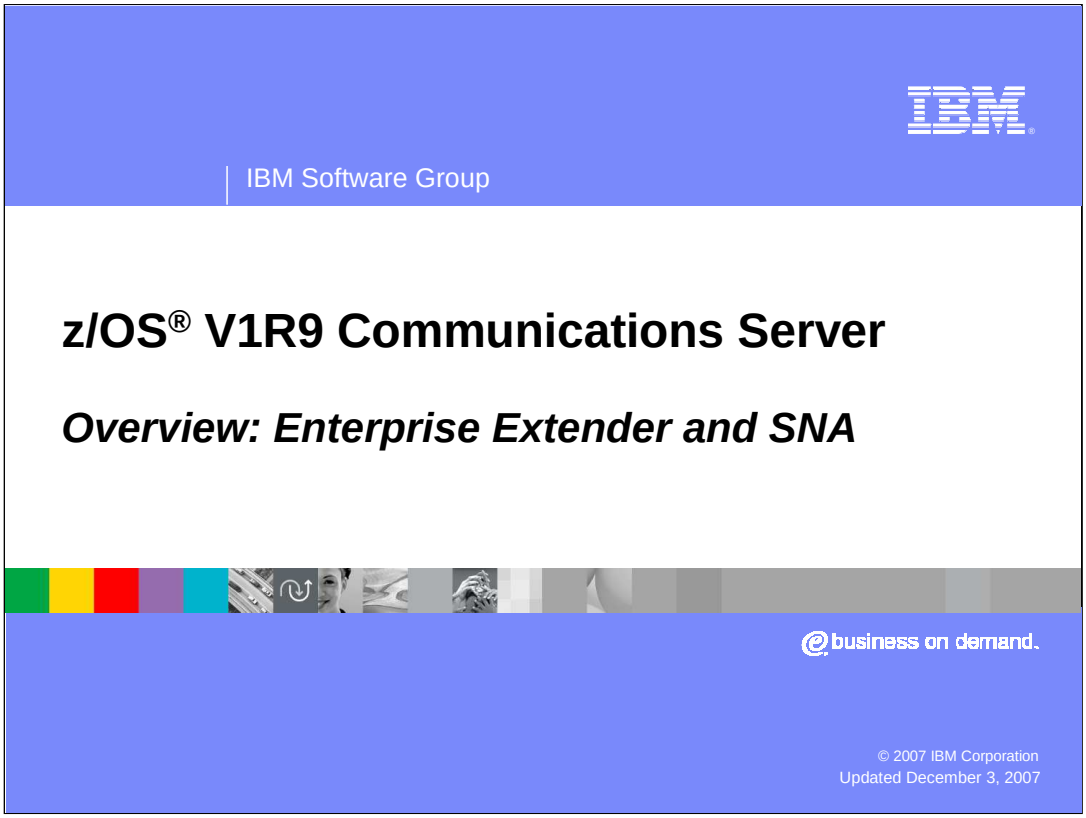

This presentation discusses the Enterprise Extender and SNA enhancements in z/OS V1R9 Communications Server.

Local MTU discovery for Enterprise Extender

- In previous releases, an Enterprise Extender connection obtained the maximum MTU size permitted for packets being transmitted to a remote Enterprise Extender endpoint when the connection is initially established.
  - MTU size obtained from TCP/IP when VTAM® gets route information
  - MTU size unchanged for the life of the connection
    - Even if route changes
  - If Enterprise Extender initializes early during start of a TCP/IP stack, the available route from which Enterprise Extender chooses the MTU size, may not be the optimal route
- Enterprise Extender monitors local routing information and dynamically changes the MTU size for that Enterprise Extender connection
  - Note: RTP pipes routed over an Enterprise Extender connection will only learn of changes in the MTU size when the RTP endpoint and the Enterprise Extender connection reside in the same node. Enterprise Extender is not aware of MTU sizes within the network.

In prior releases an Enterprise Extender connection would obtain the maximum MTU size permitted for packets being transmitted to a remote Enterprise Extender endpoint at the time that the connection is initially established. This MTU size is obtained from TCP/IP when VTAM obtains route information for a new Enterprise Extender connection and never changes for the life of the connection. During the life of this Enterprise Extender connection the TCP/IP stack could change the actual route that is associated with the Enterprise Extender connection, but Enterprise Extender never learns about that and never modifies the initially chosen MTU size. If Enterprise Extender initializes early during start of a TCP/IP stack, the available route from which Enterprise Extender chooses the MTU size, may not be the optimal route. OMPROUTE may not have started yet and the only routes available might be some default route with an MTU size that doesn't match the MTU size of the actual network the Enterprise Extender traffic later will be sent over. Currently, the only way for Enterprise Extender to learn about a better route and MTU size is to stop and restart the Enterprise Extender connection.

It is important that VTAM understand the maximum MTU size for the Enterprise Extender connection to avoid fragmentation by the TCP/IP stack which will greatly increase the path length for transmitted Enterprise Extender packets in the TCP/IP stack. Therefore today if a new interface is being utilized by TCP/IP for an existing Enterprise Extender connection (previous route obtained could have been deleted or became inactive) and the new MTU size is smaller than was previously reported to VTAM (at Enterprise Extender 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 Enterprise Extender 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 Enterprise Extender connection. Currently if TCP/IP obtains a route handle for an Enterprise Extender 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 Enterprise Extender connection. Therefore there have been a number of users that start their VTAM and TCP/IP connection where an Enterprise Extender connection is initiated by VTAM prior to TCP/IP learning about all of the potential routes from OMPROUTE and therefore the Enterprise Extender connection ends up utilizing the default route (which in many cases is not the optimal route).

Enterprise Extender will now 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 utilize the current interface associated with an Enterprise Extender connection. Furthermore TCP/IP will allow for the determination of a more optimal route for an existing Enterprise Extender connection when new routes are learned.

HPR path switch enhancements

- While HPR path switch works fine in simple environments with a small number of pipes, there has been concern voiced over the impact of a large number of pipes entering path switch state at approximately the same time, such as what occurs during major network failure scenarios.
- The major focus areas that were addressed are:
  - Optimizing the handling of HPR's path switch timers
  - Reducing amount of storage allocation and deallocation occurring during path switch scenarios
  - Implementing Path Switch Message Summarization to reduce the number of path switch messages flooding the console

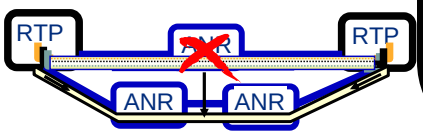


```
IST2191I HPR PATH SWITCH SUMMARY FROM 04/05/06 AT 09:45:14
IST924I -------------------------------------------------------------
IST2192I STARTED   =    12
IST2193I   TGINOP  =    12    SRQTIMER =     0    PSRETRY     =     0
IST2194I   PARTNER =     0    MNPS     =     0    UNAVAILABLE =     0
IST2195I   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
```


While HPR path switch works fine in simple environments with a small number of pipes, there has been concern voiced over the impact of a large number of pipes entering path switch state at approximately the same time, such as what occurs during major network failure scenarios.

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

The major focus areas that were addressed are:

- Optimizing the handling of HPR's path switch timers.
- Reducing the amount of storage allocation and storage deallocation occurring during path switch scenarios.
- Implementing path switch message summarization to reduce the number of path switch messages flooding the console leading to WTO buffer shortages. VTAM will output a path switch summarization display to document all the associated path switch events which occurred during a given time interval.

HPR Path Switch Summarization reduces the number of path switch message groups VTAM issues across a 60-second interval. At the end of the 60-second interval, a summarization report is issued on the total path switch activity during the interval. This helps to avoid WTO buffer shortages.

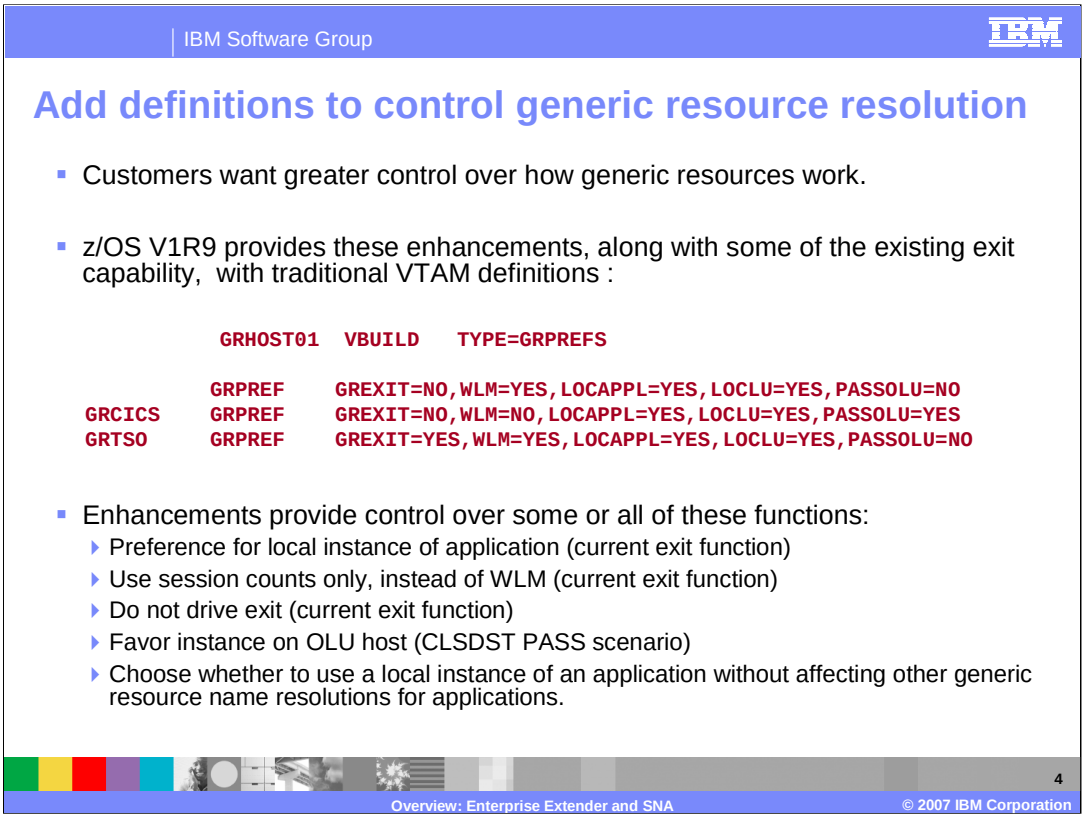

Generic resource resolution preferences are used to control the distribution of sessions to generic resources. These preferences previously could only be set globally in the Generic Resource Exit.

Recently we have received a number of requirements for greater control over how generic resources work. Since the only real control available today is through the Generic Resources exit routine (ISTEXCGR) , those requirements invariably request some additional capability for the exit. However, many of our users have provided feedback that they do NOT like to write and maintain exit code in assembler, even simple exit code that just sets/clears a bit to enable/disable a function. Therefore, it seems likely that any functional capability added to the exit will be used by very few users.

In addition since the GR exit flag settings only affect the next GR resolution, they can not be used to differentiate GR resolution behavior between different GR applications (for example GRCICS vs GRTSO). Users will still have the ability to code a GR exit to perform complex GR resolution using any criteria they choose, however the generic resource preferences will now only be settable using a generic resource preference table.

This function allows you to use VTAM definitions to customize generic resource resolution preferences for individual generic resources. In addition a new generic resource resolution preference is also being introduced that will allow a generic resource resolution during 3rd party initiated (CLSDST-PASS) sessions to favor a generic resource on the origin host of a session.

A new VBUILD type GRPREFS has been created to identify the generic resource preferences table. A new definition statement GRPREF has been defined within the GRPREFS table to define GR resolution preferences. A GRPREF statement can be defined for each GR name. A nameless GRPREF statement can be used to define default GR preferences.

You can activate a GRPREFS table using the VARY NET,ACT,ID= command where the name of the table is the VTAMLST member name that contains the generic resource preferences definitions.

You can also start the GRPREFS table using the VTAM Config List using the same name. Since a table cannot be inactivated, to effectively inactivate a table activate a generic resource preferences table with a nameless entry and no operands.

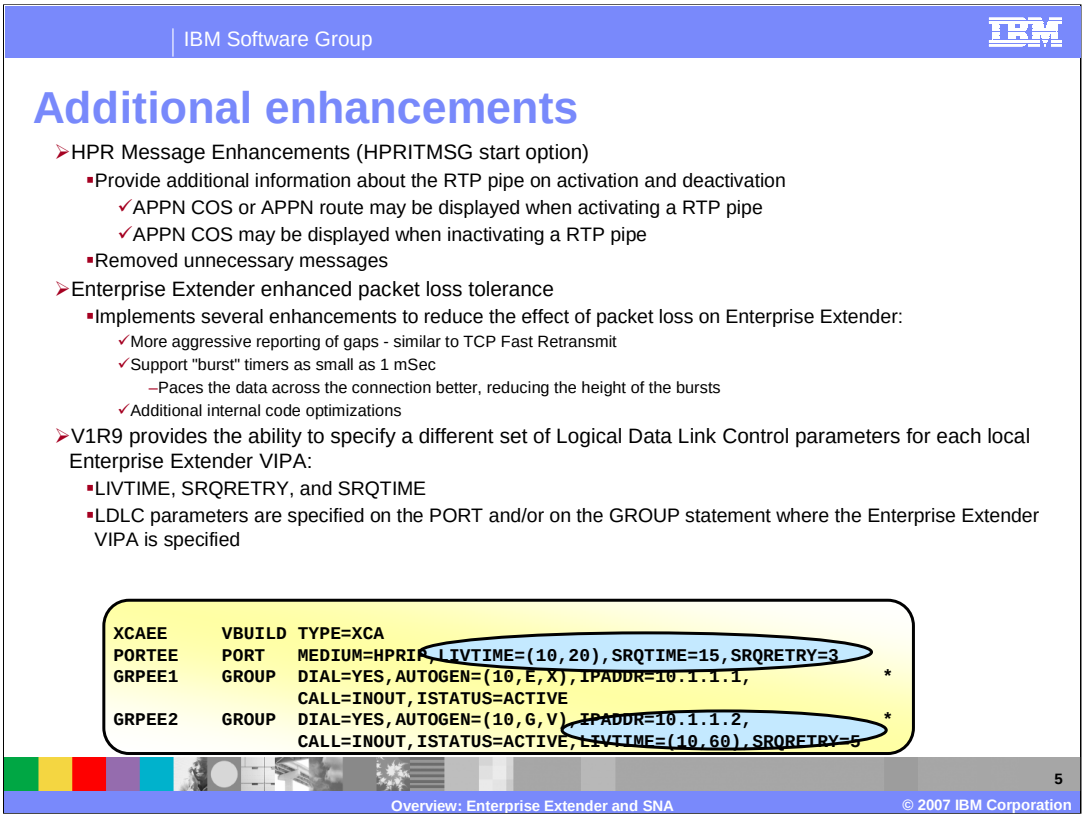

During RTP pipe activation, VTAM does not identify the APPN COS or APPN route associated with the RTP pipe. Without this information, you cannot identify the priority of the RTP pipe or verify the correct APPN route has been selected. Also VTAM does not identify the associated APPN COS during RTP pipe inactivation. Therefore it is difficult to identify which priority RTP pipe is cleaning up. Many of the dynamic RTP PU cleanup messages issued during RTP inactivation are unnecessary. If large number of RTP pipes are cleaning up, this may lead to hundreds or thousands of unnecessary messages being issued to the system console.

In z/OS V1R9 Communications Server, VTAM may now display the APPN COS or APPN route associated with the RTP pipe when activating the 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. VTAM may now display the associated APPN COS during RTP inactivation. 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. This behavior is only available when the new HPR message enhancement start option, HPRITMSG, is set to a value of ENHANCED.

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 CPU overhead, higher storage utilization and significantly reduced throughput for an RTP pipe suffering from packet loss.

In z/OS V1R9 Communications Server, HPR has been changed to be more tolerant of packet loss. 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 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.
The VTAM Enterprise Extender Logical Data Link Control (LDLC) layer monitors the Enterprise Extender connection by testing for the remote partner availability. During periods of inactivity on the Enterprise Extender connection, when the liveness timer expires, LDLC polls the partner with an LDLC TEST request. This verifies that the Enterprise Extender partner is still available. The LDLC inactivity trigger is controlled by Enterprise Extender timer parameters LIVTIME, SRQTIME and SRQRETRY on the PORT statement. The LDLC timer operands apply to the whole Enterprise Extender network. They are not unique to each Enterprise Extender connection. Network conditions may vary between internal communications (LAN) or external communications (WAN). The LDLC timer operands may be optimal for one Enterprise Extender connection, but it may be way off for another Enterprise Extender connection. In z/OS V1R9 Communications Server, VTAM allows LDLC liveness and short request timer values to be specified for each local static VIPA address defined for Enterprise Extender. This is accomplished by allowing the definition of Enterprise Extender LDLC liveness and short request timer operands on the GROUP statements in the XCA major node. VTAM allows the dynamic update of the LDLC timer parameters LIVTIME, SRQTIME and SRQRETRY on GROUP macro with V NET,ACT,UPDATE command.

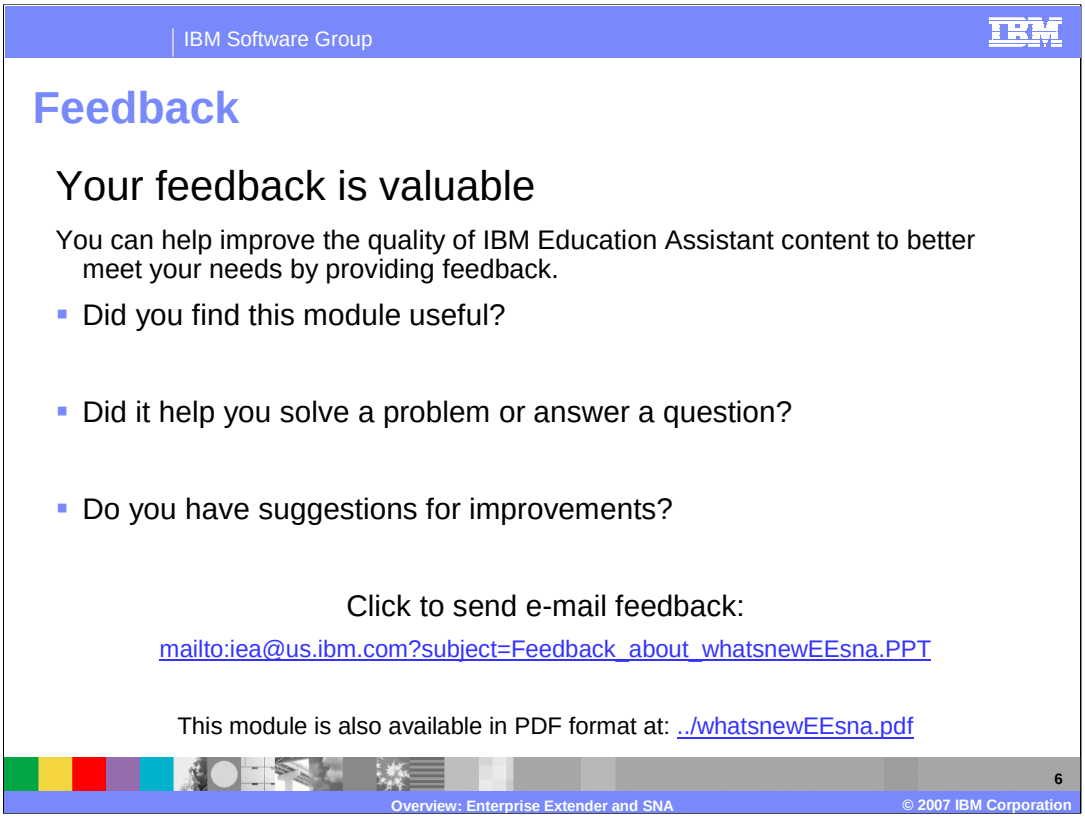

You can help improve the quality of IBM Education Assistant content by providing feedback.

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.