I DISPLAY ACCEPTED
```

```
IST350I DISPLAY TYPE = RTPS
```

```
IST1695I PU NAME          CP NAME          COSNAME SWITCH CONGEST STALL SESS
```

```
IST1960I CNR00004 NETA.SSCP2A          #INTER      NO      NO      NO      5
```

```
IST2084I 1 OF 1 MATCHING RTP PIPES DISPLAYED
```

```
IST314I END
```

- Example: Clear the RTP diagnostic counters

```
D NET,RTPS,CLEAR=ALL
```

```
IST097I DISPLAY ACCEPTED
```

```
IST350I DISPLAY TYPE = RTPS
```

```
IST2248I ALL DIAGNOSTIC COUNTERS CLEARED FOR 4 RTP PIPES
```

```
IST314I END
```



---

# **Enterprise Extender with z/OS V1R11 CS**

Copyright International Business Machines Corporation 2009. All rights reserved.


---



# HPR Performance Enhancements - Progressive Mode ARB



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



APAR OA26490  
is strongly  
recommended  
for V1R8, V1R9,  
and V1R10 for  
compatibility

## HPR Performance Enhancements - Progressive Mode ARB...

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

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

- Progressive Mode ARB is also available in Distributed CS V6.4.0 (for Windows, AIX, and Linux (both Intel and System z))
  - The architectural definition of progressive mode ARB is available for other platforms and vendors to implement

# HPR Performance Enhancements - Path Switch Delay



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

A large red circle with a diagonal slash, indicating a prohibition or warning, is overlaid on the console output.

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

# HPR Performance Enhancements - Path Switch Delay...




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

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

## Reduction in ECSA Requirements for RTP Pipes

- Prior to V1R11, each RTP pipe is represented by a control block in ECSA
- In V1R11, a large portion of the RTP control block was moved to an extension control block in VTAM private storage, resulting in a significant ECSA savings for installations with a large number of RTP pipes
- Preliminary estimates of the reduction in required RTP pipe ECSA storage for various RTP counts:

RTP Pipes	ECSA Reduction
4000	11.5%
12000	23.5%
20000	29.5%



# EE/IPSec Performance Enhancements



- Improved performance for EE over IPSec
  - The “bursty” nature of HPR traffic can cause significant performance degradation when it is carried over IPSec tunnels
    - Smaller bursts frequently get encrypted and sent before larger bursts, resulting in out-of-order segments that will be dropped at the other end of the IPSec tunnel, forcing retransmits.
  - In V1R11, large bursts over EE (with IPSec enabled) are broken into small batches
    - Allows them to be encrypted and sent to the partner node before subsequent smaller bursts, thereby avoiding the out-of-order problem at the partner node
- Improved support for EE/IPSec over zIIP
  - Support for offloading outbound EE/IPSec traffic to a zIIP processor. Previously only inbound traffic would be processed on the zIIP.
  - Care taken to avoid redispach to a zIIP or general CPU if benefit would not be worth the context switch

To be available  
on V1R10 via  
APAR PK93190



## Additional EE-Related Items in V1R11

- RTP will now free packets for fully-acknowledged PIUs, even after a gap
  - Previously, VTAM held the packets until the gap was filled, which was not required by the architecture (or necessary)
  - This helps reduce storage spikes associated with dropped packets
- The REXMIT operand on the DISPLAY EEDIAG command now allows the specification of fractional percentages

```
D NET,EEDIAG,REXMIT=0.3,IP=197.51.153.1,LIST=DETAIL
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = EEDIAG
IST2065I ENTERPRISE EXTENDER CONNECTION REXMIT INFORMATION
IST2067I EEDIAG DISPLAY ISSUED ON 02/25/09 AT 13:31:05
IST924I -----
IST1680I LOCAL IP ADDRESS 197.51.153.1
IST1910I LOCAL HOSTNAME HOVIP1
IST1680I REMOTE IP ADDRESS 197.11.115.1
IST1909I REMOTE HOSTNAME SLVIP1
IST2024I CONNECTED TO SWITCHED PU EHO2SL
.
.
IST924I -----
IST2035I TOTALS FOR ALL PORT PRIORITIES
IST2036I   NLPS SENT           =           9200   ( 009K )
IST2038I   NLPS RETRANSMITTED =             67   ( 000K )
IST2068I   NLP RETRANSMIT RATE =           0.73%
IST2069I REXMIT COUNTERS LAST CLEARED ON 02/25/09 AT 13:20:42
IST2042I 1 OF 1 EE CONNECTIONS DISPLAYED
IST314I END
```



# For More Information...

URL	Content
<a href="http://www.twitter.com/IBM_Commserver">http://www.twitter.com/IBM_Commserver</a> 	IBM Communications Server Twitter Feed
<a href="http://www.facebook.com/IBMCommserver">http://www.facebook.com/IBMCommserver</a> 	IBM Communications Server Facebook Fan Page
<a href="http://www.ibm.com/systems/z/">http://www.ibm.com/systems/z/</a>	IBM System z
<a href="http://www.ibm.com/systems/z/hardware/networking/index.html">http://www.ibm.com/systems/z/hardware/networking/index.html</a>	IBM System z Networking
<a href="http://www.ibm.com/software/network/commserver/zos/">http://www.ibm.com/software/network/commserver/zos/</a>	IBM 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>	IBM Communications Server for Linux on zSeries
<a href="http://www.ibm.com/software/network/ccl/">http://www.ibm.com/software/network/ccl/</a>	IBM Communication Controller for Linux on System z
<a href="http://www.ibm.com/software/network/commserver/library">http://www.ibm.com/software/network/commserver/library</a>	IBM Communications Server Library - white papers, product documentation, etc.
<a href="http://www.redbooks.ibm.com">http://www.redbooks.ibm.com</a>	IBM Redbooks
<a href="http://www.ibm.com/software/network/commserver/support">http://www.ibm.com/software/network/commserver/support</a>	IBM 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 (RFCs)
<a href="http://publib.boulder.ibm.com/infocenter/ieduasst/stgv1r0/index.jsp">http://publib.boulder.ibm.com/infocenter/ieduasst/stgv1r0/index.jsp</a>	IBM Education Assistant

- Recommended Redbooks:

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