CA NetMaster® Network Management for TCP/IP

Best Practices Guide

Release 12.1

technologies

This Documentation, which includes embedded help systems and electronically distributed materials, (hereinafter referred to as the “Documentation”) is for your informational purposes only and is subject to change or withdrawal by CA at any time.

This Documentation may not be copied, transferred, reproduced, disclosed, modified or duplicated, in whole or in part, without the prior written consent of CA. This Documentation is confidential and proprietary information of CA and may not be disclosed by you or used for any purpose other than as may be permitted in (i) a separate agreement between you and CA governing your use of the CA software to which the Documentation relates; or (ii) a separate confidentiality agreement between you and CA.

Notwithstanding the foregoing, if you are a licensed user of the software product(s) addressed in the Documentation, you may print or otherwise make available a reasonable number of copies of the Documentation for internal use by you and your employees in connection with that software, provided that all CA copyright notices and legends are affixed to each reproduced copy.

The right to print or otherwise make available copies of the Documentation is limited to the period during which the applicable license for such software remains in full force and effect. Should the license terminate for any reason, it is your responsibility to certify in writing to CA that all copies and partial copies of the Documentation have been returned to CA or destroyed.

TO THE EXTENT PERMITTED BY APPLICABLE LAW, CA PROVIDES THIS DOCUMENTATION “AS IS” WITHOUT WARRANTY OF ANY KIND, INCLUDING WITHOUT LIMITATION, ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NONINFRINGEMENT. IN NO EVENT WILL CA BE LIABLE TO YOU OR ANY THIRD PARTY FOR ANY LOSS OR DAMAGE, DIRECT OR INDIRECT, FROM THE USE OF THIS DOCUMENTATION, INCLUDING WITHOUT LIMITATION, LOST PROFITS, LOST INVESTMENT, BUSINESS INTERRUPTION, GOODWILL, OR LOST DATA, EVEN IF CA IS EXPRESSLY ADVISED IN ADVANCE OF THE POSSIBILITY OF SUCH LOSS OR DAMAGE.

The use of any software product referenced in the Documentation is governed by the applicable license agreement and such license agreement is not modified in any way by the terms of this notice.

The manufacturer of this Documentation is CA.

Provided with “Restricted Rights.” Use, duplication or disclosure by the United States Government is subject to the restrictions set forth in FAR Sections 12.212, 52.227-14, and 52.227-19(c)(1) - (2) and DFARS Section 252.227-7014(b)(3), as applicable, or their successors.

CA Technologies Product References

This document references the following CA Technologies products:

- CA NetMaster® File Transfer Management (CA NetMaster FTM)
- CA NetMaster® Network Automation (CA NetMaster NA)
- CA NetMaster® Network Management for SNA (CA NetMaster NM for SNA)
- CA NetMaster® Network Management for TCP/IP (CA NetMaster NM for TCP/IP)
- CA NetSpy™ Network Performance (CA NetSpy)
- CA OPS/MVS® Event Management and Automation (CA OPS/MVS)
- CA SYSVIEW® Performance Management (CA SYSVIEW)

Contact CA Technologies

Contact CA Support

For your convenience, CA Technologies provides one site where you can access the information that you need for your Home Office, Small Business, and Enterprise CA Technologies products. At http://ca.com/support, you can access the following resources:

- Online and telephone contact information for technical assistance and customer services
- Information about user communities and forums
- Product and documentation downloads
- CA Support policies and guidelines
- Other helpful resources appropriate for your product

Providing Feedback About Product Documentation

If you have comments or questions about CA Technologies product documentation, you can send a message to techpubs@ca.com.

To provide feedback about CA Technologies product documentation, complete our short customer survey which is available on the CA Support website at http://ca.com/docs.

Best Practices Guide Process

These best practices represent years of product experience, much of which is based on customer experience reported through interviews with development, technical support, and technical services. Therefore, many of these best practices are truly a collaborative effort stemming from customer feedback.

To continue and build on this process, we encourage users to share common themes of product use that might benefit other users. Please consider sharing your best practices with us.

To share your best practices, contact us at techpubs@ca.com and preface your email subject line with "Best Practices for *product name*" so that we can easily identify and categorize them.

Contents

Chapter 1: Introduction 7

Purpose of this Guide 7
Audience 7
Mainframe 2.0 Overview 7
Mainframe 2.0 Features 8

Chapter 2: Installation and Configuration Best Practices 11

Installation 11
Address Space Sharing 12
Security Considerations 12
    UAMS VSAM Data Set Sharing 13
    Background Users 14
Configuration for Optimal Performance 15
    zIIPs 15
    Express Setup 16
    Transient Logs 17
Online Help 17
Interfaces and Integration Points 18
Multisystem Deployment 19
    How Deployment Works 20
    Software Changes 21
    Create Generic Initialization File and RUNSYSIN Member for Multiple Regions 22
    Data Set Deployment 25
    Started Task Deployment 26
    Software Changes on Target Systems 26
    Multisystem Configuration 27

Chapter 3: Business Application Workload Monitoring Best Practices 29

Business Applications 29
    Business Application Names 30
    Business Applications and Address Spaces 31
    FTP-related Business Applications 31
    Telnet-related Business Applications 32

Chapter 4: Enterprise Extender Best Practices 33

Enterprise Extender Management 33
- Monitoring 33
- PALNK EE Line Activation Failures 34
- PGAIN and NEVAC EE Line Activation Failures 35
- Link Activation Failure 36
- Improved Throughput 37
- VTAM CPU Optimization 38
- Connection Networks....39
- Throughput When Multipath Routing Is Enabled 40
- EE and Firewalls 41
- EE Connection Timeouts 42

Chapter 5: Real-time Connection Monitoring Best Practices 45

Long Running Connections....45

Failing Connections 46

Chapter 6: Network Monitoring Best Practices 47

Local Physical Network Interfaces....47

Chapter 7: Network Planning Best Practices 49

IP Growth Tracker 49

Chapter 8: Non-network Personnel Best Practices 51

DB2 Network Information Center 51

IP Network Security Center 52

Index 53

Chapter 1: Introduction

This section contains the following topics:

Purpose of this Guide (see page 7)
Audience (see page 7)
Mainframe 2.0 Overview (see page 7)
Mainframe 2.0 Features (see page 8)

Purpose of this Guide

The guide provides a brief introduction to the CA Technologies mainframe management strategy and features, and describes the best practices for installing and configuring CA NetMaster NM for TCP/IP.

Audience

The intended audience of this guide is systems programmers and administrators who install, configure, deploy, and maintain CA NetMaster NM for TCP/IP.

Mainframe 2.0 Overview

Mainframe 2.0 is our strategy for providing leadership in the mainframe operating environment. We intend to lead the mainframe marketplace for customer experience, Out-Tasking solutions, and solution innovation. After listening to customer needs and requirements to keep the mainframe operating environment viable and cost-effective, we are providing new tools to simplify usage and to energize this operating environment for years to come.

CA Mainframe Software Manager™ (CA MSM) is an important step in realizing the Mainframe 2.0 strategy. CA MSM simplifies and standardizes the delivery, installation, and maintenance of mainframe products on z/OS systems. CA MSM has a web-based interface with a modern look and feel for managing those solutions. As products adopt Mainframe 2.0 features and CA MSM services, you can acquire, install, and manage your software in a common way.

We follow the IBM z/OS packaging standards using SMP/E, with some additional CA Technologies qualities of service added, to make installation simple and consistent. Additionally, through the synchronization of product releases and the use of common test environments, we will declare a yearly mainframe software stack that includes many new releases with enhanced functionality. This stack is certified for interoperability across the CA Technologies mainframe product portfolio and the base IBM z/OS product stack.

Mainframe 2.0 Features

Mainframe 2.0 has the following main features:

CA Mainframe Software Manager (CA MSM)

Delivers simplified acquisition, installation, and deployment capabilities using a common z/OS-based web application delivered through a browser-based UI. CA MSM includes the following services:

Product Acquisition Service (PAS)

Facilitates the acquisition of our mainframe products and services, including product base installation packages and program temporary fixes (PTFs). This service integrates the inventory of products available on your system with CA Support, providing a seamless environment for managing and downloading software and fixes onto your system.

Software Installation Service (SIS)

Facilitates the installation and maintenance of our mainframe products in the software inventory of the driving system. This service enables you to browse and manage the software inventory using a web interface, and automates tasks for products that use SMP/E to manage installation. You can browse downloaded software packages, and browse and manage one or more consolidated software inventories (CSIs) on the driving system.

Software Deployment Service (SDS)

Facilitates the deployment of CA Technologies mainframe products from the software inventory of the driving system. This service enables you to deploy installed products that are policy-driven with a set of appropriate transport mechanisms across a known topology. The enterprise system topology can include shared DASD environments, networked environments, and z/OS systems. Policies represent a combination of metadata input and user-supplied input. Metadata input identifies the component parts of a product. User-supplied input identifies the deployment criteria, such as where it goes and what it is named.

Software Configuration Service (SCS)

Facilitates the mainframe products configuration from the software inventory of the driving system to the targeted z/OS mainframe operating system. The SCS guides you through the configuration creation process, and through the manual steps to implement the configuration. In addition, the SCS includes an address space communications service running on each targeted z/OS system.

Electronic Software Delivery (ESD)

Enables you to get our products from an FTP server. We have improved this process so that you no longer need to build a tape to install the product.

Best Practices Management

Integrates with IBM Health Checker for z/OS to verify that deployed software follows our best practices. The health checks continually monitor the system and software to provide feedback on whether the software continues to be configured optimally.

Best Practices Guide

Provides best practices for product installation and configuration.

Active and Heartbeat Event Management through CA OPS/MVS EMA

CA Technologies mainframe products can automatically communicate both active status events and heartbeat events to CA OPS/MVS in a consistent manner. The enabling technology for this feature is through a generic event API call that CA OPS/MVS provides to the other products so that they can communicate events to CA OPS/MVS.

Two versions of this API call are provided to support this initiative:

- An active status event API call that allows other products to generate events for the CA OPS/MVS EMA System State Manager (SSM) component when they are starting, up, stopping, or down.
- A heartbeat API call that allows other CA Technologies products to communicate a normal, warning, or problem overall health status and reasoning to CA OPS/MVS EMA on a regular interval.

After a CA Technologies product begins generating heart beat events for CA OPS/MVS, CA OPS/MVS can also react to the lack of a heart beat event from another CA Technologies product address space, treating this as an indication that there is either a potential problem with the CA Technologies product address space, or there is a larger system-level problem.

SSM is a built-in feature of CA OPS/MVS that uses an internal relational data framework to proactively monitor and manage started tasks, online applications, subsystems, JES initiators, and other z/OS resources including your CA Technologies mainframe products. SSM compares the current state of online systems, hardware devices, and the other resources with their desired state, and then automatically makes the necessary corrections when a resource is not in its desired state. This provides proactive and reactive state management of critical resources. As previously noted, SSM is particularly interested in receiving active status events consistently from all CA Technologies products when they are starting, up, stopping, or down. Without this consistent type of events, SSM must maintain separate rules in CA OPS/MVS for each product unique messages that are associated with starting and stopping.

Note: For additional information about the CA Mainframe 2.0 initiative, see http://ca.com//mainframe2.

Chapter 2: Installation and Configuration Best Practices

This section contains the following topics:

Installation (see page 11)
Address Space Sharing (see page 12)
Security Considerations (see page 12)
Configuration for Optimal Performance (see page 15)
Online Help (see page 17)
Interfaces and Integration Points (see page 18)
Multisystem Deployment (see page 19)

Installation

Use CA MSM to acquire, install, and maintain your product.

Business Value:

CA MSM provides a web interface, which works with ESD and standardized installation, to provide a common way to manage CA mainframe products. You can use it to download and install CA NetMaster NM for TCP/IP.

CA MSM lets you download product and maintenance releases over the Internet directly to your system from the CA Support website. After you use CA MSM to download your product or maintenance, you use the same interface to install the downloaded software packages using SMP/E.

Additional Considerations:

After you install the product, use the product's Install Utility to set it up. CA MSM can continue to help you maintain your product.

Note: If there is maintenance for VSAM data sets, you must use the Install Utility to update those data sets for each region you have set up.

More Information:

For more information about CA MSM, see the *CA Mainframe Software Manager Product Guide*. For more information about product setup, see the *Installation Guide*.

Address Space Sharing

If your site uses multiple CA Mainframe Network Management products, share the address space with the following products for performance and usability optimization:

- CA NetMaster NM for SNA
- CA NetMaster NA

It can also share address space with the following products:

- CA NetMaster FTM
- CA SOLVE:FTS

Business Value:

Sharing an address space has the following values:

- You require only a single logon to access multiple products from one interface.
- You have better integration between products. You can have a single integrated configured address space instead of having to configure multiple address spaces.
- The multiple products can share resources.

Security Considerations

Implement the NMSAF solution. The NMSAF solution is built around a partial security exit. The solution uses the product's User Access Maintenance Subsystem (UAMS) data set to store information for your product region, and uses your installed security product to perform user validation and password checking (through the IBM-defined system authorization facility (SAF) interfaces).

Business Value:

This setup is ideal for organizations that want the flexibility of allowing the administrator to control specific region authorities, while still ensuring that access to the region is secured by their security product.

More Information:

For more information about the NMSAF solution and UAMS, see the *Security Guide*.

UAMS VSAM Data Set Sharing

Implement record-level sharing (RLS), and include the XOPT=RSLU parameter in the SYSIN member for each product region sharing the UAMS VSAM data set.

Business Value:

Multiple users on multiple systems can update a UAMS VSAM data set at the same time. The standard VSAM share options do not guarantee data set integrity with simultaneous updates from multiple systems. Using RLS, the UAMS VSAM data set can be shared without the possibility of corruption, which reduces the possibility of region outage.

Additional Considerations:

The implementation of RLS requires the proper configuration and availability of SMSVSAM. Some SMS rules for the RLS-managed data sets are also required on the systems using RLS.

More Information:

The *Security Guide* contains more information about the sharing of UAMS data set using RLS. The IBM DFSMS guides describe the implementation of RLS for VSAM data sets. For a comprehensive overview of RLS, see the chapter "VSAM Record Level Sharing" in the IBM Redbooks publication *VSAM Demystified* (SG24-6105).

Background Users

In a multisystem environment, reduce the number of background user IDs you add to security by specifying the same value for NMSUP in all regions.

Business Value:

Particularly in large complexes, this practice assists in simplifying the administration of internal background user IDs and reduces the possibility of outages associated with nonexistent, or incorrectly defined user IDs.

Additional Considerations:

CA NetMaster NM for TCP/IP uses background users to perform various tasks. By default, the NMSAF solution checks the background user IDs in advanced program-to-program communications (APPC). You must add them to your installed security product.

Note: The following NMSAF SXCTL parameters set the user ID checking: APPCCHECK and SYSCHECK.

The following list identifies the background user IDs:

- *xxxx*AOMP
- *xxxx*BLOG
- *xxxx*BMON
- *xxxx*BSVR
- *xxxx*BSYS
- *xxxx*LOGP

xxxx

Is the prefix specified by the NMSUP region job control language (JCL) parameter.

By specifying the same value for NMSUP in all regions, you only have to add one background user to security. For example, if you set NMSUP to MFNM in all regions, then the user ID for the *xxxx*BSYS background users in those regions is MFNMBSYS.

To use NMSUP, add the following statement to the TESTEXEC(RUNSYSIN) members for the regions, using the same *xxxx* value:

```
PPREF='NMSUP=xxxx'
```

More Information:

For information about SXCTL, see the *Security Guide*.

Configuration for Optimal Performance

As a performance pattern develops for your product, tune the relevant controls. You probably never have to tune many of the controls.

Business Value:

Reviewing the configuration and tuning parameters helps ensure that you are not performing unnecessary processing, such as collecting and logging data that your organization does not require, thus saving CPU cycles. As you become more familiar with the capabilities of the product, you can make informed decisions on what functions are desirable and therefore only incur overhead where there are obvious benefits.

Additional Considerations:

A product with the breadth and capability of CA NetMaster NM for TCP/IP supports many external tuning controls. Configuring every last aspect of its operation can seem like a large task. However, you can set up an effective environment by simply using the default settings.

If you have a newly implemented region, a basic configuration is created during setup and initial startup. Essential parameter groups are updated, and Express Setup have run.

More Information:

For more information about product setup and initial startup, see the *Installation Guide*.

zIIPs

If IBM System z Integrated Information Processors (zIIPs) are available, elect to use zIIPs when you set up your regions.

Business Value:

Using zIIPs provides the following benefits:

- Reducing the execution time on the normal central processing unit (CPU), providing savings in billable CPU time
- Freeing up processing cycles from the CPU to other work
- Exploiting the processing power of zIIPs

More Information:

The following JCL parameters control the usage of zIIPs: PAEXMODE for the SOLVE Subsystem Interface and XM for the region. For information about the parameters, see the *SOLVE Subsystem Interface Guide* and *the Reference Guide*.

Express Setup

Express Setup defines a collection of resources, which is a static snapshot at the time of discovery. The collection is not updated dynamically even if a new resource appears a minute later. To capture new resources, update the collection. Add new resources to the collection (or system image) manually.

Review your monitoring environment over time. Add local knowledge to the monitoring setup. Modify the discovered resources so that you are not monitoring things that are not critical to your business. Modify the monitoring of a resource from the IP resource monitor (/IPMON) using the UM line command.

Aim to get your region to satisfy the following objectives:

- Everything that is monitored is useful so that each alert must be taken notice of and each status change is significant.
- Only the things you want are monitored, but in depth.
- Automated actions are targeted and valuable.

Business Value:

Monitoring everything wastes system resources and is distracting—inconsequential alerts can distract you while causing you to miss the important ones. Removing unnecessary monitoring also reduces the processing that the region has to do.

Additional Considerations:

Express Setup is a process that you run during your first logon to a new region. The process uses rules to discover the IP resources that are present at the time it runs. These resources include active address spaces, IP nodes within a certain number of hops, defined Open Systems Adapter (OSA), virtual IP address (VIPA), and Enterprise Extender (EE) devices, and more. If requested, Express Setup also defines simple business applications.

What Express Setup discovers can only be a starting point for your monitoring environment. Express Setup does not know your business or your wider network.

Everything Express Setup discovers is placed in a system image. You can have multiple system images, each containing a particular collection of resources. System images are given unique names—generally the system ID and a version number (for example, SYS1-0001). A region can have only one system image active at a time, which is usually loaded at region startup.

You can rerun Express Setup, but only advanced users should do this and only if a major reconfiguration has occurred since the last time it was run. Rerunning Express Setup creates another system image, which has to be reconfigured from scratch. The changes you made to the existing system image are not propagated to the new image.

More Information:

For more information, see the *Implementation Guide*.

Transient Logs

Tune your transient logs to reduce storage. Disable all logging initially, and then implement logging for business critical resources (applications).

Business Value:

Mainframe storage costs money and should be used only if there is a business requirement. Tuning the size of transient logs enables you to set storage at a level appropriate to your business requirements.

Additional Considerations:

Transient logs provide a snapshot history of activities at the resource level. From a resource monitor, you can use the SETTLOG command to disable logging or reset the log size for one or more monitored resources.

Online Help

Use online help to find out more about the interface in context.

Business Value:

CA NetMaster NM for TCP/IP has many features and can be overwhelming to new users. However, you have access to substantial online help at both the 3270 and WebCenter interfaces, usually by pressing F1 or clicking the Help link. You are encouraged to request online help, to promote product understanding, save time on issue resolution, and potentially save money on basic product training.

Additional Considerations:

The IBM standard code page for accessing a z/OS mainframe with US English is 037. If your language of choice is English, set your TN3270 emulators and mainframe terminals to code page 037.

Interfaces and Integration Points

Integrate with other CA products to help you manage your business.

Business Value:

CA NetMaster NM for TCP/IP integrates with the following CA products:

- **Other CA NetMaster products**—All CA NetMaster products can share the same address space. They can also communicate with each other using their multisystem capabilities. The use of common monitors, such as the alert monitor and the status monitor, supports the combined monitoring and control of network events and resources irrespective of whether they are IP, SNA, or file transfer related.
- **CA NetSpy**—You can display CA NetSpy information in CA NetMaster NM for TCP/IP, including TN3270 response time information. You can also issue CA NetSpy commands from and display their responses in CA NetMaster NM for TCP/IP.
- **CA Service Desk**—CA NetMaster NM for TCP/IP supports the automatic creation of trouble tickets in CA Service Desk, facilitating problem notification and resolution.
- **CA OPS/MVS**—CA OPS/MVS can forward system events programmatically to CA NetMaster NM for TCP/IP for display on the alert monitor. The WebCenter presentation of the alert monitor is an ideal consolidation point for all mainframe network and system events. The integration facilitates the flow of information between CA management products and users.
- **CA SYSVIEW**—CA NetMaster NM for TCP/IP comes predefined with CA SYSVIEW event detectors. This enables a system performance related event to initiate actions in CA NetMaster NM for TCP/IP, including displaying CA SYSVIEW events on the alert monitor.

Multisystem Deployment

If you have multiple systems, deploy CA NetMaster NM for TCP/IP in a multisystem environment to provide a consolidated view of your enterprise.

The following diagram shows an example of a multisystem environment. Logging on to Console A allows visibility to all resources. Logging on to Console B allows visibility to the subordinate system image only.

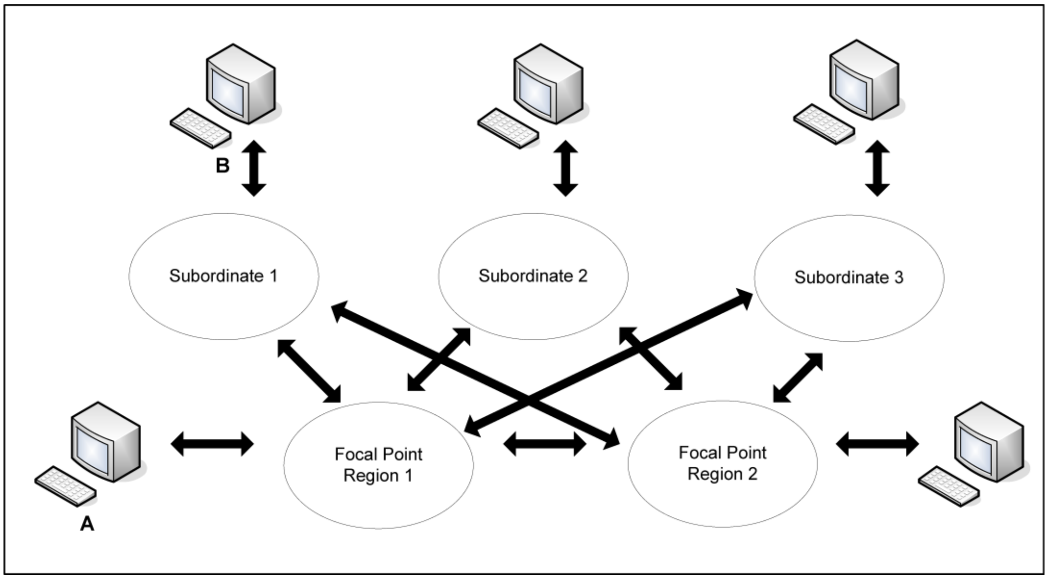

Business Value:

Particularly in large multisystem environments, deployment can be both arduous and time consuming. Following an effective and proven process has the following values:

- Reduce the time taken to migrate to new releases, and therefore enable access to new functions more readily.
- Free key resources to perform other tasks, such as the exploitation of product functions.
- Reduce the likelihood of errors and subsequent outages associated with poor deployment processes.

Additional Considerations:

We recommend that you use the CA MSM SDS to deploy the product SMP/E target libraries to the remote systems.

Note: For more information about SDS, see the *CA Mainframe Software Manager Product Guide*.

You set up and configure the product once, typically on a test system, which becomes the deployment system. After the product is configured, you create backup data sets for the configuration files (see page 25). You can then use SDS to deploy the product target libraries and at the same time, the backup data sets as custom data sets. On the remote system, you can restore the configuration files from the deployed backup data sets.

How Deployment Works

Before you proceed to perform multisystem deployment, you should have one properly configured region.

Typically, deployment consists of the following stages:

1. Create a generic RUNSYSIN member and a generic initialization file for sharing between regions.
2. Copy the required data sets to and allocate them on the target systems.
3. Deploy started task members on target systems.
4. APF authorize load libraries on target systems.
5. Determine focal and subordinate regions.
6. Link regions to create the multisystem environment.

More information:

Multisystem Configuration (see page 27)

Software Changes

Changes are required to set up subsystem IDs, load libraries, and VTAM:

- Two subsystem IDs are required for the initialization of the required subsystems. The IDs have the following default values:

  - SOLV for the SOLVE Subsystem Interface (SSI), which enables a region to communicate with other software on the system

  - Domain ID of the region for the region interface that enables a region to issue operating system commands and receive messages

  The SOLVE SSI started task and the region automatically identify these IDs to the system. If you want to set the IDs permanently, you can set them in the SYS1.PARMLIB(IEFSSN*xx*) member. Add the ID for the region interface first (after the job entry subsystem (JES)) in the list of subsystem names.

- The following load libraries for CA NetMaster NM for TCP/IP must be APF-authorized:

  - CC2DLOAD

  - CC22DPLD (if SSL is installed)

- A VTAM major node member, which contains application definition statements for all ACBs required by your product region, must be created and added to SYS1.VTAMLST. You can use the Create VTAM Definitions and Table option of the product's Install Utility to perform this task.

Note: For more information, see the *Installation Guide*.

Create Generic Initialization File and RUNSYSIN Member for Multiple Regions

Create a generic RUNSYSIN member that points to a generic initialization file so you can use the member for all the regions deployed in your enterprise.

To create generic initialization file and RUNSYSIN member

1. Generate the initialization file for a properly configured region.
2. Replace specific information in the file by product variables and system symbols.

   A generic initialization file is created.
3. Replace specific information in RUNSYSIN by system symbols.
4. Update RUNSYSIN with the following statement:

   ```
   PPREF='INIFILE=xxxxINI'
   ```

   ***xxxx*INI**

   Is the name of the generic initialization file.

   A generic RUNSYSIN member is created.
5. Start the region using the generic RUNSYSIN member to verify that it is free of errors.

   If initialization errors occur, review RUNSYSIN and the initialization file to correct the errors.
6. Repeat the previous step until the region initializes without error.

   The generic RUNSYSIN member is ready for use by other regions.

Region Initialization File

Region customization parameters are stored in a virtual file system (VFS) data set, which is a virtual storage access method (VSAM) data set and is not easy to update outside of CA NetMaster NM for TCP/IP. However, a RUNSYSIN member can point to an initialization file member in TESTEXEC by using the INIFILE parameter.

An initialization file is a Network Control Language (NCL) procedure that contains the parameter information.

When an initialization file is in use, the region gets the parameter information from the file at startup and updates the VFS data set. Because the region uses the initialization file each time it starts up, any changes you make manually using the /PARMS panel shortcut are not retained. To keep the changes, regenerate the file using the /CUSTOM.G panel path.

Even if you do not use the initialization file for region startup, you can use it as a backup of the parameters in the VFS data set by generating it before updating the parameters using /PARMS.

The initialization file is also useful during rollout to other systems because it is relatively simple to update for different systems. Through the use of product variables and system symbols, the file can be made generic enough for all the regions you plan to deploy.

Note: For more information about how to use a region initialization file, see the *Administration Guide*.

Generic Initialization File

You can modify an initialization file to use system symbols to support its use throughout your enterprise.

Example: Initialization File With System Symbols

The following sample code shows statements for the IPFILES parameter group using the &SYSNAME system symbol:

```
.IPFILES
    &$IAP$IPDBDSN = &STR NETW.NM.NETM&SYSNAME.IPFILE
    &$IAP$IPLDSN = &STR NETW.NM.NETM&SYSNAME.IPLOG
    &$IAP$IPDTDSN = &STR NETW.NM.NETM&SYSNAME.IPDETAIL
    &$IAP$IPTRDSN = &STR NETW.NM.NETM&SYSNAME.IPTREND
    &$IAP$IPLSDSN = &STR NETW.NM.NETM&SYSNAME.IPLOGSEQ
```

Generic RUNSYSIN

By building a RUNSYSIN member using system symbols, you can create a generic RUNSYSIN member that can be deployed throughout your enterprise.

You can identify the symbols defined to your system from the response to the following system command:

```
D SYMBOLS
```

To tell the NETMASTR program to perform symbol substitution, include the following statement in RUNSYSIN:

```
SUBS=YES
```

Example: RUNSYSIN with System Symbols

The following sample code shows RUNSYSIN statements using the &SYSNAME and &SYSCLONE system symbols:

```
SUBS=YES -* Required to invoke system symbols
PGM=NM001
ERROR=U0001
PPREF='PRI=NETM&SYSNAME' -* if &SYSNAME = "ABCD", PRI=NETMABCD
PPREF='NMDID=&SYSCLONE.NW'
PPREF='INIT=NMINIT'
PPREF='READY=NMREADY'
PPREF='SSID=NMSS'
PPREF='DSNQLCL=NETW.NM.NETM&SYSNAME'
PPREF='DSNQLNV=NETW.NM.VSAM.NETM&SYSNAME'
```

Note: A symbol used in the middle of the name must be defined with two periods (..), for example:

```
DD=VFS,DISP=SHR,DSN=NETW.NM.NETM&SYSNAME..VFS
```

Data Set Deployment

During deployment, you copy data sets to the target systems. However, in a shared DASD environment, you do not need to copy shared data sets.

Note: When deploying CA NetMaster NM for TCP/IP with other CA NetMaster products, such as CA NetMaster NM for SNA, additional data sets specific to these products may also be required. For information, see the *Best Practices Guide* specific to those products.

You must copy the following data sets:

- ALERTH
- ICOPANL
- IPFILE
- IPLOG
- IPLOGSEQ
- IPMIBX
- MODSUSR
- MSDB
- NMLOG01
- NMLOG02
- NMLOG03
- OBEYFILE
- PANLUSR
- PSPOOL
- RAMDB
- RAMDBST
- RAMBDWK
- REXXAN
- REXXREP
- SSIDB
- TESTEXEC
- VFS

The following data sets can be shared:

- CC11EXEC (shareable)
- CC2DEXEC (shareable)
- CC2DLOAD (shareable)
- CC2DPLD (shareable)
- MODSDIS (shareable)
- OSCNTL (shareable)
- PANLDIS (shareable)
- PARMLIB (shareable)
- SSIPARM (shareable)
- UAMS (shareable)

One method to distribute data sets is to use a backup utility, such as DFDSS, to create a single data set that can be transferred to the target systems and restored.

When the Install Utility sets up a region, the utility creates the following data set members:

S10DUMP

Creates backup data sets that include the configuration files for the region. These backup data sets are *dsnpref*.DFDSS.LOCAL (containing files specific to the region) and *dsnpref*.DFDSS.SHARED (containing files that multiple regions can share).

S11REST

Restores the configuration files from the backup data sets.

After you submit the S10DUMP job, you use SDS to deploy the created backup data sets to the target system. Also, you copy the S11REST job to the target system. On the target system, you submit S11REST to restore the configuration files.

Started Task Deployment

During deployment, you copy the region and SOLVE SSI started task members to SYS*x*.PROCLIB on the target systems.

Software Changes on Target Systems

During deployment, you add the subsystem IDs and ACBs, and APF-authorize the load libraries on the target systems.

More information:

Software Changes (see page 21)

Multisystem Configuration

Regions can be linked together into a complex. Within a complex, you can have two types of regions: focal and subordinate.

A focal region has visibility to, and command and control capabilities over, every region in the complex, including other focal regions.

A subordinate region only sends data to the focal regions. A subordinate does not receive data from other regions in the complex.

To reduce network traffic, focal regions only receive status information if someone is actually using one of the various monitors.

In a multisystem environment, operators can log on to one focal region and monitor the entire complex.

The Resource Automation Monitor database (RAMDB) for a focal region contains copies of the system images for all regions within the complex. The RAMDB for a subordinate region contains only the system images for itself.

In general, you configure regions on communication management configuration (CMC) systems (hosts) as focal, and all the others as subordinate.

How You Prepare RAMDB Before Linking

You can use the following methods to prepare RAMDB before you link your regions to set up the multisystem environment:

Important! When you link two regions, one region has the database you want and the other region will have its database overwritten. Linking must always be initiated from the region whose database is to be overwritten.

- You can create the system images for the individual systems on which the regions are deployed. You then assign one region as focal, transmit the images from the other regions to it, and then link the other region to it.

  The linking must be done from the new region where the database is deleted and rebuilt to mirror that in the focal region.

- You can create the system images for all the required systems in the complex in the main focal region. Then each new region can be deployed with the default RAMDB provided during setup. The default RAMDB is deleted and rebuilt with the required system images when the region is linked to the focal region.

After the regions are linked, their RAMDBs are kept synchronized automatically.

Note: For more information about how to set up a multisystem environment, see the *Administration Guide*.

Chapter 3: Business Application Workload Monitoring Best Practices

This section contains the following topics:

Business Applications (see page 29)

Business Applications

Group IP connections into meaningful business application names to help you understand the network activity and workload associated with key business applications.

Business Value:

The primary goal of a mainframe-based network is to provide reliable access to critical data and applications that reside on z/OS systems. This data and these applications underpin many of your business applications and services. Being able to view network activity and workload in terms of your key business applications and services enables you to better understand their well-being, prioritize network events, and assure service to the business.

Additional Considerations:

You define business application names using connection criteria.

- If the Create Application Name Definitions? field is set to YES when Express Setup is run, a business application name is created for every active address space discovered. The name is the address space name, and the ports are the active ports when the address space was discovered.
- Define meaningful names that are useful to report on. Do not confuse flexibility with granularity—too much granularity can be expensive. Large numbers of unnecessary application names or historical connections can use up space in the SOLVE SSI database and can increase CPU usage.

More Information:

For more information about defining and monitoring business applications, see the *Implementation Guide* and *User Guide*.

Business Application Names

Business application names can be anything you like—often unrelated to the address space or task name. You can set separate performance and alerting criteria for each name. Names can have optional dynamic suffixes (for example, local port, job name, remote address, or remote port).

Examples: Business Application Names

- The following application names group connections by ports:
  - WebSphereMQ as an application name for all connections to local port 1414
  - Tomcat as an application name for all connections to local ports 8080 and 8443
  - FredUpd as an application name for all connections to task name FRED, port 12300
  - FredAdm as an application name for all connections to task name FRED, port 12301
- The following application name groups all connections between task ABC and local port 12345, and remote address *x*: BANK1TRANS.
- The following application names group connections by functions:
  - Web as an application name for all connections to local ports 80, 8080, 443, and 8443
  - HTTPS as an application name for all connections to local ports 443 and 8443
  - HTTP as an application name for all connections to local ports 80 and 8080
- The following application name groups all connections to remote address *x* and port *y*: PerthOffice.
- The following application names group connections to tasks:
  - CICSSYD as an application name for all connections to CICS* from remote addresses *w* through *x*
  - CICSMELB as an application name for all connections to CICS* from remote addresses *y* through *z*
  - CICS*addr* as an application name for all connections to CICS* from remote address *addr*
- You can use the same name for a business application on different systems, even if the address space names are different.

  CICSACCT can refer to address space CICSSYS1 on SYS1 and CICSSYS2 on SYS2.

Business Applications and Address Spaces

Although you can optionally specify an address space name, port, and stack as criteria in a business application, you can name the application independently of address spaces. Business application setup is flexible and is not bound by address spaces. A business application does not need to know what address spaces a connection uses.

Examples: Business Application for Address Spaces

- A business application can be a subset of connections to one address space (such as those connections from specific foreign hosts or ports). One address space can have many associated business applications for different subsets of its connections.
- A business application can be all connections to a single address space. You can even call the application the same as the address space, though this implementation is a waste of business application flexibility.

  However, you can map an address space and its backup to the same business application. For example, you can map address spaces AS1 and AS1BKUP to business application APPL123, enabling you to keep continuous performance statistics even when the underlying address space changes.
- A business application can be a subset of the connections to multiple address spaces (for example, traffic from any address spaces to a specific printer).
- A business application can be all connections to multiple address spaces. For example, you have several CICS subsystems and want to group all CICS traffic.

FTP-related Business Applications

Express Setup groups all FTP connections under the business application name FTP. You can define other names to make the monitoring of your FTP transfers more granular, for example:

- Name groups of FTP transfers based server or location.
- Name groups of FTP transfers based on who is at the other end.

You can specify separate performance alerting criteria for each name. One group of FTP transfers can have much stricter controls than others.

Telnet-related Business Applications

Express Setup groups all Telnet connections under the business application name Telnet. You can define other names to make the monitoring of your Telnet connections more granular, for example:

- Names to group Telnet applications
- Names to group logical function or geographical location of the users
- Names to group type or location of the Telnet servers
- Names to group security requirements of connections

Chapter 4: Enterprise Extender Best Practices

This section contains the following topics:

Enterprise Extender Management (see page 33)

Enterprise Extender Management

Use these procedures to manage your EE installation.

Business Value:

EE is complex to debug and manage. These procedures are gained through experience and research, which show you how to use CA NetMaster NM for TCP/IP to improve the reliability and efficiency of your EE installation.

Additional Considerations:

EE enables legacy SNA applications and clients to continue to work unchanged by transparently sending SNA APPN/HPR traffic over an IP network. From an SNA view, EE is a logical link defined as an XCA (external communications adapter) major node and a switched major node. From an IP view, EE is UDP traffic over the IP backbone.

Monitoring

Restrict monitoring in *one* of the following ways:

- Define a filter to restrict the monitoring to specific remote control points (CPs).
- Limit the number of monitored EE connections.

Monitoring everything wastes system resources and is distracting—inconsequential alerts can distract you while causing you to miss the important ones. Removing unnecessary monitoring also reduces the processing that the region has to do.

You define your filter or limit the number of monitored connections on the EE Monitoring Definition panel of the definition.

PALNK EE Line Activation Failures

A line status of PALNK indicates that VTAM's connection to the stack using the IUTSAMEH device is not active.

Use CA NetMaster NM for TCP/IP to check EE line activation

1. From the command prompt, enter the following panel shortcut:

   ```
   /EEXCA
   ```

   The EE XCA Major Node Summary is displayed. The stack name appears in the top right of the page header. Line status is the second column.

2. Note the stack name of any line with a status of PALNK.
3. From the command prompt, enter the following panel shortcut:

   ```
   /STACK.M
   ```

   The IP Resource Monitor filtered for STACK is displayed.

4. Verify that the stack is active for the system being checked.
5. Enter **DL** against the required stack.

   The Device Links List is displayed.

6. Verify that the IUTSAMEH device has status of READY (on IPv6 systems, scroll right).
7. Verify that the source VIPA address is specified:
   - In the IPADDR VTAM start option
   - In the XCA GROUP definition statement

PGAIN and NEVAC EE Line Activation Failures

A line status of PGAIN indicates that name resolution is in progress. This process can time out, and the status changes to NEVAC.

Use CA NetMaster NM for TCP/IP to check EE line activation

1. From the command prompt, enter the following panel shortcut:

   ```
   /EEXCA
   ```

   The EE XCA Major Node Summary is displayed. Line status is the second column.

2. Note the host name of any line with a status of PGAIN or NEVAC.
3. From the command prompt, enter the following command:

   `CMD NSLOOKUP` *`hostname`*

   The command entry panel is displayed.

4. Press Enter.

   The name lookup executes and returns IPGP12*xx* messages.

   IPGP1218 indicates that name lookup failed. IPGP1214 indicates a successful lookup.

5. Verify that the specified source VIPA address host name can be resolved from the host name specified in the HOSTNAME VTAM start option or the XCA GROUP definition statement.

Note: The NSLOOKUP result is an indicative test when CA NetMaster NM for TCP/IP's SOCKETS interface uses the same stack as VTAM, and DNR is configured to use the vendor's interface. If SHOW DNR indicates DNR=SOLVE, the NSLOOKUP result is an indicative test when the configuration file setup is the same as the stack's. If CA NetMaster NM for TCP/IP is using a different stack, or the DNR configuration file setup is different to the stack, the NSLOOKUP command does not necessarily return the same result that VTAM is receives.

Link Activation Failure

Link activation failures occur when VTAM does not receive responses to XID requests. CA NetMaster NM for TCP/IP retains the inactive PU in the EE XCA major node summary.

Use CA NetMaster NM for TCP/IP to investigate link activation failures

1. From the command prompt, enter the panel shortcut:

   ```
   /EEXCA
   ```

   The EE XCA Major Node Summary is displayed.

2. Do *one* of the following:

   - For a known static connection with a status of INACT, enter **A** against the connection. If the activation fails, enter CT against the connection.
   - For a known static connection with a status of RESET, enter **CT** against the connection.
   - For a new connection:

     a. From the command prompt, enter the command:

     ```
     /EE
     ```

     The Enterprise Extender Management menu is displayed.

     b. Enter **CT** and the IP address.

   The EE Connectivity Test panel is displayed. If the destination:

   - Cannot be reached, the following causes are possible:
     - The IP connectivity is lost within the network.
     - The firewall does not allow UDP traffic for EE ports 12000 through 12004.
   - Can be reached but is not responding, check that EE is enabled on the remote end point.

Improved Throughput

Frequent path switches caused by setting PSWEIGHT to EQUAL or SAMEROUT can lead to reduced throughput.

Use CA NetMaster NM for TCP/IP to improve throughput

1. From the command prompt, enter the following panel shortcut:

   ```
   /RTP
   ```

   The RTP Pipes List is displayed.

2. From the command prompt, enter the following command:

   ```
   SORT SW# D
   ```

   The RTPs are sorted by the number of path switches in descending order.

3. Scroll right to review the number of path switches and the reason for the last path switch.

   A high number of switches with the reason AUTO PATH SWITCH FOR PSRETRY indicates a possible problem.

4. From the command entry panel, enter the following command:

   ```
   D VTAMOPTS,OPTION=PSRETRY
   ```

   The VTAM options are displayed. Values of 0 indicate that there is no timer-based switching of RTPs.

5. From the command entry panel, enter the following command:

   ```
   D VTAMOPTS,OPTION=PSWEIGHT
   ```

   The VTAM options are displayed. Review the START option for the weight comparison between an old route and a newly calculated route before a path switch is attempted.

VTAM CPU Optimization

If you have many EE connections kept indefinitely active, the following functions can reduce VTAM CPU use:

HPR alive timer optimization

This function is enabled by default and is controlled by the HPREELIV operand on the XCA major node.

LDLC Keep-Alive reduction

This function requires you to specify a maximum time value on the LIVTIME operand for the EE PORT statement.

Use CA NetMaster NM for TCP/IP to determine VTAM CPU use

1. From the command prompt, enter the following panel shortcut:

   ```
   /EEXCA
   ```

   The EE XCA Major Node Summary is displayed.

2. Press F4 (Display).

   The results of the command D EE,LIST=DETAIL are displayed in Command Entry.

3. Scroll down, and review the IST2004I message for each line group.

   LIVTIME=(*default*,*max*) is displayed.

   - A *max* value of 0 indicates that the time between TEST sends does not increase when the connection is idle. LDLC Keep-Alive reduction is not active.
   - A nonzero value for *max* indicates that the default time is doubled during idle periods until the max value is reached.

4. From the command prompt, enter the following panel shortcut:

   ```
   /EERTP
   ```

   The RTP pipes currently using an EE connection are listed.

5. Enter **S** next to a pipe.

   The RTP panel is displayed.

6. Check that the Liveness Timer field has a value of '0s'.

Connection Networks

During an IP network outage, if the EE connection network path has the lowest weight of any available path to the partner node, continuous attempts to redial the partner node are made over this virtual routing node (VRN) path. This results in failures.

Use CA NetMaster NM for TCP/IP to avoid EE connection retries

1. From the command prompt, enter the following panel shortcut:

   ```
   /EEXCA
   ```

   The EE XCA Major Node Summary is displayed.

2. Enter **UN** against a line group associated with a VRN.
3. Determine how many partners have been marked as unreachable.
4. If the IP problems are fixed, access the /EEXCA display and enter **UNC** against the VRN.

   The unreachable partner information is cleared.

5. If the IP problems are not fixed, enter the command:

   ```
   /F TOPO ID=?virtualnode,FUNCTION=QUIESCE,SCOPE=LOCAL/NETWORK
   ```

   The VRN is disabled for new APPN connections.

6. When the IP problem is resolved, enter the command:

   ```
   /F TOPO ID=?virtualnode,FUNCTION=NORMAL,SCOPE=LOCAL/NETWORK
   ```

   The VRN is enabled for APPN connections.

Throughput When Multipath Routing Is Enabled

If multipath routing is enabled and multiple equal-cost routes exist to the partner EE node, TCP/IP sends EE packets across each of these routes. If one route cannot reach the partner EE node, EE might not activate or performance might be impacted.

Use CA NetMaster NM for TCP/IP to determine the effect of multipath routing

1. From the command prompt, enter the following panel shortcut:

   ```
   /EEUDP
   ```

   A list of EE UDP Connections is displayed.

2. Enter **S** against the required connection.

   The EE UDP Connection Information is displayed.

3. Scroll down to the Interface Usage, and review the number of interfaces used for a connection and the Packets Out volume for each interface.

Use CA NetMaster NM for TCP/IP to determine if multipath routing is enabled

1. From the command prompt, enter the following panel shortcut:

   ```
   /CONFIG
   ```

   The Stack Configuration Information menu is displayed.

2. Enter **DP** against the stack used by VTAM for EE.

   The Stack Configuration Information panel is displayed.

3. Enter **S** against a profile TCP/IP.

   The Browse PROFILE Dataset panel is displayed.

4. Review the profile options to see if MULTIPATH is set.

EE and Firewalls

The firewalls must allow both inbound and outbound UDP traffic on all five EE ports. If this condition is not satisfied, the EE connection is established but sessions cannot be established.

Use CA NetMaster NM for TCP/IP to test a firewall for an existing connection

1. From the command prompt, enter the following panel shortcut:

   ```
   /EEXCA
   ```

   The EE XCA Major Node Summary is displayed.

2. Enter **CT** against the required connection.

   The connectivity test is invoked. If the destination:

   - Cannot be reached, the following causes are possible:
     - The IP connectivity is lost within the network.
     - The firewall does not allow UDP traffic for EE ports 12000 through 12004
   - Can be reached but is not responding, check that EE is enabled on the remote end point.

Use CA NetMaster NM for TCP/IP to test a new host

1. From the command prompt, enter the following panel shortcut:

   ```
   /EE
   ```

   The Enterprise Extender Management menu is displayed.

2. Enter **CT** and a value for the Remote Host Name/Addr.

   (You may be prompted for a local VIPA.)

   The connectivity test is invoked. If the destination:

   - Cannot be reached, the following causes are possible:
     - The IP connectivity is lost within the network.
     - The firewall does not allow UDP traffic for EE ports 12000 through 12004
   - Can be reached but is not responding, check that EE is enabled on the remote end point.

EE Connection Timeouts

The EE connection terminations are caused by exchange identification (XID) or Logical Data Link Control (LDLC) timeouts.

The LDLC layer monitors the EE connection with the LIVTIME, SRQRETRY, and SRQTIME parameters on the PORT statement. The connection is terminated if no activity or response occurs for a duration (in seconds) approximated by the following expression:

```
LIVTIME + (SRQRETRY+1) * SRQTIME
```

LIVTIME

Specifies the amount of inactivity time that can lapse before LDLC tests the connection.

SRQTIME

Specifies the amount of time LDLC waits for a response to its test.

SRQRETRY

Specifies the number of times the test is retried.

The actual duration of the outage detection is plus or minus one SRQTIME interval of the calculated amount.

Some methods to avoid timeouts are:

- Lengthen the HPR path switch timers (HPRPST) to ensure that all four timers are longer than the LDLC timeout interval.
- Ensure that the values of the LDLC parameters are consistent between the end points of the connection.
- For predefined EE PUs, specify DISCNT=NO.

Use CA NetMaster NM for TCP/IP to avoid connection timeouts

1. From the command entry panel, enter the following command:

   ```
   D VTAMOPTS,OPTION=HPRPST
   ```

   The HPR path switch timers are displayed. The shortest time is typically for the network priority and is 60 seconds.

2. From the command prompt, enter the following panel shortcut:

   ```
   /EEXCA
   ```

   The EE Major Node Summary is displayed.

3. Enter **S** against a specific PU from the connection, and press F4 (Display).

   The LIVTIME settings are displayed.

   Review message IST2114I for the current LIVTIME setting.

4. Press (F3) Exit.

   The EE Major Node Summary is displayed.

5. Press F4 (Display).

   The ENTERPRISE EXTENDER GENERAL INFORMATION messages are displayed.

6. Review the message IST2004I to calculate the timeout duration:

   ```
   livtime + (srqretry+1) * srqtime
   ```

Chapter 5: Real-time Connection Monitoring Best Practices

This section contains the following topics:

Long Running Connections (see page 45)
Failing Connections (see page 46)

Long Running Connections

Identify any applications that rely on one or more long running connections, and define an event detector to identify issues associated with these connections.

Business Value:

Critical business applications sometimes rely on specific IP connections being always in existence. The loss of one or more of these connections can result in serious impact to the business application and the service it provides. The quick identification of a long running IP connection ending can allow for automated or operator assisted recovery to minimize the business impact.

Additional Considerations:

The appropriate event detector to use depends on the monitoring requirements, for example:

- When a specific number of connections must always be present, use a CONNSTAT event detector, which enables you to use the minimum number of active connections as a criterion. This detector verifies that the minimum number of connections are present at startup and highlights any subsequent failures that lead to the minimum number not being available.
- When there are not a specific number of connections required but any started connections must remain active, use a TCPEND event detector, which supports the detection of any connection ending. If you are only interested in abnormal terminations, you can specify a termination reason code.

  For TCPEND event detectors, we recommend that you specify alerting by server to consolidate all failures for an application under a single alert.

More Information:

For information about how to define event detectors, see the *Implementation Guide*.

Failing Connections

Define event detectors to monitor for connection failures for all critical applications.

Business Value:

Critical business applications sometimes rely on specific IP connections being always in existence. The loss of one or more of these connections can result in serious impact to the business application and the service it provides. The quick identification of an IP connection failing can allow for automated or operator assisted recovery to minimize the business impact.

Additional Considerations:

The appropriate event detector to use depends on the monitoring requirements, for example:

- You can use an SVRRESET event detector to detect a connection failing because the server resets the connection. If the server is resetting several connections, it can highlight a critical condition with the application that must be addressed.
- You can use a TCPEND event detector to detect any failing connection. Specify a range of termination reason codes as appropriate.

For SVRRESET and TCPEND event detectors, we recommend that you specify alerting by server to consolidate all failures for an application under a single alert.

More Information:

For information about how to define event detectors, see the *Implementation Guide*.

Chapter 6: Network Monitoring Best Practices

This section contains the following topics:

Local Physical Network Interfaces (see page 47)

Local Physical Network Interfaces

Identify the local physical network interface in a ping request.

Business Value:

Identifying a local physical network interface in a ping request enables you to check the health of that interface.

Additional Considerations:

Consider defining an IP node for each business critical interface and using the interface name to identify the definition. To minimize network traffic, use a nearby router or host for the monitored IP address or host name. This node provides monitoring of the operational status of each of these interfaces.

In the associated monitor group, add alerting on the PING attribute to notify a network operator if the ping is unsuccessful. Corrective action can then be quickly taken.

More Information:

For information about how to configure IP node monitoring, see the *Implementation Guide*.

Chapter 7: Network Planning Best Practices

This section contains the following topics:

IP Growth Tracker (see page 49)

IP Growth Tracker

Use the WebCenter IP Growth Tracker to review the growth in IP traffic demand and the distribution of connection durations.

Business Value:

The IP Growth Tracker helps you identify trends and characteristics of your network usage and activity. This information supports the making of informed and timely network planning decisions, for example:

- Provisioning for additional workload
- Balancing workload across available resources to maximize the benefits from your existing infrastructure
- Identifying long-term inconsistencies that can indicate problem situations

Additional Considerations:

The IP Growth Tracker uses column charts to show the growth in IP traffic on a system over a period. The page also shows the distribution of TCP connections by the time over which the connections are active.

You access the IP Growth Tracker page from the WebCenter login dialog (if enabled by the WEBCENTER parameter group) or from Performance Center after you log in.

Chapter 8: Non-network Personnel Best Practices

This section contains the following topics:

DB2 Network Information Center (see page 51)
IP Network Security Center (see page 52)

DB2 Network Information Center

Use the DB2 Network Information Center to help you understand your DB2 network.

Business Value:

The center consolidates the functions to access DB2 network information in one place. The DB2 staff has a single place to go to get information without the distraction of other product features.

Additional Considerations:

The DB2 Network Information Center provides a single point of access for DB2 staff to find out about DB2 network activities:

- You can find out about and diagnose Distributed Data Facility (DDF) connections.
- You can display statistics on DB2 address space activities.
- You can trace packets for defined ASMON resources of Type DB2.

You can access the menu for the DB2 Network Information Center using the /DB2 panel shortcut, or the D option on the Address Space and Port Management menu. To learn more about the center, see the tutorial on the menu.

IP Network Security Center

Use the IP Network Security Center to help you understand the security of your IP network.

Business Value:

The center consolidates the functions to access security information about your IP network in one place. The security staff has a single place to go to get information without the distraction of other product features.

Additional Considerations:

The IP Network Security Center provides a single point of access for you to find out about and manage the security of your IP network:

- You can find out about and diagnose problems for secured connections.
- You can find out about IPSec configuration and manage tunnels.

You can access these functions from the IP Security menu, using the /SECURE panel shortcut or the SEC option on the Stack Management menu.

Index

A

activation failures • 34, 35, 36
address spaces
- business application names • 31
- sharing • 12

B

background users • 14
business application names • 29
- address spaces • 31
- examples • 30
- FTP • 31
- Telnet • 32

C

CA MSM (CA Mainframe Software Manager) • 8, 11
CA OPS/MVS integration • 9
configuration • 15
connection networks • 39
connections
- durations • 49
- terminations • 42, 45

CPUs (central processing units)
- offloading work from • 15
- use, EE • 38

D

DB2 Network Information Center • 51
deployment
- data sets • 25
- multisystem • 19

E

EE (Enterprise Extender) • 33
ESD (Electronic Software Delivery) • 8
event detectors • 45, 46
Express Setup • 16

F

failures
- connections • 46
- EE • 34, 35, 36

firewalls, EE • 41
FTP (File Transfer Protocol), business application names • 31

H

health checks • 8
help • 17

I

initialization • 22
installation • 8
integration • 18
IP Growth Tracker • 49
IP Network Security Center • 52

L

load library authorization • 21

M

Mainframe 2.0 • 7
monitoring
- EE • 33

multipath routing, EE • 40
multisystem support • 19

N

network information
- DB2 • 51
- security • 52

O

online help • 17

P

PAS (Product Acquisition Service) • 8
path switches • 37

R

RAMDB • 28
regions
- UAMS data set sharing • 13

remote CPs, EE • 33

S

SDS (Software Deployment Service) • 8
security • 12
setup • 12
SIS (Software Installation Service) • 8
size, transient logs • 17
SOLVE SSI • 21
SSM (System State Manager) • 9
subsystem IDs • 21

T

TCP traffic growth • 49
Telnet, business application names • 32
throughput, EE • 37, 40
transient logs • 17

U

UAMS data set • 13
UDP traffic
    firewalls • 41
    growth • 49

V

VRNs (virtual routing nodes) • 39
VTAM CPU use, EE • 38

Z

zIIPs • 15