APPN™ Management Using NetView™ (Presenter Notes)

February 3, 1994

Steven M. Golberg

Network Management Products Field Support
A86/502
Research Triangle Park, NC

VNET: SGOLBERG(RALVM14)
IBMMAIL: USIB237D(IBMMAIL)
Internet: smg@vnet.ibm.com

PHONE: (919) 254-7159 (Tie 444-7159)

IBM Internal Use Only

Notice

The information contained in this document has not been submitted to any formal IBM test, and is distributed on an ″As Is″ basis without any warranty, either expressed or implied. The use of this information or the implementation of any of these techniques is a customer responsibility and depends on the customer′s ability to evaluate and integrate them into the customer′s operational environment. While each item may have been reviewed by IBM for accuracy in a specific situation, there is no guarantee that the same or similar results will be obtained elsewhere. Customers attempting to adapt these techniques to their own environments do so at their own risk.

In this document, any references made to an IBM licensed program are not intended to state or imply that only IBM′s licensed program may be used; any functionally equivalent program may be used instead.

Any performance data contained in this document was determined in a controlled environment, and therefore the results which may be obtained in other operating environments may vary significantly. Users of this document should verify the applicable data for their specific environment.

It is possible that this material may contain references to, or information about IBM products (machines and programs), programming or services that are not announced in your country. Such references or information must not be construed to mean that IBM intends to announce such IBM products, programming or services in your country.

© Copyright International Business Machines Corporation 1994

Preface

Trademarks

The following are trademarks or registered trademarks of the International Business Machines Corporation

- IBM
- NetView
- AS/400
- OS/400
- VTAM
- APPN
- Advanced peer-to-peer networking
- MVS/ESA
- OS/2
- SystemView
- AIX
- APPC

Purpose of Presentation

The purpose of this presentation is to give both an overview and a fairly detailed look into what network management capabilities are present from S/390 NetView in an APPN environment. The presentation only covers products and services that are available as of 1/94. This presentation is not intended to be a presentation concerning IBM′s directions in the management of APPN networks, nor is it intended to cover how to manage an APPN environment from any of IBM′s other management platforms (such as AIX SystemView NetView/6000, OS/400, or LAN NetView).

This presentation can be given in many different manners. Normally, it can be given in 1.5-2 hours without many of the detailed examples. In such a manner, only parts A, B, C, and D will be covered. Parts E, F, and G are present for additional background information and to give the customer some concrete examples of the implementation of a few different remote operations techniques in an APPN environment. Additionally, by skipping many of the examples of part C, the presentation may be shortened to 1 hour or so. Last, for those audiences wanting to see every detail, part D should be presented last. Part E may be presented after the Remote Operations charts in section C. Presenting this presentation in this manner may very well take 3-4 hours.

It is critical that the audience have some background in APPN and NetView before the presentation. If the audience does not have any APPN background, then a 10-15 minute discussion of various APPN terminology (such as network nodes, end

nodes, LEN nodes, topology databases, and APPN route selection) should be done first. If the audience has no NetView background, then another introductory presentation on NetView should be performed first (possibly along with a demonstration, either using the SuperSystem or resources locally available to the presenter - contact SUPERSYS @ RALVM12 or see MKTTOOLS for more information on the SuperSystem).

I would recommend that the presenter of these FOILS have a good working knowledge of at least APPN and NetView. Otherwise, even with these Presenter's notes, it may prove difficult to properly handle questions that are brought up by the audience.

The release information contained in this document is the minimum release of a product needed to perform the specified function. Unless otherwise stated, the specified function is supported by any follow-on releases of that product that were known about as of 1/94. For example, when NetView V2R2 is mentioned as supporting a function, that function is also supported by later releases of NetView (such as V2R3 and V2R4) unless otherwise stated.

Acknowledgements

I would like to thank all of the people who reviewed the charts for this presentation, as well as those that provided assistance and input to this presentation. I would like to thank Roger Adema, Mike Allen, Pam Buckoski, Jim Fletcher, Mark McKelvey, Stuart Riordan, David Schmidt, David Shogren, and Richard Szulewski. My apologies to anyone who I have mistakenly forgotten to list.

My thanks to all of the above for their help. All errors are of course my own.

Feedback and questions may be sent to me either on-line or by phone via the addresses listed on the title page.

Contents

Part B - Agenda . . . 1
Agenda (B.2) . . . 1

Part C - APPN Management from NetView . . . 3
Overview (C.2) . . . 3
Transporting APPN management data . . . 3
First Page (C.3) . . . 3
Second Page (C.4) . . . 4
Focal Point and Sphere of Control . . . 5
First Page (C.5) . . . 5
Second Page (C.6) . . . 6
Third Page (C.7) . . . 6
Fourth Page (C.8) . . . 6
Fifth Page (C.9) . . . 7
Pages 6-10 (C.10-C.14) . . . 7
Problem reporting to NetView (C.15) . . . 7
Remote Operations for APPN nodes . . . 8
First Page (C.16) . . . 8
Second Page (C.17) . . . 9
Third Page (C.18) . . . 9
Fourth Page (C.19) . . . 10
Session Level PD in APPN . . . 10
First Page (C.20) . . . 11
Second Page (C.21) . . . 11
Third Page (C.22) . . . 11
Fourth Page (C.23) . . . 12
Fifth Page (C.24) . . . 12
Session PD Configuration (C.25) . . . 12
Session PD Example 1 . . . 13
First Page (C.26) . . . 13
Second Page (C.27) . . . 13
Session PD Example 2 . . . 13
First Page (C.28) . . . 14
Second Page (C.29) . . . 14
Session PD Example 3 . . . 14
First Page (C.30) . . . 14
Second Page (C.31) . . . 15
Session PD Example 4 (C.32) . . . 15
Session PD Example 5 . . . 16
First Page (C.33) . . . 16
Second Page (C.34) . . . 16
Session PD Example 6 . . . 16
First Page (C.35) . . . 16
Second Page (C.36) . . . 17
Graphical display of an APPN network . . . 17
First Page (C.37) . . . 17
Second Page (C.38) . . . 18
Third Page (C.39) . . . 18
Fourth Page (C.40) . . . 19
Fifth Page (C.41) . . . 19
Sixth Page (C.42) . . . 19

Seventh Page (C.43) . . . 19
APPC Accounting Collection (C.44) . . . 20
APPNTAM Accounting Function . . . 20
First Page (C.45) . . . 20
Second Page (C.46) . . . 21
Third Page (C.47) . . . 21
Fourth Page (C.48) . . . 21
Fifth Page (C.49) . . . 22
Sixth Page (C.50) . . . 22
Seventh Page (C.51) . . . 23
Agent-Manager Communications (C.52) . . . 23

Part D - Summary . . . 25
Summary (D.2) . . . 25

Part E - Remote Operations Examples . . . 27
Remote Operations Examples - VTAM . . . 27
First Page (E.2) . . . 27
Second Page (E.3) . . . 27
Third Page (E.4) . . . 27
Fourth Page (E.5) . . . 28
Fifth Page (E.6) . . . 28
Remote Operations Examples - AS/400 . . . 28
First Page (E.7) . . . 28
Second Page (E.8) . . . 28
Third Page (E.9) . . . 29
Fourth Page (E.10) . . . 29
Fifth Page (E.11) . . . 29
Pages 6-8 (E.12-E.14) . . . 29
Ninth Page (E.15) . . . 29
Tenth Page (E.16) . . . 29
Pages 11-15 (E.17-E.21) . . . 29
Pages 16-17 (E.22-E.23) . . . 30
Eighteenth Page (E.24) . . . 30
Nineteenth Page (E.25) . . . 30
Twentieth Page (E.26) . . . 30
Twenty-first Page (E.27) . . . 30
Twenty-second Page (E.28) . . . 30
Pages 23-27 (E.29-E.33) . . . 31
Pages 28-29 (E.34-E.35) . . . 31
Remote Operations Examples - RS/6000 . . . 31
First Page (E.36) . . . 31
Second Page (E.37) . . . 31
Third Page (E.38) . . . 32
Fourth Page (E.39) . . . 32
Fifth Page (E.40) . . . 32
Sixth Page (E.41) . . . 32
Seventh Page (E.42) . . . 33
Eighth Page (E.43) . . . 33
Ninth Page (E.44) . . . 33
Tenth Page (E.45) . . . 33
Eleventh Page (E.46) . . . 33
Twelfth Page (E.47) . . . 33
Thirteenth Page (E.48) . . . 33
Fourteenth Page (E.49) . . . 34

Fifteenth Page (E.50) 34
Sixteenth Page (E.51) 34
Seventeenth Page (E.52) 34
Remote Operations Examples - CM/2 34
First Page (E.53) 34
Pages 2-5 (E.54-E.57) 34
Sixth Page (E.58) 34
Pages 7-8 (E.59-E.60) 35
Pages 9-10 (E.61-E.62) 35
Pages 11-12 (E.63-E.64) 35
Thirteenth Page (E.65) 35
Fourteenth Page (E.66) 35

Part F - NetView Function/Release Information 37
NetView Function/Release Information (F.2-F.4) 37

Part G - Further Information 41

Part B - Agenda

Agenda (B.2)

IBM has been adding APPN management support to NetView since NetView V2R2 (even before there was an APPN implementation on the S/390).

This presentation will investigate and cover the various APPN management functions that are present in the NetView product. The presentation will cover the functions present in NetView

- from V2R2 through V2R4 NetView (which became available on 12/17/93).
- including the APPN Topology and Accounting Management (APPNTAM) feature that became available on 1/14/94, and which requires NetView V2R4 to run.
- including the IBM NetView Remote Operations Manager and IBM NetView Remote Operations Agent/400 products which became available on 12/17/93.

The various functions in this presentation were introduced in various releases of NetView and of the products mentioned above. Many of the charts in this presentation do not specifically list which release the function was introduced in. Rather, to simplify the bulk of the presentation, part F (NetView Function/Release Information) discusses where each of the various functions was introduced.

Please note, this presentation assumes some prior knowledge of NetView, and also some prior knowledge of APPN. Please see the preface for comments on the level of background that the audience should have, and some suggestions if the audience does not have this background.

Part C - APPN Management from NetView

Overview (C.2)

NetView has been implementing various APPN management functions since NetView V2R2. Therefore, NetView has been capable of some APPN management functions, even while residing on a subarea node (VTAM V3 or VTAM V4 configured only as a subarea node).

This presentation will cover various network (and systems) management functions that NetView can perform in an APPN environment. The functions that will be examined are:

- How NetView transports APPN management data to and from APPN nodes, even if NetView happens to reside in the subarea network. In other words, how APPN management data gets from point A to point B (even when one of the points is a NetView running on a subarea node).
- How NetView supports the APPN Focal Point and Sphere of Control architectures. These architectures determine how entry point-focal point relationships are established, who maintains them, and other related functions.
- How problems are reported to NetView in an APPN environment (remember, there are (or may be) no SSCP-PU sessions in this environment).
- How commands are sent to APPN nodes, and how the responses are sent back. This is basically remote operations, where you are controlling another node from NetView.
- How session level problem determination (PD) can be performed in an APPN environment, and in which cases session level PD data is present.
- The graphical display of the topology of your APPN network. Also, as part of this function, we will cover what control you have over the APPN network from this graphical display.
- The collection of APPC (which is also known as LU 6.2 - both terms will be used interchangeably in this presentation) accounting information from an APPN network, and how it can be gathered from more than one location.

Transporting APPN management data

First Page (C.3)

Summarize the key points of this page. See the SNA MS Reference Manual (listed in Part G) for more information about the MDS and high performance transports mentioned on these two pages.

Some important things to note -

- The MDS transport is used in an APPN environment just as SSCP-PU sessions are used in a subarea environment - to transport network management data through the network.

- The MDS transport architecture can be used even in a subarea environment if both endpoints support the architecture. For instance, you can use the MDS transport architecture between two NetViews or between a NetView and a CM/2 in a subarea environment. The MDS sessions have a couple of advantages over SSCP-PU sessions.
  - 31K of data can be sent per send (vs. 256-512 bytes for SSCP-PU sessions - depending on the direction the data is being sent).
  - Multiple requests that are waiting on a reply can be outstanding at once vs. just one outstanding request at a time for SSCP-PU sessions.
- Point out that many products have implemented the MDS transport.
- Also point out that four of the major IBM platforms have APIs for customers and vendors to use in writing applications that send/receive data over the MDS transport - NetView, OS/400, CM/2 and AIX SNA Server/6000.

Second Page (C.4)

NetView has implemented this function independent of whether it is running on a VTAM that has APPN capabilities. However, when you are using this function (in NetView V2R4) on a VTAM that is running as an APPN node (either EN or NN specified for NODETYPE - VTAM V4R1 or above), then you can tell NetView to try to utilize the MDS transport function that is built into VTAM as an APPN node.

When VTAM is a network node (NN), NetView will notify all of served end nodes (ENs) that are attached to this NN who the NN′s focal point is. This is the domain focal point function. That way, if the EN has no focal point, it can forward data to its domain focal point. This function is only available when NetView (V2R4) is set up to use the enhanced MS transport (via VTAMCP USE definition in DSIDMN).

If you have OS/400 V2R1 or S/36s with APPN installed on them, and you want to receive alerts from them via LU 6.2, then you also need to use NetView V2R4 with the enhanced MS transport. This is because the VTAM MDS implementation also supports the Migration (APPN) Node support that is implemented in OS/400 V2R1 and the S/36s with the APPN feature.

Even when the enhanced MS transport is being used (and NetView is trying to utilize the MDS transport function that is present in VTAM as an APPN node), the MDS transport function is still present in NetView. Therefore, programs that have been defined to send data over the MDS transport to NetView (via NetView′s LU name) do not need to be changed immediately. They can continue to send data to NetView′s LU, and NetView will be able to receive and reply (if necessary) to that data.

NetView also provides support for the High Performance transport. This transport is similar to MDS except that it does not use all of the confirmation flows. This makes it perform better and quicker at the cost of not getting a guaranteed error notification if a piece of data does not reach the other endpoint. To my knowledge, NetView is the only implementer of this transport so far, and uses it for connections that may have high traffic rates or that may be unsuitable to the MDS transport.

Focal Point and Sphere of Control

First Page (C.5)

Use the diagram at the bottom of this page (also provided in the freelance presentation charts (.PRS file)) to illustrate the various points on the next three pages. In fact, you could just use the diagram and talk to the points on these three pages using the diagram as your example.

It is important to get across a few different terms using the diagram and the first page. The terms that can be introduced are: entry point, focal point, category, nested focal point, nesting focal point, primary focal point, backup focal point implicit focal point, and explicit focal point. See the glossary in the SNA MS Reference for the official definitions of these terms. Depending on the amount of time that you have for this presentation, and the level of knowledge in the audience, you may want to skip or quickly go over some of these terms. The definition of entry point, focal point and category are the most important.

Entry point - this is a node (or an application on a node) that provides distributed network management support for a particular function set or category. It may reside in an APPN node or a subarea PU.

Focal point - this is a node (or an application on the node) that provides centralized management and control for other entry points within its sphere of control for a given function set or category. It may resides on an SSCP in the subarea network, or in an APPN NN.

Category - This is a NetView term that refers to the collection of services that together provides a given management services function for a node (whether it is the focal point or entry point services). For example, management services for problem management (SNA alert data) falls in the **alert** category.

Nested focal point - This is a node that provides focal point function for a category, but which, in turn, has a focal point itself. For example, if PC 3 forwards alert data to NetView 3, but NetView 3 has an alert focal point itself (NetView 1), then NetView 3 is a nested focal point. This is also sometimes referred to as an intermediate focal point.

Nesting focal point - This is a focal point node that is providing focal point function for other (nested) focal points. Using the prior example, NetView 1 would be the nesting focal point for NetView 3 for the alert category.

Primary focal point - This is whom you should normally be forwarding data to for a given category.

Backup focal point - This is whom you forward data to if the primary focal point is unavailable. For example, let us say that PC 1 has a primary focal point of NetView 1, and a backup focal point of NetView 2. If NetView 1 is down (or unreachable for some reason), PC 1 would forward its data to NetView 2 (the backup focal point).

Implicit focal point - This is when the focal point-entry point relationship is established by the entry point asking the focal point to provide focal point services for it. Normally, this is done by defining the focal point to the entry point and the entry

point contacting the focal point. However, backup focal point to entry point relationships are always implicit also (since the entry point knows who the backup focal point is, but the backup focal point doesn't necessarily know it is a backup focal point). For example, when PC 1 above establishes the relationship to NetView 2 as its backup focal point, then NetView 2 becomes an implicit focal point for PC 1.

Explicit focal point - This is when the focal point tells the entry point that it is going to provide focal point services for it, and that the entry point should forward its data to that focal point. In this case the focal point-entry point relationship is established by the focal point.

Second Page (C.6)

Sphere of Control is closely related to the focal point function (in fact, in NetView they are implemented in the same task). It lets you manage the focal point-entry point relationships from the focal point's perspective. For example, the sphere of control function in NetView (V2R4) would let you issue a command on NetView 1 and see what nodes are forwarding alert data to it via the alert focal point function, and to see whether the focal point-entry point relationships are explicit or implicit, and what state the relationships are in (active, inactive but being retried, etc.).

NetView also provides a DEFFOCPT statement (in DSI6INIT) that allows NetView as an entry point (V2R2) or nested focal point (V2R4) to have defined to it who its focal point should be (i.e., lets you set up an implicit focal point relationship).

Third Page (C.7)

This page lists some of the various focal point commands in NetView that allows you to control your focal point-entry point relationship, display its status, and control and display your sphere of control. Summarize the various points.

Fourth Page (C.8)

This page discusses why a customer would ever care about these functions. The prior three pages may be a bit esoteric to many customers, and they may not be all that interested in what architected functions NetView has implemented. This page (and the next) try to address why they may be interested that NetView has implemented these architectures.

Currently, the only application that actively uses the focal point function to control where data is sent is the alert function (Hardware Monitor in NetView). NetView (V2R4) currently only supports being an alert focal point (and not being an alert entry point) via the MDS transport. It uses a private (LUC) protocol to forward problem information from one NetView to another.

If you have an APPN environment, you would use the focal point and sphere of control architectures (and implementations) to control where alerts are sent in your network. It is important to know (when setting such an environment up) what alert focal point functions your entry points have - do they support having an implicit focal point? A backup focal point? Do they support reacquiring the primary focal point after the focal point entry point relationship has been lost? The sphere of control function in NetView will allow you to define all of your entry points to NetView and to have NetView establish the focal point-entry point relationships.

Another item to note is that if an EN supports having a domain focal point, then you only need to establish the focal point-entry point relationship with the NN. The NN will tell its served ENs about its focal point, and if they do not have a focal point, they will use the domain focal point as their focal point (allows you to only define NNs in a sphere of control, and not all of the NNs and ENs).

Fifth Page (C.9)

This page gives a specific example for where a sphere of control table may be useful. Please note, the sphere of control configuration file in NetView (V2R4) is dynamic (via the FOCALPT REFRESH command). Also, prior to NetView V2R4, if a focal point-entry point relationship could not be established from a NetView focal point (via the FOCALPT CHANGE command), then NetView would not attempt to retry establishing the relationship. With the addition of the sphere of control function to NetView V2R4, this has changed and it is periodically retried until the relationship is successfully established.

Pages 6-10 (C.10-C.14)

These pages are examples of the various NetView FOCALPT commands in a real environment. They are meant to show how to use some of the commands, and what sort of information they return. The text of these charts as well as the information from the screen captures should be fairly self-explanatory.

These examples were done on a NetView V2R4 system.

Page six (C.10) shows a FOCALPT QUERY command for the alert category. Note, for categories where NetView also supports proprietary forwarding schemes (such as alerts), this command will show any focal points, whether they are via MDS, or via the private NetView method.

Page seven (C.11) shows a FOCALPT DISPSOC (display sphere of control) command where there is nothing in NetView′s sphere of control for the alert category.

Page eight (C.12) shows the addition of 5 nodes to the sphere of control configuration file, along with the issuance of the FOCALPT REFRESH command to have that change take affect immediately on NetView N2412.

Page nine (C.13) shows the FOCALPT DISPSOC command again, to show the results of what happened after the new sphere of control file took affect. Only the entries in the SOC configuration file that had this NetView as the primary focal point were processed and acted upon by this NetView.

Page ten (C.14) is an example of a FOCALPT CHANGE command for the category of alerts that attempts to establish a focal point-entry point relationship (an explicit focal point) to node A19CMMCP.

Problem reporting to NetView (C.15)

Just as in a subarea environment, SNA alerts are used in an APPN environment to report problems. However, in an APPN environment, the alerts travel over the MDS transport. Also, since there is no automatic boss in an APPN environment (such as the SSCP in the subarea), the focal point architecture is used to determine where to forward alerts to.

It is important to note that this architecture (the focal point architecture and MDS transport) can be used to transport alerts even in a subarea environment, if both endpoints have implemented the architecture. This allows you to forward alerts to someone other than the owning SSCP of a node in the subarea environment (providing you with more flexibility in managing your network and systems).

Various devices support having an alert focal point. One thing to note about 6611′s support of having an alert focal point. The APPN support in 6611 is provided with the Multiprotocol Network Program (MPNP). In 6611′s first release of APPN support (MPNP V1R1.2), they only supported receiving notification of an alert focal point (explicit focal point) and forwarding alerts to that focal point. All OEM products that are based off of that level of 6611′s code may have the same limitation. 6611′s MPNP V1R2 (GA of 3/25/94) has expanded this support to include the ability to recover a focal point when the focal point-entry point relationship is lost (extremely important, as the focal point may not detect the loss or the entry point may be responsible for recovery of the relationship), and the ability to support the held alert function.

The held alert function (supported on various devices such as 6611 using MPNP R2 and CM/2 for example) lets you hold a specified number of alerts in case you generate an alert and don′t currently have a focal point to forward it to. You will hold that number of alerts, and when you reacquire a focal point (either primary or backup) this function lets you forward your held alerts to that focal point.

NetView has also supported the automation of the alerts themselves since NetView V2R2. This provides customers with the ability to take actions based on alerts received in NetView, and to be able to take that action based on the complete information present in the alert. Also, by not needing a message to automate the alert, it lets you save CPU time by not generating the BNJ146I message in NetView when alerts are received.

One item not mentioned in the presentation at this time is 3174′s support for an APPN alert focal point. This support is in 3174 C5 level of microcode, but it is not mentioned in this presentation since it will not GA until 6/24/94. A future update of this presentation will add the 3174 support. Currently (with the C3 or C4 level of microcode), 3174 can be an APPN NN, but they can only forward network management data via SSCP-PU sessions. The C5 level of the microcode will remove that restriction.

Remote Operations for APPN nodes

First Page (C.16)

This page is used to introduce remote operations in an APPN environment, along with why it is important. If you have the Freelance color charts, you may wish to simply show the Freelance chart and not the text on any of these four pages (and just talk to the points using the color chart). The color chart simply shows some APPN devices connected to each other via an arbitrary APPN network (the cloud).

Since APPN is a peer networking architecture, there is no preset hierarchy (other than the fact that NNs provide certain networking services for their served ENs). Therefore, if you want to take an action concerning a node, you need to take that action at the node itself (similar to TCP/IP which is also a peer networking architecture - you need to do a GET/SET on that node to display/change some information

about it). This makes remote operations even more important in an APPN environment than in a subarea environment where you can do much of this from an SSCP.

You can also use remote operations to perform systems management functions, since generally, if you can issue any command on that node, you can issue both networking and system oriented commands. This allows you to display or control system based functions on the remote nodes.

Second Page (C.17)

This page shows the method used by NetView to send commands to remote S/370 and S/390 nodes. The remote node must be running NetView and VTAM in order for this to work, as the NetView RMTCMD function (introduced in NetView V2R2) is used to send the command to the remote node and receive the replies back.

Any command that does not result in a full screen output can be sent using this command (RMTCMD). This means that VTAM, MVS, VM and VSE commands, as well as other application commands can be sent from NetView to a remote S/390. The RMTCMD response can be correlated, so you can use it in PIPEs (very powerful). See the NETVPIPE presentation on MKTTOOLS for many PIPEs examples, including the use of RMTCMD with PIPEs.

Also, the RMTCMD can be issued using the WINDOW command (sample in NetView V2R3, a supported command in V2R4), which presents the results in a full-screen view panel. This allows the customer to scroll through the results, or to re-issue the command by just pressing a PF key.

If NetView V2R4 is on both the local and remote system, you can also use the NETLOG browse and member browse functions remotely. NetView V2R4 supports remote NETLOG/member browse over LU 6.2 sessions, so this gives you an added measure of function for the remote operations of your network.

Part E contains some examples of remote operations commands that pertain to this environment.

Third Page (C.18)

If you have AS/400s in your network, you can remotely issue commands and receive responses from them. In order to do this, you need at least NetView V2R3 (along with pre-req APAR OY66689), NetView Remote Operations Manager, NetView Remote Operations Agent/400, and OS/400 V2R2. These products support the OPSMGMT command from NetView to AS/400, along with a correlated response in return. The OPSMGMT command uses the operations management architecture (support for which was introduced in NetView V2R2), and it transports the commands and responses via the MDS transport.

The OPSMGMT command allows any AS/400 command that can be entered in batch mode to be entered from NetView with the response being returned to NetView.

This function lets you avoid having to use HCF (Host Command Facility) and manually logging onto and off of the AS/400. It lets you automate your control of the remote AS/400s and automate the responses. It also allows you to issue the OPSMGMT command within a NetView PIPE.

The NetView Remote Operations Manager ships a CommandTree/2 (CT/2) library with some of the AS/400 commands. It also ships some samples where you can have AS/400s added to RODM and displayed via NGMF dynamically as they come up (if they are running the NetView Remote Operations Agent/400), and have them removed from RODM when the agents go down.

Fourth Page (C.19)

If you wish to remotely control PCs in an APPN environment, you can do so by sending RUNCMD (or any other SPCS command) over LU 6.2 to them.

NetView V2R3 and CM/2 1.0 support the sending of RUNCMDs over LU 6.2, as well as sending the responses to the commands back over LU 6.2 (using the MDS transport).

Note: You must have the CM/2 SPA Router and Remote Operations functions installed and active in order to do this for generic commands to the PC. This is important, since SPA Router and Remote Operations are not installed by default when you install CM/2. You need to add Remote Operations to the list of installed functions when installing CM/2 in order to get the Remote Operations and SPA Router functions.

RUNCMD responses can be correlated (base V2R4, maintenance needed for V2R3), so you can use RUNCMD within PIPEs also.

Any command that can be issued in linemode on the PC (CM/2 commands, OS/2 commands like DIR, PATH, SET PATH, etc.) can be issued using this method.

If you wish to remotely control workstations (e.g., RS/6000s), AIX NetView Service Point V1R2 and NetView/6000 (in combination) support receiving RUNCMDs via LU 6.2 (MDS Transport) and sending the replies back to NetView (again, NetView V2R3 is required for this function).

This function allows you to issue any linemode based RS/6000 command from NetView and to get the response back from NetView. This allows you to kick off shell scripts on the remote workstation, including AIX SystemView NetView/6000 commands (like rping) and scripts, and AIX SNA Server/6000 (V2R1) commands. AIX SNA Server/6000 V2R1 is the level of SNA server that provides the RS/6000 with APPN support.

NetView/6000 is required for the RUNCMD function, since the service point application daemon (spappld) ships as part of NetView/6000 and not as part of AIX NetView Service Point.

This function is very similar to CM/2′s support, with no differences from NetView′s perspective.

Session Level PD in APPN

First Page (C.20)

The Session Monitor function in NetView V2R4 has been enhanced to handle APPN session information. With the advent of VTAM V4R1 and NCP V6R2, APPN can now be implemented on the mainframe and NCP (traditionally where only subarea networking resided). NetView has been enhanced to collect network management data from VTAM and NCP regardless of whether the session information is for a pure subarea session, a pure APPN session, or a mixed subarea-APPN session. NetView can also display APPN information from remote NetViews (they must be V2R4 also) via the SDOMAIN and implicit cross domain data retrieval functions in Session Monitor.

Class of Service (COS) and logmode information pertaining to the APPN session can be retrieved. The session path for both the subarea and APPN portions (of a mixed subarea-APPN session) can be retrieved, depending upon the session configuration and the placement of VTAMs and NetViews. Also, flow control data (pacing information) can be retrieved from VTAM and NCP for APPN TGs. This is in addition to the pacing information and VR testing/analysis that Session Monitor can perform today for subarea TGs.

The biggest drawback in the APPN session PD function is that it is very dependent upon VTAM being part of the session setup path, or the session information is not collected. Currently, there is no architecture defined for the remote collection and forwarding of session information (other than accounting information) in an APPN environment. Therefore, if VTAM does not see the session, it (and NetView) will have no knowledge of it. However, some of the APPN devices provide you with commands that can query LU 6.2 sessions, links, etc., for that device (like the CM/2 DISPLAY command). If the device supports remote operations in an APPN environment, and it supports these types of display commands, then you still have the ability to do some level of remote session PD via remotely issuing the native commands from NetView - it just won't be available as part of Session Monitor.

OS/400 also has session level PD function for their sessions. The commands listed here (WRKPRB and ANZPRB) are not batch mode commands, so they cannot be invoked via the OPSMGMT command (remote operations). However, I did want to mention the functions here so that customers do not think that there is no session PD support on the AS/400 platform.

Second Page (C.21)

This page shows a Session Configuration screen for a session that has both subarea and APPN components. The additional APPN information on this screen is the APPN transmission priority, and the APPN COS. Also, this screen indicates that A12M can be acting either as an SSCP or as a CP.

The AR selection on the bottom of this display is also new and allows you to proceed to the APPN route information display (if any APPN route information has been collected for this session).

Third Page (C.22)

This page shows an APPN session route information display. This display contains the information that was present in the Route Services Control Vector (RSCV) for this session. This information shows which APPN nodes the session goes through. This display is similar to the ER display for subarea sessions, except that it shows APPN nodes and their connecting TGs.

The PAR and SAR selections on the bottom are to allow you to see other RSCVs if the path of this session was such that multiple RSCVs were used for the session. This is possible either in cross-network configurations, when the session traverses APPN Border nodes (and goes through multiple APPN subnetworks), or when the APPN session goes through one or more subarea nodes before going into another APPN network on the other side. See the Interpreting Session Data appendix in the NetView V2R4 Operations manual for more information about this type of situation.

Fourth Page (C.23)

This page is the VR display in the Session Monitor. The only change to it is the addition of OFC and DFC for selections on the bottom (the OFC/DFC are session specific selections, so they will only show up if you select VR from the session configuration screen). OFC is to request flow control information (for APPN TGs) from the origin node on this display, and DFC is to request flow control information (for APPN TGs) from the destination node on this display. Note, this will only retrieve data for APPN TGs, and not for subarea TGs.

Fifth Page (C.24)

This display shows the APPN TG flow control information from my NCP for the session I was looking at. Note, there is only secondary stage information present. This is because only in the direction from the NCP to the SLU is there an APPN TG off of this NCP. The TG in the direction of the PLU from this NCP is a subarea TG. If the TG in both directions was an APPN TG, then there would be flow control data here for both primary and secondary session stages (an additional column would be present for the primary session stage information). In such a case (where the APPN TG exists in both directions), this would be a VR within a single node. Therefore, since the origin and destination nodes would be the same, there would just be an FC selection (and not an OFC and a DFC) to retrieve the flow control data.

The data in this display shows you the most recent PIUs sent and received as well as pacing information for this TG.

Session PD Configuration (C.25)

This chart shows the configuration that is going to be used for the six session PD examples contained in this presentation. If you have the color Freelance chart, you can just show that and talk to the configuration.

For these examples, unless otherwise noted, assume the following:

- Both VTAMs are running VTAM V4R1 set up as an interchange node.
- Both NetViews are V2R4, and their Session Monitors are set up to communicate with each other.
- All three NCPs are at V6R2.
- VTAM 1 owns NCP 1, and VTAM 2 owns NCP 2 and NCP 3.
- There are two composite network nodes in this network - identified on the chart as CNN 1 and CNN 2, but in the real network the node name would be that of the VTAM CP/SSCP. Since NCP cannot be an APPN node on its own, the NCP along with its owning VTAM appear as a single APPN node to the APPN network, and the node is known by its VTAM name.

- All of the TGs between the VTAMs and NCPs are subarea TGs (FID4), with subarea VRs used to communicate over them.
- The TGs from the NCPs to the APPN devices (AS/400s, 3174s, and PCs) are all set up as APPN TGs.
- The PCs are all running CM/2 1.1.
- The AS/400s have OS/400 V2R2 on them.
- The 3174 is running the C3 level of microcode (this microcode provides the APPN NN function for 3174).

Session PD Example 1

If you have the Freelance color chart, you can just use that and talk to the points made for this example. In the Freelance chart, the session path is in red. In both charts (Freelance and LIST3820), the session path is marked by an asterisk (*) preceding the node names of the nodes that the session traverses.

First Page (C.26)

For the first example, assume that a CICS LU 6.2 session is going from PC 11 to CICS 1 which is running on VTAM 1.

The session will have an APPN component consisting of the path from PC 11 to the composite network node represented by VTAM 1/NCP 1 (CNN 1). Since NCP cannot be an APPN node on its own, the NCP along with its owning VTAM appear as a single APPN node to the APPN network, and the node is known by its VTAM name.

The session will also have a subarea component consisting of the ER/VR for the path from NCP 1 to VTAM 1/CICS 1.

Second Page (C.27)

If your operator is on NetView 1, then they will be able to see both the APPN route information (from the RSCV built when creating the APPN portion of the session route), and the ER information for the subarea portion of the session path.

In this configuration and for this session, VTAM 1 as the interchange node has the RSCV for the APPN route information. That is because the session goes from being an APPN session to a subarea session when it hits CNN 1 (which is VTAM 1). VTAM 1 also has the subarea session information since it owns both endpoints of the subarea portion of the session (NCP 1 and CICS 1).

If your operator is on NetView 2, then they will need to SDOMAIN to NetView 1 to see anything relating to this session. This is because VTAM 2 has nothing at all to do with this particular session.

Session PD Example 2

If you have the Freelance color chart, you can just use that and talk to the points made for this example. In the Freelance chart, the session path is in red. In both charts (Freelance and LIST3820), the session path is marked by an asterisk (*) preceding the node names of the nodes that the session traverses.

First Page (C.28)

For this example, assume that a CICS LU 6.2 session is going from PC 11 to CICS 2 which is running on VTAM 2.

The session will have an APPN component consisting of the path from PC 11 to the composite network node represented by VTAM 1/NCP 1 (CNN 1). Since NCP cannot be an APPN node on its own, the NCP along with its owning VTAM appear as a single APPN node to the APPN network, and the node is known by its VTAM name.

The session will also have a subarea component consisting of the ER/VR for the path from NCP 1 to VTAM 2/CICS 2.

Second Page (C.29)

If your operator is on NetView 1, then they will be able to see both the APPN route information (from the RSCV built when creating the APPN portion of the session route), and the ER information for the subarea portion of the session path.

Note: The Session Monitors in the two NetViews must be communicating with each other (via LUC sessions) in order for NetView 1 to get the complete information about the session (even the complete subarea information). This is because VTAM 1 knows about parts of the session, and VTAM 2 knows about other parts. The LUC communications between Session Monitors is established via CDRMDEF and CNMTARG definitions (see the NetView Administration Reference Manual for more information).

In this configuration and for this session, VTAM 1 as the interchange node has the RSCV for the APPN route information. That is because the session goes from being an APPN session to a subarea session when it hits CNN 1 (which is VTAM 1). VTAM 1 has part of the subarea session information since it owns one endpoint of the subarea portion of the session (NCP 1), and VTAM 2 has the other part of the subarea session information since it owns CICS 2.

If your operator is on NetView 2, then they should also be able to see the complete information (subarea and APPN), as long as the session monitors in NetView 1 and NetView 2 are communicating with each other.

Session PD Example 3

If you have the Freelance color chart, you can just use that and talk to the points made for this example. In the Freelance chart, the session path is in red. In both charts (Freelance and LIST3820), the session path is marked by an asterisk (*) preceding the node names of the nodes that the session traverses.

First Page (C.30)

For this example, assume that an LU 6.2 session is going from PC 10 (PLU) to AS/400 1 (SLU).

The session will have two APPN components consisting of the path from PC 10 to the composite network node represented by VTAM 1/NCP 1 (CNN 1), and consisting of the path from the composite network node represented by VTAM 2/NCP 2 (CNN 2) to AS/400 1. These two APPN components are represented by two separate RSCVs.

The session will also have a subarea component consisting of the ER/VR for the path from NCP 1 to NCP 2.

Second Page (C.31)

If your operator is on NetView 1, then they will be able to see the APPN route information for the session path. The APPN route information that will be presented from the AR screen will be the RSCV from PC 10 to CNN 1. Then, you will need to use the SAR to see the second RSCV for CNN 2 to AS/400 1 (since that RSCV is stored in NetView 2). The ER information for the subarea portion of the session path will also be present.

Note: The Session Monitors in the two NetViews must be communicating with each other (via LUC sessions) in order for NetView 1 to get the complete information about the session (even the complete subarea information). This is because VTAM 1 knows about parts of the session, and VTAM 2 knows about other parts. The LUC communications between Session Monitors is established via CDRMDEF and CNMTARG definitions (see the NetView Administration Reference Manual for more information).

In this configuration and for this session, VTAM 1 as the interchange node has the RSCV for the APPN route information for the path from PC 10 to CNN 1. That is because the session goes from being an APPN session to a subarea session when it hits CNN 1 (which is VTAM 1). VTAM 1 has part of the subarea session information since it owns one endpoint of the subarea portion of the session (NCP 1), and VTAM 2 has the other part of the subarea session information since it owns NCP 2. VTAM 2 as an interchange node has the other part of the APPN route information - namely the RSCV for the session path from CNN 2 to AS/400 1. That is because the session goes from being a subarea session (in NCP 2) to an APPN session when it leaves CNN 2 and proceeds to the AS/400.

If your operator is on NetView 2, then they should also be able to see the complete information (subarea and APPN), as long as the session monitors in NetView 1 and NetView 2 are communicating with each other.

Session PD Example 4 (C.32)

If you have the Freelance color chart, you can just use that and talk to the points made for this example. In the Freelance chart, the session path is in red. In both charts (Freelance and LIST3820), the session path is marked by an asterisk (*) preceding the node names of the nodes that the session traverses.

For this example, assume that an LU 6.2 session is going from PC 1 to PC 3. This is a pure APPN session that neither VTAM ever sees or knows about. Therefore, since VTAM never sees the session at all, it does not have any session awareness information that it can pass on to NetView. So the Session Monitor function in NetView will not know about this session.

In a real configuration, if you needed to find out more information about this session, you should use the remote operations functions in NetView and CM/2 and issue a RUNCMD (via LU 6.2) to display various pieces of information on that PC that may be helpful to performing session PD functions (display the active links, the LU definitions and partner LU definitions, the LU 6.2 sessions, the directory information, etc.).

Session PD Example 5

If you have the Freelance color chart, you can just use that and talk to the points made for this example. In the Freelance chart, the session path is in red. In both charts (Freelance and LIST3820), the session path is marked by an asterisk (*) preceding the node names of the nodes that the session traverses.

First Page (C.33)

For this example, assume that a CICS LU 6.2 session is going from PC 11 to CICS 1 which is running on VTAM 1. In this example, assume that the connection from NCP 1 to AS/400 2 is a LEN boundary function connection, and that VTAM 1 is set up as a subarea only node (or that it is running V3 VTAM which does not support APPN on the mainframe).

The session will have an APPN component consisting of the path from PC 11 to AS/400 2.

The session will also have a subarea component consisting of the ER/VR for the path from what VTAM views as the LEN node (AS/400 2) to CICS 1 on VTAM 1.

Second Page (C.34)

If your operator is on NetView 1, then they will only be able to see the ER information for the subarea portion of the session path. They will not see the APPN session route information, since VTAM when acting as a pure subarea node does not receive RSCVs.

If your operator is on NetView 2, then they will need to SDOMAIN to NetView 1 to see anything relating to this session. This is because VTAM 2 has nothing at all to do with this particular session. For this example, they will only see the subarea session information.

Session PD Example 6

If you have the Freelance color chart, you can just use that and talk to the points made for this example. In the Freelance chart, the session path is in red. In both charts (Freelance and LIST3820), the session path is marked by an asterisk (*) preceding the node names of the nodes that the session traverses.

First Page (C.35)

For this example, assume that an LU 6.2 session is going from PC 10 to PC 1. In this example, assume that the connection from NCP 2 to AS/400 1 is a LEN boundary function connection, and that VTAM 2 is set up as a subarea only node (or that it is running V3 VTAM which does not support APPN on the mainframe).

The session will have two APPN components consisting of the path from PC 10 to the composite network node represented by VTAM 1/NCP 1 (CNN 1), and consisting of the path from the AS/400 1 to PC 1. These two APPN components are represented by two separate RSCVs.

The session will also have a subarea component consisting of the ER/VR for the path from NCP 1 to NCP 2 (and really to what VTAM 2 thinks is a LEN node - AS/400 1).

Second Page (C.36)

If your operator is on NetView 1, then they will be able to see the APPN route information for the session path from PC 10 to CNN 1 (VTAM 1/NCP 1). The APPN route information that will be presented from the AR screen will be the RSCV from PC 10 to CNN 1. The APPN route information represented by the second RSCV (from AS/400 1 to PC 1) will not be able to be displayed since VTAM 2 as a subarea node will never receive that RSCV. The ER information for the subarea portion of the session path will also be present.

Note: The Session Monitors in the two NetViews must be communicating with each other (via LUC sessions) in order for NetView 1 to get the complete subarea information about the session. This is because VTAM 1 knows about parts of the session, and VTAM 2 knows about other parts. The LUC communications between Session Monitors is established via CDRMDEF and CNMTARG definitions (see the NetView Administration Reference Manual for more information).

In this configuration and for this session, VTAM 1 as the interchange node has the RSCV for the APPN route information for the path from PC 10 to CNN 1. That is because the session goes from being an APPN session to a subarea session when it hits CNN 1 (which is VTAM 1). VTAM 1 has part of the subarea session information since it owns one endpoint of the subarea portion of the session (NCP 1), and VTAM 2 has the other part of the subarea session information since it owns NCP 2. VTAM 2 does not have the other part of the APPN route information represented by the second RSCV since it is a subarea SSCP only (as explained above).

If your operator is on NetView 2, then they should also be able to see the complete subarea information as well as the partial APPN information, as long as the Session Monitors in NetView 1 and NetView 2 are communicating with each other.

Graphical display of an APPN network

First Page (C.37)

You can graphically display your APPN network from an NGMF workstation using NetView. However, in order to do this, you need to have the APPN Topology and Accounting Management (APPNTAM) Feature installed. The topology manager portion of the feature runs on top of NetView V2R4. Ordered with the APPNTAM feature is the APPNTAM topology agent. This agent is a CM/2 application that requires CM/2 1.1. Currently, the CM/2 application is the only agent available. However, IBM has announced that it intends to produce a VTAM topology agent.

The topology manager and topology agent communicate via CMIP (X.700) flows. See the Agent-Manager Communications chart for more information about the manager-agent flows.

The topology manager gives you the capability to let an operator monitor various APPN resources from an NGMF workstation. What you can monitor depends upon what type of topology information that you are collecting from your agents. If you are only monitoring network topology (which is the APPN routing topology - similar to the subarea backbone topology), then you will be able to monitor the network nodes and their interconnected TGs. You only need one agent active (that

is residing on a NN) in order to monitor your APPN network topology (assuming that you only have one APPN subnetwork). However, if your APPN subnetwork may be split when a TG goes down, you probably want to have one agent being monitored in each of the two sections. Also, for a larger (or more critical) APPN subnetwork, you may want more than one agent running so that you have more locations that you are collecting topology from (in case one of the agents goes down or becomes unreachable from NetView). Last, you are going to need an agent in every APPN subnetwork if you want to see the topology of that subnetwork.

If you want to see more than just your APPN routing network, then you need to collect local topology from your agents. Local topology lets you see any nodes that are adjacent to this node (including LEN nodes), as well as letting you see the ports and links (the real resources that make up the TGs).

Second Page (C.38)

You may also issue commands from NGMF to your APPNTAM agents.

Note: You need to have an agent actively running on the node in order to send commands down to it. From an NGMF workstation, an operator can issue a command for an APPN resource by selecting the resource and issuing the Activate, Inactivate, or Recycle commands from the command pull-down. Also, the Build commands function can be used in order to pull up a CT/2 command tree that will let you do the appropriate actions for that resource in addition to allowing you to issue a start or stop monitoring command for the APPNTAM agent (you can start/stop monitoring local or network topology).

The resources that you can issue commands against are listed in the chart. In addition, you can select a TG circuit and issue a start, stop, or recycle command against it (assuming you are collecting local topology for that node). APPNTAM will use RODM to determine the proper underlying resources, and it will issue the command against that resource (the link).

Note: There are no commands to recycle a node - that is because in APPN, there is no such thing as activating a node. If the software/hardware that gives it the capability to act as an APPN node is up, and there is a link active to that node, then that node is available to the APPN network. There is no concept of APPN node activation (vs. the concept of activating a node from an SSCP in the subarea network).

Now, all of the topology mentioned above is stored in RODM. For a customer to use the topology management function, they need to just load a SystemView consistent data model into RODM (which APPNTAM provides and uses), and start collecting topology from the agents. APPNTAM will dynamically add the appropriate resources and connections to RODM (using the SystemView consistent data model) automatically - the customer will not have to manually load RODM with the various APPN resources in their network.

Third Page (C.39)

This chart is an example of what the APPN network topology looks like for my test network on an NGMF workstation. The nodes depicted in the diagram are either network nodes, or virtual routing nodes. Virtual routing nodes are connection networks (in our network some of the nodes have the LAN defined as a connection network).

Fourth Page (C.40)

This chart is an example of what the domain topology looks like for an APPN domain. The resources that will show up in this display are the NN from the prior chart that you asked for more detail about (more detail on a NN domain aggregate gives you the domain topology), the links from that NN, and any served ENs of that NN (since the ENs contribute to the aggregation status of their serving NN′s domain aggregate). Please note, if local topology has not been collected at some point in time, then APPNTAM will have no way of knowing about the served ENs from a NN, and they will not show up on this display.

Fifth Page (C.41)

This chart (and the two after it) are all based on local topology being collected from the agent (A19SG2CP in this case). This first chart is the more detail view of the local topology node. If you start with the prior chart (Fourth Page), and double click with the left mouse button (to get the more detail function), you will get two views up (assuming local topology is being collected). You will get a view with just a gray square for the local topology (it is present just for view navigation purposes), and the ports and links view (next chart). If you select the local topology node (gray square) and ask for more detail, you will get this display.

This display contains the node itself, as well as any adjacent node - whether it is a NN node, EN node, LEN node, or even a virtual routing node (connection network). It also shows the connections (TG circuits) between the nodes. All of this information is based off of the local topology database (vs. the network topology database that is shared among all NNs) that is present on the node that we are collecting local topology from.

Sixth Page (C.42)

This chart is the ports and links view for a node. It is gotten by selecting a node from a domain topology view, and by asking for more detail of that node. This view shows the ports (adapters) that are present on that node (and are being used for APPN), and the links that exist for those adapters. Dynamically defined links are shown too (NETA.A19SG2CP.@AAAAAAA in this view is a dynamic link).

Again, this information is not available unless you are collecting local topology from the node.

Seventh Page (C.43)

You can also issue commands from your NGMF workstation against the APPN resources. If you follow the installation procedures documented in the APPNTAM Feature Implementation Guide, then the generic Activate, Inactivate, and Recycle commands will work for GMFHS resources as well as for APPNTAM resources. This chart shows an example of a link being selected, and the generic activate command being used for that link to try to activate the link. Again, the node must have an agent active on it at the time the command is issued for the command to be successfully sent and received at the node.

APPC Accounting Collection (C.44)

NetView can collect accounting information for APPC (LU 6.2) connections from two different sources.

1. If you wish to just collect accounting information for LU 6.2 sessions, and the sessions go through or end in a VTAM (i.e., VTAM knows about the session), then the Session Monitor can collect accounting data for that session in the same manner that it collects accounting information for LU 0, LU 2, etc., sessions today. This is not a new function (even for NetView V2R2).
2. If you wish to collect conversation accounting information, or if you wish to collect accounting information for a session that VTAM has no knowledge of, then you will need to use the APPNTAM Accounting Manager function and collect the accounting information from an APPNTAM accounting agent. The APPNTAM accounting manager runs on NetView V2R4, and it can run on a central or distributed NetView. The APPNTAM accounting agent runs as a CM/2 1.1 application.

   Note: You can order the APPNTAM Accounting Management Feature without the Topology Management feature if you wish to do accounting and are not interested in using the topology manager features mentioned above.

The rest of the accounting information in this presentation will concentrate on the APPNTAM Accounting function, since that is a new function for APPN, while the Session Monitor accounting function has been present since the NLDM product (before NetView).

APPNTAM Accounting Function

First Page (C.45)

You can remotely collect APPC accounting information from your APPN network using the APPN Accounting Management Feature. The accounting manager runs on top of NetView V2R4. Ordered with the APPNTAM feature is the APPNTAM agent. The accounting agent is a CM/2 application that requires CM/2 1.1. The CM/2 accounting agent application is the only agent available.

The accounting manager and accounting agent communicate via CMIP (X.700) flows. See the Agent-Manager Communications chart for more information about the manager-agent flows.

The APPNTAM Accounting Feature lets you collect both session and conversation level accounting information for APPC (LU 6.2) connections. The accounting data is collected and kept at the agent (on a PC). The manager has to issue a request for it to be sent to the manager. When the manager (NetView) receives it, it writes it to the SMF log (or using the external logging function in NetView, it can be sent to another log).

The ability to write these accounting records to SMF (using the same record type, just a different record subtype, as the Session Monitor accounting records) lets you collect all of your accounting information in the same SMF database. Your NetView that collects (and writes to SMF) subarea accounting information can also

be the accounting manager and collect and write the APPC accounting information to SMF.

Second Page (C.46)

Session level accounting can be collected from an agent that is installed and running either on a node at the session endpoint (one or both of the session endpoints), or you can collect session level accounting from a node that is performing intermediate routing for the session (the agent is running on a node that is part of the session path). This gives you a great deal of flexibility in collecting APPC session data, since you do not have to collect it from the session endpoint.

The chart then mentions some of the various pieces of information that are collected for APPC session accounting. Please note, if customers are processing their accounting data, the piece of information in the session level accounting record that should be used to eliminate duplicate information is the fully qualified PCID for the session, as that is the unique identifier for the session. You can not simply use the PLU and SLU (even with the start time), as APPC supports parallel sessions.

Third Page (C.47)

Conversation level accounting may only be collected from an agent that is installed and running at either end of the conversation (and session since they are the same place) endpoints. Conversation level accounting cannot be collected from a node that is performing intermediate routing for the session. This is because intermediate routing nodes only see the session, they do not notice conversations (since it is not necessary in order to perform intermediate routing and would add path length and overhead).

The chart then mentions some of the various pieces of information that are collected for APPC conversation accounting. Please note, if customers are processing their accounting data, the pieces of information in the conversation level accounting record that should be used to eliminate duplicate information are the fully qualified PCID for the session plus the conversation correlator. Both, are needed to uniquely identify any given conversation. This is because the fully qualified PCID uniquely identifies a session, and the conversation correlator uniquely identifies a conversation between any two LUs. You can′t just use one of these, as then you may eliminate records that are not really duplicates of each other.

There are a few other pieces of data that are included in the APPNTAM accounting records, but I did not feel that they were important enough to list in these two charts. For the complete accounting records, see the APPNTAM Feature Implementation Guide.

Fourth Page (C.48)

An agent has to be told to start collecting accounting records. It does not automatically do so. An agent can be told to collect accounting records in a couple of different ways. A command can be issued at the manager to start collecting accounting data from the agent - this will result in a flow down to the agent to start collecting accounting data. Also, you can specify for the accounting agent to automatically start collecting accounting data (so nothing is lost) via an initialization file on the agent (BNTACCT.INI). Last, if the accounting agent is being restarted (it was active but was terminated), then the accounting agent will automatically start collecting any accounting data that it was collecting when it was terminated.

The accounting data will stay on the agent until the manager sends a flow to the agent asking the agent to send it to the manager. The accounting data will also remain on the agent until it receives an acknowledgement that the data has been received by the manager. This minimizes the chance of losing accounting information due to problems in transmitting the data to the manager. If more than one manager is collecting a particular type of accounting data, then the agent will not delete that accounting data until all managers have acknowledged receipt of it. The only exception to this is if the accounting buffers fill up and you have told the agent to wrap the accounting data when such a condition occurs (see the next few charts).

Fifth Page (C.49)

If you want to start collecting accounting information from an agent, you can issue the ACCTSNA START command on the NetView that is acting as the accounting manager for the APPNTAM accounting (of course the APPNTAM accounting feature needs to be installed also). You must issue one ACCTSNA START command per agent that you want to collect data from, and per each type of data that you want collected from that agent (conversation, endpoint session, and intermediate session data). So if you wanted to collect all session accounting data from a node (endpoint and intermediate), you would issue two ACCTSNA START commands to that agent.

In the START command you can specify various parameters. You can tell the manager what times of day or at what time intervals (or both) that you want to collect the accounting data from the agent. You can specify either at what threshold or after what number of accounting records (or both) the agent should send a notification to the manager that it hit the threshold/number of buffers and the manager should collect the accounting data. So you could specify that after 1000 records, or if the buffer is 80% full, the agent should send a notification to the manager so that the manager will retrieve the accounting data. Please note that the thresholds only count the number of closed accounting records (terminated sessions or conversations). They do not include the number of records still open (active sessions or conversations). However, the active records will still be sent to the manager when the manager requests the accounting data from the agent - it is just that their records will indicate that they are still active vs. being complete.

Sixth Page (C.50)

Another parameter that can be specified on the ACCTSNA START command is how many records should the manager receive before sending an acknowledgement (and thereby letting the agent delete those records) to the agent. Also, you can specify what action the agent should take if its accounting buffers fill up before the accounting data retrieval has completed (and been acknowledged). You can have the agent either stop recording until the buffers are cleared, or you can have it start wrapping the accounting data.

When the manager gets a notification from the agent that the agent has hit a threshold or collected a number of buffers (as specified in the ACCTSNA START command), it will send a flow down to the agent to retrieve the accounting data. Also, when a time interval or if the time of day specified in the ACCTSNA START command passes, the manager will also generate a flow to the agent to retrieve the accounting data. You can also tell the manager to retrieve the data immediately via the ACCTSNA RETRIEVE command. This is useful if you are going to take the manager down or break the connections between the manager and agent (this way you collect the latest data and will hopefully get things back up and running before the agent fills its buffers).

You can stop the collection of accounting data via the ACCTSNA STOP command. Using this command you can specify whether accounting data should be retrieved one last time or not.

Seventh Page (C.51)

Where you place your accounting agents depends a great deal on what type of accounting data you are trying to collect, and for what purposes. The kinds of questions that customers need to decide before placing their accounting agents are: Whether session accounting data will be sufficient, or do I need conversation accounting data? If I need just session accounting data, am I trying to monitor just access to a file server or database server for capacity purposes, or am I trying to monitor the use of the network by my end users? Do all of the sessions that I want to monitor go through a common point (like a PC acting as a LAN gateway either to the company network or to a service provider)? The reasons that these questions are important is because the answers help (along with the customer's network topology) the customer determine if they can collect all the data that they wish to, and if so, where do they need to collect it from.

Use the diagram at the bottom of the chart (or the Freelance color chart if you have been able to print the .PRS presentation off) to help explain these concepts. For instance, in that chart, the CM/2 PC attached to the APPN network cloud might be a LAN gateway PC to a service provider or a corporate network. If you want to monitor usage for that (for LU 6.2 sessions - remember, LU 2 sessions can still be monitored at the host where the application endpoint is via the Session Monitor function), then you could collect intermediate LU 6.2 session accounting data from an agent running on that PC. If you just wanted to monitor usage of a database server, you could collect endpoint session data from the DB Server PC (bottom of the picture). If you want to monitor all of your users' traffic, then you would possibly/probably need agents running on all of the PCs (and this could be quite a bit of data to collect!).

Agent-Manager Communications (C.52)

This chart discusses how the managers and agents (both topology and accounting) for APPNTAM communicate.

The data between the managers and the agents is in CMIP (X.700) format. However, the transport architecture between the two is not OSI, it is SNA (MDS transport). If you wish to see more details of the CMIP implementation and CMIP over SNA (CMOS), see the APPNTAM documents and the SystemView documents listed in the part "Further Information".

Also, the APPNTAM Feature conforms to the OMNIPoints 1 standard. See the NVOMNIPT package on MKTTOOLS for information on how the APPNTAM Feature conforms to the OMNIPoints 1 standard.

Another important point is that many managers can be collecting (topology and/or accounting) data from a single agent at the same time.

Part D - Summary

Summary (D.2)

This chart is a short summary of what you have been able to (hopefully) cover with your audience. It recaps the functions that have been covered in this presentation. Use it to wrap up the session. Again, if you are giving the long version (and doing all of the examples in Part E, etc.) move this to the end of the presentation.

While the APPN management functions may not be exactly at the same level as the subarea management functions are today, there are a few important points to note. First, APPN is not exactly the same as subarea, and some functions that are present in subarea from a management perspective may not be necessary, or may need to work differently (peer vs. hierarchical). Second, the subarea architecture (and its network management architecture) has been around a very long time - much longer than the corresponding APPN architectures. And last, NetView can provide a lot more management in an APPN environment than most people realize. This is especially true with the latest releases of the product.

Part E - Remote Operations Examples

If you are using these examples in your presentation, you may want to first check who in the audience has S/370s or S/390s (hopefully almost everyone, since this is a NetView presentation!), who has AS/400s, who has PCs that can run CM/2 (even if they don't have CM/2), and who has RS/6000s. This way, you can decide which sets of examples to use, and which sets will not be of any interest.

Another important point about the remote operations examples. While the examples show the use of these commands in operator displays, a large advantage of these functions is in being able to automate these remote operations. NetView PIPEs is a powerful extension to NetView's REXX/automation functions and can be used with the RMTCMD, OPSMGMT, and RUNCMD as documented below. For more information on the NetView PIPEs function, see the NetView product documentation and the NETVPIPE package on MKTTOOLS (which contains examples).

Remote Operations Examples - VTAM

First Page (E.2)

In order to send VTAM (and MVS, VM, or VSE) commands to another node, you can use the NetView RMTCMD function to send these commands from one system to another via LU 6.2 (High Performance Transport). The only requirement is that VTAM and NetView (V2R2 or later) are active on both nodes, and that an LU 6.2 session can be established between the two nodes.

The first command that is shown here as an example is displaying the VTAM topology database for the VTAM, and listing the nodes that are adjacent to this VTAM. Please remember, since this uses RMTCMD, you can perform these commands in PIPEs, or you can put the RMTCMD in a WINDOW command so that you get the results in a full-screen view window.

The display here shows the list of adjacent nodes, their nodetype, the route resistance of routing through that node (used for APPN route selection processing), whether the node is congested (again, used for APPN route selection), and whether the link to them supports CP-CP sessions.

Second Page (E.3)

The next display shows what's in VTAM's directory database concerning itself (its own node).

Third Page (E.4)

This example shows what's in VTAM's directory database for some other node (a PC in this case). It shows that under this CP is a dynamic LU. If I wanted to modify the contents of VTAM's directory database I could do so through the F procname,DIRECTRY command. I could also use this command to delete information from the directory.

Fourth Page (E.5)

I can also remotely start a checkpoint of VTAM's topology database and/or directory database. This is useful when you need to recycle VTAM, it will come back up with a checkpoint database and does not need to start learning things from scratch. Of course, if VTAM is going to be down for a long period of time, you probably want to delete the checkpoint datasets (or start VTAM without them), since everything will probably need to be relearned anyway (unless your network is very stable).

Fifth Page (E.6)

You may also use this technique to dynamically modify how a VTAM is running (by modifying its start options). In this example I am setting the VTAM start options so that connections will by default support CP-CP sessions, and that connections to other PU type 2.1 nodes are APPN connections by default.

Remote Operations Examples - AS/400

First Page (E.7)

In order to send AS/400 commands to from NetView to an AS/400 (without logging on to the AS/400), I can use the OPSMGMT command supported by the NetView Remote Operations Manager. This command is sent via the MDS transport (LU 6.2). It requires that NetView V2R3 with the GMFHS Feature (or higher) is running on the host, and that OS/400 V2R2 (or higher) and the NetView Remote Operations Agent/400 is running on the AS/400.

The OPSMGMT command can be issued inside of a PIPE, or you can also use the FNAWIND command (NetView Remote Operations Manager version of the WINDOW command) to issue the OPSMGMT command and get the response back in a full-screen view panel. I would recommend the latter (or looking at the response in the log), as extra information is returned in the response for automation purposes, but which most people will not want to see. You can also issue OPSMGMT commands via CT/2, either from CT/2 itself, or from an NGMF view that has the AS/400 in it. See the NetView Remote Operations Manager User's Guide for instructions on setting these functions up.

In order to use OPSMGMT in a PIPE, from CT/2 within NGMF, or following an FNAWIND command, you will need NetView APAR OY66689 if you are running NetView V2R3.

Even though these examples will show the commands being issued by an operator, this function allows you to remotely control your AS/400s via automation, and to automate responses to commands.

Second Page (E.8)

AS/400 has a very large list of commands (3 volumes worth!). The majority of these commands can be issued in batch mode, and therefore, you can use them with the OPSMGMT command. Over the next few pages I have listed some commands that I have found of interest in reviewing the available commands. Most of these commands listed are used for network management, although I have also listed some commands that would be useful in systems management.

Go through the commands listed on this page.

Third Page (E.9)

Go through the commands listed on this page.

Fourth Page (E.10)

Go through the commands listed on this page.

Fifth Page (E.11)

Go through the commands listed on this page.

Following this page are examples of some AS/400 commands being issued (from a NetView command line). For these examples, some of the responses will be shown as they come back to the NetView console, and for some of the examples I will issue the OPSMGMT command as part of an FNAWIND command, so the response will be in a full-screen view panel. In some cases (due to the volume of data being returned), I will omit some of the data and just show parts of the response that are of interest.

The examples for the AS/400 commands were performed on a NetView V2R4 system, but they were done from a different test system (and different NetView) than all of the other examples in this presentation. This information is for those of you that wonder why the resource names and NetView name are different in these examples.

Pages 6-8 (E.12-E.14)

This is an example of the DSPNETA (display network attributes) command on the AS/400, with the response being returned to the NetView operator's screen. It shows items like the system name, CP name and NETID, where alerts should be forwarded, etc. Pick out a few points that may be of interest to your audience. Given the examples that follow, make sure to point out that there is no backup alert focal point currently defined (ALRBCKFP on lower screen of E.13).

Ninth Page (E.15)

In this example I went ahead and changed the network attributes of the AS/400. I set the backup alert focal point for the AS/400 to be another NetView in the same network as the AS/400.

Tenth Page (E.16)

This page shows the display network attributes being issued again, with just a subset of the returned information. In the shaded section is the backup alert focal point information, which has changed due to the CHGNETA command that was issued above.

Pages 11-15 (E.17-E.21)

This example shows the DSPAPPNINF (display APPN information) command being issued. Since no parameters are specified, this command defaults to displaying the complete topology database for all of the nodes in the topology database of the AS/400. The response to this command will be a list of all of the nodes in the topology database, plus, for each node, what links do we know about for that node (originating in that node), what node the link connects to, and what are the characteristics of each of those links.

This example just has a subset of the information present in it - after the first node and link characteristic displayed for a link from that node, I only displayed information pertaining to this AS/400 itself (CON403), or to the VTAM that my NetView is running on (NEM01). The command returned information for both active and inactive resources.

Pages 16-17 (E.22-E.23)

This example shows the results of a DSPAPPNINF INFTYPE(*LCLNODE) being issued. This command will retrieve the local directory of the AS/400. So, for instance (on page E.23, upper screen), it shows the directory entry for NEM01 (the VTAM my NetView is running on). This directory entry has the VTAM CP itself, and the NetView LU (since the OPSMGMT commands are going (in this case) from the NetView LU to the AS/400 directly).

Eighteenth Page (E.24)

This example shows a display of the AS/400′s sphere of control (AS/400 only supports the alert category as a focal point, so the category is not specified on the DSPSOCSTS command). Currently, there are no nodes in the AS/400′s sphere of control.

Nineteenth Page (E.25)

I then used the ADDSOCE (add sphere of control entry) command to add two entry points to the AS/400′s sphere of control. I added NE1PU08 and NE1PU09 (both with the same NETID as the AS/400) to the AS/400′s SOC.

Even though the resource names end in ′PU′, both resources are control point names (CPs).

Twentieth Page (E.26)

In this example, I issued another display for the sphere of control to show that the two control points were added to the sphere of control. Since no alerts have been forwarded from those control points to this AS/400, their current status is listed as never active. Once an alert is sent from those nodes to the AS/400, the status changes to active.

Twenty-first Page (E.27)

I then removed the two control points from the sphere of control using the RMVSOCE command. This resulted in the deletion of those nodes from the sphere of control table in the AS/400.

Twenty-second Page (E.28)

I then displayed the sphere of control table again to show the removal of the two nodes. Since both nodes had a state of never active when they were removed from the sphere of control table, they do not appear in the results of the DSPSOCSTS command. If the control points were in an active state, then they would still be present in this table (unless they had sent the AS/400 a focal point revocation flow for any reason - such as being acquired by another focal point). Their status would be displayed as remove pending to show that they had been removed but that the AS/400 was still providing focal point services to them.

Pages 23-27 (E.29-E.33)

These pages show the results of a DSPLCLHDW command, which displays the list of hardware resources for the AS/400 (including serial numbers). This information can be used in asset management or in inventory reports.

Pages 28-29 (E.34-E.35)

These pages show a subset of what is returned for a DSPSFWRSC command. This command displays the software resources on the AS/400, including any options and features. Page 28 (E.34) shows what the response looks like when it comes back to the NetView operator's screen, and page 29 (E.35) shows what it would look like in the full screen view panel that is returned when using FNAWIND. Of interest on these two pages is the level of the OS/400 operating system (V2R3), and the level of the NetView Remote Operations Agent/400 code (V1R1).

Remote Operations Examples - RS/6000

First Page (E.36)

Using the NetView RUNCMD command, you can send commands to a RISC machine via LU 6.2 sessions (MDS transport). NetView V2R3 on the host, and AIX NetView Service Point V1R2 and NetView/6000 on the RISC machine are required. In addition, if you want your RISC machine to be an APPN node, you must have AIX SNA Server/6000 V2R1 installed. AIX SNA Server/6000 also provides sna -d (SNA display) commands similar to some of the CM/2 DISPLAY commands. A fix for AIX NetView Service Point V1R2 is required if you are going to run it with AIX SNA Server/6000 V2R1.

Since AIX SystemView NetView/6000 will also be running on that RISC machine, you can use the RUNCMD to kick off NetView/6000 commands or shell scripts (such as rping to perform ping commands, findroute to find the routing between two IP nodes, etc.).

Second Page (E.37)

The RUNCMD command can be issued inside of a PIPE, or you can use the WINDOW command to issue the RUNCMD command and get the response back in a full-screen view panel.

NetView V2R4 added the ability to use NETVASIS as a command prefix from the NetView command line. This tells NetView not to automatically convert what was entered on the command line into upper case. In prior releases, you had to issue the command from a REXX exec after setting ADDRESS NETVASIS if you wanted to issue a command using mixed case. This is important as AIX is case sensitive.

When the NetView/6000 RUNCMD receiver daemon (spappld) receives a RUNCMD, it automatically translates all the characters to lower case. This is because it realizes that NetView usually sends RUNCMDs in upper case, and most AIX commands are in lower case. However, some commands (or their operands) need to be in mixed case, and NetView now has the ability to support mixed case commands (see above). If spappld sees that the first four characters of the command are 'asis' (in lower case), it will remove those characters from the command, and it will know not to translate the rest of the command into lower

case. It will try to invoke the command as it was submitted. The use of NETVASIS in combination with 'asis' allows you to enter mixed case commands, without NetView/6000 converting the entire command into lower case.

Note: When using 'asis', you should not leave a blank between 'asis' and the start of the actual AIX command.

Third Page (E.38)

There are a number of various SNA Server/6000, AIX SystemView NetView/6000, and AIX commands that may be of interest to customers in performing network and systems management. Over the next few pages I have listed just a few of these commands that may be of interest. Most of these commands are used for network management, although I have also listed some commands that would be useful in systems management.

Go through the commands listed on this page. These are all SNA Server/6000 commands.

Fourth Page (E.39)

Go through the commands listed on this page. These are all SNA Server/6000 commands.

Fifth Page (E.40)

Go through the commands listed on this page.

Sixth Page (E.41)

Following this page are examples of some RS/6000 commands being issued (from a NetView command line). For these examples, some of the commands will be issued using the WINDOW command, and some will be issued from the NetView console.

Currently (as of 1/94, and using NetView V2R4), the WINDOW command converts all input to upper case - even if it is prefixed by NETVASIS. Therefore, for commands or operands that needed to be entered in mixed case (using NETVASIS and the 'asis' prefix), I could not use WINDOW. Development is looking into removing this restriction.

In these examples, the RS/6000 is set up as a LEN node. This is due to the fact that I did not have the AIX NetView Service Point fix that will allow it to work with AIX SNA Server/6000 V2R1. A19RSCCP is set up as an independent LU, so we are using LU 6.2 to send these commands. And you can attach your RS/6000s as LEN nodes in your APPN network. It just means that you need to predefine to them who they are going to talk to (my NetView in this case), and you need to predefine to its NN what LUs are on the RS/6000.

AIX SNA Server/6000 V2R1 provides the full APPN support for the RS/6000. When AIX NetView Service Point is enhanced to work with SNA Server/6000, then you will be able to attach your RS/6000 as an APPN node and to send the RUNCMDs over LU 6.2 sessions. Currently you can run AIX NetView Service Point on SNA Server/6000, but you need an SSCP-PU session to communicate with NetView (for alert flows and RUNCMD).

Since AIX SNA Server/6000 V2R1 was not installed, I did not show any examples of the sna display commands mentioned above.

Seventh Page (E.42)

For the first command, I issued an ls /usr/lpp/nvix which looked at what was in that directory (similar to OS/2 dir command).

Please note that the appl= parameter in my RUNCMD is NV6000. That matches the application name that I set up when I configured spappld in the NetView/6000 set-up. The spname= parameter is A19RSCCP, which is my independent LU for that RS/6000. AIX NetView Service Point is configured to use that LU for its MDS support.

Eighth Page (E.43)

This example is an rping to pamb.raleigh.ibm.com. Since no source node was provided, the ping originated at our NetView/6000′s node (codebust.raleigh.ibm.com). You could have had the ping originate at another node by supplying a source node in addition to the destination node for the ping.

Ninth Page (E.44)

This example is a findroute between codebust.raleigh.ibm.com to pamb.raleigh.ibm.com. This returns the first route found between the two nodes, and the addresses of all the nodes in the route. In this case, both nodes where within the same IP subnet (they could communicate without first going through a router).

Tenth Page (E.45)

This example is an nmdemandpoll for mercury.raleigh.ibm.com, which is an RS/6000 acting as a router. The nmdemandpoll gets the status of the network interfaces plus their IP addresses and their polling intervals.

Eleventh Page (E.46)

I issued an snmpget for the router 6611apn7 (6611apn7.raleigh.ibm.com) to retrieve various system MIB variables. snmpget defaults to using .iso.org.dod.internet.mgmt.mib-2. as the default prefix for retrieving MIB variables. In this case, I am querying the system name, up time, description, contact, and location. Note, only 6611apn7 was defined as the destination node. The rest of the address defaulted to be the same as the local node (.raleigh.ibm.com).

I issued this from the command line using NETVASIS and the ′asis′ prefix since the MIB variable names are in mixed case. Otherwise, this command would not work.

Twelfth Page (E.47)

I issued an snmpget to determine the number of interfaces present in the router.

Thirteenth Page (E.48)

I then issued an snmpget to get more information on the fourth network interface for that router. I retrieved the speed of the interface, the physical address (MACADDR in this case), the operational state, and the number of packets being sent into the router that were not able to be routed to a higher layer protocol because of errors. These packets were discarded.

Fourteenth Page (E.49)

I then issued an snmpget for the router to get the total number of IP datagrams that were requested to be sent outbound from the router, and the number of IP datagrams that were discarded because no route could be found to deliver them to their destination.

Fifteenth Page (E.50)

I issued a rnetstat -i command to the router to display the network interfaces present on the router, and the attributes and statistics associated with those interfaces.

Sixteenth Page (E.51)

I issued a rnetstat -A to the router to display the router's cache entries, including the link address mappings.

Seventeenth Page (E.52)

I issued a rnetstat -r to the router to display its IP routing table. The display includes the destination address, the next node in the route, and whether the routing is direct or remote.

Remote Operations Examples - CM/2

First Page (E.53)

Using the NetView RUNCMD command, you can send commands to a PC over LU 6.2 sessions (MDS transport). NetView V2R3 on the host, and CM/2 1.0 (or later) on the PC are required. In addition, SPA Router and Remote Operations (parts of CM/2) need to be up and running for this function to work.

The RUNCMD command can be issued inside of a PIPE, or you can also use the WINDOW command to issue the RUNCMD command and get the response back in a full-screen view panel.

Any CM/2 command or any other OS/2 command that can be issued in line mode (via an OS/2 window or the OS/2 full-screen) and which returns just text output (i.e., it doesn't result in a presentation manager display) can be sent to the PC using this method (RUNCMD over LU 6.2).

The first example (which starts on this page and continues over the next four pages) shows the display of what is in the APPN topology database for this PC (CM/2).

Pages 2-5 (E.54-E.57)

These pages are just more of the information in the topology database on the PC. They are a continuation of the example from the first page of the CM/2 examples.

Sixth Page (E.58)

This page shows what the previous example would look like if you prefixed the RUNCMD with WINDOW (i.e., issued the RUNCMD in the WINDOW command). The response is a full-screen view window that you can scroll. You may also re-issue the command simply by using PF2.

Pages 7-8 (E.59-E.60)

These pages are an example of displaying the CM/2 directory database (vs the topology database in the prior example).

Pages 9-10 (E.61-E.62)

These pages are an example of displaying the management services information in this CM/2 node. This shows the focal points for this node, what type of focal points they are, what the status is, and other information about the focal points. It will also show any active MS transactions that are outstanding and have not been completed yet.

Pages 11-12 (E.63-E.64)

These pages are an example of displaying the LU 6.2 LUs on this CM/2 node. This shows what partner LUs are present, what logmodes are being used for the LU, the max RU size, the session limits, and the numbers of active and pending sessions.

Thirteenth Page (E.65)

This is an example of displaying the active links for the PC. It contains information such as whether it supports CP-CP sessions over that link, the partner node name and what type of node it is, its capacity, etc.

Fourteenth Page (E.66)

This page lists some other types of functions that you can perform on a PC remotely from NetView. You can cause files to be uploaded or downloaded (using TSO, VM or CICS file transfer facilities) - this is separate from the functions present in NetView Distribution Manager.

You can start or stop CM/2 traces. You can change the default configuration name for CM/2 - this controls what set of definitions CM/2 starts up with. You can also use the CM/2 DISPLAY command to display other pieces of information.

Also, as mentioned in the beginning of this set of examples, you can issue any other command (.EXE, .BAT, .CMD, etc.) that may be issued from an OS/2 command line and that will not result in a presentation manager display for output. This includes commands like DIR, PATH, DPATH, SET PATH, etc.

Part F - NetView Function/Release Information

The purpose of this part of the presentation is to have one place in the presentation where a customer can look for the function that they want to use and see what products and what releases of NetView they need in order to use that function.

The release information below is the minimum release needed to perform the specified function. Unless otherwise stated, the specified function is supported by any follow-on releases that were known about as of 1/94. For example, when NetView V2R2 is mentioned as supporting a function, that function is also supported by later releases of NetView (such as V2R3 and V2R4) unless otherwise stated.

NetView Function/Release Information (F.2-F.4)

- For transporting management data
  - **MDS Transport** - if you want to use the MDS transport architecture you need NetView V2R2. In addition, OS/400 V2R2, CM/2 1.0, VTAM V4R1, AIX NetView Service Point V1R2 (and AIX SNA Server/6000 V2R1), 6611 (w/MPNP V1R2), and other nodes all support this transport.
  - **Domain Focal Point** - If you are running NetView on a VTAM configured as a NN, and there are served ENs off of that VTAM, then you may need to be running NetView V2R4 with the enhanced MS transport function turned on (if you want a couple of very specific functions). If you want to have the ENs notified of who their domain focal point is, then you need to have NetView V2R4. Also, if you want data sent by the ENs to be routed to their domain focal point, you also need NetView V2R4 with the enhanced MS transport.
  - **Migration (APPN) Node support** - If you have OS/400 V2R1 or S/36s with APPN installed on them, and you want to receive alerts from them via LU 6.2, then you need to use NetView V2R4 with the enhanced MS transport. This is because the VTAM MDS implementation also supports the Migration (APPN) Node support that is implemented in OS/400 V2R1 and the S/36s with the APPN feature.
  - **High Performance Transport** - NetView V2R2 is required for this.
- Focal Point and Sphere of Control functions
  - **FOCALPT CHANGE command** - used to become the focal point for a specified entry point. This was provided in NetView V2R2.
  - **FOCALPT ACQUIRE command** - used to acquire a focal point (for when NetView is an entry point or nested focal point). This was provided in NetView V2R2, and expanded in NetView V2R4 to handle multiple backup focal points.
  - **FOCALPT QUERY command** - used to display who your focal point is. This was provided in NetView V2R2 and expanded in NetView V2R4 to handle multiple backup focal points and APPN domain focal points.
  - **FOCALPT DELETE command** - used to drop a focal point (stop reporting to it or not use it as a backup focal point). This was provided in NetView V2R2 and expanded in NetView V2R4 to handle multiple backup focal points.

  - **FOCALPT DISPSOC command** - used to display your sphere of control. This was provided in NetView V2R4.
  - **FOCALPT DELETE command** - used to delete an entry point from your sphere of control. This was provided in NetView V2R4.
  - **FOCALPT REFRESH command** - used to dynamically reload or refresh the sphere of control table. This was provided in NetView V2R4.
  - **Nested Focal Point support** - NetView V2R4 provided the ability to support NetView being a nested focal point (i.e., being a focal point and having a focal point also; or, in other words, being an intermediate focal point).
  - **DEFFOCPT definition** - used to define who your focal points should be. This was provided in NetView V2R2, and expanded in NetView V2R4 to handle multiple backup focal points and user defined categories.
  - **DSI6SCF** - this is the sphere of control configuration file used to determine what entry points should be in our sphere of control. This was provided in NetView V2R4.

- **Alert focal point support** - NetView V2R2 supported being an alert focal point via the MDS transport, and it also supported automating alerts.
- Remote Operations
  - **Remote Operations to VTAM and S/390** - NetView V2R2 is required for the RMTCMD function.
  - **Remote Operations to AS/400** - NetView V2R3 with the GMFHS Feature, the NetView Remote Operations Manager, OS/400 V2R2, and the NetView Remote Operations Agent/400 are required.
  - **Remote Operations of RS/6000** - NetView V2R3, AIX NetView Service Point V1R2, and AIX SystemView NetView/6000 are required. In addition, if you want the RS/6000 to be an APPN node, AIX SNA Server/6000 V2R1 is required.

    A fix is necessary if you would like to use AIX NetView Service Point and RUNCMDs over LU 6.2 with AIX SNA Server/6000 V2R1. At this time I do not know the method or time frame in which the fix will be available.
  - **Remote Operations to PCs** - NetView V2R3 and CM/2 1.0 are required. In addition, the SPA Router and Remote Operations (pieces of CM/2) need to be installed and active.

  Note: The NetView PIPEs function first shipped as part of the enabling PTFs for the NetView V2R3 GMFHS Feature. The WINDOW command was a sample in the NetView V2R3 level of PIPEs, and was shipped as a NetView command in NetView V2R4. I mention this since WINDOW and PIPEs were mentioned extensively in the Remote Operations part of the presentation.

- **Session level PD for APPN sessions** - NetView V2R4, VTAM V4R1, and NCP V6R2 are required. NCP V6R2 works in conjunction with VTAM V4R1 to provide a composite network node appearance to the rest of the APPN network.
- **Graphical Display of APPN networks** - NetView V2R4, the NetView APPN Topology and Accounting Management Feature, and at least one PC running CM/2 1.1 are required.

- **APPC Accounting information** - For sessions not going through VTAM (and NetView), NetView V2R4, the NetView APPN Accounting Management Feature, and at least one PC running CM/2 1.1 are required.

Part G - Further Information

These charts are simply sources of more information that people can use when trying to obtain more information on the topics covered in this presentation, or on APPN management.