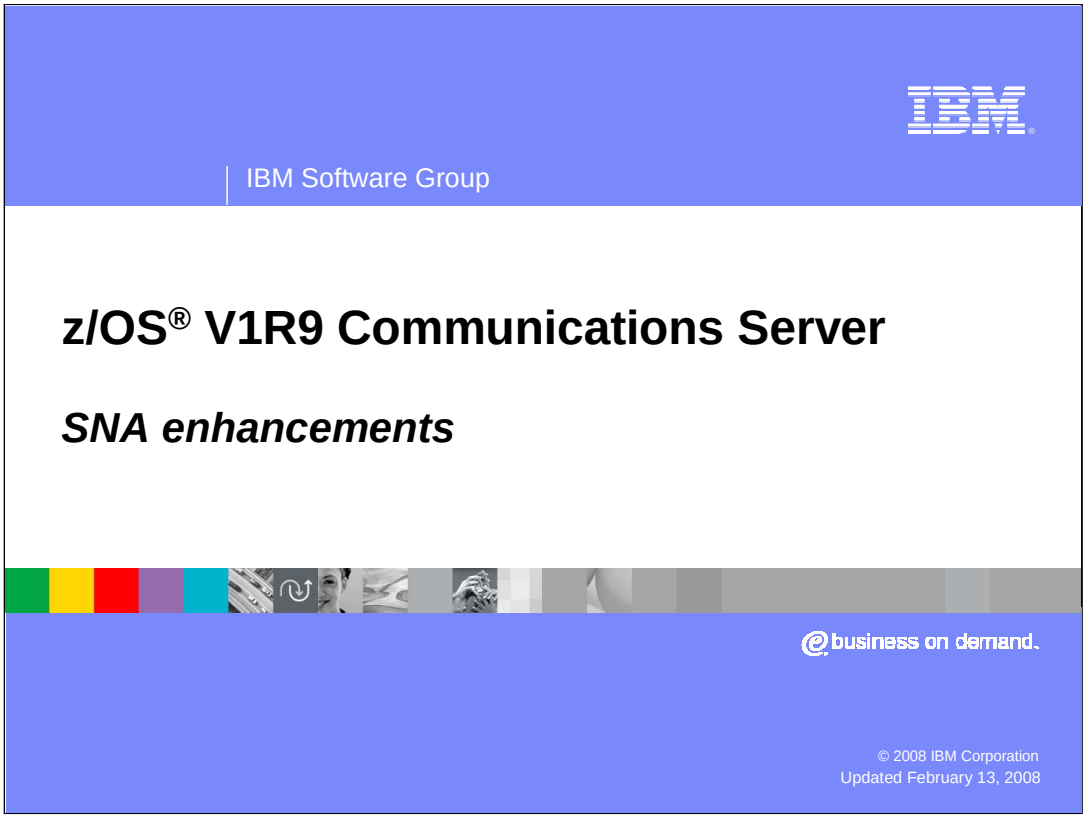

This presentation describes other SNA enhancements for z/OS V1R9 Communications Server.

Agenda

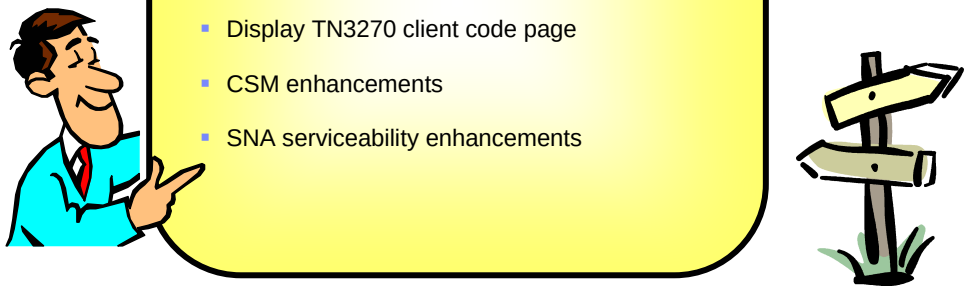

- Adjacent cluster table enhancements
- Increase maximum CAPACITY value
- Improve performance of SNA session encryption
- Display TN3270 client code page
- CSM enhancements
- SNA serviceability enhancements

z/OS V1R9 Communications Server contains multiple SNA enhancements. Each of these enhancements is discussed in this presentation.

Section

Adjacent cluster table enhancements

This section describes enhancements to the Adjacent Cluster table definitions.

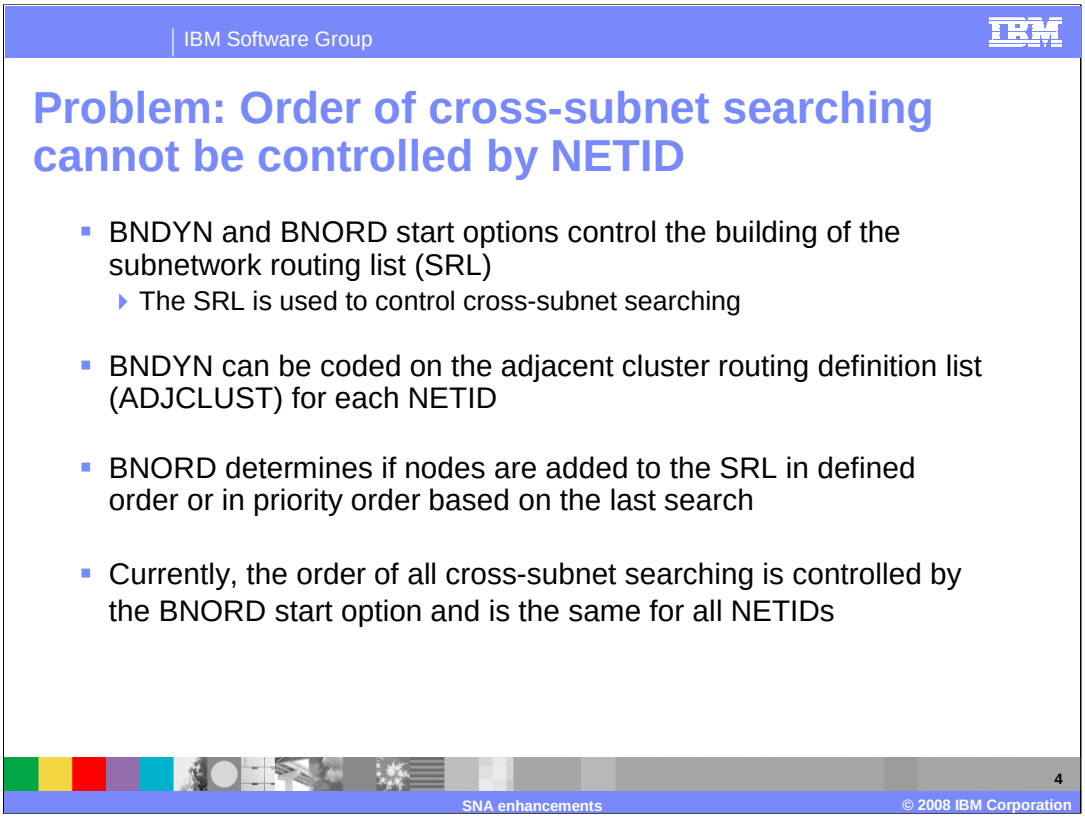

The BNDYN and BNORD start options control how the subnetwork routing list (SRL) is built. The SRL is used by a border node to control cross-subnet searching and is built for each search. The adjacent cluster routing definitions allow you to customize the building of the SRL. The BNDYN start option controls the amount of dynamics used in building the SRL and can be coded on the adjacent cluster routing definition list (ADJCLUST) to customize routing between subnetworks for each NETID. The BNORD start option determines if nodes are added to the SRL in defined order from the adjacent cluster definitions or in priority order based on the last search for a NETID.
Currently, the order of all cross-subnet searching is controlled by the BNORD start option and is the same for all NETIDs. This does not allow you to control the order of cross-subnet searching by NETID.

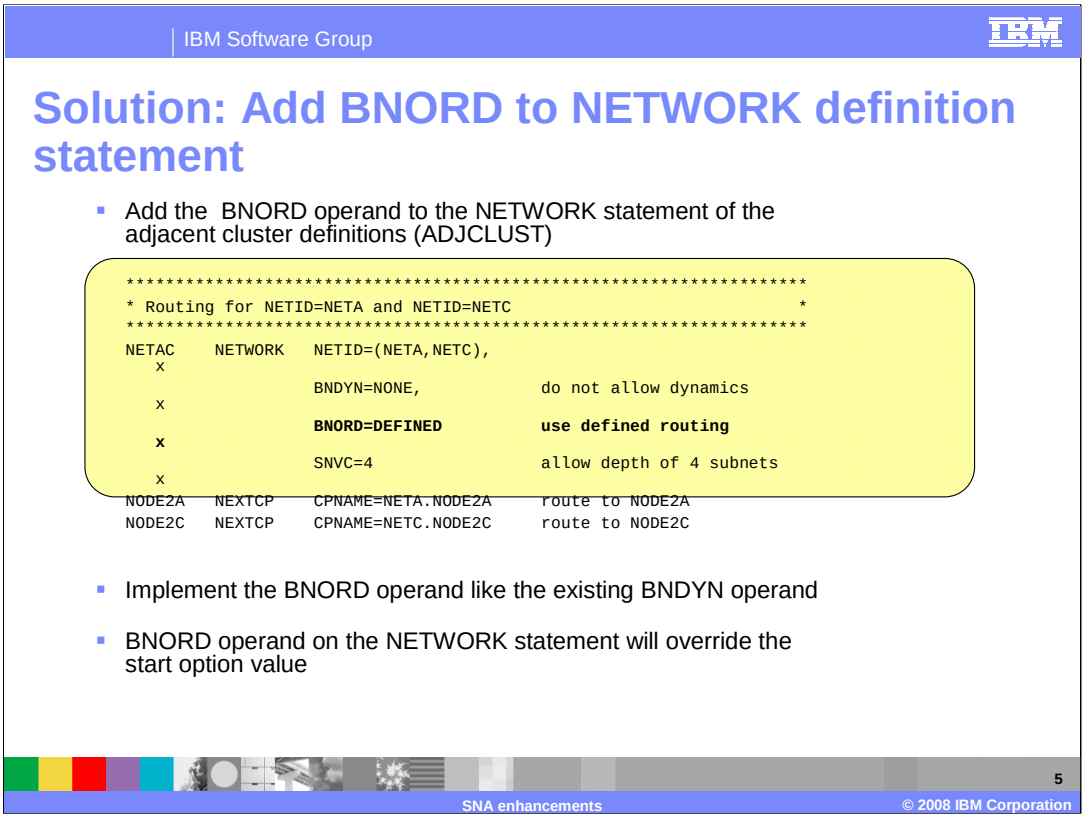

In the z/OS V1R9 Communications Server, the BNORD operand has been added to the NETWORK statement of the adjacent cluster definitions. The BNORD operand is implemented like the existing BNDYN operand on the NETWORK statement. If the BNORD operand is coded on the NETWORK statement, it will override the start option value when building a subnet routing list. Adding the BNORD operand to the NETWORK statement allows you to specify the order of cross-subnet searching for each NETID coded in the adjacent cluster routing definitions.

The values for BNORD are PRIORITY and DEFINED with PRIORITY being the default start option value. Priority routing indicates that preference is given to nodes for which the most recent search was successful. Defined routing indicates that searches are done in the order specified in the adjacent cluster definition list.

Display command example

- **New messages IST2207I, IST2208I, and IST2209I to display ADJCLUST Table values.**
- **IST2207I replaces IST1325I.**

```
d net,adjclust,netid=neta
     IST097I DISPLAY ACCEPTED
     IST350I DISPLAY TYPE = ADJACENT CLUSTER TABLE
     IST2207I DEFINED TABLE FOR NETA
     IST2208I BNDYN = LIMITED FROM START OPTION
     IST2209I BNORD = DEFINED FROM ADJCLUST TABLE
     IST1326I CP NAME           TYPE    STATE      STATUS
SNVC
     IST1327I NETA.BN3          DEFINED ACTIVE     FOUND         003
     IST1327I NETA.BN2          DEFINED NOT ACTIVE NOT SEARCHED 003
     IST1327I NETA.BN1          DYNAMIC ACTIVE     NOT SEARCHED N/A
     IST314I END
```


New messages IST2208I and IST2209I are added to the DISPLAY ADJCLUST to display the BNORD and BNDYN values. The new messages are modeled after the IST1704I and IST1705I messages for SORDER and SSCPORD for the ADJSSCP TABLE. Either IST2208I or IST2209I is issued to display the value of both the BNDYN and BNORD search control options. IST2208I is issued when the value is obtained from the START OPTION. IST2209I is issued when the value is obtained from the adjacent cluster definition table. Message IST2207I replaces message IST1325I because the border node dynamics information was expanded and moved to IST2208I and IST2209I.
New messages IST2208I and IST2209I are also issued when detailed locate search failure information is displayed to aid in problem diagnosis. The new messages will indicate the values and origin of BNDYN and BNORD that were used for the search. Detailed locate search information can be displayed by setting the LSIRFMSG and FSIRFMSG START OPTIONS.

Problem: Unable to easily restrict searches to nodes

- NEXTCP statement on Adjacent cluster definitions specifies nodes to be searched in cross-subnetwork searches
- BNDYN=FULL includes all border nodes in subnet routing list for cross-subnetwork searching
- BNDYN=NONE includes only nodes defined by CPNAME on NEXTCP statement will be included in the subnet routing list
- Adjacent cluster definitions do not provide a method to selectively restrict searches to nodes during cross-subnetwork searches
- Currently you can only restrict cross-subnetwork searches by coding BNDYN=NONE and then only listing the border nodes that are to be searched
- In networks with a large number of border nodes, users want to code BNDYN=FULL, then only list the small number of border nodes that are not to be searched

The NEXTCP statement on the adjacent cluster routing definitions specifies nodes to be searched during cross-subnet searching. Border node dynamics determines if additional nodes are added to the subnetwork routing list for searching. BNDYN=FULL specifies that all nodes valid for cross-subnet searching are included in the subnetwork routing list for cross-subnetwork searching. BNDYN=NONE specifies that only nodes defined on the NEXTCP statement will be included in the subnetwork routing list for cross-subnetwork searching.

Adjacent cluster definitions do not provide a method to selectively restrict searches to individual nodes during cross-subnetwork searching. Currently you can only restrict cross-subnetwork searches by coding BNDYN=NONE and then only listing the border nodes that are to be searched. In networks with a large number of border nodes users want to code BNDYN=FULL, so that all possible border nodes are included for cross-subnet searching, then only list the small number of border nodes that are not to be searched. Another use for this option would be during planned outages to easily restrict searches to individual border nodes where a desirable path is not available.

Solution: Add OMITCP to the NEXTCP statement

- Add OMITCP operand to the NEXTCP statement of the adjacent cluster routing definitions

```
*********************************************************************
* Routing for NETID=NETZ with OMITTED nodes                         *
*********************************************************************
NETZ     NETWORK   NETID=(NETZ),                                    x
                   BNDYN=FULL,            allow full dynamics       x
                   BNORD=PRIORITY,        use priority routing      x
                   SNVC=4                 allow depth of 4 subnets
NODE2A   NEXTCP    CPNAME=NETA.NODE2A,                               x
                   OMITCP=YES             do not route to NODE2A
NODE2C   NEXTCP    CPNAME=NETC.NODE2C,                               x
                   OMITCP=YES             do not route to NODE2C
```

- OMITCP=YES prevents the CPNAME from being included in a subnet routing list for searching

The solution is to add a new OMITCP operand to the NEXTCP statement of the adjacent cluster routing definitions. The OMITCP operand will have a value of YES or NO with NO being the default. If OMITCP=YES is coded, the node specified on the CPNAME operand would not be included in a subnet routing list built for the NETID specified on NETWORK statement. The node would not be added as a defined or dynamic entry to the SRL. This will allow you to selectively restrict searches to a specific node. The operand will work with all levels of dynamics. With BNDYN=NONE it is the same as not adding the node to the list of NEXTCPs.

The OMITCP operand can be used during planned outages to restrict searches to border nodes where a desirable path is not available. The border node selection function of the DSME allows more control cross-subnetwork searching. The DSME can customize cross-subnet searching based on search information other than the NETID, such as the OLU or DLU names.

Display command example: OMITCP

- New STATE added to the IST1327I message to display the OMITCP state value

```
d net,adjclust,netid=neta
IST097I DISPLAY ACCEPTED
IST350I DISPLAY TYPE = ADJACENT CLUSTER TABLE
IST2207I DEFINED TABLE FOR NETA
IST2208I BNDYN = FULL FROM START OPTION
IST2208I BNORD = PRIORITY FROM START OPTION
IST1326I CP NAME           TYPE    STATE       STATUS       SNVC
IST1327I NETA.SSCP2A       DEFINED OMITTED     NOT SEARCHED 003
IST1327I NETA.SSCP1A       DEFINED ACTIVE      NOT SEARCHED 003
IST1327I NETA.SSCPAA       DEFINED NOT ACTIVE  NOT SEARCHED 003
IST1327I NETB.SSCPBA       DYNAMIC ACTIVE      *** N/A ***  N/A
IST314I END
```


To display the new OMITCP operand value, a new state variable was added for message IST1327I. If OMITCP=YES is coded the state of OMITTED is displayed. If OMITCP=YES is not coded the existing states of ACTIVE or NOT ACTIVE are displayed. This display was done from NETB.SSCPBA. SSCPBA was not defined in the adjacent cluster definitions but was added because border node dynamics is set to full.

Section

Increase maximum CAPACITY value

The maximum limit of 1000M for CAPACITY has been increased in z/OS V1R9 Communications Server.

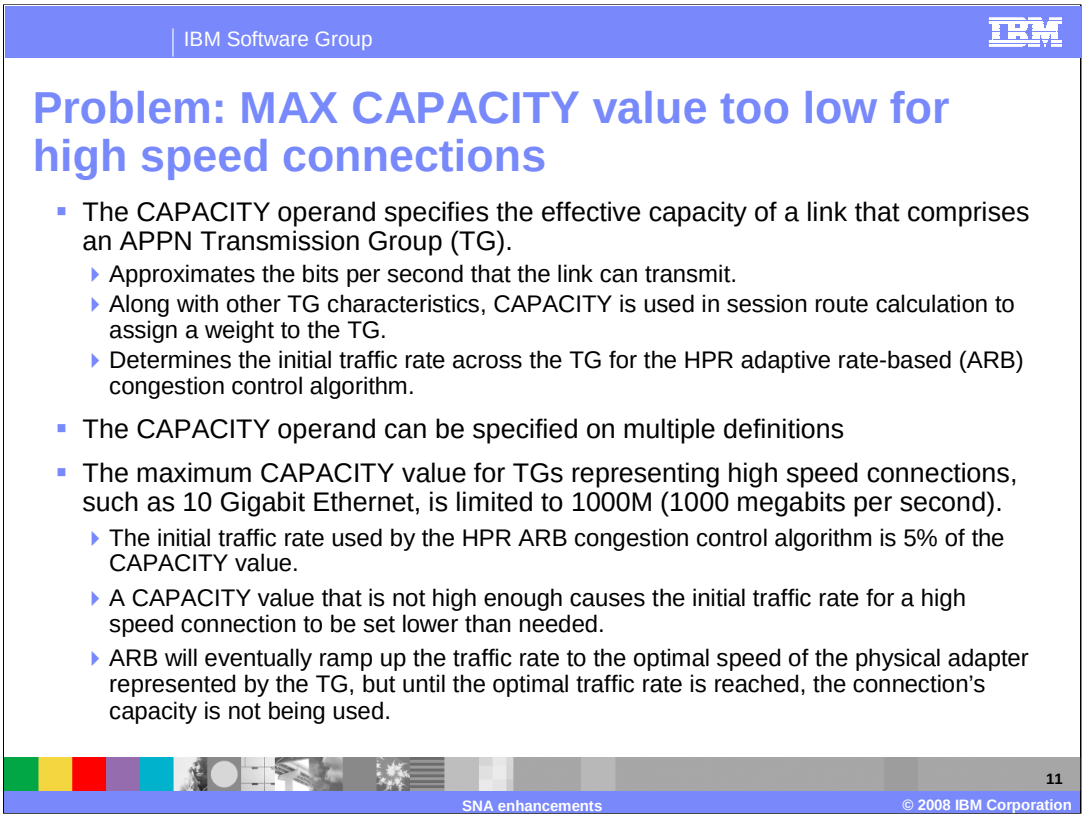

The CAPACITY operand can be specified in VTAM® major nodes where the definition statement defines an APPN TG, and on APPN TG Profiles and APPN CoS definitions. The value specifies the effective capacity of a link, approximating the bits per second that the link can transmit. Along with other TG characteristics assigned to the TG, CAPACITY is used in session route calculation to assign a weight to the TG to determine the optimal route through an APPN network. CAPACITY is also used to determine the initial traffic rate across the TG for the HPR ARB congestion control algorithm.

The maximum limit of 1000M for the CAPACITY value of a TG representing a high speed connection is not high enough to set the optimal initial traffic rate for that connection.

The CAPACITY operand can be specified on the Cross Domain Resource Manager (CDRM) major node, the External Communications Adapter (XCA) major node, the Local SNA major node, the Model major node, the Network Control Program (NCP) major node, the Switched major node, the APPN TG Profile (TGP) definitions, and the APPN Class of Service (CoS) definitions.

The initial data rate across a TG for the HPR ARB congestion control algorithm is 5% of the CAPACITY value. This is the rate at which date is initially sent across the physical adapter represented by the TG. If the physical adapter can handle larger amounts of data, ARB ramps up the value gradually until it reaches the optimal traffic rate for the adapter. However, if the data rate is initially too low, the connection's capacity is not being used until ARB increases the data rate to the connection's optimal speed. Some physical connection types that can currently benefit from a higher CAPACITY are 10 Gigabit Ethernet, FICON® Express, and Hipersockets™. These connection types are only supported by Enterprise Extender.

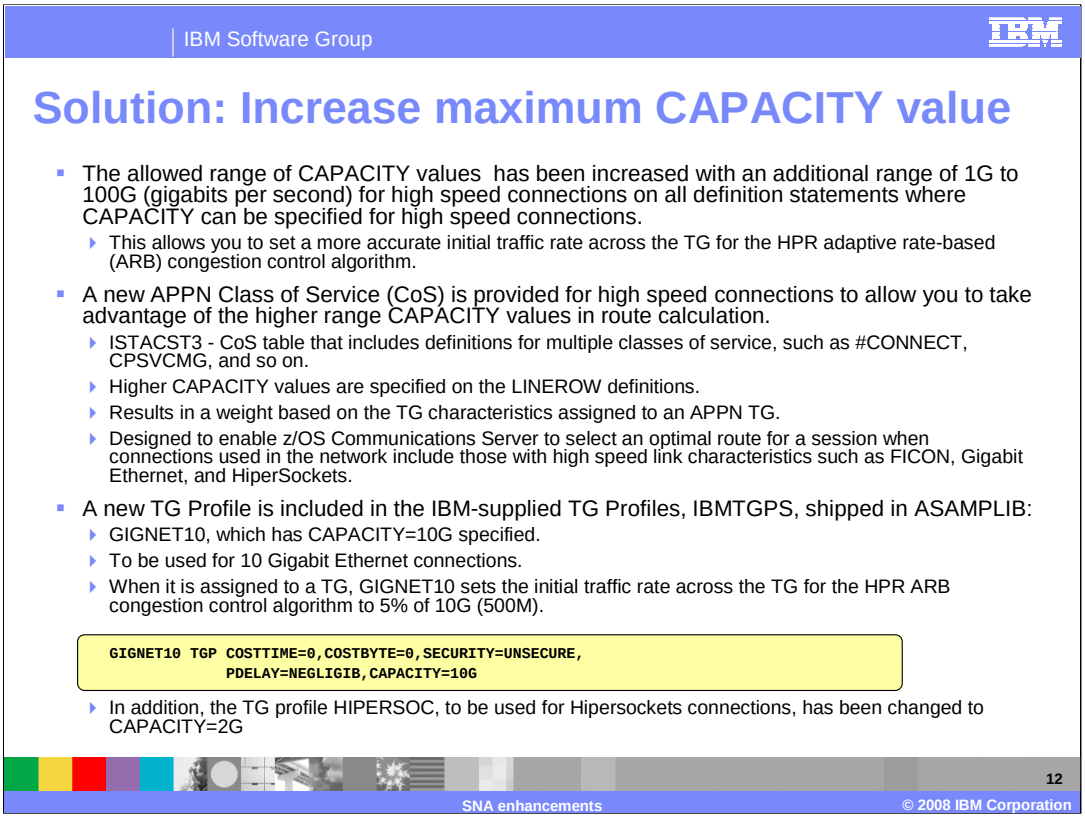

An additional range of 1G to 100G is now allowed on the CAPACITY operand on all definition statements where CAPACITY can be specified. This range allows you to specify a higher CAPACITY value for a TG than was previously available. The primary advantage of the higher CAPACITY value is to set a more accurate initial traffic rate across a high speed connection for the HPR adaptive rate-based (ARB) congestion control algorithm. ARB increases the traffic rate from the initial rate set by the CAPACITY value, so a higher initial traffic rate allows the algorithm to ramp up to an optimal traffic rate faster.

Also, for session route calculation, a new APPN Class of Service is provided to take advantage of the higher range CAPACITY values for high speed connections. ISTACST3, a new set of 12-row APPN CoS definitions is shipped in z/OS V1R9 Communications Server in ASAMPLIB. This is a table that includes definitions for multiple classes of service, such as #CONNECT, CPSVCMG, and so on. Unlike the other IBM-supplied Cos tables, COSAPPN and ISTACST2, these definitions use the new higher CAPACITY values on the LINEROW statements. These CoS definitions are designed to enable z/OS Communications Server to select an optimal route for a session when connections used in the network include those with high speed link characteristics. Some of these high speed connections are FICON, Gigabit Ethernet, and HiperSockets.

To use ISTACST3, you must copy the CoS definitions into SYS1.VTAMLST and then activate the member in which the definitions reside. You can have only one set of CoS definitions active at any time. COSAPPN is automatically activated when z/OS Communications Server is initialized. If you choose to use ISTACST3 you can activate it in one of several ways:

You can use the VARY ACT command to activate it. You can place the ISTACST3 member in the configuration list to automatically activate it at z/OS Communications Server initialization. You can rename the ISTACST3 member to COSAPPN and rename COSAPPN to something else. It will then be automatically activated at z/OS Communications Server initialization.

IBM-supplied TG Profiles are shipped in member IBMTGPS in ASAMPLIB. A new TGP for 10 Gigabit Ethernet connections, GIGNET10 is now included in IBMTGPS. This TGP sets the initial traffic rate across the TG for the HPR adaptive rate-based (ARB) congestion control algorithm to 5% of CAPACITY. GIGNET10 specifies CAPACITY=10G and will result in an initial data rate of 500M. In addition, the HIPERSOC TGP has been changed to CAPACITY=2G.

The IBMTGPS TG profiles are automatically activated at z/OS Communications Server initialization. All that is needed to activate them is to copy the definitions from ASAMPLIB into a member in VTAMLST. After IBMTGPS is activated, you can then assign the group of TG characteristics defined in a specific TG profile (for example GIGNET10) to a TG using one of these methods:

1.Specify the TGP=GIGNET10 operand on the PU definition statement

2.Assign the TGP to an already existing APPN TG with the MODIFY TGP,TGPNAME=GIGNET10,ID=adjacent_node,TGN=tg_number command.

Displaying CAPACITY and TG Weight

- You can display the CAPACITY and the weight assigned to a TG using a specific CoS with this command:

```
D NET,TOPO,ORIG=SSCP1A,DEST=SSCP2A,APPNCOS=#CONNECT

IST350I DISPLAY TYPE = TOPOLOGY
IST1299I TRANSMISSION GROUPS ORIGINATING AT CP NETA.SSCP1A
IST1357I                                                  CPCP
IST1300I DESTINATION CP    TGN      STATUS   TGTYPE      VALUE WEIGHT
IST1301I NETA.SSCP2A       21       OPER     INTERM      YES   30
IST1579I                   ------------------------------------------
IST1163I                   RSN              HPR         TIME LEFT
IST1164I                   58               YES         15
IST1579I                   ------------------------------------------
IST1302I                   CAPACITY PDELAY   COSTTIME    COSTBYTE
IST1303I                   23G      NEGLIGIB 0           0
IST1579I                   ------------------------------------------
IST1304I                   SECURITY UPARM1   UPARM2      UPARM3
IST1305I                   UNSECURE 128      128         128
IST1579I                   ------------------------------------------
IST1736I                   PU NAME
IST1737I                   AHHCPU1
```


To display the weight assigned to a TG using a specific set of CoS definitions, use the DISPLAY TOPO command with the APPNCOS operand.

Because the CAPACITY value is a one byte floating-point number, the value coded can results in a different value being displayed with the DISPLAY TGP or the DISPLAY TOPO commands. The z/OS Communications Server Resource Definition Reference contains a table, under the CAPACITY operand description for APPN transmission group profiles, that maps the value coded to the value that will be displayed with the DISPLAY TGP or the DISPLAY TOPO commands.

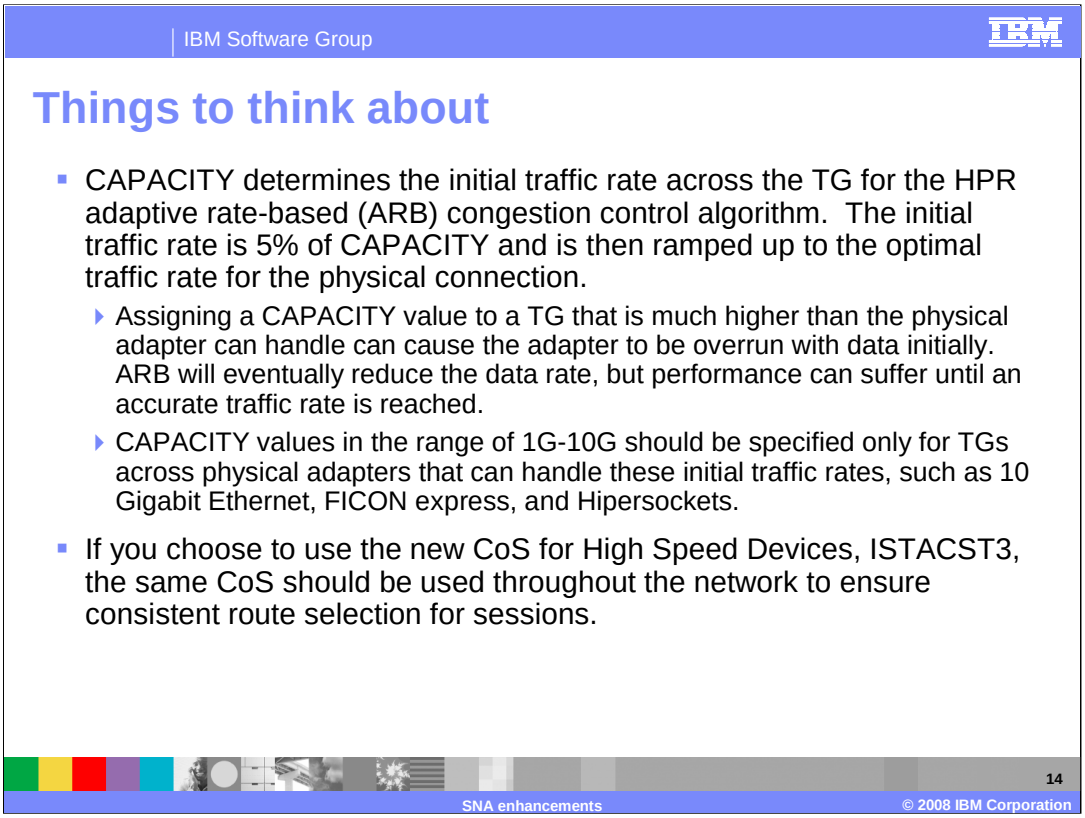

The initial data rate across a TG for the HPR ARB congestion control algorithm is 5% of the CAPACITY value. This is the rate at which date is initially sent across the physical adapter represented by the TG. If the physical adapter can handle larger amounts of data, ARB ramps up the value gradually until it reaches the optimal traffic rate for the adapter. Therefore, CAPACITY values in the range of 1G-10G should be specified only for TGs across physical adapters that can handle the initial data rate. Assigning a CAPACITY value to a TG that is much higher than the physical adapter can handle can cause the adapter to be overrun. The data rate will eventually be reduced by ARB, but assigning an initial traffic rate that is too high can cause performance to suffer until an accurate traffic rate is reached. Some physical connection types that can currently benefit from a CAPACITY in the range of 1G-10G are 10 Gigabit Ethernet, FICON Express, and Hipersockets. These connection types are only supported by Enterprise Extender.

When any Class of Service (CoS) is used in a network, whether COSAPPN, ISTACST2, or ISTACST3, it is important to use the same CoS on all network nodes in the network. If you do not use the same CoS, different session routes can be selected, depending on the TG characteristics specified in the different Classes of Service. This is true even across network boundaries when you are using border node configurations.

Section

Improve performance of SNA session encryption

“Improve performance of SNA session encryption” is a functional enhancement introduced in z/OS V1R9 Communications Server.

Problem: Session encryption impacting performance

- SNA session level encryption
  - z/OS Communications Server attempts to interface with an external cryptographic facility for each session encryption request.
  - A subtask is created for the each encryption request to allow other processing to continue while the session waits for encryption to complete.
  - The number of subtasks that can run concurrently is controlled by the DLRTCB and MAXHNRES start options.
- The creation and termination of a subtask for each session encryption request can impact performance when many sessions require encryption.
  - The default start option value for session encryption is ENCRYPTN=YES
- Other functions that are required to run under a subtask must wait for session encryption to complete if many session requests are queued. For example:
  - NCP dump, load, or restart
  - Enterprise Extender HOSTNAME resolution
  - Messages with reply requested

SNA session level encryption requires a subtask be created for each session encryption request. This results in an attempt to interface with an external cryptographic facility.

The number of subtasks (TCB structures) allowed concurrently is controlled by the DLRTCB and MAXHNRES start options. These structures use below the line storage, so limits are necessary. For example, if DLRTCB=32 and MAXHNRES=20 are specified, then a total of 52 subtasks can be attached at one time. However, once requests are received by ISTINCDP, the subtasks attached for HOSTNAME resolution are not limited to the value specified for MAXHNRES, nor are the subtasks attached for other functions limited to the value specified for DLRTCB. MAXHNRES does limit the number of DISPLAY EE and DISPLAY EEDIAG commands requiring HOSTNAME resolution that VTAM will accept at once, however. This restriction is policed in the NOS component and is meant to prevent a user CLIST from overwhelming ISTINCDP.

There are two symptoms of this problem that are commonly seen:

1. Thousands of sessions appear to be hung because they are queued up waiting for subtask resources needed for encryption requests. This often happens because the default for the ENCRYPTN start option is YES, causing ALL session requests to be sent to ISTINCDP even when encryption is not needed.

2. Because ISTINCDP processes requests in a first in first out order, functions other than encryption can't proceed if many encryption requests were queued first. There is a limit to the number of subtasks that can be attached concurrently. The functions, such as NCP loads, appear to be hung.

The overhead for creating and subsequently deleting the control block structure for each subtask can impact performance, causing many sessions to wait pending encryption. This impact can be severe, especially when the ENCRYPTN start option is allowed to default and ENCR=OPT is allowed to default on all APPL definitions. The default start option value of ENCRYPTN=YES causes ALL sessions with an application to request encryption processing, unless ENCR=NONE is specified on the APPL definition statement for that application.

Since there are a limited number of subtasks allowed to run concurrently (that number is the total of the values specified for the DLRTCB and MAXHNRES start options), many session encryption requests can cause other functions to wait until resources are available. This is because all requests for subtask creation are processed in the order that they are received.

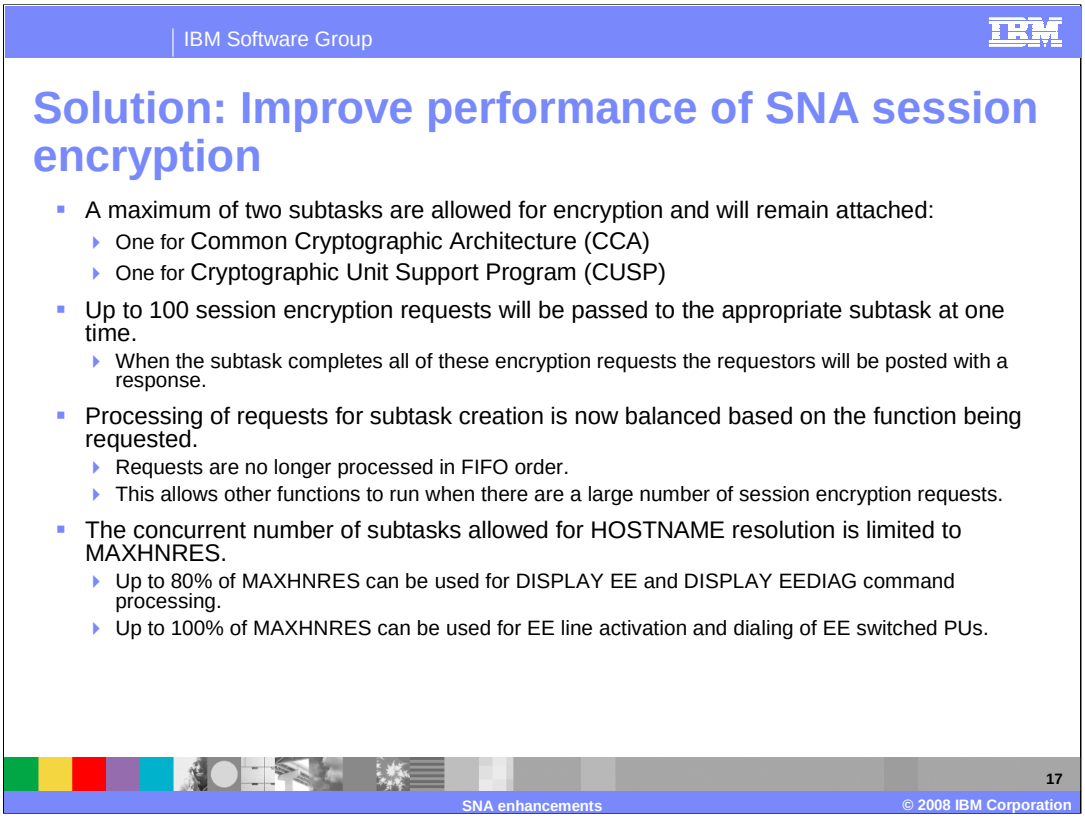

Instead of attaching and detaching a subtask for each session encryption request (represented by a DLRPL control block), one subtask is created when the first CCA request is received and one subtask will be created when the first CUSP request is received. Each of these subtasks will remain attached and waiting for work until VTAM termination. The subtask will interface with the appropriate cryptographic facility for each request. When all of the requests passed to the subtask (up to 100 at a time) have been completed, the DLRPL control blocks will be returned to the requestors so that session establishment can continue.

This should improve the performance of session level encryption because it eliminates the overhead of the ATTACH/DETACH processing for each encryption request and substitutes WAIT/POST processing, which has far less performance impact. Since only up to two subtasks are dedicated to encryption, an additional benefit is that it frees system resources for other functions supported by ISTINCDP.

Rather than creating subtasks in the order in which the request are received, requests for each functional area are load balanced to allow all functions a chance to used the limited resources needed to attach a subtask. This prevents these functions from having to wait for a large number of session encryption requests to be completed.

In addition, the concurrent number of subtasks allowed for HOSTNAME resolution is limited to the value specified for the MAXHNRES start option.

There are four functions that can result in HOSTNAME resolution:

• EE line activation with the HOSTNAME operand on the GROUP

• Dialing an EE switched PU with HOSTNAME operand

• DISPLAY EE command with HOSTNAME operand

• DISPLAY EEDIAG command with HOSTNAME operand.

Before this, HOSTNAME resolution could use all subtask resources, including those reserved by the DLRTCB start option. MAXHNRES subtask resources concurrently being used for DISPLAY EE and DISPLAY EEDIAG command processing is further limited to 80% of MAXHNRES. This will allow EE line activation with HOSTNAME resolution to proceed even if a large number of the DISPLAY commands have been entered. However, Subtask requests for functions other than HOSTNAME resolution will be allowed to use more resources than those reserved by the DLRTCB start option if all of the resources reserved by the MAXHNRES start option are not in use.

Performance results

- V1R9 shows 20.6% lower CPU cost for VTAM encryption processing compared to V1R8 with 1000 VTAM session requests.

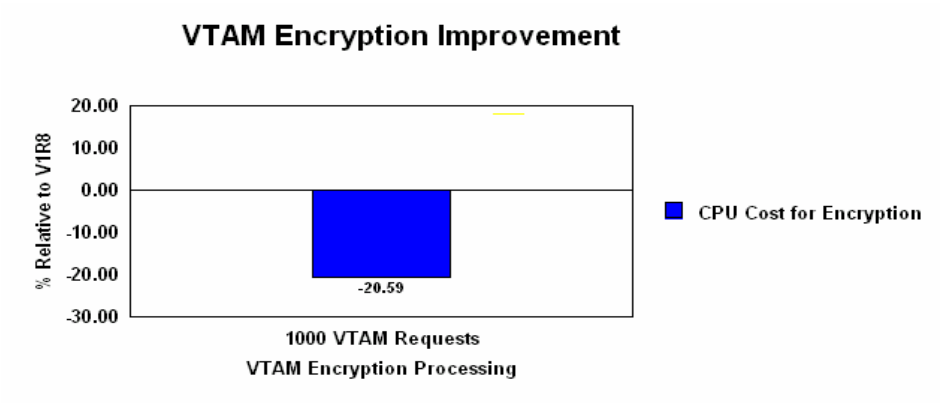

In performance testing, CPU cost was lowered by 20.6% for VTAM encryption processing with 1000 session requests from z/OS Communications Server V1R8 to V1R9.

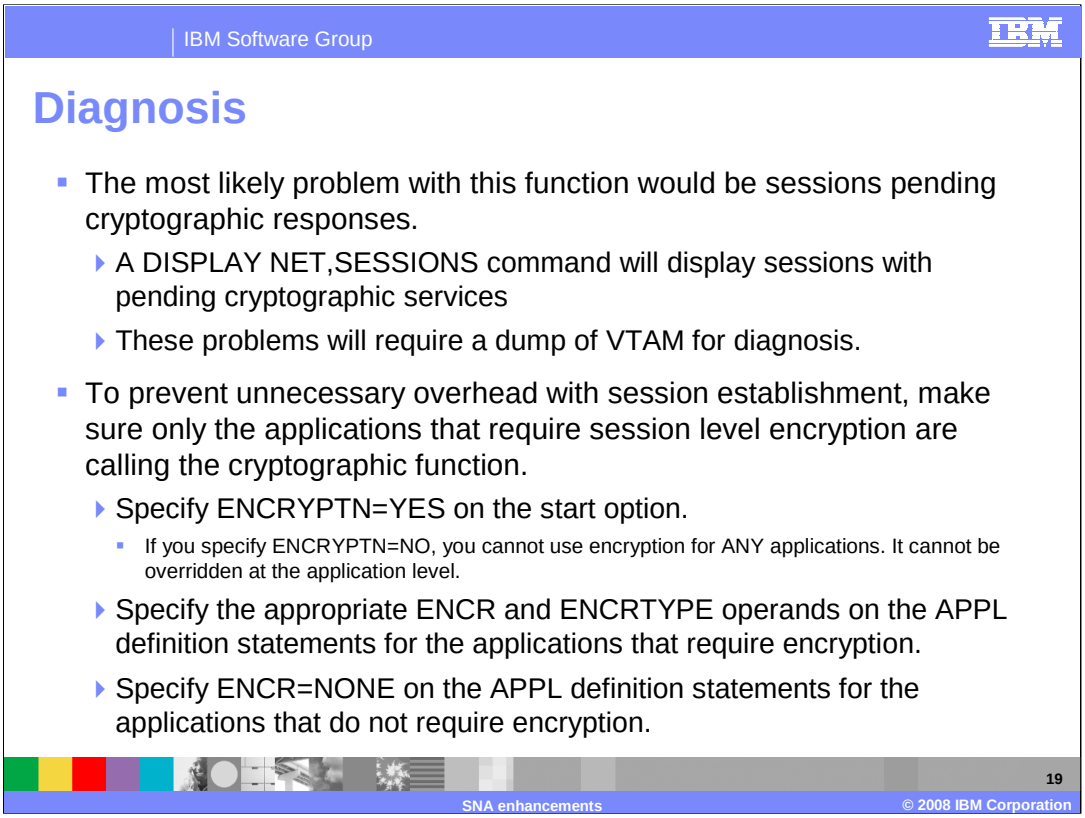

A dump of VTAM should be taken if one or more sessions are pending cryptographic responses for a period of time. You can determine that sessions are waiting for encryption with the DISPLAY SESSIONS command. There are several pending cryptographic session initiation states. See z/OS Communications Server IP and SNA Codes for a description of these states.

Often only a few applications require session encryption. If the ENCRYPTN start option and the ENCR operand on all the APPL definition statements are allowed to default, requests for encryption will be made for all sessions, even when it is not necessary. If none of your applications require session level encryption, specify ENCRYPTN=NO on the start option. However, this cannot be overridden at the application level. If only a few applications require encryption, do the following:

•Specify or allow the ENCRYPTN start option to default to YES. If you specify ENCRYPTN=NO on the start option, you cannot use session level encryption for any applications.

•Specify the appropriate ENCR and ENCRTYPE operands on the APPL definition statements for only those applications that require session level encryption.

•Specify ENCR=NONE on the APPL definition statements for all applications that do not require encryption.

Section

Display TN3270 client code page

This section describes the enhancement to display TN3270 Client Code Page information.

Problem: Application / code page compatibility

- Character Set and Code Page
  - Character Set and Code Page combination is commonly referred to as a Coded Graphic Character Set Global Identifier (CGCSGID)
  - Optionally is set by the terminal or emulator for use in a TSO session
- Some applications experience data corruption when an inappropriate Character Set and Code Page combination is used for a TSO session

A Character Set and Code Page combination is commonly referred to as a Coded Graphic Character Set Global Identifier (CGCSGID) . The CGCSGID values are set by the terminal or emulator and used for a TSO session. However, not all terminals or emulators include the CGCSGID information. Consult the documentation of the applicable terminal or emulator to see if CGCSGID information is supported.

The Character Set and Code Page combination in use for a TSO session may be inappropriate for some applications and cause data corruption.
This has been identified as a problem in some user DB2 environments where a TSO user with an incompatible client code page performs some processing that updates a DB2 database. As a result of incompatible code page, the data appears to be corrupted when stored back into DB2.

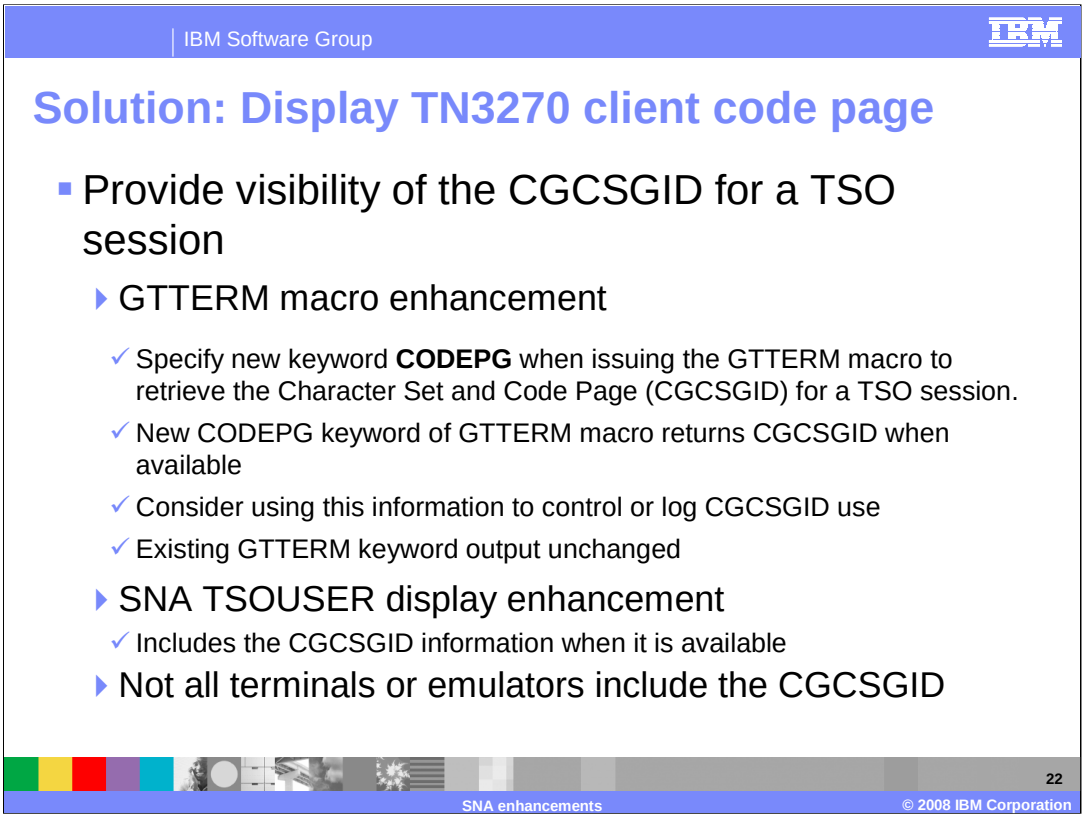

This new function provides visibility of the CGCSGID for a TSO session. Users may want to use this information as a criteria to permit or deny access to an application through use of a logon exit.

TSO/VTAM supports a GTTERM macro that the user can use to acquire information about a terminal. A new keyword, CODEPG, has been added to the GTTERM processing to allow the user to retrieve the CGCSGID information for a TSO session.

When the GTTERM macro is issued with the CODEPG keyword, the following information is returned to the issuer: Terminal name, Network ID, IP address, Port number, Character Set, if available, and Code Page, if available. Note that the Character Set and Code Page (CGCSGID) information is only available when the terminal or emulator includes it when the session is established. Consider updating your logon exit to log using a message what Code Page a client is using or enforce a certain set of Code Pages that users can use for a specific application. The existing GTTERM keyword output is unchanged. See the TSO/E Programming Services publication for information on the GTTERM macro.

The SNA TSOUSER display has been enhanced to report the CGCSGID in use for a TSO user. The TSOUSER Display will show Code Page and Character Set information when it is available.

Not all terminals or emulators include the CGCSGID information. Consult the applicable terminal or emulator documentation to see if CGCSGID information is supported.

Display command example

- **Display TSOUSER will also include the CGCSGID information when it is available.**

```
D net,tsouser,id=user1
IST097I DISPLAY ACCEPTED
IST075I NAME = USER1, TYPE = TSO USERID 949
IST486I STATUS= ACTIV, DESIRED STATE= N/A
IST576I TSO TRACE = OFF
IST262I ACBNAME = TSO0001, STATUS = ACT/S
IST262I LUNAME = TCPM1011, STATUS = ACT/S----Y
IST1727I DNS NAME: VIC127.TCP.RALEIGH.IBM.COM
IST1669I IPADDR..PORT 9.67.113.83..1027
IST2203I CHARACTER SET 0065  CODE PAGE 0025
IST314I END
```


The output of the Display TSOUSER command will also include the CGCSGID information when it is available.

Section

CSM enhancements

This section describes the enhancement made to CSM for z/OS V1R9 Communications Server.

Problem: CSM needs informative messages

- The Communications Storage Manager (CSM) adjusts the specified maximum ECSA value when it exceeds 90% of the ECSA available on the z/OS system, but no message is issued indicating that the maximum ECSA was changed.
- CSM sets the constrained level indicator when ECSA or FIXED storage reaches the constrained level, but no message is issued indicating that CSM ECSA or FIXED storage reached the constrained level.

CSM needs to issue a message when it adjusts MAX ECSA value. It will clarify the new value to the user.

CSM needs to issue a message when it sets the constrained level indicator for CSM ECSA or fixed storage. This will allow the operator to take some actions to relieve the situation.

Solution: CSM message enhancements

- CSM is enhanced to issue message IVT5590I when the requested maximum ECSA value has been adjusted to 90% of the ECSA on the z/OS system.
- CSM is enhanced to issue a message when ECSA or FIXED storage is constrained.
  - Message IVT5591I is issued when ECSA storage usage is above 80% of the MAX ECSA value and approaching 85% of the MAX ECSA value.
  - Message IVT5592I is issued when FIXED storage usage is above 80% of the MAX FIXED value and is approaching 85% of the MAX FIXED value.
- CSM is also enhanced to activate the Dynamic CSM Monitor function when the current ECSA storage usage reaches 80% or higher of the MAX ECSA value or the current fixed storage usage reaches 80% or higher of the MAX FIXED value.
- CSM will also inactivate the Dynamic CSM Monitor function when the current ECSA storage usage goes below 75% of the MAX ECSA value and the current fixed storage usage goes below 75% of the MAX FIXED value.
- CSM sets the ECSA and FIXED storage constrained indicator sooner in z/OS V1R9 Communication Server than the earlier releases of z/OS Communication Server.
  - Increase the values of MAX ECSA and MAX FIXED by 5%.

Message IVT5590I can be issued during: CSM initialization when the ECSA MAX value specified on the IVTPRM00 parmlib member is larger than 90% of the ECSA on the system. DISPLAY CSM command processing when the maximum ECSA value in effect has been adjusted by CSM. And MODIFY CSM command processing when the maximum ECSA requested is larger than 90% of the ECSA on the system.
CSM changed the definition of the ECSA and FIXED storage constrained level. CSM now sets the ECSA storage at the constrained level When ECSA storage usage is above 80% of the MAX ECSA value and approaching 85% of the MAX ECSA value. CSM sets the fixed storage at the constrained level When fixed storage usage is above 80% of the MAX FIXED value and approaching 85% of the MAX FIXED value.
CSM changed the definition of the ECSA and FIXED storage normal level. CSM sets the ECSA storage at the normal level when ECSA storage usage goes below 80% of the MAX ECSA value. CSM sets the fixed storage at the normal level when fixed storage usage goes below 80% of the MAX FIXED value.
CSM issues the message IVT5564I ECSA storage shortage relieved when the current ECSA storage usage goes below 80% of the MAX ECSA value. CSM issues the message IVT5565I fixed storage shortage relieved when the current fixed storage usage goes below 80% of the MAX FIXED value.
The CSM Monitor function is available to monitor CSM buffers between many components of z/OS for Communication Server. This function can be controlled using the Modify CSM command with the MONITOR operand. The valid options are MONITOR=ON, MONITOR=OFF and MONITOR=DYNAMIC. The default value of the MONITOR option is DYNAMIC. If you choose the option MONITOR=DYNAMIC, CSM Buffer Monitoring is dynamically activated and inactivated.
In prior releases, CSM activated dynamically CSM buffer Monitoring when CSM storage usage reached 85% or higher of the MAX ECSA value or the current fixed storage usage reached 85% or higher of the MAX FIXED value. CSM inactivated the Dynamic CSM Monitor function when the current ECSA storage usage went below 80% of the MAX ECSA value and the current fixed storage usage went below 80% of the MAX FIXED value.
In z/OS V1R9, the threshold for activating the Dynamic CSM Monitor function is when the storage usage is 80% or higher. The threshold for inactivating the Dynamic CSM Monitor function is when the storage usage goes below 75%.

Section

SNA serviceability enhancements

This section describes some serviceability enhancements made for SNA in z/OS V1R9 Communications Server.

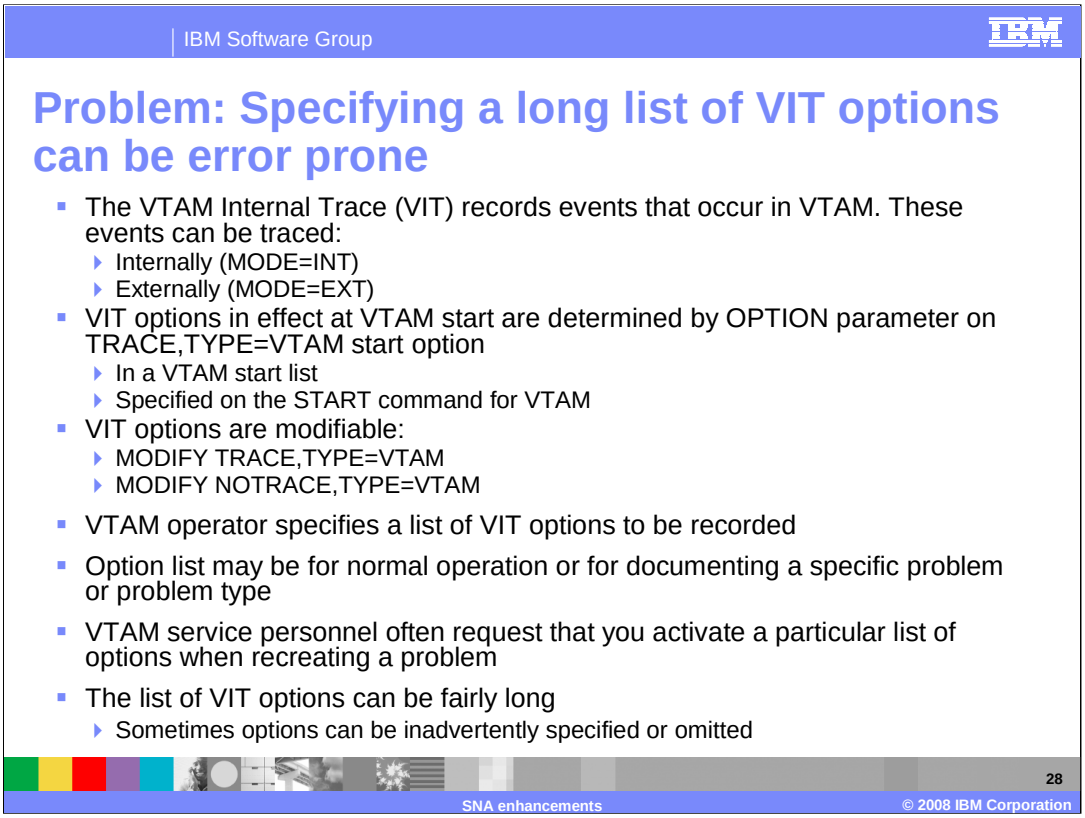

There are two wraparound tables in storage for internal VIT recording: The ECSA table is from 100 to 999 pages in size and is used to record the most recent events. The optional data space table (in 'net'.ISTITDS1) is from 10 to 50 megabytes in size. Entries are copied to it from the ECSA table periodically to preserve older event records.

The external trace can be much larger and therefore is recommended for documenting problems where a substantial amount of history is needed. GTF must be active for VTAM external tracing. External VIT tracing will only occur if explicitly requested. No VIT options are traced by default fore external tracing. Any combination of VIT options can be turned on or off.

On the other hand, VTAM always records certain VIT entries to an internal trace table. The user can expand this trace table, use the optional data space table, and specify that many more options be traced. But the user cannot completely turn off internal tracing.

Events traced fall into categories called VIT options. Each option consists of one or more individual VIT entry types. The SNA Resource Definition Reference describes how to code the TRACE,TYPE=VTAM start option. The VIT options and entries are described in detail in SNA Diagnosis Volume 2. SNA Operation describes how to modify the VIT options by turning individual options on or off. SNA Diagnosis Volume 1 has a detailed treatment of VIT option modification.

Specifying a long list of VIT options can be error prone. It may be difficult to remember the right VIT options to specify to document a particular type of problem. Forgetting an option might lead to another re-create request! So the choice for the user was to specify ALL to be sure everything needed was traced, or to ask for, look up, divine, or remember the best list of VIT options for the situation at hand. However, specifying ALL when not required fills the internal and external VIT tables quicker, making lost entries due to wrapping more likely.

Solution: VIT option group names

- VIT group options have been added to z/OS V1R9
  - Each group option represents a list of individual group options that are pertinent to tracing one type of problem area
  - This makes it easier for the operator to correctly specify the list of VIT options to be recorded during normal operation, or when diagnosing a particular type of problem.
  - The new VIT group options:
    - APIOPTS – diagnose non-LU 6.2 application program problems
    - APPCOPTS - diagnose LU 6.2 application program problems
    - CPCPOPTS - diagnose CP-CP session problems
    - CSMOPTS - diagnose communications storage manager (CSM) problems
    - DLUROPTS - diagnose dependent LU requester (DLUR) problems
    - EEOPTS - diagnose Enterprise Extender (EE) problems
    - HPDTOPTS - diagnose high performance data transfer (HPDT) problems
    - HPROPTS – diagnose high performance routing (HPR) problems
    - LCSOPTS - diagnose LAN channel station (LCS) problems
    - QDIOOPTS - diagnose queued direct I/O (QDIO) problems
    - STDOPTS - diagnose problems related to high CPU, session services, storage, Open/Close ACB, and DLCs such as MPC and CTC
    - TCPOPTS -diagnose problems related to TCP/IP
    - XCFOPTS - diagnose cross-system coupling facility (XCF) problems
- VIT group options containing HPR as a component option can be used with MODIFY NOTRACE to inactivate HPR subtrace option ARBP – but the HPR option will remain active.
- Any VIT group option can be used with MODIFY NOTRACE to inactivate SSCP subtrace option TGVC and TREE – but the SSCP option will remain active.

Before z/OS V1R9 Communications Server, the only group option available was ALL. It could be used to turn all of the VIT options on or off. The z/OS V1R9 Communications Server provides 13 new VIT group options that will make it easier to get exactly the right set of VIT options activated. The name of each group option is intended to convey its meaning. Each option is applicable to tracing a particular type of problem.

The MODIFY TRACE command will add the OPTIONs specified to the currently active list of options for the specified MODE (internal or external). It doesn't replace the currently active list of options with the ones specified. The MODIFY NOTRACE command will subtract the OPTIONs specified from the currently active list of options for the specified MODE (internal or external). Just as multiple options can be specified on for TRACE,TYPE=VTAM, multiple group options can be specified, even though they overlap. And a mixture of group options and individual options can be specified as well. VTAM will sort it out!

Two of the VIT options, HPR and SSCP, have associated subtrace options. The subtrace options are inactive by default. The HPR option has an ARBP subtrace option. The SSCP option has two subtrace options: TGVC and TREE.

Subtrace options can be turned on or off with a MODIFY TRACE or MODIFY NOTRACE command. The associated VIT option must be included in the command for this to be accepted. For example:

•F net,TRACE,TYPE=VTAM,OPTION=SSCP,SUBTRACE=TGVC is valid.

•F net,TRACE,TYPE=VTAM,OPTION=CIO,SUBTRACE=TGVC is not valid.

With this new function, any VTAM group option containing HPR as a component option can be used to activate or inactivate HPR subtrace option ARBP.

For example, F net,TRACE,TYPE=HPROPTS,SUBTRACE=ARBP will activate HPR subtrace ARBP in addition to the HPR option and the other component options of HPROPTS.

And F net,NOTRACE,TYPE=QDIOOPTS,SUBTRACE=ARBP will inactivate subtrace option ARBP and all component options of QDIOOPTS **except for HPR**! That is because F net,NOTRACE,TYPE=HPR,SUBTRACE=ARBP inactivates subtrace option ARBP but not option HPR.

All the group options contain SSCP as a component option, so any group option can be used to activate or inactivate SSCP subtraces TGVC and TREE. But such an inactivation will leave the SSCP option itself active.

It's simpler, and recommended, to use the appropriate individual VIT option to turn off subtraces.

VIT group option equivalencies Part 1

Group Option	Individual option equivalent
APIOPTS	API,MSG,NRM,PIU,PSS,SMS,SSCP
APPCOPTS	API,APPC,MSG,NRM,PIU,PSS,SMS,SSCP
CPCPOPTS	API,APPC,MSG,NRM,PIU,PSS,SMS,SSCP
CSMOPTS	API,APPC,CIO,CSM,MSG,NRM,PIU,PSS,SMS,SSCP,XBUF
DLUROPTS	API,APPC,HPR,MSG,NRM,PIU,PSS,SMS,SSCP
EEOPTS	CIA,CIO,HPR,MSG,NRM,PIU,PSS,SSCP,SMS,TCP
HPDTOPTS	CIA,CIO,HPR,MSG,PIU,PSS,SMS,SSCP

Specifying **APIOPTS** is equivalent to specifying all of these VIT options: API, MSG, NRM, PIU, PSS, SMS and SSCP.

Specifying **APPCOPTS** is equivalent to specifying all of these VIT options: API, APPC, MSG, NRM, PIU, PSS, SMS and SSCP.

Specifying **CPCPOPTS** is equivalent to specifying all of these VIT options: API, APPC, MSG, NRM, PIU, PSS, SMS and SSCP.

Specifying **CSMOPTS** is equivalent to specifying all of these VIT options: API, APPC, CIO, CSM, MSG, NRM, PIU, PSS, SMS, SSCP and XBUF.

Specifying **DLUROPTS** is equivalent to specifying all of these VIT options: API, APPC, HPR, MSG, NRM, PIU, PSS, SMS and SSCP.

Specifying **EEOPTS** is equivalent to specifying all of these VIT options: CIA, CIO, HPR, MSG, NRM, PIU, PSS, SMS, SSCP and TCP.

Specifying **HPDTOPTS** is equivalent to specifying all of these VIT options: CIA, CIO, HPR, MSG, PIU, PSS, SMS and SSCP.

VIT group option equivalencies Part 2

Group Option	Individual option equivalent
HPROPTS	API,APPC,CIA,CIO,HPR,MSG,NRM,PIU,PSS,SMS,SSCP
LCSOPTS	CIO,LCS,MSG,NRM,PIU,PSS,SMS,SSCP
QDIOOPTS	CIA,CIO,HPR,MSG,NRM,PIU,PSS,SMS,SSCP
STDOPTS	API,CIO,MSG,NRM,PIU,PSS,SMS,SSCP
TCPOPTS	CIA,CIO,MSG,NRM,PIU,PSS,SMS,SSCP,TCP
XCFOPTS	CIA,CIO,HPR,MSG,NRM,PIU,PSS,SMS,SSCP,XCF

Specifying **HPROPTS** is equivalent to specifying all of these VIT options: API, APPC, CIA, CIO, HPR, MSG, NRM, PIU, PSS, SMS and SSCP.

Specifying **LCSOPTS** is equivalent to specifying all of these VIT options: CIO, LCS, MSG, NRM, PIU, PSS, SMS and SSCP.

Specifying **QDIOOPTS** is equivalent to specifying all of these VIT options: CIA, CIO, HPR, MSG, NRM, PIU, PSS, SMS and SSCP.

Specifying **STDOPTS** is equivalent to specifying all of these VIT options: API, CIO, MSG, NRM, PIU, PSS, SMS and SSCP. These are the options traced internally by default.

Specifying **TCPOPTS** is equivalent to specifying all of these VIT options: CIA, CIO, MSG, NRM, PIU, PSS, SMS, SSCP, and TCP.

Specifying **XCFOPTS** is equivalent to specifying all of these VIT options: CIA, CIO, HPR, MSG, NRM, PIU, PSS, SMS, SSCP and XCF.

Problem: Cannot get immediate dump of VTAM

- XCF links connect VTAM hosts in a sysplex
- When an XCF link INOPs,
  - Dumps of involved VTAMs can be requested
  - Timely dump of local VTAM is possible with MODIFY CSDUMP,MESSAGE=IST1504I
- No current means to get an immediate dump of the VTAM on the other end of the INOPing XCF link
- Manual dump of remote VTAM host will likely be too late

XCF links are used to connect VTAM hosts in a sysplex. When an XCF link INOPs, existing VTAM facilities can be used to obtain a timely dump of the local VTAM for problem diagnosis. The operator can do this by setting a trigger on message IST1504I, which is only issued at the time of an XCF link INOP.

However, there is no current means to get an immediate dump of the VTAM on the other end of the XCF link. By the time a dump of the other VTAM is requested by the operator, it may be far too late to determine anything useful from it.

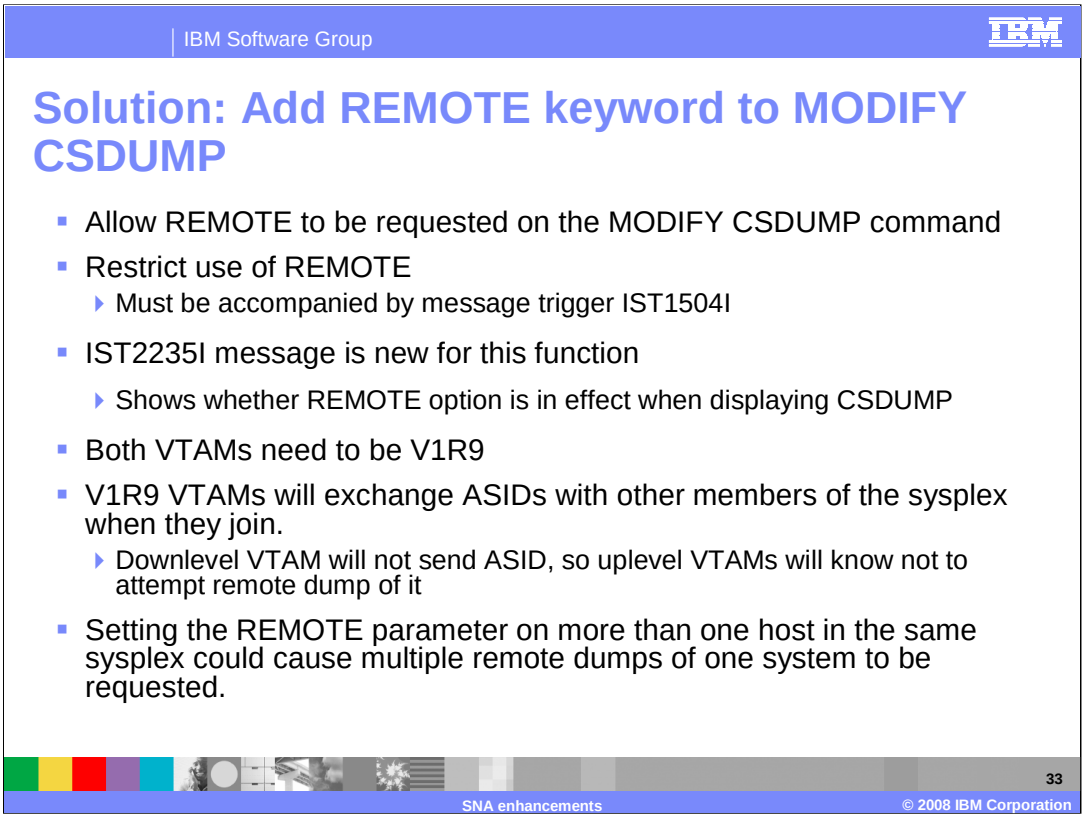

For MODIFY CSDUMP, VTAM issues an SDUMPX request to the system. If VTAM is connected to other VTAMs in a sysplex, the SDUMPX request by VTAM can include the existing REMOTE parameter to dump another VTAM in the sysplex. VTAM will only attempt this when message IST1504I is issued, and only when the operator has specifically requested it using the new REMOTE parameter on MODIFY CSDUMP.

New message IST2235I will show whether REMOTE is in effect for CSDUMP. It is added to message group IST1871I and only displayed if MESSAGE=IST1504I trigger is set.

On the SDUMPX request, VTAM needs to specify the ASID of the VTAM on the remote host. Otherwise, VTAM can't be dumped in the remote host. If REMOTE is active for CSDUMP and an XCF link INOPs, VTAM will check for the ASID of the partner VTAM. If it was not received, no remote dump attempt will be made.

It should only be necessary to set the REMOTE parameter on in one system per sysplex.

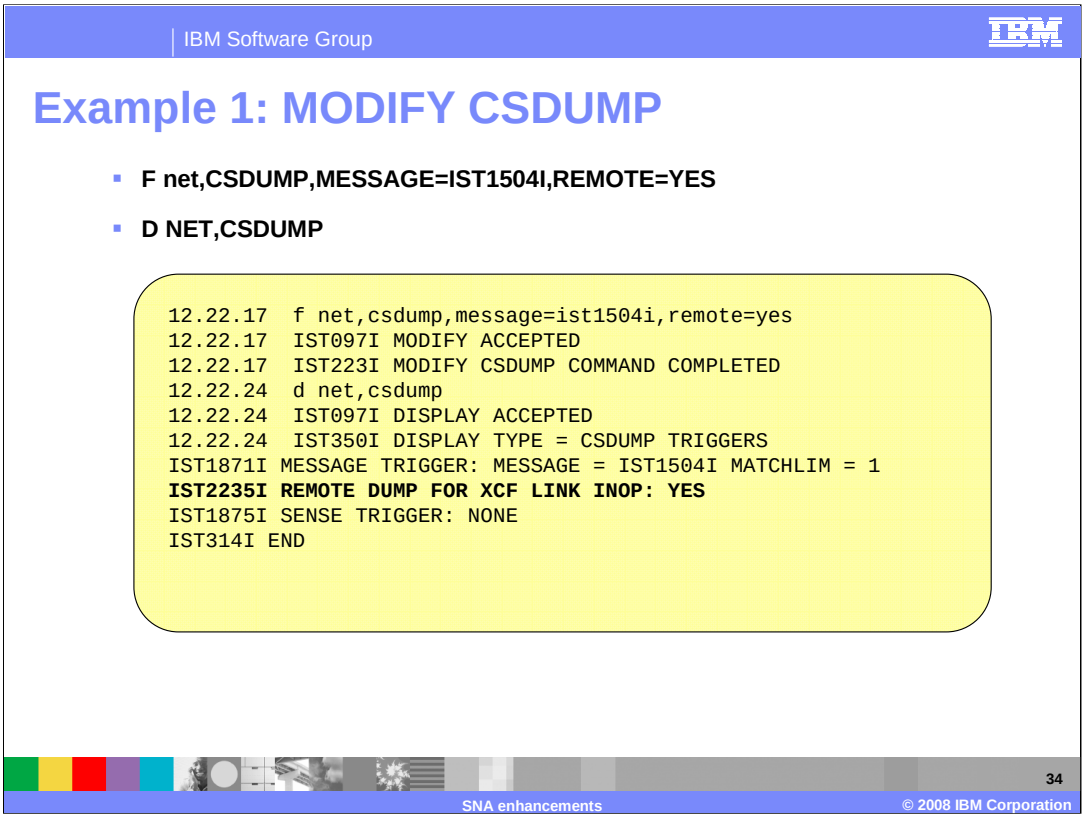

In this example, a local dump and a dump of the remote host will be attempted if the XCF link to any other VTAM connected in the sysplex should INOP. The output from the D NET,CSDUMP command tells you that the REMOTE=YES was specified on the MODIFY CSDUMP command.

Example 2: MODIFY CSDUMP

- **F net,CSDUMP,MESSAGE=(IST1504I,SSCP2A),REMOTE=YES**
- **D NET,CSDUMP**

```
12.40.08  f net,csdump,message=(ist1504i,sscp2a),remote=yes
12.40.08  IST097I MODIFY ACCEPTED
12.40.08  IST223I MODIFY CSDUMP COMMAND COMPLETED
12.40.16  d net,csdump
12.40.16  IST097I DISPLAY ACCEPTED
12.40.16  IST350I DISPLAY TYPE = CSDUMP TRIGGERS
IST1871I MESSAGE TRIGGER: MESSAGE = IST1504I MATCHLIM = 1
IST1872I   VALUE 1 = SSCP2A
IST2235I REMOTE DUMP FOR XCF LINK INOP: YES
IST1875I SENSE TRIGGER: NONE
IST314I END
```


In this second example, the message trigger includes variable text to restrict it to an IST1504I message identifying a specific system (by CP Name).

The trigger will match, and a local dump and a remote dump will be requested, only if the XCF link to the named system INOPs. If an XCF link to another system INOPs, no local or remote dump will be attempted.

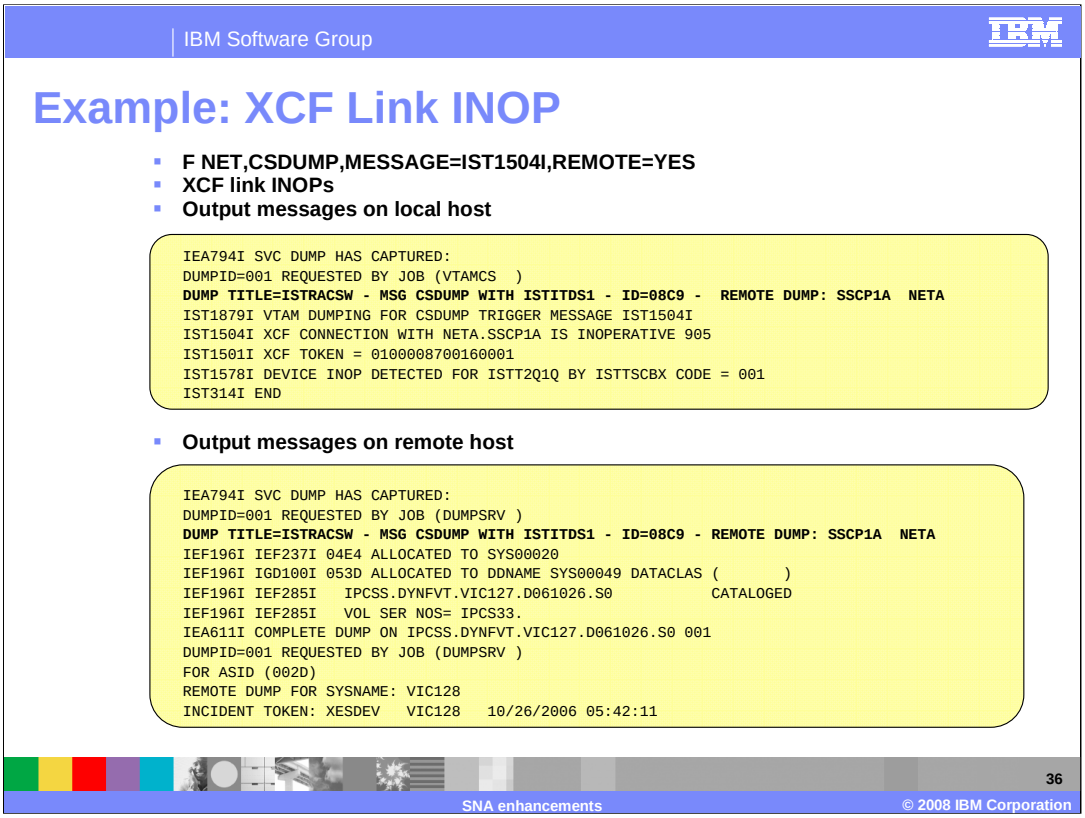

The top set of messages are seen on the VTAM host where the F net,CSDUMP,MESSAGE=IST1504I,REMOTE=YES command has been issued, and the XCF link INOPs. Note that the dump title includes the name of the remote host on which a dump is also requested.

The bottom set of messages are seen on the remote VTAM host when the XCF link INOPs. Note that the dump title is the same on both hosts. The IEA611 message identifies the host that requested that this remote dump be taken.

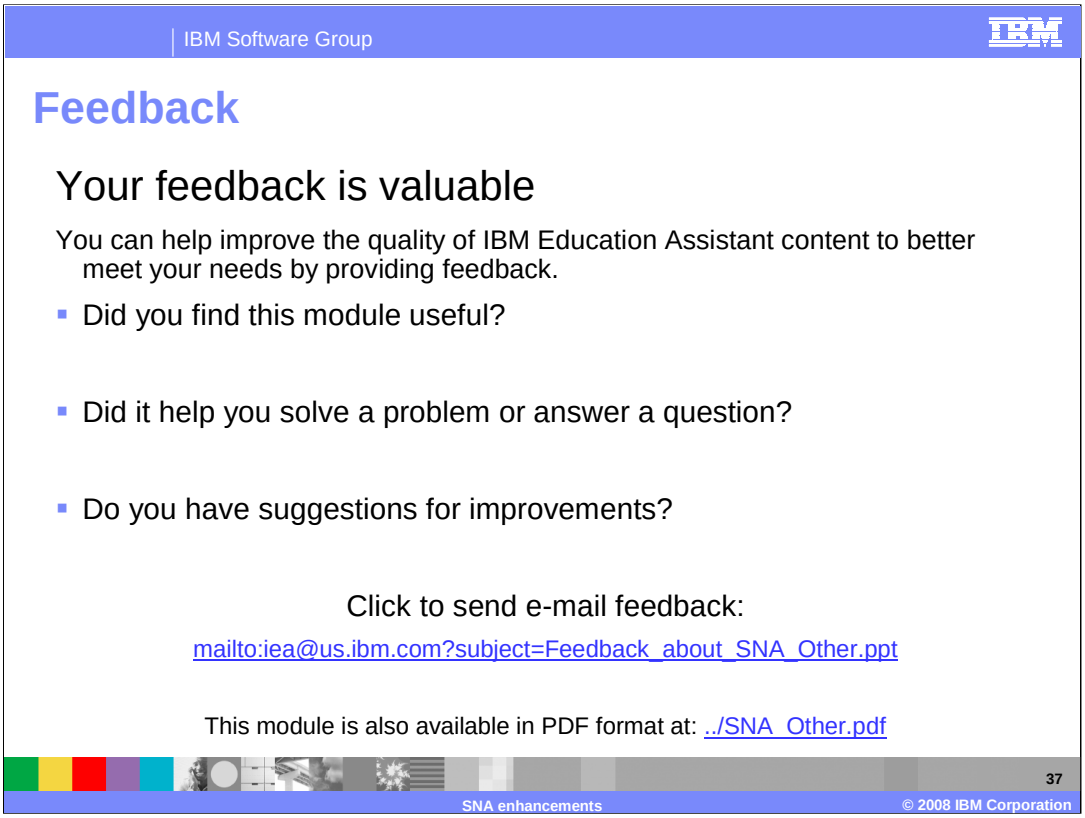

You can help improve the quality of IBM Education Assistant content by providing feedback.

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:

FICON HiperSockets 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.

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.