




IBM eServer™

Enterprise Extender and SNA: Configuration and control

@business on demand software

© 2007 IBM Corporation

Agenda - SNA and Enterprise Extender



- 1 Entry Border Node Alias Search Control
- 2 Directory Services Management Exit Interchange Node Selection Function
- 3 Dynamic update of VTAM® application major nodes

Entry Border Node Alias Search Control

Alias searches across NETIDs

- **An alias search in APPN is one that does not have an authentic NETID for the target LU.**
- **An authentic NETID allows APPN border nodes to control multiple network searching using NETID-specific adjacent cluster(ADJCLUST) tables.**
- **Alias searches may result in broadcast searching of multiple networks.**
- **Many applications do not support network qualified names and rely on alias name translation or CDRSCs to convert the alias NETID to a real NETID.**
 - The conversion must be done in the OLU network to be effective.
- **Many customers have experienced problems with alias searches entering their network due to inadequate controls in adjacent customer networks.**
- **These alias searches can result in broadcast searches of their network and attached networks.**
- **The only method available to prevent these alias searches from entering a network is the Initial Authorization function of the DSME.**

New option to control reception of alias searches

➤ **z/OS® V1R8 provides a simple method of preventing alias searches from entering an APPN network over an intersubnetwork link.**

➤ **This is done by adding a new ALIASRCH operand to the Adjacent Control Point minor node.**

- The operand is only allowed on ADJCP definitions activated by entry border nodes.
- The operand can only be coded on ADJCP definitions for adjacent non-native nodes.
- If an alias search is received from an ADJCP with ALIASRCH=NO specified, the search is rejected with a unique sense code x'087F0009' before any searching is performed.
- The adjacent non-native node is allowed to continue searching.

➤ **This solution could completely eliminate the need for a DSME for most customers.**

➤ **Syntax for the adjacent control point major node**

```

VBUILD
>>-----V-BUILD--TYPE=ADJCP-----<
  '-name-'

ADJCP
>>-name--ADJCP----->
  .-ALIASRCH=YES-----
>----->
  |           .-YES-. |
  '-ALIASRCH=-----'
  '-NO--'

```

Things to think about

- **Setting ALIASRCH=NO may cause search failures for adjacent customers.**
- **Online help for border node searching is available:**
 - EBN Searching in APPN/HPR Border Node networks
 - <http://www.ibm.com/support/docview.wss?uid=swg27006998>

Directory Services Management
Exit (DSME) Interchange Node
(ICN) Selection Function

ICN selection for subarea searches

- **Subarea nodes are dynamic end nodes in the DS database.**
 - Network Node Server of a subarea node is ICN where search enters/leaves the APPN network.
- **Database can be updated with each search sent or received.**
 - Not just for the target LU, but for every LU on that same subarea node.
 - May result in undesirable routes being cached, if multiple ICNs provide connectivity to a subarea node.
 - Relief requires periodically issuing MODIFY DIRECTRY.
- **Customers do not have a way to reliably control through which Interchange Node (ICN) a resource that resides "in subarea" is found.**
- **In most situations and for most customers this is not a problem since any ICN that provides the necessary subarea connectivity will do.**
- **For APPN networks with geographically separated ICNs, more control over ICN searching may be necessary.**

Notes about ...ICN Selection

NOTES

- Customers do not have a way in which to reliably control through which Interchange Node (ICN) a resource that resides "in subarea" is found. In most situations and for most customers this is not a problem since any ICN that provides the necessary subarea connectivity will do.
- Where this lack of control is shown to be a big problem is for customers who have "regional data centers." For example, it is not desirable to use a "west coast" ICN if both the requester of the session and the target of the session are on the "east coast", as this would result in a session path that traverses the continent twice (once in each direction).
- It is also a problem for multiple subnetwork configurations with border node and SNI connections. SNI-connected resources are treated as native APPN resources when located through native ICNs, and this limits the effectiveness of the border node selection function.
- The most common way in which VTAM could end up using an undesirable (wrong) ICN for sessions occurs when any resource on the subarea node is the ORIGINATOR of another session/search. If this INBOUND search from that resource is received through an undesirable ICN, it can cause other NNs in the network to cache the undesirable ICN and use it for subsequent OUTBOUND searches/sessions for any resource on the subarea node. In cases such as this, it is not always possible for the customer to prevent this from happening because they are not always in control of how searches are sent INTO their network (it is under the control of an adjacent customer).
- There is currently no convenient way to control the order in which ICNs are searched. The only "relief" that has been available to customers are the following two methods which are not very desirable:
 - Periodically issue the `MODIFY DIRECTRY` command on their VTAM NNs (using a CLIST) to modify the DS resource hierarchy for these subarea LUs to prevent GOOD info from being updated/replaced by BAD info,
 - Configure all of their ICNs as EBNs and code a DSME to do some unorthodox stuff (like disregarding any CACHED info and using EBN selection to force ICNs/EBNs to be searched in a customer defined order).

New ICN selection function in the DSME exit routine

- **The solution is to give customers a method of re-ordering ICNs to the preferred order for routing locate searches to cross domain subarea resources.**
- **This is accomplished by adding a new ICN selection function to the DSME exit.**
- **The new ICN selection function will have a similar design as the DSME border node selection function.**
- **When the serial ICN search task is called, a list of ICN names is passed to the exit.**
- **The exit can make the following changes to the ICN list:**
 - Reorder the entries in the list
 - Remove entries from the list
- **In addition to modifying the list, the exit can indicate that no interchange node search procedure should be attempted.**
 - To prevent DS from using non-optimal information stored in the database, the existing bit in the border node options will be used.
 - To prevent DS from using database information for cross domain subarea resources, the bit should be turned off during begin processing.

Things to think about

- **Turning off the normal cross domain subarea searching will result in cached information being ignored and more broadcast searching.**
- **When normal cross domain subarea searching is disabled a network node will not utilize the Central Directory Server.**

Dynamic update of VTAM
application major nodes

Changing application major nodes

➤ **It is currently not possible to modify application major nodes, which is a source of concern:**

-Adding application

- To prevent outage - add a new APPL major node
- Adding application major nodes results in operational complexity
- To eliminate operational complexity major nodes are merged during IPL of system
- Merge process during IPL can be difficult and error prone, especially during unplanned outages

➤ **z/OS V1R8 provides the ability to modify an application major node by allowing the UPDATE operand on the vary activate command. This change allows:**

-Adding APPL resources

- specify "update=add" to add resources

-Deleting APPL resources

- specification of "update=all" is required to delete APPL resources

-Changing APPL resources

- specification of "update=all" is required to change existing APPL values

Adding applications

➤ Adding APPL resources

- Not necessary to inactivate major node

- Procedure:
 - Edit VTAMLST member to add new applications
 - Issue "v net,act,id=appl_major_node,update=add"
 - "update=add" or "update=all" will add resources
 - Resource added after a model in VTAMLST will be added after clones in storage

Deleting applications

➤ Deleting APPL resources

- Not necessary to inactivate major node
- Procedure:
 - Edit VTAMLST member to remove applications
 - Inactivate APPL to be deleted
 - If MODEL - inactivate all clone applications
 - Issue "v net,act,id=appl_major_node,update=all"
- Note that the "all" value is required for delete operations
- Application to be deleted must be inactive
- Clone application deleted when inactivated
 - Never considered for deletion by the UPDATE process

Changing applications

➤ Changing APPL resources

- Not necessary to inactivate major node
- Procedure:
 - Edit VTAMLST member to change application definitions
 - modify APPL / GROUP definition
 - inactivate APPL resources to be changed
 - Issue "v net,act,id=appl_major_node,update=all"
- Any definition parameter may be changed
- Requirement to change APPL parameters
 - APPL resource must be inactive
 - Exception! Table changes do not require the APPL to be inactive
 - Tables for APPL are ASLTAB, MODETAB, MDLTAB, USSTAB
- GROUP definition changes
 - Only resources in required state are changed
 - Error message is issued for each resource not in required state in the GROUP
- Clone APPL resources can not be changed



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:

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.

© Copyright International Business Machines Corporation 2007. All rights reserved.

Note to U.S. Government Users - Documentation related to restricted rights-Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract and IBM Corp.