

z/OS® V1R10 Communications Server

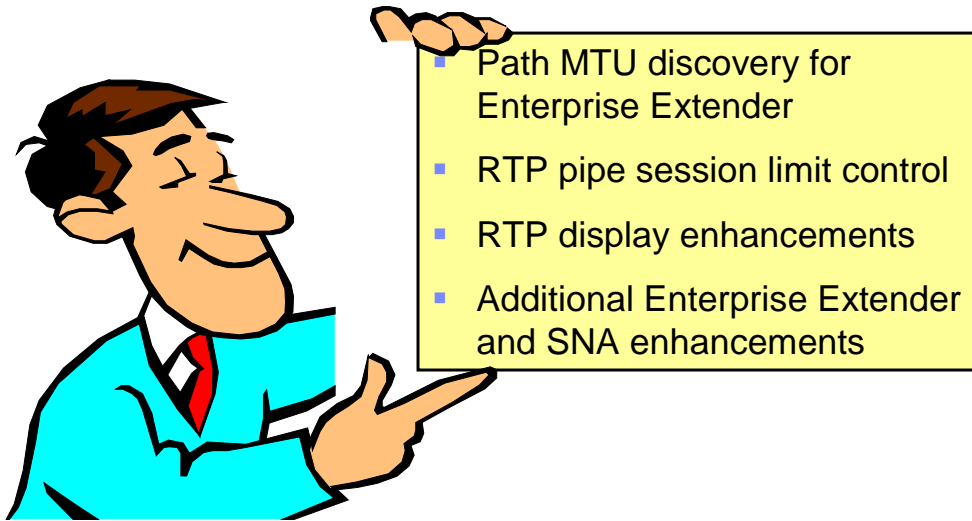
Overview : EE and SNA

@business on demand software

© 2008 IBM Corporation


This presentation is an overview of EE and SNA enhancements in the Communications Server for z/OS V1R10.

Agenda



This presentation covers several Enterprise Extender and SNA enhancements for CS z/OS V1R10. These include improvements to Path MTU discovery for Enterprise Extender, RTP pipe session limit control and RTP display information..

Path MTU discovery for Enterprise Extender



- How does EE know what packet size to use?
 - Default is to use the configured interface MTU size
 - A routing table entry with an overriding MTU size for the destination

But what if you do not know what the minimum packet size in the network is, or it changes dynamically because of IP network topology changes?

- In V1R10, EE added support for detecting excessive retransmissions
- Also, if TCP/IP is configured with path MTU discovery, information learned by path MTU discovery is available to EE

Before V1R9 the interface MTU size was used during Enterprise Extender (EE) connection setup and was not changed during the life of the connection, even if the local interface MTU size changed. This slide shows a size of 8992 being selected. This will result in packet loss.

In V1R9, support was added to inform EE if the local MTU size changes, for example due to dynamic routing entries to the destination. However, the Interface MTU size only provides the first hop MTU size and has no knowledge of MTU size beyond first-hop router.

There are several things to remember when using a routing table entry to control interface MTU size. Before V1R9, this had to be a static route entry (can be configured to 1500 and would work). However, this would be used for all traffic to that destination, not just EE.

V1R9 added support for Policy-Based Routing and an EE-specific policy route with EE-specific MTU size could be configured (could be configured to 1500 and would work without affecting non-EE users of the interface).

V1R10 adds support for multiple VLANs per OSA port. One such VLAN could be dedicated for EE traffic (using a PBR rule) with a configured MTU size of 1500.


In V1R10 EE can use the existing path MTU discovery TCP/IP function, and can now dynamically adapt to network MTU size.

Excessive retransmissions are defined as more than six. If retransmissions are excessive, EE will now reduce packet size to 768.

RTP pipe session limit control



Example: 150,000 sessions over a single RTP pipe



Example: 150,000 sessions over five parallel RTP pipes (30,000 over each pipe)

- Large numbers of SNA sessions between EE nodes are becoming more common
- A very high number of sessions on a single RTP pipe can have a negative impact on performance
- z/OS V1R10 implements an optional VTAM® start option to specify a threshold for number of sessions on an RTP pipe

An example of a case where there can be a large number of SNA sessions is a front-end z/OS TN3270 server LPAR using EE to back-end z/OS application LPARs.

A very high number of sessions over one pipe can cause negative impacts to queue sizes, and send/receive processing. Since each RTP pipe is a separate process in VTAM, VTAM cannot exploit parallel processors when processing work for this single RTP.

A new configuration parameter, HPRSESLM, is provided. When the session count exceeds HPRSESLM, a new parallel RTP pipe will be set up to carry the next session.

Display RTP enhancements

- The Display RTPS command is enhanced with options to easier identify poorly performing RTP pipes

```
D NET,RTPS,REXMIT=5
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
```

- A new message is added to the HPRDIAG=YES output section

```
ISTHM02I NLP RETRANSMIT RATE          = 0.0099%
```

VTAM now allows the operand REXMIT with the retransmission rate percentage value on D NET,RTPS command to display RTPs. The D,NET,RTPS,REXMIT command output will display all RTPs that meet or exceed the retransmission rate percentage value specified on REXMIT operand.

The DISPLAY NET,RTPS command is enhanced to clear all diagnostic counters with new CLEAR operand.

The DISPLAY NET,ID command with HPRDIAG=YES is enhanced to clear all diagnostic counters with new CLEAR operand.

Additional EE/SNA enhancements

➤ Start option for CSDUMP triggers

➤ Stalled HPR pipe termination

```
IST1955I STALL DETECTED FOR RTP puname TO cpname
IST1956I STALL CONTINUES FOR RTP puname TO cpname
IST1957I STALL ALLEVIATED FOR RTP puname TO cpname
```

➤ TGN Parameter for EE model PUs

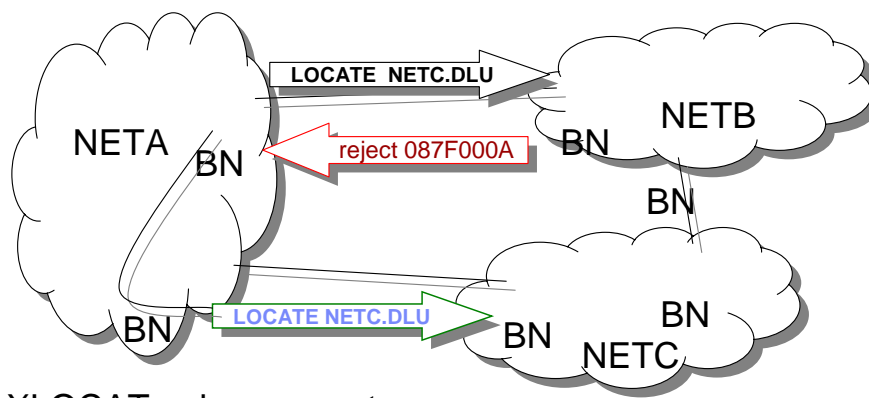
A CSDUMP start option, modeled partially on the current MODIFY CSDUMP command, will help customers provide documentation for problem determination. It will allow the same message and sense code triggers currently set by the MODIFY CSDUMP command to be set at VTAM startup time. Unlike the MODIFY CSDUMP command, it will not be used to take an immediate dump.

Improvements were made to handle stalled HPR pipes. z/OS V1R6 implemented enhanced messaging to indicate when pipes were hung. IST956I will be repeated every 30 seconds if the pipe remains stalled. z/OS V1R10 added new support to automatically terminate a pipe that has been in stalled state for more than a certain period of time. This is configurable using a new VTAM start option, `HPRSTALL=time_limit`.

The TG number for a predefined EE connection can be specified using the TGN parameter on the switched PU definition. However, for dynamically defined EE TGs, there is currently no way to specify the TG number to be used for the connection. The architectural default of 21 will be used unless the partner negotiates a different value. z/OS V1R10 allows the specification of a suggested TG number by coding the TGN parameter on a DYNTYPE=EE PU entry in the VTAM model major node. This simplifies network management and problem determination.

Additional EE/SNA enhancements

➤ Extended Border Node searching enhancement



➤ MAXLOCAT enhancement

7

Overview : EE and SNA

© 2008 IBM Corporation

This function extends the ALIASRCH function that was added in z/OS V1R8. ALIASRCH as it was implemented in z/OS V1R8 can be used to reject alias searches being received by your network, such as a search without a NETID. In z/OS V1R10, a new configuration option enables the rejection of an incoming cross-network search that includes a target NETID that is not one from the configured list of allowed target NETIDs. This is implemented using the AUTHNETS operand on the ADJCP minor node definition for an adjacent non-native CP. This support further reduces the need for DSME exit coding.

This diagram shows a LOCATE from NETA to NETB for a NETC resource. The entry border node in NETB has AUTHNETS specified on the adjacent control point minor node for the sending BN, NETC is not in the list. This LOCATE will be rejected with 087F000A sense code. This allows NETA to continue searching for the DLU.

A change has also been made in the way MAXLOCAT and IOPURGE parameters affect directory services requests. In V1R9, if an outbound CP-CP session is hung, but the inbound CP-CP session is not, new directory services requests will not be sent. However, if any requests are received, the replies will be queued to the hung CP. Storage might eventually fill up and lead to outages. In z/OS V1R10, once MAXLOCAT initiates, Directory Services will start monitoring the amount of replies being queued towards that same CP. When the number of Requests plus Replies exceeds MAXLOCAT*2 for IOPURGE amount of time (at least 180 seconds) the CP-CP session will be unbound with a retryable sense code to clean up storage.

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_wneesna.ppt

This module is also available in PDF format at: [../wneesna.pdf](http://wneesna.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:

VTAM z/OS

A current list of other IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

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.

© Copyright International Business Machines Corporation 2008. All rights reserved.

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