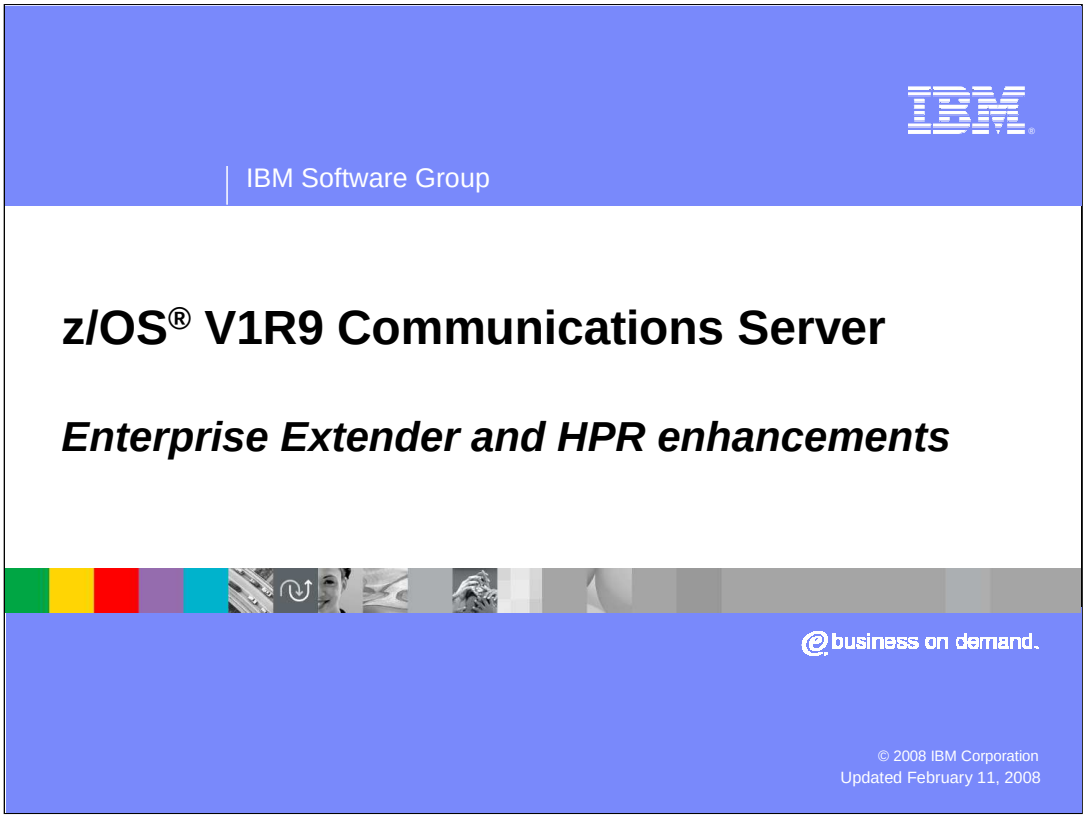

This presentation describes the Enterprise Extender and HPR enhancements for z/OS V1R9 Communications Server.

Agenda

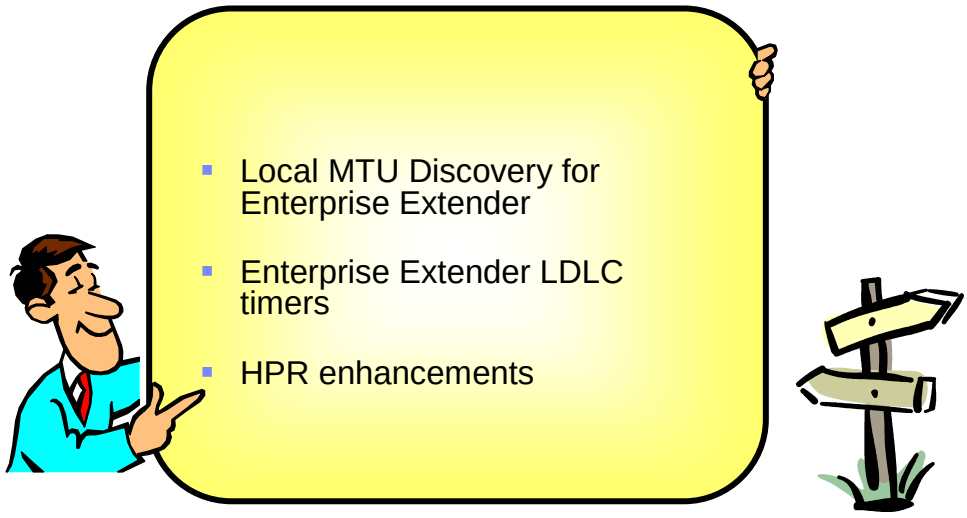

z/OS V1R9 Communications Server contains multiple Enterprise Extender and HPR enhancements. Each of these enhancements is discussed in this presentation.

Section

Local MTU discovery for Enterprise Extender

This section describes the new capabilities to detect changes in an EE connection's Maximum Transmission Unit (MTU) size and underlying IP route.

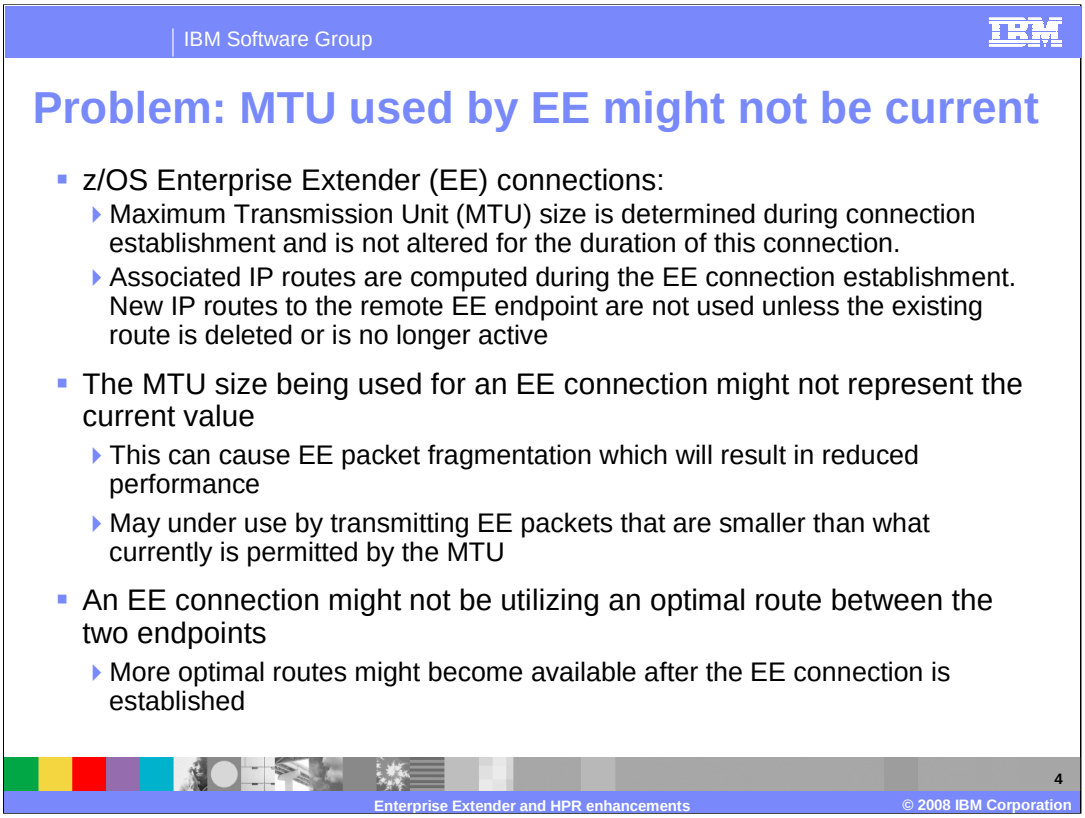

In prior releases an Enterprise Extender (EE) connection will obtain the minimum MTU size permitted for packets being transmitted to a remote EE endpoint at the time that the connection is **initially** established. It is important that VTAM® understand the maximum MTU size for this connection to avoid fragmentation by the TCP/IP stack which will greatly increase the path length for a transmitted EE packets in the TCP/IP stack. Therefore today if a new interface is being used by TCP/IP for an existing EE connection (previous route obtained might have been deleted or became inactive) and the new MTU size is smaller than was previously reported to VTAM (at EE connection initialization) then there are negative performance implications.

Additionally there are changes needed to allow TCP/IP to attempt to obtain a new route for an existing EE connection when updates have been made to the IP routing table (by OMPROUTE, policy changes and so on.). When changes have been made to the IP routing table a more optimal route can perhaps be determined for an EE connection. Currently if TCP/IP obtains a route handle for an EE connection and that route is associated with a default route then there is no way to ever move from this default route without SNA terminating the EE connection. Therefore there have been several users that start their VTAM and TCP/IP connection where an EE connection is initiated by VTAM before TCP/IP learns about all of the potential routes from OMPROUTE and therefore the EE connection ends up using the default route (which in many cases is not the optimal route).

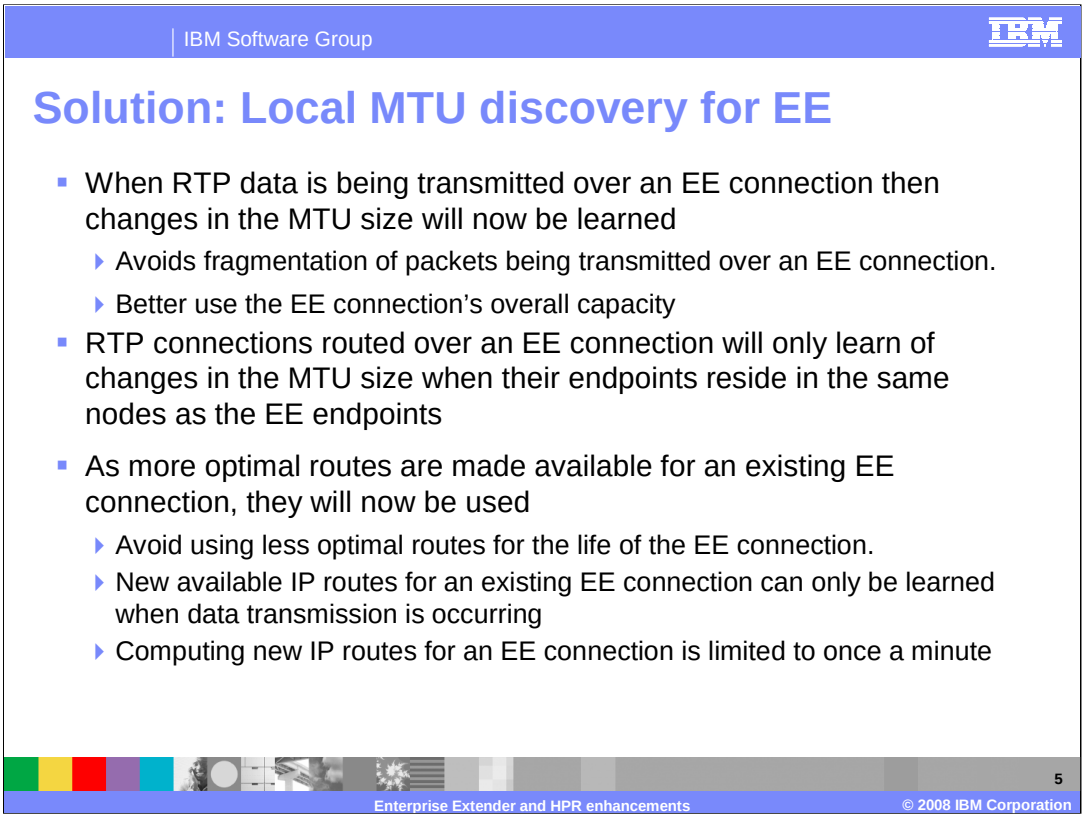

Enterprise Extender autonomics is associated with two potential performance enhancements:

1) Allow for VTAM to learn of changing MTU sizes associated with an Enterprise Extender connection. This permits the avoidance of packet fragmentation when the MTU size is decreased. And in some rare cases for VTAM to pass larger packets to TCP/IP to better use the current interface associated with an EE connection.

2) When new routes are learned by TCP/IP allow for the determination of a more optimal route for an existing Enterprise Extender connection.

An RTP connection routed over an Enterprise Extender (EE) connection will learn of changes in the MTU size only when their endpoints reside in the same nodes as the EE endpoints.

The RTP connection's network layer packet (NLP) size can only be increased to the maximum packet size returned by the Route_Setup signal exchanged during RTP initialization or path switch. The RTP connection's NLP size is negotiated to be a value no smaller than 1492 for the traversal of data over an EE connection. If the EE connection MTU size is less than 768 bytes, VTAM sets the maximum NLP packet size to 768 (this is the smallest maximum packet size allowed by VTAM for HPR packets). This limitation can cause TCP/IP to fragment but exists because the RTP layer cannot allow the HPR header to be segmented in the RTP layer.

Display NET,EE

- Display NET,EE,ID=name,DETail to display information about an Enterprise Extender connection

```
D NET,EE,ID=SWEE42AI,DETAIL
IST097I DISPLAY ACCEPTED
     .
IST924I -------------------------------------------------------------
IST2030I PORT PRIORITY = SIGNAL
IST2029I   MTU SIZE =   572
     .
IST924I -------------------------------------------------------------
IST2031I PORT PRIORITY = NETWORK
IST2029I   MTU SIZE =   972
     .
IST924I -------------------------------------------------------------
IST2032I PORT PRIORITY = HIGH
IST2029I   MTU SIZE =  1472
     .
IST924I -------------------------------------------------------------
IST2033I PORT PRIORITY = MEDIUM
IST2029I   MTU SIZE =   1472
     .
IST924I -------------------------------------------------------------
IST2034I PORT PRIORITY = LOW
IST2029I   MTU SIZE =  1472
```


When VTAM detects this condition the EE connection's MTU size during the transmission of an NLP then the MTU size is altered (this change can be seen on the message IST2029I when you issue the DISPLAY NET,EE command). Policy-based routing is obviously being used in this example. Policy-based routing tables have been installed which directs the traffic for some of the ports over different interfaces which have different MTU sizes.

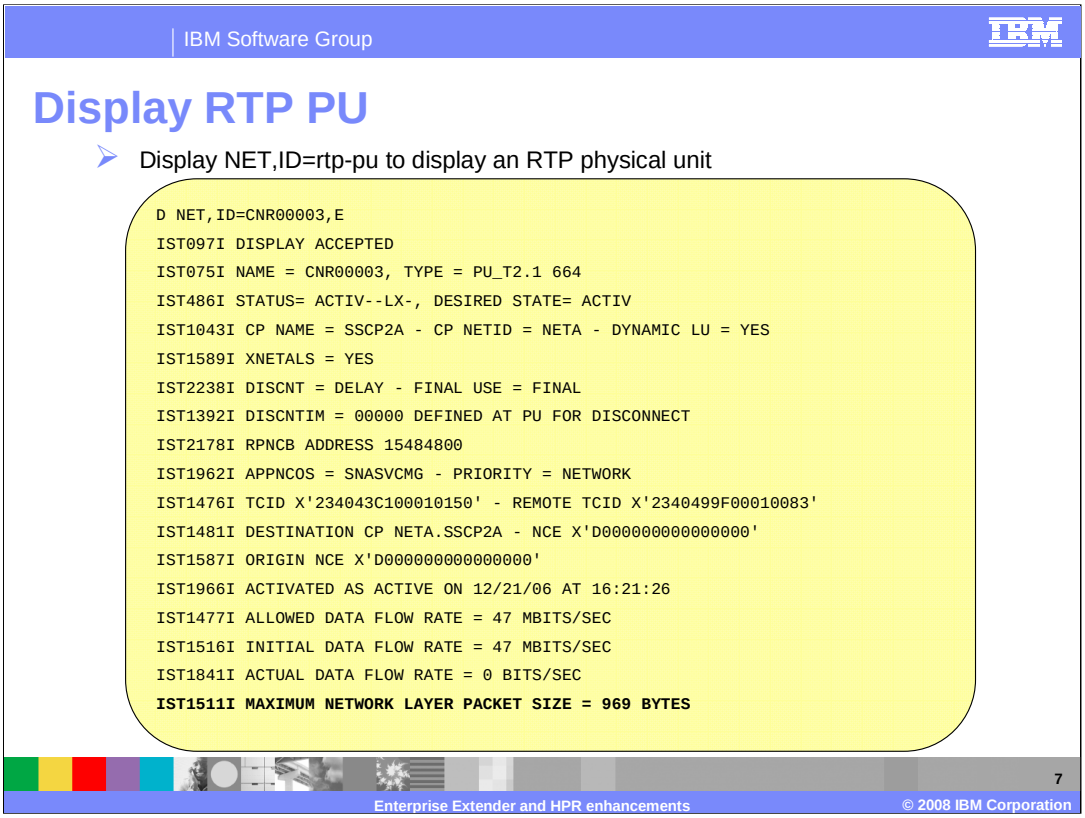

If a change in the EE MTU size also alters the permitted NLP size (NLP size cannot be increased beyond the originally negotiated value for the RTP connection) then this change can be viewed in the IST1511I message which is displayed on the DISPLAY ID=rtp-pu command.

Section

Enterprise Extender LDLC timers

This section describes the changes to the Enterprise Extender LDLC timers function.

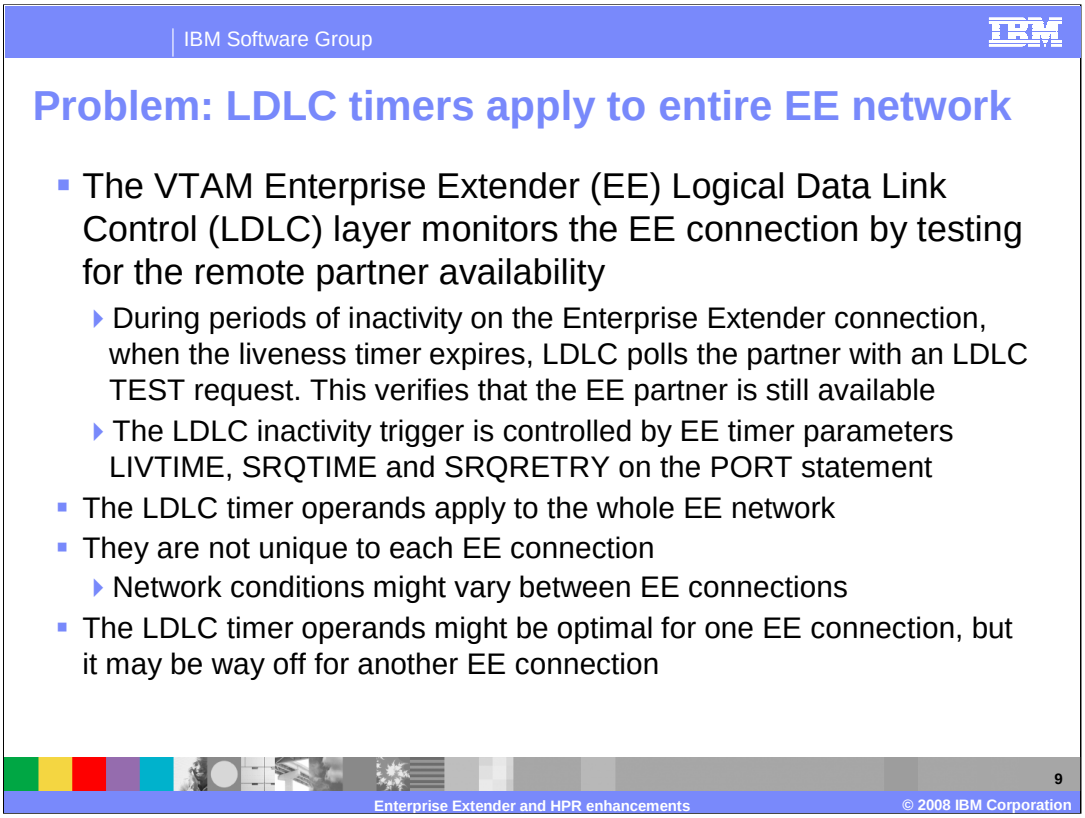

LIVTIME specifies the Enterprise Extender logical data link control liveness timer interval range, in seconds. Two values can be specified on the LDLC liveness timer (LIVTIME). These values are optional with the first being the initial LIVTIME value (init_value) and the second the maximum LIVTIME value (max_value). Specifying a max_value larger than the init_value enables the EE LDLC Keep-Alive Reduction Function. This function enables the current LIVTIME window to expand and contract based on current network conditions. Expanding the current LIVTIME window reduces the number of LDLC test flows that occur during periods of inactivity. SRQTIME specifies the Enterprise Extender logical link control short request timer interval in seconds. SRQRETRY specifies the number of times the short request timer is retried before the port becomes inoperative.
The LDLC layer monitors the EE connection, sending a test frame if no inbound activity is detected for the number of seconds specified by the LIVTIME operand as coded (or defaulted). If no response is received for the number of seconds specified by the SRQTIME operand, another test frame is issued. As long as no response is received, LDLC retries SRQRETRY. If no response is received after the last retry, the EE link will be disconnected.

z/OS Communications Server currently does not provide flexibility to EE LDLC timers. The LDLC timer operands apply to the entire EE network. They are not unique to each EE connection. Since network conditions can vary between connections, the operands can be optimal for one EE connection and not optimal for another EE connection.

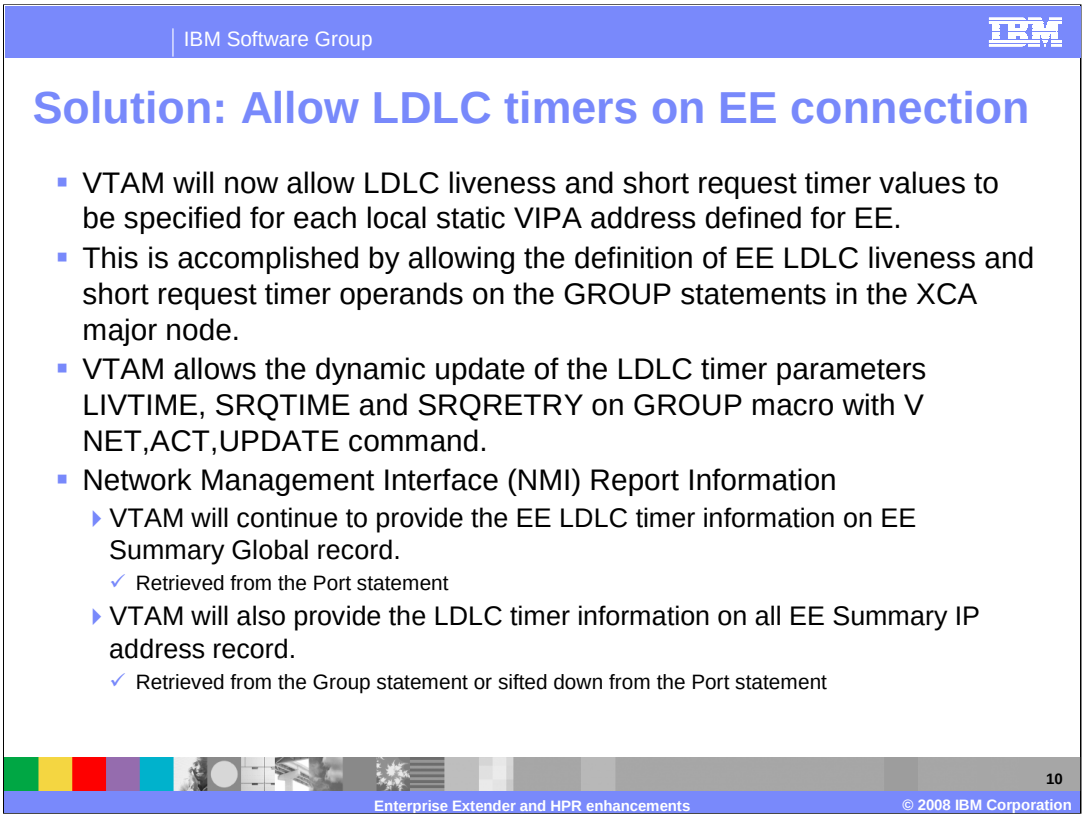

In z/OS V1R9 Communications Server, VTAM will provide the LDLC timers for each static VIPA by allowing the timer operands on the GROUP statement in an XCA major node. VTAM will allow the dynamic update of the LDLC timer parameters on GROUP statement with V NET,ACT,UPDATE command. The PORT statement defines the system wide EE LDLC timer parameters. These values are used if they are not specified on the GROUP statement. If the GROUP statement has one or more EE timer parameters, they will override the EE timer parameters of the PORT statement for this GROUP. If the GROUP statement has only one LDLC timer parameter specified, VTAM will sift down the other two LDLC timer parameters from PORT statement.
If two or more GROUPs are using the same static VIPA and they have the different EE timer parameters, VTAM will use the EE timer parameters of the first activated GROUP for that VIPA. Other GROUPs will use the EE timer parameters specified on the first activated GROUP.

VTAM provides EE LDLC timer parameters LIVTIME, SRQTIME and SRQRETRY information on EE Summary Global record. VTAM will also provide EE LDLC timer parameters LIVTIME, SRQTIME and SRQRETRY information on EE Summary IP address record. The EE Summary IP address record contains these fields: EESumIP_Timer_LIVTIME, EESumIP_Timer_SRQTIME, and EESumIP_Timer_SRQRETRY.

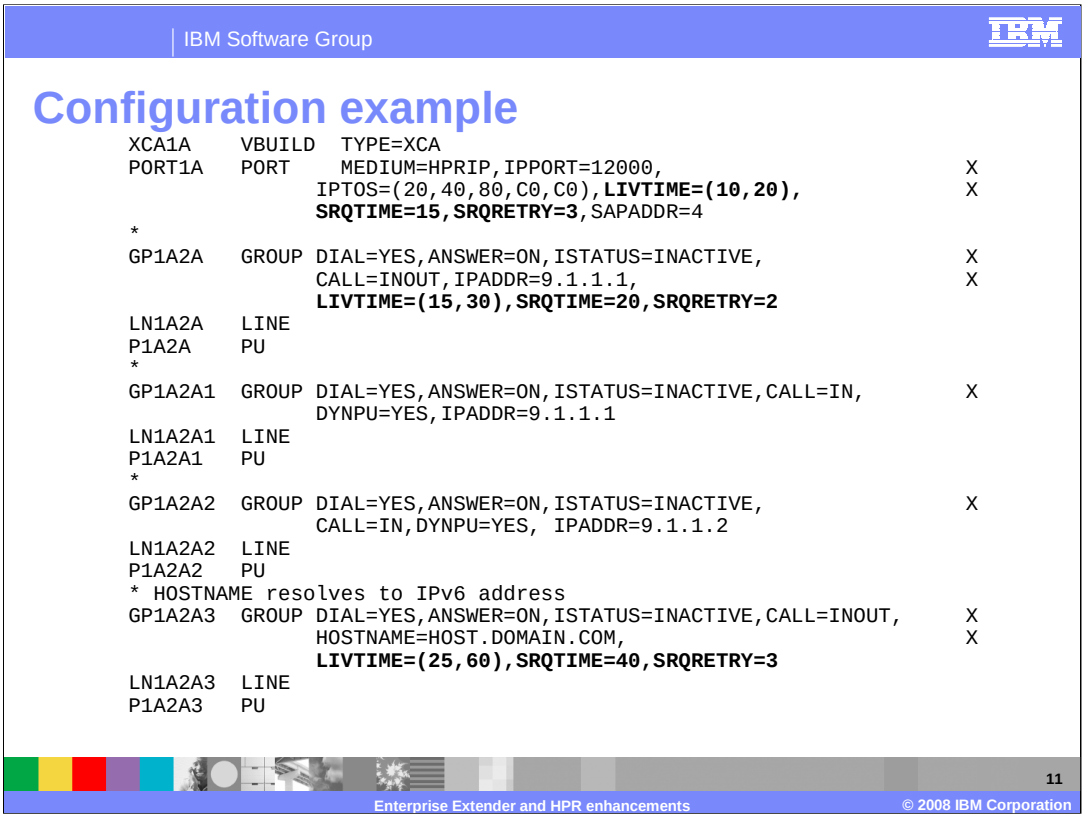

This is the example of sample XCA Major node definition. Two GROUPs, GP1A2A and GP1A2A1, have the same static VIPA IP address. Group GP1A2A has LDLC timer parameters specified and they are different than the LDLC timer parameters specified on PORT definition statement. Group GP1A2A1 does not have LDLC timer parameters specified so it will inherit the LDLC timer parameters value from the PORT definition statement (sift down effect). VTAM uses the LDLC timer parameters of the first activated group, GP1A2A or GP1A2A1, for this static VIPA. When the second GROUP is activated, it will receive the error message indicating that it is using the LDLC timer parameters of the first activated GROUP.

Group GP1A2A2 does not have LDLC timer parameters specified so it will inherit the LDLC timer parameters value from the PORT definition statement (sift down effect). Group GP1A2A3 has LDLC timer parameters specified and they are different than the LDLC timer parameters specified on PORT definition statement. Both GROUPs have different static VIPA address. So, they will use their own LDLC timer parameters.

Display EE detail example

```
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 = (10,20) 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 9.67.1.5
IST2004I LIVTIME = (10,20) SRQTIME = 15 SRQRETRY = 3
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 -------------------------------------------------------------
IST1680I LOCAL IP ADDRESS 9.67.1.3
IST2004I LIVTIME = (10,20) SRQTIME = 20 SRQRETRY = 4
IST2009I RTP PIPES =          2      LU-LU SESSIONS      =          1
IST2010I INOPS DUE TO SRQRETRY EXPIRATION                 =          0
IST1324I VNNAME = IP.GVRN3           VNGROUP = GPIP3     (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
```


An excerpt of a Display EE Detail output is shown on this slide. The first IST2004I message shows the LDLC timer values from the PORT definition statement. Subsequent IST2004I messages show the LDLC timers definitions coded on GROUP definition statement or sifted down from PORT definition statement.

The D NET,EE,HOSTNAME= and the D NET,EEDIAG,HOSTNAME= reports also contain message IST2004I which shows the LDLC timer parameters defined on GROUP statement or sifted down from PORT definition statement.

The D NET,EE,LIST=SUMMARY report contains message IST2004I which shows the LDLC timers definitions coded or defaulted on PORT definition statement.

Section

HPR enhancements

This section describes the HPR enhancements made in the z/OS V1R9 Communications Server.

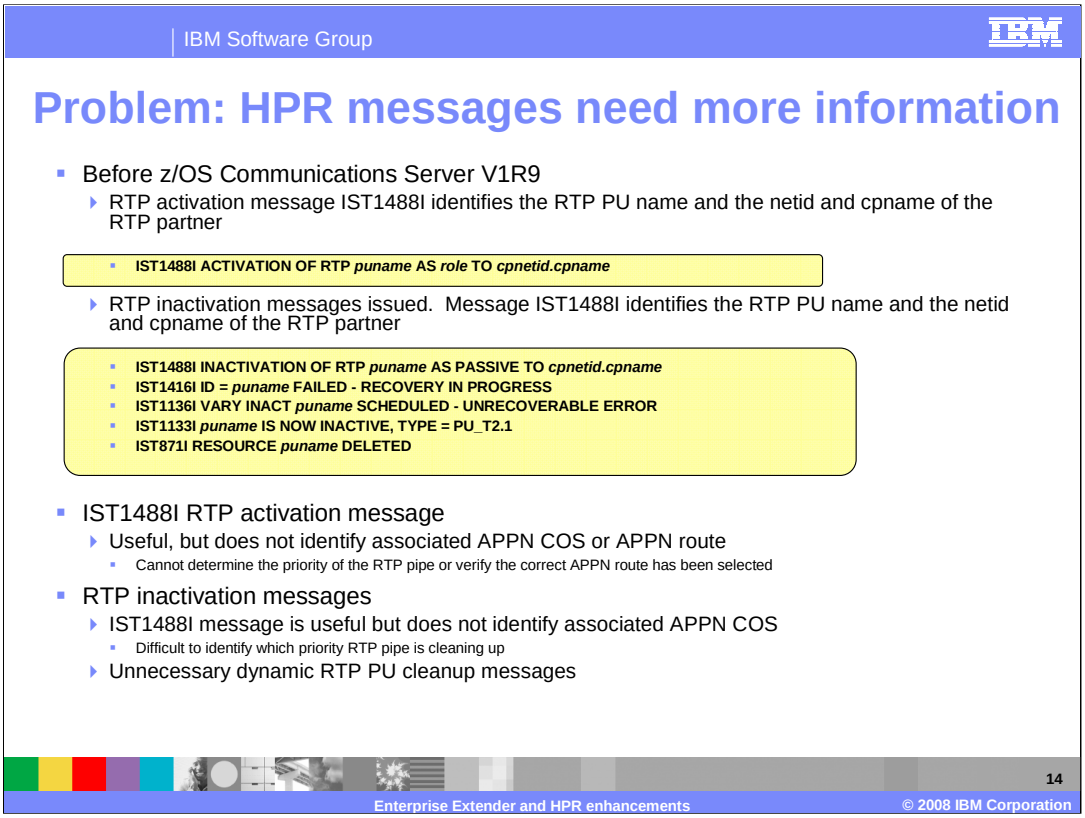

In previous releases, when an RTP pipe is activated a single message, IST1488I, is issued by VTAM to identify the RTP PU name and the associated NETID and CPNAME of the RTP partner. When an RTP pipe is inactivated, VTAM also issued message IST1488I to identify the RTP PU name and the associated NETID and CPNAME of the RTP partner. In addition to the IST1488I message, VTAM issued a few other dynamic PU cleanup messages including IST1416I, IST1136I, ISt1133I ad IST871I.

During RTP pipe activation, VTAM does not identify the APPN COS or APPN route associated with the RTP pipe. Without this information, you cannot identify the priority of the RTP pipe or verify the correct APPN route has been selected.

During RTP pipe inactivation, VTAM does not identify the associated APPN COS. Without this information, it is difficult to identify which priority RTP pipe is cleaning up. Also, many of the dynamic RTP PU cleanup messages are unnecessarily issued. If a large number of RTP pipes are cleaning up, this may lead to hundreds or thousands of unnecessary messages being issued to the system console.

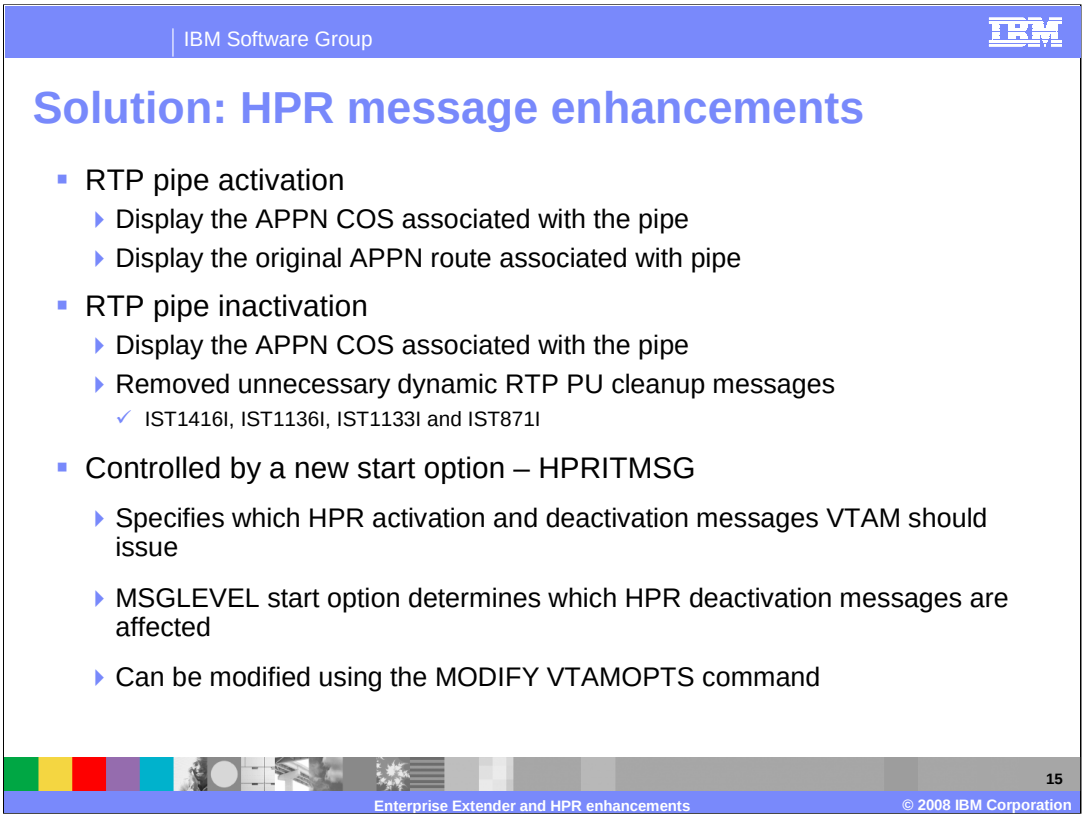

During RTP pipe activation, VTAM may now display the APPN COS or APPN route associated with the RTP pipe. With this new information, you can easily identify the priority of the RTP pipe which is activating. You can also verify the correct APPN route has been selected.

During RTP pipe inactivation, VTAM may now display the associated APPN COS. This new information makes it easier for the operator to understand which priority traffic is ending. Also, the dynamic RTP PU cleanup messages are no longer issued. Removing these unnecessary messages helps cleanup the system log so the system operator can focus on more important messages.

The new HPR inactivation and deactivation message enhancements are controlled by the new start option HPRITMSG. This start option may be specified as BASE or ENHANCED. The default is set to BASE which means VTAM will issue the base RTP pipe activation and inactivation messages. When this start option is set to ENHANCED, VTAM will issue the enhanced versions of the RTP pipe activation and inactivation messages.

Since some of the messages affected by this start option have different versions, the exact message affected in the output depends on the value specified on the MSGLEVEL start option.

If the current value of the HPRITMSG is not appropriate for your system, you can modify the HPRITMSG value by using the MODIFY VTAMOPTS command.

The HPRITMSG start option is only valid if VTAM provides RTP level HPR support.

HPR message enhancements: Activation message example

- Enhanced RTP activation message group:

```
IST1488I ACTIVATION OF RTP puname AS role TO cpnetid.cpname
[IST1962I APPNCOS = appncos_name- PRIORITY = NETWORK]
[IST1963I APPNCOS = appncos_name- PRIORITY = HIGH]
[IST1964I APPNCOS = appncos_name- PRIORITY = MEDIUM]
[IST1965I APPNCOS = appncos_name- PRIORITY = LOW]
[IST1480I RTP END TO END ROUTE - RSCV PATH]
[IST1460I TGN  CPNAME             TG TYPE      HPR]
[IST1461I tgn  cpname             tgtype       hpr]
.
.
IST314I END
```


When the HPRITMSG is set to ENHANCED, VTAM will now display the APPN COS and APPN route when an RTP pipe is activated.

HPR message enhancements: Deactivation message example

- Enhanced RTP deactivation message group:

```
IST1488I INACTIVATION OF RTP puname AS role TO cpnetid.cpname
[IST1962I APPNCOS = appncos_name- PRIORITY = NETWORK]
[IST1963I APPNCOS = appncos_name- PRIORITY = HIGH]
[IST1964I APPNCOS = appncos_name- PRIORITY = MEDIUM]
[IST1965I APPNCOS = appncos_name- PRIORITY = LOW]
IST314I END
```

- V NET,INACT,ID=rtppuname,F command
  - Message IST105I or IST1133I will still be issued back to the console.

When the HPRITMSG is set to ENHANCED, VTAM will now display the APPN COS when an RTP pipe is inactivated. VTAM also no longer issues the dynamic PU cleanup messages when an RTP pipe is inactivated. One exception is when a vary inactivate command is issued against an RTP pipe. In this case, message IST105I or IST1133I is still issued to the console so the operator receives a response to the vary inactivate command.

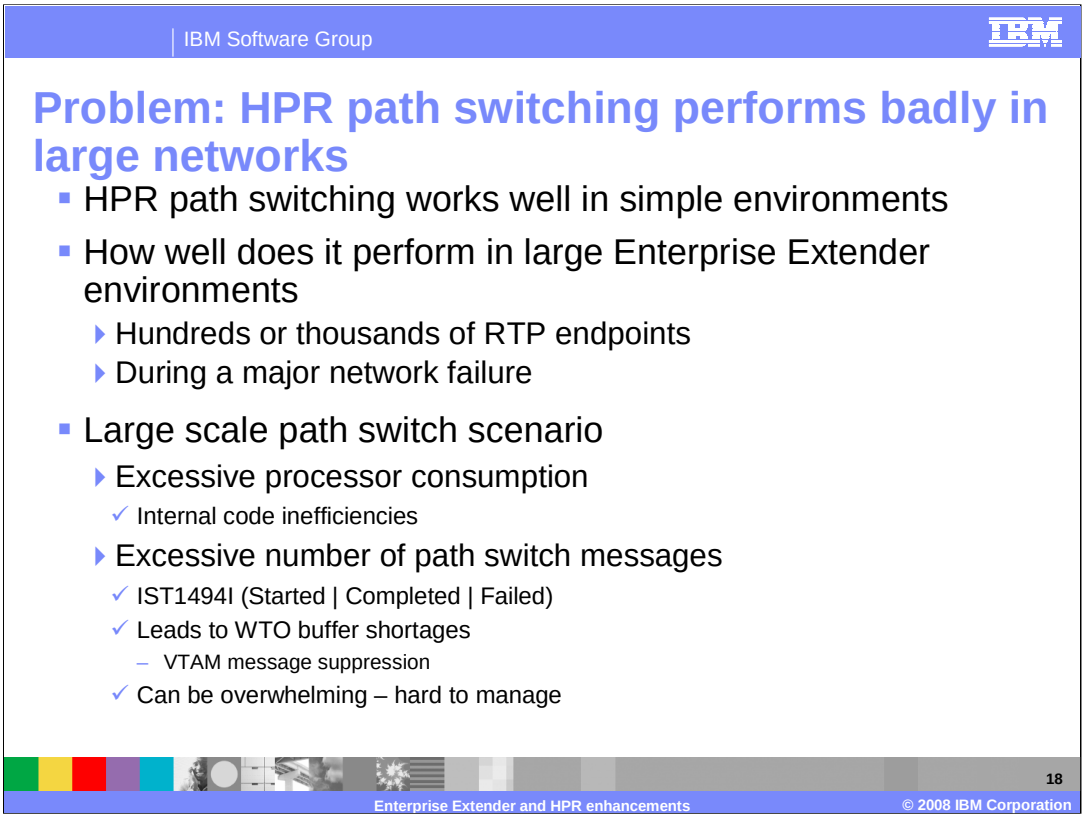

HPR path switching works well in simple environments, but there have been concerns raised on how well it performs during a large network failure.

During a large scale path switch scenario, VTAM consumes too much processor and issues too many path switch messages. In addition, the HPR path switch messages can be overwhelming and hard to manage.

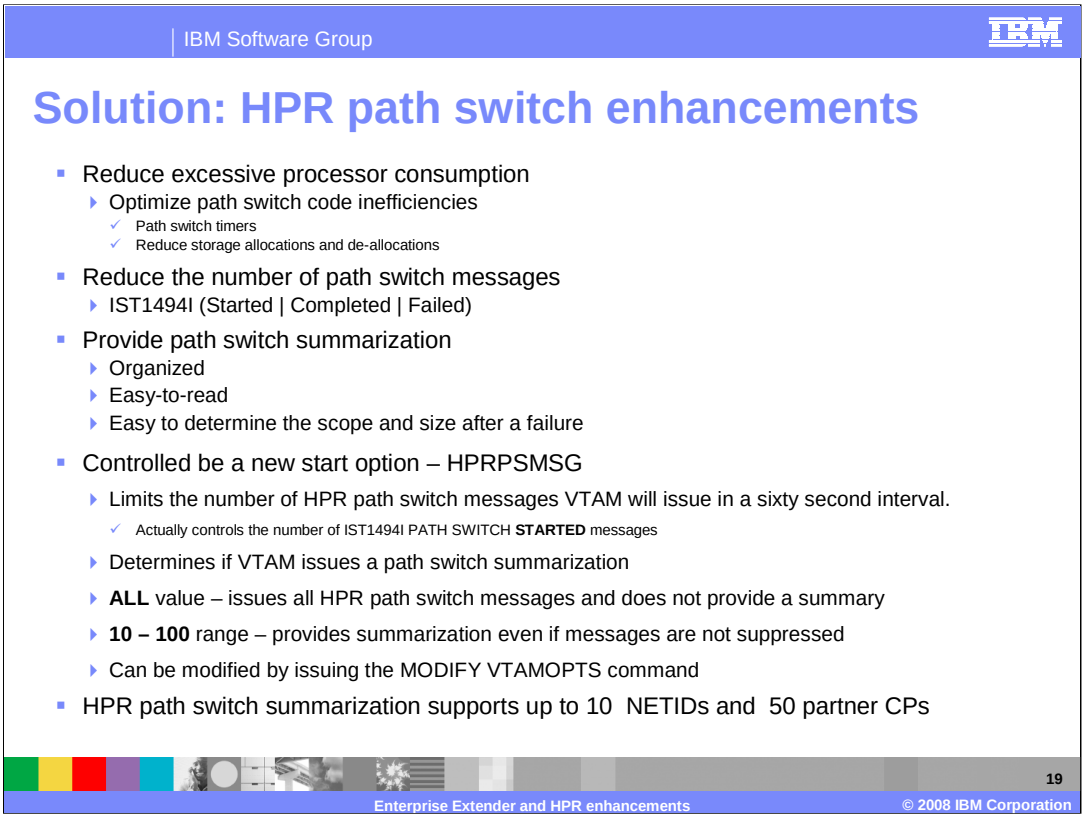

Changes have been made to the internal HPR path switch code to optimize path switch code to reduce processor usage. Changes have also been made to reduce the number of path switch messages issued to the console. In addition, VTAM will also output a path switch summarization display to document all the associated path switch events which occurred during a given time interval.

The HPR path switch enhancements are controlled by a new start option, HPRPSMSG. The default of ALL issues all HPR path switch message and does not provide a summary display of the path switch events. Specifying a value between 10 and 100 allows VTAM to limit the number of HPR path switch messages issued to the console in a sixty second interval. If a STARTED message is issued for a pipe, the COMPLETED or FAILED message is always issued. The Sixty second interval starts when a path switch event occurs.

The path switch message summary is always issued to the console at the end of the time interval, whether or not messages were suppressed. The summary will include all path switch events which occurred during the given path switch event time interval, including path switch event information that was issued to the console.

The HPRPSMSG start option is only valid when VTAM provides RTP level HPR support. If the current value of the HPRPSMSG is not appropriate for your system, you can modify the HPRPSMSG value by using the MODIFY VTAMOPTS command.

HPR path switch summarization supports 10 NETIDs and 50 partner CPs. If a large outage exceeds these limits, the report will not contain all NETID and CP specific counts. Message IST2206I will show greater than 10 NETIDs if there was a NETID overflow and greater than 50 CPs if there was a CP overflow. The path switch started (IST2192I), completed (IST2196I), and failed (IST2197I) message counts are always accurate. Even if overflow occurs.) The NETID and CP specific counts are limited to 999. IST2200I and IST2201I will display 999 when count is 999 or higher. You can have multiple path switch summaries per outage because HPRPST may have different values for each pipe priorities.

HPR path switch enhancements: Summarization example 1

```
IST2191I HPR PATH SWITCH SUMMARY FROM 04/05/06 AT 09:45:14
IST924I -------------------------------------------------------------
IST2192I STARTED   =     2
IST2193I   TGINOP  =     0    SRQTIMER =     2    PSRETRY     =     0
IST2194I   PARTNER =     0    MNPS     =     0    UNAVAILABLE =     0
IST2195I   NETWORK =     1  HIGH =     1  MEDIUM =     0  LOW =     0
IST924I -------------------------------------------------------------
IST2196I COMPLETED =     2
IST2195I   NETWORK =     1  HIGH =     1  MEDIUM =     0  LOW =     0
IST924I -------------------------------------------------------------
IST2197I FAILED    =     0
IST2195I   NETWORK =     0  HIGH =     0  MEDIUM =     0  LOW =     0
IST924I -------------------------------------------------------------
IST2198I NETID          STARTED          COMPLETED         FAILED
IST2199I   CPNAME   NET  HI MED LOW  NET  HI MED LOW  NET  HI MED LOW
IST2205I ---------  ---------------  ---------------  ---------------
IST2200I NETA         1   1   0   0    1   1   0   0    0   0   0   0
IST2201I   SSCP1A     1   1   0   0    1   1   0   0    0   0   0   0
IST2206I 4 PATH SWITCH EVENTS FOR 1 CPS IN 1 NETIDS
IST314I END
```


This is an example of an HPR path switch summarization display. The output is organized into four basic sections. The first section display the number of RTP pipes which entered path switch during this interval. This section also displays counts by path switch reason and by the associated pipe priority. The second section display the number of RTP pipes which successfully completed path switch. This section also display the path switch counts by the associated pipe priority. The third section displays the number of RTP pipes which unsuccessfully path switched. This section also displays the path switch counts by the associated pipe priority. The fourth section organizes the path switch information by NETID and CPNAME.

The path switch summarization shows all the path switch activity from 09:45:14 until the current time (approximately 09:46:14). During this interval, 2 RTP pipes path switched due to timeouts (Short Request Timer expiration). These 2 pipes consisted of 1 network priority pipe (CP-CP) and 1 high priority pipe. During this interval, the same two RTP pipes successfully completed path switch. During this interval, no RTP pipes failed path switch. The display shows a breakdown of path switch events by each CP within each NETID. In this example, the output clearly identifies the problem is isolated to the connectivity specific to this one partner node.

HPR path switch enhancements: Summarization example 2

```
IST2191I HPR PATH SWITCH SUMMARY FROM 04/27/06 AT 06:22:11
IST924I -------------------------------------------------------------
IST2192I STARTED   =    20
IST2193I   TGINOP  =    20    SRQTIMER =     0    PSRETRY     =     0
IST2194I   PARTNER =     0    MNPS     =     0    UNAVAILABLE =     0
IST2195I   NETWORK =     5  HIGH =     5  MEDIUM =     5  LOW =     5
IST924I -------------------------------------------------------------
IST2196I COMPLETED =     0
IST2195I   NETWORK =     0  HIGH =     0  MEDIUM =     0  LOW =     0
IST924I -------------------------------------------------------------
IST2197I FAILED    =     0
IST2195I   NETWORK =     0  HIGH =     0  MEDIUM =     0  LOW =     0
IST924I -------------------------------------------------------------
IST2198I NETID          STARTED          COMPLETED         FAILED
IST2199I   CPNAME   NET  HI MED LOW  NET  HI MED LOW  NET  HI MED LOW
IST2205I ---------  ---------------  ---------------  ---------------
IST2200I NETA         5   5   5   5    0   0   0   0    0   0   0   0
IST2201I   SSCP3A     1   1   1   1    0   0   0   0    0   0   0   0
IST2201I   SSCP7A     1   1   1   1    0   0   0   0    0   0   0   0
IST2201I   SSCP99     1   1   1   1    0   0   0   0    0   0   0   0
IST2201I   SSCP7B     1   1   1   1    0   0   0   0    0   0   0   0
IST2201I   SSCP2AB    1   1   1   1    0   0   0   0    0   0   0   0
IST924I -------------------------------------------------------------
IST2206I 20 PATH SWITCH EVENTS FOR 5 CPS IN 1 NETIDS
IST314I END
```


In some cases, when path switches do not complete quickly, the path switch information associated with a given RTP pipe may be spread across multiple path switch summary message groups. For example, this may occur when HPRPSMSG=10 and HPRPST =(4M,4M,2M,2M) are specified, and twenty RTP pipes (5 low_priority, 5 medium_priority, 5 high_priority, 5 network_priority) enter path switch state. The path switch message event time interval is set when the first pipe enters path switch state. The first ten RTP pipes will have IST1494I (Started) message groups issued to the console. The next ten RTP pipes will not have IST1494I(Started) message groups issued to the console. In this example, there is not an alternate route available. The RTP pipes will stay in path switch state until their respective HPRPST timers expire, at which point the path switches will fail. Since the HPRPST values for the various priorities are specified as different values, the path switches will fail at different times. When the current path switch message event interval ends, a summary of twenty path switch started events is issued to the console.

HPR path switch enhancements: Summarization example 2 - Continued

```
IST2191I HPR PATH SWITCH SUMMARY FROM 04/27/06 AT 06:23:11
IST924I -------------------------------------------------------------
IST2192I STARTED   =     0
IST2193I   TGINOP  =     0    SRQTIMER =     0    PSRETRY     =     0
IST2194I   PARTNER =     0    MNPS     =     0    UNAVAILABLE =     0
IST2195I   NETWORK =     0  HIGH =     0  MEDIUM =     0  LOW =     0
IST924I -------------------------------------------------------------
IST2196I COMPLETED =     0
IST2195I   NETWORK =     0  HIGH =     0  MEDIUM =     0  LOW =     0
IST924I -------------------------------------------------------------
IST2197I FAILED    =    10
IST2195I   NETWORK =     5  HIGH =     5  MEDIUM =     0  LOW =     0
IST924I -------------------------------------------------------------
IST2198I NETID          STARTED          COMPLETED         FAILED
IST2199I   CPNAME   NET  HI MED LOW  NET  HI MED LOW  NET  HI MED LOW
IST2205I ---------  ---------------  ---------------  ---------------
IST2200I NETA         0   0   0   0    0   0   0   0    5   5   0   0
IST2201I   SSCP3A     0   0   0   0    0   0   0   0    1   1   0   0
IST2201I   SSCP7A     0   0   0   0    0   0   0   0    1   1   0   0
IST2201I   SSCP99     0   0   0   0    0   0   0   0    1   1   0   0
IST2201I   SSCP7B     0   0   0   0    0   0   0   0    1   1   0   0
IST2201I   SSCP2AB    0   0   0   0    0   0   0   0    1   1   0   0
IST924I -------------------------------------------------------------
IST2206I 10 PATH SWITCH EVENTS FOR 5 CPS IN 1 NETIDS
IST314I END
```


At the end of two minutes, the five network and five high priority RTP pipes will fail to path switch successfully. At this point, a new sixty second path switch interval is started. Again, the IST1494I(Failed) messages will only be issued for the RTP pipes which had IST1494I(Started) messages issued earlier. When this interval expires, the IST2191I path switch summary message group is issued.

HPR path switch enhancements: summarization example 2 - Continued

```
IST2191I HPR PATH SWITCH SUMMARY FROM 04/27/06 AT 06:25:11
IST924I -------------------------------------------------------------
IST2192I STARTED   =     0
IST2193I   TGINOP  =     0    SRQTIMER =     0    PSRETRY     =     0
IST2194I   PARTNER =     0    MNPS     =     0    UNAVAILABLE =     0
IST2195I   NETWORK =     0  HIGH =     0  MEDIUM =     0  LOW =     0
IST924I -------------------------------------------------------------
IST2196I COMPLETED =     0
IST2195I   NETWORK =     0  HIGH =     0  MEDIUM =     0  LOW =     0
IST924I -------------------------------------------------------------
IST2197I FAILED    =    10
IST2195I   NETWORK =     0  HIGH =     0  MEDIUM =     5  LOW =     5
IST924I -------------------------------------------------------------
IST2198I NETID          STARTED          COMPLETED         FAILED
IST2199I   CPNAME   NET  HI MED LOW  NET  HI MED LOW  NET  HI MED LOW
IST2205I ---------  --------------  --------------  --------------
IST2200I NETA         0   0   0   0    0   0   0   0    0   0   5   5
IST2201I   SSCP3A     0   0   0   0    0   0   0   0    0   0   1   1
IST2201I   SSCP7A     0   0   0   0    0   0   0   0    0   0   1   1
IST2201I   SSCP99     0   0   0   0    0   0   0   0    0   0   1   1
IST2201I   SSCP7B     0   0   0   0    0   0   0   0    0   0   1   1
IST2201I   SSCP2AB    0   0   0   0    0   0   0   0    0   0   1   1
IST924I -------------------------------------------------------------
IST2206I 10 PATH SWITCH EVENTS FOR 5 CPS IN 1 NETIDS
IST314I END
```


The same occurs at the four minute mark with the final ten RTP pipes fail to path switch successfully.

HPR Path switch improvements

- How well does it perform?

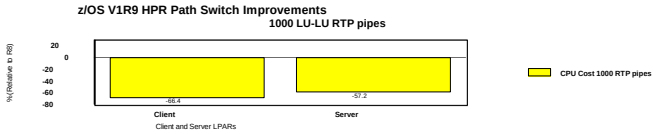

  - V1R8 versus V1R9 path switch performance comparison
    - ✓ Network Node – Network Node configuration over one-hop EE (1Gb OSA)
      - Two EE TGs defined between the network nodes
    - ✓ 1000 LU-LU RTP pipes path switched from one TG to the other
    - ✓ V1R8
      - Message IST1494I added to the VTAM message-flooding prevention table
    - ✓ V1R9
      - Message IST1494I removed from the VTAM message-flooding prevention table
      - HPRPSMSG = 10
    - ✓ 57% - 66% processor savings for this scenario
- If you use the VTAM message-flooding prevention table
  - IST1494I should be removed before enabling the HPR path switch message reduction and summarization function
  - If not, this will affect the number of path switch messages you receive on the system console

Performance measurements were taken for this line item to verify the changes made in V1R9 are beneficial. The configuration used consisted of two network nodes connected to one another using two Enterprise Extender transmission groups. One thousand LU-LU RTP pipes were established over the first transmission group. For V1R8, message IST1494I was added to the VTAM message-flooding prevention table. For V1R9, message IST1494I was removed from the VTAM message-flooding prevention table. Also, the new HPRPSMSG start option was coded to 10. During this scenario, the first EE TG was inactivated, causing the one thousand RTP pipes to path switch to the other EE transmission group. This test measures the processor costs when performed in a V1R8 versus V1R9 environment. In the end, processor savings ranged from 57 to 66 percent in the V1R9 scenario. These processor savings will vary depending on the configuration.

If the HPR path switch message reduction and summarization function is enabled, message IST1494I should not be specified in the VTAM message-flooding prevention table. If it is not removed and this new function is enabled, you will not receive the expected number of path switch messages on the system console.

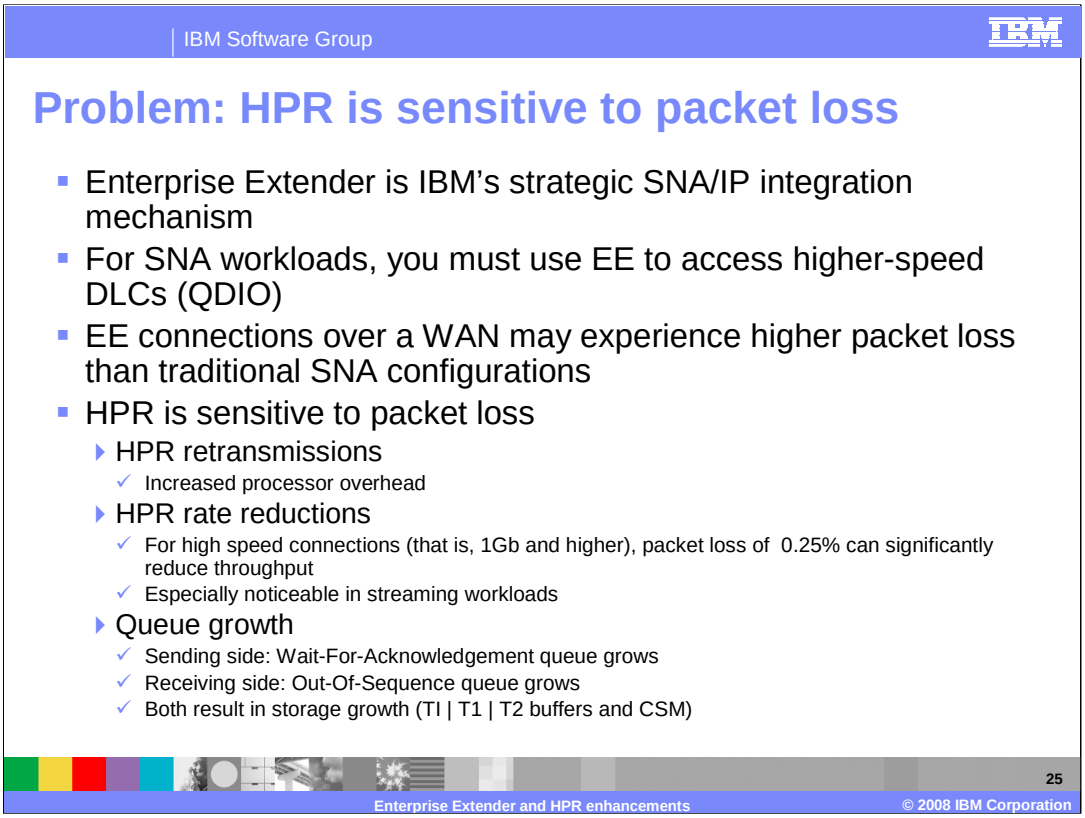

Enterprise Extender is IBM's strategic SNA over IP integration mechanism. Depending on the reliability of the IP backbone, some Enterprise Extender connections may experience higher packet loss than traditional SNA configurations.

HPR is sensitive to packet loss. If packet loss occurs this may cause HPR retransmissions, rate reductions and queue growth. As a result, you may see increased processor overhead, higher storage utilization and significantly reduced throughput for an RTP pipe suffering from packet loss.

Solution: EE improved packet loss tolerance

- HPR needs to be more reactive
  - Receiving side
    - ✓ Report “gaps” sooner
    - ✓ More aggressive REFIFO timer formula
  - Sending side
    - ✓ When partner reports gaps
      - – Allow burst timer to run as small as 1ms
      - – Paces data across EE connection more evenly
- These changes require a more granular HPR clock
  - Clock generally runs at 25ms intervals
  - Clock can now run at 1ms intervals, when necessary
  - New start option – HPRCLKRT
    - ✓ Controls the rate at which the HPR clock runs
    - ✓ Only applies when RTP pipes run over EE with a defined capacity of 1 Gigabit or higher
    - ✓ Standard – HPR clock only runs in standard mode (25ms mode)
    - ✓ Adaptive – allows HPR clock to change modes (standard or high) based on network conditions
    - ✓ Cannot be modified
- Optimized HPR “Liveness” timer processing
  - Beneficial when DISCNT=NO specified for dynamic RTP PUs
  - Allows the HPR clock to stop when all RTP pipes are idle
  - processor savings as a result

HPR has been changed to be more tolerant of packet loss. To begin with, a more aggressive REFIFO timer formula has been implemented to allow the receiver to report gaps sooner to the partner. The REFIFO timer is used by RTP pipes to delay reporting missing packets to the partner to avoid unnecessary transmissions. If a missing packet is detected, the RTP pipe will set the REFIFO timer. When the timer expires and the packet is still missing it is reported to the partner so it can be sent again.

The BURST timer is used by an RTP pipe to pace the data across the connection at specific intervals. Depending on the speed of the RTP connection, the amount of data which can be sent in a burst interval varies. Generally, the BURST timer runs at 25ms intervals. Now, the sending side has been changed to allow the BURST timer to run as small as one millisecond. This will allow the RTP pipe to better pace the data across the connection. When necessary, the HPR clock must now be allowed to run at a one millisecond rate to support these new timer changes.

To support these changes, a new start option, HPRCLKRT, has been introduced to control this function. Specifying STANDARD requires the HPR clock always run in standard mode or twenty-five millisecond mode. When specifying ADAPTIVE, the HPR clock is allowed to change modes based on network conditions. This start option is only valid when RTP pipes are directly connected to an Enterprise Extender link with a defined capacity of one gigabit or higher.

One last change was made to optimize the HPR liveness timer processing. The HPR liveness timer is used to send keep-alive signals to the partner to see if they are still active. If you set DISCNT to no for dynamic RTP physical units, this means the RTP pipe will not inactivate when the last session ends. Instead the RTP pipe stays idle until a new session uses it. While it is active, liveness timer processing may still occur. In an environment where there is little to no activity at certain times of the day, there is still HPR clock overhead necessary to perform the liveness timer processing. This liveness timer processing has been optimized for this type of environment to allow the HPR clock to stop when all RTP pipes are idle. If all RTP pipes are idle, this change may save processor overhead as a result. The liveness optimization change is in the base code and is not controlled by any VTAM start option.

EE improved packet loss tolerance Display NET,EE changes

```
D NET,EE

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 = XCAEE2
IST2004I LIVTIME = (10,0) SRQTIME = 15 SRQRETRY = 3
IST2005I IPRESOLV = 0
IST2231I CURRENT HPR CLOCK RATE = STANDARD
IST2232I HPR CLOCK RATE LAST SET TO HIGH ON 11/14/06 AT 22:58:41
IST2233I HPR CLOCK RATE LAST EXITED HIGH ON 11/14/06 AT 22:58:45
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 -------------------------------------------------------------
IST2017I TOTAL RTP PIPES =          4     LU-LU SESSIONS =          3
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
```


This display has been enhanced to display new message IST2231I to indicate the current mode of the HPR clock rate. Message IST2232I displays the last time the HPR clock entered high mode. Message IST2233I displays the last time the HPR clock exited high mode. In this case, since message IST2232I and IST2233I are present you know that the HPRCLKRT start option has been set to adaptive mode, but the current mode of the HPR clock is standard mode.

EE improved packet loss tolerance HPRDIAG changes (1 of 4)

```
D NET,ID=CNR00005,HPRDIAG=YES

IST097I DISPLAY ACCEPTED
IST075I NAME = CNR00005, TYPE = PU_T2.1 059
IST1392I DISCNTIM = 00010 DEFINED AT PU FOR DISCONNECT
IST486I STATUS= ACTIV--LX-, DESIRED STATE= ACTIV
IST1043I CP NAME = NS24 - CP NETID = NETA - DYNAMIC LU = YES
IST1589I XNETALS = YES
IST2238I DISCNT = NO - FINAL USE = *NA*
IST231I RTP MAJOR NODE = ISTRTPMN
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1500I STATE TRACE = OFF
IST2178I RPNCB ADDRESS 2241E800
IST1965I APPNCOS = #BATCH - PRIORITY = LOW
IST1476I TCID X'35056090000100E3' - REMOTE TCID X'350564F7000100E6'
IST1481I DESTINATION CP NETA.NS24 - NCE X'D000000000000000'
IST1587I ORIGIN NCE X'D000000000000000'
IST1967I ACTIVATED AS PASSIVE ON 11/14/06 AT 22:57:22
IST2237I CNR00005 CURRENTLY REPRESENTS A LIMITED RESOURCE
IST1479I RTP CONNECTION STATE = CONNECTED - MNPS = NO
IST1959I DATA FLOW STATE = NORMAL
IST1855I NUMBER OF SESSIONS USING RTP = 20
IST1480I RTP END TO END ROUTE - RSCV PATH
IST1460I TGN  CPNAME             TG TYPE      HPR
IST1461I  21  NETA.NS24          APPN         RTP
IST875I ALSNAME TOWARDS RTP = SWIP25
IST1738I ANR LABEL               TP           ER NUMBER
IST1739I 800100D701000000        *NA*         *NA*
```


Message IST2237I is issued to identify this RTP pipe as a limited resource. For an HPR PU, this really means the underlying DLC has the DISCNT parameter specified as DELAY or YES.

EE improved packet loss tolerance HPRDIAG changes (2 of 4)

```
IST924I -------------------------------------------------------------
IST1968I ARB INFORMATION:
IST1844I ARB MODE = GREEN
IST1697I RTP PACING ALGORITHM = ARB RESPONSIVE MODE
IST1477I ALLOWED DATA FLOW RATE = 1505 MBITS/SEC
IST1516I INITIAL DATA FLOW RATE = 47 MBITS/SEC
IST1841I ACTUAL DATA FLOW RATE = 160 MBITS/SEC
IST1969I MAXIMUM ACTUAL DATA FLOW RATE = 907 MBITS/SEC
IST1862I ARB MAXIMUM SEND RATE = 944 MBITS/SEC
IST1846I CURRENT RECEIVER THRESHOLD = 36998 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 = 0 SECONDS
IST1851I SMOOTHED ROUND TRIP TIME = 10 MILLISECONDS
IST1972I SHORT REQUEST TIMER = 250 MILLISECONDS
IST2229I REFIFO TIMER = 7 MILLISECONDS
IST924I -------------------------------------------------------------
```


New message IST2229I displays the current value of the HPR REFIFO timer. This is the amount of time this end of the HPR pipe waits before reporting missing packets (gaps) to the partner.

Notice that message IST1862I indicates an ARB maximum send rate of 944 MBITS/SEC. When CAPACITY is coded at 1000M or 1G on the underlying PU definition, the connection actually displays 944 MBITS/SEC.

EE improved packet loss tolerance HPRDIAG changes (3 of 4)

```
IST924I ---------------------------------------------------------
IST1973I OUTBOUND TRANSMISSION INFORMATION:
IST1974I NUMBER OF NLPS SENT = 12150570 ( 12M )
IST1975I TOTAL BYTES SENT = 16154988112 ( 16G )
IST1849I LARGEST NLP SENT = 1377 BYTES
IST1980I SEQUENCE NUMBER = 2878146315 (X'AB8D070B')
IST1842I NUMBER OF NLPS RETRANSMITTED = 9
IST2236I LAST NLP RETRANSMITTED ON 11/14/06 AT 23:00:43
IST1976I BYTES RETRANSMITTED = 11034 ( 11K )
IST1478I NUMBER OF UNACKNOWLEDGED BUFFERS = 19
IST1958I NUMBER OF ORPHANED BUFFERS = 0
IST1843I NUMBER OF NLPS ON WAITING-TO-SEND QUEUE = 0
IST1847I NUMBER OF NLPS ON WAITING-FOR-ACKNOWLEDGEMENT QUEUE = 19
IST1977I MAXIMUM NUMBER OF NLPS ON WAITING-FOR-ACK QUEUE = 639
IST1978I WAITING-FOR-ACK QUEUE MAX REACHED ON 11/14/06 AT 23:02:03
IST2085I NUMBER OF NLPS ON OUTBOUND WORK QUEUE = 0
IST2086I MAXIMUM NUMBER OF NLPS ON OUTBOUND WORK QUEUE = 153
IST2087I OUTBOUND WORK QUEUE MAX REACHED ON 11/14/06 AT 23:02:03
IST1511I MAXIMUM NETWORK LAYER PACKET SIZE = 1469 BYTES
IST924I ---------------------------------------------------------
IST1979I INBOUND TRANSMISSION INFORMATION:
IST2059I NUMBER OF NLPS RECEIVED = 12038943 ( 12M )
IST1981I TOTAL BYTES RECEIVED = 16372334887 ( 16G )
IST1850I LARGEST NLP RECEIVED = 1377 BYTES
IST1980I SEQUENCE NUMBER = 3101845908 (X'B8E26994')
IST1853I NUMBER OF NLPS ON OUT-OF-SEQUENCE QUEUE = 0
IST2230I MAXIMUM NUMBER OF NLPS ON OUT-OF-SEQUENCE QUEUE = 163
IST1854I NUMBER OF NLPS ON INBOUND SEGMENTS QUEUE = 6
IST1982I NUMBER OF NLPS ON INBOUND WORK QUEUE = 0
IST1983I MAXIMUM NUMBER OF NLPS ON INBOUND WORK QUEUE = 724
IST924I ---------------------------------------------------------
```


New message IST2236I displays the date and time of when the last NLP was retransmitted. Message IST2230I displays the high-water-mark for the RTP out-of-sequence queue. In this example 163 network layer packets (NLPs) have been on the out-of-sequence queue since the HPR pipe was started.

EE improved packet loss tolerance HPRDIAG changes (4 of 4)

```
IST924I -------------------------------------------------------------
IST1984I PATH SWITCH INFORMATION:
IST1985I PATH SWITCHES INITIATED FROM REMOTE RTP = 0
IST1986I PATH SWITCHES INITIATED FROM LOCAL RTP = 0
IST1987I PATH SWITCHES DUE TO LOCAL FAILURE = 0
IST1988I PATH SWITCHES DUE TO LOCAL PSRETRY = 0
IST924I -------------------------------------------------------------
IST1857I BACKPRESSURE REASON COUNTS:
IST1858I PATHSWITCH  SEND QUEUE MAX  STORAGE FAILURE  STALLED PIPE
IST2205I ----------  --------------  ---------------  ------------
IST1859I      0             3                0              0
IST2211I ACK QUEUE MAX
IST2205I -------------
IST2212I      0
IST2213I LAST BACKPRESSURE APPLIED ON 11/14/06 AT 23:00:24
IST2215I BACKPRESSURE REASON: SEND QUEUE MAXIMUM REACHED
IST924I -------------------------------------------------------------
IST314I END
```


The HPR backpressure section now displays the date and time of when back pressure was applied to this RTP pipe. If also displays the latest reason for the backpressure. This section also displays a new backpressure count titled “ACK QUEUE MAX”. This HPR backpressure is applied when the RTP waiting for acknowledgement queue reaches a depth of ten thousand elements. The backpressure is relieved when the queue depth returns to five thousand or less elements.

Message IST12212I displays the number of times that this Rapid Transfer Protocol (RTP) went into backpressure (holding up outbound data transmission) since the HPR PU was activated.

Feedback

Your feedback is valuable

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

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

Click to send e-mail feedback:

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

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

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

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 VTAM z/OS

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 or changes in the products or programs 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 (for example, 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 multiprogramming 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.

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.