CA Mainframe Application Tuner

Best Practices Guide

Version 9.0.00

CA Technologies Product References

This document references the following CA Technologies products:

- CA Endevor® Software Change Manager
- CA InterTest™ Batch
- CA InterTest™ for CICS
- CA JCLCheck™ Workload Automation
- CA Mainframe Application Tuner (CA MAT)
- CA Mainframe Software Manager ™ (CA MSM)
- CA Optimizer®/II
- CA SymDump® Batch
- CA SymDump® for CICS

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.

Documentation Changes

The following documentation updates have been made since the last release of this documentation:

Following are the updates made to many topics:

- Changed the release number from r8.5 to V9.0
- Removed the text <variable>

Following are the changes made to selected topics:

- DB2 Configuration Best Practices—Removed the following list items: Collecting Information from the Plan Table, Extracting SQL Statement Text from the DB2 Catalog, and Defining the Size of the Internal Table. Removed the Additional Considerations heading and text.
- Removed the following topics from the chapter "Configuration Best Practices": Getting Information about the Acces Path Chosen by DB2 and Extracting SQL Statement Text from the DB2 Catalog.

Contents

Chapter 1: Introduction

This section contains the following topics:

Purpose of This Guide

The guide describes the best practices for installing and setting up CA Mainframe Application Tuner, and how to use it to acquire, install, maintain, and deploy your CA Technologies mainframe products.

Audience

The intended audience of this guide is systems programmers and administrators who install, configure, deploy, and maintain CA Mainframe Application Tuner.

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 Mainframe Software Manager) is an important step in realizing the Mainframe 2.0 strategy. CA Mainframe Software Manager simplifies and standardizes the delivery, installation, and maintenance of mainframe products on z/OS systems. CA Mainframe Software Manager has a web-based interface with a modern look and feel for managing those solutions. As products adopt Mainframe 2.0 features and CA Mainframe Software Manager 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.

Chapter 2: Installation Best Practices

This section contains the following topics:

Use CA Mainframe Software Manager for Installation

Use CA Mainframe Software Manager (CA MSM) to acquire, install, deploy, and maintain your product. Although CA MSM is the preferred method for installing CA Technologies mainframe products, some sites may decide to use the Electronic Software Delivery (ESD) installation method instead.

Business Value:

CA MSM provides a web interface, which works with ESD and standardized installation, to provide a common way to manage CA Technologies mainframe products. You can use it to download, install, and deploy your product.

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. After installation, deployments allow system objects to be deployed across your enterprise's different environments.

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

Use Electronic Software Delivery

Although CA MSM is the preferred method for installing your CA Technologies mainframe products, some sites may decide to use the Electronic Software Delivery (ESD) method instead. For sites who have decided to use ESD, download the installation files from ca.com/support and install directly from your disk.

Business Value:

Using electronic software delivery (ESD) avoids ordering, shipping, and processing physical tape media to install the application. It is more timely, more cost-effective, and environmentally friendly. It uses standard z/OS utilities to prepare the product installation image on your system.

For information about how to download your CA Technologies products from the CA Support Online web site for installation using the enhanced ESD pax process, see the *Installation Guide.*

CA Common Services

Make sure you have installed the most current release of CA Common Services.

Business Value:

The latest release of CA Common Services contains the most current infrastructure updates, allowing you to successfully use the latest features and preventing potential errors that can occur from using out-of-date services.

For more information about CA Common Services, see the *Installation Guide*.

IBM APARs

Install all of the most recent IBM APARs appropriate for your environment.

Business Value:

If pertinent APARs are missing from your environment, it may impact the operation or performance of your CA Technologies products. We recommend that you review our current list of IBM APARs and apply only those that are appropriate to your environment.

Additional Considerations:

Be sure to review the system requirements in the *Installation Guide* and the upgrade information at support.ca.com. These locations explain any known conflicts or additional requirements you need to know when applying an APAR.

Installation in a Test Environment

Perform your installation and initial evaluations of a new release of the product and its components on a test system.

Business Value:

New releases of CA Technologies testing tools can always co-exist with previous releases, letting you test a new release on a test system while still running the older version on a production system. Evaluating the product in a test environment lets you detect any possible problems before you roll out the product to a production system.

Additional Considerations:

Always be sure to review any migration considerations in the *Installation Guide* before upgrading your CA Technologies product.

Common Symbolic Component

When installing more than one of the CA Technologies mainframe testing tools, we recommend that you install and maintain only one version of CA Technologies common symbolic component that is shared by all of the testing tool products.

Several testing tool products (CA InterTest Batch, CA InterTest for CICS, CA SymDump for CICS, CA SymDump Batch, CA Optimizer/II, and CA Mainframe Application Tuner) each distribute the common symbolic component in the VH FMID. This FMID was named CVHrr00 in previous versions and is now named CAVHrr0. The *rr* portion of the FMID contains the version number. For example, CAVH900 contains the symbolic common component provided with Version 9.0.

When installing any of these CA Technologies testing tool products, first query your Consolidated Software Inventory (CSI) to determine whether you already have a version of CVHrr00 or CAVHrr0 installed. If you use one CSI for all of your testing tool products then you need to perform only one query for these FMIDs. If you have one CSI for each testing tool product then you will need to query each CSI.

If you locate an installed version of CVHrr00 or CAVHrr0 during your query, compare the rr in your installed version of that FMID with the *rr* in the FMID provided with the product you are installing. Replace your existing FMID only if the *rr* in the product you are installing is higher than the *rr* in your existing function.

Business Value:

By installing and maintaining a single version of the CA Technologies common symbolic component, you reduce your maintenance effort, save disk space, and eliminate the possibility of executing symbolic utilities that are not up to date with the latest maintenance.

Additional Considerations:

The APPLY job of the installation places an SMP/E copy of the CAVH*rr*0 function in a library; however, many companies do not execute the testing tool products from this library. Therefore, always be sure to make the executables for the common symbolic component available for the testing tools products to use.

There are many ways to make the executables for the common symbolic component available, including:

- LNKLIST

  Always update your LNKLIST library after applying maintenance to CAVH*rr*0 if you have copied the executables for the common symbolic component from CAVH*rr*0 into a separate load library that you have added to your LNKLIST for all products to share.

- STEPLIBs
- CLISTs (for ISPF interfaces)

We recommend using the LNKLIST method because it has the specific advantage of making the executables available to all CA Technologies products with no additional changes to any JCL or CLISTs. If you cannot use LNKLIST for some reason (for example, your company has issues with updating LNKLIST or has specific rules prohibiting or limiting the use of LNKLIST), it is acceptable to add these executables to your STEPLIB or CLIST allocations. However, using STEPLIB or CLIST requires more initial setup and can be difficult to maintain if library names change.

Chapter 3: Configuration Best Practices

This section contains the following topics:

General Configuration Best Practices

The following topics describe general best practices to follow when configuring CA MAT.

- Allocating Data Sets
- Security Considerations
- Global Area Access

Allocating Data Sets

We recommend that you create a high-level qualifier in DSNPREFX. This parameter gives you a single high-level qualifier that is appended to each monitor data set as the first qualifier. Security can be defined with a universal access (UACC) of READ.

Business Value:

Defining this high-level qualifier helps simplify security tasks and the maintenance of the monitor data sets.

Additional Information:

Verify that the resulting data set is within the MVS limit of 44 characters. You can define SMS rules to migrate data sets beginning with this high-level qualifier. The MATUNER STC needs Create/Allocate authority.

Security Considerations

We recommend that you grant READ access to the MATUNER started task for load libraries. This authority enables information to be gathered about the programs that are observed during a measurement. The SECURITY parameter indicates which security points have to be tested and can have the following values:

NONE

Specifies that all security is turned off.

Note: If SECURITY is set to NONE, no user information is passed to other servers within the sysplex.

LOCAL

Specifies that security checks are performed only on requests within the same z/OS LPAR.

REMOTE

Specifies that security checks are performed for requests originating from other z/OS images through XCF.

BOTH

Specifies that security checks are performed regardless of where the request originates.

We also recommend that you set the parameter SECCRES to TARGET where your security environment permits application programs READ access to load libraries. Setting SECCRES to TARGET specifies that SECCRES inherits the access rights of the measured address space.

SECCRES specifies the security environment to be inherited by the CSECT resolution subtask. When CA MAT is measuring a job, all CSECT level resolution is performed in the CA MAT Server started task (STC) by a subtask. The CSECT level resolution subtask opens up load libraries (including libraries from LPALIST and LINKLIST) for READ access to obtain the CSECT level information for observed load modules.

Business Value:

The security configurations recommended in this best practice eliminate the occurrence of security violation messages during measurement and the need for creating specific security rules for CA MAT.

Additional Information:

SECCRES is only used for the global resolution of observed CSECTs. If internal CA MAT security is activated, a generic profile secprefx.** (for example, MATUNER.**) must be defined with a universal access (UACC) of READ. If the SAF facility class for CSVDYNL has been activated, this security environment must include READ access to CSVDYNL.

Global Area Access

Define security access to the global area of CA MAT using the UPDATE_GT/DS_CONTENTS system parameter. Define this access by the rules set in the security definitions (.SAF) or by listing specific user IDs.

Business Value:

Using this best practice protects the global area of CA MAT by limiting access to the area.

Reporting Best Practices

The following topics describe reporting best practices to follow when configuring CA MAT:

- Running Batch Reports on Measurement Data
- Defining Report Granularity
- Defining the Detail of Transaction Information in Batch Reports

Running Batch Reports on Measurement Data

Provide a valid job card and an optional JCLLIB statement that defines the library with the needed TUNBATCH procedure. A valid job card is required to run batch reports on measurement data. The default values are as follows:

```
JOB1=//JOBNAME JOB (ACCT),'CAMAT BATCH REPORT',
JOB2=//        CLASS=A,MSGCLASS=H,MSGLEVEL=(1,1),REGION=0M,
JOB3=//        NOTIFY=&SYSUID
JOB4=//TUNJCL JCLLIB ORDER=CAMAT_prefix.UTRSAMP
```

CA MAT stores the job card values in every measurement profile it creates.

Business Value:

This best practice enables you to request batch reports and create CSV files. These reports are a useful way to analyze measurement data.

Additional Information:

The JCLLIB in parameter JOB4 defines the UTRSAMP library that contains the member TUNBATCH. If TUNBATCH has been copied to a system procedure (PROCLIB) library, comment out the last line of the job card statement.

Defining Report Granularity

We recommend setting the values for ACTIVE_THRESHOLD, WAIT_THRESHOLD, and TRANSACTION_THRESHOLD to 1. This value specifies the minimum percentage of activity that will be provided in the batch reports.

Business Value:

This best practice provides more detailed information in batch report or CSV file compared to the default values. More detailed information is useful when analyzing measurement data.

Defining the Detail of Transaction Information in Batch Reports

We recommend setting the value for TRANCODE, TRANDELAY, and TRANDATA to ALL to provide detailed information about all transactions in the batch reports. These three parameters define the content in the report for transactions. If set to YES, information for the top five transactions is provided. If set to ALL full details for all transactions are provided.

Business Value:

This best practice provides reports with more detail, so you can perform more in-depth analysis on your CICS and IMS environments.

Symbolic Information Best Practices

The following topics describe best practices to follow when organizing symbolic information that is stored in your PROTSYM files and configuring PROTSYM parameters:

- Symbolic Files
- Dynamic Symbolic Support for CA Endevor SCM
- Retrieving and Registering Program Listings from PROTSYM
- Defining PROTSYM Library Data Set Names
- Designating the PROTSYM Library to Receive DSS Auto-Population

Symbolic Files

Before configuring your CA Technologies testing tool product, carefully consider how to best organize your symbolic information that is stored in your PROTSYM files. Well organized symbolic information allows the most effective and efficient access to your symbolics during testing.

There is no one correct way to organize symbolic information that works best for everyone. How you organize your symbolic information will always depend on your environments, your application systems, and the number of programs and programmers involved in your testing. Although there is no best organization that works for everyone, you should consider the following guidelines when planning for symbolic support:

- In a *production environment*, consider maintaining a separate PROTSYM file for each application system that can be shared by all of your programmers. Because only one copy of each program is typical in a production environment, the same symbolic information could be shared by any programmer debugging in that environment.

  In production environments, symbolic information is often limited to a single version of each program, and testing or debugging is typically less frequent and often limited to critical, real-time failures. Separate, shared PROTSYM files allow you to maintain a single set of symbolic information for each production application system that would be ready and available for access in the event of a failure.

- In a *test environment*, multiple versions of the same program are more likely to exist, and shared executable libraries are complimented by *sandbox* libraries that house test fixes or various works in progress. For this environment, a single shared symbolic repository is adequate for the shared executables, but additional symbolic repositories would be required for each individual programmer to use with their own testing environment. Allowing each programmer to own their own PROTSYM file provides them with the freedom to test using the correct symbolic information without impacting other members of the team.

Business Value:

Well organized symbolic information lets your application programmers easily locate, access, and share symbolic information for their programs. As a result, you could see better performance for your CA Technologies testing tools, optimal use of your DASD space, and the best possible productivity for your application programmers.

Additional Considerations:

As a general rule, think of your PROTSYM files as nothing more than a match for your load libraries, remembering that for each unique version of a load module you may require symbolic information for testing or debugging. Like the load libraries that exist at each level or stage of your development environment, the corresponding symbolic files should vary in size depending on the maximum number of programs expected at each level, with the individual PROTSYM files being the smallest by far.

If you are also licensed for CA Endevor Software Change Manager (CA Endevor SCM), CA testing tools provide an additional level of automation through dynamic symbolic support. This feature lets you populate your PROTSYM file dynamically on demand with symbolic information for programs built using CA Endevor SCM.

More Information:

For more information about creating, populating, maintaining, and using your PROTSYM symbolic repositories to enhance your testing, see the *CA Application Quality and Testing Tools Symbolic Guide*, which is provided with each CA Technologies testing tool product.

For more information about dynamic symbolic support, see the *Installation Guide.*

Dynamic Symbolic Support for CA Endevor SCM

If you are licensed for CA Endevor Software Change Manager (CA Endevor SCM), CA Technologies recommends using the dynamic symbolic support feature provided by the CA Technologies testing tools. This feature allows you to populate your PROTSYM file dynamically on demand with symbolic information for programs built using CA Endevor SCM.

CA Technologies recommends setting DYNREG_ENDEVOR to YES to enable dynamic symbolic support for CA Endevor SCM if the listing is not already registered or if there is a date/time stamp mismatch.

Because there are many ways to maintain the content of your PROTSYM files, you may encounter the scenario where you have multiple versions of the same program for which you need to locate and select the correct symbolic information. Dynamic symbolic support gives you the option to create your symbolic information only when it is needed, rather than always reloading your PROTSYM file after each compilation or assembly. This can save time, CPU utilization, and even DASD space.

Business Value:

The dynamic symbolic support feature is available with many of CA Technologies testing tool products including CA InterTest Batch, CA InterTest for CICS, CA SymDump Batch, and CA SymDump for CICS. This feature simplifies the setup requirements for using symbolic support, while simultaneously eliminating the majority of potential setup errors associated with it. By improving the accuracy of the content of your symbolic file and eliminating common setup errors, dynamic symbolic support can greatly increase the productivity of your programmers.

Setting DYNREG_ENDEVOR to YES activates the dynamic registering of compile listings from Endevor SCM. The manual registration of activities becomes unnecessary with the implementation of this feature.

Additional Considerations:

Dynamic symbolic support requires that you compile or assemble your applications using a CA Endevor SCM process with the footprint feature active. When symbolic information is required by one of the CA Technologies testing tools, your PROTSYM files are inspected first for the matching symbolic information. When no matching information is found, the executable is inspected for a CA Endevor SCM footprint. If this footprint is found, that footprint is used to locate the matching listing that is dynamically loaded into your specified PROTSYM file. The matching symbolic information is then used by your test session.

More Information:

For specific instructions related to setting up and using dynamic symbolic support, see the *Installation Guide* and the *User Guide*.

Retrieving and Registering Program Listings from PROTSYM

Set DYNREG_PROTSYM to YES to allow CA MAT to dynamically retrieve and register program listings from PROTSYM, if the listing is available.

Business Value:

Following this best practice makes computing addresses, displacements, monitoring address changes, and the manual registration of programs after recompilation unnecessary.

With a single preprocessing step, the same symbolic listings are dynamically available to CA InterTest, CA SymDump, CA Optimizer/II, and CA MAT users. Following this best practice automatically correlates the program source code to the CA MAT report.

Additional Information:

Manual registration of programs is still available. This best practice also applies to modules that consist of separately compiled programs that are link-edited together.

Defining PROTSYM Library Data Set Names

We recommend that you enter from one to eight fully qualified PROTSYM library data set names if DYNREG_PROTSYM=YES is specified.

Example: DYNREG_PROTSYM_LIB1='CAI.PROTSYM.LIBRARY'

Business Value:

This best practice may increase the productivity of CA Mainframe Application Tuner users because they do not need to spend time determining the names of data sets.

Designating the PROTSYM Library to Receive Dynamic Symbolic Support Auto-Population

When using dynamic registration (DYNREG_PROTSYM=YES and DYNREG_ENDEVOR=YES are both specified), designate which PROTSYM library name is the Dynamic symbolic support Auto-Population receiver.

Business Value:

Following this best practice makes computing addresses, displacements, monitoring address changes, and the manual registration of programs after recompilation unnecessary.

With a single preprocessing step, the same symbolic listings are dynamically available to CA InterTest, CA SymDump, CA Optimizer/II, and CA MAT users. Following this best practice automatically correlates the program source code to the CA MAT report.

Additional Information:

Manual registration of programs is still available. This best practice also applies to modules that consist of separately compiled programs that are link-edited together. Enter a number from one to eight corresponding to the PROTSYM library names to designate it as the receiver of Dynamic symbolic support Auto-Population.

DB2 Configuration Best Practices

The following topics contain best practice information for DB2 parameters that are set in member TUNSSP00 in the library UTRPARM:

- Using DB2 Harvester to Collect Detailed DB2 Information
- Setting the Granularity of SQL Statements
- Displaying Information about SQL Statements by Transaction

Using DB2 Harvester to Collect Detailed DB2 Information

We recommend using the DB2 Harvester (DB2HRVST=YES) to collect detailed DB2 information.

Business Value:

CA Mainframe Application Tuner is a sampling-based data collector for z/OS-based applications that measures delays incurred by an application program. You can use the DB2 Harvester technology to gather additional DB2 run-time statistics such as Total CPU, Total Response Time, Get Pages, and so forth. This extra DB2 data can provide useful information about the SQL activity, allowing you to make more informed decisions.

Setting the Granularity of SQL Statements

We recommend setting the values for DB2HVLOC, DB2HVIID, and DB2HVTHD to NO unless detail information is needed (for example, for measuring DDF address spaces). The parameters control the granularity of SQL statement information that is detected by the DB2 Harvester. If the parameters are set to YES, the location, operator, and thread IDs are collected and multiple entries for the same SQL statement with different values for these metrics are created. This configuration can result in a high number of different entries in the CA Mainframe Application Tuner DB2 view. If you set the value of these parameters to NO, no detail is provided.

Business Value:

This best practice enables you to see the source of your statements in one place, which saves you the time it would otherwise take to locate the source information.

Displaying Information about SQL Statements by Transaction

We recommend setting DB2HVCOR to YES to provide detailed information about the use of the SQL statement by transaction. In this context, the correlation ID is the transaction name.

Business Value:

This configuration provides additional granularity and is useful for measuring online regions. Implementing this feature correlates the execution of SQL statement metrics to transactions. Having additional granularity in your reports enables you to perform more detailed analysis on your SQL statements.

Web Server Configuration Best Practices

An Apache Tomcat web application server provides the interface used by the GUI to access CA Mainframe Application Tuner data which resides on the z/OS system.

The default settings are configured to run the server as a 31-bit Java application and this is the recommended setting. The server also can run as a 64-bit application, however it will have a larger memory footprint in 64-bit mode.

Your Workload Manager (WLM) policy should place the server in a service class which receives high priority, similar to that of an OMVS daemon. The server will spawn OMVS child processes on behalf of the GUI user to perform functions such as Analyze, Invoke and Stop. Although it is a common practice to define OMVS child processes in a multi-period service class, this may result in long response times, particularly when analyzing large monitor samples. Instead, we recommend that OMVS child processes for GUI users be defined in a single period service class with high importance.

If you have configured CA Mainframe Application Server to run on multiple LPARs in a sysplex, you need one Apache Tomcat web application server running on one of those LPARs. For more information about using sysplex, see the chapter "Post Installation and Customization Instructions" in the *Installation Guide*.

One of the requirements of using CA Mainframe Application Tuner in a sysplex environment is a shared Global Information File (GIF). In this release, the GUI works exclusively with Monitors defined and/or invoked in the GIF.

Once the sysplex environment is available, you must decide which CA Mainframe Application Tuner server on which LPAR you want GUI users to connect to. As long as each of the monitor profiles (that you want to use) are defined with an '*' in the System parameter for targeting systems in a Sysplex, CA Mainframe Application Tuner will monitor the job; the job may be monitored on any LPAR in the sysplex that has a CA Mainframe Application Tuner server running that participates in that sysplex group.

Business value

This best practice allows GUI users connected to a single CA Mainframe Application Tuner server to have access to all monitors in the sysplex. For more information, see Specify Target Systems for Parallel Sysplex in the chapter "Setting up a Monitor Definition" in the *CA Mainframe Application Tuner User Guide*.

Install the CA MAT Bookshelf on the Web Server

CA Technologies provides online access to the CA MAT product documentation at support.ca.com. However, this documentation is restricted to registered users of the CA support website. To make this documentation readily available to all users of a web site, the documentation bookshelf can be downloaded from the web site and installed on the CA MAT web server. This procedure requires you to have knowledge in using USS (Unix System Services).

Follow these steps:

1. Select the CA MAT documentation at support.ca.com
2. Download the bookshelf. This is a link available on the page. Click "Download the Bookshelf".

   The downloaded bookshelf is a ZIP file.
3. Rename the ZIP file to mathelp.zip
4. Determine the location of the CA MAT web server directory in USS. This is the value of SERVER_DIR used in the server start up options file CEESUOPT. The default is /cai/mat/server
5. Copy this ZIP file to the /cai/mat/server directory using any file transfer method, for example ftp.
6. Switch to OMVS or telnet into USS.
7. Change to /cai/mat/server/webapps:

   ```
   cd /cai/mat/server/webapps
   ```
8. Create a new directory named mathelp:

   ```
   mkdir mathelp
   ```
9. Check if a directory name ROOT exists:

   ```
   ls –l
   ```

   ROOT may be on the list.

   If ROOT does not exist, create it:

   ```
   mkdir ROOT
   ```
10. Move the ZIP file to the ROOT directory:

    ```
    mv /cai/mat/server/mathelp.zip /cai/mat/server/webapps/ROOT
    ```
11. Change to /cai/mat/server/webapps/mathelp directory:

    ```
    cd /cai/mat/server/webapps/mathelp
    ```

Now you must unzip the mathelp.zip file into the mathelp directory.

As USS does not provide an unzip program, you must use the java tool named jar. To enable this tool to work, the java settings must be properly set. A valid version of java must be in your PATH, and JAVA_HOME must be set. To test if this is done, simply type jar. You may either get a list of help output from jar or see the message *FSUM7351 not found*. If the jar is not found you cannot proceed.

12. Use the following command to unzip the file:

    ```
    jar xf ../ROOT/mathelp.zip
    ```

    The mathelp.zip file is now unzipped into the mathelp directory.

13. Check the contents of the mathelp directory using the following command:

    ```
    ls –l
    ```

    The contents of the mathelp directory may appear in the same way as the following text:

    ```
    -rw-r--r--   1 MATUSER  OMVSDFG    11670 Nov  5  2011 Bookshelf.html
    drwxr-xr-x   4 MATUSER  OMVSDFG     8192 Mar 29 11:07 Bookshelf_Files
    -rw-r--r--   1 MATUSER  OMVSDFG     4670 Mar 29 11:07 Bookshelf_Metadata.xml
    -rw-r--r--   1 MATUSER  OMVSDFG       56 Mar 29 11:07 message.log
    -rw-r--r--   1 MATUSER  OMVSDFG     5233 May 27  2011 readme.html
    ```

14. Test the install using the CA MAT web server. Use the following URL:

    ```
    http://host:port/mathelp/Bookshelf.html
    ```

    The online documentation appears successfully.

Chapter 4: Parameter Recommendations When Using Performance Management Assistant

This section contains the following topics:

Include All Samples

The recommended value for NORMBAT is NORMAL. This value ensures that all samples with a status of ACTIVE or WAIT are extracted for use in the measurement reports imported into the Performance Management Assistant (PMA).

Business Value:

Configuring this parameter as recommended imports all measurement information gathered for the target application into the PMA performance database. This configuration can affect the data mining feature in PMA, which allows you to search and cross-reference stored measurement information. The most current measurement information of a job step search is accumulated and made available for cross-referencing.

Including Detailed CPU Use Information

We recommend setting TACTPCT to 0.05 to include detailed information about CPU usage.

Business Value:

Using this best practice keeps visibility to the information about components using CPU resources because of the default thresholds. This configuration can affect the data mining feature in PMA, which allows you to search and cross-reference stored measurement information. The most current measurement information of a job step search is accumulated and made available for cross-referencing.

Providing Detailed CPU Information

We recommend setting TWAITPCT to 0.1 to provide details about wait situations.

Business Value:

Configuring this parameter as recommended imports all measurement information gathered for the target application into the PMA performance database. This configuration can affect the data mining feature in PMA, which allows you to search and cross-reference stored measurement information. The most current measurement information of a job step search is accumulated and made available for cross-referencing.

Providing Information about CPU Resources and Waits

We recommend setting TDETPCT to 0.05. This value provides detailed information about the combined percent values for CPU and Wait.

Business Value:

Configuring this parameter as recommended imports all measurement information gathered for the target application into the PMA performance database. This configuration can affect the data mining feature in PMA, which allows you to search and cross-reference stored measurement information. The most current measurement information of a job step search is accumulated and made available for cross-referencing.

Providing Transaction Details

We recommend setting TRANPCT to 0.05. This value provides details about executed transactions.

Business Value:

Configuring this parameter as recommended imports all measurement information gathered for the target application into the PMA performance database. This configuration can affect the data mining feature in PMA, which allows you to search and cross-reference stored measurement information. The most current measurement information of a job step search is accumulated and made available for cross-referencing.

Displaying Detailed Information in Code View Chapters

We recommend setting TCODPCT to 0.05 to display detailed information in the Code View chapters.

Business Value:

Configuring this parameter as recommended imports all measurement information gathered for the target application into the PMA performance database. This configuration can affect the data mining feature in PMA, which allows you to search and cross-reference stored measurement information. The most current measurement information of a job step search is accumulated and made available for cross-referencing.

Creating Detailed Code Use Information

We recommend setting TRANDET to A to produce detailed reports for all transactions. The default value of Y creates detailed code usage information for only the top five transactions within a measurement.

Business Value:

Configuring this parameter as recommended imports all measurement information gathered for the target application into the PMA performance database. This configuration can affect the data mining feature in PMA, which allows you to search and cross-reference stored measurement information. The most current measurement information of a job step search is accumulated and made available for cross-referencing.

Creating Delay Categories

We recommend setting TRANDLY to A to produce reports for all transactions. The default value of Y creates delay categories for only the top five transactions within a measurement.

Business Value:

Configuring this parameter as recommended imports all measurement information gathered for the target application into the PMA performance database. This configuration can affect the data mining feature in PMA, which allows you to search and cross-reference stored measurement information. The most current measurement information of a job step search is accumulated and made available for cross-referencing.

Additional Considerations:

The default value of Y creates delay categories for only the top five transactions within a measurement.

Displaying Data Use Information

We recommend setting TRANDAT to A to produce reports for all transactions. The default value of Y creates data usage information for the top five transactions within a measurement.

Business Value:

Configuring this parameter as recommended imports all measurement information gathered for the target application into the PMA performance database. This configuration can affect the data mining feature in PMA, which allows you to search and cross-reference stored measurement information. The most current measurement information of a job step search is accumulated and made available for cross-referencing.

Chapter 5: Graphical User Interface Customization Best Practices

This section contains the following topics:

Managing Software Installation

The Graphical User Interface (GUI) tools provided for CA Mainframe Application Tuner (CA MAT) is provided as a Microsoft Windows Installer file (MSI) published on the mainframe based web server deployed as part of CA Mainframe Application Tuner. Any subsequent changes to the GUI are provided as a Windows Installer patch file (MSP). This image is available by accessing a web page containing all available product install images and the latest patch (PTF), if any. The URL for this web page is http://host:port/software.

The files are downloaded as ZIP files that contain a single executable (.EXE). You must double-click the .EXE file to execute (run) it. In the case of an install, this .EXE file is a self-extracting program that extracts the necessary Microsoft Installer file (MSI) and command script needed to invoke the install for the product Part of this script checks if an existing version of the product exists. If so, it gets uninstalled without any user intervention and then the program install dialog begins. In the case of a patch, this .EXE is a self-extracting program that extracts the necessary Microsoft Installer patch file (MSP) and command script needed to patch an installed product. If the program is installed, the patch will be applied.

For users who want to integrate the CA MAT GUI components into an existing Eclipse platform IDE, the web site of the server acts as an Eclipse update site. The URL for this web page is http://host:port/software.

Your site can deploy this software using alternate software delivery technologies and may also prefer to inhibit access to the web server based install software. The web server provides a document containing details on how to do this action at this location: http://host:port/software/MATInstallBP.exe. Download this file to a Windows PC and execute the file. A detailed document with descriptions about the various options is available. This file is provided as .EXE file as it requires the reader to agree to CA's End User License Agreement (EULA).

Customized Cheat Sheets

The CA MAT GUI is a program developed using the Eclipse plugin framework provided by the Eclipse Foundation (http:/eclipse.org). The CA MAT GUI contains tutorials that explain the use of the GUI. These tutorials are provided as Eclipse Cheat Sheets. Users are free to develop and distribute their own customized cheat sheets using the Eclipse development platform (Eclipse IDE). The IDE can be downloaded from http:/eclipse.org.

After the cheat sheets have been developed, they can be deployed to users as standalone files, shared files, or installed on a webserver.

The information necessary to implement a cheat sheet is provided as part of the Eclipse IDE and is beyond the scope of this document.

Customized Online Help

The online help provided with the GUI best utilizes the online help capability of the Eclipse Framework. Users can extend the help available by developing help using the Eclipse framework. However, distribution of help files is more complex than that for cheat sheets. The online help must be packaged as an eclipse plugin for distribution.

The information necessary to implement a plugin is provided as part of the Eclipse IDE and is beyond the scope of this document.

Index

U

W