CA SOLVE:Operations® Automation

Best Practices Guide

Release 11.9

CA Technologies Product References

This document references the following CA Technologies products:

- CA Mainframe Connector for Linux on System z (CA Mainframe Connector)
- CA Mainframe Software Manager™ (CA MSM)
- CA SOLVE:Operations® Automation
- CA SOLVE:Operations® Automation for CICS

Contact CA Technologies

Contact CA Support

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

- 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.

If you would like 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

This section contains the following topics:

Purpose of this Guide

The guide provides a brief introduction to CA's Mainframe 2.0 strategy and features, and describes the best practices for installing and configuring CA SOLVE:Operations Automation.

Audience

The intended audience of this guide is systems programmers and administrators who install, configure, deploy, and maintain CA SOLVE:Operations Automation.

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 browser-based user interface (UI) 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.

CA MSM provides software acquisition and installation that make it easier for you to obtain and install CA mainframe products, and apply the recommended maintenance. The services within CA MSM enable you to manage your software easily based on industry accepted best practices. The common browser-based UI makes the look and feel of the environment friendly and familiar.

We follow the IBM z/OS packaging standards using SMP/E, with some additional CA 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 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 that identifies the component parts of a product and user-supplied input that identifies the deployment criteria, such as where it will go and what it will be called.

Software Configuration Service (SCS)

Facilitates the configuration of mainframe products 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 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 CA mainframe 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 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 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 product's address space, treating this as an indication that there is either a potential problem with the CA product's 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 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 products when they are starting, up, stopping, or down. Without these consistent type of events, SSM must maintain separate rules in CA OPS/MVS for each product's unique messages 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

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 SOLVE:Operations Automation.

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*.

Deployment

Use CA MSM to deploy your product.

Business Value:

CA MSM provides a web interface to provide a common way to manage CA mainframe products. After you use it to download and install CA SOLVE:Operations Automation, it helps you deploy the product using the same interface.

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*.

How You Deploy the Product

After you install the product, you deploy the regions to the target systems. The process includes the following stages:

1. Configure your regions using the Install Utility. The following command starts the utility:

   `EXEC 'dsnpref.OPB9.CC2DJCL(INSTALL)'`

   During this stage, the utility generates a series of jobs in the *dsnpref*.OPB9.*rname*.JCL data set, including the following jobs that help with deployment: S10DUMP and S11REST.

2. For each region you configure, submit the S10DUMP job. The job creates the following backup data sets that include the configuration files for the region: *dsnpref*.DFDSS.LOCAL (containing files specific to a region) and *dsnpref*.DFDSS.SHARED (containing files that multiple regions can share).

3. Deploy each region to the target system using CA MSM. Include the data sets created in Step 2 and the *dsnpref*.OPB9.*rname*.JCL data set as custom data sets.

4. For each region you deploy, submit the S11REST job on the target system. The job restores the region files in the backup data sets.

How You Deploy Product Maintenance

After you install product maintenance, you deploy updated software to the regions on the target systems. The process includes the following stages:

1. (Optional) If maintenance exists for VSAM data sets (as indicated by HOLDDATA), apply the maintenance using the Install Utility. Select option 8, V.

   Note: If no VSAM maintenance exists, Option V is not selectable.

2. (Optional) If maintenance exists for VSAM data sets, submit the F22DUMP job.
3. Shut down the regions you want to maintain.
4. Deploy product maintenance to each target system using CA MSM. If there is VSAM maintenance, include the *dsnpref*.DFDSS.SHARED and *dsnpref*.OPB9.*rname*.JCL data sets as custom data sets.

   Note: We recommend that you back up the existing SHARED data set on a target system before you perform the deployment.

5. (Optional) If maintenance exists for VSAM data sets, review the F23REST job on a target system and customize as required. Submit the job to restore the files in the *dsnpref*.DFDSS.SHARED data set only.
6. Restart the regions.

Address Space Sharing

If your site also uses CA SOLVE:Operations Automation for CICS, share the address space with CA SOLVE:Operations Automation for CICS for performance and usability optimization. To share the address space, set up a region that includes the products.

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 SOLVE:Operations Automation 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 SOLVE:Operations Automation 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 with some essential parameters updated during setup. Further customization can be performed progressively.

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 automation and monitoring environment over time. Add local knowledge to the setup. Modify the discovered resources so that you are not monitoring things that are not critical to your business.

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 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 resources that are present at the time it runs.

What Express Setup discovers can only be a starting point for your automation and monitoring environment. Express Setup does not know your business.

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 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 *Administration Guide*.

Standardized Object Naming

Name resources in a standard way.

Business Value:

Presentation, management, and control of CA SOLVE:Operations Automation objects (resources, rule sets, and rules) is better when objects are organized and named in a structured and standard way. A naming standard makes the product easier to use (less training with less user error). Additionally, generic automation is easier when dealing with objects that conform to a standard, reducing the effort and maintenance required for building automation.

Example: EventView Rule Set Names

In EventView, you can implement multiple rule sets and include them in a primary rule set. The following naming organizes the rule sets in a structured way. The standard uses the first part of the name to identify the type of resource and the second part of the name to identify the rule action. For example, the CICSAUTO rule set automates actions in response to CICS messages, the ZOSENH rule set enhances the presentation of z/OS messages, and the ZOSSUPP suppresses z/OS messages.

```
PRIMARY
   CICSAUTO
   CICSEHN
   CICSSUPP
   ZOSAUTO
   ZOSENH
   ZOSSUPP
```

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.

Audit

If your site has a requirement for government regulatory accounting and auditing (Sarbanes-Oxley) compliance, recordkeeping, and corporate transparency, enable the Audit feature and generate the Audit events as system management facilities (SMF) records.

Business Value:

In addition to delivering regulatory compliance, activity tracking, and reporting, the Audit feature assists in documenting the level of benefit that CA SOLVE:Operations Automation is providing to your organization such as:

- Number of messages suppressed and modified
- Number of EventView rules that are hit
- Number of EventView timers that are triggered
- Number of system commands issued
- Number of resource state changes seen by automation
- Number of alerts of each severity seen by automation

Additional Considerations:

The AUDIT parameter group enables auditing; the SMF parameter group lets you generate Audit events as SMF records. If you have special requirements, you can add site-specific auditing using the $NMAUAPI Audit application program interface (API).

More Information:

For information about how to enable auditing, see the *Administration Guide*. For information about the API, see the *Reference Guide*.

Online Help

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

Business Value:

CA SOLVE:Operations Automation 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 SOLVE:Operations Automation integrates with the following CA products:

- **CA SOLVE:Operations Automation for CICS**—CA SOLVE:Operations Automation and CA SOLVE:Operations Automation for CICS 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 monitoring and control of events and resources irrespective of whether they are CICS or system related.
- **CA Service Desk**—CA SOLVE:Operations Automation supports the automatic creation of trouble tickets in CA Service Desk, facilitating problem notification and resolution.

Multisystem Deployment

If you have multiple systems, deploy CA SOLVE:Operations Automation 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.

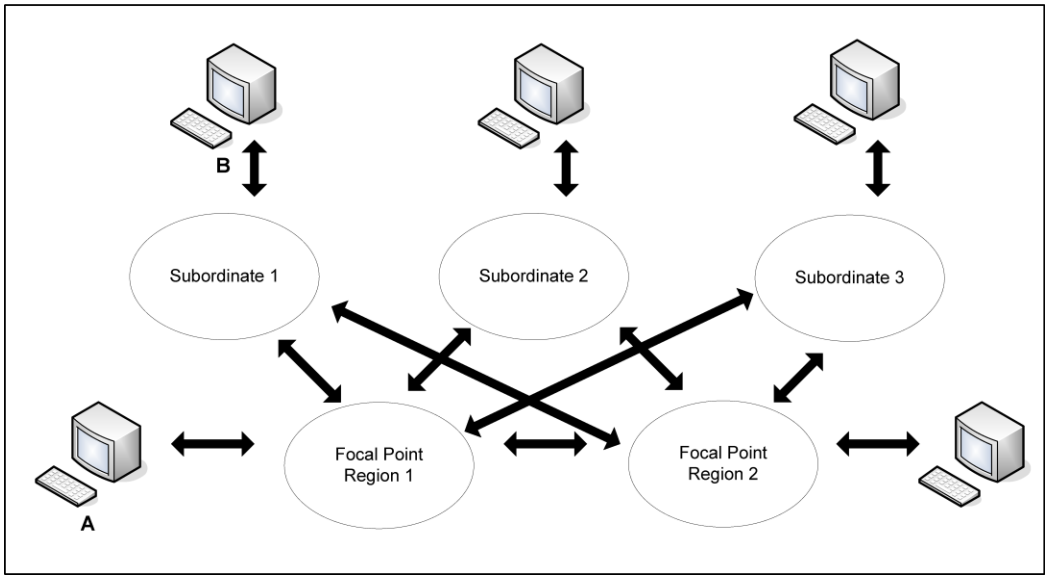

Business Value:

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 28). 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 29)

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 SOLVE:Operations Automation 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=`*`xxxx`*`INI'`

   ***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 SOLVE:Operations Automation. 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 LOGFILES parameter group using the &SYSNAME system symbol:

```
.LOGFILES
   &$IAPLOGPR    = &STR $LOPROC
   &$IAPLOGVW    = &STR $LOBROW
   &$IAPLOGF1    = &STR NMLOG01
   &$IAPLOGD1    = &STR NETW.NM.NETM&SYSNAME.NMLOG01
   &$IAPLOGO1    = &STR LSR SIS DEFER
   &$IAPLOGI1    = &STR SHR
   &$IAPLOGF2    = &STR NMLOG02
   &$IAPLOGD2    = &STR NETW.NM.NETM&SYSNAME.NMLOG02
   &$IAPLOGO2    = &STR LSR SIS DEFER
   &$IAPLOGI2    = &STR SHR
   &$IAPLOGF3    = &STR NMLOG03
   &$IAPLOGD3    = &STR NETW.NM.NETM&SYSNAME.NMLOG03
   &$IAPLOGO3    = &STR LSR SIS DEFER
   &$IAPLOGI3    = &STR SHR
```

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 SOLVE 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 those data sets that are 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 23)

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: Automation Best Practices

This section contains the following topics:

Resource Definition Templates

Automate resource management by defining resources using templates.

Business Value:

Templates help you to define your resources consistently. A template also helps you to apply changes to all resources using that template quickly. These features help ensure that your resources are automated and managed correctly for your business requirements.

Additional Considerations:

The product comes with a set of templates that you can use. You can also create your own templates.

When you use Express Setup Facility or the Auto Populate Facility, resources are defined automatically using templates.

More Information:

The *Administration Guide* describes the Express Setup Facility, the Auto Populate Facility, and how you can define your own templates.

Management of Linux Resources

Note: This topic applies if you are using CA SOLVE:Operations Automation and CA Mainframe Connector to manage your Linux resources on z/VM.

Enable logging of z/VM and Linux messages to help you learn about your Linux resources.

Business Value:

Knowing the messages generated by your Linux resources helps ensure that you define the resources appropriately for your business requirements.

Additional Considerations:

The following process helps you implement your Linux management environment:

1. In the LINUXCONNECT parameter group, enable message logging and dynamic discovery, but disable the loading of z/VM system images.
2. After the resources are discovered, review the Linux and z/VM messages in the activity log.
    a. If you find a message you want to use for the discovered resources, update the template used for the resources and apply the template.
    b. Look for messages relevant to the Linux applications you want to manage. Use the messages to help you define your application resources.
3. When your resource definitions are ready, assign a home system to the z/VM system image definition to load the system image.
4. (Optional) You can disable the logging of Linux and z/VM messages.
5. (Optional) You can enable the loading of discovered z/VM system images. When your site configures a Linux system on another z/VM system for this region to manage, the region defines and loads the new image.

More Information:

The *Linux Management Guide* describes how to configure the Linux management environment, and work with the defined images and resources.

Management of Linux Applications That Do Not Write to Syslog

Note: This topic applies if you are using CA SOLVE:Operations Automation and CA Mainframe Connector to manage your Linux resources on z/VM.

If you want to manage a Linux application that does not write to syslog, run the application from a shell script that includes message logging.

Business Value:

With message logging, you can define resources to capture the logged messages and take appropriate actions. Being able to see what your applications are doing helps ensure the availability of your business processes.

Additional Considerations:

To see the startup and termination of the application, include the following commands in the shell script:

```
logger application_name starting
application_name
logger application_name terminated
```

Through CA Mainframe Connector, CA SOLVE:Operations Automation sees the following messages:

```
LXLOG001I linux_system vm_host_node application_name starting
```

```
LXLOG001I linux_system vm_host_node application_name terminated
```

You can define a resource definition for the application to automate actions in response to these messages. Use the $MN- special message prefix in the definition to capture the messages.

z/VM System Shutdown

Note: This topic applies if you are using CA SOLVE:Operations Automation and CA Mainframe Connector to manage your Linux resources on z/VM.

Define a resource for each z/VM system that hosts the Linux systems you want to manage. Define parent-child relationships between the z/VM system and the guest Linux systems.

Business Value:

The implementation enables you to use CA SOLVE:Operations Automation to perform an orderly shutdown of a z/VM system. The relationships prevent the shutdown of the z/VM system while a guest Linux system is still active.

Additional Considerations:

In the z/VM system image, define a VMGST-class, HOST-type z/VM system resource using the VMHOST template. Relate the Linux system resources in the image as children of the z/VM system resource. The VMHOST inactivation processing uses the CP SHUTDOWN z/VM command to shut down the z/VM system.

To perform an orderly shutdown of the z/VM host and the Linux systems, you can issue the SHUTSYS command against the system image. To restart the z/VM system, you perform the action manually outside of CA SOLVE:Operations Automation.

If you want CA SOLVE:Operations Automation to maintain control over the shutdown and restarting process, use the CP SHUTDOWN REIPL z/VM command instead for the inactivation processing. Also, add the following state change exits:

State Type	Change From	Change To	Process	Parameters
ACTUAL	ANY	STOPPING	$NCL	$NCL=$RMCONS LCMD=MM OBJID='&ZRMDBOBJID'
DESIRED	INACTIVE	ACTIVE	$NCL	$NCL=$RMCONS LCMD=MR OBJID='&ZRMDBOBJID'

You can then use STARTSYS to return the resources in the system image to their desired states.

System Hardware Message Processing

Note: This topic applies if you are using CA SOLVE:Operations Automation and Hardware Interface Service to process system hardware messages.

Capture and review RMHIS* messages to help you learn about your system hardware.

Business Value:

Knowing the messages generated by your system hardware helps you automate responses that are appropriate for your business requirements.

Additional Considerations:

By default, RMHIS* messages do not go to the activity log.

Follow these steps:

1. Update the Log HISRV Messages? field in the HISRV parameter group to enable logging.

   If you find too many messages are logged to the activity log, use an alternate method that allows more control of what is logged:

   - Leave Log HISRV Messages? set to NO.
   - Define an RMHIS EventView message rule to log RMHIS* messages in the activity log. On the Message Actions panel, specify the following OCS command: **LOG &ZMSGTEXT**.
   - If you want to ignore certain RMHIS* messages, define specific rules that do nothing for those messages.

2. Review the RMHIS* messages in the activity log periodically for possible automation:

   - Where a system hardware event indicates a problem and a standard (manual) process exists to address the problem, then this process can be automated.
   - Events associated with an LPAR becoming disabled can be of use. Events of this nature depend on the way you implement your LPAR hierarchy.
   - Events associated with an IPL can be used in the same way.
   - Environmental alerts associated with critical hardware (CP, zIIP, and so on) can be analyzed and then intercepted to provide serviceability reporting for specific hardware items.

3. For identified automation opportunities, define specific message rules to perform the required actions.

4. (Optional) Disable logging when you are satisfied you have captured all messages of interest.

Message Burst Protection

If you have many multiline messages with the same major line but different minor lines, and you want to use some of those messages to trigger actions, disable burst protection for multiline messages.

Business Value:

Multiline messages with the same major line are not suppressed because of burst protection. Important messages are detected and acted on to help ensure the smooth running of your business.

Additional Considerations:

Message burst protection prevents a region from being overloaded by the same message arriving many times over a short period. The AOMQUEUES parameter group specifies the protection criteria. For multiline messages, the region uses the major line to determine whether the received messages are the same. Therefore, burst protection can suppress multiline messages with the same major line but different minor lines, causing the region to miss relevant messages.

Chapter 4: Monitoring Best Practices

This section contains the following topics:

Filters and Formats

Apply filters and formats to the Alert Monitor and the various status monitors to minimize the noise caused by unwanted information.

Business Value:

Filters and formats enable operators and system administrators to focus on their area of responsibility without being distracted by extraneous data. This leads to a rapid and focused response to any problem condition that arises.

Additional Considerations:

To apply a filter or a format to a monitor, enter **FILTER** or **FORMAT** in the primary command field, then select the appropriate filter or format from the list. Administrators can define new filters and formats for the Alert Monitor from the Alert Monitor Administration Menu (/ALADMIN), and for the status monitors from the Automation Services Administration Menu (/ASADMIN).

More Information:

For information about how to define filters and formats, see the *Administration Guide* and the *Reference Guide*.

Index

T

U

Z