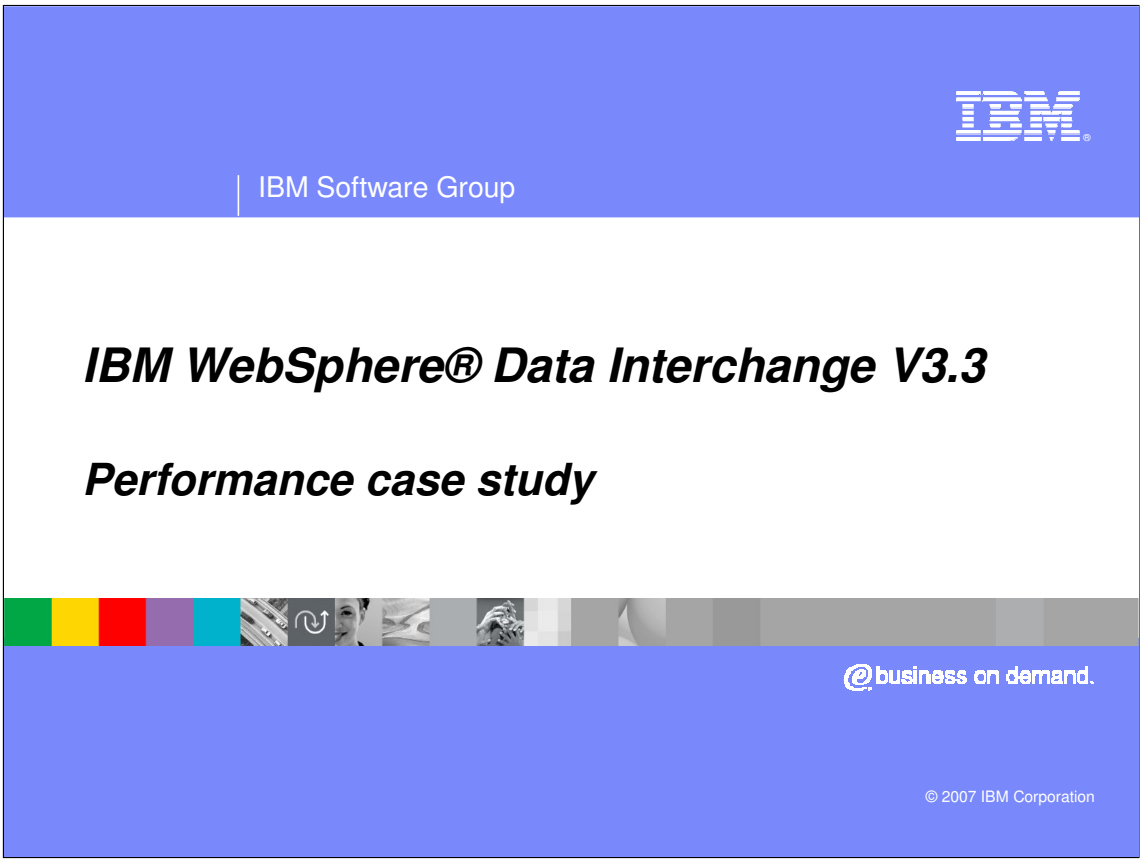

This presentation will present a case study that focuses on performance within WebSphere Data Interchange.

Topics

- Pageable AMM results
- Throughput Case Study
- Send / Receive Case Study

Topics to be discussed are
Pageable AMM results,
A Throughput Case Study, and
A Send / Receive Map Case Study.

Without PAMM

For 25+ MB input file :
real 5m42.59s
user 5m34.30s
sys 0m4.89s

For 75+ MB input file :
real 18m45.12s
user 18m29.38s
sys 0m14.69s

For 100+ MB input file:
real 9h29m51.52s
user 9h28m41.99s
sys 0m44.25s

Base line numbers for DT translations without the Pageable AMM feature.

Pageable AMM on z/OS

10K Claim input file:
Before PAMM enhancement:
System area used: 364K 10M
Virt storage used: 712K 354M
Step completion code: 0000
Total CPU time used: 00:04:20.91

With PAMM enhancement:
System area used: 360K 10M
Virt storage used: 728K 122M
Step completion code: 0000
Total CPU time used: 00:04:18.02

Note: Virt storage comparison shows a 230MB improvement.

60K Claim input file:
Before PAMM enhancement:
System area used: 364K 10M
Virt storage used: 8820K 1517M
Step completion code: 0012
Total CPU time used: 04:24:01.75

With PAMM enhancement:
System area used: 360K 10M
Virt storage used: 728K 515M
Step completion code: 0000
Total CPU time used: 00:29:25.59
SRB CPU time used: 00:00:03.96

Note: Did not complete successfully before PAMM enhancement -- ran out of virtual storage. Nonetheless, virtual storage comparison shows at least 1GB improvement.

These are some statistics for two different size executions with and without PAMM. The sample data is for Health care claim messages.

Pageable AMM and PAGETHRESHOLD effect

Pageable AMM feature with PAGETHRESHOLD(500) in the PERFORM command and the 60MB test case without the PAMM feature. The following are the CPU time when both test cases were executed on z/OS

With PAMM feature
TOTAL CPU TIME= 20.53 TOTAL ELAPSED TIME=113.73
Without PAMM feature
TOTAL CPU TIME= 40.46 TOTAL ELAPSED TIME=117.99

Pageable AMM feature with PAGETHRESHOLD(500) in the PERFORM command and the 60MB test case without the PAMM feature. The slide shows the CPU time when both test cases were executed on z/OS.

Pageable AMM Results

Memory high water mark

With change:

System area used:	352K	10M
Virtual storage used:	672K	106M

Before change:

System area used:	352K	10M
Virtual storage used:	672K	892M

Elapsed and CPU time

With change:

Time elapsed: 00:27:52.57 Total CPU time used: 00:10:46.84

Before change:

Time elapsed: 00:22:06.51 Total CPU time used: 00:13:29.94

Results from another Pageable AMM test.
The memory high water mark went down dramatically.
Elapsed time increased, but this could be due to system load.

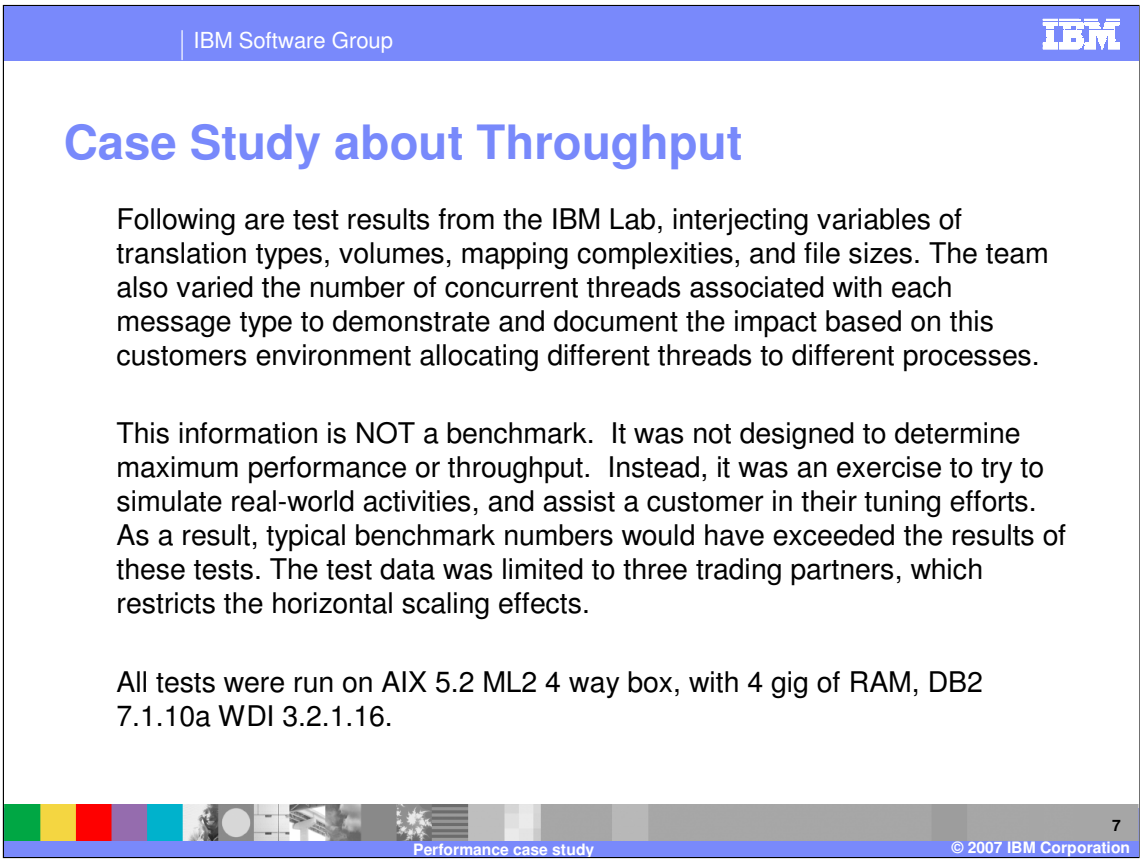

Following are test results from the IBM Lab, interjecting variables of translation types, volumes, mapping complexities, and file sizes. The team also varied the number of concurrent threads associated with each message type to demonstrate and document the impact based on this customers environment allocating different threads to different processes.

This information is NOT a benchmark. It was not designed to determine maximum performance or throughput. Instead, it was an exercise to try to simulate real-world activities, and assist a customer in their tuning efforts. As a result, typical benchmark numbers would have exceeded the results of these tests. The test data was limited to three trading partners, which restricts the horizontal scaling effects.

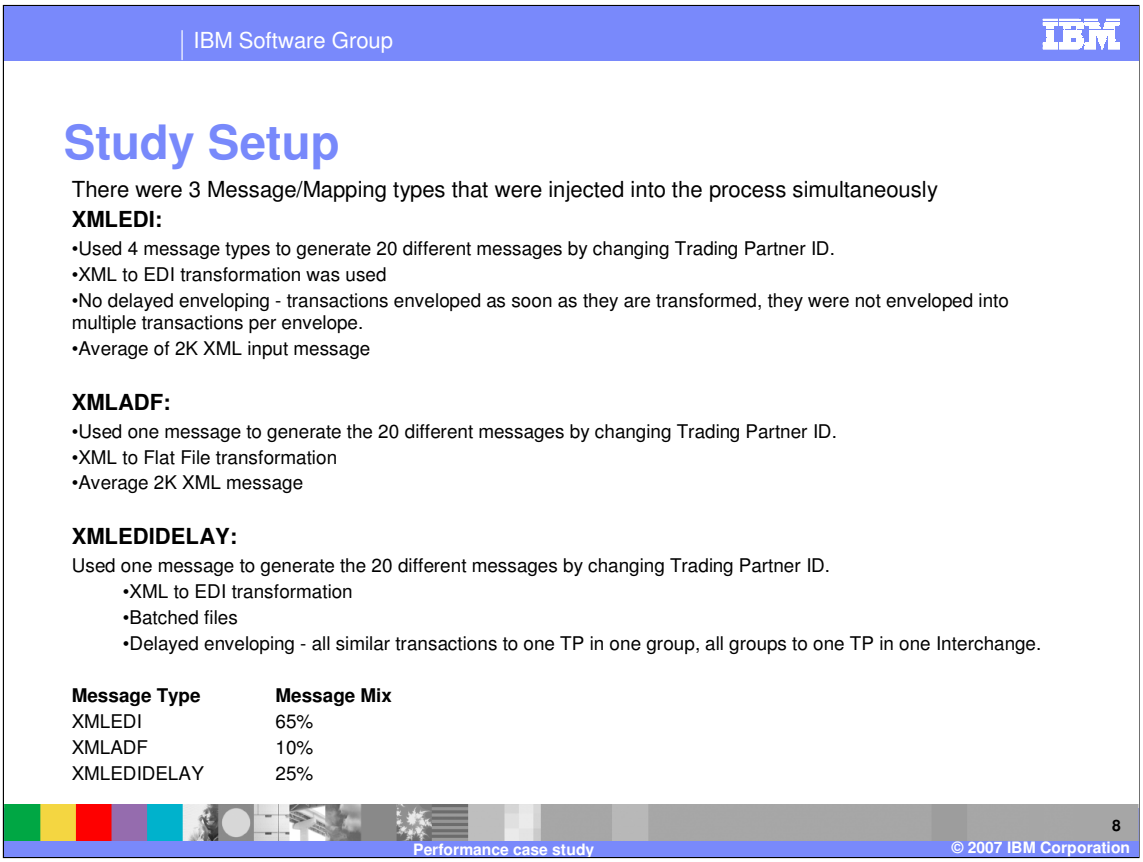

There were 3 Message/Mapping types that were injected into the process simultaneously

XMLEDI:

Used 4 message types to generate 20 different messages by changing Trading Partner ID.

XML to EDI transformation was used

No delayed enveloping - transactions enveloped as soon as they are transformed, they were not enveloped into multiple transactions per envelope.

Average of 2K XML input message

XMLADF:

Used one message to generate the 20 different messages by changing Trading Partner ID.

XML to Flat File transformation

Average 2K XML message

XMLEDIDELAY:

Used one message to generate the 20 different messages by changing Trading Partner ID.

XML to EDI transformation

Batched files

Delayed enveloping - all similar transactions to one TP in one group, all groups to one TP in one Interchange.

Message Type	Message Mix
XMLEDI	65%
XMLADF	10%
XMLEDIDELAY	25%

Execution results

The wait time is about 80% and CPU usage is very low.
Memory is steady at 12MB per translator.

Message Type	Number of Threads	15 Minute Total	1 Hour Total
XMLEDI	4	7,743	30,972
XMLADF	2	3,387	13,548
XMLEDIDELAY	1	1,844	7,376
GRAND TOTAL		12,951	51,808

Execution 1: Total of 5 threads, achieving highest throughput
Execution 2: Total of 7 threads, minimizing unused threads
Execution 3: Total of 7 threads, held one thread to zero transactions to process

Results of executing the same throughput mix with differing number of processing threads is shown.

Case Study 2 - Send / Receive Performance

The test runs on a z/OS machine.

The input is a file of 235 Transactions. There are 2 maps being used.

The maps are complex.

Functional Acknowledgments are being generated.

FF0588 Command: PERFORM DEENVELOPE AND TRANSLATE WHERE FILEID(INPUT) DUPENV(Y)
FF0588 Command: RAWDATA(Y) RECOVERY(E) PURGINT(-1) FUNACKFILE(FACK)

A comparison of a Send / Receive execution with WDI 3.1 and WDI 3.2.

Send / Receive Performance

Send / Receive Map on DataInterchange Version 3.1

```
*-------------------------------------------------------------------------------------------------------------------------*
*       SYSTEM NAME - SCND                 S T E P   C O M P L E T I O N   R E P O R T            SYSTEM LEVEL - MVS/SP 7.0.4  *
*-------------------------------------------------------------------------------------------------------------------------*
*     JOB NAME : RPOPEA62      REGION        <16MB        >16MB                          CIO: 00:00:00.08                    *
*     STEP NAME: RUNDI         AVAILABLE:    10,216K   1,597,440K   PRV MOUNT: 0       WLM: BATCH                          *
*     STEP NO  : 3             REQUESTED:    10,216K           0K   SCR MOUNT: 0       SCN: BATCH3                         *
*     PGM NAME : IKJEFT01      USED USER:     1,628K      45,888K                       GRN:                                *
*     COND CODE: 0002          USED SYS :       516K      10,232K                       RCN:                                *
*-------------------------------------------------------------------------------------------------------------------------*
*   DDNAME    UCB   VOLUME    I/O TIME    MAX BLKSZ   EXCP COUNT    DDNAME    UCB   VOLUME    I/O TIME    MAX BLKSZ   EXCP COUNT  *
*                                                                                                                         *
*   STEPLIB   762E  SN0210  00:00:00.00     19,069            2    STEPLIB   721E  S60312  00:00:00.03     19,069           36  *
*   STEPLIB   7519  S60301  00:00:00.00     19,069           60    STEPLIB   741E  S60371  00:00:00.00     19,069          248  *
*   STEPLIB   7619  S60347  00:00:00.16     19,069          274    STEPLIB   7000  S60100  00:00:00.00     19,069           31  *
*   SYSTSIN   7040  S60438  00:00:00.00      6,160            2    FFSWORK   7B24  S60503  00:00:00.68     32,760        2,156  *
*   TTABLE01  7621  S60451  00:00:00.00     27,920            2    FACK      721A  S60296  00:00:00.12     23,440          156  *
*   INPUT     7519  S60301  00:00:00.46     27,000          265    INPUT     7719  S60306  00:00:00.46     27,000          267  *
*   IN811     7621  S60451  00:00:03.82     27,000        2,191                                                               *
*-------------------------------------------------------------------------------------------------------------------------*
IEF373I STEP/RUNDI   /START 2005145.0946
IEF374I STEP/RUNDI   /STOP  2005145.0957 CPU    2MIN 20.98SEC SRB    0MIN 00.15SEC VIRT  1628K SYS   516K EXT   45888K SYS   10232K
*-------------------------------------------------------------------------------------------------------------------------*
*       SYSTEM NAME - SCND                 J O B   C O M P L E T I O N   R E P O R T            SYSTEM LEVEL - MVS/SP 7.0.4  *
*-------------------------------------------------------------------------------------------------------------------------*
*             HIGHEST CONDITION   JOB START   JOB START    JOB END      JOB END     JOB ELAPSED TIME        PRIVATE     SCRATCH  *
*  JOBNAME          CODE             DATE         TIME        DATE         TIME      (HHHH:MM:SS.TH)         MOUNTS      MOUNTS   *
*                                                                                                                         *
*  RPOPEA62         0002         05/25/2005  09:46:15.91  05/25/2005  09:57:25.19      00:11:09.28               0           0   *
*-------------------------------------------------------------------------------------------------------------------------*
```


The WDI 3.1 results.

Send / Receive Performance

Send / Receive Map on WebSphere Data Interchange Version 3.2.1

```
*-------------------------------------------------------------------------------------------------------------------------------*
*       SYSTEM NAME - SCND                  S T E P   C O M P L E T I O N   R E P O R T             SYSTEM LEVEL - MVS/SP 7.0.4  *
*-------------------------------------------------------------------------------------------------------------------------------*
*     JOB NAME : RPOPEA51     REGION        <16MB        >16MB                            CIO: 00:00:00.06                      *
*     STEP NAME: RUNDI        AVAILABLE:    10,216K   1,597,440K   PRV MOUNT: 0          WLM: BATCH                            *
*     STEP NO  : 3            REQUESTED:    10,216K           0K   SCR MOUNT: 0          SCN: BATCH3                           *
*     PGM NAME : IKJEFT01     USED USER:     1,580K      46,032K                          GRN:                                  *
*     COND CODE: 0002         USED SYS :       484K       9,836K                          RCN:                                  *
*-------------------------------------------------------------------------------------------------------------------------------*
*   DDNAME    UCB   VOLUME    I/O TIME    MAX BLKSZ   EXCP COUNT    DDNAME    UCB   VOLUME    I/O TIME    MAX BLKSZ   EXCP COUNT  *
*                                                                                                                               *
*   STEPLIB   762E  SN0210  00:00:00.00     19,069            2    STEPLIB   721E  S60312  00:00:00.05     19,069           45  *
*   STEPLIB   7414  S60299  00:00:00.35     19,069          541    EDITSIN   7040  S60438  00:00:00.00      6,160            2  *
*   SYSTSIN   7040  S60438  00:00:00.00      6,160            2    FFSWORK   7FFF  VIO     00:00:00.00     32,760          888  *
*   FACK      7A43  S60586  00:00:00.13     23,440          159    INPUT     7519  S60301  00:00:00.46     27,000          265  *
*   INPUT     7719  S60306  00:00:00.46     27,000          267    IN811     7621  S60451  00:00:03.42     27,000        1,955  *
*-------------------------------------------------------------------------------------------------------------------------------*
IEF373I STEP/RUNDI   /START 2005144.1930
IEF374I STEP/RUNDI   /STOP  2005144.1937 CPU    1MIN 26.19SEC SRB    0MIN 00.11SEC VIRT  1580K SYS   484K EXT   46032K SYS    9836K
*-------------------------------------------------------------------------------------------------------------------------------*
*       SYSTEM NAME - SCND                  J O B   C O M P L E T I O N   R E P O R T             SYSTEM LEVEL - MVS/SP 7.0.4  *
*-------------------------------------------------------------------------------------------------------------------------------*
*            HIGHEST CONDITION   JOB START   JOB START    JOB END      JOB END     JOB ELAPSED TIME        PRIVATE      SCRATCH  *
*  JOBNAME         CODE             DATE         TIME        DATE         TIME      (HHHH:MM:SS.TH)         MOUNTS       MOUNTS   *
*                                                                                                                               *
*  RPOPEA51         0002         05/24/2005  19:30:01.61  05/24/2005  19:37:14.82      00:07:13.21              0            0    *
*-------------------------------------------------------------------------------------------------------------------------------*
```


WDI 3.2.1 results.

Send / Receive Performance

	CPU Time	Elapsed Time
Send / Receive Map V3.1	2 Min 21 Secs	11 Min 09 Secs
Send / Receive Map V3.2.1	1 Min 26.19 Sec	7 Min 13 Secs

Conclusion: Send / Receive maps run at least as well in V3.2 as in V3.1, and could be significantly better.

Conclusion: Send / Receive maps run at least as well in V3.2 as in V3.1, and could be significantly better.

Summary

- Pageable AMM is a significant benefit with very large messages
- Send / Receive performance for 3.1 and 3.2 are statistically the same

Pageable AMM is a significant benefit with very large messages
Send / Receive performance for 3.1 and 3.2 are statistically the same

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	CICS	IMS	WMQ	Tivoli
IBM(logo)	Cloudscape	Informix	OS/390	WebSphere
e(logo)business	DB2	iSeries	OS/400	xSeries
AIX	DB2 Universal Database	Lotus	pSeries	zSeries

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a registered trademark of Linus Torvalds.

Other company, product and service names may be trademarks or service marks of others.

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 and/or changes in the product(s) and/or program(s) 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 (e.g., 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.