



Communications Server z/OS® V1R5 and V1R6 Technical Update

Enterprise Extender and SNA

in z/OS V1R5 CS

© Copyright International Business Machines Corporation 2004. All rights reserved.



@server

➤ z/OS V1R5

- ⌈ EE IPv6 and NAT GVRN support
- ⌈ Multiple VRN support
- ⌈ EE model PU
- ⌈ Sift-down in model major nodes
- ⌈ EE network management
- ⌈ D RTPS command
- ⌈ Concurrent APINGS
- ⌈ EE performance enhancements
- ⌈ Improved APPN search diagnostics
- ⌈ SSCPORDER flexibility
- ⌈ Non-Sysplex NNs for generic resource
End Nodes
- ⌈ HPDT packing
- ⌈ XCA slowdown monitoring

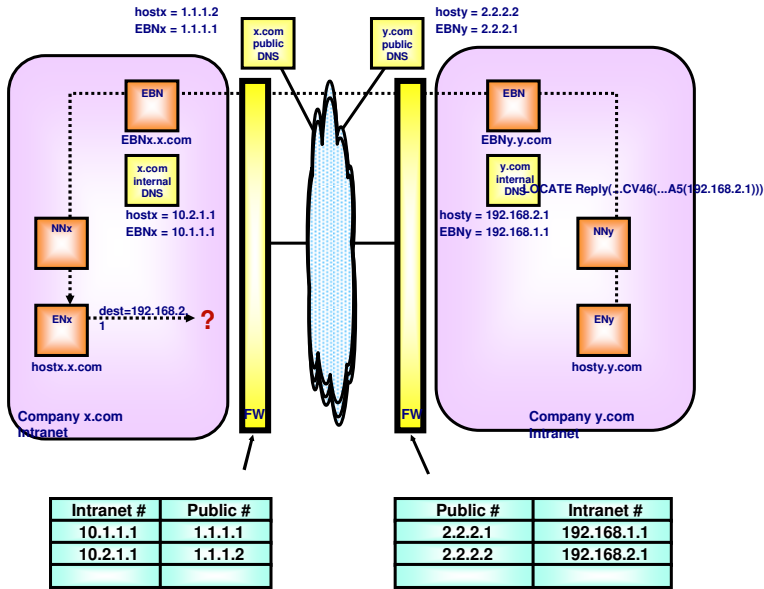


Enterprise Extender and SNA in z/OS V1R5 CS

Copyright International Business Machines Corporation 2004. All rights reserved.



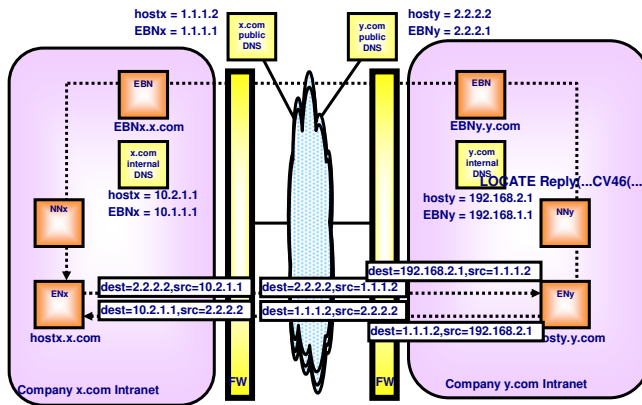
NAT Compatibility with EE VRN



➤ The control vectors representing an EE connection network carry the IP address of the node that owns the vectors.

➤ Applications that carry IP addresses in their data stream are fundamentally incompatible with Network Address Translation (NAT).h

Enhanced HOSTNAME Support

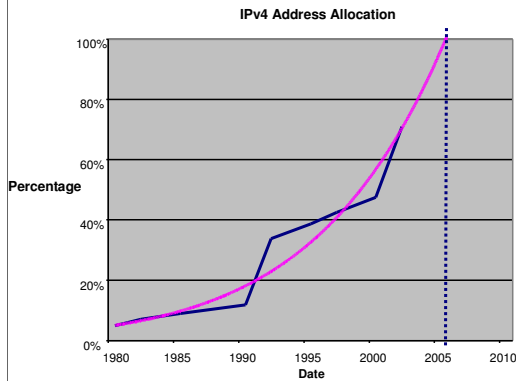


- EE architecture has been updated to allow the EE connection network control vectors to carry the hostname corresponding to the EE VIPA.
 - LOCATE Reply...CV46(...A5(hosty.y.com))
- IPv4 connection network control vectors will continue to carry IP address as well as hostname to ensure compatibility with downlevel nodes.
- Administrative requirement of coordinating NAT tables and public DNS entries is a known administrative procedure to installations that use NAT.
- There will be recommended maximum fully-qualified hostname lengths due to limited space in the route selection control vector.

Intranet #	Public #
10.1.1.1	1.1.1.1
10.2.1.1	1.1.1.2

Public #	Intranet #
2.2.2.1	192.168.1.1
2.2.2.2	192.168.2.1


EE IPv6 Support



➤ It is projected that IPv4 address demand will exceed supply and that IPv6 will be required for parts of the Internet within the next few years.

f 128-bit IPv6 addresses provide an address space with 340 billion billion billion billion addresses (as compared to 1 billion for IPv4)

- Some projections state that by 2005 half of the nodes on the Internet will be IPv6 nodes.
- EE needs the capability to communicate over an IPv6 network.

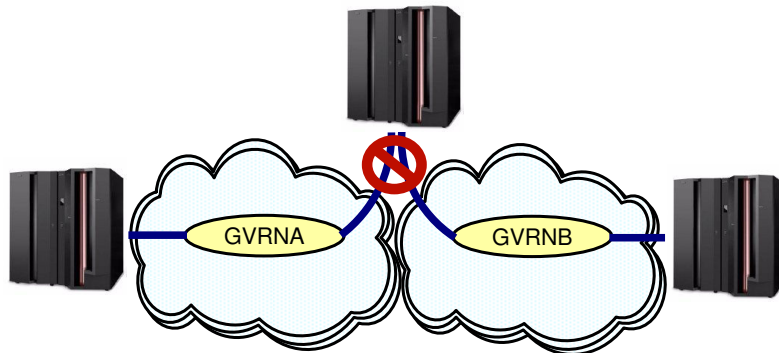


EE IPv6 Support...



- The Enterprise Extender architecture has been updated to allow for the exploitation of an IPv6-enabled network.
- The architectural/definitional requirements overlap with those for the Enhanced Hostname Support (for EE/NAT compatibility).
- Existing IPADDR keywords (start option, path definition in switched major node) are IPv4-only
- IPv6 support requires use of the HOSTNAME keyword (start option, GROUP, path definition)
- EE Connection networks must be IPv4-only or IPv6-only
- Nodes supporting both IPv4 and IPv6 will define an IPv4 VRN (local and/or global) and an IPv6 VRN
- Architectural changes have been submitted to other products supporting EE

Multiple Virtual Routing Node Limitation



- Prior to V1R2, only one EE connection network could be defined.
- V1R2 allows the definition of one local and one global connection network.

Multiple VRN Limitation...




➤ A single virtual routing node is not sufficient in some cases:

⌘ There are two or more disjoint IP networks to which a given VTAM is connected. All nodes connected to a given VRN must, by definition, be able to connect directly to any other node also connected to that VRN. This restriction means that only one of the two IP networks can use a connection network, with the other network having to manually define all connections from VTAM to all other nodes in that network.

⌘ Different VRNs may need to support different link characteristics. For instance, a subset of the interfaces out of a S/390 may need to run over secure links, while others may be unsecure. Depending on the requirements of the session, users may need to connect to the S/390 using the appropriate session characteristics. Restricting the number of connection networks to one GLOBAL and one LOCAL does not satisfy this requirement.

Multiple VRN/VIPA Support



- V1R5 will allow the specification of multiple local and/or multiple global EE connection networks.
 - ┆ In the diagram above, Node B defines 2 local VRNs (both IPv4) and 2 global VRNs (one IPv4 and one IPv6)
- EE will allow multiple (static) VIPAs, defined on a GROUP basis in the EE XCA major node.
 - ┆ All EE VIPAs must still belong to a single TCP/IP stack

Example EE XCA Major Node



```
XCAEE1A VBUILD TYPE=XCA
PORTEE  PORT  MEDIUM=HPRIP
* Pre-Defined EE
GPEE1   GROUP ANSWER=ON, *
        AUTOGEN=(1,LNEE1,PUEE1), *
        CALL=INOUT, *
        DIAL=YES, *
        IPADDR=10.1.1.1, *
        ISTATUS=INACTIVE
* Pre-Defined EE
GPEE4   GROUP ANSWER=ON, *
        AUTOGEN=(1,LNEE4,PUEE4), *
        CALL=INOUT, *
        DIAL=YES, *
        HOSTNAME=VIPA24.SSCP1A, *
        ISTATUS=INACTIVE
* Pre-Defined EE
GPEE6   GROUP ANSWER=ON, *
        AUTOGEN=(1,LNEE6,PUEE6), *
        CALL=INOUT, *
        DIAL=YES, *
        HOSTNAME=VIPA26.SSCP1A, *
        ISTATUS=INACTIVE
```

Display of EE XCA Major Node



```
d net,id=xcae1a,e
IST097I DISPLAY ACCEPTED
IST075I NAME = XCAE1A, TYPE = XCA MAJOR NODE
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST1679I MEDIUM = HPRIP
IST1685I TCP/IP JOB NAME = TCPCS
IST924I -----
IST1902I GROUP = GP EE1
IST1680I LOCAL IP ADDRESS 10.1.1.1
IST924I -----
IST1902I GROUP = GP EE4
IST1680I LOCAL IP ADDRESS 10.1.2.1
IST1910I LOCAL HOSTNAME VIPA24.SSCP1A
IST924I -----
IST1902I GROUP = GP EE6
IST1680I LOCAL IP ADDRESS 10.:1:2:1
IST1910I LOCAL HOSTNAME VIPA26.SSCP1A
IST924I -----
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1656I VTAMTOPO = REPORT, NODE REPORTED - YES
IST170I LINES:
IST1901I LINES UNDER GROUP: GP EE1
IST232I LNEE1000 ACTIV
IST1901I LINES UNDER GROUP: GP EE4
IST232I LNEE4000 ACTIV
IST1901I LINES UNDER GROUP: GP EE6
IST232I LNEE6000 ACTIV
IST314I END
```

Enterprise Extender Model PUs



- Currently, unless the Configuration Services XID exit is used, Enterprise Extender does not allow the user to code a model for non-connection-network dynamic PUs created on the dial-in side.
- If no predefined PU matching the caller is available and DYNPU=YES is coded on an EE switched line group in the XCA major node, EE creates a "vanilla" dynamic PU.
- The characteristics of this "vanilla" PU are all of the default Switched Major Node PU characteristics, with the following exceptions:
 - ƒ MAXOUT=8
 - ƒ ANS=CONT
 - ƒ DISCNT=(YES,F)
 - ƒ DYNADJCP=YES
 - ƒ CPCP=YES
- This approach does not allow a customer to customize specific characteristics for non-connection network PUs without using the Configuration Service XID exit.

Enterprise Extender Model PUs...



- A new model type, DYNTYPE=EE, is added to the PU statement in the Model Major Node.
- An EE model PU allows for customizing dynamic non-connection network PUs in various manners:
 - ⌘ Coding the DISCNT operand so that a disconnect delay time can be specified.
 - ⌘ Specifying the DWINOP, REDIAL, and REDDELAY operands to specify whether to drive redial attempts, and how often and how many times to attempt redial, when INOPs occur for the connections using these dynamic PUs.
 - ⌘ Overriding the default TG characteristics that in the past were used for these dynamic PUs.

Example EE Model PU



```
*****
*
*   NAME: MODEL1A   EXAMPLE MODEL MAJOR NODE
*
*****
MODEL1A VBUILD  TYPE=MODEL
*
EEMODEL  PU      DYNTPYE=EE,
                  CAPACITY=100K,
                  COSTTIME=5,
                  CPCP=YES,
                  DISCNT=NO,
                  DWINOP=YES,
                  REDIAL=30,
                  REDDELAY=60
```

Sift-down support for model major nodes



- Rather than coding the same keywords on all the LU definitions in a model major node, customers would like a way to code common LU keywords at a single place to sift down to all the underlying LU definitions. This will make coding many LUs in a model major node easier and less error prone.
- To accomplish this, the capability to code a GROUP statement in the model major node has been added in z/OS V1R5. The system programmer will be allowed to code LU keywords on the GROUP statement and those keywords will sift down to the LU statements below, unless they are overridden on the LU statement or another GROUP statement is encountered.
 - ┆ The GROUP definition statement is now allowed to be specified on the model major node.
 - ┆ LU keywords are allowed to be specified on the GROUP definition statement.
 - ┆ LU keywords specified on the GROUP definition statement will sift down to the underlying LU definition statements unless overridden by the keywords coded on the LU definition statement itself.
 - ┆ The sift-down of the LU keywords specified on the GROUP definition statement will remain in effect until another GROUP definition statement is specified or until the end of the model major node.
 - ┆ PU keywords are not allowed on the GROUP definition statement.
 - ┆ LU keywords are not allowed on the PU definition statement.
 - ┆ Recall that there is no relationship between the PU definition statements and the LU definition statements in the model major node.

Sift-down support for model major nodes - example



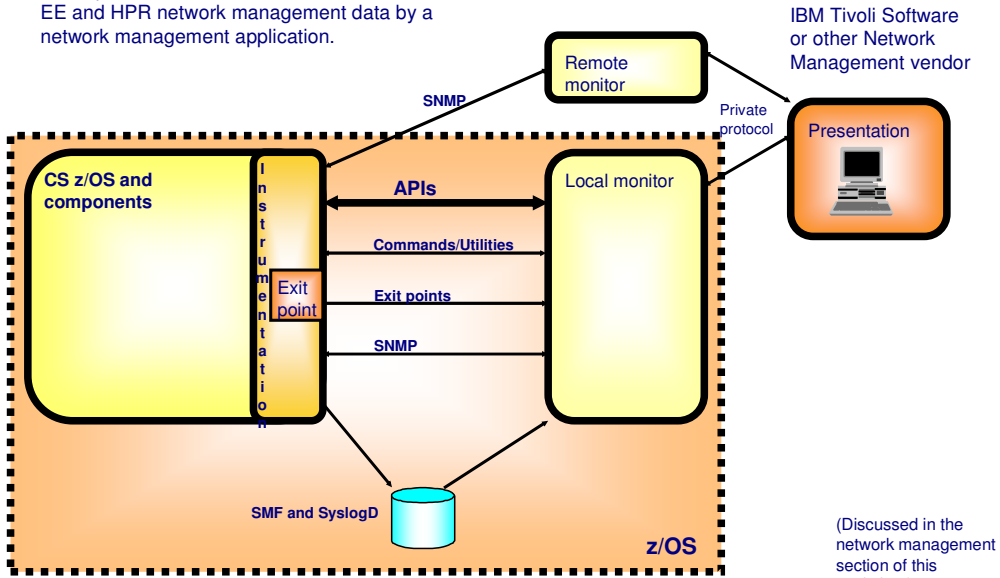
```
VBUILD  TYPE=MODEL
GROUP1  GROUP  CERTIFY=YES , CLRSESSQ=YES , ENCR=OPT
PU2     PU      ADDR=2
LU1     LU      LOCADDR=1
LU2     LU      LOCADDR=2
PU3     PU      ADDR=3
LU3     LU      LOCADDR=3
GROUP2  GROUP  CERTIFY=NO , CLRSESSQ=YES
PU4     PU      ADDR=4
LU4     LU      LOCADDR=4
LU5     LU      LOCADDR=5 , CLRSESSQ=NO
```

- In the example above, CERTIFY=YES, CLRSESSQ=YES, and ENCR=OPT sift down to LU1, LU2, and LU3.
- CERTIFY=NO and CLRSESSQ=YES sift down to LU4, but only CERTIFY=NO sifts down to LU5 because CLRSESSQ=NO is explicitly coded on LU5, thus overriding the GROUP2 coding of CLRSESSQ=YES.
- ENCR=NONE is the default that is applied to LU4 and LU5.
- None of the PU definition statements are affected by the GROUP, nor do they have any affect on the sifting of the LU keywords from GROUP definition statement to LU definition statement.

EE Network Management API



➤ V1R5 provides a new API to allow retrieval of EE and HPR network management data by a network management application.



(Discussed in the network management section of this workshop)

D RTPS Enhancement



- D NET,RTPS displays HPR pipes that match the input criteria. Prior to z/OS V1R5, several keywords (APPNCOS, CPNAME, CONGEST, SWITCH, and ID) could be used to determine the matching criteria.
- To address the need to display the HPR pipes traversing a particular first hop, we added the following keywords to the D NET,RTPS command: FIRSTTG, FIRSTCP, and ALSNAME. They can be used alone, in conjunction with each other, or with the other already existing keywords.
- To address the need to display HPR pipes going to a particular network, we added the capability of specifying the CPNAME keyword with a network qualified name, where the netid is specified, but the name is an asterisk (*). This will indicate that the netid portion of the CP name must match, but the name portion of the CP name will not be a matching criteria. For example, if CPNAME=NETB.* is specified, then IST1697I will only be displayed for HPR pipes with a destination CP in the network known as NETB. No other format other than netid.* will be allowed, when using the *.
- The new FIRSTCP operand can use the asterisk in the same way as CPNAME.
- The customer can use the new keywords individually, together, and in conjunction with the current keywords:
 - /D NET,RTPS,FIRSTTG=21
 - /D NET,RTPS,FIRSTCP=SSCP1A
 - /D NET,RTPS,FIRSTCP=NETA.SSCP1A
 - /D NET,RTPS,ALSNAME=AHHCPU1
 - /D NET,RTPS,FIRSTTG=21,FIRSTCP=SSCP1A
 - /D NET,RTPS,ALSNAME=AHHCPU1,CONGEST=YES
- The customer can use the netid.* format for the CPNAME and FIRSTCP keywords:
 - /D NET,RTPS,CPNAME=NETA.*
 - /D NET,RTPS,FIRSTCP=NETB.*

D RTPS Examples



```
d net, rtps, firsttg=21, firstcp=sscp2a
```

```
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = RTPS
IST1695I PU NAME      CP NAME      COS NAME SWITCH CONGEST  SESSIONS
IST1696I CNR00005 NETA.SSCP2A      #INTER   NO      NO      1
IST1696I CNR00004 NETA.SSCP2A      #BATCH   NO      NO      1
IST1454I 2 RTP(S) DISPLAYED
IST314I END
```

```
d net, rtps, alsname=ahhcpu1
```

```
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = RTPS
IST1695I PU NAME      CP NAME      COS NAME SWITCH CONGEST  SESSIONS
IST1696I CNR00001 NETA.SSCP1A      CPSVCMG  NO      NO      1
IST1454I 1 RTP(S) DISPLAYED
IST314I END
```

```
d net, rtps, cpname=neta. *
```

```
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = RTPS
IST1695I PU NAME      CP NAME      COS NAME SWITCH CONGEST  SESSIONS
IST1696I CNR00005 NETA.SSCP2A      #INTER   NO      NO      1
IST1696I CNR00004 NETA.SSCP2A      #BATCH   NO      NO      1
IST1696I CNR00001 NETA.SSCP1A      CPSVCMG  NO      NO      1
IST1454I 3 RTP(S) DISPLAYED
IST314I END
```

Multiple concurrent APING support



- The DISPLAY APING function is implemented using two Transaction Programs (TPs). They are ISTAPING and APINGDTP.
 - ┆ The ISTAPING TP is more commonly referred to as APINGTP. This is the TP that initiates the transaction. When a VTAM operator issues a DISPLAY APING command, the APINGTP will issue an allocate to start a conversation over an LU 6.2 session.
 - ┆ The APINGD TP behaves like a server. It receives the transaction from the ISTAPING TP. It completes allocation of the conversation and responds by sending reply data.
 - The MODIFY APINGDTP command allows you to change the number of APINGD transaction programs permitted to run concurrently to respond to APING requests from other nodes.
 - The DISPLAY APINGDTP command displays the number of APINGD transaction programs permitted to run concurrently to respond to APING request from other nodes.

- Customers would like to have the ability to issue multiple DISPLAY APING commands simultaneously.

- Customers want to use DISPLAY APING under programmatic control to gather performance data. Being limited to a single DISPLAY APING at a time greatly increases the time to gather such information for the entire network.

- Many customers are primarily concerned about error scenarios, where the DISPLAY APING target is entered incorrectly, or is not findable. In those cases, the operator must wait for the searching to complete and the previous DISPLAY APING to fail before being able to try again.

- In a case where the DISPLAY APING inadvertently connects to an APPC application which does not understand the APING transaction protocols, the target application might not respond and the conversation can hang until cancelled by the operator.

Multiple concurrent APING support - command changes



- The DISPLAY APING command will be enhanced to allow multiple requests to be outstanding at the same time.
- Additional enhancements will be made to further support the multiple DISPLAY APING function. These include:
 - ⌘ A new MODIFY APINGTP command will allow the instance limit for the ISTAPING transaction program to be updated by the operator. The command is similar to the existing MODIFY APINGDTP command, which allows modification of the instance limit for the APINGD transaction program.
 - ⌘ A new DISPLAY APINGTP command will allow the operator to display the current instance limit of the ISTAPING transaction program. This command is similar to the existing DISPLAY APINGDTP command.
 - ⌘ A new LIST=(ALL | SUMMARY) operand on the DISPLAY APING command gives the operator the ability to get only summary data about the APING transaction.
 - ⌘ A new LIST=(ALL | COUNT | ONLY) option has been added on the DISPLAY APINGDTP command.
 - ⌘ A new MAX=msglimit operand on the DISPLAY APINGDTP and DISPLAY APINGTP commands limits the number of response messages.
- New or changed commands
 - ⌘ F procname,APINGTP,INSTANCE=10 (new command)
 - ⌘ D NET,APINGTP,LIST=[**ONLY** | ALL | COUNT] (new command)
 - ⌘ D NET,APING,ID=SSCP2A,LIST=[**ALL** | SUMMARY]
 - ⌘ D NET,APINGDTP,LIST=[**ONLY** | ALL | COUNT]

Multiple concurrent APING support LIST=SUMMARY



D NET, APING, ID=SSCP2A, LIST=SUMMARY

```
IST097I DISPLAY ACCEPTED
IST1488I ACTIVATION OF RTP CNR00006 AS ACTIVE TO NETA.SSCP2A
IST1489I APING SESSION INFORMATION
IST1490I DLU=NETA.SSCP2A SID=EAABEEC351AAC19A
IST933I LOGMODE=#INTER , COS=*BLANK*
IST875I APPNCOS TOWARDS SLU = #INTER
IST1460I TGN CPNAME TG TYPE HPR
IST1461I 21 NETA.SSCP2A APPN RTP
IST314I END
IST1457I VTAM APING VERSION 2R33 (PARTNER TP VERSION 2R33)
IST1490I DLU=NETA.SSCP2A SID=EAABEEC351AAC19A
IST1462I ECHO IS ON
IST1463I ALLOCATION DURATION: 130 MILLISECONDS
IST1464I PROGRAM STARTUP AND VERSION EXCHANGE: 54 MILLISECONDS
IST1465I DURATION DATA SENT DATA RATE DATA RATE
IST1466I (MILLISECONDS) (BYTES) (KBYTE/SEC) (MBIT/SEC)
IST1468I TOTALS: 80 400 5 0
IST1469I DURATION STATISTICS:
IST1470I MINIMUM = 36 AVERAGE = 40 MAXIMUM = 44
IST314I END
```

IST1467I is not displayed for each iteration on LIST=SUMMARY.

Multiple concurrent APING support LIST=ALL



```
D NET,APING, ID=SSCP2A, LIST=ALL
IST097I DISPLAY ACCEPTED
IST1488I ACTIVATION OF RTP CNR00006 AS ACTIVE TO NETA.SSCP2A
IST1489I APING SESSION INFORMATION
IST1490I DLU=NETA.SSCP2A SID=EAABEEC351AAC19A
IST933I LOGMODE=#INTER , COS=*BLANK*
IST875I APPNCOS TOWARDS SLU = #INTER
IST1460I TGN CPNAME          TG TYPE          HPR
IST1461I 21 NETA.SSCP2A      APPN            RTP
IST314I END
IST1457I VTAM APING VERSION 2R33 (PARTNER TP VERSION 2R33)
IST1490I DLU=NETA.SSCP2A SID=EAABEEC351AAC19A
IST1462I ECHO IS ON
IST1463I ALLOCATION DURATION: 130 MILLISECONDS
IST1464I PROGRAM STARTUP AND VERSION EXCHANGE: 54 MILLISECONDS
IST1465I          DURATION      DATA SENT  DATA RATE  DATA RATE
IST1466I          (MILLISECONDS) (BYTES)    (KBYTE/SEC) (MBIT/SEC)
IST1467I              44           200           4           0
IST1467I              36           200           5           0
IST1468I TOTALS:          80           400           5           0
IST1469I DURATION STATISTICS:
IST1470I MINIMUM = 36 AVERAGE = 40 MAXIMUM = 44
IST314I END
```


Multiple concurrent APING support New APINGTP related commands



```
F NET,APINGTP,INSTANCE=10
  IST097I MODIFY ACCEPTED
  IST223I MODIFY APINGTP COMMAND COMPLETED

D NET,APINGTP,LIST=ONLY
  IST097I DISPLAY ACCEPTED
  IST1474I APING TP CONCURRENT INSTANCE LIMIT = 10

D NET,APINGTP,LIST=ALL
  IST097I DISPLAY ACCEPTED
  IST1474I APING TP CONCURRENT INSTANCE LIMIT = 10
  IST1870I APING HAS 1 ACTIVE SESSIONS
  IST1888I DLU = NETA.SSCP2A SID = F6ABEEC354AAC29B
  IST314I END

D NET,APINGTP,LIST=COUNT
  IST097I DISPLAY ACCEPTED
  IST1474I APINGD TP CONCURRENT INSTANCE LIMIT = 10
  IST1870I APINGD HAS 1 ACTIVE SESSIONS
  IST314I END
```


EE performance overview



- > There has been significant effort put into EE performance enhancements
 - f Increased throughput
 - f Reduced CPU Utilization
- > Most of the enhancements have been delivered via PTFs
 - f Throughput: Apply OA02213 & PQ69398, and upgrade OSA Microcode Level to 3.26 (z/Series 2064 GA3) or 4.28 (G5/G6)
 - f CPU utilization: Some customer environments will benefit from OA04393 (Inactivity Timer Optimization)
- > Monitor the EE Informational APAR (II12223) for news on further enhancements

APAR	Purpose	PTF	Notes
OW53393	ARB Enhancements	UW94491	V1R2 only, base in V1R4
OW56896	LAN Idle	UA00067	V1R2 only, base in V1R4
OW52291	EE Packing & QDIOSTG Option	UA00131 - V1R2 UA00132 - V1R4	
OW53978	EE Outbound Data Ordering	UA00131 - V1R2 UA00132 - V1R4	Coreq: PQ69398
PQ69398	Fast UDP Outbound Ordering	UQ73923	V1R2 only, base in V1R4
OW57459 OW56893	HPR Resequencing	UA00131 - V1R2 UA00132 - V1R4	
OA02213	Send SRB Optimization	UA01999 - V1R2 UA02000 - V1R4	Will prereq all VTAM APARs in table

CPU consumption during roll-out of EE



© Copyright International Business Machines Corporation 2004. All rights reserved.

Improved APPN directory search diagnostics



- Customers would like the IST663I message groups to be issued under more conditions for console diagnoses of network problems.
- Today the IST663I message group is issued for DSRLSTs which failed to locate the destination resource. However, the message group is issued only if VTAM had sent the DSRLST for searches performed during the course of session establishment. The message group is not issued when the DSRLST fails to locate the destination resource for searches which may not result in a session being established (such as INQUIRE OPTCD=APPSTAT). The IST663I group is desirable for console diagnoses of network routing problems.
 - f A new start option - DSIRFMSG - has been included in this release to control the issuing of the IST663I for search only type of RU flows.
 - f Setting DSIRFMSG (to ALLSSCP, OLUSSCP or NONE) for search only RUs is functionally equivalent to setting SIRFMSG for session setup RUs.
 - f Adjacent SSCP failure messages IST894I, IST895I, IST1704I and/or IST1705I may be included in the IST663I group depending on the setting of the FSIRFMSG Start Option. (This has always been true for SIRFMSG. It is now also true for DSIRFMSG.)
 - f The value of DSIRFMSG can be changed with the MODIFY VTAMOPTS command while VTAM is running.
 - f The existing SIRFMSG operand on CDRSC definitions will now enable or disable both the SIRFMSG and DSIRFMSG functions for that CDRSC.

Improved APPN directory search diagnostics - configuration



- The DSIRFMSG start option can be coded with the following values:
 - ⌘ ALLSSCP - Specifies that this SSCP will always issue messages.
 - ⌘ OLUSSCP - Specifies that this SSCP will only issue messages when it owns the originating logical unit (OLU).
 - ⌘ NONE (default) - Specifies that this SSCP will never issue messages.

- The Modify VTAMOPTS command can be used to change the value of the DSIRFMSG start option.
 - ⌘ F procname,VTAMOPTS,DSIRFMSG=ALLSSCP
 - ⌘ F procname,VTAMOPTS,DSIRFMSG=OLUSSCP
 - ⌘ F procname,VTAMOPTS,DSIRFMSG=NONE

- The DSIRFMSG start option may be overridden at a resource level (CDRCS definition)
 - ⌘ Coding the SIRFMSG operand on a CDRSC will now override both the SIRFMSG and DSIRFMSG start options for that resource.

Vocabulary lesson:

- SIRFMSG stands for Session Initiation Request Failure MeSsaGes
 - ⌘ The F in FSIRFMSG means Failing Sense Code
 - ⌘ The E in ESIRFMSG means Extended Sense Data
 - ⌘ The D in DSIRFMSG means DSRLST (as opposed to CDINIT searches)
 - ⌘ The L in LSIRFMSG means APPN Locate searches [new in z/OS V1R6]

Improved APPN directory search diagnostics - example



EXAMPLE OF RESULTING MESSAGES

```
IST663I CD DSRLST REQUEST FAILED, SENSE=087D0001
IST664I REAL OLU=NETA.APPL1 ALIAS DLU=NETA.LOSTAPPL
IST889I SID = EAABEEC34D8ACCB0
IST1705I SORDER = APPN FROM START OPTION
IST1705I SSCPORD = PRIORITY FROM START OPTION
IST894I ADJSSCPS TRIED FAILURE SENSE ADJSSCPS TRIED FAILURE SENSE
IST895I ISTAPNCP 08420001 SSCP2A 087D0001
```

NOTES: Messages IST1705I, IST894I, and IST895I are included in the output when the FSIRFMSG start option is set to OLUSSCP (in the OLU host) or ALLSSCP (in any host).

Extended Sense Data (optionally displayed on IST891I, IST892I and IST893I) is not available for search-only RUs. Therefore, the ESIRFMSG start option (which is used to enable or disable these messages) has no affect on the output that is generated by the new DSIRFMSG start option.

Improved APPN directory search diagnostics APPN locate searches in z/OS V1R6



- z/OS V1R6 implements LSIRFMSG function for APPN Locates to display failure information. The valid values are NONE, OLUNNS, and ALLNNS.
 - ┆ The LSIRFMSG start option is only valid on Network Nodes (NNs)
 - ┆ SIRFMSG and/or DSIRFMSG can be used to capture search failure information on ENs.
 - ┆ The default will be NONE, because of the large amount of messages that will very likely be generated.
- The LSIRFMSG start option can be coded with the following values:
 - ┆ ALLNNS - Specifies that this NN will always issue messages.
 - ┆ OLUNNS - Specifies that this NN will only issue messages when it is the NNS of the OLU or a Central Directory Server.
 - ┆ NONE (default) - Specifies that this NN will never issue messages.
- The Modify VTAMOPTS command can be used to change the value of the LSIRFMSG start option.
 - ┆ F procname,VTAMOPTS,LSIRFMSG=ALLNSS
 - ┆ F procname,VTAMOPTS,LSIRFMSG=OLUNSS
 - ┆ F procname,VTAMOPTS,LSIRFMSG=NONE
- When FSIRFMSG and/or ESIRFMSG are used in conjunction with LSIRFMSG:
 - ┆ ALLSSCP is equivalent to ALLNNS
 - ┆ OLUSSCP is equivalent to OLUNNS
- The LSIRFMSG start option may be overridden at a resource level (CDRCS definition). Coding the SIRFMSG operand on a CDRSC will now override the LSIRFMSG start option for that resource (if CPNAME is also coded), as well as the SIRFMSG and DSIRFMSG start options.
- New messages for APPN Search Tasks.

Same flexibility for SSCPORD as for SORDER



- SORDER controls the order in which the APPN and subarea networks are searched when a network search request is received from the subarea network.
- SSCPORD specifies whether VTAM searches an adjacent SSCP table in priority order or in the order in which the table is defined.
- SSCPORD does not have the same flexibility of coding that SORDER has. Customers would like to have the ability to override the SSCPORD start option in the appropriate sections of the ADJSSCP tables, just as they can do now with SORDER.
- z/OS V1R5 will allow SSCPORD= to be specified on NETWORK and/or CDRM statements within an ADJSSCP Table definition.
- When coded on a NETWORK statement, the SSCPORD value sifts down to all subsequent ADJSSCP tables under that NETWORK statement (unless SSCPORD is overridden on the CDRM statement for one or more of the subsequent tables).
- If SSCPORD= is NOT coded for (and does not sift down to) one or more ADJSSCP tables, then the current SSCPORD Start Option value will be used for those ADJSSCP tables.
- This support will, most importantly, allow installations to use a defined order for SNI searches but use priority order for intra-network searches.

Same flexibility for SSCPORD as for SORDER - details on specifying

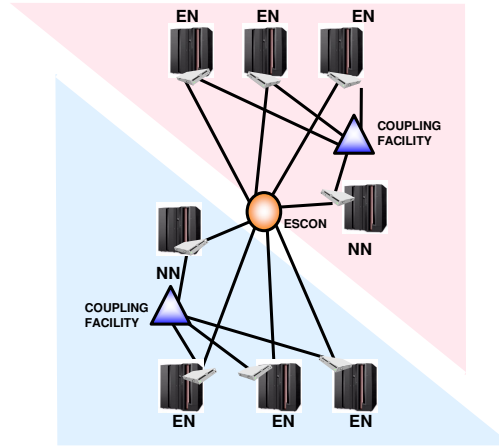


- When coding SSCPORD on NETWORK or CDRM statements in ADJSSCP definitions, the same values are allowed as for the SSCPORD start option (DEFINED/PRIORITY).
- An additional value (STARTOPT) is also allowed, to indicate that this ADJSSCP table should use the current value of the SSCPORD start option. This value can be used on CDRM statements to override the sift down value from the preceding NETWORK statement. When coded on a NETWORK statement, this value works the same way as when SSCPORD is not coded on the NETWORK statement.
- Currently, the value of SORDER is displayed in several message groups to show the search order that was in effect at the time these message groups were displayed. The SORDER value is displayed using one of two messages: IST1705I (indicating the SORDER start option value was used) or IST1704I (indicating the SORDER operand was coded on the chosen ADJSSCP table). These messages will be reused to also display the SSCPORD value by replacing the first token (SORDER) with SSCPORD (changing this token from "fixed text" to a "substitutable field"). The SORDER and SSCPORD values are displayed in this manner for the following situations:
 - ┆ When displaying ADJSSCP tables using DISPLAY ADJSSCPS
 - ┆ When displaying outstanding searches using DISPLAY SRCHINFO,SID=
 - ┆ When searches fail to find the target resource and the FSIRFMSG start option is in effect

Non-SYSPLEX Network Node Server for generic resources End Nodes



- Currently the network node server for end nodes running generic resource applications must be connected to the same coupling facility structure as the served end nodes. To avoid a single point of failure this requires two network node servers in each sysplex configuration.
- Customers would like to have the flexibility of having a backup network node server that is not connected to the same sysplex as the served end nodes but continues to support the generic resource function, including session level load balancing.



Non-SYSPLEX Network Node Server for generic resources End Nodes - how to specify



- For the generic resource function to work when the network node is not connected to the generic resource structure, the end node must allow searching for unknown resources. This will allow the LOCATE, with the generic name, to be forwarded to the end node so the resolution of the generic name can be done on the end node.
- APPN architecture does not allow for end nodes to reply to LOCATEs with a different owning CP name. Therefore only the end node that owns the real instance can return the positive found reply to the LOCATE. Generic name resolution will be done by the first end node to receive the search and access the generic resource structure. This will create an affinity for that session set up. But if the resolved real instance does not reside on this end node, the affinity will be used by the owning end node to confirm the resolved name (affinity) and return the positive reply.
- Customers will control this function by coding the new ENBCAST operand on the network node server list major node.
 - ‡ The default value is ENBCAST=NO.
 - ‡ Coding ENBCAST=YES for the backup network node server in the network node server list allows a backup network node server (outside the SYSPLEX) to search the end node for the generic resource name (which is considered an unknown resource) as part of the end node broadcast search phase.
 - ‡ This function relies on the backup network node server to find the real instance during the domain broadcast and redirect the search correctly. This means that if any of the generic resource end nodes are using the backup network node server, then all of the other generic resource end nodes must be using that same backup network node server as well.
 - ‡ Due to the possible performance implications, we do not recommend that customers run their Generic Resources configuration in this "backup mode" of operation for a long time. The "backup mode" should only be used during a temporary outage of the primary network node server...which should still be attached to the same SYSPLEX as the GR ENs.
- Message IST1253I will be updated to include the ENBCAST operand on the display of the network node server list (DISPLAY NETSRVR command).

Non-SYSPLEX Network Node Server for generic resources End Nodes - usage notes



NOTES

> This backup mode of operation only works when all end nodes connected to the same generic resource structure are using the same backup node as their network node server. For this reason, it is necessary to configure all of these end nodes such that they will only switch to the backup-mode network node server in the event of a failure of the primary network node server (rather than, for example, when a link failure occurs). This can best be achieved by following these guidelines:

- 1 Configure all end nodes attached to the same generic resource structure with the same network node server (NETSRVR) list. The first entry in the network node server list should define the primary network node server (that is, the network node that is attached to the same generic resource structure as the end nodes) and should specify (or default to) ENBCAST=NO. The second (and last) entry in the list should define the backup-mode network node server and should specify ENBCAST=YES. This is the only way to guarantee that all end nodes switch to the backup-mode network node server at the same time.
- 2 Code ORDER=FIRST in the network node server (NETSRVR) list, so all end nodes will always attempt to use the primary network node server (the one attached to the same generic resource structure) first.
- 3 Code the NNSPREF start option on all end nodes with the name of the primary network node server (the one attached to the same generic resource structure). This will cause all end nodes to automatically terminate the backup-mode of operation by switching back to the primary network node server as soon as it becomes available again.
- 4 To reduce the possibility of a link failure causing only some end nodes to switch to the backup-mode network node server, it is recommended that redundant links (which support CP-CP sessions) be defined between each end node and its primary network node server. XCF links are a very good choice for this, because they provide the necessary connectivity as long as there is connectivity to the coupling facility. (In fact, because connectivity to the coupling facility is required to support generic resources, defining XCF links may eliminate the need to define redundant links specifically for this purpose.)

Non-SYSPLEX Network Node Server for generic resources End Nodes - display commands



NOTES

D NET,NETSRVR

```
JOB 2 IST097I DISPLAY ACCEPTED
JOB 2 IST350I DISPLAY TYPE = NETWORK NODE SERVER LIST
IST1252I DEFINED NETWORK NODE SERVER LIST, NAME = NNSLCA
IST1253I NETA.SSCP1A SLUINIT = REQ ENBCAST = NO
IST1253I NETA.SSCP2A SLUINIT = REQ ENBCAST = YES
IST924I -----
IST1254I SERVER LIST PROCESSED ORDER = FIRST
IST924I -----
IST1255I OTHER NETWORK NODES ALLOWED AS SERVERS
IST1253I NONE
IST924I -----
IST1256I CURRENT NETWORK NODE SERVER
IST1253I NETA.SSCP1A SLUINIT = REQ ENBCAST = NO
IST924I -----
IST1677I PREFERRED NETWORK NODE SERVER
IST1253I NETA.SSCP1A SLUINIT = REQ ENBCAST = NO
IST314I END
```

Non-SYSPLEX Network Node Server for generic resources End Nodes - migration concerns



- The implementation is limited to the generic resource end nodes and requires no changes to either the primary or backup network node servers.
- The backup-mode of operation requires that all of the generic resources end nodes allow themselves to be searched for unknown resources. This will not be limited to searches for generic resource names; rather all network searches for any unknown resource will be sent to all of the generic resource end nodes while they are running in this backup-mode, which may result in a noticeable increase in CPU utilization on these end nodes.

HPDT packing



- High Performance Data Transfer (HPDT) Multi-Path Channel (MPC) was originally designed as a large data transport and in that configuration it performs well.
- However, many configurations use predominately small data units or a mixture of large and small data units, which results in less than optimal performance.
 - ┆ This is due to fragmentation of the data stream, causing a small amount of real data crossing the media, with a large amount of overhead required.
- Throughput of small outbound SNA or Enterprise Extender data packets across HPDT MPC media can be significantly improved by enabling HPDT packing.
- HPDT packing prevents fragmentation by packing small outbound data units into larger buffers to be transported.
- A new PACKING operand on the TRLE definition statement controls HPDT packing.
 - ┆ A new PACKING operand is provided on the TRLE definition statement to allow for control of HPDT packing
 - ┆ PACKING=OFF (default) sets the packing size limit to zero which, in effect, disables the function
 - ┆ PACKING=ON enables packing by setting a packing size limit of 2K.
 - ┆ PACKING=max_size enables packing by setting a packing size limit of max_size. The valid range for max_size is 1024 to 8192 inclusive.

HPDT Packing Notes



NOTES

- For point-to-point connections using the HPDT MPC protocol, throughput of small SNA or Enterprise Extender data packets can be significantly improved by enabling HPDT packing. This solution provides for better utilization of the HPDT MPC data stream by eliminating all of the alignment bytes transmitted in the HPDT data segment.
- HPDT packing is a compromise between an increase in storage and CPU consumption in order to increase throughput by improving channel utilization.
- Data is eligible to be packed when
 - f The packet has not already been packed by TCP
 - f Its total size does not exceed the packing size limit
 - f Each physically contiguous piece of data does not exceed 2K
- Testing using an Enterprise Extender workload shows benefits in a host-to-router or host-to-channel extender configuration. Testing in a host-to-host configuration showed limited benefit and therefore HPDT packing is not recommended in the host-to-host configuration.
- In the host-to-router or host-to-channel extender configuration, where bottlenecks may be the channel bandwidth or adjacent link station capacity, it is likely the benefit of the increase in throughput would exceed the cost of additional host storage and CPU. In other configurations, the cost of the packing buffers and/or CPU resource will exceed the throughput benefits.

HPDT Packing Notes (cont.)



**N
O
T
E
S**

- > It is recommended HPDT packing be disabled for the following conditions:
 - f* On systems that are storage or CPU constrained.
 - f* When the majority of the PDUs are ineligible (ex. TCP/IP native traffic), which will cause the packing buffers to be allocated but under-utilized.
 - f* For MPC groups containing multiple write devices. The storage cost could be excessive and therefore likely to outweigh the throughput benefit.

HPDT packing results



➤ Based on analysis of z/OS V1R5 CS MPC HPDT Packing sizes of 1K, 2K, and 4K relative to z/OS V1R4 CS with no packing. As always, results may vary due to network and application specifics.

➤ Results with primarily EE (or native SNA) interactive traffic with MPC HPDT packing:

- ┆ Significant throughput improvement of 25-27%.
- ┆ CPU usage relatively flat.

➤ The cost of HPDT packing outweighs any throughput benefit in the following workloads:

- ┆ Mixed interactive workload (small datagrams) and stream workload (large datagrams).
- ┆ Mixed MPC EE interactive workload and TCP non-EE interactive workload.

➤ **RECOMMENDATION:**

In some cases a customer will benefit if EE (or native SNA) interactive traffic is isolated to a specific MPC group, with TCP/IP and/or stream traffic going across a different MPC group (i.e. where the HPDT packing improvements on the SNA MPC group represents a savings that justifies the additional ESCON channels).

- ┆ EE (or native SNA) interactive traffic should deploy HPDT packing with PACKING=ON (default 2K) specified on the TRLE definition statement.
- ┆ For all other traffic, HPDT packing should be disabled.

XCA slowdown monitoring



- When an XCA device goes into slowdown at the subchannel level, there is no external indication.
- A subchannel slowdown can have the following results:
 - ⌘ All sessions using the XCA connection hang.
 - ⌘ If the XCA device remains in slowdown for long periods of time, VTAM can accumulate large amounts of outbound data, which can eventually lead to ECSA storage shortages.
- New slowdown monitoring and operator awareness is being provided for XCA subchannels
- When an XCA subchannel enters slowdown, a display of the XCA major node will now include an indication that the device is in slowdown.
- After a specified amount of time, if the XCA subchannel remains in slowdown the operator will be notified with a new message.
- The function will be shipped enabled with a default value of 180 seconds for detecting an extended XCA subchannel slowdown period.
- A second MAXSLOW subparameter can be specified on the XCA PORT definition statement to allow the user to change the slowdown detection time limit from the default.
- A new subparameter is available on the existing MAXSLOW parameter on the XCA PORT definition statement.
- The original format: MAXSLOW=slowdown_inop
 - ⌘ Where slowdown_inop specifies the number of seconds that VTAM will allow an XCA LINE to be in SLOWDOWN state before INOPing the link station.
- A second MAXSLOW parameter is now introduced with a new format: MAXSLOW=(slowdown_inop,slowdown_msg)
 - ⌘ Where slowdown_msg specifies the number of seconds that an XCA subchannel (PORT) remains in slowdown before the operator is notified.

© Copyright International Business Machines Corporation 2004. All rights reserved.

XCA slowdown monitoring Messages and displays



- A new message is issued on the DISPLAY of an XCA major node to indicate if the subchannel is in slowdown and the number of start I/O operations on the subchannel (from the time the node was last activated):
 - f IST1885I SIO = %%%5 SLOWDOWN = %3
- A new unsolicited message is issued to notify the operator if the XCA subchannel level slowdown condition exists for a period of time longer than the MAXSLOW value specified for the subchannel:
 - f IST1886I SLOWDOWN TIME EXCEEDS MAXSLOW = %3 SECONDS FOR DEVICE %4
 - f If the XCA device fails to exit slowdown, IST1886I will be repeated every three minutes to notify the operator of this serious problem.
- If the XCA device exits slowdown after exceeding the subchannel MAXSLOW time period, a new unsolicited message will be issued to notify the operator:
 - f IST1887I DEVICE %4 EXITED SLOWDOWN MODE

```
d net,id=lan1a1,e
IST097I DISPLAY ACCEPTED
IST075I NAME = LAN1A1, TYPE = XCA MAJOR NODE
IST486I STATUS= ACTIV, DESIRED STATE= ACTIV
IST1021I MEDIUM=BUS,ADAPNO= 1,CUA=0500,SNA SAP= 4
IST1885I SIO = 1234 SLOWDOWN = YES
IST654I I/O TRACE = OFF, BUFFER TRACE = OFF
IST1656I VTAMTOPO = REPORT, NODE REPORTED - YES
IST170I LINES:
IST232I L1A111 NEVAC
IST232I L1A112 NEVAC
IST232I L1A113 NEVAC
IST314I END
```

Additional SNA enhancements



➤ Serviceability/Problem Determination

- ƒ Usability improvements to CSDUMP command (dump for sense code or message)
 - Display settings of CSDUMP command and allow for the deletion of triggers
- ƒ APPN Trace enhancement
 - New SUBTRACE option TGVC will provide TG Vectors in appropriate trace records
- ƒ Track CSM Buffers internally through components
- ƒ Enable dumping for VTAM Inoperative conditions with new Modify INOPCode command
 - Works in conjunction with existing Modify INOPDump command

➤ New/Enhanced Functionality

- ƒ SWNORDER and DLRODER
 - Enhanced to allow greater control over PU selection during connection processing. On start option or in XCA or NCP major nodes.

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	MQSeries	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.

© Copyright International Business Machines Corporation 2005. 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.