The slide features a blue header and footer with a pattern of glowing blue circles and a grid. The IBM logo is in the top right, and the text 'IBM eServer™' is in the top left. The word 'Miscellaneous' is centered in the white middle section. The footer contains the slogan '@business on demand software' and the copyright notice '© 2007 IBM Corporation'.


IBM eServer™

# Miscellaneous

@business on demand software

© 2007 IBM Corporation

## Agenda - Miscellaneous



- 1 **RPCBIND - IPv6 Support for RPC**
- 2 **Prevent the IP Address of a Dynamic VIPA from becoming OSPF Routerid**
- 3 **Packet Trace Formatting Enhancements**
- 4 **SMTP Support for dynamic timezone setting**
- 5 **Resolver Timeout Values below one Second**
- 6 **Recent APAR activity**

RPCBIND  
IPv6 Support for RPC

## RPC Standards

- **RFC 1831 RPC: Remote Procedure Call Protocol Specification Version 2**
  - Defines Open Network Computing (ONC™) Remote Procedure Call Version 2
  - Defines the RPC model
- **RFC 1832 XDR: External Data Representation Standard**
  - Defines how data is encoded over the network for RPC
- **RFC 1833 Binding Protocols for ONC RPC Version 2**
  - Defines binding protocol Version 2, 3 and 4.
  - Binding protocol Version 2 is known as the Portmap program protocol
  - Binding protocol Version 3 & 4 are known as the RPCBIND program protocol
- **z/OS® PORTMAP only supports RPC Binding Protocol 2**
  - Portmap binding protocol is compatible with IPv4 TCP and UDP
  - It does not support IPv6
- **z/OS NFS uses RPC protocol**
  - RFC 1833 support on z/OS is needed to support IPv6 networked mapped disks

## RPCBIND

- **RPCBIND is a new server in z/OS V1R8**
  - udp, udp6, tcp and tcp6 transports only
- **Allows NFS disks to be used in a IPv6 Network**
- **No application change required to move to RPCBIND**
- **RPCBIND can functionally replace PORTMAP**
  - RPCBIND supports a superset of PORTMAP's APIs
  - Archive library used to create RPC applications has not changed
  - Existing RPC client and server applications run as-is
  - Messaging and logging are different in RPCBIND
  - PORTMAP continues to be shipped for migration purposes
- **Supports multiple TCP/IP stacks**
  - RPCBIND allows RPC servers (like NFS) that can handle multiple stacks seamless migration
  - RPC servers that don't handle multiple stacks will continue to receive requests from clients over their original TCP/IP stack
- **Improved tracing and logging**
  - Logging shows which RPC APIs used
  - Useful for customers to see RPCBIND activity
  - Dynamically change tracing

## Changes for RPCBIND

- **Create a proclib member named RPCBIND similar to PORTMAP**
- **Change TCP/IP profile to autolog RPCBIND instead of PORTMAP**
- **Change TCP/IP profile to reserve port 111 (tcp and udp) for RPCBIND instead of PORTMAP**
- **Update any security programs like RACF®**
  - Sample SEZAINST(EZARACF) shows security server commands
- **RPCBIND has some enhanced startup options over PORTMAP**
  - i option to specify the directory for the rpcbind.pid file
  - d turns on all debug statements that go to the syslog daemon
    - -dl turns on logging for each RPC API issued
    - -df turns on flow logging generally for IBM Service
    - -dx turns on XDR logging generally for IBM Service
  - n option for installations that want to run non-swappable
- **That's it!**

## Migration

### ➤ **Both RPCINFO and ORPCINFO remain unchanged**

- These shipped applications only query IPv4 information
- To query IPv6 information use rpcinfo applications from other platforms in the network

### ➤ **Archive library unchanged**

- The API library shipped only supports IPv4 APIs
- Both RPCGEN and ORPCGEN remain unchanged

### ➤ **Registering applications**

- When registering an application, use IPv4 and IPv6 addresses
- It is not recommended to use IPv4 mapped IPv6 addresses
- If a server is willing to accept applications over any IPv4 and IPv6 addresses,
  - Register with INADDR6\_ANY and INADDR\_ANY

### ➤ **Size limitations of IPv6 networks**

- RPCBPROC\_GETADDRLIST API can use UDP packets
- For UDP packets size limits the number of addresses returned
- TCP may be more useful for this API
  - Especially if many interfaces and
  - Many servers registered with INADDR6\_ANY and INADDR\_ANY

Prevent the IP Address of a  
Dynamic VIPA from becoming  
OSPF Routerid



## OMPROUTE router ID

- **If a Dynamic VIPA is explicitly coded as a Routerid (which is NOT recommended), a new messages EZZ8134I will be issued.**
  - EZZ8134I OSPF ROUTERID ipaddr IS A DYNAMIC VIPA. THIS IS NOT A RECOMMENDED CONFIGURATION FOR jobname
  
- **If no specific Routerid is coded in the OMPROUTE Configuration file, then OMPROUTE will now cycle through all the OSPF interfaces, skipping over the DVIPAs until an OSPF Interface is found which is not a DVIPA.**
  - OMPROUTE will now use this non-DVIPA Interface as its Routerid.
  - If the only interfaces known to OMPROUTE are DVIPAs, then use one for the Routerid, but issue the EZZ8134I message as a warning.
  
- **It should be extremely rare and definitely not recommended that OMPROUTE be started with only DVIPA interfaces coded as OSPF Interface statements, which would now be the only way where a DVIPA should be chosen as the Routerid at startup.**
  
- **Also a rare situation, if OSPF was not enabled at the startup of OMPROUTE but an OSPF\_Interface wildcard had been defined in the OMPROUTE Configuration file which matches a DVIPA which has been vary obeyed into the home list, and if a DVIPA is the first match on this wildcard, it will be chosen as Routerid, but now the EZZ8134I message will be issued as a warning.**

## What to do if message EZZ8134I is seen

- If the EZZ8134I message is received, check the output of
  - D TCP/IP,TCPCS,OMP,OSPF,STATS to determine which DVIPA is being used as the Routerid:

```

00- 12.24.58 f omprout1,ospf,stats
12.24.58 EZZ7856I OSPF STATISTICS 176 C
          OSPF ROUTER ID:      192.168.99.130
          EXTERNAL COMPARISON: TYPE 2
          AS BOUNDARY CAPABILITY: NO

ATTACHED AREAS:                1  OSPF PACKETS RCVD:                0
OSPF PACKETS RCVD W/ERRS:      0  TRANSIT NODES ALLOCATED:      2
TRANSIT NODES FREED:           1  LS ADV. ALLOCATED:            1
LS ADV. FREED:                  1  QUEUE HEADERS ALLOC:         32
QUEUE HEADERS AVAIL:           32  MAXIMUM LSA SIZE:             1428
# DIJKSTRA RUNS:                1  INCREMENTAL SUMM. UPDATES:    0
INCREMENTAL VL UPDATES:         0  MULTICAST PKTS SENT:         0
UNICAST PKTS SENT:              0  LS ADV. AGED OUT:            0
LS ADV. FLUSHED:                0  PTRS TO INVALID LS ADV:      0
INCREMENTAL EXT. UPDATES:       0

```

- Then explicitly code a Routerid which is not a Dynamic VIPA in the OMPROUTE Configuration file.
- If you are sharing OMPROUTE configuration files across multiple stacks and therefor cannot code the Routerid explicitly, ensure that there is at least one non-DVIPA interface coded as an OSPF Interface.
- OMPROUTE will need to be restarted for the change to take effect.

## Packet Trace Formatting Enhancements

## Packet tracing of encapsulated packets, such as GRE

- **With the advent of protocols like GRE we are seeing more packets that are encapsulated in other packets.**
- **GRE packets cause a particular problem: When used with a VIPA only the input packets use GRE. The outbound packets do not.**
- **Protocols that use encapsulation:**
  - GRE
  - AH
  - IPv4 over IPv6
  - and IPv6 over IPv4
  - ICMP and ICMPv6
  - ESP
- **The Packet Trace formatter is modified to set which IP header is to be used for selection**
  - FIRST - Use the first IP header for selection (the default)
  - LAST - Use the last IP header for selection
- **The Packet Trace formatter is modified to specify which IP header is formatted.**
  - FORMAT(ALL) - Format all IP headers of a packet (the default)
  - FORMAT(FIRST) - Format the first IP header of a packet
  - FORMAT(LAST) - Format the last IP header of a packet

## Select packets with a specific security parameters ID (SPI)

### ➤ OPTIONS( SPID( x'hhhhhhh' ) )

-Find packets with matching security parameters ID value

```

COMPONENT TRACE SUMMARY FORMAT
SYSNAME(VIC034)
COMP(SYSTCPDA)
OPTIONS((SPID(X'A0654CA1'))))
z/OS TCP/IP Packet Trace Formatter, (C) IBM 2000-2006, 2006.065
DSNAME('MWS.ISAKMP5.PTRACE')

**** 2005/08/16
I - Inbound packet
O - Outbound packet

DP  Nr hh:mm:ss.mmmmm IpId  Seq_num  Ack_num  Wndw  Flags           DatLn Source/Destination
OU  32 15:40:22.416255 001D protocol=UDP          44 192.168.21.2-4500
      A0654CA1 00000001 *..<..... .eL.....*
      10.0.20.2-4503
OU  34 15:45:05.680990 001E protocol=UDP          60 192.168.21.2-4500
      A0654CA1 00000002 *..<..... .eL.....*
      10.0.20.2-4503
OU  37 15:45:22.113186 001F protocol=UDP          60 192.168.21.2-4500
      A0654CA1 00000003 *..<..... .eL.....*
      10.0.20.2-4503
OU  39 15:45:30.149967 0020 protocol=UDP          60 192.168.21.2-4500
      A0654CA1 00000004 *..<..... .eL.....*
      10.0.20.2-4503

All packets check summed correctly
No packets required reassembly

=====
SYSTCPDA Trace Statistics
71 ctrace records processed

```

## Various new packet selection flag options

### > **FLAGS(ALL)**

- The packet must meet all the requirements specified in the FLAGS to be selected

### > **FLAGS(ANY)**

- The packet must meet at least one of the requirement specified in the FLAGS to be selected.

### > **FLAGS(BAD)**

- Select packets that have internal format errors.

### > **FLAGS(NTA)**

- Select packets that are form the OSA NTA function

SMTP Support for dynamic  
timezone setting

## SMTP/NJE and time zone settings

- **Originally the Date and Time Specification for mail was documented in RFC 822 showed the timezone as a character string occurring after the time value. For example, the character string "EST" represented Eastern Standard Time.**
- **Both SMTP started task and SMTPNOTE supported a configuration parameter called Timezone whose value represented the printable name of the local time zone which was appended to the time stamp in the RFC 822 headers that were generated by the code. It was to be set manually to whatever the customer needed for their environment.**
- **The timezone value used by the SMTP started task (SMTPPROC) and SMTPNOTE had to be hardcoded in the configuration parameters used by the code.**
  - Customers were required to stop and restart the SMTP started task with the new timezone value twice a year.
  - Customers were required to change all their copies of SMTPNOTE manually to update the timezone value to accurately reflect their system timezone.



## SMTP/NJE and dynamic time zone adjustments

- **The value of "SYSTZ" in the timezone parameter for the SMTP configuration file and the SMTPNOTE exec will allow the code to dynamically retrieve the value of the timezone from the system Communication Vector Table (CVT) control block.**
  - This value is then converted to a string with the format of + or - and four digits (+HHMM).
- **Everytime the SMTP start task or SMTPNOTE creates a new RFC822 date and time stamp header for mail it will use what is currently in the CVT.**
  - Therefore when the timezone is changed for the system using MVS SET CLOCK = hh.mm.ss or MVS SET TIMEZONE = {W|E}.hh.mm - the new value will be reflected immediately in the new date and time stamp headers that SMTP creates.
- **Note that existing date and time stamp values are NOT altered in the mail headers.**
- **This support was routed back to z/OS V1R4 and higher releases via APAR PK12000 in the following PTFs:**
  - UK07843 HIP6140
  - UK07844 HIP6150
  - UK07845 HIP6160
  - UK07846 HIP6170

Resolver Timeout Values below  
one Second

## APAR PK29603 - ResolverTimeOut value below one second

- The current minimum value for ResolverTimeOut is one second.
- In some environments that rely heavily on name resolution activity, this doesn't allow for a timeout value low enough to meet the overall response time objectives.
- This APAR allows the TCPIP.DATA ResolverTimeOut statement to specify other than whole seconds.
- The fraction of seconds that can be specified is in milliseconds.

```

      _RESOLVERTIMEOUT 30 _____
>> | _____ | ____<
    | _____ _RESOLVERTIMEOUT__time_out_value_|
    | _system_name: _|
  
```

- The time can be specified in whole seconds, milliseconds or a combination of both.
  - Values below one second may be specified with or without a leading zero before the decimal point:
    - \*.010
    - \*.010
  - The default timeout is 30 seconds (unchanged)
  - The minimum value is 10 milliseconds (0.010)
- A time\_out\_value less than 10 milliseconds will be set to 10 milliseconds (0.010).
  - For example, RESOLVERTIMEOUT 0.005 will be processed as if it were RESOLVERTIMEOUT 0.010
- A time\_out\_value of 0 is equivalent to RESOLVERTIMEOUT 1. This is unchanged.
- Specifying more than three decimal positions is considered a parse error and is ignored.
  - For example, RESOLVERTIMEOUT 0.0100 is a parse error and is not processed.

## Understanding how the z/OS Communications Server IP resolver queries DNS for a host name to IP address resolution

➤ **This example is for a GetHostByName API call to obtain an IPv4 address.**

1 Assume that the default of OPTIONS ndots:1 is used and that the input host name contains no dots.

➤ **The resolver builds a DNS query containing the fully-qualified domain name (FQDN) for the requested host name. Based on our assumptions, it does that in the following manner:**


- 2 The resolver adds a domain name to the application input. The domain name is either specified by the DomainOrigin statement value or by the list of domain names specified by the SEARCH statement. If SEARCH is specified, the resolver starts with the first domain name specified by the search list, then uses the remaining domain names until the name is resolved by a DNS or until the list of domain names is exhausted.
- 3 The FQDN query is then sent using a UDP sendto to the first DNS IP address configured.
  - If the name server replies NOERROR, the resolution is complete and the resolution attempts terminate.
  - If the name server replies NXDOMAIN, the resolution attempt of this FQDN terminates. If there are other domain names in the SEARCH list, a new FQDN is created and another UDP sendto is performed.
  - If the query times out or a SERVFAIL reply is received, the next name server, if any, is tried with the current FQDN. If the name servers all time out, the current query is tried again with the same list of DNS until the ResolverUDPRetries value is reached. (Note that ResolverUDPRetries is the total number of attempts to contact a DNS, not just the number of retries.) If the ResolverUDPRetries value is reached and all the name servers have timed out, the resolution is terminated, even if there are additional domain names in the SEARCH list specification.
- 4 If after trying all FQDN and the name is not resolved, the resolver tries one last DNS query specifying the original name as presented across the API as the FQDN.
- 5 If none of the DNS queries resolve the name, the resolver tries local host tables as determined by the API environment of the call (for example, Native MVS™ or z/OS UNIX®).

## Resolver example 1 - all name servers time out

Local application program


```
gethostbyname(hostx)
```

```
RESOLVETIMEOUT 10
RESOLVERUDPRETRIES 2
SEARCH a.com
SEARCH b.com
OPTIONS ndots:1
NAMESERVER 9.99.99.98
NAMESERVER 9.99.99.99
LOOKUP DNS LOCAL
```



If NXDOMAIN response from any query, continue with B.COM  
 If all queries for a.com time out, do not try any other suffixes (b.com)

(Skipping b.com since all queries for a.com timed out)



If NXDOMAIN response from any query or all queries time out, continue with local hosts tables

Local HOSTS tables  
 hostx.a.com  
 hostx.b.com  
 hostx

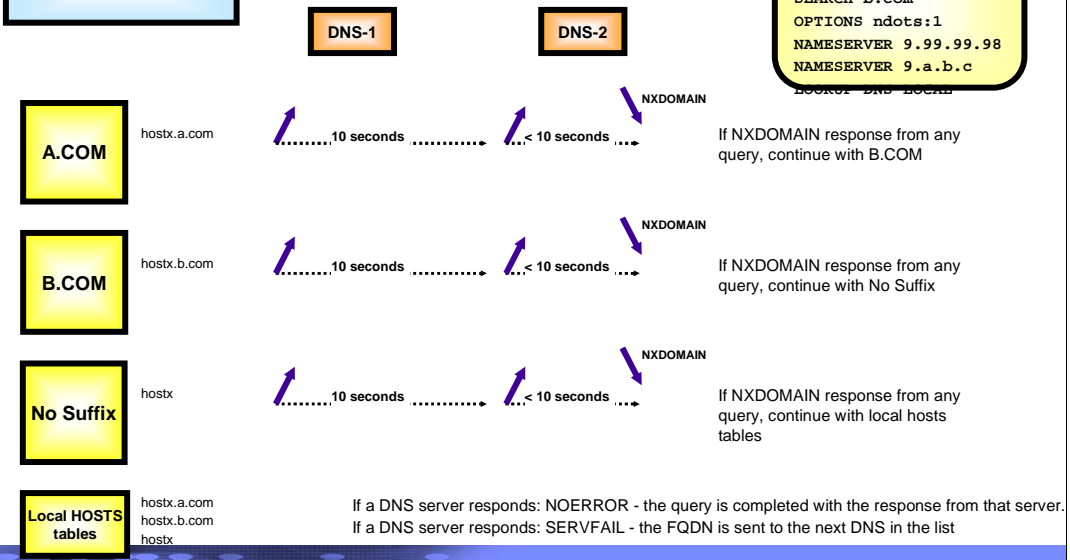
If a DNS server responds: NOERROR - the query is completed with the response from that server.  
 If a DNS server responds: SERVFAIL - the FQDN is sent to the next DNS in the list

## Resolver example 2 - DNS-2 responds NXDOMAIN

Local application program

```
gethostbyname(hostx)
```

```
RESOLVETIMEOUT 10
RESOLVERUDPRETRIES 2
SEARCH a.com
SEARCH b.com
OPTIONS ndots:1
NAMESERVER 9.99.99.98
NAMESERVER 9.a.b.c
```



## Things to think about

- **A RESOLVERTIMEOUT value that is too small can increase the overall time to perform a name or address resolution by causing the resolver to unnecessarily cycle through the DNSs in the NSINTERADDR list.**
- **A RESOLVERTIMEOUT value that is too large can increase the overall time to perform a name or address resolution by causing the resolver to unnecessarily delay cycling through the DNSs in the NSTINTERADDR list.**
- **To significantly influence the amount of time taken to resolve the name, choose:**
  - The number of domain names in the SEARCH list
    - Don't make it longer than absolutely necessary
  - The number of DNS to query (NameServer/NSInterAddr)
  - The number of "dots" in a name before it is considered a fully-qualified name (OPTIONS ndots:)
  - The time to wait for a DNS response (ResolverTimeOut)
  - The number of times to try the list of DNS (ResolverUDPRetries)

## Recent APAR Activity



## A few selected recent APARs that added new functions

### ➤ **PQ98005**

#### -Description:

- The FTP client implements a new FTP.DATA statement `SECURE_HOSTNAME` (REQUIRED or OPTIONAL) that will allow the client to verify the hostname in the server's digital certificate as per RFC 2818. The hostname that the client is connecting to will be verified against the server's certificate. Either the common name or the subject alternate name contained in the Server's X.509 certificate will be used to validate the hostname. If the verification fails, the connection is terminated.

-PTFed back to z/OS V1R4

### ➤ **PQ89672 and PK15174**

#### -Description:

- The FTP client sends AUTH TLS, but some older servers require an AUTH SSL instead. The FTP client has been enhanced to support a new option on the `SECURE_MECHANISM` FTP.DATA statement for the client. The client will now support `SECURE_MECHANISM SSL` as well as `SECURE_MECHANISM TLS` and `SECURE_MECHANISM GSSAPI`.

-PTFed back to z/OS V1R5

### ➤ **PK16540**

#### -Description:

- New SMTP/NJE configuration parameter (`MAXMSGSENT`) limiting the number of messages sent on a single TCP/IP connection is needed for compatibility reasons with vendor software. These servers have implemented a limit on the number of messages they will accept and when they reach that limit a reply code of '552' is given. Causing the client to send an undeliverable notice back to the sender.

-PTFed back to z/OS V1R6

## A few selected recent APARs that added new functions - continued

### ➤PK17858

- Description:
  - New function in VTAM® to respond to an LDLC probe request.
- PTFed back to z/OS V1R4

### ➤PK21685

- Description:
  - New Function APAR to provide a configuration option to disable segmentation offload. TCP segmentation offload is enabled by default if all of the requirements outlined in PK02490 are met. The option is on GLOBALCONFIG and is [NO]SEGMENTATIONOFFLOAD.
- PTFed back to z/OS V1R6

### ➤PK24752

- Description:
  - The default query rate of WLM by TCP/IP is once per minute. In certain environments this is not rapid enough as the systems the values are for are constrained and changing at a faster rate. However, in other nvironments increasing the rate may be excessive as the values do not change as quickly. A new GLOBALCONFIG option has been added: SYSPLEXWLM POLL nn - where nn defaults to 60 seconds. Valid range is 1 to 180 seconds.
- PTFed back to z/OS V1R6

## A few selected recent APARs that added new functions - continued

### ➤ **PK25763 and PK21468**

#### -Description:

- Beginning in z/OS V1R5, JES3 implemented a change in their design such that data placed on the spool included trailing spaces. In previous releases, these spaces had been truncated. Currently, FTP--when using a FILETYPE of JES--will simply return the data exactly as it exists on the spool, with no alteration or truncation. As such, with this new JES3 implementation, it is possible that trailing spaces will be returned with any data retrieved. Depending on the number of FTP transfers occurring, as well as the volume of data transferred, this can introduce performance impacts. To prevent this, an FTP.DATA configuration statement (JESTRAILINGBLANKS [TRUE/FALSE]) will be added in order for users to have the option of allowing FTP to truncate these spaces.

- PTFed back to z/OS V1R5

### ➤ **PK26746**

#### -Description:

- The FTP client should allow the CCC command to flow to clear the control connection after the initial andshake. The data connection can remain private.

- Will be PTFed back to z/OS V1R4 - PTFs not yet available (August 2006)



## 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            MVS            RACF            VTAM            z/OS

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

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

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 2007. 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.