

RS/6000 Cluster Technology



# Event Management Programming Guide and Reference



RS/6000 Cluster Technology



# Event Management Programming Guide and Reference

**Note!**

Before using this information and the product it supports, be sure to read the general information under "Notices" on page 227.

**Second Edition (April 2000)**

This edition applies to:

- Version 3 Release 2 of the IBM Parallel System Support Programs for AIX (PSSP) Licensed Program (product number 5765-D51) and to all subsequent releases and modifications until otherwise indicated in new editions, and
- The Enhanced Scalability feature of Version 4 Release 3 of the IBM High Availability Cluster Multi-Processing for AIX (HACMP) Licensed Program (product number 5765-D28) and to all subsequent releases and modifications

Order publications through your IBM representative or the IBM branch office serving your locality. Publications are not stocked at the address below.

IBM welcomes your comments. A form for readers' comments may be provided at the back of this publication, or you may address your comments to the following address:

International Business Machines Corporation  
Department 55JA, Mail Station P384  
2455 South Road  
Poughkeepsie, NY 12601-5400  
United States of America

FAX (United States & Canada): 1+914+432-9405

FAX (Other Countries):

Your International Access Code +1+914+432-9405

IBMLink (United States customers only): IBMUSM10(MHVRCFS)

IBM Mail Exchange: USIB6TC9 at IBMMAIL

Internet e-mail: mhvrdfs@us.ibm.com

If you would like a reply, be sure to include your name, address, telephone number, or FAX number.

Make sure to include the following in your comment or note:

- Title and order number of this book
- Page number or topic related to your comment

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1998, 2000. All rights reserved.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

<b>Figures</b> . . . . .	v
<b>About This Book</b> . . . . .	vii
Checking Which Level(s) of Code You Have . . . . .	vii
PSSP . . . . .	vii
HACMP . . . . .	viii
RSCT . . . . .	viii
Who Should Use This Book . . . . .	viii
Typographic Conventions . . . . .	viii
<b>Chapter 1. Understanding Event Management</b> . . . . .	1
Introducing Event Management . . . . .	1
The Event Management Components . . . . .	1
Communicating across Nodes . . . . .	3
Event Management and Domains . . . . .	3
How Can You Use Event Management Services? . . . . .	4
Writing Resource Monitors . . . . .	5
Defining the Resource Data . . . . .	5
Choosing the Resource Monitor Type . . . . .	10
Configuring the Resource Data and Resource Monitor . . . . .	12
Coding and Testing the Resource Monitor . . . . .	13
Writing Event Management Clients . . . . .	17
Finding Out What Resource Data Is Available . . . . .	17
Using Expressions to Define Events . . . . .	20
Coding and Testing the EM Client . . . . .	22
Event Management Performance Considerations . . . . .	25
Resource Monitors and PTX Shared Memory . . . . .	25
DDS Shared Memory . . . . .	26
Releasing Shared Memory Manually . . . . .	26
<b>Chapter 2. Event Management Configuration Data Reference</b> . . . . .	27
Event Management Configuration Data (emcdb) . . . . .	28
<b>Chapter 3. Event Management Subroutine Reference</b> . . . . .	39
EMAPI Subroutine Summary . . . . .	39
RMAPI Subroutine Summary . . . . .	39
ha_em_end_session Subroutine . . . . .	41
ha_em_get_ecgid Subroutine . . . . .	43
ha_em_receive_response Subroutine . . . . .	44
ha_em_restart_session Subroutine . . . . .	53
ha_em_send_command Subroutine . . . . .	56
ha_em_start_session Subroutine . . . . .	63
ha_rr_add_var Subroutine . . . . .	66
ha_rr_del_var Subroutine . . . . .	69
ha_rr_end_session Subroutine . . . . .	72
ha_rr_get_ctrlmsg Subroutine . . . . .	74
ha_rr_get_interval Subroutine . . . . .	78
ha_rr_init Subroutine . . . . .	80
ha_rr_makserv Subroutine . . . . .	83
ha_rr_reg_var Subroutine . . . . .	85
ha_rr_rm_ctl Subroutine . . . . .	88
ha_rr_send_val Subroutine . . . . .	91
ha_rr_start_session Subroutine . . . . .	93

ha_rr_terminate Subroutine . . . . .	96
ha_rr_touch Subroutine . . . . .	98
ha_rr_unreg_var Subroutine . . . . .	100
<b>Chapter 4. Event Management Files Reference . . . . .</b>	<b>103</b>
EMAPI Errors (err_emapi) . . . . .	104
ha_emapi.h File. . . . .	110
Expressions (haemexpr) . . . . .	123
Resource Variables (haemrvars) . . . . .	126
RMAPI Errors (err_rmap) . . . . .	131
ha_rmap.h File. . . . .	136
<b>Chapter 5. Using the RMAPI: Some Resource Monitor Examples . . . . .</b>	<b>141</b>
The rmap_smpcmd.c Sample Program . . . . .	141
The rmap_smpdae.c Sample Program . . . . .	148
The rmap_smpsig.c Sample Program . . . . .	161
The rmap_smp.msg Message File. . . . .	173
The rmap_smp.loadsdr Shell Script . . . . .	174
The rmap_smp.unloadsdr Shell Script . . . . .	178
<b>Chapter 6. Using the EMAPI: Some Event Management Client Examples . . . . .</b>	<b>181</b>
The emapi_v02_ex01.c Sample Program . . . . .	181
The emapi_v02_ex02.c Sample Program . . . . .	194
The emapi_v02_ex03.c Sample Program . . . . .	206
The emapi_v02_ex04.c Sample Program . . . . .	219
<b>Notices . . . . .</b>	<b>227</b>
Trademarks . . . . .	228
Publicly Available Software . . . . .	229
Software Update Protocol . . . . .	229
<b>Glossary of Terms and Abbreviations. . . . .</b>	<b>231</b>
<b>Bibliography . . . . .</b>	<b>239</b>
Information Formats . . . . .	239
Finding Documentation on the World Wide Web. . . . .	239
Accessing PSSP Documentation Online. . . . .	239
Manual Pages for Public Code . . . . .	240
RS/6000 SP Planning Publications. . . . .	240
RS/6000 SP Hardware Publications . . . . .	240
RS/6000 SP Switch Router Publications. . . . .	240
RS/6000 SP Software Publications . . . . .	241
AIX and Related Product Publications . . . . .	243
DCE Publications . . . . .	243
Red Books . . . . .	243
Non-IBM Publications . . . . .	243
<b>Index . . . . .</b>	<b>245</b>
<b>Readers' Comments — We'd Like to Hear from You . . . . .</b>	<b>253</b>

---

## Figures

1.	An Overview of the Event Management Components . . . . .	3
2.	Examples of Structured Byte String Definitions . . . . .	7
3.	Examples of Resource Variable Names and Resource IDs . . . . .	8
4.	Examples of Resource Variable Classes . . . . .	10
5.	A Pseudocode Outline of a Command-Based Resource Monitor . . . . .	14
6.	A Pseudocode Outline of a Daemon-Based Resource Monitor . . . . .	16
7.	Show Resource Variable Details Dialog Box . . . . .	20



---

## About This Book

This book contains conceptual, guidance, and reference information to help you write programs that use the Event Management APIs that are part of RS/6000 Cluster Technology (RSCT). RSCT is provided by IBM Parallel System Support Programs for AIX (PSSP) and by the Enhanced Scalability feature of IBM High Availability Cluster Multi-Processing for AIX (HACMP/ES). Specifically, this book contains information to help you write the following types of programs:

- Resource monitors, using the Resource Monitor Application Programming Interface (RMAPI)
- Event Management client programs, using the Event Management Application Programming Interface (EMAPI).

In addition, this book contains information about how to configure the Event Management subsystem.

For a list of related books and information about accessing online information, see the bibliography in the back of the book.

---

## Checking Which Level(s) of Code You Have

This book applies to PSSP Version 3 Release 2 and to HACMP/ES at the HACMP Version 4 Release 3 level.

It also applies to RSCT Version 1.2.

To find out which level(s) of code you have running on your system, follow the instructions in the sections below.

### PSSP

To find out what version of PSSP is running on your control workstation (node 0), enter the following:

```
sp1st_versions -t -n0
```

In response, the system displays something similar to:

```
0 PSSP-3.2
```

If the response indicates **PSSP-3.2**, this book applies to the version of PSSP that is running on your system.

To find out what version of PSSP is running on the nodes of your system, enter the following from your control workstation:

```
sp1st_versions -t -G
```

In response, the system displays something similar to:

```
1 PSSP-3.2  
2 PSSP-3.2  
7 PSSP-3.1.1  
8 PSSP-2.4
```

If the response indicates **PSSP-3.2**, this book applies to the version of PSSP that is running on your system.

If you are running mixed levels of PSSP, be sure to maintain and refer to the appropriate documentation for whatever versions of PSSP you are running.

## HACMP

To find out which version of HACMP is running on a particular node, enter the following:

```
lslpp -L | grep cluster.es.server.rte
```

## RSCT

To find out which code version of IBM RS/6000 Cluster Technology is running on a particular node, enter the following:

```
lslpp -L rsct.basic.rte
```

---

## Who Should Use This Book

This book is intended primarily for programmers of applications that manage system resources (which may or may not be subsystems) who want to use Event Management services to make their applications highly available. This book contains information for programmers who want to write new resource monitors, write new clients that use the EMAPI, or add the use of Event Management services to existing programs.

It assumes that you are an experienced C programmer and have a thorough understanding of RS/6000 SP (SP) hardware, IBM Parallel System Support Programs for AIX (PSSP) software and/or the Enhanced Scalability feature of IBM High Availability Cluster Multi-Processing for AIX (HACMP/ES) software, and AIX operating system fundamentals. In particular: if you are using RSCT on PSSP, you should be familiar with the operation of the SP control workstation, PSSP system partitions, and the PSSP System Data Repository (SDR); if you are using RSCT on HACMP/ES, you should be familiar with HACMP/ES user-defined event processing. It also assumes that you are familiar with the management of system resources and recovery concepts.

The Event Management subsystem requires that the **perfagent.tools** fileset be installed on the Control Workstation and on each node of the RS/6000 SP where PSSP Version 3.1 is installed, or on each node of an HACMP/ES cluster where the HACMP Version 4.3 level of HACMP/ES is installed. The **perfagent.tools** fileset is shipped with AIX. Refer to *PSSP: Installation and Migration Guide*, GA22-7347 and *HACMP: Installation Guide*, SC23-1940 for more information.

When you write a new resource monitor where the associated resource variable values are also to be supplied to Performance Toolbox Parallel Edition (PTPE), this book assumes that you have a thorough understanding of Performance Toolbox for AIX (PTX) programming concepts. These concepts are discussed in *Performance Toolbox for AIX Guide and Reference*, SC23-2625.

System administrators will need to use the information about configuring the Event Management subsystem.

The commands and interfaces described in this book require that you have the appropriate privileges and authorizations. For example, you may need to be running with an effective user ID of **root**. The specific security requirements for each subroutine and command are listed in the reference material.

---

## Typographic Conventions

This book uses the following typographic conventions:

Typographic	Usage
<b>Bold</b>	<ul style="list-style-type: none"> <li>• <b>Bold</b> words or characters represent system elements that you must use literally, such as commands, flags, and path names.</li> </ul>
<i>Italic</i>	<ul style="list-style-type: none"> <li>• <i>Italic</i> words or characters represent variable values that you must supply.</li> <li>• <i>Italics</i> are also used for book titles and for general emphasis in text.</li> </ul>
Constant width	Examples and information that the system displays appear in constant width typeface.
[ ]	Brackets enclose optional items in format and syntax descriptions.
{ }	Braces enclose a list from which you must choose an item in format and syntax descriptions.
	A vertical bar separates items in a list of choices. (In other words, it means “or.”)
< >	Angle brackets (less-than and greater-than) enclose the name of a key on the keyboard. For example, <b>&lt;Enter&gt;</b> refers to the key on your terminal or workstation that is labeled with the word Enter.
...	An ellipsis indicates that you can repeat the preceding item one or more times.
<b>&lt;Ctrl-x&gt;</b>	The notation <b>&lt;Ctrl-x&gt;</b> indicates a control character sequence. For example, <b>&lt;Ctrl-c&gt;</b> means that you hold down the control key while pressing <b>&lt;c&gt;</b> .
\	The continuation character is used in coding examples in this book for formatting purposes.



---

# Chapter 1. Understanding Event Management

This chapter introduces you to the concepts and operation of the Event Management subsystem. It also describes the use of the Event Management application programming interfaces (APIs) to write resource monitors and Event Management client programs. It briefly discusses some performance considerations to keep in mind as you write resource monitors and EM clients. Finally, it provides some information about resource monitors and shared memory.

---

## Introducing Event Management

The Event Management subsystem is a distributed subsystem of IBM RS/6000 Cluster Technology (RSCT) on the RS/6000 system. RSCT provides a set of high availability services to the IBM Parallel System Support Programs (PSSP) and to the Enhanced Scalability feature of the IBM High Availability Cluster Multi-Processing (high-availability cluster multiprocessing (HACMP)/ES) program. By matching information about the state of system resources with information about resource conditions that are of interest to client programs, it creates events. Client programs can use events to detect and recover from system failures, thus enhancing the availability of RS/6000 systems.

## The Event Management Components

Event Management operates through the cooperation of three software components:

- Resource monitors, which supply information about the current state of system resources.
- Event Management clients (EM clients), which supply information about which of those resource states are of interest.
- The Event Management subsystem, which puts the two together. It compares a condition of interest to a current state of a resource. If the state meets the test, the Event Management subsystem generates an **event** and notifies the EM client that the event has occurred.

Let us take a closer look at the roles of each of these components.

### Resource Monitors

The first Event Management component is the resource monitor. A **resource monitor** supplies information about the resources in the system. A **system resource** is an object in the system that provides services to system users. Examples of system resources are nodes, memory, disks, application programs, data bases, and file systems.

Resource monitors keep track of system resources and supply information about the characteristics or attributes of system resources to the Event Management subsystem. For example, a resource monitor might report the following:

- Whether a node is up or down
- Percentage of free memory for a particular node
- Percentage of space on a particular disk on a particular node that is free
- Whether a particular program or process is running

.

For each attribute of a resource, the resource monitor defines a **resource variable**.

Of course, there may be more than one copy of a resource in the system. For example, a system may have multiple nodes, a node may have multiple disks, a program may have multiple processes, and a database may have multiple partitions. Accordingly, the resource monitor reports the value of each copy of the resource variable in the system. Each copy is called an **instance** of the resource variable. For example, if a node has three disks, the resource monitor that tracks the percentage of free disk space reports the values of three instances of that resource variable for that node.

The resource monitor reports the value of each resource variable instance to the Event Management subsystem at defined times. The reporting occurs either periodically at configurable intervals or every time the attribute changes, depending on the type of resource variable.

A resource monitor can be a command, a daemon, or part of an application or subsystem that manages any type of system resource. It uses the Resource Monitor Application Programming Interface (RMAPI) to report its data to and receive control information from the Event Management subsystem.

### **Event Management Clients**

The second Event Management component is the Event Management client. An **Event Management client (EM client)** acts on information about the condition of system resources. Typical actions include recovery from system failures and reporting system changes.

Different conditions are of interest to different EM clients. For example, one EM client may want to know when the percentage-of-free-disk-space for any disk on any node in the system falls below 10%. Another EM client may want to know when a node in the system fails. The disk space client might allocate more space or send a note to a database administrator about the condition. The node health client might initiate a program that causes a backup node to take over the work of the failed node.

A condition in which an EM client is interested is called an **expression**. An EM client indicates that it wishes to be notified when the state of a resource instance matches a certain condition by registering an expression for particular resource variable instances.

An EM client is an application program or a subsystem. It uses the Event Management Application Programming Interface (EMAPI) to register interest in particular events, to receive information when the events occur, and to query resource variable values and definitions.

### **The Event Management Subsystem**

The last component, and the one that ties the other two together, is the Event Management subsystem. The **Event Management subsystem** communicates with the resource monitors and EM clients. It receives the values of resource-variable instances as they are reported by the resource monitors. It keeps track of the resource variables and expressions for which EM clients have registered.

At configurable intervals (or every time that values are reported, depending on the type of resource variable), the Event Management subsystem **observes** the values of resource variable instances. This means that it applies an expression to the value of the associated resource variable instances. When an expression evaluates to true (that is, when the value of a resource variable instance matches the condition expressed by the expression), the Event Management subsystem generates an event and notifies the appropriate EM client. The EM client then takes action appropriate to the event.

Figure 1 on page 3 shows the relationship between the three basic Event Management components.

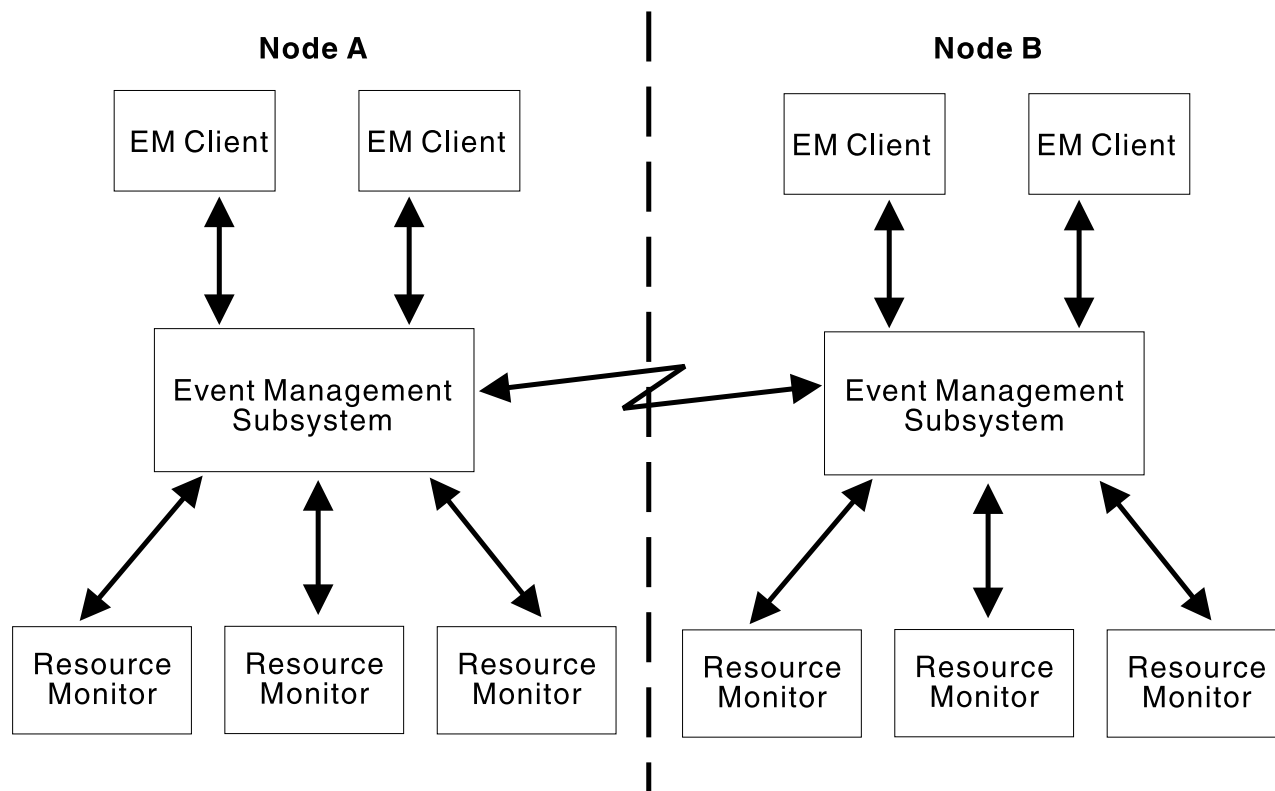


Figure 1. An Overview of the Event Management Components

## Communicating across Nodes

As Figure 1 shows, a resource monitor provides data to an instance of the Event Management subsystem that is running on the same node as the resource monitor. An EM client communicates with the instance of the Event Management subsystem that is running on the same node as the client. An instance of the Event Management subsystem on one node communicates with instances of the Event Management subsystem on other nodes to provide events to its local clients. Therefore, an EM client can register for and receive events about any resource in an RS/6000 SP or HACMP/ES system. In addition, an EM client that is executing outside of the SP system can remotely communicate with the Event Management subsystem that is located on the control workstation.

## Event Management and Domains

The Event Management subsystem operates in a **domain**. A domain is a set of RS/6000 machines upon which the Event Management subsystem executes and, exclusively of other machines, provides its services. On the RS/6000 SP a domain is a system partition. On an HACMP/ES cluster a domain is the entire cluster. Note that a machine may be in more than one domain. If an SP node is also in an HACMP/ES cluster, then the node is a member of both the SP domain and the HACMP/ES domain. In this case there are two instances of the Event Management subsystem executing on the node. The RS/6000 SP control workstation is considered to be a node in each SP system partition. An instance of the Event Management subsystem runs on the control workstation for each system partition.

Instances of the Event Management subsystem communicate only with other instances of the Event Management subsystem that are running in the same domain.

By default, an EM client establishes a session with the Event Management subsystem in the current domain. The current domain is the domain containing the node on which the client is running. Since a node may be in both an SP and an HACMP/ES cluster, a domain type must be specified. Optionally, the EM client can also specify the name of the domain with which the session is to be established. If the

domain is an SP system partition, the domain name is the system partition name. If the domain is an HACMP/ES cluster, the domain name is the cluster name.

If the domain type is HACMP/ES, the domain name must be the name of the HACMP/ES cluster containing the node on which the client is executing. If the domain type is SP, the domain name may be the name of any of the defined system partitions.

The EM client may be running on the SP control workstation or on a workstation outside of the SP. In this situation, if a domain name is not specified when a session is established, then the **SP\_NAME** environment variable is used to determine the domain name.

An EM client may establish multiple sessions. The client can then receive events from, and issue queries to, multiple domains. Multiple sessions can also be established with a single domain. Because each session provides a unique communication path, multiple sessions are useful if the client wishes to ignore certain responses from a domain while accepting others. For example, a client can register events in one session and issue queries in another.

## How Can You Use Event Management Services?

RSCT supports the following interfaces for using event management services:

- The RMAPI, an application programming interface that resource monitors use to communicate information about resource states to the Event Management subsystem
- The EMAPI, an application programming interface that EM clients use to get information about resource state changes from the Event Management subsystem
- On an RS/6000 SP, the SP Perspectives, a graphical user interface (GUI) that you can use to perform SP system management tasks, one of which is managing system events.

### Using the RMAPI

Programmers can use the RMAPI to write standalone C programs that monitor resources, or programmers can use the RMAPI to incorporate resource monitoring function into existing subsystems or applications. For example, a distributed database program can use the RMAPI to supply information about the state of its resources.

This book provides guidance and reference information for using the RMAPI. However, in this release of RSCT, you can only write resource monitors to be executed in the SP domain.

### Using the EMAPI

Programmers can use the EMAPI to write standalone C programs that take appropriate action when monitored resources change state. For example, an EM client that is notified when a resource becomes unavailable can initiate actions to recover the resource or circumvent the failure.

This book provides guidance and reference information to help you use the EMAPI.

### Using SP Perspectives to Monitor Events

System administrators and operators can use the Event Management portion of SP Perspectives to be notified when resources of interest change state. They then can run shell scripts that take recovery action. Detailed panels help you use the SP Perspectives graphical user interface to perform event management and other system management tasks.

For information on using SP Perspectives, see the *PSSP Administration Guide*.

---

## Writing Resource Monitors

To write a resource monitor to be used in an SP domain, you must perform the following steps:

- Define the resource data you want to monitor
- Choose the type of resource monitor you are going to write
- Configure the resource data and the resource monitor to the system
- Code and test the resource monitor.

This section provides you with an overview of each of these steps. For complete details on how to specify any of the data or subroutines, see the reference information in the following chapters:

- “Chapter 2. Event Management Configuration Data Reference” on page 27
- “Chapter 3. Event Management Subroutine Reference” on page 39
- “Chapter 4. Event Management Files Reference” on page 103.

For details on any commands you need to issue, see *PSSP Command and Technical Reference*.

For listings of several sample resource monitor programs, see “Chapter 5. Using the RMAPI: Some Resource Monitor Examples” on page 141.

## Defining the Resource Data

The first task in writing a resource monitor is to define the data that it is to collect and send to the Event Management subsystem. To do that, you must understand the following terms as they are used by the Event Management subsystem:

- Resources
- Resource variables
- Resource-variable names and descriptions
- Resource-variable value types and data types
- Resource-variable instances and resource IDs
- Dynamically instantiable resource variables
- The location of a resource-variable instance
- Resource-variable classes.

### Resources

A **resource** is an entity in the system that provides a set of services. Examples of resources include hardware entities such as processors, disk drives, memory, and adapters, and software entities such as database applications, processes, and file systems. Each resource in the system has one or more attributes that define the state of the resource. The number of attributes and the semantics of each attribute are defined by the resource.

The purpose of the Event Management subsystem is to inform interested programs whenever the state of a resource changes in some way. The resource state change is reflected by a change in a resource attribute. Each attribute is represented by a variable. The use of this variable removes the need to understand the semantics of resource attributes in order to perform the function.. This variable is called a resource variable.

### Resource Variables

A **resource variable** is the representation of an attribute of a resource in the Event Management subsystem. Associated with each resource variable is a name, a description, a value type, a data type, a resource ID, and, optionally, a location. Resource variables are grouped together into classes.

## Resource-Variable Names

A resource variable is identified by a **resource variable name**, which is a string that consists of a resource name followed by a period and the resource attribute.

By convention, the **resource name** consists of two or more components that include the name of the vendor and the name of the product that supplies the subsystem or application that manages the resource. The resource name may include additional components as needed to further specify the resource.

The **resource attribute** is a single component that follows the resource name.

The specification of components must follow certain rules, such as the type of character that is allowed in each component. For complete details on the specification of resource-variable names, see “Event Management Configuration Data (emcdb)” on page 28. For examples of resource-variable names, see Figure 3 on page 8.

## Resource-Variable Description

A **resource-variable description** provides information about the resource variable, in particular, its semantics. It should contain enough information so that writers of EM clients will be able to define meaningful expressions.

## Resource-Variable Value Types

A resource variable can have one of the following value types: Counter, Quantity, or State.

A **Counter** is a resource variable that represents a rate. It can have a data type of either long or float. The long and float formats are identical to the C language types of the same names. Typically, a Counter represents throughput. Examples include paging rates, I/O rates, and transaction rates.

The Event Management subsystem presents the value of an instance of a Counter as a rate. The rate represents the change in the actual contents of the instance of the Counter from one observation to the next divided by the time between the two observations.

For example, the resource variable called **IBM.PSSP.aixos.Mem.Virt.pagein** is a Counter that represents the rate at which pages are paged into virtual memory. Normally, its value is presented as a rate (number of pages per second). However, its raw value indicates the total number of pages that have been paged in since the Counter was initialized.

A **Quantity** is a resource variable whose value fluctuates over time. It can have a data type of either long or float. The long and float formats are identical to the C language types of the same names. Typically, a Quantity represents a level. Examples include the size of real memory or paging space, the number of processes in a queue, or the percentage of free disk space.

The Counter and Quantity value types and the statistics-object value types of the Performance Toolbox for AIX (PTX/6000) System Performance Measurement Interface (SPMI) have the same definitions. Counters and Quantities can be reported to the Performance Monitor subsystem at the same time as they are reported to the Event Management subsystem. The resource monitor can update instances of these two types of resource variable frequently to monitor performance. Meanwhile, the Event Management subsystem can observe these instances relatively infrequently without losing the ability to generate useful events.

A **State** is a variable whose value fluctuates over time. In addition, every fluctuation is significant and must be observed; otherwise, meaningful information would be lost. A State variable typically represents any attribute other than throughput or a level. A simple example of a State is an indicator of whether a node is up or down. For more complicated examples of State resource variables, see Figure 2 on page 7.

Values of a State variable can be reported only to the Event Management subsystem. They cannot be reported to the Performance Monitor subsystem. The Event Management subsystem observes instances

of a State variable when they are reported. While a resource monitor can update a State variable instance rapidly, on average, instances of State variables are expected to change at a relatively slow rate.

### Resource-Variable Data Types

Each of the value types (Counter, Quantity, and State) can have a data type of long or float. The long and float formats are identical to the C language types of the same names.

In addition, a State can have a data type of **structured byte string (SBS)**. An SBS is a string of bytes that consists of an SBS length field followed by one or more structured fields. The SBS length specifies the total length of the structured fields that follow it.

A structured field consists of a header followed by a value. The structured field header contains the length and data type of the structured field value and a serial number. A structured field may have a data type of long, float, null-terminated character string, or byte string.

The serial number is a unique value that identifies the structured field. Within an SBS, the serial numbers are numbered sequentially, starting with 0.

A structured field name is associated with each structured field in an SBS, but it is not included in the actual SBS. This name contains a short description of the structured byte field.

Figure 2 shows some examples of structured byte string definitions.

```
-----
SBS associated with resource variable IBM.PSSP.Prog.pcount

SBS Length = variable
SBS Fields
  Field Name      Field Length  Field Type      Field Serial Number
  CurPIDCount    4             long            0
  PrevPIDCount   4             long            1
  CurPIDList     variable      cstring         2
-----

SBS associated with resource variable IBM.PSSP.pm.Errlog

SBS Length = variable
SBS Fields
  Field Name      Field Length  Field Type      Field Serial Number
  sequenceNumber  variable     cstring         0
  errorID         variable     cstring         1
  errorClass      variable     cstring         2
  errorType       variable     cstring         3
  alertFlagsValue variable     cstring         4
  resourceName    variable     cstring         5
  resourceType    variable     cstring         6
  resourceClass   variable     cstring         7
  errorLabel      variable     cstring         8
```

Figure 2. Examples of Structured Byte String Definitions

### Resource-Variable Instances and Resource IDs

Most resources in the system have multiple copies. For example, there is more than one logical volume per node. There is more than one processor per node if the node is an SMP (Symmetric Multiprocessing) node. There are more than one pool of kernel memory buffers, more than one partition in a parallel database, and more than one node in the system partition. Each of these copies is an **instance** of the resource.

The resource variables that represent the states of these resources also have multiple copies. Each of these copies is an **instance** of the resource variable.

To uniquely identify each copy of a resource and all of its variables, each resource in the system has one, and only one, associated **resource ID**.

A resource ID is a list of elements, where each element is a name/value pair. A name/value pair consists of a resource ID element name followed by an equal sign followed by the value. In the EMAPI, the elements are separated by semicolons. In the RMAPI, the elements are separated by commas.

A **resource ID element name** is a string that describes the element.

For the EMAPI, a resource ID element value is a single value, a range of values, a comma-separated list of single values, or a comma-separated list of ranges. A range takes the form *a-b* and is valid only for resource ID elements of type integer.

For the RMAPI, a **resource ID element value** is a single value.

The resource ID of a resource specifies a particular instance of a resource and all of its resource variables within the system. To identify an instance, a resource ID can contain up to 4 elements, as defined by the **HA\_EM\_RSRC\_ID\_SIZE** constant.

For example, a variable that represents the number of used blocks in the **/tmp** file system requires a resource ID with three elements: the node number, the volume group name, and the logical volume name. A variable that represents the power state of a node requires a single element: the node number.

A resource ID element may be wildcarded in a manner that varies by subroutine.

The set of values in the resource ID uniquely identifies the copy of the resource in the system. By extension, they also uniquely identify the copy of the resource variable in the system. If there is only one copy of the resource in the system (for example, the control workstation), its resource ID is null.

The semantics of a resource ID are defined by the semantics of the resource variable with which it is associated. The names of the resource ID elements must be unique within a given resource variable's resource ID.

Figure 3 shows some examples of resource variable names and resource IDs.

Resource Variable Name	Resource ID
IBM.PSSP.aixos.CPU.gluser	NodeNum=5
IBM.PSSP.aixos.cpu.kern	NodeNum=5;Cpu=cpu0
IBM.PSSP.aixos.Mem.Real.size	NodeNum=5
IBM.PSSP.aixos.Mem.Virt.pagein	NodeNum=5
IBM.PSSP.aixos.Mem.Virt.pageout	NodeNum=5
IBM.PSSP.aixos.Mem.Kmem.inuse	NodeNum=5;Type=mbuf
IBM.PSSP.aixos.PagSp.%totalfree	NodeNum=5
IBM.PSSP.aixos.Disk.busy	NodeNum=5;Name=hdisk0
IBM.PSSP.aixos.Disk.busy	NodeNum=5;Name=hdisk1
IBM.PSSP.aixos.VG.free	NodeNum=5;VG=rootvg
IBM.PSSP.aixos.VG.free	NodeNum=5;VG=spdata
IBM.PSSP.aixos.FS.%totused	NodeNum=5;VG=rootvg;LV=hd4
IBM.PSSP.aixos.FS.%totused	NodeNum=5;VG=spdata;LV=lv00
IBM.PSSP.SP_HW.Node.powerLED	NodeNum=5
IBM.PSSP.SP_HW.Frame.frACLEd	FrameNum=1
IBM.PSSP.SP_HW.Switch.powerLED	SwitchNum=1
IBM.PSSP.Membership.Node.state	NodeNum=5
IBM.DB2.Part.size	DBname=foo;Part=3

Figure 3. Examples of Resource Variable Names and Resource IDs

In Figure 3, although most resource IDs contain a node number, some resource variables are not associated with a particular node.

Although components of a resource name can be used in more than one name, resource names must uniquely identify a resource.

## Dynamically Instantiable Resource Variables

In most cases, the instances of a resource that a resource monitor is responsible for tracking and reporting on are static or relatively so. Thus, a resource monitor knows from the time it starts running, after any necessary initialization occurs, all of the instances of the resources it monitors. For example, after any necessary initialization, a hardware monitor, a node health monitor, and a disk monitor know what frames, nodes, and disks, respectively, are in the system.

However, in some cases, a resource monitor knows how to track and report on a type of resource, but the resource monitor does not know all of the possible instances of the resource that may occur. For example, the program resource monitor knows how to track and report on a program that is running in a domain, but it does not know in advance the names of all possible programs that can run. For a resource like this, you can define the resource variable as **dynamically instantiable**.

If a resource variable is dynamically instantiable, the resource monitor does not register all of the instances of its resource variables. Instead, it waits for the Event Management subsystem to tell it what instances to register. As the Event Management subsystem receives requests from EM clients for events based on instances of dynamically instantiable variables, the Event Management subsystem passes the instances to the resource monitor in a control message with an “instantiate variables” command.

## Location of a Resource-Variable Instance

The **location** of a resource variable instance is the node on which the *resource monitor* that supplies the instance resides. Note that the resource monitor may or may not reside on the same node as the instance of the resource.

For example, in Figure 3 on page 8, the resource variables that begin with **IBM.PSSP.SP\_HW** represent hardware resources that are monitored by the hardware monitor subsystem. A resource monitor that connects to the **hardmon** daemon can reside on any node of the SP. It can even reside on an RS/6000 outside of the SP, such as the control workstation. Thus, a resource monitor on a single node supplies all instances of **IBM.PSSP.SP\_HW** resource variables. Accordingly, the location fields in all of the instances specify that node.

As another example, in Figure 3 on page 8, the resource variable named **IBM.DB2.Part.size** represents the size of a partition in a DB2 parallel database. In this case, the resource-variable instance originates on the node that contains the database partition. The database partition itself may move from one node to another if recovery of the database is required.

If a resource ID implies the location of a resource variable, the description of the resource variable should so specify.

## Resource-Variable Classes

The **class** of a resource variable indicates the subsystem or application that manages the associated resource. By convention, class names follow the same naming conventions that are used for resource names: the first component is a vendor name, the second component is a product name, and subsequent components, if any, further specify the resource class. Class names must be unique among all of the class names defined for a domain.

Grouping resource variables by class can also be used to set the observation interval for resource variables of value type Counter and Quantity. All instances of all variables of these value types in the same class are observed at the same intervals.

For examples of class names, see Figure 4 on page 10.

Resource Variable Class	Resource Variables in the Class
IBM.PSSP.aixos.CPU	IBM.PSSP.aixos.CPU.gluser IBM.PSSP.aixos.CPU.glkern IBM.PSSP.aixos.CPU.glwait IBM.PSSP.aixos.CPU.glidle IBM.PSSP.aixos.cpu.user IBM.PSSP.aixos.cpu.kern IBM.PSSP.aixos.cpu.wait IBM.PSSP.aixos.cpu.idle
IBM.PSSP.aixos.Mem	IBM.PSSP.aixos.Mem.Real.size IBM.PSSP.aixos.Mem.Real.numfrb IBM.PSSP.aixos.Mem.Real.%free IBM.PSSP.aixos.Mem.Real.%pinned IBM.PSSP.aixos.Mem.Virt.pagein IBM.PSSP.aixos.Mem.Virt.pageout IBM.PSSP.aixos.Mem.Virt.pgspgin IBM.PSSP.aixos.Mem.Virt.pgspgout IBM.PSSP.aixos.Mem.Virt.pagexct IBM.PSSP.aixos.Mem.Kmem.inuse IBM.PSSP.aixos.Mem.Kmem.calls IBM.PSSP.aixos.Mem.Kmem.failures IBM.PSSP.aixos.Mem.Kmem.memuse
IBM.PSSP.aixos.Disk	IBM.PSSP.aixos.Disk.busy IBM.PSSP.aixos.Disk.xfer IBM.PSSP.aixos.Disk.rblk IBM.PSSP.aixos.Disk.wblk

Figure 4. Examples of Resource Variable Classes

## Name Spaces, System Partitions, and Domains

With respect to RS/6000 SP system partitioning, the name space for node numbers, whether contained in a resource ID element or in a location field, are global across all system partitions of a single SP system—that is, across all SP type domains in a single SP. The name space for node numbers in an HACMP/ES domain are local to the domain. The name spaces for other element values of resource IDs are local to a domain, of any type. The name space for resource variable names and classes are also local to a domain, of any type.

## Choosing the Resource Monitor Type

After defining the data, the next step is to choose the type of resource monitor you are going to write.

The way in which a resource monitor operates depends upon the way in which it is implemented. A resource monitor can be implemented as one of the following:

- A daemon
- A command
- Routines you add to an existing program.

### Daemon-Based Resource Monitors

A resource monitor that is implemented as a daemon has the following characteristics:

- It is long-lived.
- It can be started automatically by the Event Management subsystem whenever an EM client queries or registers for events that are generated from resource variables it supplies.
- It can supply values for Counter, Quantity, and State resource variable instances.

After the resource monitor daemon initializes itself and starts a server session, the Event Management subsystem connects to it. The Event Management subsystem then begins to request values for the Counter and Quantity resource variables that the resource monitor has defined. A daemon-based resource monitor can also supply values for State resource variables to the Event Management subsystem.

- All values of Counter and Quantity variable instances can be supplied to both the Event Management subsystem and the Performance Monitor subsystem.

To configure a daemon-based resource monitor, specify a connection type of server and, optionally, the fully-qualified path and file name of the daemon. The server connection type indicates that the Event Management subsystem and the Performance Monitor subsystem initiate the connection and that the program is to act as a daemon. If specified, the path and file name tell the Event Management subsystem which daemon to start. If a path and file name are not specified, it is assumed that the daemon is started by some other means, independent of the Event Management subsystem.

### **Command-Based Resource Monitors**

A resource monitor that is implemented as a command is more limited but also less complex. It can be used to supply resource variable values from a shell script. A command-based resource monitor has the following characteristics:

- It is relatively short-lived.
- It is not started automatically by the Event Management subsystem; instead, it initiates the connection to the Event Management subsystem. It supplies values whether or not any Event Management client has registered interest in events that are generated from resource variables it is monitoring.
- It can supply values only for State resource variables and only to the Event Management subsystem. It sends the values as the states of the monitored resources change. A command-based resource monitor cannot supply values for Counter or Quantity variables and cannot respond to requests for variables.
- It cannot supply resource variable values to the Performance Monitor subsystem.

If your resource monitor is command-based, specify a connection type of client when you configure it. The client connection type indicates that the resource monitor initiates the connection and that the program is to act as a command.

### **Resource Monitors that Are Incorporated in Other Programs**

A resource monitor that is incorporated in another program can act as either a daemon-based resource monitor or a command-based resource monitor.

If it is to act as a daemon-based resource monitor, specify a connection type of server and a null path name when you configure it. The server connection type indicates that the Event Management subsystem and the Performance Monitor subsystem initiate the connection and that the resource monitoring function will behave as if it were a daemon. However, the null path name indicates that there is no daemon to be started.

If the resource monitoring function incorporated in another program is to act as a command-based resource monitor, specify a connection type of client. The client connection type indicates that the resource monitor initiates the connection and that the program is to act as a command.

### **Resource Monitor Type Considerations**

In making your choice, consider the following questions:

- Do you already have an existing program that manages a resource?  
If so, you may want to incorporate resource monitoring function into it rather than write a stand-alone program.
- What kind of data do you need to handle?  
If your resource monitor deals only with State resource variables, your resource monitor can act as either a daemon or a command. However, if you have Counter or Quantity resource variables, your resource monitor must act as a daemon, although you can incorporate the function in another program.
- Do you want to supply data to the Performance Monitor subsystem as well as the Event Management subsystem?  
If you do, you must write a resource monitor (or function) that acts as a daemon. Keep in mind, however, that only the values of Counter and Quantity resource variables are supplied to the Performance Monitor subsystem.
- How complicated does your resource monitor need to be?  
A command-based resource monitor is simpler to code than a daemon-based resource monitor.

When making your choice, also keep in mind this recovery characteristic of the Event Management subsystem: If any of its components fail, after the failed components have been restarted it is assumed that any resource variable values that have been lost can be refreshed from the resource monitor. If the resource monitor is command-based, it may take some time before it once again provides the resource variable values. A resource monitor that is a daemon, or behaves as a daemon, can provide resource variable values immediately.

## Configuring the Resource Data and Resource Monitor

Once you have determined the data your resource monitor is going to supply and the type of resource monitor you are going to write, you must make the information available to the Event Management subsystem as a set of formal definitions. This is called **Event Management configuration**.

As part of Event Management configuration, you must supply information about:

- The resource variables that your program will monitor
- Any structured byte string definitions that are required
- The resource IDs for each resource variable
- The classes in which the resource variables are grouped
- Your resource monitor.

Event Management configuration updates the SDR and “production” files for each system partition. Accordingly, it is advisable to create a separate system partition for testing purposes before you alter your production environment.

Event Management configuration requires several steps. First, from the control workstation, you must store the resource definitions for a given system partition as a set of objects in the System Data Repository (SDR). There are two ways to do this.

The preferred way to load your data for a new resource monitor into the SDR is to create a loadlist file and run the **haemloadcfg** command. For details on how to do this, see the Event Management subsystem chapter in *PSSP Administration Guide*.

Another way of loading your data into the SDR is to use such commands as the **SDRCreatObjects** command. You can do this most easily by creating a shell script with the appropriate commands. For an example, see “The `rmapi_smp.loadsd` Shell Script” on page 174.

You place the source information in the following partitioned classes in the SDR:

- Class = EM\_Resource\_Variable
- Class = EM\_Structured\_Byte\_String
- Class = EM\_Resource\_ID
- Class = EM\_Resource\_Class
- Class = EM\_Resource\_Monitor

Once you have stored the information in the SDR, you must issue the **haemcfg** command. This command compiles the Event Management data in the SDR for a system partition. It also creates a staging version of the system partition’s Event Management Configuration Data Base (EMCDB) file. The data in the EMCDB file, not the SDR, is the data that is used by the Event Management subsystem.

In order to make the new EMCDB file active, you must shut down and restart the Event Management subsystem in the system partition. To do so, shut down all of the Event Management daemons for the system partition (on the control workstation and all of the nodes) and then restart them all. For details, see *PSSP Administration Guide*.

Finally, you must repeat this process for each of the system partitions in which Event Management services are required.

Each time you update the EMCDB for a system partition, the **haemcfg** command modifies a version number. The version number allows each instance of the Event Management subsystem and each resource monitor in a system partition to verify that they have access to the correct version of configuration information. The **haemcfg** command maintains the version number in both the EMCDB and the SDR.

Each database may be updated at any time without interfering with the operation of any Event Management subsystem, client, or resource monitor. However, once a database has been updated in a system partition, the Event Management subsystem in the system partition must be restarted to refresh its configuration information from the newly updated database.

## Coding and Testing the Resource Monitor

Having defined and configured the resource monitor and its data, the next task is to code the resource monitor.

In general, all resource monitors perform certain functions. They initialize themselves with the RMAPI, manage RMAPI sessions, update resource variables, and terminate. However, the structure and functions of the resource monitor (or resource monitoring function) depend on the implementation you have chosen: daemon-based or command-based. Because the coding of a command-based resource monitor is simpler, we'll describe that type first.

Also, "Chapter 5. Using the RMAPI: Some Resource Monitor Examples" on page 141 contains several listings of sample resource monitor programs.

To test your new resource monitor with the test version of the EMCDB, run it in the separate system partition that you created for that purpose.

### Coding a Command-Based Resource Monitor

A command-based resource monitor includes the following functions:

- Initializing  
The resource monitor calls the **ha\_rr\_init** subroutine to inform the RMAPI of the identity of the resource monitor.
- Starting a session  
After the **ha\_rr\_init** subroutine completes successfully, the command-based resource monitor calls the **ha\_rr\_start\_session** subroutine to start a session with the Event Management subsystem.
- Registering resource variables and their instances  
After starting a session, the resource monitor calls the **ha\_rr\_reg\_var** subroutine to register one or more resource variable instances with the RMAPI.
- Adding resource variables  
After registering resource variables and their instances, the resource monitor must call the **ha\_rr\_add\_var** subroutine to add the resource variables to the list of those known to the RMAPI and supply their current values.
- Sending updated values of resource variable instances  
Whenever it is time to send updated values of resource variable instances to the RMAPI, the resource monitor calls the **ha\_rr\_send\_val** subroutine to send them to the Event Management subsystem.
- Deleting resource variables  
The resource monitor must call the **ha\_rr\_del\_var** subroutine to delete any resource variable that it is no longer sending from the list of those known to the RMAPI.
- Ending a session

When the command-based resource monitor is ready to end, it calls the **ha\_rr\_end\_session** subroutine to end the session.

- Terminating

Once the session has ended, the command-based resource monitor calls the **ha\_rr\_terminate** subroutine and then exits.

Figure 5 shows a pseudocode outline of a command-based resource monitor that uses RMAPI subroutines.

```
Call ha_rr_init

Call ha_rr_start_session

Call ha_rr_reg_var with all known variable instances.

Call ha_rr_add_var with the variables that were specified to
the ha_rr_reg_var routine and supply current
values.

Call ha_rr_send_val with the State variable values. Do this
as frequently as required by the operation
of the resource monitor.

Call ha_rr_del_var to delete variables added.

Call ha_rr_end_session

Call ha_rr_terminate

Exit
```

*Figure 5. A Pseudocode Outline of a Command-Based Resource Monitor*

## **Coding a Daemon-Based Resource Monitor**

A daemon-based resource monitor includes the following functions:

- Initializing

The resource monitor calls the **ha\_rr\_init** subroutine to inform the RMAPI of the identity of the resource monitor.

- Making a server session

After the **ha\_rr\_init** subroutine completes successfully, the daemon-based resource monitor calls the **ha\_rr\_makserv** subroutine to start a server session so that one or both of the resource monitor managers can connect to it. The resource monitor managers are the Event Management subsystem and the Performance Monitor subsystem.

- Registering resource variables and their instances

After starting a server session, the resource monitor calls the **ha\_rr\_reg\_var** subroutine to register one or more resource variable instances with the RMAPI.

- Starting a session

Whenever the file descriptor returned by the **ha\_rr\_makserv** subroutine is ready for reading, the daemon-based resource monitor calls the **ha\_rr\_start\_session** subroutine to start a session with the resource monitor manager.

- Getting a control message from the Event Management subsystem

When the resource monitor determines that there is data to be read for the session, it calls the **ha\_rr\_get\_ctrlmsg** subroutine to read the control message. The control message contains a command from the resource monitor manager.

- Processing the control message  
Once the resource monitor has received the control message, it processes the command it contains. For information on the commands a control message can contain and the processing that the resource monitor performs, see “Processing Control Messages”.
- Ending a session  
When a resource monitor manager has disconnected, the daemon-based resource monitor calls the **ha\_rr\_end\_session** subroutine to end the session.
- Terminating.  
When the daemon-based resource monitor no longer has a session with either resource monitor manager, it should call the **ha\_rr\_terminate** subroutine and then exit, if the resource monitor was configured to be startable by a resource monitor manager.

**Processing Control Messages:** A control message contains one of the following commands:

#### **HA\_RR\_CMD\_INSTV**

The “instantiate variables” command asks the resource monitor to create instances of specified resource variables. This command is sent only for variables that have been configured as dynamically instantiable.

To process the command, the resource monitor creates actual instances for each specified instance of a resource variable. Once the instances have been created, the resource monitor calls the **ha\_rr\_reg\_var** subroutine to register the instances.

#### **HA\_RR\_CMD\_ADDV**

The “add variables” command tells the resource monitor to add the specified resource variable instances to the RMAPI.

To process the command, the resource monitor calls the **ha\_rr\_add\_var** subroutine to add the variables and supply their current values. Then, it calls the **ha\_rr\_send\_val** subroutine to start sending updated values for the resource variable instances that have been added.

To find out how often to send values for each resource-variable class of type Counter or Quantity, it calls the **ha\_rr\_get\_interval** subroutine to get the configured reporting interval.

#### **HA\_RR\_CMD\_ADDALL**

The “add all variables” command tells the resource monitor to add all of its Counter and Quantity resource variable instances to the RMAPI.

To process the command, the resource monitor calls the **ha\_rr\_add\_var** subroutine to add the variables and supply their current values. Then, it calls the **ha\_rr\_send\_val** subroutine to start sending updated values for the resource variable instances that have been added.

To find out how often to send values for each resource-variable class of type Counter or Quantity, it calls the **ha\_rr\_get\_interval** subroutine to get the configured reporting interval.

#### **HA\_RR\_CMD\_DELV**

The “delete variables” command tells the resource monitor to delete the specified resource variable instances from the RMAPI.

To process the command, the resource monitor calls the **ha\_rr\_del\_var** subroutine to delete the resource variables. Then, it stops sending values for the resource-variable instances that have been deleted.

#### **HA\_RR\_CMD\_DELALL**

The “delete all variables” command tells the resource monitor to delete all of its Counter and Quantity resource variable instances from the RMAPI.

To process the command, the resource monitor calls the **ha\_rr\_del\_var** subroutine to delete the variables. Then, it stops sending values for the resource variable instances that have been deleted.

Figure 6 shows a pseudocode outline of a daemon-based resource monitor that uses RMAPI subroutines.

```
Call ha_rr_init

Call ha_rr_makserv
  Add the file descriptor (fdc) to the select mask.

Call ha_rr_reg_var with all known variable instances.

Loop on select
  If fdc is ready
    Call ha_rr_start_session
    Add the file descriptor (fdN) to the select mask.

  If fdN is ready
    Call ha_rr_get_ctrlmsg with fdN

    Process the control message

    If the session is disconnected
      Call ha_rr_end_session

end Loop

Call ha_rr_terminate.
```

Figure 6. A Pseudocode Outline of a Daemon-Based Resource Monitor

### Alternative Testing Method

Although it is recommended that a new resource monitor be tested in a non-production system partition, it is possible to test a new resource monitor on one or more nodes in an existing system partition. However, the Event Management subsystem on the test node(or nodes) would no longer be in the domain for the duration of the test. Also, the definitions for the new resource monitor must be added to the SDR in the system partition that contains the test nodes.

During this procedure, do not execute the **haemcfg** command on the control workstation. Specifically, while the test resource monitor definitions are in the SDR, do not create a new Event Management Configuration Database (EMCDB) in the staging area on the control workstation.

Perform the following steps:

1. On the node (or nodes) where the new resource monitor is to be tested, execute the command:  
**/usr/sbin/rsct/bin/haemctrl -k**
2. On one of the test nodes, create a loadlist file containing the definitions for the new resource monitor. See the **haemloadlist** man page in the *PSSP: Command and Technical Reference*.
3. Execute the following command, specifying the name of the loadlist file created in the prior step:  
**/usr/sbin/rsct/install/bin/haemloadcfg filename**  
See the **haemloadcfg** man page in the *PSSP: Command and Technical Reference*.
4. In the **/etc/ha/cfg** directory, execute the command:  
**/usr/sbin/rsct/bin/haemcfg -n**  
Ensure that the **-n** flag is specified. Without this flag, the version number of the production EMCDB in this system partition will be updated in the SDR. See the **haemcfg** man page in the *PSSP: Command and Technical Reference*.  
If you are testing on more than one node, copy the EMCDB to the **/etc/ha/cfg** directory on the other nodes. The name of the EMCDB is **em.domain\_name.cdb**, where *domain\_name* is the name of the domain (system partition name) that contains the node.

5. On each test node, execute the command:  
**chssys -s haem -a "-T test\_rm"**  
On each test node, execute the **/usr/sbin/rsct/bin/haemctrl -s** command. One or more Event Management subsystems on one or more test nodes can now be used to test the new resource monitor.
6. When testing is complete, execute the following command on each test node:  
**/usr/sbin/rsct/bin/haemctrl -k**
7. On each test node, execute the command:  
**chssys -s haem -a ""**
8. Using the loadlist file created earlier, execute the command:  
**/usr/sbin/rsct/install/bin/haemloadcfg -d filename**  
This will remove the definitions for the new resource monitor.
9. On each test node, execute the command:  
**/usr/sbin/rsct/bin/haemctrl -s**  
The Event Management subsystem on the test nodes will now be back in the domain.

---

## Writing Event Management Clients

To write an Event Management client (EM client), you must perform the following steps:

- Find out what resource data is available
- Define expressions to create events
- Code and test the EM client.

This section provides you with an overview of each of these steps. For complete details on how to specify any of the data or subroutines, see the reference information in the following chapters:

- “Chapter 2. Event Management Configuration Data Reference” on page 27
- “Chapter 3. Event Management Subroutine Reference” on page 39
- “Chapter 4. Event Management Files Reference” on page 103.

For details on any commands you need to issue, see *PSSP Command and Technical Reference*.

For listings of several sample Event Management clients, see “Chapter 6. Using the EMAPI: Some Event Management Client Examples” on page 181.

## Finding Out What Resource Data Is Available

As described in “Defining the Resource Data” on page 5, resource data is collected and sent to the Event Management subsystem by resource monitors.

The first task in writing an EM client is to determine what resource data is being collected in the system.

There are two ways to do this:

- You can use commands.  
You can use the RSCT **haemqvar** command to list information about resource variables and resource IDs that have been defined in the system.
- On an RS/6000 SP you can use the SP Perspectives graphical user interface.  
You can use the Event Management portion of SP Perspectives to list information about the resource variables and resource IDs that have been defined in the system.

## Getting Resource-Variable Information Using the haemqvar Command

By default, the **haemqvar** command produces a listing of all defined resource variables in the current SP domain. For each resource variable the following information is included:

- Variable name
- Value type
- Data type
- SBS format (if data type is Structured Byte String)
- Initial value
- Class
- Locator
- Variable description
- Resource ID and its description
- Default expression (if defined) and its description

Since the amount of information produced can be quite large, the output of this command should be redirected to a file:

**haemqvar > var.list**

This command can also take arguments requesting information about a particular resource variable. The following example produces information about the **IBM.PSSP.aixos.FS.%totused** resource variable:

```
haemqvar "" IBM.PSSP.aixos.FS.%totused "*" 
```

Refer to the *PSSP: Command and Technical Reference* for more information.

## Getting Resource-Variable Information from SP Perspectives

You can use the Event Management portion of SP Perspectives to get information about the resource variables and resource IDs that have been defined in the system. To do this:

1. Start the Event Management portion of SP Perspectives. To do so, enter:

**spevent**

In response, the system displays the Event Perspective window.

2. If you have more than one system partition defined:
  - a. Click on the icon that represents the system partition in which you are interested.
  - b. From the Menu bar, click on **Actions**, then click on **Set Current System Partition**.
3. If the Conditions pane is not already loaded:
  - a. From the Menu bar, click on **View**, then click on **Add Pane...**
  - b. Select **Conditions** in the "Pane type" field, then click on **OK**.

In response, the system displays the Conditions pane on the Event Perspective window.

4. Click on the Conditions pane to make it active.
5. Open the Create Condition notebook. To do so:

From the Menu bar, click on **Actions**, then select **Create....**

**Or**

From the tool bar, click on the **Create a Condition** (pencil) icon.

In response, the system displays the Create Condition notebook. On the Create Condition notebook, the Resource variable names field displays a list of all of the resource variables that are defined in the system partition you selected.

6. To get information about any particular resource variable, select that variable and click on **Show Details...**

In response, the system displays the Show Resource Variable Details dialog box. This dialog box contains the following fields, each of which provides details about the selected resource variable:

- Resource-variable name
- Resource-variable description
- Resource-variable value type
- Resource-variable data type
- SBS (Structured Byte String) format

**Note:** This field is displayed only if the resource variable data type is **sbs**.

- Resource ID format
- Resource ID description

Figure 7 on page 20 provides an example of the Show Resource Variable Details dialog box—in this case, showing details about the **IBM.PSSP.pm.Errlog** resource variable.

For complete information on using SP Perspectives, see *PSSP: Administration*, the *IBM PSSP for AIX: SP Perspectives Comprehensive Guide* redbook, and the SP Perspectives help panels.

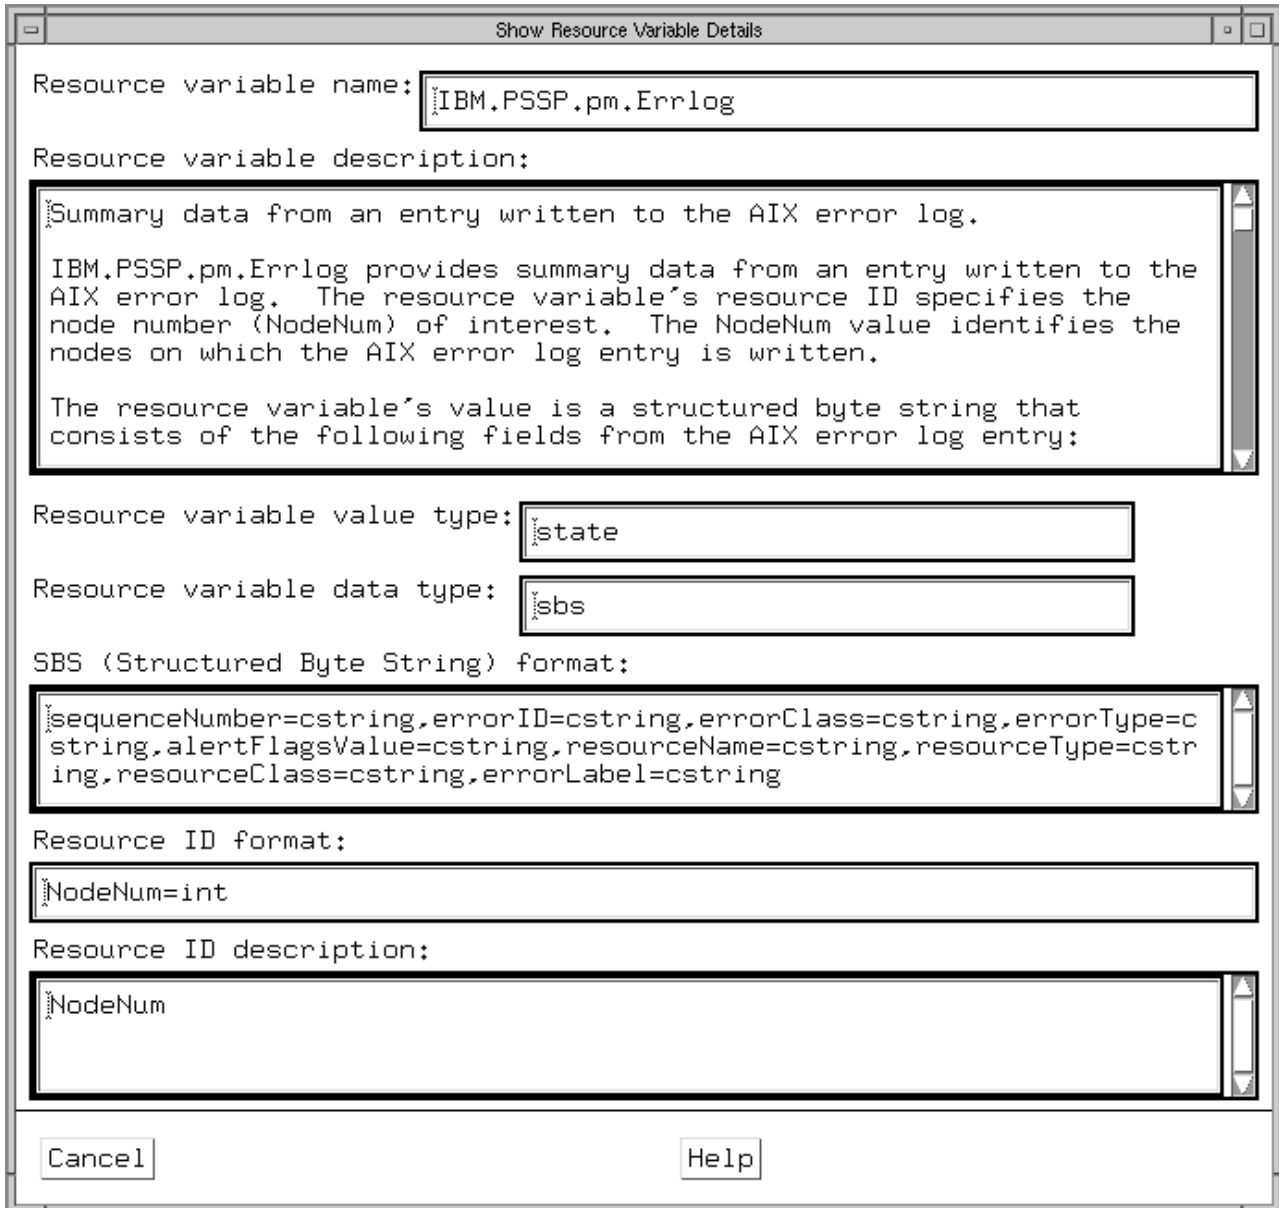


Figure 7. Show Resource Variable Details Dialog Box

## Using Expressions to Define Events

Once you know what resource data is being collected in the system, you can start defining expressions that will cause your program to be notified when conditions of interest occur. This notification is called an event. Before we get into defining expressions, though, let's take a look at how the Event Management subsystem generates an event.

### How the Event Management Subsystem Generates Events

As discussed in "Defining the Resource Data" on page 5, a resource is represented by one or more variables that are defined and updated by a resource monitor. The Event Management subsystem observes each resource variable instance at a configurable interval (for Counter and Quantity resource variables) or as the resource variable instance changes (for State resource variables). At each observation of a resource variable instance, the Event Management subsystem applies one or more expressions that were previously supplied by EM clients.

An expression is a relational condition between a resource variable and other elements, such as a constant or the value of a variable from the previous observation. The relationship can be an arithmetic one, a logical one, or a combination.

The expression is applied to each instance of the resource variable as it is observed. If the expression is true, an event is generated. More than one expression may be applied to an instance of the same resource variable at the same observation.

Here's an example of an event that could be generated. Suppose a resource monitor is keeping track of whether a particular node is up or down. The resource monitor represents the node by the **IBM.PSSP.Membership.Node.state** integer variable. As the state of the node changes, it sets the variable to a value of 1 (when the node is up) or a value of 0 (when the node is down).

An EM client is interested in knowing when the node goes down. Accordingly, it asks the Event Management subsystem to observe instances of the variable and apply the following expression to it:

```
X == 0
```

where **X** represents the **IBM.PSSP.Membership.Node.state** variable.

At each observation, the Event Management subsystem applies the expression to the resource variable instance. If the Event Management subsystem observes that the value of an instance of the **IBM.PSSP.Membership.Node.state** variable is 0, then the expression is true. Consequently, the Event Management subsystem generates an event, thereby notifying the client.

Another example of a simple expression is **X < 10**, where **X** is a resource variable that represents the percentage of free space in a file system. This expression would generate an event whenever the file system's free space was observed to be less than 10%.

To summarize, then, an **expression** is a relational condition between a resource variable and other elements, such as a constant or the value of a variable from the previous observation. An **event** is the notification that an expression that contains a resource variable is true.

## Defining Expressions

To define an expression, you must specify the resource variable and its relationship to at least one other element.

To specify the resource variable, you use the letter "X". Unmodified, "X" represents the value of the latest observation of an instance of the resource variable.

You can also modify the resource variable name to indicate any, or a combination, of the following:

- The value of the previous observation of the instance (**P**)
- The raw value of the Counter rather than the default rate (**R**)
- The value of a specific structured byte field (the serial number of the field).

The relationship you specify can be arithmetic, logical, or a combination.

For the complete definition and syntax of an expression, see "Expressions (haemexpr)" on page 123. Here are some examples of valid expression definitions:

**X == 0** The value of the resource variable instance is equal to zero.

**X < 20 || X > 80**

The value of the resource variable instance is less than 20 or greater than 80.

**!(X < 20 || X > 80)**

The value of the resource variable instance is neither less than 20 nor greater than 80.

### **X@R > X@PR**

The current raw value of the variable instance is greater than the raw value of the variable instance from its previous observation.

### **X >= X@P + 5**

The current value of the variable instance is greater than or equal to the value of the variable instance from its previous observation plus 5.

### **X@2 != 100**

For a resource variable instance that is defined as an SBS, the value of structured field number 2 is not equal to 100.

### **X@0 != X@P0 || X@1 != X@P1**

For a resource variable instance that is defined as an SBS, the value of either structured field number 0 or structured field number 1 has changed since its previous observation.

## **Defining Rearm Expressions**

When an EM client registers for an event, it can specify either one or two expressions.

By default, it specifies a single expression, the **event expression**. When the event expression is true, the EM client receives an event.

However, an EM client can also specify a second expression called the **rearm expression** in addition to the original event expression. When an EM client specifies a rearm expression, the two expressions are used alternately, as follows: the event expression is used until it is true, then the rearm expression is used until it is true, and so on.

Rearm expressions can be used in a couple of different ways. One common way is to define the rearm expression as the inverse of the event expression. For example, if the event expression tests whether a resource variable value is “on”, the rearm expression tests whether it is “off”.

A rearm expression can also be used with the event expression to define an upper and lower boundary for a condition of interest. For example, suppose an EM client is interested in how much disk space is used in the system. It defines a normal disk space condition to be in the range between 80% and 90% of the total space. However, if the amount of disk space used rises above 90% or falls below 80%, the EM client wants to be notified.

In this case, the event expression would be: “the amount of disk space used is greater than 90%,” and the rearm expression would be: “the amount of disk space used is less than 80%”. The EM client would also register for events from both expressions.

The first time disk space usage rises above 90%, the Event Management subsystem notifies the EM client and switches to the rearm expression. The EM client is not notified again until disk space usage decreases to an amount less than 80%, and then the Event Management subsystem switches to the event expression.

## **Using a Default Expression**

When the writer of a resource monitor defines a resource variable, he or she may also optionally define for it a single **default expression**. As the writer of an EM client, you can use the default expression (if it is defined) in addition to, or instead of, expressions you define.

## **Coding and Testing the EM Client**

Having determined what resource data is available and the conditions your program will be interested in, the next task is to code the EM client.

In general, the EM client starts one or more EMAPI sessions, registers for events with the Event Management subsystem, processes events from the Event Management subsystem, and ends its EMAPI sessions. It can also query resource variable information and unregister for events. The following sections describe each of these activities in more detail.

Also, “Chapter 6. Using the EMAPI: Some Event Management Client Examples” on page 181 contains several listings of sample EM client programs.

From an Event Management point of view, you do not need to set up a separate domain just for testing an EM client. However, if your EM client is taking actions that might adversely affect a production environment, you may wish to set up a separate test domain or other testing environment.

## Starting a Session

To start an event management session, the EM client calls the **ha\_em\_start\_session** subroutine.

By default, an EM client establishes a session with the Event Management subsystem in the current domain. The current domain is the domain containing the node on which the client is running. Since a node may be in both an SP and an HACMP/ES cluster, a domain type must be specified. Optionally, the EM client can also specify the name of the domain with which the session is to be established. If the domain is an SP system partition, the domain name is the system partition name. If the domain is an HACMP/ES cluster, the domain name is the cluster name.

If the domain type is HACMP/ES, the domain name must be the name of the HACMP/ES cluster containing the node on which the client is executing. If the domain type is SP, the domain name may be the name of any of the defined system partitions.

If the EM client is running either on the SP control workstation or on a workstation outside of the SP and a domain name is not specified when a session is established, then the **SP\_NAME** environment variable is used for the domain name. In this situation the EM client may only establish a session with an SP domain type.

An EM client may establish multiple sessions. The client can then receive events from, and issue queries to, multiple domains. Multiple sessions can also be established with a single domain. Because each session provides a unique communication path, a client can use multiple sessions to ignore certain responses from a domain while accepting others. For example, a client can register events in one session and issue queries in another.

If the communication path for a session is lost, the EM client calls the **ha\_em\_restart\_session** subroutine to restart an existing session.

## Registering for Events

Once it has started a session, the EM client calls the **ha\_em\_send\_command** subroutine to register for events.

The EM client may specify either of two register commands:

- **HA\_EM\_CMD\_REG**, which the EM client uses to register for an event from a single expression, the event expression. If a rearm expression is defined, it does not generate an event.
- **HA\_EM\_CMD\_REG2**, which the EM client uses to register for an event from both an event expression and a rearm expression.

For each event, the EM client supplies the name of a resource variable, a resource ID, and the expression or expressions that are to be used to generate the event. For each successfully registered event, the EMAPI returns an event ID to the EM client.

The resource ID identifies the particular instance or instances of the resource variable that the Event Management subsystem should use to generate the event. To receive events from more than one instance

of the resource variable, the EM client can use wildcard characters in the resource ID to specify multiple instances. The EM client can supply its own expression or expressions or use the default expression (if any) that was defined when the resource variable was configured.

If the EM client registers for more than one event with a single command, the events belong to an **event command group**. The EM client can call the **ha\_em\_get\_ecgid** subroutine to get the identifier of the event command group to which an event belongs.

By default, the Event Management subsystem does not generate an event until the expression is true. If, however, the EM client specifies the **HA\_EM\_SCMD\_REVAL** subcommand on either register command, then the Event Management subsystem generates an event when an instance of the resource variable is first observed,, whether or not the expression is true. In this way, the EM client can obtain the initial state of the resource variable. Then the Event Management subsystem generates events as usual.

### Unregistering for Events

When the EM client no longer wishes to receive an event, it calls the **ha\_em\_send\_command** subroutine to unregister for the event, specifying the **HA\_EM\_CMD\_UNREG** command. For each event to be unregistered, the EM client supplies the event ID that was returned when the event was registered. The EM client can unregister one or more events with a single command.

### Sending Queries to the Event Management Subsystem

The EM client can also get certain kinds of information by calling the **ha\_em\_send\_command** subroutine, specifying the **HA\_EM\_CMD\_QUERY** command. Using query subcommands, the EM client can get the following information:

- The current value of one or more resource variables, using the **HA\_EM\_SCMD\_QCUR** subcommand
- The resource IDs of one or more resource variables, using the **HA\_EM\_SCMD\_QINST** subcommand
- A list of the resource variables and default expressions that are defined in the Event Management Configuration Database (EMCDB), using the **HA\_EM\_SCMD\_QDEF** subcommand.

The scope of each query is defined by the specification of a resource-variable name and a resource ID. These specifications may be wildcarded.

### Receiving Responses from the Event Management Subsystem

Once the EM client has sent a command, it must be ready to receive a response.

First, the EM client checks the session's file descriptor, using either the **poll** or the **select** system call. When the file descriptor is ready for reading, the EM client calls the **ha\_em\_receive\_response** routine.

By default, the Event Management subsystem returns the response in a response block. The response block contains one of the following:

- One or more events
- A notification that one or more events have been unregistered
- A response to a previously issued query command
- An indicator that there was an error in a previously issued register or query command.

If you prefer to program using callback routines, your EM client may specify one or more callback routines instead. A callback routine can handle one or more events, responses to unregister commands, and responses to query commands.

You specify callback routines on input to the **ha\_em\_send\_command** subroutine. You can specify a different callback routine for each event or the same callback routine for multiple events on a single register command. For queries, you can specify only one callback routine on a single query command.

When there is no error in the event or response, the Event Management subsystem delivers it directly to the callback routine. The Event Management subsystem passes information about the event or response as input to the callback routine.

For some error conditions, the Event Management subsystem always returns from the **ha\_em\_receive\_response** subroutine with a response block that contains the error, even if the EM client specified a callback routine. For details on error handling in response to commands, see “ha\_em\_receive\_response Subroutine” on page 44.

### Ending a Session

When the EM client no longer needs Event Management services, it ends the session by calling the **ha\_em\_end\_session** subroutine.

---

## Event Management Performance Considerations

The amount of overhead that the Event Management subsystem incurs is directly proportional to the number of resource variable instances observed and the number of expressions evaluated at each observation. Accordingly, you can improve performance by specifying longer reporting and observation intervals.

Also, the Event Management subsystem does not start certain daemon-based resource monitors until at least one EM client has registered interest in the resources they are monitoring.

---

## Resource Monitors and PTX Shared Memory

A resource monitor that uses Quantity and Counter variables can send the values of those variable to the Performance Monitor as well as the Event Management subsystem. Such a resource monitor uses shared memory to pass the variable values to the Performance Monitor and, in PTX terms, is a “dynamic data-supplier program.” The Performance Monitor is a PTX “local data-consumer program”. The shared memory is called “DDS Shared Memory”; it is created and managed by the RMAPI and is not directly referenced by the resource monitor program.

If your resource monitor uses shared memory in this way, it is important to make sure that the shared memory is released when the resource monitor exits, either normally or abnormally. The best way to make sure this happens is by calling the **ha\_rr\_terminate** subroutine before exiting. The **ha\_rr\_terminate** subroutine causes the RMAPI to release the shared memory. In addition to a routine that handles a normal exit, you should establish a signal handler to catch termination signals; for example, the signal handler should handle the following signals: **SIGTERM**, **SIGHUP**, and **SIGINT**.

If the resource monitor abnormally terminates without calling the **ha\_rr\_terminate** subroutine, you may need to release the shared memory manually. The following sections, adapted from PTX documentation, explain how to do this.

You may have to use this procedure when any program that uses PTX shared memory fails to release it, due to abnormal termination or any other cause. Such programs are not limited to resource monitors, but also include the Performance Monitor subsystem and any other dynamic data-supplier programs that may be running in the domain.

Note that Quantity and Counter variables are passed to the Event Management subsystem using private shared memory, distinct from the DDS shared memory. The private shared memory is created, managed, and released by the Event Management subsystem.

## DDS Shared Memory

If a dynamic data-supplier program is terminated in a way that cannot be detected from the program itself, the DDS shared memory is not released and subsequent attempts to start the dynamic data-supplier program fail. If this happens, you must release the DDS shared memory manually, as described in "Releasing Shared Memory Manually".

To avoid the situation, never kill a dynamic data-supplier program with the option **-9** That would terminate the program with a **SIGKILL** signal, which is not detectable.

## Releasing Shared Memory Manually

When one or more data-supplier or local data-consumer programs terminate without releasing their shared memory allocations, then the shared memory segments should be released from the command line before an attempt is made to restart the programs. It is recommended that all data-supplier and local data-consumer programs, including **xmservd**, be killed before you attempt to release shared memory. Clearing of all shared memory segments can be done through the following steps:

1. Identify all data-supplier and local data-consumer programs that are running. Use the **ps** command and your knowledge of the programs in use on your system to locate all of them. For each of the running data-supplier or local data-consumer programs, note their process IDs.
2. Kill all processes associated with data-supplier and local data-consumer programs without using a command line flag.
3. Verify that all data-supplier and local data-consumer processes have been killed. If not, use the **kill -9** command to kill them.
4. List the shared memory segments in use with the **ipcs -m** command. This produces a list like the following:

```
IPC status from /dev/mem as of Fri Dec 31 07:54:44 CST 1993
T ID      KEY          MODE         OWNER GROUP
Shared Memory:
m      0      0x0d050296  --rw-----  root system
m 20481    0x5806188b  --rw-rw-rw-  nchris system
m 28674    0x780502ea  --rw-rw-rw-  root system
m 12292    0x780502e3  --rw-rw-rw-  root system
m 20485    0x780502d1  --rw-rw-rw-  root system
```

5. Identify all shared memory segments with a KEY that begins with '0x78'. All shared memory allocated by data-supplier and local data-consumer programs has this key. Now use the **ipcrm** command to remove the shared memory segments, specifying as command arguments the IDs of the segments you want to remove. To remove all three data-supplier and local data-consumer segments that are listed above, your command would be:

```
ipcrm -m 28674 -m 12292 -m 20485
```

6. Restart the dynamic data-supplier and data-consumer programs as required. To start **xmservd** and any dynamic data-supplier programs started by it, enter the **xmpeek** command.

---

## Chapter 2. Event Management Configuration Data Reference

This chapter contains the reference material for the Event Management configuration data that is stored as a set of objects in the System Data Repository (SDR).

These objects are created from a loadlist file (**haemloadlist**) by the **haemloadcfg** utility command. The SDR objects are compiled using the **haemcfg** utility command. This produces a binary file that is called the Event Management Configuration Database (EMCDB).

The man page in this chapter describes the SDR classes and attributes for the Event Management objects.

For information about the **haemloadcfg** utility command, the **haemloadlist** load list file and its format, and the **haemcfg** utility command, see *PSSP Command and Technical Reference*.

**Note:** The material in this chapter applies only to SP domains. For the Event Management subsystem in an high-availability cluster multiprocessing (HACMP)/ES domain, a pre-compiled EMCDB is shipped in the RSCT install image.

---

## Event Management Configuration Data (emcdb)

### Purpose

Event Management Configuration Data (SDR) – System Data Repository (SDR) classes and attributes for Event Management resource variables, resource IDs, resource classes, and resource monitors

### Description

Before you can use Event Management services through the Event Management Application Programming Interface (EMAPI), the system resources being monitored must be defined and configured to the Event Management subsystem. This information includes the following:

- Resource variables from which events may be generated
- Resource ID for each resource variable
- Classes in which the resource variables are grouped
- Resource monitors that supply the variables

For resource variables that are structured byte strings (SBSs), the SBS structured fields must also be defined. A distinct database is created for each SP-type domain (system partition) on the IBM RS/6000 SP.

The configuration process for each domain consists of the following steps:

1. The programmer that creates the resource monitor defines the Event Management resource information and stores it in an Event Management load list file.
2. The programmer or system administrator issues the **haemloadcfg** command to load the resource definitions from the load list file into the System Data Repository (SDR).
3. The programmer or system administrator uses the **haemcfg** command to compile the SDR objects, producing a binary file called the Event Management Configuration Database (EMCDB). The new EMCDB is placed in a staging directory.
4. The programmer or system administrator activates the new EMCDB by using **haemctrl** commands to stop all of the domain's Event Manager daemons and restart them. The commands must stop and restart the domain's daemon on the control workstation and the daemon on each of the domain's nodes. This has the effect of replacing the run-time version of the EMCDB with the new EMCDB from the staging directory.

For more details on the procedures for configuring Event Management, see the Event Management subsystem chapter in *PSSP Administration Guide*.

To place information in the SDR, you must have the appropriate authorization. If you do not have the appropriate authorization, consult with your system administrator. For more information about the requirements for access to the SDR, see *PSSP Administration Guide*.

You can use SDR commands to work with objects in the Event Management classes directly. However, IBM recommends that you use load list files and the **haemloadcfg** command to work with objects in these classes.

You cannot use System Management Interface Tool (SMIT) panels to work with objects in these classes.

### EMCDB Source Information

Source information for the EMCDB is kept in the following partitioned classes in the SDR:

- Class = EM\_Resource\_Variable
- Class = EM\_Structured\_Byte\_String
- Class = EM\_Resource\_ID
- Class = EM\_Resource\_Class

- Class = EM\_Resource\_Monitor

The tables below summarize the SDR classes for Event Management and their attributes. The value in the **Type** column is the datatype of the attribute: S=string, I=integer, F=floating point.

Detailed explanations of each class and attribute, as well as examples showing their use, follow the tables.

**Class = EM\_Resource\_Variable**

<i>Attribute Name</i>	<i>Type</i>	<i>Description</i>
<b>rvName</b>	S	The name of a resource variable.
<b>rvDescription</b>	I	The ID of the message that contains the description of the variable, including its semantics.
<b>rvValue_type</b>	S	The value type of the variable.
<b>rvData_type</b>	S	The data type of the variable.
<b>rvInitial_value</b>	S	The initial value of the resource variable.
<b>rvClass</b>	S	The name of a resource variable class.
<b>rvPTX_name</b>	S	The name that is used to read and write the variable in the Performance Toolbox (PTX) shared memory.
<b>rvPTX_description</b>	S	A comma-separated list, with no blanks, of message IDs.
<b>rvPTX_min</b>	I	The lower range for plotting this variable by the Performance Monitor.
<b>rvPTX_max</b>	I	The upper range for plotting this variable by the Performance Monitor.
<b>rvExpression</b>	S	The default expression that is to be applied to the resource variable.
<b>rvEvent_description</b>	I	The ID of the message that contains a short description of the event
<b>rvLocator</b>	S	A resource ID element name.
<b>rvDynamic_instance</b>	I	If nonzero (true), instances of this variable are created dynamically by its resource monitor whenever the instance is referenced in an EMAPI command.
<b>rvIndex_element</b>	S	A resource ID element name.

**Class = EM\_Structured\_Byte\_String**

<i>Attribute Name</i>	<i>Type</i>	<i>Description</i>
<b>sbsVariable_name</b>	S	The name of the SBS resource variable for which this object defines one of its structured fields.
<b>sbsField_name</b>	S	A structured field name, which provides a short description of the structured byte field.
<b>sbsField_type</b>	S	The data type of this field.
<b>sbsField_SN</b>	I	The serial number of the structured field.
<b>sbsField_init_val</b>	S	The initial value of the structured field, with a format suitable to the field's type.

**Class = EM\_Resource\_ID**

<i>Attribute Name</i>	<i>Type</i>	<i>Description</i>
<b>riResource_name</b>	S	The name of the resource (and all of its resource variables) for which this resource ID element is defined.
<b>riElement_name</b>	S	A resource ID element name, which describes the element.
<b>riElement_description</b>	I	The ID of the message that describes the resource ID element, including its semantics.

#### **Class = EM\_Resource\_Class**

<i>Attribute Name</i>	<i>Type</i>	<i>Description</i>
<b>rcClass</b>	S	The name of a resource variable class.
<b>rcResource_monitor</b>	S	The name of the resource monitor definition for the resource monitor that supplies the resource variables in this class.
<b>rcObservation_interval</b>	I	The observation interval, which is the amount of time, in seconds, between each observation by the Event Management subsystem of instances of any variable of value type Counter or Quantity in this class.
<b>rcReporting_interval</b>	I	The reporting interval, which is the amount of time, in seconds, between each update of any variable of value type Counter or Quantity in this class by the corresponding resource monitor.
<b>rcInstance_limit</b>	I	The maximum number of resource variable instances the Event Management subsystem will accept from the resource monitor for this class.

#### **Class = EM\_Resource\_Monitor**

<i>Attribute Name</i>	<i>Type</i>	<i>Description</i>
<b>rmName</b>	S	A resource monitor definition name.
<b>rmPath</b>	S	The executable absolute path name of the resource monitor if it is a daemon that can be started by the Event Management subsystem.
<b>rmArguments</b>	S	A list of optional arguments that the Event Management subsystem should pass when it starts the resource monitor daemon.
<b>rmMessage_file</b>	S	The name of the message catalog that contains all of the descriptions for the resource variables that are supplied by this resource monitor.
<b>rmMessage_set</b>	I	The set ID of the descriptions for the resource variables that are supplied by this resource monitor.
<b>rmConnect_type</b>	S	The type of connection.
<b>rmPTX_prefix</b>	S	The PTX name prefix.
<b>rmPTX_description</b>	S	A comma-separated list, with no blanks, of message IDs.
<b>rmPTX_asnno</b>	I	The ASN.1 number equal to the SNMP Assigned Enterprise Number for the vendor that supplies this resource monitor.
<b>rmNum_instances</b>	I	The maximum number of resource monitor instances.

#### **Class = EM\_Resource\_Variable**

The **EM\_Resource\_Variable** class contains one object for each resource variable that is defined in the database.

Some of the attributes in this class are used to meet PTX requirements. For more information on PTX, see *Performance Toolbox for AIX Guide and Reference*.

The attributes of this class are defined as follows:

#### **rvName**

A string that contains the name of a resource variable. It consists of a resource name that is followed by a period and the resource attribute. For details on specifying the components of a resource-variable name, see “Resource Variables (haemrvars)” on page 126.

Resource variables that represent different attributes of the same resource must contain the same resource name.

#### **rvDescription**

An integer that identifies the message that contains the description of the variable, including its semantics. The first line of the message must be no more than 63 bytes, excluding the newline character at the end of the line. It should be a brief abstract of the description.

The description provides information about the resource variable, in particular its semantics. It must provide enough information for a programmer or administrator to understand the following:

- Information contained in the resource variable
- Motivation for the default expression
- Under what circumstances additional expressions could be provided through the EMAPI

If the variable is of value type Counter or Quantity, the first line of the message is used when the variable is viewed through the Performance Monitor.

To find the message, use the message ID with the message set and message catalog attributes of the **EM\_Resource\_Monitor** class object for the resource monitor that defines the resource variable.

#### **rvValue\_type**

A string that indicates the value type of the variable. It may be one of the following: Counter, Quantity, or State. Case is significant.

A **Counter** is a variable whose value increases monotonically. By default, a counter is interpreted as a rate. The change in the value of an instance of the counter from one observation to the next is divided by the time, in seconds, between the two observations. However, the actual contents of the instance of the counter from the latest observation, called the raw value, is also available.

A Counter is typically used to represent an attribute of a resource that is an indicator of throughput. Instances of a Counter are observed by the Event Management subsystem at an interval, in seconds, that is configurable for the class to which the Counter belongs. Instances of a Counter may be updated oftener than they are observed without losing meaningful information for generating events.

If the value of the Counter is also to be passed to the Performance Monitor, a PTX name is required for this variable and all other Counter and Quantity variables defined for the resource monitor. See below.

A **Quantity** is a variable whose value fluctuates over time. It is typically used to represent an attribute of a resource that is a level, that is, an indicator of how many. Instances of a Quantity are observed by the Event Management subsystem at an interval, in seconds, that is configurable for the class to which the Quantity belongs. Instances of a Quantity may be updated more often than they are observed without losing meaningful information for generating events.

If the value of the Quantity is also to be passed to the Performance Monitor, a PTX name is required for this variable and all other Counter and Quantity variables defined for the resource monitor. See below.

A **State** is a variable whose value fluctuates over time, like a Quantity. However, every time an instance of a State resource variable is updated, it must be observed so that potential events are not missed.

A State variable can be used to represent an attribute of a resource that indicates anything other than throughput or a level. A resource monitor passes instances of a State variable to the Event Management subsystem as messages. Thus, an instance of a State variable is observed whenever it is received by the Event Management subsystem, rather than at regular intervals.

#### **rvData\_type**

A string that indicates the data type of the variable. It may be one of the following: long, float, or SBS. Case is significant.

The long and float formats are identical to the C language types of the same names.

The SBS data type specifies a structured byte string and is permitted only for a variable of value type State. If specified, the structured fields that make up the structured byte string must be defined as objects in the **EM\_Structured\_Byte\_String** class.

#### **rvInitial\_value**

A string that contains the initial value of the resource variable.

If the data type of the variable is long, the initial value is an integer constant in a format suitable for conversion by the **strtol** subroutine (with a base parameter of 0).

If the data type of the variable is float, the initial value is a floating point constant in a format suitable for conversion by the **strtof** subroutine.

If the data type of the variable is structured byte string, this attribute value must be the null string.

#### **rvClass**

A string that contains the name of a resource variable class.

The **class** of a resource variable indicates the application or subsystem that manages the associated resource and can also be used to group variables by observation interval.

For details on specifying the name of a resource variable class, see “Resource Variables (haemrvars)” on page 126.

#### **rvPTX\_name**

A string that contains the name that is used to write the variable into the Performance Toolbox (PTX) DDS shared memory. This attribute is required for a variable with a value type of Counter or Quantity that is to be supplied to the Performance Monitor. If the variable is not to be supplied to the Performance Monitor this attribute is omitted. Note that either all Counter and Quantity variables defined for a resource monitor must specify this attribute, or all must omit this attribute.

The **PTX name** used to place a variable into DDS shared memory must follow the PTX naming convention. If the PTX name contains components that are instances, these component values must be taken from one or more elements of the variable’s resource ID. This implies that the PTX name and the resource ID for the resource variable are designed in concert.

Where a PTX name contains a component that is an instance, that part of the name is specified by the name of the corresponding resource ID element, prepended with the dollar sign. For example, the following would have the corresponding PTX names defined in the configuration file:

Resource Variable Name	Resource ID
IBM.PSSP.CSS.ibytes_lsw	NodeNum=5
IBM.PSSP.CSS.oerrors	NodeNum=7
IBM.PSSP.VSD.blocks_rw	NodeNum=5;L1=4;L2=5; VSD=vsdn2
IBM.PSSP.VSDdrv.rejected_requests	NodeNum=5

```
SomeDatabase.transactions      NodeNum=5;DBname=foo
VendorA.Resource1.Server.counter NodeNum=8;Group=gapg33;
                               ServerID=bkg21
```

An example of a configuration file follows:

```
PTX Name

CSS/ibytes_lsw
CSS/oerrors
VSD/$L1/$L2/$VSD/blocks_rw
VSDdrv/rej_requests
SomeDatabase/$DBname/transactions
Resource1/Server/$Group/$ServerID/counter
```

The Event Management subsystem replaces PTX name components that begin with a dollar sign by the value of the corresponding resource ID element.

### **rvPTX\_description**

A string that contains a comma-separated list, with no blanks, of message IDs. Each ID specifies a message that contains a short description (no more than 63 bytes) of a PTX context.

The first ID in the list corresponds to the context specified by the first component of the name specified by the **rvPTX\_name** attribute. The second ID in the list corresponds to the context specified by the first two components of the name. The third ID corresponds to the context specified by the first three components of the name, and so on. However, there is one fewer ID in this string than the number of components in **rvPTX\_name**. (The last component in the PTX name is a statistic, not a context.)

This attribute is required if the **rvPTX\_name** attribute is specified. Otherwise, it is omitted.

To find the message, use the message ID with the message set and message catalog attributes of the **EM\_Resource\_Monitor** class object for the resource monitor that defines the resource variable.

### **rvPTX\_min**

An integer that is equal to the lower range for plotting this variable by the Performance Monitor.

This attribute is required if the **rvPTX\_name** attribute is specified. Otherwise, it is omitted.

### **rvPTX\_max**

An integer that is equal to the upper range for plotting this variable by the Performance Monitor.

This attribute is required if the **rvPTX\_name** attribute is specified. Otherwise, it is omitted.

### **rvExpression**

A string that contains the default expression that is to be applied to the resource variable.

If you do not wish to specify a default expression, omit this attribute.

For details on how to specify an expression, see “Expressions (haemexpr)” on page 123.

### **rvEvent\_description**

An integer that equals the ID of the message that contains a short description of the event generated when the default expression is applied to the variable. If a default expression is not specified, this attribute is omitted.

### **rvLocator**

A string that contains a resource ID element name.

If the resource ID for this resource implies the resource’s location, the value of this item is the name of the resource ID element whose value is the number of the node that contains the resource monitor which supplies the resource instance. Values of the resource ID element that is specified by the **rvLocator** attribute must be of type integer.

If this attribute is omitted, the Event Management subsystem must determine the location of the resource monitor.

#### **rvDynamic\_instance**

An integer that, if nonzero (true), indicates that instances of this variable are created dynamically by its resource monitor whenever the instance is referenced in an EMAPI command.

#### **rvIndex\_element**

A string that contains a resource ID element name.

If a resource ID element for this resource variable can be used as an index into an array that represents all of the variable's instances, specify its name. However, you cannot specify the same element for both the **rvLocator** and the **rvIndex\_element** fields.

Values of the resource ID element that is specified by the **rvIndex\_element** attribute must be of type integer.

Specifying an index element can improve the performance of the Event Management subsystem.

If more than one of the resource ID elements can be used as an index, specify the one that can take on the larger number of possible values. For example, if one resource ID element represents a node number in a 128-node system and another represents a frame number in the same system, choose the element that represents the node number. This is because there are more nodes than frames.

### **Class = EM\_Structured\_Byte\_String**

The **EM\_Structured\_Byte\_String** class contains one object for each structured field that is defined in a structured byte string.

The attributes of this class are defined as follows:

#### **sbsVariable\_name**

A string that contains the name of the SBS resource variable for which this object defines one of its structured fields.

#### **sbsField\_name**

A string that contains a structured field name, which provides a short description of the structured byte field. The name must consist of only alphanumeric characters and underscores; the first character must be alphabetic. This name is available through the Event Management API.

#### **sbsField\_type**

A string that indicates the data type of this field. It can be one of the following values: **long**, **float**, **cstring**, or **bstring**, indicating a structured field of data type long, float, character string, or byte string, respectively. The long and float data types are the same as in the C language. A character string consists of one or more nonzero bytes terminated by a null byte; the null byte is included in the structured field value length. A byte string consists of one or more bytes, where each byte may have any value from 0 through 255.

#### **sbsField\_SN**

An integer that is the serial number of the structured field. The serial number is a one-byte value that uniquely identifies the structured field within the structured byte string. This serial number is defined for each structured field by the resource monitor that supplies the SBS resource variable. However, the set of serial numbers for the structured byte string starts with 0 and continues sequentially, incrementing by 1.

#### **sbsField\_init\_val**

A string that contains the initial value of the structured field, with a format suitable to the field's type.

If the type of the structured field is long, the initial value is an integer constant in a format suitable for conversion by the **strtol** subroutine (with a base parameter of 0).

If the type of the structured field is float, the initial value is a floating point constant in a format suitable for conversion by the **strtof** subroutine.

Otherwise, the initial value is a string constant. To indicate nondisplayable characters, a string constant may contain `\ooo`, where `ooo` is an octal constant that represents the character. If the type of the structured field is **cstring**, a terminating null character is appended automatically.

If this attribute is omitted, default initial values are used. The default values are 0 for long fields, 0.0 for float fields, the null string for character strings and, for byte strings, a string consisting of a single byte with a value of 0.

### **Class = EM\_Resource\_ID**

The **EM\_Resource\_ID** class contains one object for each resource ID element that is defined for a resource and all of its resource variables.

The attributes of this class are defined as follows:

#### **riResource\_name**

A string that contains the name of the resource for which this resource ID element is defined.

#### **riElement\_name**

A string that contains a resource ID element name, which describes the element. The name must consist of only alphanumeric characters and underscores; the first character must be alphabetic.

#### **riElement\_description**

An integer that is the ID of the message that describes the resource ID element, including its semantics.

To find the message, use the message ID with the message set and message catalog attributes of the **EM\_Resource\_Monitor** class object for the resource monitor that defines the resource variable of which this is a resource ID.

### **Class = EM\_Resource\_Class**

The **EM\_Resource\_Class** class contains one object for each resource variable class that is defined in the database.

The attributes of this class are defined as follows:

#### **rcClass**

A string that contains the name of a resource variable class.

The **class** of a resource variable indicates the application or subsystem that manages the associated resource and can also be used to group variables by observation interval.

For details on specifying the name of a resource variable class, see “Resource Variables (haemrvars)” on page 126.

#### **rcResource\_monitor**

A string that contains the name of the resource monitor definition for the resource monitor that supplies the resource variables in this class.

#### **rcObservation\_interval**

An integer that is the observation interval, which is the amount of time, in seconds, between each observation by the Event Management subsystem of instances of any variable of value type Counter or Quantity in this class. This value must be greater than or equal to the reporting interval.

#### **rcReporting\_interval**

An integer that is the reporting interval, which is the amount of time, in seconds, between each update of any variable of value type Counter or Quantity in this class by the corresponding resource monitor.

The value of the reporting interval may be 0 if the design of the resource monitor fixes the interval between variable updates. It may also be 0 if the resource monitor is incorporated into a subsystem and the subsystem updates the variable as part of its normal execution.

#### **rcInstance\_limit**

An integer that is the limit of the number of resource variable instances for this class that the Event Management subsystem will accept from a resource monitor. If this attribute is omitted, the default value is **HA\_EM\_MAX\_INSTS**. If specified, the attribute value should be a lower value than **HA\_EM\_MAX\_INSTS**. A higher value is ignored and the default value is used.

**HA\_EM\_MAX\_INSTS** is defined in the **ha\_emcommon.h** header file. See “ha\_emapi.h File” on page 110.

Class = EM\_Resource\_Monitor

The **EM\_Resource\_Monitor** class contains one object for each resource monitor that is defined in the database.

Some of the attributes in this class are used to meet PTX requirements. For more information on PTX, see the *Performance Toolbox for AIX Guide and Reference*.

The attributes of this class are defined as follows:

#### **rmName**

A string that contains a resource monitor definition name. This name must be unique in the database.

By convention, resource monitor names follow the same hierarchical naming conventions that are used for resource names except that they may not contain percent signs. For more information, see “Resource Variables (haemrvars)” on page 126.

#### **rmPath**

A string that contains the executable absolute path name of the resource monitor if it is a daemon that can be started by the Event Management subsystem. If the resource monitor cannot be started by the Event Management subsystem, this attribute is omitted.

#### **rmArguments**

A string that contains a list of optional arguments that the Event Management subsystem should pass when it starts the resource monitor daemon. If there are no arguments, this attribute is omitted.

#### **rmMessage\_file**

A string that contains the name of the message catalog that contains all of the descriptions for the resource variables that are supplied by this resource monitor.

For information about message catalogs, see the “Message Facility Overview for Programming” topic in *AIX Version 4.3 General Programming Concepts: Writing and Debugging Programs*.

#### **rmMessage\_set**

An integer that is the set ID of the descriptions for the resource variables that are supplied by this resource monitor.

For information about message sets, see the “Message Facility Overview for Programming” topic in *AIX Version 4.3 General Programming Concepts: Writing and Debugging Programs*.

#### **rmConnect\_type**

A string that indicates the type of connection. It may be one of the following: **client** or **server**.

If the resource monitor is to connect to the Event Management subsystem, specify **client**. The resource monitor is a command.

If the Event Management subsystem and the Performance Monitor subsystem are to connect to the resource monitor, specify **server**. In this case, the resource monitor is either a daemon or incorporated within a subsystem.

#### **rmPTX\_prefix**

A string that contains the PTX name prefix. This prefix is prepended to each PTX name for resource variables of value type Counter or Quantity that are supplied by this resource monitor to the Performance Monitor. If no variables are to be supplied to the Performance Monitor, then this attribute is omitted.

Resource variables of value type Counter or Quantity that are supplied by a resource monitor to the Performance Monitor are also implemented as PTX variables and must therefore conform to the PTX naming architecture. The PTX name for either of these types of variables is formed by concatenating the following strings:

- The string **DDS/**
- The string specified by the **rmPTX\_prefix** attribute of the **EM\_Resource\_Monitor** object that defines the resource monitor.
- The string specified by the **rvPTX\_name** attribute of the **EM\_Resource\_Variable** object.

The PTX name prefix is expected to consist of at least one component but probably not more than two or three. The prefix is used to group all of the resource variables of value type Counter or Quantity that are supplied by a resource monitor to the Performance Monitor.

If specified, each resource monitor must have a unique PTX name prefix. For PSSP resource monitors, the prefix is **IBM/PSSP.res\_mon\_name**. For vendor-supplied resource monitors, IBM recommends specifying the PTX name prefix as *vendor\_name.res\_mon\_name*.

#### **rmPTX\_description**

A string that contains a comma-separated list, with no blanks, of message IDs. Each ID specifies a message that contains a short description (no more than 63 bytes) of a PTX context.

The first ID in the list corresponds to the context specified by the first component of the name specified by the **rmPTX\_prefix** attribute, the second ID in the list corresponds to the context specified by the first two components of the name, the third ID corresponds to the context specified by the first three components of the name, and so on.

This attribute is required if the **rmPTX\_prefix** attribute is specified. Otherwise, it is omitted.

To find the message, use the message ID with the message set and message catalog attributes of the **EM\_Resource\_Monitor** class object for this resource monitor.

#### **rmPTX\_asnno**

An integer that is the ASN.1 number equal to the SNMP Assigned Enterprise Number for the vendor that supplies this resource monitor. For IBM supplied resource monitors, this value is 2.

This attribute is required if the **rmPTX\_prefix** attribute is specified. Otherwise, it is omitted.

#### **rmNum\_instances**

An integer that is the maximum number of executing instances of the resource monitor. If this attribute is omitted the default value is 1.

A **server** resource monitor program may be executed in multiple processes, as long as each process supplies unique resource variable instances. If a resource monitor supplies variables to the Performance Monitor, it must be instance number 0 (resource monitor instances are numbered starting at 0). If a resource monitor defines the **rmPath** attribute, the Event Management subsystem will only start one resource monitor process. It is assumed that this resource monitor process will start other instances. Furthermore, the Event Management subsystem will not start a resource monitor instance as long as it has connections to any other instances of the resource monitor.

## Related Information

See the following PSSP commands and files in *PSSP Command and Technical Reference*.

Commands: **haemloadcfg**, **haemcfg**, **haemctrl**

Files: **haemloadlist**

For more information about configuring and managing the Event Management subsystem, see *PSSP Administration Guide*.

---

## Chapter 3. Event Management Subroutine Reference

This chapter contains the reference material for each of the subroutines in the Event Management Application Programming Interface (EMAPI) and the Resource Monitor Application Programming Interface (RMAPI). Following the next two sections, which summarize each API's subroutines by function, the subroutines are listed alphabetically.

---

### EMAPI Subroutine Summary

The EMAPI consists of the following subroutines:

**Subroutine**

**Action**

**ha\_em\_start\_session**

Establish a session with the Event Management subsystem

**ha\_em\_end\_session**

End a session with the Event Management subsystem

**ha\_em\_restart\_session**

Reconnect a session to the Event Management subsystem

**ha\_em\_send\_command**

Send a command to the Event Management subsystem

Commands are provided to register for events, unregister events, and query the Event Management subsystem for information about resource variables.

**ha\_em\_receive\_response**

Receive a response from the Event Management subsystem

**ha\_em\_get\_ecgid**

Get an event command group ID

---

### RMAPI Subroutine Summary

The RMAPI consists of the following subroutines:

**Subroutine**

**Action**

**ha\_rr\_rm\_ctl**

Set or get attributes of the RMAPI

**ha\_rr\_init**

Initialize the Resource Monitor Application Programming Interface (RMAPI)

**ha\_rr\_makserv**

Establish the resource monitor as a server to enable resource monitor managers to connect

**ha\_rr\_start\_session**

Start a session with a resource monitor manager

**ha\_rr\_get\_ctrlmsg**

Get a control message from a resource monitor manager

**ha\_rr\_reg\_var**

Register a resource variable instance with the RMAPI

**ha\_rr\_add\_var**

Add registered variables to a resource monitor manager session

**ha\_rr\_get\_interval**

Get the reporting interval for a class of resource variables

**ha\_rr\_send\_val**

Send variable values to the RMAPI

**ha\_rr\_touch**

Meet "send" frequency requirements when there are no resource variable values to send

**ha\_rr\_del\_var**

Delete resource variables from a resource monitor manager session

**ha\_rr\_unreg\_var**

Unregister a variable instance

**ha\_rr\_end\_session**

End an RMAPI session

**ha\_rr\_terminate**

Free resources and terminate use of the RMAPI

---

## ha\_em\_end\_session Subroutine

### Purpose

**ha\_em\_end\_session** – End an EM client session with the Event Management subsystem

### Library

EMAPI Thread-Safe Library (**libha\_em\_r.a**)

EMAPI Library (not thread-safe) (**libha\_em.a**)

### Syntax

```
#define HA_EM_VERSION 2
#include <ha_emapi.h>

int
  ha_em_end_session(
    int          session_fd,
    struct ha_em_err_blk *em_errb)
```

### Parameters

*session\_fd*

The file descriptor of the session that is to be ended. This file descriptor was returned by a previous call to either the **ha\_em\_start\_session** or **ha\_em\_restart\_session** subroutine.

*em\_errb*

A pointer to an error block structure.

### Description

The **ha\_em\_end\_session** subroutine is used by an EM client to end a session.

### Restrictions

Using the **ha\_em\_end\_session** subroutine to end a session causes all event registrations for the session to be lost. To continue to use an existing session to which the connection has been lost, use the **ha\_em\_restart\_session** subroutine to reestablish the connection.

When you use the EMAPI in a threaded program, IBM recommends that an Event Management session be used by only one thread at a time. If multiple threads try to use the same Event Management session at the same time, EMAPI subroutines may return the **HA\_EM\_EBUSY** error code. The thread-safe version of the EMAPI library was designed assuming that each thread using the EMAPI would have its own Event Management session or sessions.

### Return Values

If the **ha\_em\_end\_session** subroutine is successful, it returns a value of 0 and the session is ended. The file descriptor that was specified on input must no longer be used as an argument to the **select** or **poll** system call.

### Error Values

If the **ha\_em\_end\_session** subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

The EMAPI error block and error numbers are defined in the **ha\_emapi.h** header file. For more information on EMAPI errors, see “EMAPI Errors (err\_emapi)” on page 104.

For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

The file descriptor that was provided on input in the `session_fd` argument is still valid and can be used in another attempt to end the session.

## Examples

For examples of using EAPI subroutines, see the programs in the RSCT product samples directory, `/usr/sbin/rsct/samples/haem/emapi`.

## Files

`ha_emapi.h`.

## Prerequisite Information

“Chapter 1. Understanding Event Management” on page 1.

## Related Information

Structures in the `ha_emapi.h` header file: `ha_em_err_blk`

Subroutines: `ha_em_start_session`, `ha_em_restart_session`

---

## ha\_em\_get\_ecgid Subroutine

### Purpose

**ha\_em\_get\_ecgid** – Get an event command group ID

### Library

EMAPI Thread-Safe Library (**libha\_em\_r.a**)

EMAPI Library (not thread-safe) (**libha\_em.a**)

### Syntax

```
#define HA_EM_VERSION 2
#include <ha_emapi.h>

ha_em_ecgid_t
  ha_em_get_ecgid(
    ha_em_eid_t event_id)
```

### Parameters

*event\_id*

An event ID. This event ID was returned when the caller registered for events by a previous call to the **ha\_em\_send\_command** subroutine using the **HA\_EM\_CMD\_REG** or **HA\_EM\_CMD\_REG2** command.

### Description

The **ha\_em\_get\_ecgid** subroutine is used by an EM client to get the ID of the event command group to which the event specified on input belongs. All events for which the caller registered through a single **HA\_EM\_CMD\_REG** or **HA\_EM\_CMD\_REG2** command belong to the same event command group.

### Return Values

The **ha\_em\_get\_ecgid** subroutine returns the event command group ID. This event command group ID is valid until all of the events that were registered by the command are unregistered. Once these events are unregistered, the EMAPI may reuse the event command group ID.

### Error Values

None.

### Examples

For examples of using EMAPI subroutines, see the programs in the RSCT product samples directory, **/usr/sbin/rsct/samples/haem/emapi**.

### Files

**ha\_emapi.h**.

### Prerequisite Information

“Chapter 1. Understanding Event Management” on page 1.

### Related Information

Subroutine: **ha\_em\_send\_command**

---

## ha\_em\_receive\_response Subroutine

### Purpose

**ha\_em\_receive\_response** – Receive a response from the Event Management subsystem

### Library

EMAPI Thread-Safe Library (**libha\_em\_r.a**)

EMAPI Library (not thread-safe) (**libha\_em.a**)

### Syntax

```
#define HA_EM_VERSION 2
#include <ha_emapi.h>

int
  ha_em_receive_response(
    int          session_fd,
    struct ha_em_rsp_blk **em_rsp_blk,
    struct ha_em_err_blk *em_errb)
```

### Parameters

*session\_fd*

The file descriptor of the Event Management session for which the response is to be received.

*em\_rsp\_blk*

A pointer to the address of a buffer containing a response block structure. The response block contains information returned by the Event Management subsystem in response to a previously issued command.

*em\_errb*

A pointer to an error block structure.

### Description

The **ha\_em\_receive\_command** subroutine is used by an EM client to receive information returned in response to a previously issued command.

When an EM client starts or restarts a session, it obtains a file descriptor. The EM client uses this file descriptor as an argument to the **select** or **poll** system call. When the system call that is used indicates that there is data to be read from the file descriptor, the EM client calls the **ha\_em\_receive\_response** subroutine to read the data and thereby receive the response from the Event Management subsystem.

The *session\_fd* argument contains the file descriptor of the Event Management session that has data to be read.

The **ha\_em\_rsp\_blk** response block has the following definition:

```
struct ha_em_rsp_blk {
  int          em_rsp_blk_len;
  int          em_rsp_num_resp;
  short       em_cmd;
  short       em_subcmd;
  ha_em_qid_t em_qid;
  int         em_qend;
  union ha_em_resp_blk {
    struct ha_em_rpb_event em_rpb_event[1];
    struct ha_em_rpb_qcur  em_rpb_qcur[1];
    struct ha_em_rpb_qdef  em_rpb_qdef[1];
    struct ha_em_rpb_rerr  em_rpb_rerr[1];
    struct ha_em_rpb_qerr  em_rpb_qerr[1];
  } em_resp_blk;
}
```

In an actual response block, the arrays in the **em\_resp\_blk** union may contain any number of elements; the **em\_rsp\_num\_resp** field contains the actual number that are returned. The **em\_cmd** and **em\_subcmd** fields contain the values of the command and subcommands to which this is a response. If the response block contains information in response to a query and the **em\_qend** field is nonzero, this response is the last piece of information to be returned for the query specified by the **em\_qid** field. The remaining fields in the response block are used as specified in the response descriptions that follow.

## Responses to Register Event Commands

When the buffer returned by the **ha\_em\_receive\_response** subroutine contains a response block in which the value of the **em\_cmd** field is **HA\_EM\_CMD\_REG** or **HA\_EM\_CMD\_REG2**, the response block contains one or more events. The **em\_rpb\_event** array is used in the response block.

The **ha\_em\_rpb\_event** structure has the following definition:

```
struct ha_em_rpb_event {
    union ha_em_errnum em_error;
    ha_em_eid_t em_event_id;
    unsigned long em_event_flags;
    struct timeval em_timestamp;
    int em_location;
    char *em_name;
    char *em_rsrc_ID;
    enum ha_emData_Type em_data_type;
    union ha_em_val {
        long em_vall;
        float em_valf;
        void *em_valsubs;
    } em_val;
}
```

The **ha\_em\_errnum** union contains information about any error that may have occurred. It has the following definition:

```
union ha_em_errnum {
    unsigned int em_error_number;
    unsigned short em_error_codes[2];
};
```

```
#define em_errnum em_error.em_error_number /* for quick tests */
#define em_generr em_error.em_error_codes[0] /* general error code */
#define em_specerr em_error.em_error_codes[1] /* specific error code */
```

The **em\_error\_number** field indicates whether there is an error. If there is no error, the **em\_error\_number** field contains a value of 0 and the remainder of **ha\_em\_rpb\_event** information should be used.

If the **em\_error\_number** field is nonzero, there is an error and the validity of the remaining fields in the **ha\_em\_rpb\_event** structure depend on the error. For more information, see the section on error responses to register and unregister commands later in this man page.

The **em\_event\_id** field contains the ID of the event. The value can be used to match this event to the information that was used to register for the event.

The **em\_event\_flags** field contains one of the following bit flags:

### HA\_EM\_EVENT\_RE\_ARM

The event was generated from the rearmed expression.

### HA\_EM\_EVENT\_EXPR\_FALSE

The expression is false. False events can be generated when the **HA\_EM\_SCMD\_REVAL** subcommand is specified on event registration. They can also be generated once the value of the resource variable instance is known after it had been unknown due to an error.

The **em\_timestamp** field is the time that the event was generated by the Event Management subsystem. The **em\_location** field contains the number of the node where the event was generated.

The **em\_name** field is a pointer to a string that contains the name of the resource variable that was used to register the expression that generated the event. The **em\_rsrc\_ID** field is a pointer to the resource ID of the variable instance from which the event was generated. The **em\_val** field is the current value of the variable instance at the time of the event. The **em\_data\_type** field indicates whether the value is a long, a float, or a structured byte string. For information about resource variable definitions, see “Resource Variables (haemrvs)” on page 126.

If, when registering for this event, the EM client supplied a pointer to a callback routine, then instead of returning the **ha\_em\_rpb\_event** structure in a buffer, the **ha\_em\_receive\_response** subroutine invokes the callback routine. The session file descriptor, a pointer to the **ha\_em\_rpb\_event** structure, and the value of the **em\_cb\_arg** field, which was specified when registering for this event, are all passed to the callback routine. The memory addressed by the pointer to the **ha\_em\_rpb\_event** structure must not be freed or modified by the callback routine. If all of the events received by the invocation of the **ha\_em\_receive\_response** subroutine have callback routines, the **ha\_em\_receive\_response** subroutine returns a value of 0. Otherwise, it returns a value greater than zero and also returns in the response block all of the events that were not processed by callback routines.

### Responses to Unregister Event Commands

When the buffer returned by the **ha\_em\_receive\_response** subroutine contains a response block in which the value of the **em\_cmd** field is **HA\_EM\_CMD\_UNREG**, the response block contains one or more unregister event responses. The **em\_rpb\_event** array is used in the response block, which was described earlier in this man page.

The **em\_event\_id** field contains the ID of the event that has been unregistered. The **em\_event\_flags** field contains the **HA\_EM\_EVENT\_UNREG** flag. This flag is used so that the callback routine, if any is specified, can distinguish between an event and an unregister response. The **em\_timestamp** field is the time that the response was generated by the Event Management subsystem. All other fields in the **em\_rpb\_event** array element are undefined.

If the EM client supplied a callback routine for an event that is specified in one of the responses, then instead of returning the **ha\_em\_rpb\_event** structure in a buffer, the **ha\_em\_receive\_response** subroutine invokes the callback routine. The session file descriptor, a pointer to the **ha\_em\_rpb\_event** structure, and the value of the **em\_cb\_arg** field, which was specified when registering for the event, are passed to the callback routine. The memory addressed by the pointer to the **ha\_em\_rpb\_event** structure must not be freed or modified by the callback routine. If all of the events specified in the responses have callback routines, the **ha\_em\_receive\_response** subroutine returns a value of 0. Otherwise, it returns a value greater than zero and also returns in the response block all of the unregistered event responses that were not processed by callback routines.

### Responses to Queries about Current Values

When the buffer returned by the **ha\_em\_receive\_response** subroutine contains a response block in which the value of the **em\_cmd** field is **HA\_EM\_CMD\_QUERY** and the value of the **em\_subcmd** field is **HA\_EM\_SCMD\_QCUR**, the response is a reply to a query for the current values of the resource variables that were targeted in the query. The **em\_qid** field contains the unique value that was returned to the EM client when the query command was issued. The **em\_rpb\_qcur** array is used in the response block.

The **ha\_em\_rpb\_qcur** structure has the following definition:

```
struct ha_em_rpb_qcur {
    union ha_em_errnum {
        int          em_error;
        int          em_location;
    }
    char          *em_name;
    char          *em_rsrc_ID;
    int          em_data_type;
}
```

```

union ha_em_val {
    long      em_val1;
    float     em_valf;
    void      *em_valsubs;
} em_val;

```

The **ha\_em\_errnum** union contains information about any error that may have occurred. It has the following definition:

```

union ha_em_errnum {
    unsigned int    em_error_number;
    unsigned short  em_error_codes[2];
};

#define em_errnum    em_error.em_error_number    /* for quick tests */
#define em_generr    em_error.em_error_codes[0] /* general error code */
#define em_specerr   em_error.em_error_codes[1] /* specific error code */

```

The **em\_error\_number** field indicates whether there is an error. If there is no error, the **em\_error\_number** field contains a value of 0 and the remainder of **ha\_em\_rpb\_qcur** information should be used.

If the **em\_error\_number** field is nonzero, there is an error and the validity of the remaining fields in the **ha\_em\_rpb\_qcur** structure depend on the error. For more information, see the section on error responses to query commands later in this man page.

The **em\_location** field contains the number of the node running the resource monitor supplying the resource variable instance. The **em\_name** field is a pointer to a string that contains the name of the resource variable. The **em\_rsrc\_ID** field is a pointer to a string that contains the resource ID of the variable instance. The **em\_val** field is the current value of the variable. The **em\_data\_type** field indicates whether the value is a long, a float, or a structured byte string. For information about resource variable definitions, see “Resource Variables (haemvars)” on page 126.

If, when issuing the query command that resulted in this response, the EM client supplied a pointer to a callback routine, then instead of returning the **ha\_em\_rsp\_blk** structure in a buffer, the **ha\_em\_receive\_response** subroutine invokes the callback routine. The session file descriptor, a pointer to the **ha\_em\_rsp\_blk** structure, and the value of the **em\_qcb\_arg** field, which was specified when the query command was issued, are passed to the callback routine. The memory addressed by the pointer to the **ha\_em\_rsp\_blk** structure or by any of the pointers contained within it must not be freed or modified by the callback routine. The **ha\_em\_receive\_response** subroutine then returns a value of 0.

## Responses to Queries about Instances

When the buffer returned by the **ha\_em\_receive\_response** subroutine contains a response block in which the value of the **em\_cmd** field is **HA\_EM\_CMD\_QUERY** and the value of the **em\_subcmd** field is **HA\_EM\_SCMD\_QINST**, the response is a reply to a query for the instances of the resource variables that were targeted in the query.

The **em\_qid** field contains the unique value that was returned to the EM client when the query was issued. The **em\_rpb\_qcur** array is used in the response block.

The definition of the **ha\_em\_rpb\_qcur** structure is shown in the section “Responses to Queries about Current Values” on page 46, and the fields have the same meanings as described in that section, with the exception of the **em\_val** field. The **em\_val** field is the last known value of the variable. It may or may not be the current value.

## Responses to Queries about Defined Variables and Expressions

When the buffer returned by the **ha\_em\_receive\_response** subroutine contains a response block in which the value of the **em\_cmd** field is **HA\_EM\_CMD\_QUERY** and the value of the **em\_subcmd** field is **HA\_EM\_SCMD\_QDEF**, the response is a reply to a query for a list of resource variables, targeted in the query, that are defined in the Event Management Configuration Database (EMCDB). The **em\_qid** field contains the unique value that was returned to the EM client when the query command was issued. The **em\_rpb\_qdef** array is used in the response block.

The **ha\_em\_rpb\_qdef** structure has the following definition:

```
struct ha_em_rpb_qdef {
    union ha_em_errnum {
        em_error;
    }
    char      *em_name;
    char      *em_descr;
    enum ha_emValue_Type  em_value_type;
    enum ha_emData_Type  em_data_type;
    char      *em_sbs_format;
    char      *em_init_value;
    char      *em_class;
    char      *em_rsrc_ID;
    char      *em_rsrc_ID_descr;
    char      *em_ptx_name;
    char      *em_dflt_expr;
    char      *em_event_descr;
    char      *em_locator;
    char      *em_order_group;
}
```

The **ha\_em\_errnum** union contains information about any error that may have occurred. It has the following definition:

```
union ha_em_errnum {
    unsigned int  em_error_number;
    unsigned short em_error_codes[2];
};
```

```
#define em_errnum    em_error.em_error_number    /* for quick tests */
#define em_generr    em_error.em_error_codes[0] /* general error code */
#define em_specerr   em_error.em_error_codes[1] /* specific error code */
```

The **em\_error\_number** field indicates whether there is an error. If there is no error, the **em\_error\_number** field contains a value of 0 and the remainder of **ha\_em\_rpb\_qdef** information should be used.

If the **em\_error\_number** field is nonzero, there is an error and the validity of the remaining fields the **ha\_em\_rpb\_qdef** structure depend on the error. For more information, see the section on error responses to query commands later in this man page.

The **em\_name** field is a pointer to a string that contains the name of the resource variable. The **em\_descr** field is a pointer to a string that contains the text of the resource variable description.

The **em\_value\_type** field is one of Counter, Quantity, or State. The **em\_data\_type** field indicates whether the variable value is a long, a float, or a structured byte string. If the variable is a structured byte string, the **em\_sbs\_format** field points to a string that contains a description of its format. This description is a comma-separated list in serial-number order of name/value pairs, where the name is the SBS field name and the value is the SBS field type. The name/value pair is in the form *name=value*.

The **em\_init\_value** field points to a string that contains the initial value of the variable. For a structured byte string, the initial value is specified as a list of name/value pairs, separated by commas, in serial number order. The **em\_class** field is a pointer to a string that contains the variable's resource class.

The **em\_rsrc\_ID** field points to a string that contains the variable's resource ID definition. This definition is a semicolon-separated list of name/value pairs, where the value is the data type of the resource ID element. The **em\_rsrc\_ID\_descr** field points to a string that contains the text of each resource ID element's description; the descriptions are separated by the newline character. For information about resource variable definitions, see "Resource Variables (haemrvars)" on page 126.

The **em\_ptx\_name** field points to a string that contains the name that is used to read and write the resource variable in Performance Toolbox (PTX) shared memory. If the resource variable is of value type State, the **em\_ptx\_name** field points to a null string.

The **em\_dflt\_expr** field is a pointer to a string that contains the default expression. If no default expression is defined, the **em\_dflt\_expr** field points to a null string. For information about defining expressions, see “Expressions (haemexpr)” on page 123.

The **em\_event\_descr** field is a pointer to a string that contains the text of the event description for the event specified by the default expression.

The **em\_locator** field is a pointer to a string that contains the Locator value specified for the variable. If no Locator value is defined, the **em\_locator** field points to a null string.

The **em\_order\_group** field is reserved for IBM use.

If, when issuing the query command that resulted in this response, the EM client supplied a pointer to a callback routine, then instead of returning the **ha\_em\_rsp\_blk** structure in a buffer, the **ha\_em\_receive\_response** subroutine invokes the callback routine. The session file descriptor, a pointer to the **ha\_em\_rsp\_blk** structure, and the value of the **em\_qcb\_arg** field, which was specified when the query command was issued, are passed to the callback routine. The memory addressed by the pointer to the **ha\_em\_rsp\_blk** structure or by any of the pointers contained within it must not be freed or modified by the callback routine. The **ha\_em\_receive\_response** subroutine then returns a value of 0.

## Error Responses to Register and Unregister Event Commands

Errors can be returned in response to the registration commands **HA\_EM\_CMD\_REG** and **HA\_EM\_CMD\_REG2** in one of two ways, as follows:

1. When the buffer returned by the **ha\_em\_receive\_response** subroutine contains a response block in which the value of the **em\_cmd** field is either **HA\_EM\_CMD\_REG** or **HA\_EM\_CMD\_REG2**, then the response block contains one or more events that may have an error. The **em\_rpb\_event** array is used in the response block.

If a callback routine has been specified for an event, it is invoked if the event has an error.

If the event has an error, the **em\_error\_number** field of the **em\_error** union is nonzero.

An error returned in this way indicates that the corresponding event has been successfully registered but an error was encountered after registration. The event remains registered.

2. When the buffer returned by the **ha\_em\_receive\_response** subroutine contains a response block in which the value of the **em\_cmd** field is either **HA\_EM\_CMD\_RERR** or **HA\_EM\_CMD\_R2ERR**, then an error may have been encountered while processing the registration command **HA\_EM\_CMD\_REG** or **HA\_EM\_CMD\_REG2**, respectively. This type of response is also received for event registrations that specified the **HA\_EM\_SCMD\_RACK** subcommand regardless of whether an error occurred. The **em\_rpb\_rerr** array is used in the response block.

Callback routines are never invoked for this type of response.

An error returned in this way indicates that the corresponding event has not been registered.

The **ha\_em\_rpb\_rerr** structure has the following definition:

```
struct ha_em_rpb_rerr {
    union ha_em_errnum em_error;
    char               *em_name;
    char               *em_rsrc_ID;
    char               *em_expr;
    char               *em_raexpr;
    short              em_errinfo0;
    unsigned short     em_errinfo1;
    ha_em_eid_t        em_event_id;
};
```

The **ha\_em\_errnum** union, in either the **ha\_em\_rpb\_event** structure or the **ha\_em\_rpb\_rerr** structure, contains information about the error that occurred. It has the following definition:

```
union ha_em_errnum {
    unsigned int    em_error_number;
    unsigned short  em_error_codes[2];
};

#define em_errnum    em_error.em_error_number    /* for quick tests */
#define em_generr    em_error.em_error_codes[0] /* general error code */
#define em_specerr   em_error.em_error_codes[1] /* specific error code */
```

The **em\_error\_number** field indicates whether there is an error. When the field is nonzero, there is an error. When the field is zero, this is an acknowledgement of a successful event registration, as requested by the specification of the **HA\_EM\_SCMD\_RACK** subcommand.

The **em\_error\_codes** array contains further information about the error. The first element of the array is a general error code; the second element is a more specific error code. For information about error codes, see “EMAPI Errors (err\_emapi)” on page 104.

The **em\_name** field is a pointer to a string that contains the name of the resource variable specified in the original registration command. The **em\_rsrc\_ID** field is a pointer to the resource ID that was specified in the original registration command. The **em\_expr** and **em\_raexpr** fields point to strings that contain the expression and rearm expression, respectively, specified in the original registration command. For information about defining expressions, see “Expressions (haemexpr)” on page 123.

The **em\_errinfo0** and **em\_errinfo1** fields may contain additional error information, depending on the specific error.

The **em\_event\_id** field contains the ID of the event related to the error. The value can be used to match this event to the information that was used to register for the event.

Errors can also be returned in response to the unregister event command **HA\_EM\_CMD\_UNREG**. When the buffer returned by the **ha\_em\_receive\_response** subroutine contains a response block in which the value of the **em\_cmd** field is **HA\_EM\_CMD\_UNREG**, then the response block contains one or more unregistration events that may have an error. The setting of **HA\_EM\_EVENT\_UNREG** bit in the **em\_event\_flags** field indicates an unregistration event. The **em\_rpb\_event** array is used in the response block.

If a callback routine has been specified for an event, it is invoked if the event has an error.

If the event has an error, the **em\_error\_number** field of the **em\_error** union is nonzero.

An error returned in this way indicates that the corresponding event has not been unregistered.

## Error Responses to Query Commands

Errors can be returned in response to the query command **HA\_EM\_CMD\_QUERY** in one of two ways, as follows:

1. When the buffer returned by the **ha\_em\_receive\_response** subroutine contains a response block in which the value of the **em\_cmd** field is **HA\_EM\_CMD\_QUERY** and the value of the **em\_subcmd** field is **HA\_EM\_SCMD\_QCUR** or **HA\_EM\_SCMD\_QINST**, then the response block contains one or more responses that may have an error. The **em\_rpb\_qcur** array is used in the response block.

If a callback routine has been specified for the query, it is invoked if the query response has an error.

If the query response has an error, the **em\_error\_number** field of the **em\_error** union is nonzero.

- When the buffer returned by the **ha\_em\_receive\_response** subroutine contains a response block in which the value of the **em\_cmd** field is **HA\_EM\_CMD\_QERR**, then an error was encountered while processing the query command. This can occur while processing any of the query subcommands **HA\_EM\_SCMD\_QCUR**, **HA\_EM\_SCMD\_QINST**, or **HA\_EM\_SCMD\_QDEF**. The **em\_subcmd** field in the response block contains the original subcommand value. The **em\_rpb\_qerr** array is used in the response block.

Callback routines are never invoked for this type of response.

The **ha\_em\_rpb\_qerr** structure has the following definition:

```
struct ha_em_rpb_qerr {
    union ha_em_errnum em_error;
    char *em_class;
    char *em_name;
    char *em_rsrc_ID;
    unsigned short em_errinfo;
};
```

The **ha\_em\_errnum** union in either the **ha\_em\_rpb\_qcur** structure or the **ha\_em\_rpb\_qerr** structure contains information about the error that occurred. It has the following definition:

```
union ha_em_errnum {
    unsigned int em_error_number;
    unsigned short em_error_codes[2];
};

#define em_errnum em_error.em_error_number /* for quick tests */
#define em_generr em_error.em_error_codes[0] /* general error code */
#define em_specerr em_error.em_error_codes[1] /* specific error code */
```

The **em\_error\_number** field indicates whether there is an error. When the field is nonzero, there is an error.

The **em\_error\_codes** array contains further information about the error. The first element of the array is a general error code; the second element is a more specific error code. For information about error codes, see “EMAPI Errors (err\_emapi)” on page 104.

The **em\_class** field is a pointer to a string that contains the name of the resource variable class as specified in the original query command.

The **em\_name** field is a pointer to a string that contains the name of the resource variable as specified in the original query command.

The **em\_rsrc\_ID** field is a pointer to the resource ID of the variable instance as specified in the original query command.

The **em\_errinfo** field may contain additional error information, depending on the specific error.

## Restrictions

When you use the EMAPI in a threaded program, IBM recommends that an Event Management session be used by only one thread at a time. If multiple threads try to use the same Event Management session at the same time, EMAPI subroutines may return the **HA\_EM\_EBUSY** error code. The thread-safe version of the EMAPI library was designed assuming that each thread using the EMAPI would have its own Event Management session or sessions.

## Security

| DCE security may or may not be configured in an SP domain. When DCE security is configured, the EMAPI and the Event Management subsystem authenticate each other while establishing a session. Most

authentication failures are reported during a call to **ha\_em\_receive\_response** rather than during the call to **ha\_em\_start\_session** or **ha\_em\_restart\_session** that establishes the session.

## Return Values

If the **ha\_em\_receive\_response** subroutine returns a value that is greater than 0, it also returns a pointer to a buffer that contains a response block. The *em\_rsp\_blk* argument specified on input points to an area where the buffer pointer is to be returned. It is the responsibility of the calling routine to free the buffer. When the buffer is freed, any memory areas to which there are pointers in the response block can no longer be accessed.

If the **ha\_em\_receive\_response** subroutine returns a value that is equal to 0, the response either was intended for the EMAPI or it contained events or query responses, all of which were processed by the callback routines supplied by the EM client. No further action by the calling routine is necessary.

## Error Values

If the **ha\_em\_receive\_response** subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

If the **HA\_EM\_ECONNLOST** error is received, it is recommended that the EM client no longer use the session file descriptor that was specified on the failing call on subsequent **select** system calls. To reconnect to the Event Management subsystem, use the **ha\_em\_restart\_session** subroutine.

The EMAPI error block and error numbers are defined in the **ha\_emapi.h** header file. For more information on EMAPI errors, see “EMAPI Errors (err\_emapi)” on page 104.

When DCE security is configured in an SP domain, authentication errors may occur. Most authentication errors are reported by **ha\_em\_receive\_response**. The authentication errors that may be received are: **HA\_EM\_EAUTHENT**, **HA\_EM\_EAUTHOR**, **HA\_EM\_ESECSERV**, **HA\_EM\_ENOSECDAEMON**, **HA\_EM\_ENOSECCLIENT**, and **HA\_EM\_EAUTHENTDAE**.

For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

## Examples

For examples of using EMAPI subroutines, see the programs in the RSCT product samples directory, **/usr/sbin/rsct/samples/haem/emapi**.

## Files

**ha\_emapi.h**.

## Prerequisite Information

“Chapter 1. Understanding Event Management” on page 1.

## Related Information

Structures in the **ha\_emapi.h** header file: **ha\_em\_rsp\_blk**, **ha\_em\_err\_blk**, **ha\_em\_rb\_reg**, **ha\_em\_rb\_query**, **ha\_em\_rpb\_event**, **ha\_em\_rpb\_qcur**, **ha\_em\_rpb\_qdef**

Subroutines: **ha\_em\_send\_command**

---

## ha\_em\_restart\_session Subroutine

### Purpose

**ha\_em\_restart\_session** – Restart an EM client session to the Event Management subsystem

### Library

EMAPI Thread-Safe Library (**libha\_em\_r.a**)

EMAPI Library (not thread-safe) (**libha\_em.a**)

### Syntax

```
#define HA_EM_VERSION 2
#include <ha_emapi.h>

int
  ha_em_restart_session(
    int          session_fd,
    struct ha_em_err_blk *em_errb)
```

### Parameters

*session\_fd*

The file descriptor of the session that is to be reconnected. This file descriptor was returned by a previous call to either the **ha\_em\_start\_session** or **ha\_em\_restart\_session** subroutine.

*em\_errb*

A pointer to an error block structure.

### Description

The **ha\_em\_restart\_session** subroutine is used by an EM client to reconnect to an existing session when the connection has been lost.

If an error returned by the **ha\_em\_send\_command** or **ha\_em\_receive\_response** subroutine indicates that the connection to the Event Management subsystem has been lost, this subroutine can be called to create a new connection. If the EM client ended the existing session and established a new session, all of the event registrations for the old session would be lost. The **ha\_em\_restart\_session** subroutine lets the EM client reestablish a connection that has been lost without having to reregister all of the events.

Once a connection to the Event Management subsystem has been lost, it may not be possible to reestablish the connection immediately, because the Event Management subsystem may be recovering from the failure that caused the connection to be lost in the first place. The EM client may wish to make repeated calls to **ha\_em\_restart\_session** at some interval until the session is successfully reestablished. For more details, see the “Error Values” section later in this man page.

### Restrictions

When you use the EMAPI in a threaded program, IBM recommends that an Event Management session be used by only one thread at a time. If multiple threads try to use the same Event Management session at the same time, EMAPI subroutines may return the **HA\_EM\_EBUSY** error code. The thread-safe version of the EMAPI library was designed assuming that each thread using the EMAPI would have its own Event Management session or sessions.

## Security

DCE security may or may not be configured in an SP domain. When DCE security is configured, the EMAPI and the Event Management subsystem authenticate each other while establishing a session. Most authentication failures are reported during a call to **ha\_em\_receive\_response** rather than during the call to **ha\_em\_restart\_session**.

## Return Values

If the return value of the **ha\_em\_restart\_session** subroutine is greater than or equal to 0, the return value is a file descriptor. The **ha\_em\_restart\_session** subroutine closes the file descriptor that was provided on input in the *session\_fd* argument and returns a new one. The newly returned file descriptor replaces the file descriptor that was provided on input, and the EM client can use the new file descriptor in the same way as it used the file descriptor that was originally returned by the **ha\_em\_start\_session** subroutine.

## Error Values

If the **ha\_em\_restart\_session** subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

The EMAPI error block and error numbers are defined in the **ha\_emapi.h** header file. For more information on EMAPI errors, see “EMAPI Errors (err\_emapi)” on page 104.

For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

The file descriptor that was provided on input in the *session\_fd* argument is still valid and can be used in another attempt to restart the session.

If the **HA\_EM\_ECONNREFUSED** error is received, the Event Management subsystem may be recovering from the failure that caused the connection to be lost in the first place. The EM client may wish to make repeated calls to **ha\_em\_restart\_session** at some interval while it is receiving the **HA\_EM\_ECONNREFUSED** error until the session is successfully reestablished. The interval selected depends on the needs of the client application, but it is probably not useful for the interval to be less than 5 seconds long.

The EM client may try to reestablish the connection indefinitely, or it may give up after a number of failed attempts. How persistent the EM client is in attempting to reestablish the connection depends on the needs of the client application. The Event Management subsystem persistently tries to recover from failure conditions so that client connections can be established.

When DCE security is configured in an SP domain, authentication errors may occur. **HA\_EM\_ESECSERV** is the only authentication error that is reported by **ha\_em\_restart\_session**. Other authentication errors may be reported later during a call to **ha\_em\_receive\_response**.

## Examples

For examples of using EMAPI subroutines, see the programs in the RSCT product samples directory, **/usr/sbin/rsct/samples/haem/emapi**.

## Files

**ha\_emapi.h**.

## Prerequisite Information

“Chapter 1. Understanding Event Management” on page 1.

## Related Information

Structures in the `ha_emapi.h` header file: `ha_em_err_blk`

Subroutines: `ha_em_start_session`, `ha_em_end_session`

---

## ha\_em\_send\_command Subroutine

### Purpose

**ha\_em\_send\_command** – Send a command to the Event Management subsystem

### Library

EMAPI Thread-Safe Library (**libha\_em\_r.a**)

EMAPI Library (not thread-safe) (**libha\_em.a**)

### Syntax

```
#define HA_EM_VERSION 2
#include <ha_emapi.h>

int
  ha_em_send_command(
    int          session_fd,
    struct ha_em_cmd_blk *em_cmdb,
    struct ha_em_err_blk *em_errb)
```

### Parameters

*session\_fd*

The file descriptor of the Event Management session to which the command is to be sent. This file descriptor was returned by a previous call to either the **ha\_em\_start\_session** or **ha\_em\_restart\_session** subroutine.

*em\_cmdb*

A pointer to a command block structure. The command block contains the command to be sent to the Event Management subsystem in the specified session and any additional parameters that are needed to execute the command.

*em\_errb*

A pointer to an error block structure.

### Description

The **ha\_em\_send\_command** subroutine is used by the EM client to send a command to the instance of the Event Management subsystem that is associated with the session supplied on input. The command and any additional parameters that it requires are contained in the command block specified on input.

The **ha\_em\_cmd\_blk** command block has the following definition:

```
struct ha_em_cmd_blk {
  int          em_cmd_num_elem;
  short        em_cmd;
  short        em_subcmd;
  ha_em_qid_t  em_qid;
  void         (*em_qcb)(int, struct ha_em_rsp_blk *, void *);
  void         *em_qcb_arg;
  union ha_em_res_blk {
    struct ha_em_rb_reg      em_rb_reg[1];
    ha_em_eid_t              em_rb_unreg[1];
    struct ha_em_rb_query    em_rb_query[1];
  } em_res_blk;
}
```

In an actual command block, the arrays in the **ha\_em\_res\_blk** union may contain any number of elements; the **em\_cmd\_num\_elem** field contains the actual number that are allocated. It is the responsibility of the caller to allocate a buffer to contain the command block with the desired number of array elements.

The **em\_cmd** field must be set to one of the following values:

## HA\_EM\_CMD\_REG

Register for an event from a single expression, the event expression. If a rearm expression is defined, it does not generate an event.

## HA\_EM\_CMD\_REG2

Register for an event from both an event expression and a rearm expression.

## HA\_EM\_CMD\_UNREG

Unregister events

## HA\_EM\_CMD\_QUERY

Query for information

For the **HA\_EM\_CMD\_REG** or **HA\_EM\_CMD\_REG2** command, the **em\_subcmd** field may be set to 0, if no subcommand is required, or to one or more of the following values bitwise ORed together:

### HA\_EM\_SCMD\_RACK

Return a registration acknowledgement.

### HA\_EM\_SCMD\_REVAL

Generate an event at the first observation, even if the expression evaluates to false.

For the **HA\_EM\_CMD\_QUERY** command, the **em\_subcmd** field must be set to one of the following values:

### HA\_EM\_SCMD\_QINST

Query the current instances of resource variables.

### HA\_EM\_SCMD\_QCUR

Query the current values of resource variable instances.

### HA\_EM\_SCMD\_QDEF

Query defined resource variables and their default expressions

The remaining fields in the command block are used as specified in the command descriptions that follow.

## HA\_EM\_CMD\_REG and HA\_EM\_CMD\_REG2 Commands (Register for Events)

An EM client uses the command value **HA\_EM\_CMD\_REG** or **HA\_EM\_CMD\_REG2** to register for events. The **em\_rb\_reg** array is used in the command block. One array element is defined for each event to be registered.

The **ha\_em\_rb\_reg** structure has the following definition:

```
struct ha_em_rb_reg {
    char      *em_name;
    char      *em_rsrc_ID;
    char      *em_expr;
    char      *em_raexpr;
    ha_em_eid_t  em_event_id;
    void      (*em_cb)(int, struct ha_em_rpb_event *, void *);
    void      *em_cb_arg;
}
```

The **em\_name** field is a pointer to a string that contains the name of the resource variable from which the event is to be generated.

The **em\_rsrc\_ID** field is a pointer to a string that contains a resource ID. The event is registered for all instances of the specified resource variable that match the resource ID.

The resource ID is a semicolon-separated list of name/value pairs. A name/value pair consists of a resource ID element name followed by an equal sign followed by the value of the element. There are no blanks in the resource ID. A resource ID element is wildcarded by specifying its value as an asterisk. The resource ID string must contain each element that is defined for the variable's resource ID. An element

value may consist of a single value, a range of values, a comma-separated list of single values, or a comma-separated list of ranges. A range takes the form *a-b* and is valid only for resource ID elements of type integer.

The resource ID must contain each element that is defined for the resource variable. The resource ID may be wildcarded to match more than one instance by specifying an asterisk for the value of one or more elements, as shown in the following examples:

```
NodeNum=5;VG=rootvg;LV=hd4
NodeNum=*;VG=rootvg;LV=hd4
NodeNum=*;VG=*;LV=*
```

In the first example, the Event Management subsystem registers an event for one instance, associated with node 5, volume group **rootvg**, and logical volume **hd4**. In the second example, the Event Management subsystem registers an event for all instances that are associated with all nodes, volume group **rootvg**, and logical volume **hd4**. In the last example, the Event Management subsystem registers an event for all instances that are associated with all nodes, all volume groups, and all logical volumes.

The **em\_expr** field is a pointer to a string that contains the expression that is to be used to generate the event. If the expression string is null, the default expression, which is defined for the specified variable in the Event Management Configuration Database (EMCDB), is to be used to generate the event. For information about defining expressions, see “Expressions (haemexpr)” on page 123.

The **em\_raexpr** field is a pointer to a string that contains a rearm expression, which has the same syntax as the expression pointed to by the **em\_expr** field. If the **em\_raexpr** field points to a null string, a rearm expression is not specified.

If a rearm expression is specified, whenever a “true” event is generated by applying the event expression to the observed variable instance, the rearm expression is applied to the instance of the variable as it is observed from then on. When the rearm expression is true, the original event expression is used again. In other words, whenever either expression is true, the other expression is used for that instance until it is true.

When the **HA\_EM\_CMD\_REG** command is specified, an event is generated only from the event expression, not from the rearm expression. However, when the **HA\_EM\_CMD\_REG2** command is specified, an event is also generated when the rearm expression is TRUE. The event response indicates which expression was used to generate the event.

Rearm expressions can be used in a couple of different ways. One common way is to define the rearm expression as the inverse of the event expression. For example, if the event expression tests whether a resource variable value is “on”, the rearm expression tests whether it is “off”.

A rearm expression can also be used with the event expression to define an upper and lower boundary for a condition of interest. For example, suppose an EM client is interested in how much disk space is used in the system. It defines a normal disk space condition to be in the range between 80% and 90% of the total space. However, if the amount of disk space used rises above 90% or falls below 80%, the EM client wants to be notified.

In this case, it would define the event expression to be “the amount of disk space used is greater than 90%” and the rearm expression to be “the amount of disk space used is less than 80%”. The EM client would also register for events from both expressions.

The first time disk space usage rises above 90%, the Event Management subsystem notifies the EM client and switches to the rearm expression. The EM client is not notified again until disk space usage decreases to an amount less than 80%, and then the Event Management subsystem switches to the event expression.

The expression must contain the resource variable name, represented by the letter “X”, in any one or more of its modified or unmodified forms. Events generated by the specified expression(s) from all matching instances of the specified resource variable have the same event ID. These events are differentiated by inclusion of the resource ID value with the event.

If the **HA\_EM\_SCMD\_REVAL** subcommand is specified with either the **HA\_EM\_CMD\_REG** or **HA\_EM\_CMD\_REG2** command, an event is generated at the first observation of an instance of the variable after registration, even if the expression is false. The event response indicates whether the expression was true or false. This allows the EM client to obtain an initial state of the resource variable. Note, however, that this event does not indicate the time the resource variable assumed the returned value (it only indicates the time the event response was generated). When a subsequent event is received it is known that the resource variable assumed the subsequent returned value some time between the event generated as a result of the **HA\_EM\_SCMD\_REVAL** subcommand and the subsequent event.

When the **ha\_em\_send\_command** subroutine returns successfully, the **em\_event\_id** field has been set by the EMAPI with an event ID. The event ID is returned with each generated event for reference back to the registration information. At this point, the EM client can now expect to receive events from the registered resource variables.

All events that are registered by a single call to the **ha\_em\_send\_command** subroutine can be identified by an event command group ID. This ID can be used to relate individual event responses back to the command with which they were registered. The event command group ID can be obtained by calling the **ha\_em\_get\_ecgid** subroutine.

If the caller specifies a function pointer in the **em\_cb** field, the specified function is used as a callback routine by the EMAPI. Whenever an event that is associated with this registration is received by the EMAPI, the callback routine is called.

The first argument to the callback routine is the file descriptor of the session in which the event was registered. The second argument is a pointer to a structure that contains the event information. The third argument is the value of the **em\_cb\_arg** field. The **em\_cb\_arg** field specifies an argument, meaningful in the context of the specified event only to the EM client, that can be passed to the callback routine. The argument must be valid when the callback routine is invoked. For example, a pointer to a stack must be valid when the callback routine is invoked, or an argument must be valid in the context of the thread in which the callback routine executes. For more information about callback routines, see the man page for the **ha\_em\_receive\_response** subroutine.

If the **HA\_EM\_SCMD\_RACK** subcommand is specified with either the **HA\_EM\_CMD\_REG** or **HA\_EM\_CMD\_REG2** command, a registration acknowledgement is returned once the Event Management subsystem has validated the registration request. The acknowledgement is returned in a response specifying the **HA\_EM\_CMD\_RERR** or **HA\_EM\_CMD\_RERR2** command. Registration errors are always reported through these responses. The specification of **HA\_EM\_SCMD\_RACK** on the registration request causes successful registrations to be reported through these responses too. See “Error Responses to Register and Unregister Event Commands” in the **ha\_em\_receive\_response()** man page for further detail.

### **HA\_EM\_CMD\_UNREG Command (Unregister Events)**

An EM client uses the command value **HA\_EM\_CMD\_UNREG** to stop receiving events for which it registered. The **em\_rb\_unreg** array is used in the command block. One array element is defined for each event that is no longer to be sent to the EM client. The array element contains an event ID that was returned in the **ha\_em\_rb\_reg** structure when the EM client registered for events.

Because the command to unregister events is sent asynchronously to the schedule on which events are generated by the Event Management subsystem, the EM client may still receive the events for which it just unregistered until the Event Management subsystem can process the command. Once the command has

been processed, the Event Management subsystem generates a response for each unregistered event. This response has a format similar to that of the event that has just been unregistered.

If the EM client specified a callback routine for any of the unregistered events, then, when the unregistered event response is received by the EM API, the callback routine is invoked with the same arguments as if the response were the event. Thus, the callback routine sees some number of events followed by an indication that no more events will be forthcoming.

### **HA\_EM\_CMD\_QUERY Command (Query for Information)**

An EM client uses the **HA\_EM\_CMD\_QUERY** command to query for information. It is used with the following subcommands:

#### **HA\_EM\_SCMD\_QINST**

Obtain a list of instances of resource variables.

#### **HA\_EM\_SCMD\_QCUR**

Obtain the current value of one or more resource variable instances.

#### **HA\_EM\_SCMD\_QDEF**

Obtain a list of resource variables and default expressions that are defined in the Event Management Configuration Database (EMCDB).

Each query subcommand uses the **em\_rb\_query** array in the command block. One array element is defined for each query; a single query can specify multiple resource variables.

The **ha\_em\_rb\_query** structure has the following definition:

```
struct ha_em_rb_query {
    char    *em_class;
    char    *em_name;
    char    *em_rsrc_ID;
}
```

For all the query subcommands, the target of the operation is one or more resource variables. The **em\_class**, **em\_name**, and **em\_rsrc\_ID** fields are used to define the resource variables to be queried. Resource variables may be specified individually in separate **ha\_em\_rb\_query** structures or one **ha\_em\_rb\_query** structure may be used to specify many resource variables through the use of wildcards. Any number of **ha\_em\_rb\_query** structures may be specified in either manner in one subcommand.

If the **em\_class** field is a pointer to a string that contains the name of a resource variable class, then all variables in that class, as further limited by the **em\_name** and **em\_rsrc\_ID** fields, are the targets of the query. If the **em\_class** field points to a null string, then variables of all classes, as further limited by the **em\_name** and **em\_rsrc\_ID** fields, are the targets of the query.

The **em\_name** field points to a string that contains the resource variable name that is the target of the query. The resource variable name may be wildcarded in one of two ways:

- Specifying the name as a null string.
- Truncating the name after any component.

When the resource variable name is wildcarded in the first manner, then all resource variables, as further limited by the **em\_class** and **em\_rsrc\_ID** fields, are targets of the query. When the resource variable name is wildcarded in the second manner, all resource variables whose high-order (leftmost) components match the specified resource variable name, as further limited by the **em\_class** and **em\_rsrc\_ID** fields, are the targets of the query.

The **em\_rsrc\_ID** field is a pointer to a string that contains a resource ID. All instances, or definitions in the case of **HA\_EM\_SCMD\_QDEF**, of the resource variable(s) specified by the **em\_class** and **em\_name** fields that match the resource ID are the targets of the query.

If the query is **HA\_EM\_SCMD\_QCUR** or **HA\_EM\_SCMD\_QINST**, the resource ID is a semicolon-separated list of name/value pairs. A name/value pair consists of a resource ID element name followed by an equal sign followed by the value of the element. An element value may consist of a single value, a range of values, a comma-separated list of single values, or a comma-separated list of ranges. A range takes the form *a-b* and is valid only for resource ID elements of type integer. There are no blanks in the resource ID.

If the query is **HA\_EM\_SCMD\_QDEF**, the resource ID is a semicolon-separated list of element names only.

A resource ID element is wildcarded by specifying its value as the asterisk character. Only variables that are defined to contain the elements, and only the elements, specified in the resource ID are candidates for a match of the query. If any element of the resource ID consists of the asterisk character, rather than a name/value pair (or name, in the case of the **HA\_EM\_SCMD\_QDEF** subcommand), all variables that are defined to contain *at least* the remaining specified elements are candidates for a match of the query. The entire resource ID is wildcarded if the **em\_rsrc\_ID** field points to a string that contains a single asterisk.

Here are some examples of using wildcards for resource IDs specified on a query:

```
NodeNum=5;VG=rootvg;LV=hd4
NodeNum=*;VG=rootvg;LV=hd4
NodeNum=*;VG=*;LV=*

NodeNum=9
NodeNum=*

NodeNum=9;VG=*;*
NodeNum=*;*
```

For these examples, assume that the class and resource variable names are specified as null strings. If either the class or resource variable name or both are specified, matches for the query are restricted accordingly.

In the first three examples, all variables whose resource IDs are defined to contain the elements **NodeNum**, **VG**, and **LV**, and only those elements, are matched. The instances matched are the same for the event registration examples in the section on registering for events earlier in this man page.

In the fourth example, all variables whose resource IDs are defined to contain only the element **NodeNum** are matched. The instances matched are associated with node 9. In the fifth example, the same set of variables are matched, but all instances of each variable are matched.

In the sixth example, all variables whose resource IDs are defined to contain the elements **NodeNum** and **VG**, as well as zero or more additional elements, are matched. The instances matched are associated with node 9. In the last example, all variables whose resource IDs are defined to contain the element **NodeNum**, as well as zero or more additional elements, are matched. All instances of the variables are matched.

Given the flexibility in specifying resource variables for a query, it is possible that no resource variable instance or resource variable definition will match the specification in the **ha\_em\_rb\_query** structure. If there is no match, an appropriate error response is returned.

If the caller specifies a function pointer in the **em\_qcb** field, the specified function is used as a callback routine by the EMAPI. When the response or responses for this query are received by the EMAPI, the callback routine is called.

The first argument to the callback routine is the file descriptor of the session in which the query was made. The second argument is a pointer to a structure that contains the query information. The third argument is the value of the **em\_qcb\_arg** field. The **em\_qcb\_arg** field specifies an argument, meaningful in the context of the query response only to the EM client, that can be passed to the callback routine. The argument must be valid when the callback routine is invoked. For example, a pointer to a stack must be

valid when the callback routine is invoked, or an argument must be valid in the context of the thread in which the callback routine executes. For more information about callback routines, see the man page for the **ha\_em\_receive\_response** subroutine.

If the **ha\_em\_send\_command** subroutine returns successfully, the **em\_qid** field in the command block contains a value unique to each query; the EM client can now expect one or more responses that contain the requested information. Each response contains the value returned in **em\_qid**. Because responses are returned asynchronously, this value is used to match responses to the query command.

## Restrictions

When you use the EMAPI in a threaded program, IBM recommends that an Event Management session be used by only one thread at a time. If multiple threads try to use the same Event Management session at the same time, EMAPI subroutines may return the **HA\_EM\_EBUSY** error code. The thread-safe version of the EMAPI library was designed assuming that each thread using the EMAPI would have its own Event Management session or sessions.

## Return Values

If the **ha\_em\_send\_command** subroutine is successful, it returns a value of 0. Other information that may be returned depends on the command that is sent. For more information, see the command's description.

## Error Values

If the **ha\_em\_send\_command** subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

If the **HA\_EM\_ECONNLOST** error is received, it is recommended that the EM client no longer use the session file descriptor that was specified on the failing call on subsequent **select** system calls. To reconnect to the Event Management subsystem, use the **ha\_em\_restart\_session** subroutine.

The EMAPI error block and error numbers are defined in the **ha\_emapi.h** header file. For more information on EMAPI errors, see "EMAPI Errors (err\_emapi)" on page 104.

For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

## Examples

For examples of using EMAPI subroutines, see the programs in the RSCT product samples directory, **/usr/sbin/rsct/samples/haem/emapi**.

## Files

**ha\_emapi.h**.

## Prerequisite Information

"Chapter 1. Understanding Event Management" on page 1.

## Related Information

Structures in the **ha\_emapi.h** header file: **ha\_em\_cmd\_blk**, **ha\_em\_err\_blk**, **ha\_em\_rb\_reg**, **ha\_em\_rb\_query**

Subroutines: **ha\_em\_start\_session**, **ha\_em\_restart\_session**, **ha\_em\_receive\_response**

---

## ha\_em\_start\_session Subroutine

### Purpose

**ha\_em\_start\_session** – Start an EM client session with the Event Management subsystem

### Library

EMAPI Thread-Safe Library (**libha\_em\_r.a**)

EMAPI Library (not thread-safe) (**libha\_em.a**)

### Syntax

```
#define HA_EM_VERSION 2
#include <ha_emapi.h>

int
  ha_em_start_session(
    int          em_domain_type,
    char         *em_domain_name,
    ha_em_err_blk *em_errb)
```

### Parameters

*em\_domain\_type*

Indicates the type of the domain in which a session is to be established. Specify **HA\_EM\_DOMAIN\_SP** for an SP domain, or **HA\_EM\_DOMAIN\_HACMP** for an HACMP/ES domain.

*em\_domain\_name*

A pointer to a string that indicates the domain in which a session is to be established.

If the *em\_domain\_name* argument points to a null string, the session is established with the Event Management subsystem in the default domain.

For the **HA\_EM\_DOMAIN\_SP** type, the default domain is the current system partition. If the EM client is executing on the control workstation or on a workstation outside of the SP, the current system partition is determined from the value of the **SP\_NAME** environment variable.

For the **HA\_EM\_DOMAIN\_HACMP** domain type, the default domain is the HACMP/ES domain to which the workstation or SP node on which the EM client is executing belongs.

If the *em\_domain\_name* argument points to a string that is not null, the string is the name of the domain in which a session is to be established with the Event Management subsystem. For the **HA\_EM\_DOMAIN\_SP** domain type, the name of any system partition in the SP may be specified. For the **HA\_EM\_DOMAIN\_HACMP** domain type, the only domain that may be specified is the HACMP/ES domain to which the workstation or SP node on which the EM client is executing belongs.

*em\_errb*

A pointer to an error block structure.

### Description

The **ha\_em\_start\_session** subroutine is used by an EM client to establish a session with the Event Management subsystem. The session validates that the EM client is permitted to use Event Management services and then provides a communication path to the specified Event Management subsystem. If this communication path is lost, it can be restored without starting another session by using the **ha\_em\_restart\_session** subroutine. When the EM client no longer needs Event Management services, it ends the session by using the **ha\_em\_end\_session** subroutine.

When the **HA\_EM\_DOMAIN\_SP** domain type is specified, a session is established with the Event Management subsystem in the current system partition by default. To receive events from, and to issue queries to, other system partitions, the **ha\_em\_start\_session** subroutine must be called once for each other system partition.

Multiple sessions may be started with any single domain. This provides multiple communication paths to the Event Management subsystem in a single domain. Multiple paths are useful if the EM client wishes to ignore certain responses for a time while accepting others. For example, an EM client can register events in one session and issue queries in another.

## Restrictions

The EMAPI uses connection-oriented socket communication. An EM client may want to change the handling of the **SIGPIPE** signal so that it is ignored. If the EM client does not ignore the **SIGPIPE** signal, and the connection between the EM client and the Event Management subsystem is lost, delivery of a **SIGPIPE** signal may kill the EM client.

When you use the EMAPI in a threaded program, IBM recommends that an Event Management session be used by only one thread at a time. If multiple threads try to use the same Event Management session at the same time, EMAPI subroutines may return the **HA\_EM\_EBUSY** error code. The thread-safe version of the EMAPI library was designed assuming that each thread using the EMAPI would have its own Event Management session or sessions.

The syntax and semantics described here occur only if you define **HA\_EM\_VERSION** to the value 2 before including the **ha\_emapi.h** header file. If **HA\_EM\_VERSION** is not defined to the value 2, you may see the syntax and semantics of a prior version of the EMAPI.

## Security

DCE security may or may not be configured in an SP domain. When DCE security is configured, the EMAPI and the Event Management subsystem authenticate each other while establishing a session. Most authentication failures are reported during a call to **ha\_em\_receive\_response** rather than during the call to **ha\_em\_start\_session**.

## Return Values

If the return value of the **ha\_em\_start\_session** subroutine is greater than or equal to 0, the return value is a file descriptor. The EM client uses the file descriptor to determine when the Event Management subsystem has sent a response. The file descriptor is used as an argument in the **select** or **poll** system call. This descriptor is also used as a session handle in all other EMAPI subroutine calls.

## Error Values

If the **ha\_em\_start\_session** subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

The EMAPI error block and error numbers are defined in the **ha\_emapi.h** header file. For more information on EMAPI errors, see “EMAPI Errors (err\_emapi)” on page 104.

For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

If the **HA\_EM\_ECONNREFUSED** error is received, the Event Management subsystem may not be running or it may be initializing itself. The EM client may wish to exit or retry establishing a session with the Event Management subsystem. To retry, it would make repeated calls to **ha\_em\_start\_session** at some interval

while it is receiving the **HA\_EM\_ECONNREFUSED** error until the session is successfully established. The interval selected depends on the needs of the client application, but it is probably not useful for the interval to be less than 5 seconds long.

The EM client may try to establish the connection indefinitely, or it may give up after a number of failed attempts. How persistent the EM client is in attempting to establish the connection depends on the needs of the client application.

| When DCE security is configured in an SP domain, authentication errors may occur. **HA\_EM\_ESECSERV**  
| is the only authentication error that is reported by **ha\_em\_start\_session**. Other authentication errors may  
| be reported later during a call to **ha\_em\_receive\_response**.

## Examples

For examples of using EMAPI subroutines, see the programs in the RSCT product samples directory, **/usr/sbin/rsct/samples/haem/emapi**.

## Files

**ha\_emapi.h**.

## Prerequisite Information

“Chapter 1. Understanding Event Management” on page 1.

## Related Information

Structures in the **ha\_emapi.h** header file: **ha\_em\_err\_blk**

Subroutines: **ha\_em\_restart\_session**, **ha\_em\_end\_session**

---

## ha\_rr\_add\_var Subroutine

### Purpose

**ha\_rr\_add\_var** – Add and supply current values of registered resource variables to a resource monitor manager session

### Library

RMAPI Library (not thread-safe) (**libha\_rr.a**)

### Syntax

```
#include <ha_rmap.h>

int
  ha_rr_add_var(
    int          session_fd,
    struct ha_rr_variable *pv,
    int          numv,
    int          add_complete,
    struct ha_em_err_blk *rr_errb)
```

### Parameters

#### *session\_fd*

The file descriptor of the session from which the **HA\_RR\_CMD\_ADDV** or **HA\_RR\_CMD\_ADDALL** command was received or, if the resource monitor is command-based, the file descriptor that was returned by the **ha\_rr\_start\_session** subroutine.

*pv* A pointer to an array of **ha\_rr\_variable** structures, which contain the names, resource IDs, and values of the resource variables to be added.

*numv* The number of elements in the array pointed to by the *pv* parameter.

#### *add\_complete*

A flag that indicates when the resource monitor has completed adding variables in response to receiving an **HA\_RR\_CMD\_ADDALL** command. A value of 0 indicates that there are more variables to be added. A nonzero value indicates that, with this call, the adding of variables is complete.

More than one call of the **ha\_rr\_add\_var** subroutine may be used to add variables. When multiple calls are used, the *add\_complete* parameter is set to 0 on each call but the last.

For the **HA\_RR\_CMD\_ADDV** command, this argument is always 0.

*rr\_errb* A pointer to an error block structure.

### Description

The **ha\_rr\_add\_var** subroutine is used by the resource monitor to add resource variables, which have been previously registered with the RMAPI, to the resource monitor manager session specified by the **session\_fd** parameter. The monitor must start supplying the values of variables that have been added to one or more sessions using the **ha\_rr\_send\_val** subroutine.

This command is called by server type resource monitors in response to receiving an **HA\_RR\_CMD\_ADDALL** or **HA\_RR\_CMD\_ADDV** command from a resource monitor manager.

A command-based resource monitor calls the **ha\_rr\_add\_var** subroutine immediately after it registers its variable instances to add those instances to the list of those known to the RMAPI and supply current values for resource variable instances.

The **ha\_rr\_variable** structure has the following definition:

```

struct ha_rr_variable {
    char    *rr_var_name;
    char    *rr_var_rsrc_ID;
    union {
        int    rr_var_inst_id;
        void   **rr_var_hndl;
    } rr_varu;
#define rr_var_handle    rr_varu.rr_var_hndl
#define rr_var_iid      rr_varu.rr_var_inst_id
    void   *rr_value;
    int    rr_var_errno;
}

```

The **rr\_var\_name** field is a pointer to a string that contains the name of a resource variable.

The **rr\_var\_rsrc\_ID** field is a pointer to a string that contains the resource ID of an instance of the specified variable. The resource ID must be fully qualified; no wildcarding is permitted.

The resource ID is specified as a comma-separated list of name/value pairs. A name/value pair consists of a resource ID element name followed by an equal sign followed by the value of the resource ID element. There are no blanks in the resource ID.

The **rr\_var\_hndl** field is a pointer to a buffer that will be used to store the handle of the specified resource variable. On input, it must point to a valid location.

The **rr\_value** field is a pointer to a buffer that contains the current value of the variable.

The **rr\_var\_errno** field is used to hold the result of the operation for each variable.

## Security

The calling process must have a real or effective group ID of **haemrm**, or must have the **haemrm** group ID in its supplemental group list. If the process is not in the **haemrm** group, its effective user ID must be **root**.

If the calling process is instance 0 of the resource monitor, and the monitor is configured to supply Counter or Quantity variables to the Performance Monitor, it must have an effective user ID of **root**.

## Restrictions

A resource monitor of connection type server must call the **ha\_rr\_add\_var** subroutine only in response to receiving the **HA\_RR\_CMD\_ADDV** or **HA\_RR\_CMD\_ADDALL** command. Calling this subroutine at other times is a programming error for a server type of resource monitor, with undefined results.

If the name of one of the elements in any resource ID matches the resource ID element name specified by the Locator field in the associated resource variable definition, that resource ID name/value pair is never supplied by the resource monitor in the resource ID specified by the **rr\_var\_rsrc\_ID** field. This resource ID element is added by the Event Management subsystem when the instance is given in a response to any Event Management client.

## Return Values

If the **ha\_rr\_add\_var** subroutine is successful, it returns the number of resource variables that were added.

For each variable in the array that was successfully added, a variable handle is placed by the RMAPI in the location that is pointed to by the associated **rr\_var\_hndl** field in the **ha\_rr\_variable** structure. This variable handle is used when sending the resource variable to the RMAPI. If a variable has already been added, its variable handle is returned unchanged. The **ha\_rr\_add\_var** subroutine can be called multiple times for the same resource variable.

For each variable in the array that was successfully added, the associated `rr_var_errno` field is also set to 0.

If a variable could not be added, an error code is placed in its `rr_var_errno` field.

For all variables that are successfully added, the resource monitor should immediately start sending values at the intervals determined by the resource variable value type and the design of the resource monitor.

## Error Values

If the `ha_rr_add_var` subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

The following error number requires specific action by the resource monitor:

### HA\_RR\_EDISCONNECT

The resource monitor manager associated with the `session_fd` file descriptor has dropped its connection.

A server type resource monitor should call `ha_rr_del_var` to delete all variables, using the `session_fd` file descriptor, that have been added by the resource monitor. The monitor should then end the session, by calling `ha_rr_end_session`.

A client type resource monitor should call `ha_rr_terminate` to end the current use of the RMAPI.

The RMAPI error block and error numbers are defined in the `ha_rmapi.h` header file. For more information on RMAPI errors, see “RMAPI Errors (err\_rmapi)” on page 131.

For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

## Examples

For examples of using RMAPI subroutines, see the programs in the RSCT product samples directory, `/usr/sbin/rsct/samples/haem/rmapi`.

## Files

`ha_rmapi.h`

## Prerequisite Information

“Chapter 1. Understanding Event Management” on page 1.

## Related Information

Structures in the `ha_rmapi.h` header file: `ha_rr_variable`, `ha_em_err_blk`

Subroutines: `ha_rr_del_var`, `ha_rr_get_ctrlmsg`

---

## ha\_rr\_del\_var Subroutine

### Purpose

**ha\_rr\_del\_var** – Delete resource variable instances from a resource monitor manager session

### Library

RMAPI Library (not thread-safe) (**libha\_rr.a**)

### Syntax

```
#include <ha_rmap.h>

int
  ha_rr_del_var(
    int          session_fd,
    struct ha_rr_variable *pv,
    int          numv,
    struct ha_em_err_blk *rr_errb)
```

### Parameters

*session\_fd*

For server type monitors, the file descriptor of a resource manager session that has ended, or from which the **HA\_RR\_CMD\_DELV** or **HA\_RR\_CMD\_DELALL** command was received. If the resource monitor is command-based, the file descriptor that was returned by the **ha\_rr\_start\_session** subroutine.

*pv* A pointer to an array of **ha\_rr\_variable** structures, which contain the handles of the resource variables to be deleted.

*numv* The number of elements in the array pointed to by the *pv* parameter.

*rr\_errb* A pointer to an error block structure.

### Description

The **ha\_rr\_del\_var** subroutine is used by the resource monitor to delete variables from a resource monitor manager session.

This command is called by server type resource monitors whenever the resource monitor receives the **HA\_RR\_CMD\_DELALL** or **HA\_RR\_CMD\_DELV** command.

If a server type resource monitor detects that a resource monitor manager session has ended, and the monitor will continue executing, the **ha\_rr\_del\_var** subroutine should be called prior to calling **ha\_rr\_end\_session**. All variables that had been previously added by any call to **ha\_rr\_add\_var** should be deleted from the ending session.

A command-based resource monitor calls the **ha\_rr\_del\_var** subroutine when it has stopped sending variable values to the RMAPI.

The **ha\_rr\_variable** structure has the following definition:

```
struct ha_rr_variable {
  char      *rr_var_name;
  char      *rr_var_rsrc_ID;
  union {
    int      rr_var_inst_id;
    void     **rr_var_hdl;
  } rr_varu;
#define rr_var_handle rr_varu.rr_var_hdl
#define rr_var_iid    rr_varu.rr_var_inst_id
  void     *rr_value;
  int      rr_var_errno;
}
```

The `rr_var_name` and `rr_var_rsrc_ID` fields are not used.

The `rr_var_hndl` field points to the variable handle of each variable that is to be deleted. The variable handle was previously assigned by the RMAPI when the variable was added by the `ha_rr_add_var` subroutine.

The `rr_value` field is not used.

The `rr_var_errno` field is used to hold the result of the operation for each variable.

## Security

The calling process must have a real or effective group ID of `haemrm`, or must have the `haemrm` group ID in its supplemental group list. If the process is not in the `haemrm` group, its effective user ID must be `root`.

If the calling process is instance 0 of the resource monitor, and the monitor is configured to supply Counter or Quantity variables to the Performances Monitor, it must have an effective user ID of `root`.

## Restrictions

A resource monitor of connection type server must call the `ha_rr_del_var` subroutine only in response to receiving the `HA_RR_CMD_DELV` or `HA_RR_CMD_DELALL` command, or before calling the `ha_rr_end_session` subroutine. Calling this subroutine at other times is a programming error for a server type of resource monitor, with undefined results.

## Return Values

If the `ha_rr_del_var` subroutine is successful, it returns the number of variables that must no longer be sent to the RMAPI. If a variable is no longer to be sent, its variable handle is cleared by setting the location specified by the `rr_var_hndl` field to `NULL`. If a variable should continue to be sent, its variable handle is unchanged.

For each variable in the array that was successfully processed, the associated `rr_var_errno` field is set to 0. If a variable could not be processed, an error code is placed in its `rr_var_errno` field.

A return value of 0 indicates that values of all of the variables specified in the `pv` parameter should continue to be sent by calling `ha_rr_send_val`.

## Error Values

If the `ha_rr_del_var` subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

The following error number requires specific action by the resource monitor:

### **HA\_RR\_EDISCONNECT**

The resource monitor manager associated with the `session_fd` file descriptor has dropped its connection.

A server type resource monitor should call `ha_rr_del_var` to delete all variables, using the `session_fd` file descriptor, that have been added by the resource monitor. The monitor should then end the session, by calling `ha_rr_end_session`.

A client type resource monitor should call `ha_rr_terminate` to end the current use of the RMAPI.

The RMAPI error block and error numbers are defined in the `ha_rmapi.h` header file. For more information on RMAPI errors, see "RMAPI Errors (`err_rmapi`)" on page 131.

For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

## Examples

For examples of using RMAPI subroutines, see the programs in the RSCT product samples directory, `/usr/sbin/rsct/samples/haem/rmapi`.

`ha_rmapi.h`

## Prerequisite Information

“Chapter 1. Understanding Event Management” on page 1.

## Related Information

Structures in the `ha_rmapi.h` header file: `ha_rr_variable`, `ha_em_err_blk`

Subroutines: `ha_rr_add_var`, `ha_rr_get_ctrlmsg`, `ha_rr_end_session`

---

## ha\_rr\_end\_session Subroutine

### Purpose

**ha\_rr\_end\_session** – End a resource monitor manager session

### Library

RMAPI Library (not thread-safe) (**libha\_rr.a**)

### Syntax

```
#include <ha_rmap.h>

int
  ha_rr_end_session(
    int          session_fd,
    struct ha_em_err_blk *rr_errb)
```

### Parameters

*session\_fd*

The file descriptor of the session to be ended.

*rr\_errb* A pointer to an error block structure.

### Description

The **ha\_rr\_end\_session** subroutine ends the session between the resource monitor and the resource monitor manager. A resource monitor that is not a command calls this subroutine when the resource monitor manager associated with the session has disconnected. A command-based resource monitor calls this subroutine before it terminates.

### Security

The calling process must have a real or effective group ID of **haemrm**, or must have the **haemrm** group ID in its supplemental group list. If the process is not in the **haemrm** group, its effective user ID must be **root**.

If the calling process is instance 0 of the resource monitor, and the monitor is configured to supply Counter or Quantity variables to the Performance Monitor, it must have an effective user ID of **root**.

### Restrictions

Before it calls the **ha\_rr\_end\_session** subroutine, the resource monitor must first call the **ha\_rr\_del\_var** subroutine to delete all of the variables that were added to the RMAPI during the session.

If it was started by a resource monitor manager and if it no longer has an active session after it calls the **ha\_rr\_end\_session** subroutine, the resource monitor should terminate.

The **ha\_rr\_end\_session** subroutine must be called only as specified in this man page. Calling this subroutine in any other manner is a programming error with undefined results.

### Return Values

If the **ha\_rr\_end\_session** subroutine is successful, it returns a value of 0. The file descriptor for the session just ended must no longer be used as an argument to the **select** or **poll** system call.

## Error Values

If the **ha\_rr\_end\_session** subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

The RMAPI error block and error numbers are defined in the **ha\_rmapi.h** header file. For more information on RMAPI errors, see “RMAPI Errors (err\_rmapi)” on page 131.

For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

## Examples

For examples of using RMAPI subroutines, see the programs in the RSCT product samples directory, **/usr/sbin/rsct/samples/haem/rmapi**.

## Files

**ha\_rmapi.h**

## Prerequisite Information

“Chapter 1. Understanding Event Management” on page 1.

## Related Information

Structures in the **ha\_rmapi.h** header file: **ha\_em\_err\_blk**

Subroutines: **ha\_rr\_start\_session**

---

## ha\_rr\_get\_ctrlmsg Subroutine

### Purpose

**ha\_rr\_get\_ctrlmsg** – Get a control message from a resource monitor manager

### Library

RMAPI Library (not thread-safe) (**libha\_rr.a**)

### Syntax

```
#include <ha_rmap.h>

int
  ha_rr_get_ctrlmsg(
    int          session_fd,
    struct ha_rr_ctrl_msg **rr_ctrl_msg,
    struct ha_em_err_blk *rr_errb)
```

### Parameters

*session\_fd*

A session file descriptor that is ready to be read.

*rr\_ctrl\_msg*

A pointer to the address of a control message buffer.

*rr\_errb* A pointer to an error block structure.

### Description

When a resource monitor is notified that a file descriptor returned by the **ha\_rr\_start\_session** subroutine is ready for reading, as indicated by the **select** or **poll** system call, or **SIGIO** is caught, the resource monitor calls the **ha\_rr\_get\_ctrlmsg** subroutine to get a control message.

If **HA\_RR\_NOTIFY\_SELECT** notification is being used, *session\_fd* contains a session file descriptor that is ready to be read.

If **HA\_RR\_NOTIFY\_SIGIO** notification is being used, *session\_fd* contains a session file descriptor that may be ready for reading. The **ha\_rr\_get\_ctrlmsg** subroutine must be called once for each file descriptor that is returned by the **ha\_rr\_start\_session** subroutine.

### Security

The calling process must have a real or effective group ID of **haemrm**, or must have the **haemrm** group ID in its supplemental group list. If the process is not in the **haemrm** group, its effective user ID must be **root**.

If the calling process is instance 0 of the resource monitor, and the monitor is configured to supply Counter or Quantity variables to the Performance Monitor, it must have an effective user ID of **root**.

### Restrictions

It is the responsibility of the calling routine to free the control message buffer. When the buffer is freed, any memory areas to which there are pointers in the control message can no longer be accessed.

If the **SIGIO** notification method is being used, the returned buffer may contain more than one message. It is possible that more than one message had been received before the **SIGIO** signal could be processed.

Before calling the **ha\_rr\_get\_ctrlmsg** subroutine, the resource monitor must block the **SIGIO** signal and must establish a signal handler for it.

## Return Values

If the **ha\_rr\_get\_ctrlmsg** subroutine returns a value of 0, the response is intended for the Resource Monitor Application Programming Interface (RMAPI) itself. No action by the calling routine is necessary.

If the **ha\_rr\_get\_ctrlmsg** subroutine returns a value that is greater than 0, it also returns a pointer to a buffer that contains one or more control messages. The *rr\_ctrl\_msg* argument points to an area where the buffer pointer is returned. The return value indicates the length of the data in the control message buffer.

The control message has the following structure:

```
struct ha_rr_ctrl_msg {
    int    rr_ctrl_msg_len;
    int    rr_ctrl_cmd;
    int    rr_ctrl_cmdarg;
    int    rr_ctrl_num_vars;
    union  ha_rr_ctrlv {
        struct ha_rr_ctrl_var {
            char    *rr_ctrl_name;
            char    *rr_ctrl_rsrc_ID;
        } rr_ctrl_varn[1];
        int rr_ctrl_vari[1];
        struct ha_rr_ctrl_var2 {
            int    rr_ctrl_var_id;
            int    rr_ctrl_API_inst_id;
        } rr_ctrl_varn2[1];
    } rr_ctrlv;
#define rr_ctrl_vars    rr_ctrlv.rr_ctrl_varn
#define rr_ctrl_ids    rr_ctrlv.rr_ctrl_vari
#define rr_ctrl_vars2  rr_ctrlv.rr_ctrl_varn2
};
```

The **rr\_ctrl\_msg\_len** field contains the length of the message. If this value is less than the value that was returned by the **ha\_rr\_get\_ctrlmsg** subroutine, one or more additional messages are present in the buffer. To get a pointer to the next message in the buffer, increment the buffer pointer by the length of the message.

An actual control message may contain any number of elements in the **rr\_ctrl\_varn** or **rr\_ctrl\_vari** arrays; the **rr\_ctrl\_num\_vars** field contains the actual number.

The **rr\_ctrl\_varn2** array is not used in any control messages that are returned to the resource monitor.

The **rr\_ctrl\_cmd** field contains a command to the resource monitor. The **rr\_ctrl\_cmdarg** field contains an optional command argument.

The **rr\_ctrl\_cmd** field can contain one of the following values:

### HA\_RR\_CMD\_INSTV

This command tells the resource monitor to create instances of resource variables that match the variables listed in the **rr\_ctrl\_varn** array. Each element of the array contains two pointers: a pointer to a string that contains the name of a resource variable and a pointer to a string that contains a resource ID. If any resource ID element is omitted, the resource ID is wildcarded. If all resource ID elements are omitted, the resource ID string is null.

The resource monitor creates actual instances for each specified instance of a resource variable. If the resource ID is wildcarded, the resource monitor creates all actual instances that match the wildcarded resource ID for a specified resource variable. If no actual instances can be created that match the specified instance(s), this command is ignored.

This command is sent only for variables that are configured as being dynamically instantiable.

In response, the resource monitor should create the instances and then register them with the RMAPI using the **ha\_rr\_reg\_var** subroutine. If any instances specified by the command have

already been created and registered as a result of a previous **HA\_RR\_CMD\_INSTV** command, the instances must be registered again using the same values that were specified in the previous call to the **ha\_rr\_reg\_var** subroutine.

**Note:** If one of the resource ID element names of a resource variable matches the resource ID element name specified by the Locator field of the variable's definition, then that resource ID element is never supplied in a resource ID that is specified by the **HA\_RR\_CMD\_INSTV** command. Note that this may result in a null resource ID.

#### **HA\_RR\_CMD\_ADDV**

This command tells the resource monitor to add to the RMAPI the variables listed in the **rr\_ctrl\_vari** array. Each element of the array is an identifier for a particular instance of a resource variable. The identifier is supplied to the RMAPI when the resource monitor registers its variable instances.

In response, the resource monitor should call the **ha\_rr\_add\_var** subroutine to add the variables and supply current values for resource variable instances. Then, it should call the **ha\_rr\_send\_val** subroutine to start sending updated resource variable values for those instances that are indicated by the **ha\_rr\_add\_var** subroutine.

#### **HA\_RR\_CMD\_ADDALL**

This command tells the resource monitor to add to the RMAPI all of the variables of value type Counter or Quantity that have been registered with the RMAPI. The **rr\_ctrl\_num\_vars** field is set to a value of 0.

In response, the resource monitor should call the **ha\_rr\_add\_var** subroutine to add the variables and supply current values for resource variable instances. Then, it should call the **ha\_rr\_send\_val** subroutine to start sending updated resource variable values for those instances that are indicated by the **ha\_rr\_add\_var** subroutine.

This command can only be sent to the resource monitor if the monitor is configured to supply Counter or Quantity variables to the performance monitor.

#### **HA\_RR\_CMD\_DELV**

This command tells the resource monitor to delete from the resource monitor manager session the variables that are listed in the **rr\_ctrl\_vari** array. Each element of the array is an identifier for a particular instance of a resource variable. The identifier is supplied to the RMAPI when the resource monitor registers its variable instances.

In response, the resource monitor should call the **ha\_rr\_del\_var** subroutine to delete the variables. Then, it should stop sending resource variable values for those instances that are indicated by the **ha\_rr\_del\_var** subroutine.

#### **HA\_RR\_CMD\_DELALL**

This command tells the resource monitor to delete from the resource monitor manager session, using the **ha\_rr\_del\_var** subroutine, all of the variables of value type Counter or Quantity that have been registered with the RMAPI. The **rr\_ctrl\_num\_vars** field is set to a value of 0.

In response, the resource monitor should call the **ha\_rr\_del\_var** subroutine to delete the variables. Then, it should stop sending resource variable values for those instances that are indicated by the **ha\_rr\_del\_var** subroutine.

This command can only be sent to the resource monitor if the monitor is configured to supply Counter or Quantity variables to the performance monitor.

#### **HA\_RR\_CMD\_REFRESH**

This command tells the resource monitor to send immediately the latest possible values for the variables that are listed in the **rr\_ctrl\_vari** array. Each element of the array is an identifier for a particular instance of a resource variable. The identifier is supplied to the RMAPI when the resource monitor registers its variable instances.

This command is sent to the resource monitor only for variables of value type Counter or Quantity that have configurable reporting intervals.

In response, the resource monitor should call the **ha\_rr\_send\_val** subroutine to send the latest values for the variables.

This command is not currently sent by any resource monitor manager.

## Error Values

If the **ha\_rr\_get\_ctrlmsg** subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

The RMAPI error block and error numbers are defined in the **ha\_rmapi.h** header file. For more information on RMAPI errors, see “RMAPI Errors (err\_rmapi)” on page 131.

For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

The following error numbers are either normal or require specific action by the resource monitor:

### HA\_RR\_EAGAIN

The resource monitor specified **HA\_RR\_NOTIFY\_SIGIO** on the call to the **ha\_rr\_start\_session** subroutine and no data was read by the **ha\_rr\_get\_ctrlmsg** subroutine. Because the **ha\_rr\_get\_ctrlmsg** subroutine must be called each time **SIGIO** is caught, this error is normal.

If the resource monitor specified **HA\_RR\_NOTIFY\_SELECT**, then this error indicates that the monitor called **ha\_rr\_get\_ctrlmsg** for a session that does not have any data ready to read. The monitor may continue executing.

### HA\_RR\_EDISCONNECT

The resource monitor manager that is associated with the session file descriptor specified on input to the **ha\_rr\_get\_ctrlmsg** subroutine has dropped its connection. In response, the resource monitor should call the **ha\_rr\_del\_var** subroutine using this file descriptor and specifying all variables that have been added by the resource monitor. Finally, the resource monitor should call the **ha\_rr\_end\_session** subroutine for the specified session.

## Examples

For examples of using RMAPI subroutines, see the programs in the RSCT product samples directory, **/usr/sbin/rsct/samples/haem/rmapi**.

## Files

**ha\_rmapi.h**

## Prerequisite Information

“Chapter 1. Understanding Event Management” on page 1.

## Related Information

Structures in the **ha\_rmapi.h** header file: **ha\_rr\_ctrl\_msg**, **ha\_em\_err\_blk**

Subroutines: **ha\_rr\_start\_session**, **ha\_rr\_end\_session**

---

## ha\_rr\_get\_interval Subroutine

### Purpose

**ha\_rr\_get\_interval** – Get the reporting interval for a class of resource variables

### Library

RMAPI Library (not thread-safe) (**libha\_rr.a**)

### Syntax

```
#include <ha_rmap_i.h>

int
    ha_rr_get_interval(
        char *cname,
        struct ha_em_err_blk *rr_errb)
```

### Parameters

*cname* A pointer to a string that contains the name of the class of resource variables for which the reporting interval is to be returned.

*rr\_errb* A pointer to an error block structure.

### Description

The **ha\_rr\_get\_interval** subroutine is used by the resource monitor to get the reporting interval for a class of resource variables whose reporting interval is configurable. The reporting interval specifies how often a resource variable of value type Counter or Quantity of a particular class is to be sent to the RMAPI.

### Security

The calling process must have a real or effective group ID of **haemrm**, or must have the **haemrm** group ID in its supplemental group list. If the process is not in the **haemrm** group, its effective user ID must be **root**.

If the calling process is instance 0 of the resource monitor, and the monitor is configured to supply Counter or Quantity variables to the Performance Monitor, it must have an effective user ID of **root**.

### Return Values

If the **ha\_rr\_get\_interval** subroutine is successful, it returns the reporting interval, in seconds.

### Error Values

If the **ha\_rr\_get\_interval** subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

The RMAPI error block and error numbers are defined in the **ha\_rmap\_i.h** header file. For more information on RMAPI errors, see “RMAPI Errors (err\_rmap\_i)” on page 131.

For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

### Examples

For examples of using RMAPI subroutines, see the programs in the RSCT product samples directory, **/usr/sbin/rsct/samples/haem/rmap\_i**.

## **Files**

`ha_rmapi.h`

## **Prerequisite Information**

“Chapter 1. Understanding Event Management” on page 1.

## **Related Information**

Structures in the `ha_rmapi.h` header file: `ha_em_err_blk`

---

## ha\_rr\_init Subroutine

### Purpose

**ha\_rr\_init** – Initialize the Resource Monitor Application Programming Interface (RMAPI)

### Library

RMAPI Library (not thread-safe) (**libha\_rr.a**)

### Syntax

```
#include <ha_rmap_i.h>

int
  ha_rr_init(
    char      *name,
    struct ha_em_err_blk *rr_errb)
```

### Parameters

*name* A pointer to a string that contains the resource monitor name, as defined in the Event Management Configuration Database (EMCDB).

*rr\_errb* A pointer to an error block structure.

### Description

The **ha\_rr\_init** subroutine initializes a resource monitor, informing the Resource Monitor Application Programming Interface (RMAPI) of its identity. The RMAPI obtains configuration information for the resource monitor, initializes internal resources, determines the domain in which the monitor is executing, and creates a lock for the monitor instance as specified in the EMCDB and by any prior call to the **ha\_rr\_rm\_ctl** subroutine.

#### Determining the Event Management Domain

The **ha\_rr\_init** routine determines the **domain** in which the resource monitor is executing, and uses the domain name in forming socket and file names. If the resource monitor is started by a resource monitor manager, the manager sets environment variables to inform the RMAPI of the domain. The environment variables used to indicate the domain are:

##### HA\_DOMAIN\_TYPE

The type of domain, either **SP** for RS/6000 SP, or **HACMP** for HACMP/ES clusters.

##### HA\_DOMAIN\_NAME

The name of the RS/6000 SP system partition or HACMP/ES cluster the monitor is executing in.

Command-based monitors and monitors that are not started by resource monitor managers may need to set the environment variables prior to calling **ha\_rr\_init**. The RMAPI determines and verifies the domain it is executing in as follows:

- If **HA\_DOMAIN\_NAME** is set, but **HA\_DOMAIN\_TYPE** is not, the RMAPI returns an error.
- If neither the **HA\_DOMAIN\_TYPE** nor the **HA\_DOMAIN\_NAME** environment variable is set, the RMAPI defaults to the action taken when the **HA\_DOMAIN\_TYPE** is **SP**.
- The RMAPI attempts to determine the domain name based on the domain type:
  - If the domain type is **HACMP**, the RMAPI calls an HACMP/ES command to query the name of the HACMP/ES cluster.
  - If the domain type is **SP** and the RMAPI is executing on a node, the domain name is the system partition the node is in. If the RMAPI is executing on the control workstation, the system partition is

determined by the value of the **SP\_NAME** environment variable. If **SP\_NAME** has not been set, the system partition name used is the name of the default partition.

- If **HA\_DOMAIN\_NAME** is set, the RMAPI compares the value to the domain name it determined. If the domain type is **SP**, and the RMAPI is executing on the control workstation, the value of **HA\_DOMAIN\_NAME** is compared against all defined system partitions. If the value of **HA\_DOMAIN\_NAME** does not match a domain name value determined by the RMAPI, an error is returned.
- If the **HA\_DOMAIN\_NAME** variable is not set, the RMAPI uses the domain name that it determined.

## Resource Monitor Instance ID

A resource monitor may be configured to execute multiple, up to **HA\_EM\_MAX\_RM\_INSTS**, copies or *instances* of itself at the same time. The number of instances allowed is defined by the value of the **rmNum\_instances** attribute of the **EM\_Resource\_Monitor** SDR class. If the attribute is not present, the number of simultaneous instances allowed for the monitor defaults to 1.

Resource monitor instances are identified as having instance numbers or IDs in the range 0–**HA\_RR\_RM\_INSTID\_MAX**. When **ha\_rr\_init** is called, the RMAPI attempts to lock the lowest available resource monitor instance ID based on the monitor's configuration and instance ID set by a prior call to **ha\_rr\_rm\_ctl**. Since the number of allowed instances is determined by the definition of the monitor, it is possible to call **ha\_rr\_rm\_ctl** to set an instance ID that is not valid for the monitor. In this case, the **ha\_rr\_init** routine fails and returns the error.

## Security

The calling process must have a real or effective group ID of **haemrm**, or must have the **haemrm** group ID in its supplemental group list. If the process is not in the **haemrm** group, its effective user ID must be **root**.

If the calling process is instance 0 of the resource monitor, and the monitor is configured to supply Counter or Quantity variables to the Performance Monitor, it must have an effective user ID of **root**.

## Return Values

If the **ha\_rr\_init** subroutine is successful, it returns a value of 0.

## Error Values

If the **ha\_rr\_init** subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

The RMAPI error block and error numbers are defined in the **ha\_rmapi.h** header file. For more information on RMAPI errors, see “RMAPI Errors (err\_rmapi)” on page 131.

For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

The following error number requires specific action by the resource monitor:

### **HA\_RR\_ENOLOCK**

The RMAPI was not able to acquire a unique lock for the monitor process. This indicates that either the maximum number of instances the monitor is configured to allow is already executing, or a specific monitor instance requested by previous call to **ha\_rr\_rm\_ctl** was already executing.

### **HA\_RR\_EACCESS**

The RMAPI was unable to open or create a file required for initialization. This error could occur if the resource monitor was started before the Event Management subsystem and before the EMCDB file or directories used by the RMAPI had been created. If the resource monitor is a

server which is not configured to be started by resource monitor managers, it may periodically retry the `ha_rr_init` routine, waiting at least several seconds between attempts. The decision to retry the `ha_rr_init` routine when this error is encountered, and the number of attempts made, should be based upon the design of the resource monitor.

## Examples

For examples of using RMAPI subroutines, see the programs in the PSSP product samples directory, `/usr/lpp/ssp/samples/haem/rmapi`.

## Files

`ha_rmapi.h`

## Prerequisite Information

“Chapter 1. Understanding Event Management” on page 1.

## Related Information

Structures in the `ha_rmapi.h` header file: `ha_em_err_blk`

---

## ha\_rr\_makserv Subroutine

### Purpose

**ha\_rr\_makserv** – Establish the resource monitor as a server to enable connections from resource monitor managers

### Library

RMAPI Library (not thread-safe) (**libha\_rr.a**)

### Syntax

```
#include <ha_rmap.h>

int
  ha_rr_makserv(
    int          rr_notify_proto,
    struct ha_em_err_blk *rr_errb)
```

### Parameters

*rr\_notify\_proto*

A value that indicates the method the resource monitor will use to detect when a resource monitor manager is attempting to connect to the resource monitor. The notification protocol specified to the **ha\_rr\_makserv** must be the same value used when calling **ha\_rr\_start\_session**. The possible values are:

#### **HA\_RR\_NOTIFY\_SELECT**

The resource monitor expects to use the **select** or **poll** system call to determine when a manager is making a connection request.

#### **HA\_RR\_NOTIFY\_SIGIO**

The resource monitor expects to use **SIGIO** notification to determine when a manager connection request may be pending.

*rr\_errb* A pointer to an error block structure.

### Description

The **ha\_rr\_makserv** subroutine creates a server session for a resource monitor so that resource monitor managers can connect to it. Resource monitor managers are the Event Management subsystem or the Performance Monitor subsystem.

If the resource monitor is not a command, the **ha\_rr\_makserv** subroutine should be called after the call to **ha\_rr\_init**.

If the resource monitor is a command, do not call the **ha\_rr\_makserv** subroutine.

### Security

The calling process must have a real or effective group ID of **haemrm**, or must have the **haemrm** group ID in its supplemental group list. If the process is not in the **haemrm** group, its effective user ID must be **root**.

If the calling process is instance 0 of the resource monitor, and the monitor is configured to supply Counter or Quantity variables to the Performance Monitor, it must have an effective user ID of **root**.

## Restrictions

If the value of the `rr_notify_proto` argument is `HA_RR_NOTIFY_SIGIO`, then before calling the `ha_rr_makserv` subroutine, the resource monitor must block the SIGIO signal and must establish a signal handler for it.

## Return Values

If the `ha_rr_makserv` subroutine is successful, it returns a value greater than or equal to 0. The value is a file descriptor that is used by the resource monitor to determine when a resource monitor manager is attempting to start a session with the resource monitor.

If the `rr_notify_proto` argument is the value `HA_RR_NOTIFY_SELECT`, the resource monitor is expected to use this descriptor in the argument to the `select` or `poll` system call (as a read event).

If the value of the `rr_notify_proto` argument is the value `HA_RR_NOTIFY_SIGIO`, the resource monitor is expected to catch the `SIGIO` signal. The `ha_rr_makserv` subroutine enables the descriptor for `SIGIO` notification. Notification that the descriptor is ready for read, or potentially ready for read if `SIGIO` is caught, indicates that a resource monitor manager is, or may be, attempting to connect to the resource monitor.

Note that a resource monitor must always select (or poll) on the file descriptor that is returned by the `ha_rr_makserv` subroutine or must always catch the `SIGIO` signal. The Event Management subsystem and the Performance Monitor subsystem may come and go while the resource monitor is otherwise executing normally.

## Error Values

If the `ha_rr_makserv` subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

The RMAPI error block and error numbers are defined in the `ha_rmapi.h` header file. For more information on RMAPI errors, see “RMAPI Errors (err\_rmapi)” on page 131.

For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

## Examples

For examples of using RMAPI subroutines, see the programs in the RSCT product samples directory, `/usr/sbin/rsct/samples/haem/rmapi`.

## Files

`ha_rmapi.h`

## Prerequisite Information

“Chapter 1. Understanding Event Management” on page 1.

## Related Information

Structures in the `ha_rmapi.h` header file: `ha_em_err_blk`

---

## ha\_rr\_reg\_var Subroutine

### Purpose

**ha\_rr\_reg\_var** – Register a resource variable instance with the Resource Monitor Application Programming Interface (RMAPI)

### Library

RMAPI Library (not thread-safe) (**libha\_rr.a**)

### Syntax

```
#include <ha_rmap_i.h>

int
  ha_rr_reg_var(
    struct ha_rr_variable *pv,
    int numv,
    struct ha_em_err_blk *rr_errb)
```

### Parameters

*pv* A pointer to an array of **ha\_rr\_variable** structures, which contain the names, resource IDs, and instance identifiers of the resource variable instances to be registered.

*numv* The number of elements in the array pointed to by the *pv* parameter.

*rr\_errb* A pointer to an error block structure.

### Description

The **ha\_rr\_reg\_var** subroutine is used by the resource monitor to register resource variable instances with the RMAPI.

The resource monitor must register each instance of a resource variable known to the resource monitor at the following times:

- After it calls the **ha\_rr\_init** subroutine
- After it instantiates any resource variable upon receipt of the **HA\_RR\_CMD\_INSTV** command. If the **HA\_RR\_CMD\_INSTV** command specifies resource variable instances that have already been registered, the instances must be registered again.
- After it creates any resource variable's instances during normal operation.

The **ha\_rr\_variable** structure has the following definition:

```
struct ha_rr_variable {
  char *rr_var_name;
  char *rr_var_rsrc_ID;
  union {
    int rr_var_inst_id;
    void **rr_var_hdl;
  } rr_varu;
#define rr_var_hdl rr_varu.rr_var_hdl
#define rr_var_inst_id rr_varu.rr_var_inst_id
  void *rr_value;
  int rr_var_errno;
}
```

The **rr\_var\_name** field is a pointer to a string that contains the name of a resource variable.

The **rr\_var\_rsrc\_ID** field is a pointer to a string that contains the resource ID of an instance of the specified variable. The resource ID must be fully qualified; no wildcarding is permitted.

The resource ID is specified as a comma-separated list of name/value pairs. A name/value pair consists of a resource ID element name followed by an equal sign followed by the value of the resource ID element. There are no blanks in the resource ID.

The **rr\_var\_inst\_id** field is an instance identifier that the resource monitor assigns to the instance. The instance identifier is used to specify variable instances efficiently in commands that are sent to the resource monitor. Typically, the instance identifier is an index into a table that is maintained by the resource monitor. If the variable instance has already been registered, this value must match the value in the previous registration.

The **rr\_value** field is not used.

The **rr\_var\_errno** field is used to hold the result of the registration for each variable.

## Security

The calling process must have a real or effective group ID of **haemrm**, or must have the **haemrm** group ID in its supplemental group list. If the process is not in the **haemrm** group, its effective user ID must be **root**.

If the calling process is instance 0 of the resource monitor, and the monitor is configured to supply Counter or Quantity variables to the Performance Monitor, it must have an effective user ID of **root**.

## Restrictions

If the name of one of the elements in any resource ID matches the resource ID element name specified by the Locator field in the associated resource variable definition, that resource ID element name/value pair is never supplied by the resource monitor in the resource ID specified by the **rr\_var\_rsrc\_ID** field. This resource ID element is added by the Event Management subsystem when the instance is given in a response to any Event Management client.

## Return Values

If the **ha\_rr\_reg\_var** subroutine is successful, it returns the number of resource variables that were registered. For each variable that was successfully registered, the associated **rr\_var\_errno** field in the **ha\_rr\_variable** structure is set to 0. If a variable could not be registered, an error code is returned in its **rr\_var\_errno** field.

## Error Values

If the **ha\_rr\_reg\_var** subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

The following error number requires specific handling by the resource monitor:

### HA\_RR\_EDISCONNECT

This error is only returned to client type resource monitors and indicates that the session created by a call to **ha\_rr\_start\_session** was closed by the Event Management daemon. The resource monitor should respond by calling **ha\_rr\_terminate** to end the current use of the RMAPI.

The RMAPI error block and error numbers are defined in the **ha\_rmapi.h** header file. For more information on RMAPI errors, see “RMAPI Errors (err\_rmapi)” on page 131.

For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

## Examples

For examples of using RMAPI subroutines, see the programs in the RSCT product samples directory, `/usr/sbin/rsct/samples/haem/rmapi`.

## Files

`ha_rmapi.h`

## Prerequisite Information

“Chapter 1. Understanding Event Management” on page 1.

## Related Information

Structures in the `ha_rmapi.h` header file: `ha_rr_variable`, `ha_em_err_blk`

Subroutines: `ha_rr_unreg_var`

---

## ha\_rr\_rm\_ctl Subroutine

### Purpose

**ha\_rr\_rm\_ctl** – Set or get attributes of the RMAPI.

### Library

RMAPI Library (not thread safe) (**libha\_rr.a**)

### Syntax

```
#include <ha_rmap_i.h>
int
  ha_rr_rm_ctl(
    struct ha_rr_args  *argp,
    int                command,
    struct ha_em_err_blk *rr_errb);
```

### Parameters

*argp* A pointer to an **ha\_rr\_args** structure which is used to set or get attributes of the RMAPI.

*command*

Bit flag command to get or set attributes of the RMAPI.

*rr\_errb* A pointer to an error block structure.

### Description

The **ha\_rr\_args** structure is defined as:

```
struct ha_rr_args {
  int   rr_instance_id;
  char  *rr_domain_name;
  char  rr_reserved[56];
};
```

The **rr\_instance\_id** field is the requested or actual instance ID (number) of the resource monitor.

The **rr\_domain\_name** is the name of the domain in which the resource monitor is executing.

The **rr\_reserved** field is reserved for future use.

The **ha\_rr\_rm\_ctl** subroutine may be called prior to calling **ha\_rr\_init** to set attributes of the RMAPI for initialization. It may also be called after **ha\_rr\_init** in order to get the values of attributes initialized by the RMAPI. The following are valid values for the command parameter:

#### **HA\_RR\_RM\_ARGS\_SET\_INSTID**

Sets the resource monitor instance ID the RMAPI should use when **ha\_rr\_init** is called. The **rr\_instance\_id** field must be set to a value in the range, 0-N, where N is the number of simultaneous copies the monitor is configured to allow minus 1. One of the following symbols defined in **ha\_rmap\_i.h** may also be used:

#### **HA\_RR\_RM\_INSTID\_PERF**

Initialize only as instance 0 of the monitor.

#### **HA\_RR\_RM\_INSTID\_NOPERF**

Initialize as any instance except 0. The monitor must be configured to allow more than one instance of itself or **ha\_rr\_init** will fail.

#### **HA\_RR\_RM\_INSTID\_ANY**

Initialize as the lowest available instance number of the monitor. This is the default value if **ha\_rr\_rm\_ctl** is not used.

## HA\_RR\_RM\_ARGS\_GET

Gets attributes of the RMAPI after initialization. This command is only valid after **ha\_rr\_init** has been called. Upon successful return, the **ha\_rr\_args** parameter will be set with the following values:

*rr\_instance\_id*

The instance ID of the monitor.

*rr\_domain\_name*

A pointer to a buffer containing the name of the domain in which the monitor is executing.

## Security

### Restrictions

If **ha\_rr\_rm\_ctl** is called with the command **HA\_RR\_RM\_ARGS\_GET**, the buffer referred to by the **rr\_domain\_name** field must be freed by the caller. When setting the monitor instance ID, **ha\_rr\_rm\_ctl** only checks the instance value as being within the range **HA\_RR\_RM\_INSTID\_ANY** - **HA\_RR\_RM\_INSTID\_MAX**. Since the RMAPI does not know the resource monitor configuration prior to **ha\_rr\_init**, it is possible to set an instance ID that is not valid for the configuration of the monitor. In this case, the call to **ha\_rr\_init** will fail.

### Return Values

If the **ha\_rr\_rm\_ctl** subroutine is successful, it returns a 0. If the subroutine was called to get values, the **ha\_rr\_args** structure pointed to by the **argp** parameter will be filled with current values in use by the RMAPI. Note that if the **rr\_domain\_name** is not NULL, it refers to a buffer which was allocated by the RMAPI, but must be freed by the caller.

If **ha\_rr\_rm\_ctl** was called to set values, those values are only tested for being within the range of possible values valid for any resource monitor. The set values are further tested for their validity based on the configuration of the resource monitor specified when the RMAPI is initialized by calling **ha\_rr\_init**.

### Error Values

If the **ha\_rr\_rm\_ctl** subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message. The RMAPI error block and error numbers are defined in the **ha\_rmapi.h** header file. For more information on RMAPI errors, see "RMAPI Errors (err\_rmapi)" on page 131. For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

### Examples

For examples of using RMAPI subroutines, see the programs in the RSCT product samples directory, **/usr/sbin/rsct/samples/haem/rmapi**.

### Files

**ha\_rmapi.h**

### Prerequisite Information

"Chapter 1. Understanding Event Management" on page 1.

## Related Information

Structures in the `ha_rmapi.h` header file: `ha_rr_args`

Subroutines: `ha_rr_init`

---

## ha\_rr\_send\_val Subroutine

### Purpose

**ha\_rr\_send\_val** – Send variable values to the Resource Monitor Application Programming Interface (RMAPI)

### Library

RMAPI Library (not thread-safe) (**libha\_rr.a**)

### Syntax

```
#include <ha_rmap_i.h>

int
  ha_rr_send_val(
    struct ha_rr_val    *pv,
    int                 numv,
    int                 refresh,
    struct ha_em_err_blk *rr_errb)
```

### Parameters

*pv* A pointer to an array of **ha\_rr\_val** structures, which contain pointers to the values and handles of the resource variables that the resource monitor is sending to the RMAPI.

*numv* The number of elements in the array pointed to by the *pv* parameter.

*refresh*

A value that indicates whether the resource variable values are being sent in response to a **HA\_RR\_CMD\_REFRESH** command. A value of 0 indicates that the values are not being refreshed. A nonzero value indicates that they are being refreshed.

*rr\_errb* A pointer to an error block structure.

### Description

The **ha\_rr\_send\_val** subroutine is used by the resource monitor whenever it is time to send values of resource variables to the RMAPI.

The values of State type resource variables are sent whenever the monitor detects that they have changed. Counter and Quantity values are sent based on their reporting interval as determined by the design of the monitor or the reporting interval returned by the **ha\_rr\_get\_interval** subroutine.

The **ha\_rr\_val** structure has the following definition:

```
struct ha_rr_val {
    void *rr_value;
    void *rr_var_hdl;
}
```

The **rr\_value** field is a pointer to the value of a resource variable. The **rr\_var\_hdl** field contains the variable handle of the resource variable that was returned by the **ha\_rr\_add\_var** subroutine.

### Security

The calling process must have a real or effective group ID of **haemrm**, or must have the **haemrm** group ID in its supplemental group list. If the process is not in the **haemrm** group, its effective user ID must be **root**.

If the calling process is instance 0 of the resource monitor, and the monitor is configured to supply Counter or Quantity variables to the Performance Monitor, it must have an effective user ID of **root**.

## Restrictions

The resource monitor may send as many variable values as it desires in a single call of the **ha\_rr\_send\_val** subroutine. The actual number that it sends depends on the semantics of the variables, the required frequency of update, and the design of the resource monitor.

However, if any of the variables that were added are of value type Counter or Quantity and were configured to be supplied to the performance monitor, the resource monitor must perform some “send” activity every 500 seconds. If the resource monitor has no values to send within 500 seconds from the last call to **ha\_rr\_send\_val**, it must call the **ha\_rr\_touch** subroutine.

## Return Values

If the **ha\_rr\_send\_val** subroutine is successful, it returns a value of 0.

## Error Values

If the **ha\_rr\_send\_val** subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

The following error number requires specific handling by the resource monitor:

### HA\_RR\_EDISCONNECT

This error is only returned to client type resource monitors and indicates that the session created by a call to **ha\_rr\_start\_session** was closed by the Event Management daemon. The resource monitor should respond by calling **ha\_rr\_terminate** to end the current use of the RMAPI.

The RMAPI error block and error numbers are defined in the **ha\_rmapi.h** header file. For more information on RMAPI errors, see “RMAPI Errors (err\_rmapi)” on page 131.

For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

## Examples

For examples of using RMAPI subroutines, see the programs in the RSCT product samples directory, **/usr/sbin/rsct/samples/haem/rmapi**.

## Files

**ha\_rmapi.h**

## Prerequisite Information

“Chapter 1. Understanding Event Management” on page 1.

## Related Information

Structures in the **ha\_rmapi.h** header file: **ha\_rr\_val**, **ha\_em\_err\_blk**

Subroutines: **ha\_rr\_get\_ctrlmsg**, **ha\_rr\_touch**

---

## ha\_rr\_start\_session Subroutine

### Purpose

**ha\_rr\_start\_session** – Start a session with a resource monitor manager

### Library

RMAPI Library (not thread-safe) (**libha\_rr.a**)

### Syntax

```
#include <ha_rmap.h>

int
  ha_rr_start_session(
    int          rr_notify_proto,
    struct ha_em_err_blk *rr_errb)
```

### Parameters

*rr\_notify\_proto*

A value that indicates the method the resource monitor will use to detect when a resource monitor manager session file descriptor is ready to be read by a call to **ha\_rr\_get\_ctrlmsg**.

#### **HA\_RR\_NOTIFY\_SELECT**

The resource monitor expects to use the **select** or **poll** system call to detect when session file descriptors are ready to be read.

#### **HA\_RR\_NOTIFY\_SIGIO**

The resource monitor expects to use **SIGIO** notification to determine when a session file descriptor may be ready to be read.

**Note:** Since client type resource monitors do not receive commands from resource monitor managers, the value of the *rr\_notify\_proto* parameter is ignored.

*rr\_errb* A pointer to an error block structure.

### Description

The **ha\_rr\_start\_session** subroutine is called by a resource monitor to establish a communication path, or session, between the resource monitor and one of its managers. This communication path is used to send control messages between the resource monitor manager and the resource monitor. For a command-based resource monitor, a session establishes a communication path to the Event Management subsystem only.

If the resource monitor is command-based, it does not call the **ha\_rr\_makserv** subroutine. Instead, it calls the **ha\_rr\_start\_session** subroutine immediately after the call to the **ha\_rr\_init** subroutine completes successfully.

Otherwise, whenever the file descriptor returned by the **ha\_rr\_makserv** subroutine is ready for reading, as indicated by the notification protocol used by the monitor, the resource monitor calls the **ha\_rr\_start\_session** subroutine.

### Security

The calling process must have a real or effective group ID of **haemrm**, or must have the **haemrm** group ID in its supplemental group list. If the process is not in the **haemrm** group, its effective user ID must be **root**.

If the calling process is instance 0 of the resource monitor, and the monitor is configured to supply Counter or Quantity variables to the Performance Monitor, it must have an effective user ID of **root**.

## Restrictions

If the **ha\_rr\_makserv** subroutine is called before the **ha\_rr\_start\_session** subroutine is called, the value of the *rr\_notify\_proto* argument must be the same for both calls.

If the value of the **rr\_notify\_proto** argument is **HA\_RR\_NOTIFY\_SIGIO**, then before calling the **ha\_rr\_start\_session** subroutine, the resource monitor must block the SIGIO signal and must establish a signal handler for it.

## Return Values

If the **ha\_rr\_start\_session** subroutine is successful, it returns a value greater than or equal to 0. The value is a file descriptor that is used by the resource monitor to determine when a resource monitor manager has sent a control message to the resource monitor.

If it is not command-based, the resource monitor must be prepared to accept multiple sessions, up to a value of **HA\_RR\_MAXSESSIONS**.

If the *rr\_notify\_proto* argument is the value **HA\_RR\_NOTIFY\_SELECT**, the resource monitor is expected to use this descriptor in the argument to the **select** or **poll** system call (as a read event).

If the *rr\_notify\_proto* argument is the value **HA\_RR\_NOTIFY\_SIGIO**, the resource monitor is expected to catch the **SIGIO** signal. The **ha\_rr\_start\_session** subroutine enables the descriptor for **SIGIO** notification. Notification that the descriptor is ready for read, or potentially ready for read if **SIGIO** is caught, indicates that a control message has, or may have, arrived.

The file descriptor is also used as a session handle by other RMAPI subroutines. Because a command-based resource monitor does not receive control messages, it uses the file descriptor only as a session handle.

## Error Values

If the **ha\_rr\_start\_session** subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

Error Numbers:

### **HA\_RR\_ECONNREFUSED**

This error number applies only to client type resource monitors and indicates that a connection attempt to the Event Management daemon was refused. Since the Event Management daemon may not be running or may be initializing, the resource monitor may periodically retry the connection attempt to the Event Management daemon by waiting several seconds, then calling **ha\_rr\_start\_session** again. The monitor may otherwise end the current use of the RMAPI by calling **ha\_rr\_terminate**. Since client type resource monitors are typically implemented in a command, the decision to retry and the number of connection attempts made should be based on the design of the resource monitor.

### **HA\_RR\_EMAXSESSIONS**

Accepting the connection request would result in more than **HA\_RR\_MAX\_SESSIONS** connections to resource monitor managers. The RMAPI closes the pending connection request and the resource monitor may continue normal execution. This error only applies to server type resource monitors. To avoid this error, resource monitors should ensure that the

**ha\_rr\_end\_session** subroutine is called whenever the **HA\_RR\_EDISCONNECT** error has been detected. This allows the RMAPI to clean up closed manager sessions, enabling new connection requests to be accepted.

## **HA\_RR\_EAGAIN**

No connection request was found by the **ha\_rr\_start\_session** routine. If the notification protocol is **HA\_RR\_NOTIFY\_SIGIO**, this can occur because the **ha\_rr\_start\_session** must be called each time a SIGIO signal is caught. If the notification protocol is **HA\_RR\_NOTIFY\_SELECT**, the error indicates that the resource monitor attempted to start a session before the select or poll system call indicated the server socket returned by **ha\_rr\_makserv** had a pending connection request. This error applies only to server type resource monitors.

The RMAPI error block and error numbers are defined in the **ha\_rmapi.h** header file. For more information on RMAPI errors, see “RMAPI Errors (err\_rmapi)” on page 131.

For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

## **Examples**

For examples of using RMAPI subroutines, see the programs in the RSCT product samples directory, **/usr/sbin/rsct/samples/haem/rmapi**.

## **Files**

**ha\_rmapi.h**

## **Prerequisite Information**

“Chapter 1. Understanding Event Management” on page 1.

## **Related Information**

Structures in the **ha\_rmapi.h** header file: **ha\_em\_err\_blk**

---

## ha\_rr\_terminate Subroutine

### Purpose

**ha\_rr\_terminate** – Terminate communications with the Resource Monitor Application Programming Interface (RMAPI)

### Library

RMAPI Library (not thread-safe) (**libha\_rr.a**)

### Syntax

```
#include <ha_rmap_i.h>

int
    ha_rr_terminate(
        struct ha_em_err_blk *rr_errb)
```

### Parameters

*rr\_errb* A pointer to an error block structure.

### Description

The **ha\_rr\_terminate** subroutine is used by the resource monitor to terminate its activity with the RMAPI.

If a server type resource monitor is configured to be started by a resource monitor manager, the **ha\_rr\_terminate** subroutine is called whenever it no longer has any sessions with resource monitor managers. Upon return from this subroutine, it exits.

If the resource monitor logic is contained within a subsystem, the **ha\_rr\_terminate** subroutine is called before the exit of the subsystem.

If the resource monitor is a command, it calls the **ha\_rr\_terminate** subroutine before it exits.

The **ha\_rr\_terminate** subroutine should also be invoked whenever the resource monitor detects an unrecoverable error, from the RMAPI or from any other source, that no longer permits the successful execution of the resource monitor's function. Upon detection of the unrecoverable error, the resource monitor should not call any other RMAPI subroutine before it calls **ha\_rr\_terminate**. The **ha\_rr\_terminate** subroutine releases all resources that were acquired by previous calls to the RMAPI.

### Security

The calling process must have a real or effective group ID of **haemrm**, or must have the **haemrm** group ID in its supplemental group list. If the process is not in the **haemrm** group, its effective user ID must be **root**.

If the calling process is instance 0 of the resource monitor, and the monitor is configured to supply Counter or Quantity variables to the Performance Monitor, it must have an effective user ID of **root**.

### Return Values

If the **ha\_rr\_terminate** subroutine is successful, it returns a value of 0. The file descriptor that was returned by the **ha\_rr\_makserv** subroutine must no longer be used in an argument to the **select** or **poll** system call.

## Error Values

If the **ha\_rr\_terminate** subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

The RMAPI error block and error numbers are defined in the **ha\_rmapi.h** header file. For more information on RMAPI errors, see “RMAPI Errors (err\_rmapi)” on page 131.

For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

## Examples

For examples of using RMAPI subroutines, see the programs in the RSCT product samples directory, **/usr/sbin/rsct/samples/haem/rmapi**.

## Files

**ha\_rmapi.h**

## Prerequisite Information

“Chapter 1. Understanding Event Management” on page 1.

## Related Information

Structures in the **ha\_rmapi.h** header file: **ha\_em\_err\_blk**

Subroutines: **ha\_rr\_makserv**

---

## ha\_rr\_touch Subroutine

### Purpose

**ha\_rr\_touch** – Meet “send” frequency requirements when there are no resource variable values to send to the Resource Monitor Application Programming Interface (RMAPI)

### Library

RMAPI Library (not thread-safe) (**libha\_rr.a**)

### Syntax

```
#include <ha_rmap_i.h>

int
    ha_rr_touch(
        struct ha_em_err_blk  *rr_errb)
```

### Parameters

*rr\_errb* A pointer to an error block structure.

### Description

The **ha\_rr\_touch** subroutine is used by the resource monitor to meet “send” frequency requirements when it has no resource variable values to send to the RMAPI.

If any of the variables that were added to the list of those known to the RMAPI by a resource monitor are of value type Counter or Quantity and they are configured to be supplied to the performance monitor, then every 500 seconds the resource monitor must call either the **ha\_rr\_send\_val** subroutine (to send values) or the **ha\_rr\_touch** subroutine.

### Security

The calling process must have a real or effective group ID of **haemrm**, or must have the **haemrm** group ID in its supplemental group list. If the process is not in the **haemrm** group, its effective user ID must be **root**.

If the calling process is instance 0 of the resource monitor, and the monitor is configured to supply Counter or Quantity variables to the Performance Monitor, it must have an effective user ID of **root**.

### Return Values

If the **ha\_rr\_touch** subroutine is successful, it returns a value of 0.

### Error Values

If the **ha\_rr\_touch** subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

The RMAPI error block and error numbers are defined in the **ha\_rmap\_i.h** header file. For more information on RMAPI errors, see “RMAPI Errors (err\_rmap\_i)” on page 131.

For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

## Examples

For examples of using RMAPI subroutines, see the programs in the RSCT product samples directory, `/usr/sbin/rsct/samples/haem/rmapi`.

## Files

`ha_rmapi.h`

## Prerequisite Information

“Chapter 1. Understanding Event Management” on page 1.

## Related Information

Structures in the `ha_rmapi.h` header file: `ha_em_err_blk`

Subroutines: `ha_rr_send_val`

---

## ha\_rr\_unreg\_var Subroutine

### Purpose

**ha\_rr\_unreg\_var** – Unregister a variable instance with the Resource Monitor Application Programming Interface (RMAPI)

### Library

RMAPI Library (not thread-safe) (**libha\_rr.a**)

### Syntax

```
#include <ha_rmap_i.h>

int
  ha_rr_unreg_var(
    struct ha_rr_variable *pv,
    int numv,
    struct ha_em_err_blk *rr_errb)
```

### Parameters

*pv* A pointer to an array of **ha\_rr\_variable** structures, which contain the names and resource IDs of the resource variable instances to be unregistered.

*numv* The number of elements in the array pointed to by the *pv* parameter.

*rr\_errb* A pointer to an error block structure.

### Description

The **ha\_rr\_unreg\_var** subroutine is used by the resource monitor to unregister resource variable instances with the RMAPI. It is called when the resource variable no longer exists.

If a variable instance has been added to the RMAPI by a call to the **ha\_rr\_add\_var** subroutine, the call to the **ha\_rr\_unreg\_var** subroutine automatically deletes it. If the variable instance is successfully unregistered, the handle returned by the **ha\_rr\_add\_var** subroutine is no longer valid and must not be used in any further calls to the RMAPI.

The **ha\_rr\_variable** structure has the following definition:

```
struct ha_rr_variable {
  char *rr_var_name;
  char *rr_var_rsrc_ID;
  union {
    int rr_var_inst_id;
    void **rr_var_hdl;
  } rr_varu;
#define rr_var_handle rr_varu.rr_var_hdl
#define rr_var_iid rr_varu.rr_var_inst_id
  void *rr_value;
  int rr_var_errno;
}
```

The **rr\_var\_name** field is a pointer to a string that contains the name of a resource variable.

The **rr\_var\_rsrc\_ID** field is a pointer to a string that contains the resource ID of an instance of the specified variable. The resource ID must be fully qualified; no wildcarding is permitted.

The resource ID is specified as a comma-separated list of name/value pairs. A name/value pair consists of a resource ID element name followed by an equal sign followed by the value of the resource ID element. There are no blanks in the resource ID.

The **rr\_varu** field is not used.

The **rr\_value** field is not used.

The **rr\_var\_errno** field is used to hold the result of the operation for each variable.

## Security

The calling process must have a real or effective group ID of **haemrm**, or must have the **haemrm** group ID in its supplemental group list. If the process is not in the **haemrm** group, its effective user ID must be **root**.

If the calling process is instance 0 of the resource monitor, and the monitor is configured to supply Counter or Quantity variables to the Performance Monitor, it must have an effective user ID of **root**.

## Restrictions

If the name of one of the elements in any resource ID matches the resource ID element name specified by the Locator field in the associated resource variable definition, that resource ID element name/value pair is never supplied by the resource monitor in the resource ID specified by the **rr\_var\_rsrc\_ID** field.

The Event Management subsystem starts to cache resource variable instances once they have been registered through the RMAPI. Until the first send of an instance value occurs, the instance assumes the default value from the configuration file. Variable instances are removed from this cache only when they are unregistered; they are not removed if the resource monitor terminates without calling the **ha\_rr\_unreg\_var** subroutine. Therefore, even if a resource monitor terminates, its variable instances are still available for query through the EAPI. Typically, a command-based resource monitor does not call the **ha\_rr\_unreg\_var** subroutine.

The monitor should wait 10 seconds before reusing any instance identifiers of unregistered variables. Variable instance identifiers are supplied to the RMAPI by the resource monitor when variables are registered, in the **rr\_var\_inst\_id** field of **ha\_rr\_variable** structure.

## Return Values

If the **ha\_rr\_unreg\_var** subroutine is successful, it returns the number of resource variables that were unregistered. For each variable that was successfully unregistered, the associated **rr\_var\_errno** field in the **ha\_rr\_variable** structure is set to 0. If a variable could not be unregistered, an error code is placed in its **rr\_var\_errno** field.

## Error Values

If the **ha\_rr\_unreg\_var** subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

The following error number requires specific handling by the resource monitor:

### **HA\_RR\_EDISCONNECT**

This error is only returned to client type resource monitors and indicates that the session created by a call to **ha\_rr\_start\_session** was closed by the Event Management daemon. The resource monitor should respond by calling **ha\_rr\_terminate** to end the current use of the RMAPI.

The RMAPI error block and error numbers are defined in the **ha\_rmapi.h** header file. For more information on RMAPI errors, see “RMAPI Errors (err\_rmapi)” on page 131.

For information about error messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

## Examples

For examples of using RMAPI subroutines, see the programs in the RSCT product samples directory, `/usr/sbin/rsct/samples/haem/rmapi`.

## Files

`ha_rmapi.h`

## Prerequisite Information

“Chapter 1. Understanding Event Management” on page 1.

## Related Information

Structures in the `ha_rmapi.h` header file: `ha_rr_variable`

Subroutines: `ha_rr_reg_var`.

---

## Chapter 4. Event Management Files Reference

This chapter contains the reference material for the error numbers and header files in the Event Management Application Programming Interface (EMAPI) and the Resource Monitor Application Programming Interface (RMAPI). The files are listed alphabetically.

---

## EMAPI Errors (err\_emapi)

### Purpose

EMAPI Errors – Error numbers and the error block for the Event Management Application Programming Interface (EMAPI)

### Description

The Event Management Application Programming Interface (EMAPI) provides several classes of errors. These consist of errors that are returned:

- Synchronously by EMAPI subroutines
- Asynchronously as the result of EMAPI commands.

### Synchronous Errors That Are Returned by EMAPI Subroutines

If an EMAPI subroutine is unsuccessful, it returns a value of -1. It also returns other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

The EMAPI uses an error block that is common to both the EMAPI and the RMAPI. The block is defined in the **ha\_emcommon.h** header file that is included by the **ha\_emapi.h** header file.

The EMAPI error numbers are defined in the **ha\_emapi.h** header file.

The Event Management error block has the following definition:

```
struct ha_em_err_blk {
    int     em_errline;
    char    em_errlevel[HA_EM_MAXERRLVL];
    char    em_errfile[HA_EM_MAXERRFN];
    int     em_errno;
    char    em_errmsg[HA_EM_MAXERRMSG];
}
```

The **em\_errline**, **em\_errlevel**, and **em\_errfile** fields are reserved for IBM use in providing information for problem determination.

If an error occurs, the **em\_errno** field contains one of the values listed below. The **em\_errmsg** field contains a message text that describes the error in further detail. For information about EMAPI messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

The EMAPI error numbers are:

#### **HA\_EM\_EGETDOMAININFO**

The Event Management subsystem could not get information about the domain that is associated with a session.

#### **HA\_EM\_EGETNODNUM**

The Event Management subsystem could not determine the number of the node on which the application is running.

#### **HA\_EM\_ESDROPN**

The Event Management subsystem could not establish a System Data Repository (SDR) session.

#### **HA\_EM\_ESDRGET**

The Event Management subsystem could not obtain required information from the System Data Repository (SDR).

#### **HA\_EM\_ENOMEM**

The Event Management subsystem could not allocate required memory.

**HA\_EM\_ESYSCALL**

The Event Management subsystem received an unexpected error from a system call.

**HA\_EM\_EEXIST**

A connection to the Event Management subsystem was being created using a file descriptor that was already being used for another connection to the Event Management subsystem. This can occur if the EM client closes a file descriptor that is an Event Management session handle. An EM client should not close such a descriptor. When an Event Management session is no longer needed, it should be ended through a call to the **ha\_em\_end\_session** subroutine.

**HA\_EM\_ECONNREFUSED**

The Event Management subsystem refused a requested connection.

This error can be received during a call to either the **ha\_em\_start\_session** or the **ha\_em\_restart\_session** subroutine. It indicates that the Event Management subsystem either is not running, is recovering from a failure, or is initializing itself. The EM client may wish to exit or retry the call repeatedly at some interval until the call is successful. The interval selected depends on the needs of the client application, but it is probably not useful for the interval to be less than 5 seconds.

The EM client may try to establish the connection indefinitely, or it may give up after a number of failed attempts. How persistent the EM client is in attempting to establish the connection depends on the needs of the client application.

**HA\_EM\_ENOCONNECT**

An error occurred while the Event Management subsystem attempted a connection.

**HA\_EM\_ECONNLOST**

The connection with the Event Management subsystem was lost. To reconnect to the Event Management subsystem, use the **ha\_em\_restart\_session** subroutine.

**HA\_EM\_EAUTHENT**

The Event Management subsystem could not authenticate the user. The user of the EMAPI may need to log in to DCE (Distributed Computing Environment).

**HA\_EM\_EAUTHOR**

The user is not authorized to use the Event Management subsystem. The user of the EMAPI must be a member of the Event Management DCE group.

**HA\_EM\_ESECSERV**

An error occurred while requesting a security service.

**HA\_EM\_ENOENT**

The specified session does not exist.

**HA\_EM\_EBUSY**

The specified session is busy. This can occur when multiple threads are attempting to use the same session at the same time. It is recommended that an Event Management session be used by only one thread.

**HA\_EM\_EINVALID**

An invalid parameter was specified.

**HA\_EM\_ENOTOWNER**

The process does not own the specified session.

**HA\_EM\_EUNEXPECTED**

The Event Management subsystem encountered an unexpected condition.

**HA\_EM\_EUNUSABLE**

The specified session has become unusable. The session may have become unusable because of

an error previously reported for the session. In a multi-threaded program, a session can also become unusable in the following circumstance: A thread that is executing an EAPI subroutine is canceled with the `pthread_cancel` system call.

#### **HA\_EM\_ECONNECTED**

A restart of the specified session is not permitted, because the session's connection with the Event Management subsystem has not been lost.

#### **HA\_EM\_ENOSECCLIENT**

The system on which the EAPI client is running does not have DCE security configured, and the client is connected to an Event Management subsystem that requires the use of DCE security.

#### **HA\_EM\_ENOSECDAEMON**

The system on which the EAPI client is running requires the use of DCE security, and the client is connected to an Event Management subsystem that does not have DCE security configured.

#### **HA\_EM\_EAUTHNTDAE**

The EAPI could not authenticate the Event Management subsystem.

### **Asynchronous Errors That Result from EAPI Commands**

When an EAPI command causes an asynchronous error, the error is indicated as follows. The Event Management subsystem sets the `em_error` field of one of the structures specified in the `em_resp_blk` union of the `ha_em_rsp_blk` structure to a nonzero value. It also passes two error codes in the `em_error_codes` array. The first element of the array contains a general error code that indicates the nature of the operation that resulted in the error. The second element indicates in more detail the actual error.

The general error codes are:

#### **HA\_EM\_RSP\_EGEN\_COMMAND**

An error was detected in the specification of a command.

#### **HA\_EM\_RSP\_EGEN\_RESOURCE**

An error was caused because an internal Event Management subsystem resource was temporarily unavailable.

#### **HA\_EM\_RSP\_EGEN\_RESMON**

An error resulted from starting or communicating with a resource monitor.

#### **HA\_EM\_RSP\_EGEN\_EVALERR**

An error was detected during the evaluation of an expression.

Each general error code can have a specific error code associated with it.

Unless otherwise stated, if any of these specific error codes are returned in the `em_error` field of the `ha_em_rpb_rerr` or `ha_em_rpb_qerr` structures, then if the value of the `em_errinfo1` field in the `ha_em_rpb_rerr` structure or the value of the `em_errinfo` field in the `ha_em_rpb_qerr` structure is nonzero, the value has the following meaning. The nonzero value represents the position of the error within the input string that is in error, as specified by the specific error code.

The specific error codes associated with each general error code are listed below.

### **The HA\_EM\_RSP\_EGEN\_COMMAND General Error Code**

Specific error codes associated with the `HA_EM_RSP_EGEN_COMMAND` general error code are:

#### **HA\_EM\_RSP\_ENOVARENT**

The specified variable name does not exist in the Event Management Configuration Database (EMCDB).

**HA\_EM\_RSP\_EEXPR**

The specified expression could not be parsed correctly. In the **ha\_em\_rpb\_rerr** structure, the expression parsing error is contained in the **em\_errinfo0** field, but the position of the error within the expression is contained in the **em\_errinfo1** field.

For a list of parsing errors, see the expression parsing error section later in this man page.

**HA\_EM\_RSP\_ERAEXPR**

The specified rearm expression could not be parsed correctly. In the **ha\_em\_rpb\_rerr** structure, the expression parsing error is contained in the **em\_errinfo0** field and the position of the error within the expression is contained in the **em\_errinfo1** field.

For a list of parsing errors, see the expression parsing error section later in this man page.

**HA\_EM\_RSP\_ENODFLTEXPR**

No expression was specified, and no default expression is configured.

**HA\_EM\_RSP\_ENORIDNAME**

The specified resource ID was missing a resource ID element name.

**HA\_EM\_RSP\_ENORIDVALUE**

The specified resource ID was missing a resource ID element value.

**HA\_EM\_RSP\_ERIDSYNTAX**

The specified resource ID contained a syntax error.

**HA\_EM\_RSP\_ERIDVALLENGTH**

A resource ID element-value length, including the terminating null character, was longer than the maximum allowed (**HA\_EM\_RSRC\_ID\_ELEM\_VALUE\_LEN\_MAX**).

**HA\_EM\_RSP\_EMISSINGRID**

The specified resource ID was missing a required element.

**HA\_EM\_RSP\_EDUPRIDENT**

The specified resource ID contained a duplicate element.

**HA\_EM\_RSP\_ENORIDENT**

The specified resource ID contained an element that is not defined for the specified variable.

**HA\_EM\_RSP\_EINVALEID**

An invalid event ID was specified.

**HA\_EM\_RSP\_ECLASSMISMATCH**

The specified class does not match the class defined for the specified variable.

**HA\_EM\_RSP\_ERIDMISMATCH**

The specified resource ID does not match the resource ID defined for the specified variable or variables.

**HA\_EM\_RSP\_EWCRMISMATCH**

The class or resource ID, (or both) if specified, do not match the class or resource ID (or both) defined for any of the variables that matched the specified wildcarded variable name.

**HA\_EM\_RSP\_ENOCLASS**

The specified class name does not exist in the Event Management Configuration Database (EMCDB).

**HA\_EM\_RSP\_ERIDTOOLONG**

The specified resource ID string is too long.

**HA\_EM\_RSP\_EINVALNODENUM**

A node number implied in a resource ID is not within the range supported by the Event Management subsystem.

**HA\_EM\_RSP\_EINVALRIDVAL**

A resource ID element value of type integer is invalid. It is either less than 0, greater than 4096 or, if a range is specified, the second value is less than the first.

**Expression Parsing Error Codes**

Expression parsing error codes associated with the **HA\_EM\_RSP\_EEXPR**, **HA\_EM\_RSP\_ERAEXPR** and **HA\_EM\_RSP\_EGEN\_EVALERR** error codes are:

**HA\_EM\_EXPR\_ERVAR**

There is an internal resource-variable error. Contact the IBM Support Center.

**HA\_EM\_EXPR\_EEXPRLEN**

The parsing table overflowed.

**HA\_EM\_EXPR\_ESYNTAX**

There is an error in the syntax of the expression.

**HA\_EM\_EXPR\_ESTACK**

An internal stack overflowed. Contact the IBM Support Center.

**HA\_EM\_EXPR\_EILLEGALOP**

The specified operation is illegal between the specified data types.

**HA\_EM\_EXPR\_ESTRING**

There is an error in the syntax of a string constant.

**HA\_EM\_EXPR\_ESBSSN**

The serial number of a structured byte string field is either missing or out of range.

**HA\_EM\_EXPR\_EFIELDTYPE**

The data type of a structured byte string field is either illegal or does not match its value.

**HA\_EM\_EXPR\_EDATATYPE**

There is an internal data type error. Contact the IBM Support Center.

**HA\_EM\_EXPR\_ESBSLEN**

There is a mismatch between the data type and the length of structured byte string field.

**HA\_EM\_EXPR\_EDIVIDE**

An attempt was made to divide by zero.

**HA\_EM\_EXPR\_EOCTAL\_INVALID**

An octal value was not valid.

**HA\_EM\_EXPR\_EOCTAL\_2BIG**

An octal value was greater than the maximum allowed.

**The HA\_EM\_RSP\_EGEN\_RESOURCE General Error Code**

Specific error codes associated with the **HA\_EM\_RSP\_EGEN\_RESOURCE** general error code are:

**HA\_EM\_RSP\_EQNONE**

The Event Management subsystem could not generate a response for this query.

**HA\_EM\_RSP\_EQCOMPLETE**

The query is complete. All information was returned in previous responses.

**HA\_EM\_RSP\_EINSTNOTAVAIL**

The variable instance was not available.

**HA\_EM\_RSP\_EINSTSTALE**

The variable instance is stale; that is, the resource monitor that supplies the instance has terminated unexpectedly.

**HA\_EM\_RSP\_EDAEMONGONE**

The Event Management daemon that supplies instances that match the resource ID has terminated.

**HA\_EM\_RSP\_ESHMEM**

The variable instance value could not be obtained due to a shared memory error.

**The HA\_EM\_RSP\_EGEN\_RESMON General Error Code**

Specific error codes associated with the **HA\_EM\_RSP\_EGEN\_RESMON** general error code are:

**HA\_EM\_RSP\_ERMLOCKSTATUS**

An error was encountered while trying to obtain the lock status of the resource monitor.

**HA\_EM\_RSP\_ERMNOTRUNNING**

The resource monitor is not running, and it cannot be started by the Event Management subsystem.

**HA\_EM\_RSP\_ENOSTARTRM**

The resource monitor could not be started by the Event Management subsystem.

**HA\_EM\_RSP\_ENORMCONNECT**

The Event Management subsystem could not connect to the resource monitor.

**The HA\_EM\_RSP\_EGEN\_EVALERR General Error Code**

The specific error codes associated with the **HA\_EM\_RSP\_EGEN\_EVALERR** general error code are listed in the expression parsing error section elsewhere in this man page.

**Related Information**

EMAPI subroutines: **ha\_em\_end\_session**, **ha\_em\_receive\_response**, **ha\_em\_restart\_session**, **ha\_em\_send\_command**, **ha\_em\_start\_session**

EMAPI header file: **ha\_emoji.h**

Messages: *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*

---

## ha\_emapi.h File

### Purpose

**ha\_emapi.h** – Header file for the Event Management Application Programming Interface (EMAPI)

### Description

The **ha\_emapi.h** header file provides data types and structures for use with the Event Management Application Programming Interface (EMAPI) subroutines. The EMAPI provides two libraries, one that is thread-safe (**libha\_em\_r.a**) and one that is not (**libha\_em.a**). Any program that uses the EMAPI subroutines must include this file, which resides in the **/usr/include** directory.

The following listing shows the contents of the **ha\_emapi.h** file. Nested header files, **ha\_emapi\_base.h** and **ha\_emcommon.h**, follow the **ha\_emapi.h** file listing.

#### ha\_emapi.h File

```
/* IBM_PROLOG_BEGIN_TAG                               */
/* This is an automatically generated prolog.         */
/*                                                    */
/*                                                    */
/* Licensed Materials - Property of IBM                */
/*                                                    */
/* (C) COPYRIGHT International Business Machines Corp. 1996,1999 */
/* All Rights Reserved                                */
/*                                                    */
/* US Government Users Restricted Rights - Use, duplication or */
/* disclosure restricted by GSA ADP Schedule Contract with IBM Corp. */
/*                                                    */
/* IBM_PROLOG_END_TAG                                 */

/*=====*/
/*                                                    */
/* Module Name:  ha_emapi.h                            */
/*                                                    */
/* Description:  */
/*                                                    */
/*   Interface definitions of the Event Manager API needed by */
/*   Event Manager API clients and the Event Manager API library. */
/*                                                    */
/*   This file is formatted to be viewed with tab stops set to 4. */
/*=====*/

/* @(#)41  1.13  src/rsct/pem/emcommon/ha_emapi.h, emcommon, rsct_rmoh, rmoht3f5 1/25/99 10:13:58 */

#ifndef _HA_EMAPI_H
#define _HA_EMAPI_H

/*-----*/
/* Event Manager API Versioning.                       */
/*-----*/

/*
 * HA_EM_VERSION identifies the version of the Event Manager Application
 * Programming Interface desired by a client.  If the client has not
 * specified a particular version, version 1 is used.  The current valid
 * versions are 1 and 2.
 */

#ifndef HA_EM_VERSION
#define HA_EM_VERSION 1
#endif /* HA_EM_VERSION */

#if (HA_EM_VERSION < 1) || (HA_EM_VERSION > 2)
#error "HA_EM_VERSION MUST BE AN INTEGER BETWEEN 1 AND 2, INCLUSIVE."
#endif

/*
 * The following macros convert references to the generic names of EMAPI
 * routines, like ha_em_start_session(), to references to version specific
 * EMAPI routines, like ha_em_start_session_1().  The version used during
 * conversion is determined by the value of HA_EM_VERSION.  Note that these
 * macros can deal with the different versions of a routine having different
 * arguments.
 */

#define ha_em_start_session \
    HA_EM_ROUTINE_VERSION(ha_em_start_session, HA_EM_VERSION)
```

```

#define ha_em_restart_session          \
    HA_EM_ROUTINE_VERSION(ha_em_restart_session, HA_EM_VERSION)

#define ha_em_end_session              \
    HA_EM_ROUTINE_VERSION(ha_em_end_session, HA_EM_VERSION)

#define ha_em_send_command            \
    HA_EM_ROUTINE_VERSION(ha_em_send_command, HA_EM_VERSION)

#define ha_em_receive_response        \
    HA_EM_ROUTINE_VERSION(ha_em_receive_response, HA_EM_VERSION)

#define ha_em_get_ecgid              \
    HA_EM_ROUTINE_VERSION(ha_em_get_ecgid, HA_EM_VERSION)

#define HA_EM_ROUTINE_VERSION(routine, version) \
    HA_EM_ROUTINE_VERSION_GLUE(routine, version)

#define HA_EM_ROUTINE_VERSION_GLUE(routine, version) \
    routine ## _ ## version

/*-----*/
/* Event Manager API Include Files.                */
/*-----*/

#include <ha_emoji_base.h>

/*-----*/
/* Event Manager API constants.                    */
/*-----*/

/*
 * Event Manager error values.
 */

#if (HA_EM_VERSION == 2) || defined(HA_EM_ALL_VERSIONS)
#define HA_EM_EGETDOMAININFO 1 /* Could not get domain information */
#endif

#if (HA_EM_VERSION == 1) || defined(HA_EM_ALL_VERSIONS)
#define HA_EM_EGETPARTINFO 1 /* Could not get partition information */
#endif

#define HA_EM_EGETNODNUM 2 /* Could not determine host node number */
#define HA_EM_ESDROPN 3 /* Could not establish an SDR session */
#define HA_EM_ESDRGET 4 /* Could not obtain needed info from SDR */
#define HA_EM_ENOMEM 5 /* Could not allocate needed memory */
#define HA_EM_ESYSCALL 6 /* Unexpected system call error */
#define HA_EM_EEXIST 7 /* Unexpected re-use of file descriptor */
#define HA_EM_ECONNREFUSED 8 /* Connection to E.M. daemon refused */
#define HA_EM_ENOCONNECT 9 /* Connection error to E.M. daemon */
#define HA_EM_ECONNLOST 10 /* Connection lost with E.M. daemon */
#define HA_EM_EAUTHENT 11 /* Could not authenticate user */
#define HA_EM_EAUTHOR 12 /* User not authorized to use E.M. daemon */
#define HA_EM_ESECSERV 13 /* Security service error encountered */
#define HA_EM_ENOENT 14 /* Specified session does not exist */
#define HA_EM_EBUSY 15 /* Specified session is busy */
#define HA_EM_EINVAL 16 /* Invalid parameter specified */
#define HA_EM_ENOTOWNER 17 /* The process does not own the session */
#define HA_EM_EUNEXPECTED 18 /* Unexpected condition occurred */
#define HA_EM_EUNUSABLE 19 /* Specified session has become unusable */
#define HA_EM_ECONNECTED 20 /* Session restart not permitted */
#define HA_EM_ENOSECDAEMON 21 /* E.M. daemon does not support security */
#define HA_EM_ENOSECCLIENT 22 /* Client does not support security */
#define HA_EM_EAUTHENTDAE 23 /* E.M. daemon could not authenticate daemon */

/*
 * Event Manager Domain Type Values.
 */

#if (HA_EM_VERSION == 2) || defined(HA_EM_ALL_VERSIONS)
#define HA_EM_DOMAIN_SP 1 /* The SP domain */
#define HA_EM_DOMAIN_HACMP 2 /* The HACMP domain */
#endif

/*-----*/
/* Event Manager API type and structure definitions. */
/*-----*/

/*
 * One event registration request is made through a ha_em_rb_reg
 * (Event Manager Request Block: REGISTER) structure. This structure is
 * included in the ha_em_res_blk union, defined later in this file.
 */

struct ha_em_rb_reg {
    char *em_name; /* input: resource variable name */
    char *em_rsrc_ID; /* input: resource ID */
    char *em_expr; /* input: expression */
    char *em_raexpr; /* input: re-arm expression */
};

```

```

    ha_em_eid_t em_event_id; /* output: event identifier */
    void (*em_cb)(int, struct ha_em_rpb_event *, void *);
    void *em_cb_arg; /* input: callback routine address */
    void *em_cb_arg; /* input: callback routine argument */
};

/*
 * One query request is made through a ha_em_rb_query
 * (Event Manager Request Block: QUERY) structure. This structure is
 * actually defined in ha_emapi_base.h. It is included in the ha_em_res_blk
 * union, defined later in this file.
 */

struct ha_em_rb_query;

/*
 * Multiple requests of a particular type are made through a ha_em_res_blk
 * (Event Manager REQUEST BLOCK) union. All the requests must be of the
 * same type. The type of array used depends on the type of the requests
 * being made. The number of elements in the array depends on the number
 * of requests being made. This union is included in the ha_em_cmd_blk
 * structure, defined later in this file.
 */

union ha_em_res_blk {
    struct ha_em_rb_reg em_rb_reg[1]; /* Used to register for events */
    ha_em_eid_t em_rb_unreg[1]; /* Used to unregister events */
    struct ha_em_rb_query em_rb_query[1]; /* Used to query for information */
};

/*
 * A block of requests of a particular type are made through a ha_em_cmd_blk
 * (Event Manager COMMAND BLOCK) structure passed to the ha_em_send_command()
 * routine. All the requests must be of the same type. The type of the
 * requests are identified in the em_cmd and em_subcmd fields. The requests
 * themselves are in an array within the ha_em_res_blk union.
 */

struct ha_em_cmd_blk {
    int em_cmd_num_elem; /* input: number of requests */
    short em_cmd; /* input: type of requests */
    short em_subcmd; /* input: subtype of requests */
    ha_emqid_t em_qid; /* output: query identifier */
    void (*em_qcb)(int, struct ha_em_rsp_blk *, void *);
    /* input: query callback routine */
    void *em_qcb_arg; /* input: query callback arg. */
    union ha_em_res_blk em_res_blk; /* the array of requests */
};

/*
 * A block of responses of a particular type are delivered through a
 * ha_em_rsp_blk (Event Manager RESPONSE BLOCK) structure passed through
 * the ha_em_receive_command() routine. This structure is
 * actually defined in ha_emapi_base.h.
 */

struct ha_em_rsp_blk;

/*-----*/
/* Event Manager API compatibility definitions. */
/*
 * These definitions are provided to maintain source compatibility with
 * programs written using prior versions of this header file.
 * If these definitions result in inappropriate substitutions, then define
 * the symbol HA_EM_NO_NAME_COMPAT prior to inclusion of this header file.
 * If HA_EM_NO_NAME_COMPAT is defined, the source files that include this
 * header file must be modified to use the new symbol names if the old
 * symbol names are referenced therein.
 */
/*-----*/

#ifndef HA_EM_NO_NAME_COMPAT
#ifndef em_ivector
#define em_ivector em_rsrc_ID /* replace em_ivector by em_rsrc_ID */
#endif
#ifndef em_pred
#define em_pred em_expr /* replace em_pred by em_expr */
#endif
#ifndef em_rapred
#define em_rapred em_raexpr /* replace em_rapred by em_raexpr */
#endif
#endif /* HA_EM_NO_NAME_COMPAT */

/*-----*/
/* Event Manager API function prototypes - Version 2 */
/*-----*/

/*
 * Function prototypes for functions intended for Event Manager API clients.
 * Clients should use the generic routine names defined by macros included
 * in this header file, such as ha_em_start_session(), instead of the version

```

```

* specific routine names, like ha_em_start_session_2(). Those macros
* convert the generic names to the appropriate version specific names.
*/

#if ((HA_EM_VERSION == 2) || defined(HA_EM_ALL_VERSIONS)) && !defined(_NO_PROTO)

extern int ha_em_start_session_2( /* Start an Event Manager session */
    int em_domain_type, /* input: session domain type */
    char *em_domain_name, /* input: session domain name */
    struct ha_em_err_blk *em_errb /* output: error description block */
); /* return: session ID or -1 (error) */

extern int ha_em_restart_session_2( /* Restart an Event Manager session */
    int em_session_fd, /* input: old session ID */
    struct ha_em_err_blk *em_errb /* output: error description block */
); /* return: new ID or -1 (error) */

extern int ha_em_end_session_2( /* End an Event Manager session */
    int em_session_fd, /* input: Event Manager session ID */
    struct ha_em_err_blk *em_errb /* output: error description block */
); /* return: 0 or -1 (error) */

extern int ha_em_send_command_2( /* Send an Event Manager command */
    int em_session_fd, /* input: Event Manager session ID */
    struct ha_em_cmd_blk *em_cmb, /* input/output: Command block */
    struct ha_em_err_blk *em_errb /* output: error description block */
); /* return: 0 or -1 (error) */

extern int ha_em_receive_response_2( /* Receive an Event Manager response */
    int em_session_fd, /* input: Event Manager session ID */
    struct ha_em_rsp_blk **em_rsp_blk, /* output: pointer to response block */
    struct ha_em_err_blk *em_errb /* output: error description block */
); /* return: number of responses */
/* (may be 0), or -1 (error) */

extern ha_em_ecgid_t ha_em_get_ecgid_2( /* Get event command group identifier*/
    ha_em_eid_t em_eid /* input: event identifier */
); /* return: event command group ID */

#endif

#if ((HA_EM_VERSION == 2) || defined(HA_EM_ALL_VERSIONS)) && defined(_NO_PROTO)
extern int ha_em_start_session_2(); /* Start an Event Manager session */
extern int ha_em_restart_session_2(); /* Restart an Event Manager session */
extern int ha_em_end_session_2(); /* End an Event Manager session */
extern int ha_em_send_command_2(); /* Send an Event Manager command */
extern int ha_em_receive_response_2(); /* Receive an Event Manager response */
extern ha_em_ecgid_t ha_em_get_ecgid_2(); /* Get event command group ID */
#endif

/*-----*/
/* Event Manager API function prototypes - Version 1 */
/*-----*/

/*
* Function prototypes for functions intended for Event Manager API clients.
* Clients should use the generic routine names defined by macros included
* in this header file, such as ha_em_start_session(), instead of the version
* specific routine names, like ha_em_start_session_1(). Those macros
* convert the generic names to the appropriate version specific names.
*/

#if ((HA_EM_VERSION == 1) || defined(HA_EM_ALL_VERSIONS)) && !defined(_NO_PROTO)

extern int ha_em_start_session_1( /* Start an Event Manager session */
    char *em_part_name, /* input: SP partition of session */
    struct ha_em_err_blk *em_errb /* output: error description block */
); /* return: session ID or -1 (error) */

extern int ha_em_restart_session_1( /* Restart an Event Manager session */
    int em_session_fd, /* input: old session ID */
    struct ha_em_err_blk *em_errb /* output: error description block */
); /* return: new ID or -1 (error) */

extern int ha_em_end_session_1( /* End an Event Manager session */
    int em_session_fd, /* input: Event Manager session ID */
    struct ha_em_err_blk *em_errb /* output: error description block */
); /* return: 0 or -1 (error) */

extern int ha_em_send_command_1( /* Send an Event Manager command */
    int em_session_fd, /* input: Event Manager session ID */
    struct ha_em_cmd_blk *em_cmb, /* input/output: Command block */
    struct ha_em_err_blk *em_errb /* output: error description block */
); /* return: 0 or -1 (error) */

extern int ha_em_receive_response_1( /* Receive an Event Manager response */
    int em_session_fd, /* input: Event Manager session ID */
    struct ha_em_rsp_blk **em_rsp_blk, /* output: pointer to response block */
    struct ha_em_err_blk *em_errb /* output: error description block */
); /* return: number of responses */
/* (may be 0), or -1 (error) */

```

```

extern ha_em_ecgid_t ha_em_get_ecgid_1( /* Get event command group identifier*/
    ha_em_eid_t em_eid /* input: event identifier */
); /* return: event command group ID */

#endif

#if ((HA_EM_VERSION == 1) || defined(HA_EM_ALL_VERSIONS)) && defined(NO_PROTO)
extern int ha_em_start_session_1(); /* Start an Event Manager session */
extern int ha_em_restart_session_1(); /* Restart an Event Manager session */
extern int ha_em_end_session_1(); /* End an Event Manager session */
extern int ha_em_send_command_1(); /* Send an Event Manager command */
extern int ha_em_receive_response_1(); /* Receive an Event Manager response */
extern ha_em_ecgid_t ha_em_get_ecgid_1(); /* Get event command group ID */
#endif

#endif /* _HA_EMAPI_H */

```

## ha\_emapi\_base.h File

```

/* IBM_PROLOG_BEGIN_TAG */
/* This is an automatically generated prolog. */
/* */
/* */
/* Licensed Materials - Property of IBM */
/* */
/* (C) COPYRIGHT International Business Machines Corp. 1996,1999 */
/* All Rights Reserved */
/* */
/* US Government Users Restricted Rights - Use, duplication or */
/* disclosure restricted by GSA ADP Schedule Contract with IBM Corp. */
/* */
/* IBM_PROLOG_END_TAG */
/* vi:set tabstop=4: */
/*=====*/
/* Module Name: ha_emapi_base.h */
/* Description: */
/* Interface definitions of the Event Manager API needed by */
/* Event Manager API clients, the Event Manager API library, and the */
/* Event Manager daemon. */
/* This file is formatted to be viewed with tab stops set to 4. */
/*=====*/

/* @(#)42 1.24 src/rsct/pem/emcommon/ha_emapi_base.h, emcommon, rsct_rmoh, rmoht3f5 1/25/99 10:17:03 */

#ifndef _HA_EMAPI_BASE_H
#define _HA_EMAPI_BASE_H

#include <sys/time.h>
#include <ha_emcommon.h>

/*-----*/
/* Event Manager API constants. */
/*-----*/

/*
 * Event Manager Commands.
 */

#define HA_EM_CMD_REG 1 /* Register for events */
#define HA_EM_CMD_REG2 2 /* Register for events from both expressions */
#define HA_EM_CMD_UNREG 3 /* Unregister events */
#define HA_EM_CMD_QUERY 4 /* Query for information */

#define HA_EM_CMD_RERR 5 /* Returned for registration command errors */
#define HA_EM_CMD_R2ERR 6 /* Returned for registration 2 command errors*/
#define HA_EM_CMD_QERR 7 /* Returned for query command errors */

/*
 * Event Manager Subcommands for the HA_EM_CMD_REG and HA_EM_CMD_REG2 commands.
 * The HA_EM_SCMD_REVAL and HA_EM_SCMD_RACK subcommands may be bitwise OR'ed
 * if both functions are desired.
 */

#define HA_EM_SCMD_REVAL 1 /* Evaluate expression at first observation */
/* and return result as an event */
#define HA_EM_SCMD_RACK 2 /* Return 0 error code if registration req */
/* has no errors, as part of */
/* HA_EM_CMD_RERR or HA_EM_CMD_R2ERR */
/* response */

/*
 * Event Manager Subcommands for the HA_EM_CMD_QUERY command. These are also
 * returned with HA_EM_CMD_QERR to indicate the query subcommand that
 * resulted in errors.
 */

```

```

#define HA_EM_SCMD_QCUR    1 /* Query current resource variable values */
#define HA_EM_SCMD_QDEF    2 /* Query for defined resource variables and */
                          /* their default expressions */
#define HA_EM_SCMD_QINST  3 /* Query for resource variable instances in */
                          /* order to obtain their resource IDs */
                          /* (the returned values are last known) */

/*
 * Event Manager Error Codes, returned in em_error of the various response
 * structures. General error codes indicate the nature of the operation that
 * resulted in the error. Specific error codes detail the actual error.
 */

#define HA_EM_RSP_EGEN_COMMAND 1 /* This general error code indicates an */
                              /* error was detected in the */
                              /* specification of a command */
/*-----*/
/* Specific error codes associated with HA_EM_RSP_EGEN_COMMAND */
/*-----*/
#define HA_EM_RSP_ENOVARENT 1 /* The specified variable name does not */
                              /* exist in the EM data base */
#define HA_EM_RSP_EEXPR    2 /* The specified expression could not be */
                              /* correctly parsed. The expression */
                              /* parsing error is contained in the */
                              /* em_errinfo0 field of the */
                              /* ha_em_rpb_rerr structure; the */
                              /* position of the error within the */
                              /* expression is contained in the */
                              /* em_errinfo1 field of the */
                              /* ha_em_rpb_rerr structure. */
#define HA_EM_RSP_ERAEEXPR 3 /* The specified re-arm expression could */
                              /* not be correctly parsed. The returned */
                              /* value is encoded as for the */
                              /* HA_EM_RSP_EEXPR error. */
#define HA_EM_RSP_ENODFLTEXPR 4 /* No expression was specified and no */
                              /* default expression is configured */
#define HA_EM_RSP_ENORIDNAME 5 /* the specified resource ID was missing */
                              /* a resource ID element name */
#define HA_EM_RSP_ENORIDVALUE 6 /* the specified resource ID was missing */
                              /* a resource ID element value */
#define HA_EM_RSP_ERIDSYNTAX 7 /* the specified resource ID contained */
                              /* a syntax error. */
#define HA_EM_RSP_ERIDVLENGTH 8 /* a resource ID value length was longer */
                              /* including the terminating NULL, than */
                              /* HA_EM_RSRC_ID_ELEM_VALUE_LEN_MAX */
#define HA_EM_RSP_EMISSINGRID 9 /* the specified resource ID was missing */
                              /* a required element. */
#define HA_EM_RSP_EDUPRIDENT 10 /* the specified resource ID contained */
                              /* a duplicate element. */
#define HA_EM_RSP_ENORIDENT 11 /* the specified resource ID contained */
                              /* an element that is not defined for */
                              /* the specified variable. */
#define HA_EM_RSP_EINVALID 12 /* An invalid event ID was specified. */
#define HA_EM_RSP_ECLASSMISMATCH 13 /* The specified class does not */
                              /* match the class defined for the */
                              /* specified variable. */
#define HA_EM_RSP_ERIDMISMATCH 14 /* The specified resource ID does not */
                              /* match the resource ID defined for */
                              /* the specified variable or variables. */
#define HA_EM_RSP_EWCRMISMATCH 15 /* The class, if specified, and/or the */
                              /* resource ID, if specified, do not */
                              /* match the class and/or resource ID */
                              /* defined for any of the variables that */
                              /* matched the specified wildcarded */
                              /* variable name. */
#define HA_EM_RSP_ENOCLASS 16 /* The specified class name does not */
                              /* exist in the EM data base */
#define HA_EM_RSP_ERIDTOOLONG 17 /* The specified resource ID string is */
                              /* too long. */
#define HA_EM_RSP_EINVALNODENUM 18 /* A node number implied in a resource */
                              /* ID is not within the range supported */
                              /* by the Event Manager */
#define HA_EM_RSP_EINVALRIDVAL 19 /* A resource ID value of type 'int' is */
                              /* invalid, i.e. less than 0, greater */
                              /* than 4096 or, if the second value in */
                              /* a value range, it is less than the */
                              /* first value in the range */
#define HA_EM_RSP_EPRED HA_EM_RSP_EEXPR /* for compatibility */
#define HA_EM_RSP_ERAPRED HA_EM_RSP_ERAEEXPR /* for compatibility */
#define HA_EM_RSP_ENODFLT_PRED HA_EM_RSP_ENODFLTEXPR /* for compatibility */
#define HA_EM_RSP_ENOVECNAME HA_EM_RSP_ENORIDNAME /* for compatibility */
#define HA_EM_RSP_ENOVECVALUE HA_EM_RSP_ENORIDVALUE /* for compatibility */
#define HA_EM_RSP_EVECSYNTAX HA_EM_RSP_ERIDSYNTAX /* for compatibility */
#define HA_EM_RSP_EVECVLENGTH HA_EM_RSP_ERIDVLENGTH /* for compatibility */
#define HA_EM_RSP_EMISSINGVEC HA_EM_RSP_EMISSINGRID /* for compatibility */
#define HA_EM_RSP_EDUPVECENT HA_EM_RSP_EDUPRIDENT /* for compatibility */
#define HA_EM_RSP_ENOVECENT HA_EM_RSP_ENORIDENT /* for compatibility */
#define HA_EM_RSP_EVECMISMATCH HA_EM_RSP_ERIDMISMATCH /* for compatibility */
#define HA_EM_RSP_EWCRMISMATCH HA_EM_RSP_EWCRMISMATCH /* for compatibility */
#define HA_EM_RSP_EVECTOOLONG HA_EM_RSP_ERIDTOOLONG /* for compatibility */
#define HA_EM_RSP_EINVALVECVAL HA_EM_RSP_EINVALRIDVAL /* for compatibility */

```

```

#define HA_EM_RSP_EGEN_RESOURCE 2 /* This general error code indicates an */
/* error was due to the temporary */
/* unavailability of an internal Event */
/* Manager resource. */
/*-----*/
/* Specific error codes associated with HA_EM_RSP_EGEN_RESOURCE */
/*-----*/
#define HA_EM_RSP_EQNONE 1 /* No response could be generated for */
/* this individual query request. */
#define HA_EM_RSP_EQCOMPLETE 2 /* The query is complete; all data was */
/* returned in previous message(es) */
#define HA_EM_RSP_EINSTNOTAVAIL 3 /* Variable instance not available */
#define HA_EM_RSP_EINSTSTALE 4 /* Variable instance is stale, i.e. the */
/* resource monitor supplying the */
/* instance has terminated unexpectedly. */
#define HA_EM_RSP_EDAEMONGONE 5 /* the Event Manager daemon supplying */
/* instances matching the resource ID */
/* has terminated. */
#define HA_EM_RSP_ESHMEM 6 /* Variable instance value cannot be */
/* obtained due to a shared memory error */

#define HA_EM_RSP_EGEN_RESMON 3 /* This general error code indicates an */
/* error resulted from starting or */
/* communicating with a resource monitor */
/*-----*/
/* Specific error codes associated with HA_EM_RSP_EGEN_RESMON */
/*-----*/
#define HA_EM_RSP_ERMLOCKSTATUS 1 /* Error in obtaining resource monitor */
/* lock status. */
#define HA_EM_RSP_ERMNOTRUNNING 2 /* a resource monitor is not running and */
/* it cannot be started by Event Mgt */
#define HA_EM_RSP_ENOSTARTRM 3 /* a resource monitor could not be */
/* started by Event Management */
#define HA_EM_RSP_ENORMCONNECT 4 /* the Event Management daemon could not */
/* connect to a resource monitor. */

#define HA_EM_RSP_EGEN_EVALERR 4 /* This general error code indicates an */
/* error was detected in the evaluation */
/* of an expression */
/*-----*/
/* Specific error codes associated with HA_EM_RSP_EGEN_EVALERR are defined */
/* in ha_emcommon.h using names of the form HA_EM_EXPR_* */
/*-----*/

#define HA_EM_RSP_EGEN_SEC 5 /* This general error code indicates a */
/* security error was detected when */
/* starting/restarting a session. An */
/* EMAPI client does not see these error */
/* codes. An EMAPI client sees security */
/* errors through error codes returned */
/* by EMAPI routines. Those codes are */
/* defined in ha_emapi.h */
/*-----*/
/* Specific error codes associated with HA_EM_RSP_EGEN_SEC */
/*-----*/
#define HA_EM_RSP_ENODCECLIENT 1 /* DCE is not configured as a security */
/* method in the client but it is */
/* required by the daemon */
#define HA_EM_RSP_ENODCEDAEMON 2 /* DCE is not configured as a security */
/* method in the daemon but it is */
/* required by the client */
#define HA_EM_RSP_ENOAUTHENT 3 /* The client cannot be authenticated */
#define HA_EM_RSP_ENOAUTHOR 4 /* The client is not authorized */
#define HA_EM_RSP_ESECFUNC 5 /* Error detected in security function */

/*-----*/
/* Event Manager API type and structure definitions. */
/*-----*/

/*
 * Event command group identifiers are of type ha_em_ecgid_t
 * (Event Manager Event Command Group Identifier).
 */
typedef unsigned long ha_em_ecgid_t;

/*
 * Event identifiers are of type ha_em_eid_t (Event Manager Event Identifier).
 */
typedef unsigned long ha_em_eid_t;

```

```

/*
 * Query identifiers are of type ha_em_qid_t (Event Manager Query Identifier).
 */
typedef unsigned long ha_em_qid_t;

/*
 * One query request is made through a ha_em_rb_query
 * (Event Manager Request Block: QUERY) structure. This structure is
 * included in the ha_em_res_blk union, defined in ha_emapi.h.
 */
struct ha_em_rb_query {
    char    *em_class;          /* input: resource variable class name */
    char    *em_name;          /* input: resource variable name */
    char    *em_rsrc_ID;       /* input: resource ID */
};

/*
 * Resource variables can have one of three types of values, signed long
 * integer, float, or structured byte string. The ha_em_val (Event Manager
 * VALue) union reflects these possibilities.
 */
union ha_em_val {
    long    em_val;           /* output: long integer value */
    float   em_valf;         /* output: floating point value */
    void    *em_valsbs;      /* output: structured byte string value */
};

/*
 * Event Manager error codes consist of a general error code and a specific
 * error code, encoded in the ha_em_errnum union.
 */
union ha_em_errnum {
    unsigned int    em_error_number;
    unsigned short  em_error_codes[2];
};

#define em_errnum    em_error.em_error_number    /* for quick tests */
#define em_genererr  em_error.em_error_codes[0] /* general error code */
#define em_specerr   em_error.em_error_codes[1] /* specific error code */

/*
 * The occurrence of an event, or the unregistration of an event, is reported
 * through a ha_em_rpb_event (Event Manager ResPonse Block: EVENT) structure.
 * This structure is included in the ha_em_resp_blk union, defined later in
 * this file.
 */
struct ha_em_rpb_event {
    union ha_em_errnum  em_error;          /* output: error number */
    ha_em_eid_t         em_event_id;       /* output: event identifier */
    unsigned long       em_event_flags;    /* output: event flags. See below */
    struct timeval       em_timestamp;     /* output: time of event */
    int                 em_location;       /* output: node generating event */
    char                *em_name;          /* output: resource variable name */
    char                *em_rsrc_ID;       /* output: resource ID */
    enum ha_emData_Type em_data_Type;      /* output: resource var. data type */
    union ha_em_val     em_val;           /* output: resource variable value */
};

#define HA_EM_EVENT_RE_ARM    0x0001 /* event generated from re-arm */
/* expression. */
#define HA_EM_EVENT_EXPR_FALSE 0x0002 /* expression evaluated to FALSE. */
/* (response to HA_EM_SCMD_REVAL */
/* subcommand) */
#define HA_EM_EVENT_UNREG    0x0004 /* event has been unregistered */
#define HA_EM_EVENT_PRED_FALSE HA_EM_EVENT_EXPR_FALSE /* for compatibility */

/*
 * When the current value of a resource variable is queried, it is reported
 * through a ha_em_rpb_qcur (Event Manager ResPonse Block: Query CURrent)
 * structure. This structure is included in the ha_em_resp_blk union,
 * defined later in this file.
 *
 * This structure is also used to return the resource IDs of variables
 * when the current value is not required (although the last known value
 * is returned).
 */
struct ha_em_rpb_qcur {
    union ha_em_errnum  em_error;          /* output: error number */
    int                 em_location;       /* output: node containing variable */
    char                *em_name;          /* output: name of resource variable */
    char                *em_rsrc_ID;       /* output: resource ID */
    enum ha_emData_Type em_data_Type;      /* output: resource var. data type */
    union ha_em_val     em_val;           /* output: resource variable value */
};

```

```

/*
 * When defined resource variables and expressions are queried, they are
 * reported through ha_em_rpb_qdef (Event Manager ResPonse Block: Query
 * DEFined) structure. This structure is included in the ha_em_resp_blk union,
 * defined later in this file.
 */

struct ha_em_rpb_qdef {
    union ha_em_errnum {
        em_error; /* output: error number */
        char *em_name; /* output: name of resource variable */
        char *em_descrp; /* output: resource variable description */
        enum ha_emValue_Type em_value_type; /* output: Counter, Quantity, or
        /* State */
        enum ha_emData_Type em_data_type; /* output: long, float, or structured*/
        /* byte string (SBS) */
        char *em_sbs_format; /* output: description of SBS format */
        char *em_init_value; /* output: initial value of variable */
        char *em_class; /* output: variable's resource class */
        char *em_rsrc_ID; /* output: resource ID definition */
        char *em_rsrc_ID_descrp; /* output: resource ID description */
        char *em_ptx_name; /* output: PTX shared memory name */
        char *em_dflt_expr; /* output: default expression */
        char *em_event_descrp; /* output: event description */
        char *em_locator; /* output: location indicator */
        char *em_order_group; /* output: resource variable order group */
    };

/*
 * When the parsing of an element in an event registration request results in
 * an error, it is reported through a ha_em_rpb_rerr (Event Manager ResPonse
 * Block: Registration ERror) structure. This structure is included in the
 * ha_em_resp_blk union, defined later in this file.
 *
 * The fields in this structure are a reflection of the information in the
 * request in error.
 */

struct ha_em_rpb_rerr {
    union ha_em_errnum {
        em_error; /* output: error number */
        char *em_name; /* output: resource variable name */
        char *em_rsrc_ID; /* output: resource ID */
        char *em_expr; /* output: expression */
        char *em_raexpr; /* output: re-arm expression */
        short em_errinfo0; /* output: additional error info */
        unsigned short em_errinfo1; /* output: additional error info */
        ha_em_eid_t em_event_id; /* output: event identifier */
    };
#define em_expr_err em_errinfo0 /* output: expression parsing error */
#define em_expr_pos em_errinfo1 /* output: position of parsing error */
#define em_pred_err em_expr_err /* for compatibility */
#define em_pred_pos em_expr_pos /* for compatibility */

/*
 * When the parsing of an element in a query request results in an error, it
 * is reported through a ha_em_rpb_qerr (Event Manager ResPonse Block: Query
 * ERror) structure. This structure is included in the ha_em_resp_blk union,
 * defined later in this file.
 *
 * The fields in this structure are a reflection of the information in the
 * request in error.
 */

struct ha_em_rpb_qerr {
    union ha_em_errnum {
        em_error; /* output: error number */
        char *em_class; /* output: resource variable class name */
        char *em_name; /* output: resource variable name */
        char *em_rsrc_ID; /* output: resource ID */
        unsigned short em_errinfo; /* output: additional error info */
    };

/*
 * When Event Manager daemon server credentials are returned, they are returned
 * through a ha_em_rpb_creds (Event Manager ResPonse Block: CREdentials)
 * structure. This structure is included in the ha_em_resp_blk union,
 * defined later in this file.
 */

struct ha_em_rpb_creds {
    union ha_em_errnum {
        em_error; /* output: error number */
        long em_cred_len; /* output: length of credentials */
        char *em_cred_data[1]; /* output: the user credentials */
    };

/*
 * Multiple responses of a particular type are made through a ha_em_resp_blk
 * (Event Manager ResPonse Block) union. All the responses are of the
 * same type. The number of elements in the array depends on the number
 * of responses being made. This union is included in the ha_em_rpb_blk
 * structure, defined later in this file.
 */

union ha_em_resp_blk {

```

```

    struct ha_em_rpb_event  em_rpb_event[1];    /* Used for events          */
    struct ha_em_rpb_qcur   em_rpb_qcur[1];    /* Used for current values */
    struct ha_em_rpb_qdef   em_rpb_qdef[1];    /* Used for defined variables*/
    struct ha_em_rpb_rerr   em_rpb_rerr[1];    /* Used for reg cmd errors  */
    struct ha_em_rpb_qerr   em_rpb_qerr[1];    /* Used for query cmd errors*/
    struct ha_em_rpb_creds  em_rpb_creds;      /* Used for credentials     */
};

/*
 * A block of responses of a particular type are delivered through a
 * ha_em_rsp_blk (Event Manager ReSPonse BLock) structure passed through
 * the ha_em_receive_command() routine. All the responses are of the same type.
 * The type of the requests are identified in the em_cmd and em_subcmd
 * fields. The responses themselves are in an array within the ha_em_rsp_blk
 * union.
 */

struct ha_em_rsp_blk {
    int          em_rsp_blk_len; /* output: length of responses */
    int          em_rsp_num_resp; /* output: number of responses */
    short       em_cmd; /* output: type of responses */
    short       em_subcmd; /* output: subtype of responses */
    ha_em_qid_t em_qid; /* output: query identifier */
    int         em_qend; /* output: query response end */
    union ha_em_rsp_blk em_resp_blk; /* output: the array of responses*/
};

/*-----*/
/* Event Manager API compatibility definitions. */
/* */
/* These definitions are provided to maintain source compatibility with */
/* programs written using prior versions of this header file. */
/* If these definitions result in inappropriate substitutions, then define */
/* the symbol HA_EM_NO_NAME_COMPAT prior to inclusion of this header file. */
/* If HA_EM_NO_NAME_COMPAT is defined, the source files that include this */
/* header file must be modified to use the new symbol names if the old */
/* symbol names are referenced therein. */
/*-----*/

#ifndef HA_EM_NO_NAME_COMPAT
#ifndef em_ivector
#define em_ivector em_rsrc_ID /* replace em_ivector by em_rsrc_ID */
#endif
#ifndef em_ivector_descrp
#define em_ivector_descrp em_rsrc_ID_descrp /* replace em_ivector_descrp by */
/* em_rsrc_ID_descrp */
#endif
#ifndef em_dflt_pred
#define em_dflt_pred em_dflt_expr /* replace em_dflt_pred by em_dflt_expr */
#endif
#ifndef em_pred
#define em_pred em_expr /* replace em_pred by em_expr */
#endif
#ifndef em_rapred
#define em_rapred em_raexpr /* replace em_rapred by em_raexpr */
#endif
#endif /* HA_EM_NO_NAME_COMPAT */

#endif /* _HA_EMAPI_BASE_H */

```

## ha\_emcommon.h File

```

/* IBM_PROLOG_BEGIN_TAG */
/* This is an automatically generated prolog. */
/* */
/* */
/* Licensed Materials - Property of IBM */
/* */
/* (C) COPYRIGHT International Business Machines Corp. 1996,1998 */
/* All Rights Reserved */
/* */
/* US Government Users Restricted Rights - Use, duplication or */
/* disclosure restricted by GSA ADP Schedule Contract with IBM Corp. */
/* */
/* IBM_PROLOG_END_TAG */

/*=====*/
/* */
/* Module Name: ha_emcommon.h */
/* */
/* Description: */
/* Common definitions for the Event Management Subsystem */
/* */
/* */
/*=====*/

/* @(#)37 1.13 src/rsct/pem/emcommon/ha_emcommon.h, emcommon, comm_rmoh, rmoht3f5 5/20/98 14:55:03 */

```

```

#ifndef HA_EMCOMMON_H
#define HA_EMCOMMON_H

#define HA_EM_MAXERRMSG 256
#define HA_EM_MAXERRLVL 16
#define HA_EM_MAXERRFN 32

struct ha_em_err_blk {
    int    em_errline;
    char   em_errlevel[HA_EM_MAXERRLVL];
    char   em_errfile[HA_EM_MAXERRFN];
    int    em_errno;
    char   em_errmsg[HA_EM_MAXERRMSG];
};

#define HA_EM_CLEAR_ERR(X)      memset(X,0,sizeof(struct ha_em_err_blk))

#define HA_EM_RSRC_ID_SIZE     4      /* max number of elements in a
resource ID */
#define HA_EM_VECTOR_SIZE      HA_EM_RSRC_ID_SIZE /* for compatibility */

/* Max length of a resource ID value including null */
#define HA_EM_RSRC_ID_ELEM_VALUE_LEN_MAX 32
#define HA_EM_VECTOR_ELEM_VALUE_LEN_MAX HA_EM_RSRC_ID_ELEM_VALUE_LEN_MAX

/* Max length of a Structured Byte String (including length field) */
#define HA_EM_MAX_SBS_VALUE_LEN 2048

/* Max number of variable instances accepted per class from RM */
#define HA_EM_MAX_INSTS 10000

enum ha_emValue_Type {
    ha_emVTcounter,
    ha_emVTquantity,
    ha_emVTstate
};

enum ha_emData_Type {
    ha_emDTlong,
    ha_emDTfloat,
    ha_emDTsbs
};

enum ha_emField_Type {
    ha_emFTlong,
    ha_emFTfloat,
    ha_emFTchar,
    ha_emFTbyte
};

enum ha_emElement_Type {
    ha_emETint,
    ha_emETstring
};

/* define pathnames for lock, socket, log and config files */

/* Define type and macros used to manage a resource monitor lock file.
* This lock file is used to track a number of instances of the same
* resource monitor program.
*
* ha_em_lckreg_t is the region of the file that is locked.
* HA_EM_MAX_RM_INSTS limits the number of instances of a resource monitor
* and, hence, the number of lock regions.
* In HA_EM_RM_LOCK_FMT, first %s is resource monitor name, second %s
* is domain name.
*/
typedef int ha_em_lckreg_t;
#define HA_EM_RM_LOCK_FMT      "/var/ha/lck/haem/em.RM%s.%s"
#define HA_EM_MAX_RM_INSTS    8
#define HA_EM_LCK_SIZE        (sizeof(ha_em_lckreg_t))

/* file used in creation of key file, which holds the shared memory ID of
* a shared memory segment used by a resource monitor instance and the
* EM daemon. "Conceptually" a lock file, so put it in lck directory. First
* %s is resource monitor name, %d is resource monitor instance number and
* the last %s is the domain name.
*/
#define HA_EM_RMSHM_KEY_FMT    "/var/ha/lck/em.RM%s.%dSHM.%s"

/* Define macros useful in finding old key files from prior incarnations of
* the EM daemon. The %s is the domain name.
*/
#define HA_EM_RMSHM_KEY_DIR_FMT  "/var/ha/lck"
#define HA_EM_RMSHM_KEY_PFX_FMT  "em.RM"
#define HA_EM_RMSHM_KEY_SFX_FMT  "SHM.%s"

/* name of file used to dump the red zone of a shared memory segment.
* The red zone page is dumped when it is modified unexpectedly. The %s
* is a resource monitor name, the first %d is a resource monitor instance
* number and the last %d is a timestamp (in seconds).
*/

```

```

#define HA_EM_RMSHM_RZ_FMT          "rzdump.RM%s.%d.%d"

/* file passed as argument to SPMI routines by the RMAPI. The %s is a
 * resource monitor name.
 */
#define HA_EM_SPMI_KEY_FMT          "/var/ha/lck/em.RM%sSPMI"

/* lock file to ensure single instance of an event management daemon.
 * %s is domain name.
 */
#define HA_EM_DAEMON_LOCK_FMT      "/var/ha/lck/em.haemd.%s"

/* name of Unix domain socket used by EMAPI to connect to event management
 * daemon. %s is domain name.
 */
#define HA_EM_CL_TO_D_SOCKET_FMT   "/var/ha/soc/em.clsrv.%s"

/* name of Unix domain socket used by RMAPI to connect to event management
 * daemon. %s is domain name.
 */
#define HA_EM_RM_TO_D_SOCKET_FMT   "/var/ha/soc/em.rmsrv.%s"

/* name of Unix domain socket used by event management daemon to connect to
 * a server resource monitor instance. First %s is resource monitor name,
 * the %d is a resource monitor instance number and the last %s is a domain
 * name.
 */
#define HA_EM_D_TO_RM_SOCKET_FMT   "/var/ha/soc/haem/em.RM%s.%d.%s"

/* name of event management daemon log file for tracing. %s is domain
 * name.
 */
#define HA_EM_HAEMD_TRACE_LOG_FMT  "/var/ha/log/em.trace.%s"

/* name of event management daemon log file for tracing messages between the
 * various daemon components. %s is domain name.
 */
#define HA_EM_HAEMD_MSGTRACE_LOG_FMT "/var/ha/log/em.msgtrace.%s"

/* name of event management daemon log file for default messages. These are
 * messages produced before logging to error log is possible. %s is domain
 * name.
 */
#define HA_EM_HAEMD_DFLT_LOG_FMT   "/var/ha/log/em.default.%s"

/* name of event management configuration data base. %s is domain name. */
#define HA_EM_HAEMD_CFGDB_FMT      "/etc/ha/cfg/em.%s.cdb"

/* name of file that contains the version string of the CDB currently
 * being used by the event management daemon. This is used by the RMAPI
 * to insure it is using the same copy. %s is domain name.
 */
#define HA_EM_HAEMD_CFGDB_VERS_FMT "/etc/ha/cfg/em.%s.cdb_vers"

/* name of directory where event management instance executes; place where
 * core files are produced. %s is domain name.
 */
#define HA_EM_HAEMD_CWD_FMT        "/var/ha/run/haem.%s"

/* define the service name for the port used by Event Management for peer
 * communication. %s is domain name.
 */
#define HA_EM_SERVICE_NAME_FMT     "haem.%s"

/*
 * Expression parsing and evaluation errors.
 */

#define HA_EM_EXPR_ERVAR           1 /* Internal resource variable error. */
#define HA_EM_EXPR_EEXPRLEN       2 /* Parsing table overflow. */
#define HA_EM_EXPR_ESYNTAX        3 /* Error in expression syntax. */
#define HA_EM_EXPR_ESTACK         4 /* Internal stack overflow. */
#define HA_EM_EXPR_EILLEGALOP     5 /* Illegal operation between types. */
#define HA_EM_EXPR_ESTRING        6 /* Error in string constant syntax. */
#define HA_EM_EXPR_ESBSSN         7 /* SBS out of range or missing from
 * value. */
#define HA_EM_EXPR_EFIELDTYPE     8 /* Illegal or mismatch in SBS field
 * type. */
#define HA_EM_EXPR_EDATATYPE      9 /* Invalid internal data type. */
#define HA_EM_EXPR_ESBSLEN        10 /* Mismatch SBS data type and field
 * length. */
#define HA_EM_EXPR_EDIVIDE        11 /* Divide by zero. */
#define HA_EM_EXPR_EOCTAL_INVALID 12 /* Octal value was not valid. */
#define HA_EM_EXPR_EOCTAL_2BIG    13 /* Octal value was > maximum allowed */

#define HA_EM_PRED_ERVAR          HA_EM_EXPR_ERVAR /* for compat */
#define HA_EM_PRED_EEXPRLEN      HA_EM_EXPR_EEXPRLEN /* for compat */
#define HA_EM_PRED_ESYNTAX       HA_EM_EXPR_ESYNTAX /* for compat */
#define HA_EM_PRED_ESTACK        HA_EM_EXPR_ESTACK /* for compat */
#define HA_EM_PRED_EILLEGALOP    HA_EM_EXPR_EILLEGALOP /* for compat */
#define HA_EM_PRED_ESTRING       HA_EM_EXPR_ESTRING /* for compat */
#define HA_EM_PRED_ESBSSN        HA_EM_EXPR_ESBSSN /* for compat */

```

```
#define HA_EM_PRED_EFIELDTYPE      HA_EM_EXPR_EFIELDTYPE      /* for compat */
#define HA_EM_PRED_EDATATYPE      HA_EM_EXPR_EDATATYPE      /* for compat */
#define HA_EM_PRED_ESBSLEN        HA_EM_EXPR_ESBSLEN        /* for compat */
#define HA_EM_PRED_EDIVIDE        HA_EM_EXPR_EDIVIDE        /* for compat */
#define HA_EM_PRED_EOCTAL_INVALID HA_EM_EXPR_EOCTAL_INVALID /* for compat */
#define HA_EM_PRED_EOCTAL_2BIG    HA_EM_EXPR_EOCTAL_2BIG    /* for compat */

#endif /* _HA_EMCOMMON_H */
```

## Related Information

EMAPI subroutines: **ha\_em\_end\_session**, **ha\_em\_receive\_response**, **ha\_em\_restart\_session**, **ha\_em\_send\_command**, **ha\_em\_start\_session**

---

## Expressions (haemexpr)

### Purpose

Expressions – Expressions that are used in Event Management

### Description

In the Event Management Application Programming Interface (EMAPI), an expression is a condition that a resource variable must meet in order to generate an event. More formally, an expression is the relational condition between a resource variable and other elements, such as a constant or the value of a variable from the previous observation.

Expressions that represent conditions of interest are defined by the programmers of Event Management clients. Default expressions are defined by the designers of resource monitors when they design the data that their resource monitors report.

The expression is applied to each instance of the resource variable as it is observed. If the expression is true, an event is generated. More than one expression may be applied to an instance of the same resource variable at the same observation.

An example of a simple expression is  $X < 10$ , where  $X$  is a resource variable that represents the percentage of free space in a file system. This expression would generate an event whenever the file system's free space was observed to be less than 10%.

For information about how resource variables are defined, see "Resource Variables (haemrvars)" on page 126.

An expression is of the following form:

```
expression::  operand log_op operand
              operand rel_op operand
              operand arith_op operand
              unary_op operand

operand::     expression
              var_name
              constant

var_name::    X
              X@var_name_mod

var_name_mod:: one of
               P, R, PR, sbs_sn, Psbs_sn

unary_op::    !

rel_op::      one of
               == != < > <= >=

log_op::      one of
               && ||

arith_op::    one of
               * / % + -

sbs_sn::      a structured field serial number
constant::    one of
               long_value float_value "char constant" "byte constant"
```

Operators have the same meaning and precedence as in the C language. Parentheses may be used for grouping as in C. Constants are integer or floating point, also as in C.

Fractional constants must be in decimal format.

Character and byte string constant values are enclosed in double quotes (""). Octal values may be inserted into these constants to represent non-printable characters by including a '\ ' character followed by the octal digits. The '\ ' and octal value are converted into the one-byte value and must be between 0 and 255. The '\ ' is represented by placing two adjacent '\ ' characters in the string. Character string constants are automatically **NULL** terminated.

The variable name modifier **P** indicates that the value of the variable instance from its previous observation is used. The variable name modifier **R** indicates that the raw value of the variable instance is used; this is useful only with resource variables of type Counter. When both of these modifiers are present (**PR**), the raw value of the variable instance from its previous observation is used. If neither modifier is present for a Counter, a rate is used. The rate is calculated as follows. The raw value of the previous observation of the variable instance is subtracted from the raw value of the latest observation. The result is divided by the time, in seconds, between the two observations.

The variable name modifier *sbs\_sn* is a structured field serial number. This modifier is used only with a structured byte string (SBS) resource variable and is used to select the structured field value that is to be used in the evaluation of the expression. When used with **P**, the selected structured field is taken from the value of the variable instance from its previous observation. For more information about the format of SBS resource variables, see "Resource Variables (haemrvars)" on page 126.

Operands that are modified to select structured fields of type character string or byte string may only be used with the relational operators. Both operands must be of the same type. When the operands are character strings, the implied comparison and its result are equivalent to the C library function **strcmp(*loprn*,*roprn*)**, where *loprn* is the left operand in the expression and *roprn* is the right operand in the expression. When the operands are byte strings, the implied comparison and its result are equivalent to the C library function **memcmp(*loprn*,*roprn*,*N*)**, where *loprn* is the left operand in the expression, *roprn* is the right operand in the expression, and *N* is the length of the shortest byte string.

Expressions use the letter "X" to represent the resource-variable name. The variable name may be repeated in the expression and may be in any one of its modified forms.

Here are some examples of valid expression definitions:

**X == 0** The value of the resource variable instance is equal to zero.

**X < 20 || X > 80**

The value of the resource variable instance is less than 20 or greater than 80.

**!(X < 20 || X > 80)**

The value of the resource variable instance is neither less than 20 nor greater than 80.

**X@R > X@PR**

The current raw value of the variable instance is greater than the raw value of the variable instance from its previous observation.

**X >= X@P + 5**

The current value of the variable instance is greater than or equal to the value of the variable instance from its previous observation plus 5.

**X@2 != 100**

For a resource variable instance that is defined as an SBS, the value of the structured field number 2 is not equal to 100.

**X@0 != X@P0 || X@1 != X@P1**

For a resource variable instance that is defined as an SBS, the value of either structured field number 0 or structured field number 1 has changed since its previous observation.

## **Related Information**

EMAPI subroutines: **ha\_em\_receive\_response**, **ha\_em\_send\_command**

---

## Resource Variables (haemrvars)

### Purpose

Resource Variables and Resource IDs – Resource variables and resource IDs used in Event Management

### Description

In the Event Management API (EMAPI) and the Resource Monitor API (RMAPI), the attributes of system resources are represented by resource variables. In addition, the copy of each resource is represented by a resource ID.

Resource variables and resource IDs are defined by the programmers of resource monitors. Unless otherwise indicated, the definitions and specifications described here apply to both the EMAPI and the RMAPI.

### Resource Variables

A resource variable represents the attribute of a system resource. Associated with each resource variable is a name, a value type, a data type, a resource ID and, optionally, a location.

### Resource Variable Names

A resource variable is identified by a **resource variable name**, which is a string that consists of a resource name followed by a period followed by the resource attribute.

By convention, a **resource name** represents a hierarchical organization of components, from general to specific, similar to a filename. It consists of a series of two or more components separated by periods, as follows:

- The first component is the name of the vendor that supplies the subsystem or application that manages the resource.
- The second component is the name of the product that contains the subsystem or application.
- Any additional components that are required to name the resource appear after the second component.

The **resource attribute** is a single component that follows the resource name.

Each component of the resource-variable name is a character string that may contain only alphanumeric characters, underscores, or a percent sign. The first component (vendor name) must begin with an alphabetic character. All other components must begin with either an alphabetic character or a percent sign.

Here are some examples of resource-variable names:

```
IBM.PSSP.aixos.CPU.gluser
IBM.PSSP.aixos.Mem.Real.size
IBM.PSSP.aixos.Mem.Virt.pagein
IBM.PSSP.aixos.Mem.Virt.pageout
IBM.PSSP.aixos.Mem.Kmem.inuse
IBM.PSSP.aixos.PagSp.%totalfree
IBM.PSSP.aixos.Disk.busy
IBM.PSSP.Membership.Node.state
```

The name space for resource-variable names is local to a domain.

Although components of a resource-variable name can be used in more than one name, resource names must uniquely identify a resource.

### Resource Variable Value Types

A resource variable belongs to one of three value types: Counter, Quantity, or State.

A **Counter** is a resource variable that represents a rate. It can have a data type of either long or float. Typically, a Counter represents throughput. Examples include paging rates, I/O rates, and transaction rates.

By default, the Event Management subsystem presents the value of an instance of a Counter as a rate. The rate represents the change in the actual contents of the instance of the Counter from one observation to the next divided by the time between the two observations. However, the Event Management subsystem can also present the **raw value** of the instance of the Counter. This is the actual contents of the instance of the Counter from the latest observation. The raw value of a Counter always increases over time.

For example, the resource variable called **IBM.PSSP.aixos.Mem.Virt.pagein** is a Counter that represents the rate at which pages are paged into virtual memory. Normally, its value is presented as a rate (number of pages per second). However, its raw value indicates the total number of pages that have been paged in since the Counter was initialized.

A Counter can have a data type of either long or float. These formats are identical to the C language types of the same names.

The Event Management subsystem assumes that instances of a Counter may be observed at an interval different from that of their update without loss of meaningful information.

A **Quantity** is a variable whose value fluctuates over time. Typically, a Quantity is used to represent a level, that is, an indicator of how many.

A Quantity can have a data type of either long or float. These formats are identical to the C language types of the same names.

The Event Management subsystem assumes that instances of a Quantity may be observed at an interval different from that of their update without loss of meaningful information.

A **State** is a variable whose value fluctuates over time, like a Quantity. Unlike a Quantity, however, the semantics of a State variable are assumed to be such that every change in value of an instance of a State variable must be observed; otherwise, meaningful information would be lost. A State variable can be used to represent an attribute of a resource that indicates anything other than throughput or a level.

A State can have a data type of long, float, or structured byte string (SBS). The long and float formats are identical to the C language types of the same names. The format of an SBS is described in “Resource Variable Data Types”.

### **Resource Variable Data Types**

A Counter and a Quantity can have a data type of either long or float. The long and float formats are identical to the C language types of the same names.

A State can have a data type of long, float, or structured byte string. The long and float formats are identical to the C language types of the same names.

A **structured byte string (SBS)** is a string of bytes, where each byte may have any value from 0 through 255, that consists of a four-byte SBS length field followed by one or more structured fields. The SBS length specifies the total length of the structured fields that follow it.

A structured field consists of a four-byte header followed by a value. The first two bytes of the header are the length of the structured field value, the third byte is a structured field data type, and the fourth byte is a serial number.

The structured field type is one of long, float, character string, or byte string. These types are defined by the enumerator **ha\_emField\_Type**, which is included in the **ha\_emapi.h** header file. Long and float are the same as in the C language. A character string consists of one or more nonzero bytes terminated by a null byte; the null byte is included in the structured field value length. A byte string consists of one or more bytes, where each byte may have any value from 0 through 255.

The serial number is a unique value that identifies the structured field. This serial number is defined for each structured field by the resource monitor that supplies the SBS resource variable. However, the set of serial numbers for the structured byte string starts with 0 and is contiguous.

Associated with each structured field in an SBS, but not included in the actual SBS, is a structured field name. This name is used to provide a short description of the structured byte field and must consist of only alphanumeric characters and underscores; the first character must be alphabetic.

Here are some examples of structured-byte-string definitions that were taken from resource variables provided by RSCT resource monitors:

-----  
 SBS associated with resource variable IBM.PSSP.Prog.pcount

SBS Length = variable

SBS Fields

Field Name	Field Length	Field Type	Field Serial Number
CurPIDCount	4	long	0
PrevPIDCount	4	long	1
CurPIDList	variable	cstring	2

-----  
 SBS associated with resource variable IBM.PSSP.pm.Errlog

SBS Length = variable

SBS Fields

Field Name	Field Length	Field Type	Field Serial Number
sequenceNumber	variable	cstring	0
errorID	variable	cstring	1
errorClass	variable	cstring	2
errorType	variable	cstring	3
alertFlagsValue	variable	cstring	4
resourceName	variable	cstring	5
resourceType	variable	cstring	6
resourceClass	variable	cstring	7
errorLabel	variable	cstring	8

The Counter and Quantity types have the same definition as the value types for statistics objects that are defined by the Performance Toolbox for AIX (PTX/6000) System Performance Measurement Interface (SPMI).

## Resource Variable Instances and Resource IDs

Most resources in the system have multiple copies. Each of these copies is an **instance** of the resource.

The resource variables that represent the states of these resources also have multiple copies. Each of these copies is an **instance** of the resource variable.

To uniquely identify each copy of a resource and all of its variables, each resource in the system has one, and only one, associated resource ID. A resource ID is a list of elements, where each element is a name/value pair. In the EMAPI, the elements are separated by semicolons. In the RMAPI, the elements are separated by commas.

A name/value pair consists of a resource ID element name followed by an equal sign followed by the value of the resource ID element. There are no blanks in the resource ID. A resource ID can contain up to 4 elements, as defined by the **HA\_EM\_RSRC\_ID\_SIZE** constant.

The resource ID element name is a string that describes the element. It must consist of only alphanumeric characters and underscores; the first character must be alphabetic. In the EMAPI, a resource ID element value may consist of a single value, a range of values, a comma-separated list of single values, or a comma-separated list of ranges. A range takes the form *a-b* and is valid only for resource ID elements of type integer. In the RMAPI, only single values are allowed.

A resource ID element may be wildcarded in a manner that varies by subroutine.

The set of values in the resource ID uniquely identify the copy of the resource in the system. By extension, they also uniquely identify the copy of the resource variable in the system. If there is only one copy of the resource in the system (for example, the control workstation), its resource ID is null.

The semantics of a resource ID are defined by the semantics of the resource with which it is associated. The names of the resource ID elements must be unique within a given resource's resource ID. These names may be used in other resource IDs and may or may not have the same semantics. In other words, the name space for resource ID element names is local to each defined resource.

Here are some examples of PSSP resource variable names and resource IDs:

Resource Variable Name	Resource ID
IBM.PSSP.aixos.CPU.gluser	NodeNum=5
IBM.PSSP.aixos.cpu.kern	NodeNum=5;cpu=cpu0
IBM.PSSP.aixos.Mem.Real.size	NodeNum=5
IBM.PSSP.aixos.Mem.Virt.pagein	NodeNum=5
IBM.PSSP.aixos.Mem.Virt.pageout	NodeNum=5
IBM.PSSP.aixos.Mem.Kmem.inuse	NodeNum=5;Type=mbuf
IBM.PSSP.aixos.PagSp.%totalfree	NodeNum=5
IBM.PSSP.aixos.Disk.busy	NodeNum=5;Name=hdisk0
IBM.PSSP.aixos.Disk.busy	NodeNum=5;Name=hdisk1
IBM.PSSP.aixos.VG.free	NodeNum=5;VG=rootvg
IBM.PSSP.aixos.VG.free	NodeNum=5;VG=spdata
IBM.PSSP.aixos.FS.%totused	NodeNum=5;VG=rootvg;LV=hd4
IBM.PSSP.aixos.FS.%totused	NodeNum=5;VG=rootvg;LV=hd4
IBM.PSSP.SP_HW.Node.powerLED	NodeNum=5
IBM.PSSP.SP_HW.Frame.frACLEd	FrameNum=1
IBM.PSSP.SP_HW.Switch.powerLED	SwitchNum=1
IBM.PSSP.Membership.Node.state	NodeNum=5
IBM.DB2.Part.size	DBname=foo;Part=3

Although many resources are associated with a particular node, some are not. Therefore, not all resource IDs contain a node number.

## Dynamically Instantiable Resource Variables

In most cases, the instances of a resource that a resource monitor is responsible for tracking and reporting on are static or relatively so. Thus, a resource monitor knows from the time it starts running, after any necessary initialization occurs, all of the instances of the resources it monitors. For example, after any necessary initialization, a hardware monitor, a node-health monitor, and a disk monitor know what frames, nodes, and disks, respectively, are in the system.

However, in some cases, a resource monitor knows how to track and report on a type of resource. Nonetheless, the resource monitor does not know all of the possible instances of the resource that may occur. For example, the program resource monitor knows how to track and report on a program that is running in a domain, but it does not know in advance the names of all possible programs that can run. For a resource like this, you can define the resource variable as **dynamically instantiable**.

If a resource variable is dynamically instantiable, the resource monitor does not register all of the instances of its resource variables. Instead, it waits for the Event Management subsystem to tell it what instances to register. As the Event Management subsystem receives requests from EM clients for events

based on instances of dynamically instantiable variables, the Event Management subsystem passes the instances to the resource monitor in a control message with an “instantiate variables” command.

### Location of a Resource Variable Instance

The **location** of a resource variable instance is the node on which the *resource monitor* that supplies the instance resides. The resource monitor may or may not reside on the same node as the instance of the resource.

### Resource Variable Classes

The **class** of a resource variable is used to indicate the subsystem or application that manages the associated resource. It can also be used to group resource variables by observation interval. Instances of all resource variables of value type Counter and Quantity in the same class are observed at the same intervals.

By convention, class names follow the same hierarchical naming conventions that are used for resource names.

Each class name must be unique among class names within a domain.

Here are some examples of class names:

Resource Variable Class	Resource Variables in the Class
IBM.PSSP.aixos.CPU	IBM.PSSP.aixos.CPU.gIuser IBM.PSSP.aixos.CPU.gIkern IBM.PSSP.aixos.CPU.gIwait IBM.PSSP.aixos.CPU.gIidle IBM.PSSP.aixos.cpu.user IBM.PSSP.aixos.cpu.kern IBM.PSSP.aixos.cpu.wait IBM.PSSP.aixos.cpu.idle
IBM.PSSP.aixos.Mem	IBM.PSSP.aixos.Mem.Real.size IBM.PSSP.aixos.Mem.Real.numfrb IBM.PSSP.aixos.Mem.Real.%free IBM.PSSP.aixos.Mem.Real.%pinned IBM.PSSP.aixos.Mem.Virt.pagein IBM.PSSP.aixos.Mem.Virt.pageout IBM.PSSP.aixos.Mem.Virt.pgspgin IBM.PSSP.aixos.Mem.Virt.pgspgout IBM.PSSP.aixos.Mem.Virt.pagexct IBM.PSSP.aixos.Mem.Kmem.inuse IBM.PSSP.aixos.Mem.Kmem.calls IBM.PSSP.aixos.Mem.Kmem.failures IBM.PSSP.aixos.Mem.Kmem.memuse
IBM.PSSP.aixos.Disk	IBM.PSSP.aixos.Disk.busy IBM.PSSP.aixos.Disk.xfer IBM.PSSP.aixos.Disk.rblk IBM.PSSP.aixos.Disk.wblk

## Related Information

EMAPI subroutine: [ha\\_em\\_send\\_command](#)

---

## RMAPI Errors (err\_rmapi)

### Purpose

RMAPI Errors – Error numbers and the error block for the Resource Monitor Application Programming Interface (RMAPI)

### Description

The Resource Monitor Application Programming Interface (RMAPI) provides a class of errors that are returned synchronously by RMAPI subroutines.

### Synchronous Errors That Are Returned by RMAPI Subroutines

If an RMAPI subroutine is unsuccessful, it returns a value of -1 and other error information in the error block specified on input. The error block contains an error number and a null-terminated error message.

The RMAPI uses an error block that is common to both the EMAPI and the RMAPI. The block is defined in the **ha\_emcommon.h** header file that is included by the **ha\_rmapi.h** header file.

The RMAPI error numbers are defined in the **ha\_rmapi.h** header file.

The Event Management error block has the following definition:

```
struct ha_em_err_blk {
    int     em_errline;
    char    em_errlevel[HA_EM_MAXERRLVL];
    char    em_errfile[HA_EM_MAXERRFN];
    int     em_errno;
    char    em_errmsg[HA_EM_MAXERRMSG];
}
```

The **em\_errline**, **em\_errlevel**, and **em\_errfile** fields are reserved for IBM use in providing information for problem determination.

If an error occurs, the **em\_errno** field contains one of the values listed below. The **em\_errmsg** field contains a message text that describes the error in further detail. For information about RMAPI messages, see *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*.

In addition to indicating the reason that an RMAPI subroutine is unsuccessful, RMAPI error numbers are also used to indicate errors in individual elements of the **ha\_rr\_variable** array that was passed to the **ha\_rr\_add\_var**, **ha\_rr\_del\_var**, **ha\_rr\_reg\_var**, or **ha\_rr\_unreg\_var** subroutine. The error number is returned in the **rr\_var\_errno** field of the **ha\_rr\_variable** array element. The **ha\_rmapi.h** header file indicates which error numbers may be returned in the **rr\_var\_errno** field.

The RMAPI error numbers are:

#### HA\_RR\_ENOSDR

An error occurred while trying to obtain information from the System Data Repository (SDR). Either an SDR class had a missing or invalid attribute, or an SDR subroutine failed.

#### HA\_RR\_EACCESS

An error occurred while attempting to open or create a file used by the RMAPI.

#### HA\_RR\_EBADCDB

A problem was detected with the Event Management Configuration Database (EMCDB) file. Either the file was found to be corrupted, or the file version does not match the version of the EMCDB being used by the Event Management daemon. The latter condition could occur if the EMCDB was updated after the monitor started.

#### HA\_RR\_ESPMIFAIL

A System Performance Measurement (SPMI) routine failed.

**HA\_RR\_ENOTSERVER**

An RMAPI routine specific to server type resource monitors was called by a client type resource monitor.

**HA\_RR\_ENOHNDR**

The resource monitor requested **HA\_RR\_NOTIFY\_SIGIO** as the notification protocol, but does not have a SIGIO signal handler installed.

**HA\_RR\_ENOSERVSOCK**

An attempt to start a session was made by the server type resource monitor prior to calling **ha\_rr\_makserv** to establish a server socket.

**HA\_RR\_EBADPROTO**

The notification protocol specified to the RMAPI was not a valid value, or does not match the protocol value that was specified on a previous call to the RMAPI.

**HA\_RR\_EEXIST**

A session or server file descriptor already exists. Server resource monitors could receive this error if **ha\_rr\_makserv** is called more than once or if more than one session has been started to the same manager because **ha\_rr\_start\_session** was called before a previously closed session with **ha\_rr\_end\_session** was ended. A command-based resource monitor receives this error if **ha\_rr\_start\_session** is called more than once.

**HA\_RR\_EAGAIN**

A request was made to accept a connection on the resource monitor server socket, but no connection request was found, or a request to read data from a connection failed because there was no data ready to be read.

**HA\_RR\_EMAXSESSIONS**

A request for a session was denied because the number of active sessions would exceed the maximum number that are allowed.

**HA\_RR\_EBADSESSION**

The RMAPI did not find a session matching the session file descriptor parameter.

**HA\_RR\_EDISCONNECT**

A resource monitor manager dropped its connection to the resource monitor.

**HA\_RR\_ENOENT**

The specified entity was not found in the EMCDB file.

**HA\_RR\_EBADRIDNAME**

The value of a resource ID element was not valid. The value was either too long (greater than **HA\_EM\_RSRC\_ID\_ELEM\_VALUE\_LEN\_MAX**), contained blanks, or the element was defined to be numeric and contained characters that were not in the range '0'-'9'.

**HA\_RR\_ENORIDNAME**

The resource ID specified was not valid. One or more portions of the resource ID elements name/value pairs was missing.

**HA\_RR\_EBADRSRCID**

Undefined resource ID elements were found in the resource ID string. One possible cause for this error could occur if the only resource ID element defined for the variable is also defined to be the Locator field. In this case the monitor should not specify a resource ID string for the variable.

**HA\_RR\_EDUPRSRCID**

A resource ID element name/value pair was found more than once in the resource ID string.

**HA\_RR\_EBADCLASS**

The specified resource monitor class was not found in the EMCDB file for this resource monitor.

**HA\_RR\_ESYSCALL**

An error was returned by a system call.

**HA\_RR\_ENOMEM**

Required memory could not be allocated.

**HA\_RR\_EINTERNAL**

The RMAPI encountered an internal error.

**HA\_RR\_EPARM**

A parameter passed to an RMAPI subroutine was not valid.

**HA\_RR\_EHADDR**

The address of the variable handle location is null.

**HA\_RR\_EHANDLE**

The handle value passed to the RMAPI was **NULL** or not a valid handle.

**HA\_RR\_ENOTREG**

The specified resource variable instance was not registered with the RMAPI.

**HA\_RR\_EDUPREG**

The resource variable instance identifier is different from the value passed to the RMAPI on a prior registration of the variable. Resource variables may be registered more than once with the RMAPI, but the instance identifier specified by the monitor must remain the same.

**HA\_RR\_EGETNODNUM**

The RMAPI could not determine the number of the node on which the application is running.

**HA\_RR\_ESDRGET**

The RMAPI could not obtain required information from the System Data Repository (SDR).

**HA\_RR\_EDOMAIN**

An error occurred when trying to determine the domain the monitor is executing in.

**HA\_RR\_ENOSBS\_VALUE**

The value of the State type variable was **NULL**.

**HA\_RR\_ESBSLEN**

The length reported in an SBS value is too long (greater than **HA\_EM\_MAX\_SBS\_VALUE\_LEN**), or the actual length (sum of the SBS fields) does not match the length as specified in the SBS header.

**HA\_RR\_EBAD\_SBSSN**

An SBS field in an SBS value specified a serial number that was not valid for the definition of the resource variable.

**HA\_RR\_EDUP\_SBSFIELD**

An SBS field was found more than once in the SBS value. An SBS value must contain each field defined for the SBS once.

**HA\_RR\_EBAD\_SBSFIELDTYPE**

The SBS value contained a field type that was not valid. The field type was an unknown value, or did not match the definition of the SBS field.

**HA\_RR\_EBAD\_SBSFIELDLEN**

The length of a field in an SBS value was not valid for the data type of the field.

**HA\_RR\_EMISSING\_SBSFIELD**

One or more fields are missing from the SBS value. An SBS value must contain each field defined for the SBS once.

**HA\_RR\_EPERM**

The requested function could not be performed. This error is returned if the resource monitor attempts to call an RMAPI subroutine in a manner that is not allowed. Conditions that return this error include the following.

1. Calling **ha\_rr\_rm\_ctl** to set attributes of the RMAPI after **ha\_rr\_init** has been called.

2. An attempt by a server-type resource monitor to add variables to a resource monitor manager session that has not sent an add command request.
3. An attempt to add State type variables in response to receiving an **HA\_RR\_CMD\_ADDALL** command.

a server type resource monitor attempting t or

#### **HA\_RR\_ERMINSTID**

The requested resource monitor instance is not valid for the definition of the resource monitor. This condition could occur if **ha\_rr\_rm\_ctl** was called prior to initializing the RMAPI, and the value of the **rr\_instance\_id** field of the **ha\_rr\_args** parameter was not valid for the monitor.

#### **HA\_RR\_EAPI\_LOCKED**

The RMAPI subroutines were locked due to an unrecoverable error that was previously returned. The monitor should call the **ha\_rr\_terminate** routine before attempting to re-initialize the RMAPI.

#### **HA\_RR\_ENOCONNECT**

An error occurred when the RMAPI attempted to make a connection to a resource monitor manager.

#### **HA\_RR\_ESHM**

A shared memory segment used by the RMAPI was found to be corrupted. This could occur if the process calling the RMAPI overwrites a portion of the shared memory segment, or another process has erroneously attached and written to the segment.

#### **HA\_RR\_EERRCMD**

A resource monitor manager detected an error in the operation of the RMAPI or the resource monitor. See the accompanying error message for details of the error.

#### **HA\_RR\_EAUTH**

The calling process does not have the authority to call the RMAPI.

#### **HA\_RR\_EINIT**

The RMAPI has not been initialized by a call to **ha\_rr\_init**. This error may also be returned by the **ha\_rr\_init** subroutine, indicating that a specified instance, or the maximum numbers of instances defined for the monitor, are already executing.

#### **HA\_RR\_ECONREFUSED**

The Event Management daemon refused a requested connection. This error can occur when a client type resource monitor calls **ha\_rr\_start\_session** to connect to the Event Management daemon. The error indicates that the Event Management daemon may not be running, is recovering from a failure, or is initializing.

#### **HA\_RR\_ENOLOCK**

During initialization, the RMAPI attempted to create a unique lock for the resource monitor process. The RMAPI was unable to obtain a lock for the monitor instance. All locks in the range 0 to the maximum number of instances the monitor is configured to allow, and as specified by any previous calls to **ha\_rr\_rm\_ctl**, were attempted. This error indicates the following. Either the maximum numbers of instances of this monitor are already executing, or a specific instance ID specified by a call to **ha\_rr\_rm\_ctl** was not available.

#### **HA\_RR\_EBADMGRID**

The RMAPI received a command from a resource manager session which is used to identify the resource monitor manager. The manager ID value in the command was not a valid value known to the RMAPI.

## **Related Information**

RMAPI subroutines: **ha\_rr\_add\_var**, **ha\_rr\_del\_var**, **ha\_rr\_end\_session**, **ha\_rr\_get\_ctrlmsg**, **ha\_rr\_get\_interval**, **ha\_rr\_init**, **ha\_rr\_makserv**, **ha\_rr\_reg\_var**, **ha\_rr\_send\_val**, **ha\_rr\_start\_session**, **ha\_rr\_terminate**, **ha\_rr\_touch**, **ha\_rr\_unreg\_var**, **ha\_rr\_rm\_ctl**

RMAPI header file: **ha\_rmap.h**

Messages: *PSSP: Messages Reference* or *HACMP: Troubleshooting Guide*

---

## ha\_rmapih File

### Purpose

**ha\_rmapih** – Header file for the Resource Monitor Application Programming Interface (RMAPI)

### Description

The **ha\_rmapih** header file provides data types and structures for use with the Resource Monitor Application Programming Interface (RMAPI) subroutines, which reside in the **libha\_rr.a** library. The **libha\_rr.a** library is not thread-safe. Any program that uses the RMAPI subroutines must include this file, which resides in the **/usr/include** directory.

The following listing shows the contents of the **ha\_rmapih** file. For the contents of the nested header file, **ha\_emcommonh**, see **ha\_emapih**.

```
/* IBM_PROLOG_BEGIN_TAG                               */
/* This is an automatically generated prolog.         */
/*                                                    */
/*                                                    */
/*                                                    */
/* Licensed Materials - Property of IBM                */
/*                                                    */
/* (C) COPYRIGHT International Business Machines Corp. 1996,1998 */
/* All Rights Reserved                                */
/*                                                    */
/* US Government Users Restricted Rights - Use, duplication or */
/* disclosure restricted by GSA ADP Schedule Contract with IBM Corp. */
/*                                                    */
/* IBM_PROLOG_END_TAG                                 */
/*=====*/
/*                                                    */
/* Module Name:  ha_rmapih                             */
/*                                                    */
/* Description:                                       */
/*   Include for Event Management resource monitor api (librmapih.a) */
/*                                                    */
/*=====*/
/* static char *sccsid = "@(#)38  1.17  src/rsct/pem/emcommon/ha_rmapih, emcommon, rsct_rtro 4/28/98 10:38:28"; */

#ifndef _H_HA_RMAPI
#define _H_HA_RMAPI

#include <ha_emcommon.h>

/*****
 *
 * Good and Bad return codes
 *
 *****/

#define HA_RR_PASS          (int)0
#define HA_RR_FAIL         (int)-1

/*****
 *
 * RMAPI error codes
 *
 * A '+' in the description indicates the error may be returned in either a
 * ha_em_err_blk structure, or in the rr_var_errno field of a ha_rr_variable
 * structure.
 *
 *****/
```

```

#define HA_RR_ENOSDR (int)101 /* SDR session subroutine failed */
#define HA_RR_EACCESS (int)102 /* Unable to open or create file */
#define HA_RR_EBADCDB (int)103 /* EMCDB corrupted or back level */
#define HA_RR_ESPMIFAIL (int)104 /* libSpmi.a subroutine failed */
#define HA_RR_ENOTSERVER (int)105 /* Calling RM is not a server */
#define HA_RR_ENOHNDR (int)106 /* No signal handler for SIGIO */
#define HA_RR_ENOSERVSOCK (int)107 /* RM server socket not created */
#define HA_RR_EBADPROTO (int)108 /* Bad or mismatched notify potocol */
#define HA_RR_EEXIST (int)109 /* Socket already exists */
#define HA_RR_EAGAIN (int)110 /* No data or connection pending */
#define HA_RR_EMAXSESSIONS (int)111 /* HA_RR_MAX_SESSIONS reached */
#define HA_RR_EBADSESSION (int)112 /* Specified session not found */
#define HA_RR_EDISCONNECT (int)113 /* Session connection closed */
#define HA_RR_ENOENT (int)114 /* + Specified entity not in EMCDB */
#define HA_RR_EBADRIDNAME (int)115 /* + Bad resource ID name */
#define HA_RR_ENORIDNAME (int)116 /* + Missing resource ID element(s) */
#define HA_RR_EBADRSRCID (int)117 /* + Bad resource ID element value */
#define HA_RR_EDUPRSRCID (int)118 /* + Duplicate resource ID element */
#define HA_RR_EBADCLASS (int)119 /* Resource class not found */
#define HA_RR_ESYSCALL (int)120 /* A system call failed */
#define HA_RR_ENOMEM (int)121 /* Memory allocation failed */
#define HA_RR_EINTERNAL (int)122 /* Internal RMAPI error ocured */
#define HA_RR_EPARAM (int)123 /* Bad parameter passed to RMAPI */
#define HA_RR_EHADDR (int)124 /* + Bad handle location */
#define HA_RR_EHANDLE (int)125 /* + Handle value not valid */
#define HA_RR_ENOTREG (int)126 /* + Variable was not registered */
#define HA_RR_EDUPREG (int)127 /* + ID changed on subsequent reg */
#define HA_RR_EGETNODNUM (int)128 /* Unable to determine node number */
#define HA_RR_ESDRGET (int)129 /* SDRGet subroutine failed */
#define HA_RR_EDOMAIN (int)130 /* Error determining domain name */
#define HA_RR_ENOVALUE (int)131 /* + RV value address was NULL */
#define HA_RR_ESBSLEN (int)132 /* + Bad SBS value length */
#define HA_RR_EBAD_SBSSN (int)133 /* + Bad SBS field serial number */
#define HA_RR_EDUP_SBSFIELD (int)134 /* + Duplicate SBS field in value */
#define HA_RR_EBAD_SBSFIELDTYPE (int)135 /* + Bad SBS field type */
#define HA_RR_EBAD_SBSFIELDLEN (int)136 /* + Bad SBS field value length */
#define HA_RR_EMISISSING_SBSFIELD (int)137 /* + SBS field(s) missing in value */
#define HA_RR_EPERM (int)138 /* Operation could not be performed */
#define HA_RR_ERMINSTID (int)139 /* Monitor instance id not valid */
#define HA_RR_EAPI_LOCKED (int)140 /* RMAPI locked on previous error */
#define HA_RR_ENAMETOOLONG (int)141 /* File name/socket path too long */
#define HA_RR_ENOCONNECT (int)142 /* Error connecting to RMM */
#define HA_RR_ESHM (int)143 /* Bad private shm segment */
#define HA_RR_EERRCMD (int)144 /* RM error detected by manager */
#define HA_RR_EAUTH (int)145 /* No authority to use RMAPI */
#define HA_RR_EINIT (int)146 /* API already/not yet initialized */
#define HA_RR_ECONNREFUSED (int)147 /* Mgr not available for connect */
#define HA_RR_ENOLOCK (int)148 /* Unable to lock an RM instance */
#define HA_RR_EBADMGRID (int)149 /* Unknown resource monitor manager */

#define HA_RR_EBADVECTNAME HA_RR_EBADRIDNAME /* For compatibility */
#define HA_RR_ENOVECTNAME HA_RR_ENORIDNAME /* For compatibility */
#define HA_RR_EBADVECTOR HA_RR_EBADRSRCID /* For compatibility */
#define HA_RR_EDUPVECTOR HA_RR_EDUPRSRCID /* For compatibility */

/*****
*
* ctrl command values received from resource managers
*
*****/

#define HA_RR_CMD_ADDV ((int)0x00c9) /* (int)201 */
#define HA_RR_CMD_ADDALL ((int)0x00ca) /* (int)202 */
#define HA_RR_CMD_DELV ((int)0x00cb) /* (int)203 */
#define HA_RR_CMD_DELALL ((int)0x00cc) /* (int)204 */
#define HA_RR_CMD_REFRESH ((int)0x00cd) /* (int)205 */

```

```

#define HA_RR_CMD_INSTV      ((int)0x00ce) /* (int)206 */
#define HA_RR_CMD_MGR_ID    ((int)0x00cf) /* (int)207 */
#define HA_RR_CMD_UNREG_ACK ((int)0x00d0) /* (int)208 */
#define HA_RR_CMD_ATTACH_SHM ((int)0x00d1) /* (int)209 */
#define HA_RR_CMD_ERROR     ((int)0x00d2) /* (int)210 */
#define HA_RR_CMD_MGR_ID_EVM ((int)0x01f5) /* (int)501 */
#define HA_RR_CMD_MGR_ID_PERF ((int)0x01f6) /* (int)502 */

/*****
 *
 * Error codes returned in the rr_ctl_cmdarg field if the command
 * is HA_RR_CMD_ERROR. (Only the header is sent).
 * These errors are processed by the RMAPI and are not returned to
 * the resource monitor.
 *
 *****/

#define HA_RR_CMD_ERROR_DUP    1
#define HA_RR_CMD_ERROR_ADD    2
#define HA_RR_CMD_ERROR_DEL    3
#define HA_RR_CMD_ERROR_BADRZ  4

/*****
 *
 * ctrl msg notification method HA_RR_NOTIFY_SELECT (for select or poll),
 * HA_RR_NOTIFY_SIGIO for SIGIO signal notification
 * HA_RR_MAX_SESSIONS is the max number of client sessions for server RMs
 *
 *****/

#define HA_RR_NOTIFY_SELECT (int)301
#define HA_RR_NOTIFY_SIGIO (int)302
#define HA_RR_MAX_SESSIONS (int)8

/*****
 *
 * Resource monitor instance id values.
 *
 *****/

#define HA_RR_RM_INSTID_PERF      (int)0
#define HA_RR_RM_INSTID_NOPERF   (int)-1
#define HA_RR_RM_INSTID_ANY      (int)-2
#define HA_RR_RM_INSTID_DEFAULT  HA_RR_RM_INSTID_ANY
#define HA_RR_RM_INSTID_MAX      (HA_EM_MAX_RM_INSTS - 1)

/*****
 *
 * Commands for the ha_rr_rm_ctl() routine.
 *
 *****/

#define HA_RR_RM_ARGS_GET      (int)0x00000001
#define HA_RR_RM_ARGS_SET_INSTID (int)0x00000002

/*****
 *
 * rmap structures
 *
 *****/

struct ha_rr_args {
    int    rr_instance_id;
    char   *rr_domain_name;
    char   rr_reserved[56];
};

struct ha_rr_variable {
    char   *rr_var_name;

```

```

    char    *rr_var_rsrc_ID;
    union {
        int    rr_var_inst_id;
        void   **rr_var_hndl;
    } rr_varu;
#define rr_var_handle    rr_varu.rr_var_hndl
#define rr_var_iid      rr_varu.rr_var_inst_id
    void    *rr_value;
    int     rr_var_errno;
};

struct ha_rr_val {
    void    *rr_value;
    void    *rr_var_hndl;
};

struct ha_rr_ctrl_msg {
    int     rr_ctrl_msg_len;
    int     rr_ctrl_cmd;
    int     rr_ctrl_cmdarg;
    int     rr_ctrl_num_vars;
    union   ha_rr_ctrlv {
        struct ha_rr_ctrl_var {
            char    *rr_ctrl_name;
            char    *rr_ctrl_rsrc_ID;
        } rr_ctrl_varn[1];
        int rr_ctrl_vari[1];
        struct ha_rr_ctrl_var2 {
            int     rr_ctrl_var_id;
            int     rr_ctrl_API_inst_id;
        } rr_ctrl_varn2[1];
    } rr_ctrlv;
#define rr_ctrl_vars    rr_ctrlv.rr_ctrl_varn
#define rr_ctrl_ids     rr_ctrlv.rr_ctrl_vari
#define rr_ctrl_vars2   rr_ctrlv.rr_ctrl_varn2
};

/*****
 * Resource Monitor API compatibility definitions.
 *
 * These definitions are provided to maintain source compatibility with
 * programs written using prior versions of this header file.
 * If these definitions result in inappropriate substitutions, then define
 * the symbol HA_EM_NO_NAME_COMPAT prior to inclusion of this header file.
 * If HA_EM_NO_NAME_COMPAT is defined, the source files that include this
 * header file must be modified to use the new symbol names if the old
 * symbol names are referenced therein.
 *****/

#ifndef HA_EM_NO_NAME_COMPAT
#ifndef rr_var_ivector
#define rr_var_ivector rr_var_rsrc_ID
/* replace rr_var_ivector by rr_var_rsrc_ID */
#endif
#ifndef rr_ctrl_ivector
#define rr_ctrl_ivector rr_ctrl_rsrc_ID
/* replace rr_ctrl_ivector by rr_ctrl_rsrc_ID */
#endif
#endif /* HA_EM_NO_NAME_COMPAT */
/*****
 *
 * rmap function prototypes
 *
 *****/

#ifdef _NO_PROTO
int ha_rr_rm_ctl();

```

```

int ha_rr_init();
int ha_rr_makserv();
int ha_rr_start_session();
int ha_rr_get_ctrlmsg();
int ha_rr_end_session();
int ha_rr_add_var();
int ha_rr_del_var();
int ha_rr_get_interval();
int ha_rr_send_val();
int ha_rr_reg_var();
int ha_rr_unreg_var();
int ha_rr_terminate();
int ha_rr_touch();
#else /* NO_PROTO */
int ha_rr_rm_ctl(struct ha_rr_args *, int, struct ha_em_err_blk *);
int ha_rr_init(char *, struct ha_em_err_blk *);
int ha_rr_makserv(int, struct ha_em_err_blk *);
int ha_rr_start_session(int, struct ha_em_err_blk *);
int ha_rr_get_ctrlmsg(int, struct ha_rr_ctrl_msg **, struct ha_em_err_blk *);
int ha_rr_end_session(int, struct ha_em_err_blk *);
int ha_rr_add_var(int, struct ha_rr_variable *, int, int, struct ha_em_err_blk *);
int ha_rr_del_var(int, struct ha_rr_variable *, int, struct ha_em_err_blk *);
int ha_rr_get_interval(char *, struct ha_em_err_blk *);
int ha_rr_send_val(struct ha_rr_val *pv, int, int, struct ha_em_err_blk *);
int ha_rr_reg_var(struct ha_rr_variable *, int, struct ha_em_err_blk *);
int ha_rr_unreg_var(struct ha_rr_variable *, int, struct ha_em_err_blk *);
int ha_rr_terminate(struct ha_em_err_blk *);
int ha_rr_touch(struct ha_em_err_blk *);
#endif /* NO_PROTO */

#endif /* _H_HA_RMAPI */

```

## Related Information

RMAPI subroutines: **ha\_rr\_add\_var**, **ha\_rr\_del\_var**, **ha\_rr\_end\_session**, **ha\_rr\_get\_ctrlmsg**, **ha\_rr\_get\_interval**, **ha\_rr\_init**, **ha\_rr\_makserv**, **ha\_rr\_reg\_var**, **ha\_rr\_send\_val**, **ha\_rr\_start\_session**, **ha\_rr\_terminate**, **ha\_rr\_touch**, **ha\_rr\_unreg\_var**, **ha\_rr\_rm\_ctl**

---

## Chapter 5. Using the RMAPI: Some Resource Monitor Examples

This chapter contains the listings of three sample resource monitors and their associated files:

- “The `rmapi_smpcmd.c` Sample Program” is an example of a command-based resource monitor.
- “The `rmapi_smpdae.c` Sample Program” on page 148 is an example of a daemon-based resource monitor that uses **select** system call notification.
- “The `rmapi_smpsig.c` Sample Program” on page 161 is an example of a daemon-based resource monitor that uses **SIGIO** notification.
- “The `rmapi_smp.msg` Message File” on page 173 is the message source file for the messages used by the examples.
- “The `rmapi_smp.loadsdr` Shell Script” on page 174 and “The `rmapi_smp.unloadsdr` Shell Script” on page 178 are shell scripts that you can use to load the definitions of the resource variables used by the samples into the SDR, and unload them, respectively.

You can find the files for these programs online in the PSSP product samples directory, `/usr/sbin/rsct/samples/haem/rmapi`.

---

### The `rmapi_smpcmd.c` Sample Program

```
/* IBM_PROLOG_BEGIN_TAG                               */
/* This is an automatically generated prolog.         */
/*                                                    */
/*                                                    */
/* Licensed Materials - Property of IBM                */
/*                                                    */
/* (C) COPYRIGHT International Business Machines Corp. 1996,1998 */
/* All Rights Reserved                                */
/*                                                    */
/* US Government Users Restricted Rights - Use, duplication or */
/* disclosure restricted by GSA ADP Schedule Contract with IBM Corp. */
/*                                                    */
/* IBM_PROLOG_END_TAG                                 */
/*=====*/
/* @(#)42 1.9 src/rsct/pem/emtools/rmapi_samples/rmapi_smpcmd.c, emtools, rsct_rtrto 6/5/98 11:25:29 */

#include <unistd.h>
#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <ha_rmapi.h>

/*
 * rmapi_smpcmd.c
 *
 * This program presents an example of using the Resource Monitor Application
 * Programming Interface (RMAPI). The configuration data for this monitor
 * is in the file RmapiSample.loadsdr. This monitor is an example of a command
 * based client monitor (connects to the Event Management daemon) and uses
 * state type variables.
 *
 * Event Management objects defined for this monitor:
 *
 * Monitor:
 *   IBM.PSSP.SampleCmdMon                # client monitor
 *
 * Class:
 *   IBM.PSSP.SampleCmdClass              # variable class
 *
 * Variables:
 *   IBM.PSSP.SampleCmdMon.state          # Valuetype state, datatype long used to
 *                                         # reflect the state of the resource:
 *                                         # Configured, Unconfigured, Running.
 *                                         # Used as locator field of variable.
 *                                         # Name of the resource instance.
 *
 *   Resource Identifiers: NodeNum
 *                           NAME
 *
 *   IBM.PSSP.SampleCmdMon.call           # Valuetype state, datatype SBS used to
 *                                         # reflect the state of the last command
 *                                         # for this resource instance.
 *                                         # Used as locator field of variable.
 *                                         # Name of the resource instance.
 *
 *   SBS Field:      Action              # Datatype long, sn=0.
 *                                         # Last action attempted for this resource.
 *
 *   Options          # Datatype char, sn=1.
 */
```

## rmapi\_smpcmd.c

```
*
*                               # Dummy options string to put into SBS.
*                               # Datatype long, sn=2.
*                               # 0 if action failed, else !0.
*                               # Datatype long, sn=3.
*                               # Final state of the resource (same as the state var).
*
* What this monitor does:  rmapi_smpcmd provides an example of a command which
* controls some imaginary resource. The command is used to change the "state"
* of the resource to one of the following: Unconfigured, Configured, or Running,
* by passing an "action" argument: config, unconfig, start, or stop. One
* or both of the state variables will be updated through the RMAPI on each
* invocation of this command.
*
* Legal state transitions (action):
* Unconfigured -> Configured   (config)
* Configured   -> Unconfigured (unconfig)
* Configured   -> Running      (start)
* Running      -> Configured   (stop)
*
* If a "legal" action is requested, the command will add both of the variables
* for this resource to the Event Management client session. The IBM.PSSP.SampleCmdMon.state
* variable will be added by the call ha_rr_add_var() with a value of TRANSITION and then
* updated to the new state value by ha_rr_send_val().
*
* If the action is "illegal" for the current state, only the IBM.PSSP.SampleCmdMon.call variable
* will be updated so that the IBM.PSSP.SampleCmdMon.state variable remains unchanged.
*
* Parameters:
*
* -a <action>      Required. Possible options are config, unconfig, start or stop. This is
*                  the action requested.
*
* -c <curr_state>  Required. Possible options are Configured, Unconfigured, or Running.
*                  Since this command does no real work, this option tells the command
*                  what "state" the resource is in before attempting the action.
*
* -n <name>        Required. Used as the resource ID NAME element to uniquely
*                  identify different instances of the resource variables.
*
* -o <option>      Optional. Any string, used to demonstrate a char datatype within an SBS
*                  variable.
*
* Examples:
*
* 1. rmapi_smpcmd -a config -c Unconfigured -n test_resource -o test_opt_str
*
* Will cause both variables defined for this monitor to be created
* using "test_resource" as the NAME resource ID. Both variables will
* be registered and added to the RMAPI. The IBM.PSSP.SampleCmdMon.call
* will have the following values sent to the Event Manager daemon on
* the call to ha_rr_add():
*   action:      0           # CONFIG
*   options:     test_opt_str # -o flag parm
*   state_change: 1         # successful state change
*   state:       1           # CONFIGURED
*
* The IBM.PSSP.SampleCmdMon.state variable will have a initial value of TRANSITION
* sent to the RMAPI on ha_rr_add(). It's value will then be updated to CONFIGURED
* and sent to the RMAPI using ha_rr_send_val().
*
* 2. rmapi_smpcmd -a unconfig -c Running -n test_resource
*
* Will cause only the IBM.PSSP.SampleCmdMon.call variable to be created
* using "test_resource" as the NAME resource ID value. The "state"
* variable is not sent to the RMAPI because it was unchanged. IBM.PSSP.SampleCmdMon.call
* will have the following values sent to the Event Manager daemon on
* the call to ha_rr_add():
*   action:      1           # UNCONFIG
*   options:     No options. # default value for -o flag
*   state_change: 0         # state change failed
*   state:       2           # RUNNING
*
* This program can be compiled with the following command:
*
*   cc -O rmapi_smpcmd.c -o rmapi_smpcmd -lha_rr
*/
/* local defines
CMDNAME                - Name of this command.
RESOURCE_MONITOR_NAME  - Name of this monitor as defined in the configuration database.
STATE_VARIABLE_NAME    - Name of the "state" variable defined for this monitor.
CALL_VARIABLE_NAME     - Name of the "call" variable defined for this monitor.
RES_ID                 - Resource ID.
MAX_VARIABLES         - Number of variables defined for this monitor.
DEFAULT_OPTION_STRING  - Default string for -o opt.
USAGE_MSG              - Error msg for cmd syntax errors.
*/
#define CMDNAME          "rmapi_smpcmd"
#define RESOURCE_MONITOR_NAME  "IBM.PSSP.SampleCmdMon"
#define STATE_VARIABLE_NAME    "IBM.PSSP.SampleCmdMon.state"
#define CALL_VARIABLE_NAME     "IBM.PSSP.SampleCmdMon.call"
#define RES_ID            "NAME="
```

```

#define MAX_VARIABLES      2
#define DEFAULT_OPTION_STRING "No options."
#define USAGE_MSG         "Usage: \"CMDNAME\" -a <action> -c <curr_state> -n <name> -o <options>\n\"
                          "\t-a <action>\tResource action (config,unconfig,start or stop).\n\"
                          "\t-c <curr_state>\tCurrent \"state\" of the resource (Configured,Unconfigured,Running).\n\"
                          "\t-n <name>\tName of the resource.\n\"
                          "\t-o <options>\tResource options.\n\"

/* local functions */
void display_rmapi_err(struct ha_em_err_blk *errblk);
enum States process_action(enum States, enum Actions);
void * create_sbs(enum Actions, char *, long, enum States);
void rr_start();
void rr_reg_var();
void rr_add_var();
void rr_del_var();
void rr_end();

/* globals */
struct ha_rr_variable Variables[MAX_VARIABLES]; /* RMAPI variable array. */
struct ha_rr_val Values[MAX_VARIABLES]; /* RMAPI value array. */
int NumVariables = 0; /* Number of RMAPI variables to use. */
struct ha_em_err_blk ErrBlock; /* global error block for RMAPI calls. */
int ClientSock = HA_RR_FAIL; /* Client session socket. */
enum Actions {
    CONFIG,
    UNCONFIG,
    START,
    STOP,
    NUM_ACTIONS
};
char *ActionNames[] = {"config", "unconfig", "start", "stop"};
enum States {
    CONFIGURED,
    UNCONFIGURED,
    RUNNING,
    NUM_STATES,
    TRANSITION
};
char *StateNames[] = {"Configured", "Unconfigured", "Running"};

/*****
/* Main - Main - Main - Main - Main - Main - Main - Main - Main - Main - Main */
*****/
main(int argc, char **argv)
{
    enum States new_state;
    enum States curr_state;
    enum States trans_state = TRANSITION;
    enum Actions action;
    char *action_name = NULL;
    char *curr_state_name = NULL;
    char *res_name = NULL;
    char *res_opts = NULL;
    char *res_id;
    int c, i, rc;

    /*
     * Parse the command line arguments.
     */
    while ((c = getopt(argc, argv, "a:c:n:o:")) != EOF) {
        switch (c) {
            case 'a' : action_name = optarg; break;
            case 'c' : curr_state_name = optarg; break;
            case 'n' : res_name = optarg; break;
            case 'o' : res_opts = optarg; break;
            default :
                fprintf(stderr,USAGE_MSG);
                exit(1);
        }
    }

    /*
     * Check for required options.
     */
    if ((res_name == NULL) || (action_name == NULL) || (curr_state_name == NULL)) {
        fprintf(stderr,"%s: missing -a, -c or -n flag.\n",CMDNAME);
        fprintf(stderr,USAGE_MSG);
        exit(1);
    }

    /*
     * Check for valid action and state options.
     */
    for (i=0; (i < NUM_ACTIONS) && strcmp(ActionNames[i],action_name); i++);
    action = i;

    for (i=0; (i < NUM_STATES) && strcmp(StateNames[i],curr_state_name); i++);
    curr_state = i;

    if ((action == NUM_ACTIONS) || (curr_state == NUM_STATES)) {
        fprintf(stderr,"%s: invalid -a or -c parameter.\n",CMDNAME);
    }
}

```

## rmapi\_smpcmd.c

```
    fprintf(stderr,USAGE_MSG);
    exit(1);
}

/*
 * Make sure an option string is put into the sbs variable.
 */
if (res_opts == NULL) {
    res_opts = DEFAULT_OPTION_STRING;
}

fprintf(stdout,"%s: Current state of resource %s is %s. Requested action is %s.\n",
        CMDNAME,res_name,curr_state_name,action_name);
/*
 * Dummy resource processing. Just checks that the "action" is legal for the
 * current "state" and returns the new state (new_state == curr_state if not
 * a legal action for the current state).
 */
new_state = process_action(curr_state,action);

/*
 * Initialize the RMAPI variables.
 */
memset(Variables, 0 , sizeof(struct ha_rr_variable) * MAX_VARIABLES);
memset(Values, 0 , sizeof(struct ha_rr_val) * MAX_VARIABLES);

/*
 * Create an resource ID for the variables.
 */
resid = malloc(strlen(res_name) + strlen(RES_ID) + 1);
sprintf(resid,"%s%s",RES_ID,res_name);

Variables[0].rr_var_name = CALL_VARIABLE_NAME;
Variables[0].rr_var_rsrc_ID = resid;
Variables[0].rr_varu.rr_var_inst_id = 0;

Variables[1].rr_var_name = STATE_VARIABLE_NAME;
Variables[1].rr_var_rsrc_ID = resid;
Variables[1].rr_varu.rr_var_inst_id = 1;

if (new_state != curr_state) {
    /*
     * The "action" was successful indicated by the state change.
     * Send both the new "state" variable and the "call" variable,
     * otherwise only send the failed "call" variable.
     */
    fprintf(stdout,"%s: Action %s successful. Resource %s will change states %s->TRANSITION->%s.\n",
            CMDNAME,action_name,res_name,curr_state_name,StateNames[new_state]);
    NumVariables = 2;
} else {
    fprintf(stderr,"%s: Cannot perform action %s when in state %s. The "
            "\"state\" variable will not be updated.\n",
            CMDNAME,action_name,curr_state_name);
    NumVariables = 1;
}

/*
 * Initialize the RMAPI and client session.
 */
rr_start();

/*
 * Register variable(s) with the RMAPI.
 */
rr_reg_var();

/*
 * Set the handle addresses and initial values in the
 * RMAPI variables and add them to the client session.
 * The initial value for IBM.PSSP.SampleCmdMon.state is
 * TRANSITION. If the state changed, the new state value
 * will then be sent by the RMAPI routine ha_rr_send_val().
 * If the state did not change, IBM.PSSP.SampleCmdMon will
 * not be added.
 */
Variables[0].rr_varu.rr_var_hdl = &(Values[0].rr_var_hdl);
Variables[0].rr_value = create_sbs(action,res_opts,(new_state==curr_state),new_state);
Variables[1].rr_varu.rr_var_hdl = &(Values[1].rr_var_hdl);
Variables[1].rr_value = &(trans_state);

/*
 * Add the variable(s) to the RMAPI. This will also cause the initial
 * values (Variables[n].rr_value) to be sent to the Event Manager.
 */
rr_add_var();

if (new_state != curr_state) {
    /*
     * State change occurred. ha_rr_add_var() sent the initial value for
     * the "state" variable as being in TRANSITION state. Now call
     * the RMAPI to send the new state of the resource. First check that
     * an error didn't occur when register or adding the variable.
     */

```

```

*/
if ((NumVariables == 2) && (Variables[1].rr_var_errno == 0)) {
    /*
     * Sleep to look like were actually performing
     * an action on the resource.
     */
    sleep(2);

    /*
     * Send the new state value of variable IBM.PSSP.SampleCmdMon.state.
     */
    Values[1].rr_value = &(new_state);
    rc = ha_rr_send_val(&Values[1],1,0,&ErrBlock);
    if (rc == HA_RR_FAIL) {
        display_rmapi_err(&ErrBlock);
    }
}
}

/*
 * Delete the variable(s) that were registered and added.
 */
rr_del_var();

/*
 * Close the EM session and terminate the RMAPI.
 */
rr_end();

/*
 * Exit return code.
 */
exit (new_state==curr_state);
}

/*-----
 * Resource "processing" routine. Returns the new state if the action
 * is valid for the current state, otherwise returns the current state.
 * Legal state transitions:
 * Unconfigured -> Configured (config)
 * Configured -> Unconfigured (unconfig)
 * Configured -> Running (start)
 * Running -> Configured (stop)
 *-----*/
enum States process_action(enum States state, enum Actions action) {
enum States new_state = state;

    switch (state) {
        case CONFIGURED :
            switch (action) {
                case CONFIG :
                case STOP :
                    break;
                case UNCONFIG :
                    new_state = UNCONFIGURED; break;
                case START :
                    new_state = RUNNING; break;
            }
            break;
        case UNCONFIGURED :
            switch (action) {
                case UNCONFIG :
                case START :
                case STOP :
                    break;
                case CONFIG :
                    new_state = CONFIGURED; break;
            }
            break;
        case RUNNING :
            switch (action) {
                case CONFIG :
                case UNCONFIG :
                case START :
                    break;
                case STOP :
                    new_state = CONFIGURED; break;
            }
            break;
    }
    return new_state;
}

/*-----
 * Creates an SBS variable value for the call variable.
 *-----*/
void * create_sbs(enum Actions action, char *opts, long state_chg, enum States state) {
int buff_sz; /* size of the buffer needed for the sbs value. */
int slen; /* length of the opts sbs field. */
char *sbs, *p; /* sbs buffer pointers. */

    /*

```

## rmapi\_smpcmd.c

```
    * Get the length of the "opts" field - including terminating character.
    */
    slen = strlen(opts) + 1;

    /*
    * Calculate the size of the buffer needed:
    * length of the opts field +
    * length of the 3 long fields +
    * 1 long for each sbs field header.
    */
    buff_sz = slen + (sizeof(long) * 7);

    /*
    * Allocate an sbs buffer and prefix the buffer size.
    * Note: The actual buffer is sizeof(long) larger than the
    * calculated length (the SBS is prefixed with a long value
    * which contains the length of the SBS data that follows).
    */
    sbs = p = malloc(buff_sz + sizeof(long));
    *((long *)p) = buff_sz; p += sizeof(long);

    /*
    * Copy the first (Action) sbs field. The field prefixes are a
    * short field length, 1 character for the data type and 1
    * character for the field serial number.
    */
    *((short *)p) = (short)sizeof(long); p += sizeof(short);
    *p = (char)ha_emFTlong; p++;
    *p = 0; p++;
    *((long *)p) = (long)action; p += sizeof(long);

    /*
    * Copy the second (Options) field (sbs serial number 1).
    */
    *((short *)p) = (short)slen; p += sizeof(short);
    *p = (char)ha_emFTchar; p++;
    *p = 1; p++;
    strcpy(p,opts); p += slen;

    /*
    * Copy the third (StateChange) field (sbs serial number 2).
    */
    *((short *)p) = (short)sizeof(long); p += sizeof(short);
    *p = (char)ha_emFTlong; p++;
    *p = 2; p++;
    *((long *)p) = (long)state_chg; p += sizeof(long);

    /*
    * Copy the fourth (State) field (sbs serial number 3).
    */
    *((short *)p) = (short)sizeof(long); p += sizeof(short);
    *p = (char)ha_emFTlong; p++;
    *p = 3; p++;
    *((long *)p) = (long)state; p += sizeof(long);

    /*
    * Return the pointer to the sbs value.
    */
    return sbs;
}

/*-----
 * Initializes the RMAPI and starts the Client session with haemd.
 *-----*/
void rr_start() {
int rc;

    /*
    * Initialize the RMAPI.
    */
    rc = ha_rr_init(RESOURCE_MONITOR_NAME, &ErrBlock);
    if (rc == HA_RR_FAIL) {
        display_rmapi_err(&ErrBlock);
    } else {
        /*
        * Create a client session with the Event Manager daemon.
        */
        ClientSock = ha_rr_start_session(HA_RR_NOTIFY_SELECT, &ErrBlock);
        if (ClientSock == HA_RR_FAIL) {
            display_rmapi_err(&ErrBlock);
        }
    }
}
if ((rc == HA_RR_FAIL) || (ClientSock == HA_RR_FAIL)) {
    /*
    * RMAPI Initialize or start session failed. Set the number of
    * variable to 0 so that RMAPI calls to register, add, send,
    * and delete variables are ignored.
    */
    NumVariables = 0;
}
}
```

```

/*-----
 * Registers the variables to be added with the RMAPI.
 *-----*/
void rr_reg_var() {
int i, num_registered;

if (NumVariables) {
num_registered = ha_rr_reg_var(Variables, NumVariables, &ErrBlock);
if (num_registered != NumVariables) {
if (num_registered == HA_RR_FAIL) {
display_rmapi_err(&ErrBlock);
} else {
for (i=0; i < NumVariables; i++) {
/*
 * Loop through the variables that were attempted
 * to be registered and report any the errors.
 */
if (Variables[i].rr_var_errno) {
fprintf(stderr, "ha_rr_reg_var() error with variable(%s) resource id(%s) "
"instance id(%d) RMAPI errno=(%d).\n",
Variables[i].rr_var_name,
Variables[i].rr_var_rsrc_ID,
Variables[i].rr_var_inst_id,
Variables[i].rr_var_errno);
}
}
}
}
if ((num_registered == HA_RR_FAIL) || (num_registered == 0)) {
/*
 * Register failed - set NumVariables to 0
 * so that ha_rr_add_var() isn't attempted.
 */
NumVariables = 0;
}
} else {
fprintf(stdout, "%s: No variables to register.\n", CMDNAME);
}
}

/*-----
 * Adds variables to the Client session with haemd.
 *-----*/
void rr_add_var() {
int i, num_added;

if (NumVariables) {
num_added = ha_rr_add_var(ClientSock, Variables, NumVariables, 1, &ErrBlock);
if (num_added != NumVariables) {
if (num_added == HA_RR_FAIL) {
display_rmapi_err(&ErrBlock);
} else {
/*
 * Loop through the variables that were attempted
 * to be added and report the errors.
 */
for (i=0; i < NumVariables; i++) {
if (Variables[i].rr_var_errno != 0) {
fprintf(stderr, "Variable(%s) resource id(%s) bad errno(%d) from ha_rr_add_var().\n",
Variables[i].rr_var_name,
Variables[i].rr_var_rsrc_ID,
Variables[i].rr_var_errno);
}
}
}
}
}
if ((num_added == HA_RR_FAIL) || (num_added == 0)) {
/*
 * Add failed - set NumVariables to 0 so that ha_rr_send_var()
 * and ha_rr_del_var() are not called.
 */
NumVariables = 0;
}
} else {
fprintf(stdout, "%s: No variables to add.\n", CMDNAME);
}
}

/*-----
 * Deletes any variables added to the RMAPI for the Client session with haemd.
 *-----*/
void rr_del_var() {
int num_deleted;

if (NumVariables) {
/*
 * Delete any variables that have been added.
 */
num_deleted = ha_rr_del_var(ClientSock, Variables, NumVariables, &ErrBlock);
if (num_deleted == HA_RR_FAIL) {
display_rmapi_err(&ErrBlock);
}
}
}

```

## rmapi\_smpcmd.c

```
    } else {
        fprintf(stdout,"%s: No variables to delete.\n",CMDNAME);
    }
}

/*-----
 * Ends the Client session with haemd (if started) and terminates the RMAPI.
 *-----*/
void rr_end() {
int rc;

    if (ClientSock != HA_RR_FAIL) {
        /*
         * Call the RMAPI to close the client session with the Event Manager.
         */
        rc = ha_rr_end_session(ClientSock, &ErrBlock);
        if (rc == HA_RR_FAIL) {
            display_rmapi_err(&ErrBlock);
        }
    }

    /*
     * Terminate the RMAPI.
     */
    rc = ha_rr_terminate(&ErrBlock);
    if (rc && ErrBlock.em_errno != HA_RR_EACCESS) {
        display_rmapi_err(&ErrBlock);
    }
}

/*-----
 * Displays an RMAPI error.
 *-----*/
void display_rmapi_err(struct ha_em_err_blk *errblk) {
    fprintf(stderr,
        "RMAPI Error: File(%s) Version(%s) Line(%d) Errno(%d)\n\t%s",
        errblk->em_errfile,
        errblk->em_errlevel,
        errblk->em_errline,
        errblk->em_errno,
        errblk->em_errmsg);
}
}
```

---

## The rmapi\_smpdae.c Sample Program

```
/* IBM_PROLOG_BEGIN_TAG                                     */
/* This is an automatically generated prolog.              */
/*                                                         */
/*                                                         */
/* Licensed Materials - Property of IBM                    */
/*                                                         */
/* (C) COPYRIGHT International Business Machines Corp. 1996,1998 */
/* All Rights Reserved                                     */
/*                                                         */
/* US Government Users Restricted Rights - Use, duplication or */
/* disclosure restricted by GSA ADP Schedule Contract with IBM Corp. */
/*                                                         */
/* IBM_PROLOG_END_TAG                                     */
/*=====*/
/* @(#)43 1.9 src/rsct/pem/emtools/rmapi_samples/rmapi_smpdae.c, emtools, rsct_rtro 6/5/98 11:25:36 */

#include <unistd.h>
#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <errno.h>
#include <sys/select.h>
#include <sys/signal.h>
#include <sys/events.h>

#include <ha_rmapi.h>

/*
 * rmapi_smpdae.c
 *
 * This program presents an example of using the Resource Monitor Application
 * Programming Interface (RMAPI). The configuration data for this monitor
 * is in the file RmapiSample.loadsdr. This monitor is an example of server
 * monitor (Resource Monitor Managers connect to the monitor) function that
 * could be incorporated into a daemon or subsystem.
 *
 * Event Management objects defined for this monitor:
 *
 * Monitor:
 * IBM.PSSP.SampleDaeMon # server monitor definition
 *
 * Class:
 * IBM.PSSP.SampleDaeClass # variable class
 */
```

```

*
* Variables:
* IBM.PSSP.SampleDaeMon.StaticVars.static_var1 # valuetype quantity, datatype long
* Resource Identifiers: NodeNum # used as locator field of var
* IBM.PSSP.SampleDaeMon.StaticVars.static_var2 # valuetype quantity, datatype long
* Resource Identifiers: NodeNum # used as locator field of var
* IBM.PSSP.SampleDaeMon.StaticVars.static_var3 # valuetype quantity, datatype long
* Resource Identifiers: NodeNum # used as locator field of var
* IBM.PSSP.SampleDaeMon.InstVars.inst_var1 # valuetype quantity, datatype long
* Resource Identifiers: NodeNum # used as locator field of var
* Name # instance name of the resource
* IBM.PSSP.SampleDaeMon.InstVars.inst_var2 # valuetype quantity, datatype long
* Resource Identifiers: NodeNum # used as locator field of var
* Name # instance name of the resource
*
* What this monitor does: rmap_i_smpdae provides an example of a server monitor
* that has 3 non-instantiable and 2 instantiable Quantity variables which
* are updated with random values on a fixed interval defined in the variable class.
* When the monitor is started, it creates and registers the static and 1 set of the
* instantiable variables. To further demonstrate instantiation, new instances are
* created and registered on the 5th and 12th call to the function send_values().
*
* The monitor will continue executing, adding/deleting and sending values for the
* variables as requested by control messages sent from the Event Management daemon
* through the RMAPI. This monitor uses select notification protocol to determine when a
* connection is ready to accept on the server socket created by ha_rr_makserv() or
* when a message is ready to be read through the RMAPI from one of the monitor's
* client sessions.
*
* The monitor ends execution when it receives a terminating signal, or when all
* all clients have closed their connections.
*
* This program also provides an example of executing multiple copies of a resource
* monitor at the same time. The monitor may be started from the command line, and is
* not configured to be started by resource monitor managers. The following command
* line arguments may be used when starting the monitor:
*
* -i <number> Optional: start the requested number of monitor copies,
* where <number> is in the range 1 - HA_RR_RM_INSTID_MAX.
* -h Displays command line help.
* -H <name> Optional: Specifies the name of the HACMP domain the
* monitor is to execute in.
* -S <name> Optional: Specifies the name of the SP domain (system
* partition name) the monitor is to execute in.
*
* The [-H, -S] options are mutually exclusive. If neither is supplied, the monitor
* will assume it is executing in the default system partition of a SP environment.
*
* The static variables are intended to represent global values which are common
* to all instances of the resource being monitored. Since by definition only one
* instance of the each static variable may exist on a node, they are supplied only
* by the resource monitor whose monitor instance id is 0. This ensures that they are
* available for both performance and event monitoring. Each monitor copy contrives
* unique values for the Name resource ID element of the instantiable variables.
* The intention is to provide an example of multiple copies of a monitor, each
* monitoring a set of instances of the same resource.
*
* Examples:
*
* # rmap_i_smpdae
* - Will start one copy of the monitor in the default system partition of the
* local SP node.
*
* # rmap_i_smpdae -i 4 -P x10s
* - Will start four copies of the monitor which will execute in the system
* partition named "x10s" on the local SP node.
*
* This program can be compiled with the following command:
*
* cc -O rmap_i_smpdae.c -o rmap_i_smpdae -lha_rr
*/
/* local defines
PROGNAME - Name of this program.
MAX_INTERVAL - Max number of seconds before ha_rr_send_value or ha_rr_touch must be called.
RESOURCE_MONITOR_NAME - Name of this monitor as defined in the config data.
MONITOR_CLASS_NAME - Class name as defined in the config data.
SOCKET_TABLE_SIZE - HA_RR_MAX_SESSIONS + 1 for the RMAPI server socket.
RMAPI_SERVER_INDEX - Index in the socket call table for the RMAPI server socket.
STATIC_VAR_NAME_PREFIX - Common portion of the static vars as defined in the config data.
INST_VAR_NAME_PREFIX - Common portion of the inst vars as defined in the config data.
INST_VAR_RESID_FORMAT - Resource ID format for instantiable vars.
INITIAL_VALUE - Variable initial value used on ha_rr_add_var().
MAX_VALUE - Maximum variable value.
NUM_STATIC_VARS - Number of static variables defined for monitor.
NUM_INST_VARS - Number of instantiable variables defined for monitor.
*/
#define PROGNAME "rmap_i_smpdae"
#define MAX_INTERVAL 500
#define RESOURCE_MONITOR_NAME "IBM.PSSP.SampleDaeMon"
#define MONITOR_CLASS_NAME "IBM.PSSP.SampleDaeClass"
#define SOCKET_TABLE_SIZE (HA_RR_MAX_SESSIONS + 1)

```

## rmapi\_smpdae.c

```
#define RMAPI_SERVER_INDEX      (HA_RR_MAX_SESSIONS)
#define STATIC_VAR_NAME_PREFIX "IBM.PSSP.SampleDaeMon.StaticVars.static_var"
#define INST_VAR_NAME_PREFIX  "IBM.PSSP.SampleDaeMon.InstVars.inst_var"
#define INST_VAR_RESID_FORMAT "InstName=Resource%d"
#define INITIAL_VALUE         10
#define MAX_VALUE              500
#define NUM_STATIC_VARS        3
#define NUM_INST_VARS          2

/* local structs */
struct sock_table_entry {
    int socket_fd; /* fd for manager session or RMAPI server socket. */
    void (*sock_funcp)(int); /* function to be called for this socket fd. */
};

struct local_vars {
    char var_name[128]; /* Variable name (as in the CDB). */
    char var_resid[32]; /* Qualified resource ID (for instantiable variables). */
    long value; /* variables value. */
    void *var_handle; /* place for RMAPI to store the var handle. */
};

/* local functions */
void display_rmapi_err(struct ha_em_err_blk *errblk);
void server_socket_handler(int table_index);
void session_socket_handler(int table_index);
void end_session(int table_index);
void register_variables();
void add_variables(int sock_fd, struct ha_rr_ctrl_msg *ctrl_msg);
void del_variables(int sock_fd, struct ha_rr_ctrl_msg *ctrl_msg);
void send_values();
void control_loop();
void set_signals();
void catch_alarm_signal(int);
void catch_exit_signal(int);
void mon_exit(int);

/* globals */
int NumMgrs = 0; /* number of resource managers connected. */
struct ha_rr_variable *Variables = (struct ha_rr_variable *)0; /* RMAPI variable array. */
struct ha_rr_val *Values = (struct ha_rr_val *)0; /* RMAPI value array. */
struct sock_table_entry SocketTable[SOCKET_TABLE_SIZE]; /* RMAPI server and session sockets. */
struct ha_em_err_blk ErrBlock; /* global error block for RMAPI calls. */
struct local_vars *LocalVars = (struct local_vars *)0; /* array of LocalVars variables. */
static sigset_t SignalMask; /* mask for signals. */
timer_t IntervalTimerId; /* ID for var update interval timer. */
int Interval; /* Update interval for this monitor. */
int SampleTime = 0; /* flag set by interval alarm sig. */
int Terminate = 0; /* flag set by terminate signal. */
int SigCaught = 0; /* signal caught by term. */
int TimeOut = MAX_INTERVAL; /* time out counter to call ha_rr_touch. */
int NumVariables = 0; /* Number of variables in arrays. */
int MonitorInstanceID; /* Instance ID of the running monitor */
int Pid; /* resource monitor process id. */
char DomainHACMP[] = "HA_DOMAIN_TYPE=HACMP"; /* string for HACMP domain type. */
char DomainSP[] = "HA_DOMAIN_TYPE=SP"; /* string for HACMP domain type. */
int RmapiInit = 0; /* flag to indicate RMAPI initialized */

/*****
 * Main - Main - Main - Main - Main - Main - Main - Main - Main - Main - Main *
 *****/
main(int argc, char **argv)
{
    int c, i, rc;
    int num_rm_insts = 1;
    struct timer_struct_t itimer;
    char *domain_name = NULL, *domain_type = NULL, *bp;
    struct ha_rr_args rr_args;

    while ((c = getopt(argc, argv, "i:hH:S:")) != EOF) {
        switch (c) {
            case 'i':
                num_rm_insts = atoi(optarg);
                if ((num_rm_insts < 1) || (num_rm_insts > HA_RR_RM_INSTID_MAX)) {
                    fprintf(stderr, "%s(%d): Invalid number of monitor instances, %d. "
                        "Allowable 1-%d.\n",
                        PROGRAMNAME, Pid, num_rm_insts, HA_RR_RM_INSTID_MAX);
                }
                mon_exit(1);
                break;
            case 'H':
                if (domain_name != NULL) {
                    fprintf(stderr, "%s(%d): Specify either the -S or -H options once.\n",
                        PROGRAMNAME, Pid);
                }
                mon_exit(1);
                domain_name = optarg;
                domain_type = DomainHACMP;
                break;
            case 'S':
                if (domain_name != NULL) {
                    fprintf(stderr, "%s(%d): Specify either the -S or -H options once.\n",

```

```

        PROGNAMe,Pid);
    }
    mon_exit(1);
}
domain_name = optarg;
domain_type = DomainSP;
break;
case 'h':
default:
    fprintf(stdout, "Usage:\t%s [-h] [-i number] [-H domain_name | -S domain_name]\n"
        "\t\t-h\tDisplays this help message\n"
        "\t\t-i\tCreate the number of monitor instances specified\n"
        "\t\t-H\tUse the HACMP domain name\n"
        "\t\t-S\tUse the SP domain name\n",
        PROGNAMe);
    exit(0);
}
}

Pid = getpid();
fprintf(stdout, "%s(%d): %s resource monitor started. %d instance(s) of the monitor requested.\n",
    PROGNAMe,Pid,RESOURCE_MONITOR_NAME,num_rm_insts);

/*
 * Set up the domain type and name environment variables (if
 * supplied on the command line). Make the default domain
 * type "SP".
 */
if (domain_type == NULL) domain_type = DomainSP;
if (putenv(domain_type)) {
    fprintf(stderr, "%s(%d): putenv() failed, errno=%d.\n",PROGNAMe,Pid,errno);
    mon_exit(1);
}
if (domain_name != NULL) {
    bp = (Char *)malloc(strlen(domain_name) +
        strlen("HA_DOMAIN_NAME=") + 1);
    if (bp == NULL) {
        fprintf(stderr, "%s(%d): malloc() failed.\n",PROGNAMe,Pid);
        mon_exit(1);
    }
    sprintf(bp,"HA_DOMAIN_NAME=%s",domain_name);
    if (putenv(bp)) {
        fprintf(stderr, "%s(%d): putenv() failed, errno=%d.\n",PROGNAMe,Pid,errno);
        mon_exit(1);
    }
}
fprintf(stdout, "%s(%d): Resource monitor environment: DomainType=%s DomainName=%s.\n",
    PROGNAMe,Pid,getenv("HA_DOMAIN_TYPE"),getenv("HA_DOMAIN_NAME"));

/*
 * Setup signal handlers.
 */
set_signals();

/*
 * Initialize the socket table. Socket file descriptors are initialized to -1.
 */
for (i=0;i<SOCKET_TABLE_SIZE;i++) {
    SocketTable[i].socket_fd = -1;
    if (i == RMAPI_SERVER_INDEX) {
        SocketTable[i].sock_funcp = server_socket_handler;
    } else {
        SocketTable[i].sock_funcp = session_socket_handler;
    }
}

/*
 * Set the monitor instance number. If only 1 instance was specified,
 * request only instance number 0, so that the monitor can supply
 * variables to both the EM daemon and PTPE. If more than 1 copy is
 * to be started, specify HA_RR_RM_INSTID_ANY, so that for each copy,
 * the RMAPI attempts to lock the first available instance id within
 * the range defined for this monitor. Note that this is the default
 * action the RMAPI would take.
 */
rr_args.rr_instance_id = num_rm_insts > 1 ? HA_RR_RM_INSTID_ANY : 0;
rc = ha_rr_rm_ctl(&rr_args,HA_RR_RM_ARGS_SET_INSTID,&ErrBlock);
if (rc == HA_RR_FAIL) {
    display_rmpi_err(&ErrBlock);
    mon_exit(1);
}

fflush(stdout);
if (num_rm_insts > 1) {
    /*
     * fork() multiple copies of the monitor.
     */
    for (i=0;i<num_rm_insts-1;i++) {
        if ((rc = fork()) < 0) {
            /*
             * fork() system call failed.
             */
            fprintf(stderr, "%s(%d): Attempting to fork() monitor copy %d failed. Error: %d-%s.\n",

```

## rmapi\_smpdae.c

```
        PROGNAME,Pid,i,errno,strerror(errno));
    }
    mon_exit(1);
}
if (rc > 0) {
    /*
     * Child RM, break from the loop.
     */
    break;
}
}
}
Pid = getpid();

/*
 * Initialize the RMAPI.
 */
rc = ha_rr_init(RESOURCE_MONITOR_NAME, &ErrBlock);
if (rc == HA_RR_FAIL) {
    display_rmapi_err(&ErrBlock);
    mon_exit(1);
}
RmapiInit = 1;

/*
 * Query and report the monitor instance number and domain name.
 */
rc = ha_rr_rm_ctl(&rr_args,HA_RR_RM_ARGS_GET,&ErrBlock);
if (rc == HA_RR_FAIL) {
    display_rmapi_err(&ErrBlock);
    mon_exit(1);
}
MonitorInstanceID = rr_args.rr_instance_id;

fprintf(stdout,"%s(%d): Resource monitor instance id=%d domain=%s.\n",
        PROGNAME,Pid,rr_args.rr_instance_id,rr_args.rr_domain_name);

/*
 * Query the update interval from the RMAPI.
 */
Interval = ha_rr_get_interval(MONITOR_CLASS_NAME,&ErrBlock);
if (Interval == HA_RR_FAIL) {
    display_rmapi_err(&ErrBlock);
    mon_exit(1);
}
fprintf(stdout,"%s(%d): Interval for class %s is %d.\n",
        PROGNAME,Pid,MONITOR_CLASS_NAME,Interval);

/*
 * Register variables with the RMAPI
 */
register_variables();

/*
 * Make this Resource Monitor a server.
 */
rc = SocketTable[RMAPI_SERVER_INDEX].socket_fd = ha_rr_makserv(HA_RR_NOTIFY_SELECT, &ErrBlock);
if (rc == HA_RR_FAIL) {
    display_rmapi_err(&ErrBlock);
    mon_exit(1);
}
fprintf(stdout,"%s(%d): Assigned RMAPI server socket_fd(%d) to table index(%d).\n",
        PROGNAME,Pid,rc,RMAPI_SERVER_INDEX);

/*
 * Set up sampling timer.
 */
IntervalTimerId = gettimerid(TIMERID_REAL, DELIVERY_SIGNALS);
itimer.it_value.tv_sec = Interval;
itimer.it_value.tv_nsec = 0;
itimer.it_interval.tv_sec = Interval;
itimer.it_interval.tv_nsec = 0;

if (incinterval(IntervalTimerId, &itimer, (struct itimerstruc t *)0) < 0) {
    fprintf(stderr,"%s(%d): incinterval() failed with errno=%d.\n",
        PROGNAME,Pid,errno);
    mon_exit(1);
}

/*
 * Start the control loop - should never return.
 */
control_loop();

/*
 * Should never get here.
 */
mon_exit(0);
}

/*
 * Sets the signal handlers.
 */
```

```

void set_signals() {
struct sigaction sigactn;

sigemptyset(&SignalMask);
sigaddset(&SignalMask, SIGALRM);
sigaddset(&SignalMask, SIGTERM);

/*
 * SIGALRM caught on the interval boundary.
 * Handler sets a global flag to do sampling.
 */
sigactn.sa_handler = catch_alarm_signal;
sigactn.sa_mask = SignalMask;
sigactn.sa_flags = 0;

if (sigaction(SIGALRM, &sigactn, (struct sigaction *)0) < 0) {
    fprintf(stderr,"%s(%d): sigaction() failed for ALRM signal. errno=%d.\n",
        PROGRAMNAME,Pid,errno);
    mon_exit(1);
}

/*
 * Need terminate handler to cleanup RMAPI on exit.
 */
sigactn.sa_handler = catch_exit_signal;
sigactn.sa_mask = SignalMask;
sigactn.sa_flags = 0;

if (sigaction(SIGTERM, &sigactn, (struct sigaction *)0) < 0) {
    fprintf(stderr,"%s(%d): sigaction() failed for TERM signal. errno=%d.\n",
        PROGRAMNAME,Pid,errno);
    mon_exit(1);
}

/*
 * Set signal mask, i.e. block SIGALRM, SIGTERM
 */
if (sigprocmask(SIG_SETMASK, &SignalMask, (sigset_t *)0) < 0) {
    fprintf(stderr,"%s(%d): sigprocmask() failed. errno=%d.\n",PROGRAMNAME,Pid,errno);
    mon_exit(1);
}
}

/*
 * Sends new values to the RMAPI - called when SIGALRM is caught.
 */
void send_values() {
int i, rc;
int num_vals_to_send = 0;
static int times_called = 0;

times_called++;
for (i=0;i<NumVariables;i++) {

    /*
     * Give the variable a new random value.
     */
    LocalVars[i].value = (random() % MAX_VALUE);

    /*
     * Only need to send values for vars with non-NULL handles.
     */
    if (LocalVars[i].var_handle != (void *)0) {

        /*
         * Copy the handle and value to the next available RMAPI structure.
         */
        Values[num_vals_to_send].rr_var_hdl = LocalVars[i].var_handle;
        Values[num_vals_to_send].rr_value = &(LocalVars[i].value);

        /*
         * Increment the number to send count.
         */
        num_vals_to_send++;
    }
}

if (num_vals_to_send) {
    /*
     * Send the Value array to the RMAPI.
     */
    rc = ha_rr_send_val(Values,num_vals_to_send,0,&ErrBlock);
    if (rc == HA_RR_FAIL) {
        display_rmpi_err(&ErrBlock);
        mon_exit(1);
    }

    /*
     * Reset the TimeOut counter.
     */
    TimeOut = MAX_INTERVAL;
}
}

```

## rmapi\_smpdae.c

```
} else if ((TimeOut -= Interval) <= 0) {
    /*
     * No values have been sent in the TimeOut period. Need
     * to call ha_rr_touch to meet server requirements.
     */
    if (ha_rr_touch(&ErrBlock) == HA_RR_FAIL) {
        display_rmapi_err(&ErrBlock);
        mon_exit(1);
    }

    /*
     * Reset the TimeOut counter.
     */
    TimeOut = MAX_INTERVAL;
}

/*
 * Add some new instantiable variables
 * on the 5th and 12th time called.
 */
if ((times_called == 5) || (times_called == 12)) {
    fprintf(stdout,"%s(%d): send_values() called %d times - calling "
           "register_variables() to create new variables.\n",
           PROGNAME,Pid,times_called);
    register_variables();
}
}

/*
 * Catches SIGALRM and sets global flag SampleTime to send values.
 */
void catch_alrm_signal(int sig) {
    SampleTime = 1;
    /*
     * flush stdout/stderr once in a while in
     * case the output has been redirected.
     */
    fflush(stdout);
    fflush(stderr);
}

/*
 * Catches SIGTERM and sets global flag Terminate to exit.
 */
void catch_exit_signal(int sig) {
    Terminate = 1;
    SigCaught = sig;
}

/*
 * Handles a connection request on the RMAPI server
 * socket returned by ha_rr_makserv().
 */
void server_socket_handler(int table_index) {
    int i, mgr_sock, rc;

    mgr_sock = ha_rr_start_session(HA_RR_NOTIFY_SELECT,&ErrBlock);
    if (mgr_sock < 0) {
        display_rmapi_err(&ErrBlock);
        if ((ErrBlock.em_errno != HA_RR_EMASESSIONS) && (ErrBlock.em_errno != HA_RR_EAGAIN)) {
            /*
             * Severe error in the RMAPI - exit.
             */
            mon_exit(1);
        }
        return;
    }
    /*
     * Find an available element in the session socket table.
     */
    for (i=0;(i < HA_RR_MAX_SESSIONS)&&(SocketTable[i].socket_fd >= 0);i++);

    if (i < HA_RR_MAX_SESSIONS) {
        /*
         * Save the mgr socket fd returned by the RMAPI.
         * The socket function was initialized in main.
         */
        SocketTable[i].socket_fd = mgr_sock;
        NumMgrs++;
        fprintf(stdout,"%s(%d): New session accepted to index(%d) socket_fd(%d).\n",
               PROGNAME,Pid,i,mgr_sock);
    } else {
        /*
         * Should never get here, RMAPI handles too many sessions.
         */
        rc = ha_rr_end_session(mgr_sock,&ErrBlock);
    }
}

/*
```

```

* Handles a message being received from the manager session
* specified by [table_index]. Called following select.
*/
void session_socket_handler(int table_index) {
int rc, session_sock;
struct ha_rr_ctrl_msg *ctrl_msg;

/*
 * Get the socket file descriptor for this session.
 */
if ((session_sock = SocketTable[table_index].socket_fd) < 0) {
return;
}

/*
 * Call ha_rr_get_ctrl_msg to read the command from the manager.
 */
fprintf(stdout, "%s(%d): Calling ha_rr_get_ctrlmsg for session(%d) socket_fd(%d).\n",
PROGNAME, Pid, table_index, session_sock);
rc = ha_rr_get_ctrlmsg(session_sock, &ctrl_msg, &ErrBlock);

if (rc == 0) {
/*
 * Message was for the RMAPI or was incomplete.
 */
return;
}

if (rc == HA_RR_FAIL) {
display_rmpi_err(&ErrBlock);
if (ErrBlock.em_errno == HA_RR_EDISCONNECT) {
end_session(table_index);
} else if (ErrBlock.em_errno != HA_RR_EAGAIN) {
/*
 * Severe RMAPI error.
 */
mon_exit(1);
}
return;
}

switch (ctrl_msg->rr_ctrl_cmd) {
case HA_RR_CMD_ADDALL :
case HA_RR_CMD_ADDV :
add_variables(table_index, ctrl_msg);
break;
case HA_RR_CMD_DELALL :
case HA_RR_CMD_DELV :
del_variables(table_index, ctrl_msg);
break;
default :
/*
 * Unknown or unsupported msg - ignore it.
 */
fprintf(stderr, "%s(%d): Received an unexpected cmd(%d) from socket_fd(%d), ignoring...\n",
PROGNAME, Pid, ctrl_msg->rr_ctrl_cmd, session_sock);
break;
}

/*
 * Free the message allocated by the RMAPI.
 */
free(ctrl_msg);
}

/*
 * Adds the variables referenced in the control message to
 * the manager session specified by [table_index].
 */
void add_variables(int table_index, struct ha_rr_ctrl_msg *ctrl_msg) {
int i;
int sock_fd;
int inst_id;
int num_added, num_to_add = 0;

sock_fd = SocketTable[table_index].socket_fd;
switch (ctrl_msg->rr_ctrl_cmd) {
case HA_RR_CMD_ADDALL :
/*
 * Add all variables to this manager session.
 */
fprintf(stdout, "%s(%d): Processing cmd HA_RR_CMD_ADDALL to add all variables to session_fd %d.\n",
PROGNAME, Pid, sock_fd);
for (i=0; i<NumVariables; i++) {
/*
 * Copy the var to the next RMAPI var struct.
 */
Variables[num_to_add].rr_var_name = LocalVars[i].var_name;
Variables[num_to_add].rr_var_rsrc_ID = LocalVars[i].var_resid;
Variables[num_to_add].rr_varu.rr_var_hdl = &(LocalVars[i].var_handle);
Variables[num_to_add].rr_value = &(LocalVars[i].value);
}
}
}

```

## rmapi\_smpdae.c

```

        /*
         * Increment the count of variables to be added.
         */
        num_to_add++;
    }
    break;
case HA_RR_CMD_ADDV :
    /*
     * Add a vector of variables to the manager session.
     */
    fprintf(stdout,"%s(%d): Processing cmd HA_RR_CMD_ADDV to add %d vars for session_fd %d.\n",
            PROGNAME,Pid,ctrl_msg->rr_ctrl_num_vars,sock_fd);
    for (i=0;i < ctrl_msg->rr_ctrl_num_vars;i++) {

        /*
         * The inst id field in the ctrl message is
         * the index into the LocalVars array.
         */
        inst_id = ctrl_msg->rr_ctrlv.rr_ctrl_vari[i];

        if (inst_id < NumVariables) {
            /*
             * Copy the variable to the next RMAPI structure.
             */
            Variables[num_to_add].rr_var_name = LocalVars[inst_id].var_name;
            Variables[num_to_add].rr_var_rsrc_ID = LocalVars[inst_id].var_resid;
            Variables[num_to_add].rr_varu.rr_var_hndl = &(LocalVars[inst_id].var_handle);
            Variables[num_to_add].rr_value = &(LocalVars[inst_id].value);

            num_to_add++;
        }
    }
    break;
default :
    break;
}
if (num_to_add) {
    /*
     * Add the variables by calling ha_rr_add_var.
     */
    num_added = ha_rr_add_var(sock_fd, Variables, num_to_add, 1, &ErrBlock);

    if (num_added == HA_RR_FAIL) {
        display_rmapi_err(&ErrBlock);

        if (ErrBlock.em_errno == HA_RR_EDISCONNECT) {
            /*
             * Session closed by manager.
             */
            end_session(table_index);
        } else {
            mon_exit(1);
        }
    }
} else {
    fprintf(stdout,"%s(%d): ha_rr_add_var() added %d vars to session_fd %d.\n",
            PROGNAME,Pid,num_added,sock_fd);
    if (num_to_add != num_added) {
        /*
         * Loop through the variables that were attempted
         * to be added and report any that had errors.
         */
        for (i=0;i<num_to_add;i++) {
            if (Variables[i].rr_var_errno != 0) {
                fprintf(stderr,"%s(%d): Variable(%s) resource ID(%s) had bad errno(%d) from ha_rr_add_var().\n",
                        PROGNAME,Pid,Variables[i].rr_var_name,
                        Variables[i].rr_var_rsrc_ID,Variables[i].rr_var_errno);
            }
        }
    }
}
}
}

/*
 * Deletes the variables referenced in the ctrl_msg for the manager
 * session specified by [table_index].
 */
void del_variables(int table_index, struct ha_rr_ctrl_msg *ctrl_msg) {
    int i;
    int sock_fd;
    int ctrl_cmd;
    int inst_id;
    int num_deleted, num_to_del = 0;

    sock_fd = SocketTable[table_index].socket_fd;
    if (ctrl_msg == (struct ha_rr_ctrl_msg *)0) {
        /*
         * Called with NULL message when a session is ending. Need to delete
         * all variables from the session before calling ha_rr_end_session.
         */
    }
}

```

```

    ctrl_cmd = HA_RR_CMD_DELALL;
} else {
    ctrl_cmd = ctrl_msg->rr_ctrl_cmd;
}

switch (ctrl_cmd) {
case HA_RR_CMD_DELALL :
    fprintf(stdout,"%s(%d): Processing cmd HA_RR_CMD_DELALL to delete all vars for session_fd %d.\n",
        PROGNAME,Pid,sock_fd);
    /*
     * Delete all variables for this manager session.
     */
    for (i=0;i<NumVariables;i++) {
        if (LocalVars[i].var_handle != (void *)0) {
            /*
             * Copy the variable to the next RMAPI structure.
             */
            Variables[num_to_del].rr_varu.rr_var_hndl = &(LocalVars[i].var_handle);

            /*
             * Increment the count of variables to be deleted.
             */
            num_to_del++;
        }
    }
    break;
case HA_RR_CMD_DELV :
    fprintf(stdout,"%s(%d): Processing cmd HA_RR_CMD_DELV to delete %d vars for session_fd %d.\n",
        PROGNAME,Pid,ctrl_msg->rr_ctrl_num_vars,sock_fd);
    /*
     * Delete a vector of variables for this manager session.
     */
    for (i = 0; i < ctrl_msg->rr_ctrl_num_vars; i++) {
        inst_id = ctrl_msg->rr_ctrlv.rr_ctrl_vari[i];
        if ((inst_id < NumVariables) &&
            (LocalVars[inst_id].var_handle != (void *)0)) {
            /*
             * Copy the variable to the next RMAPI structure.
             */
            Variables[num_to_del].rr_varu.rr_var_hndl = &(LocalVars[inst_id].var_handle);

            /*
             * Increment the count of variables to be deleted.
             */
            num_to_del++;
        }
    }
    break;
default :
    break;
}

if (num_to_del) {
    /*
     * Delete the variables by calling ha_rr_del_var.
     */
    num_deleted = ha_rr_del_var(sock_fd, Variables, num_to_del, &ErrBlock);

    if (num_deleted == HA_RR_FAIL) {
        display_rmapl_err(&ErrBlock);

        if (ErrBlock.em_errno == HA_RR_EDISCONNECT) {
            /*
             * Session closed by manager.
             */
            end_session(table_index);
        } else {
            mon_exit(1);
        }
    }
}

/*
 * rc>0 from ha_rr_del_var is the number of variables that no longer need
 * their values updated. The RMAPI will have set their handles to NULL.
 */
fprintf(stdout,"%s(%d): %d variables no longer need to be updated.\n",
    PROGNAME,Pid,num_deleted);
}

/*
 * Ends the manager session indexed by the [table_index] parameter
 * in the socket table.
 */
void end_session(int table_index) {
int rc;

    fprintf(stdout,"%s(%d): Ending session index(%d) session_fd(%d).\n",
        PROGNAME,Pid,table_index,SocketTable[table_index].socket_fd);

    /*
     * Make sure this is a valid session.

```

## rmapi\_smpdae.c

```
    /*
    if ((table_index < 0) || (table_index > HA_RR_MAX_SESSIONS) ||
        (SocketTable[table_index].socket_fd < 0)) {
        return;
    }

    /*
    * Delete all variables for this manager.
    */
    del_variables(table_index, (struct ha_rr_ctrl_msg *)0);

    /*
    * Call RMAPI to end the session.
    */
    rc = ha_rr_end_session(SocketTable[table_index].socket_fd, &ErrBlock);
    if (rc == HA_RR_FAIL) {
        display_rmapi_err(&ErrBlock);
        monitor_exit(1);
    }

    /*
    * Reset the socket file descriptor in the SocketTable and
    * decrement the manager count.
    */
    SocketTable[table_index].socket_fd = -1;
    NumMgrs--;
}

/*
* Registers variables with the RMAPI.
*/
void register_variables() {
    int i, j, num_var;
    int num_registered, num_to_reg = 0;
    static int times_called = 0;

    fprintf(stdout, "%s(%d): Call number %d to register_variables().\n",
            PROGNAME, Pid, times_called);
    times_called++;
    if (times_called == 1) {
        /*
        * First time called, allocate the LocalVars, Variables and
        * Value arrays. Only register the StaticVars if this is
        * resource monitor instance 0. Since only 1 monitor will
        * supply these non-instantiable common attributes, we want it to
        * be instance 0, so that the variables are made available to
        * both the EM daemon and PTPE.
        */
        NumVariables = NUM_INST_VARS;

        if (MonitorInstanceID == 0) NumVariables += NUM_STATIC_VARS;

        fprintf(stdout, "%s(%d): Creating %d initial variables.\n",
                PROGNAME, Pid, NumVariables);

        LocalVars = (struct local_vars *)malloc(sizeof(struct local_vars) * NumVariables);
        Variables = (struct ha_rr_variable *)malloc(sizeof(struct ha_rr_variable) * NumVariables);
        Values = (struct ha_rr_val *)malloc(sizeof(struct ha_rr_val) * NumVariables);

        /*
        * Initialize the arrays.
        */
        memset(LocalVars, 0, sizeof(struct local_vars) * NumVariables);
        memset(Variables, 0, sizeof(struct ha_rr_variable) * NumVariables);
        memset(Values, 0, sizeof(struct ha_rr_val) * NumVariables);

        if (MonitorInstanceID == 0) {
            /*
            * Initialize the local variable structures and copy the name, resource ID
            * and value pointers to the RMAPI variables for registration.
            */
            for (i=0; i<NUM_STATIC_VARS; i++) {

                /*
                * Create the non-instantiable variable names. The resource IDs are NULL
                * strings because they are not instantiable. The only resource ID, NodeNum,
                * defined for these variables is used as the locator field and does not
                * need to be sent to the RMAPI. The index of the variable in the local
                * is used as the inst_id (Variables[n].rr_var iid). This id is returned by
                * resource monitor managers in control messages to identify variables.
                */
                sprintf(LocalVars[num_to_reg].var_name, "%s%d", STATIC_VAR_NAME_PREFIX, i+1);
                LocalVars[num_to_reg].var_resid[0] = '\0';
                LocalVars[num_to_reg].value = INITIAL_VALUE;

                Variables[num_to_reg].rr_var_name = LocalVars[num_to_reg].var_name;
                Variables[num_to_reg].rr_var_rsrc_ID = LocalVars[num_to_reg].var_resid;
                Variables[num_to_reg].rr_var_iid = num_to_reg;

                num_to_reg++;
            }
        }
    }
}
```

```

for (i=0;i<NUM_INST_VARS;i++) {
    /*
     * Initialize the instantiable variables. The difference here is that
     * they have a resource ID that needs to be supplied. The resource ID
     * value will be InstSet<n> where <n> is the number of the call to
     * this routine plus the instance id of the monitor times 100. There
     * is no significance to this naming scheme - it is only used to
     * generate unique instances between multiple copies of the monitor.
     */
    sprintf(LocalVars[num_to_reg].var_name,"%s%d",INST_VAR_NAME_PREFIX,i+1);
    sprintf(LocalVars[num_to_reg].var_resid,INST_VAR_RESID_FORMAT,
            times_called + (MonitorInstanceID * 100));
    LocalVars[num_to_reg].value = INITIAL_VALUE;

    Variables[num_to_reg].rr_var_name = LocalVars[num_to_reg].var_name;
    Variables[num_to_reg].rr_var_rsrc_ID = LocalVars[num_to_reg].var_resid;
    Variables[num_to_reg].rr_var_iid = num_to_reg;

    num_to_reg++;
}
} else {
    /*
     * Subsequent call to this routine. This is an example of new instantiations
     * being created during normal execution. A new set of instantiable variables
     * will be registered using the current number of times this routine was called
     * in the resource ID.
     */
    fprintf(stdout,"%s(%d): Creating %d new variables.\n",
            PROGNAME,Pid,NUM_INST_VARS);

    /*
     * Get the new number of variables.
     */
    num_var = NumVariables + NUM_INST_VARS;

    /*
     * realloc the arrays to the new size and initialize the local array.
     */
    LocalVars = (struct local_vars *)realloc(LocalVars, sizeof(struct local_vars) * num_var);
    Variables = (struct ha_rr_variable *)realloc(Variables, sizeof(struct ha_rr_variable) * num_var);
    Values = (struct ha_rr_val *)realloc(Values, sizeof(struct ha_rr_val) * num_var);
    memset(LocalVars+NumVariables,0,sizeof(struct local_vars) * NUM_INST_VARS);

    for (i=NumVariables,j=0;i<num_var;i++,j++) {
        /*
         * Create a new set of variables as above and copy them to the
         * RMAPI variable structure to be registered.
         */
        sprintf(LocalVars[i].var_name,"%s%d",INST_VAR_NAME_PREFIX,j+1);
        sprintf(LocalVars[i].var_resid,INST_VAR_RESID_FORMAT,
                times_called + (MonitorInstanceID * 100));
        LocalVars[i].value = INITIAL_VALUE;

        Variables[num_to_reg].rr_var_name = LocalVars[i].var_name;
        Variables[num_to_reg].rr_var_rsrc_ID = LocalVars[i].var_resid;
        Variables[num_to_reg].rr_var_iid = i;

        num_to_reg++;
    }
    NumVariables = num_var;
}

if (num_to_reg) {
    /*
     * Call the RMAPI to register the variables.
     */
    num_registered = ha_rr_reg_var(Variables,num_to_reg,&ErrBlock);
    if (num_registered != num_to_reg) {
        if (num_registered == HA_RR_FAIL) {
            display_rmpi_err(&ErrBlock);
            mon_exit(1);
        }
        fprintf(stdout,"%s(%d): %d variables registered with the RMAPI.\n",
                PROGNAME,Pid,num_registered);
        for (i=0; i<num_to_reg; i++) {
            if ((Variables[i].rr_var_errno) || (num_registered == HA_RR_FAIL)) {
                /*
                 * Display a message for each variable with a registration error.
                 */
                fprintf(stderr,"%s(%d): ha_rr_reg_var() error with variable(%s) instid(%d) RMAPI errno=(%d).\n",
                        PROGNAME,Pid,
                        Variables[i].rr_var_name,
                        Variables[i].rr_varu.rr_var_inst_id,
                        Variables[i].rr_var_errno);
            }
        }
    }
}
}

```

## rmapi\_smpdae.c

```
    }  
}  
  
/*  
 * Called once by main to establish the control loop.  
 */  
void control_loop() {  
    int i, rc;  
    int sock;  
    int CallSockFunctions;  
    fd_set sockfd_set;  
  
    for(;;) {  
  
        /*  
         * Initialize the socket mask for select.  
         */  
        FD_ZERO(&sockfd_set);  
        for(i=0;i<SOCKET_TABLE_SIZE;i++) {  
            if (SocketTable[i].socket_fd >= 0) FD_SET(SocketTable[i].socket_fd,&sockfd_set);  
        }  
        CallSockFunctions = 0;  
  
        /*  
         * Unblock the SIGTERM and SIGALRM signals.  
         */  
        if (sigprocmask(SIG_UNBLOCK, &SignalMask, (sigset_t *)0) < 0) {  
            /*  
             * sigprocmask sys call failed.  
             */  
            fprintf(stderr,"%s(%d): Cannot unblock signals. sigprocmask() errno=%d.\n",  
                PROGRAMME,Pid,errno);  
            mon_exit(1);  
        }  
  
        /*  
         * Only select if not terminating or SIGALRM caught to sample values.  
         */  
        if (!(SampleTime || Terminate)) {  
            sock = select(FD_SETSIZE, &sockfd_set, NULL, NULL, NULL);  
            if (sock > 0) {  
                CallSockFunctions = 1;  
            } else if (sock < 0) {  
                /*  
                 * select() call failed - exit if not an interrupt.  
                 */  
                if (errno != EINTR) {  
                    fprintf(stderr,"%s(%d): select() failed with errno=%d.\n",  
                        PROGRAMME,Pid,errno);  
                    mon_exit(1);  
                }  
                continue;  
            }  
        }  
  
        /*  
         * Block the SIGTERM and SIGALRM sigs.  
         */  
        if (sigprocmask(SIG_BLOCK, &SignalMask, (sigset_t *)0) < 0) {  
            fprintf(stderr,"%s(%d): Cannot block signals. sigprocmask() errno=%d.\n",  
                PROGRAMME,Pid,errno);  
            mon_exit(1);  
        }  
  
        if (Terminate) {  
            /*  
             * SIGTERM caught - call mon_exit() to clean up.  
             */  
            fprintf(stdout,"%s(%d): Caught signal(%d)...calling exit...\n",  
                PROGRAMME,Pid,SigCaught);  
            mon_exit(SigCaught);  
        }  
  
        if (CallSockFunctions) {  
            /*  
             * Call the socket function for each socket returned by select.  
             */  
            for (i=0;<SOCKET_TABLE_SIZE;i++) {  
                if (SocketTable[i].socket_fd >=0 &&  
                    FD_ISSET(SocketTable[i].socket_fd,&sockfd_set)) {  
                    fprintf(stdout,"%s(%d): Calling socket function for session(%d) socket_fd(%d).\n",  
                        PROGRAMME,Pid,i,SocketTable[i].socket_fd);  
                    (SocketTable[i].sock_funcp)(i);  
                }  
            }  
        }  
  
        if (SampleTime) {  
            /*  
             * SIGALRM caught - call send_values()  
             * and reset the sample flag.  
             */  
        }  
    }  
}
```

```

        send_values();
        SampleTime = 0;
    }
}

/*
 * Routine to terminate the monitor.
 */
void mon_exit(int s)
{
static int recursively_called = 0;
int i, rc;

/*
 * Check for recursive call - some routines mon_exit()
 * calls can likewise call mon_exit().
 */
if (recursively_called) return;
recursively_called = 1;

if (RmapiInit) {
/*
 * Gracefully close all sessions.
 */
for (i=0; i<HA_RR_MAX_SESSIONS; i++) {
    if (SocketTable[i].socket_fd >= 0) {
        end_session(i);
    }
}
/*
 * Terminate the RMAPI.
 */
rc = ha_rr_terminate(&ErrBlock);
if (rc) {
    display_rmapi_err(&ErrBlock);
}
}

fprintf(stdout, "%s(%d): %s resource monitor exiting.\n",
        PROGRAMNAME, Pid, RESOURCE_MONITOR_NAME);
exit(s);
}

/*
 * Display an RMAPI error.
 */
void display_rmapi_err(struct ha_em_err_blk *errblk)
{
    fprintf(stderr,
        "%s(%d): RMAPI Error: File(%s) Version(%s) Line(%d) Errno(%d)\n\t%s",
        PROGRAMNAME, Pid,
        errblk->em_errfile,
        errblk->em_errlevel,
        errblk->em_errline,
        errblk->em_errno,
        errblk->em_errmsg);
}

```

---

## The rmapi\_smpsig.c Sample Program

```

/* IBM_PROLOG_BEGIN_TAG                               */
/* This is an automatically generated prolog.         */
/*                                                    */
/*                                                    */
/* Licensed Materials - Property of IBM                */
/*                                                    */
/* (C) COPYRIGHT International Business Machines Corp. 1996,1998 */
/* All Rights Reserved                                */
/*                                                    */
/* US Government Users Restricted Rights - Use, duplication or */
/* disclosure restricted by GSA ADP Schedule Contract with IBM Corp. */
/*                                                    */
/* IBM_PROLOG_END_TAG                                   */
/*=====*/
/* @(#)44 1.9 src/rsct/pem/emtools/rmapi_samples/rmapi_smpsig.c, emtools, rsct_rtr0 6/5/98 11:25:45 */

#include <unistd.h>
#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <errno.h>
#include <signal.h>
#include <sys/types.h>
#include <sys/events.h>
#include <ha_rmapi.h>

/*
 * rmapi_smpsig.c

```

## rmapi\_smplib.c

```
*
* This program presents an example of using the Resource Monitor Application
* Programming Interface (RMAPI). The configuration data for this monitor
* is in the file RmapiSample.loadsdr. This monitor is an example of server
* monitor (Resource Monitor Managers connect to the monitor) function that could
* be incorporated into a daemon or subsystem. This example monitor demonstrates
* using SIGIO as the notification protocol and the use of dynamically instantiable
* variables.
*
* Event Management objects defined for this monitor:
*
* Monitor:
*   IBM.PSSP.SampleSigMon          # server monitor definition
*
* Class:
*   IBM.PSSP.SampleSigDynInstClass # class for dynamically instantiable vars
*   IBM.PSSP.SampleSigNonInstClass # class for non-instantiable vars
*
* Variables:
*   IBM.PSSP.SampleSigMon.DynInstVar.var # valuetype quantity, datatype long
*   Resource Identifiers: NodeNum       # used as locator field of var
*   InstName                             # instance name of variable
*   IBM.PSSP.SampleSigMon.NonInstVars.var1 # valuetype counter, datatype long
*   Resource Identifiers: NodeNum       # used as locator field of var
*   IBM.PSSP.SampleSigMon.NonInstVars.var2 # valuetype counter, datatype long
*   Resource Identifiers: NodeNum       # used as locator field of var
*   IBM.PSSP.SampleSigMon.NonInstVars.var3 # valuetype counter, datatype long
*   Resource Identifiers: NodeNum       # used as locator field of var
*
* What this monitor does: rmapi_smplib provides an example of a server monitor
* which has 1 dynamically instantiable variable and 3 non-instantiable variables.
*
* The monitor will continue to run until it either receives a signal to terminate, or
* all clients have closed their connections.
*
* This program can be compiled with the following command:
*
*   cc -O rmapi_smplib.c -o rmapi_smplib -lha_rr
*/

#ifndef MIN
#define MIN(x, y) ((x) < (y) ? (x) : (y))
#endif

/* local defines
PROGNAME          - Name of this program.
MAX_INTERVAL     - Max number of seconds before ha_rr_send_value or ha_rr_touch must be called.
RESOURCE_MONITOR_NAME - Name of this monitor as defined in the config data.
MONITOR_CLASS_NAME - Class name as defined in the config data.
SOCKET_TABLE_SIZE - HA_RR_MAX_SESSIONS + 1 for the RMAPI server socket.
RMAPI_SERVER_INDEX - Index in the socket call table for the RMAPI server socket.
STATIC_VAR_NAME_PREFIX - Common portion of the static vars as defined in the config data.
INST_VAR_NAME_PREFIX - Common portion of the inst vars as defined in the config data.
INST_VAR_RESID_FORMAT - Resource ID format for instantiable vars.
INITIAL_VALUE    - Variable initial value used on ha_rr_add_var().
MAX_VALUE        - Maximum variable value.
NUM_STATIC_VARS  - Number of static variables defined for monitor.
NUM_INST_VARS    - Number of instantiable variables defined for monitor.
*/
#define PROGNAME          "rmapi_smplib"
#define MAX_INTERVAL     500
#define RESOURCE_MONITOR_NAME "IBM.PSSP.SampleSigMon"
#define DYN_INST_CLASS_NAME "IBM.PSSP.SampleSigDynInstClass"
#define NON_INST_CLASS_NAME "IBM.PSSP.SampleSigNonInstClass"
#define NUM_NONINST_VARS 3
#define SOCKET_TABLE_SIZE (HA_RR_MAX_SESSIONS + 1)
#define RMAPI_SERVER_INDEX (HA_RR_MAX_SESSIONS)
#define DYN_INST_VAR_NAME "IBM.PSSP.SampleSigMon.DynInstVar.var"
#define NON_INST_VAR_NAME_PREFIX "IBM.PSSP.SampleSigMon.NonInstVars.var"
#define DYN_INST_VAR_RESID "InstName"
#define INITIAL_VALUE    10
#define MAX_VALUE        500
#define NULL_RESID      ""

/* local structs */
struct sock_table_entry {
    int socket_fd; /* fd for manager session or RMAPI server socket. */
    void (*sock_func)(int); /* function to be called for this socket fd. */
};

struct local_vars {
    char var_name[128]; /* Variable name (as in config data). */
    char var_resid[32]; /* Qualified resource ID (for instantiable variables). */
    short registered; /* flag set when variable is registered. */
    long value; /* variables value. */
    void *var_handle; /* place for RMAPI to store the var handle. */
};

/* local functions */
void display_rmapi_err(struct ha_em_err_blk *errblk);
void server_socket_handler(int table_index);
void session_socket_handler(int table_index);
```

```

void end_session(int table_index);
void register_variables();
void inst_variables(struct ha_rr_ctrl_msg *);
void add_variables(int sock_fd, struct ha_rr_ctrl_msg *ctrl_msg);
void del_variables(int sock_fd, struct ha_rr_ctrl_msg *ctrl_msg);
void send_values();
void control_loop();
void set_signals();
void catch_io_signal(int);
void catch_alarm_signal(int);
void catch_exit_signal(int);
void mon_exit(int);

/* globals */
int NumMgrs = 0; /* number of resource managers connected. */
struct ha_rr_variable *Variables = (struct ha_rr_variable *)0; /* RMAPI variable array. */
struct ha_rr_val *Values = (struct ha_rr_val *)0; /* RMAPI val array. */
struct sock_table_entry SocketTable[SOCKET_TABLE_SIZE]; /* RMAPI sever and session sockets. */
struct ha_em_err_blk ErrBlock; /* global error block for RMAPI calls. */
struct local_vars *LocalVars = (struct local_vars *)0; /* array of local variables. */
static sigset_t SignalMask; /* mask for signals. */
timer_t IntervalTimerId; /* ID for var update interval timer. */
int Interval; /* Update interval for this monitor. */
int SampleTime = 0; /* flag set by catching interval SIGALRM. */
int Terminate = 0; /* flag set by catching SIGTERM. */
int SigioCaught = 0; /* flag set by catching SIGIO. */
int SigCaught = 0; /* actual signal caught by term. */
int TimeOut = MAX_INTERVAL; /* time out counter to call ha_rr_touch. */
int NumVariables = 0; /* Number of variables available. */
int RmaplInit = 0; /* flag to indicate RMAPI initialized. */

/*****
/* Main - Main - Main - Main - Main - Main - Main - Main - Main - Main - Main */
/*****
main(int argc, char **argv)
{
    int interval1, interval2;
    int i, rc;
    struct itimerstruc_t itimer;

    fprintf(stdout, "%s: %s resource monitor started.\n", PROGNAME, RESOURCE_MONITOR_NAME);

    /*
     * Setup signal handlers.
     */
    set_signals();

    /*
     * Initialize the socket table. Socket file descriptors are set to -1.
     */
    for (i=0; i<SOCKET_TABLE_SIZE; i++) {
        SocketTable[i].socket_fd = -1;
        if (i == RMAPI_SERVER_INDEX) {
            SocketTable[i].sock_funcp = server_socket_handler;
        } else {
            SocketTable[i].sock_funcp = session_socket_handler;
        }
    }

    /*
     * Initialize the RMAPI.
     */
    rc = ha_rr_init(RESOURCE_MONITOR_NAME, &ErrBlock);
    if (rc == HA_RR_FAIL) {
        display_rmapl_err(&ErrBlock);
        mon_exit(1);
    }
    RmaplInit = 1;

    /*
     * Query the update intervals for each class from the RMAPI.
     */
    interval1 = ha_rr_get_interval(DYN_INST_CLASS_NAME, &ErrBlock);
    if (interval1 <= 0) {
        display_rmapl_err(&ErrBlock);
        mon_exit(1);
    }
    fprintf(stdout, "Interval for class %s is %d.\n", DYN_INST_CLASS_NAME, interval1);

    interval2 = ha_rr_get_interval(NON_INST_CLASS_NAME, &ErrBlock);
    if (interval2 <= 0) {
        display_rmapl_err(&ErrBlock);
        mon_exit(1);
    }
    fprintf(stdout, "Interval for class %s is %d.\n", NON_INST_CLASS_NAME, interval2);

    /*
     * For simplicity, the real interval will be the smallest of the two.
     */
    Interval = MIN(interval1, interval2);

    /*

```

## rmapi\_smpsig.c

```
    * Allocate the local and RMAPI arrays for the non-instantiable variables.
    */
    NumVariables = NUM_NONINST_VARS;
    LocalVars = (struct local_vars *)malloc(sizeof(struct local_vars) * NumVariables);
    Variables = (struct ha_rr_variable *)malloc(sizeof(struct ha_rr_variable) * NumVariables);
    Values = (struct ha_rr_val *)malloc(sizeof(struct ha_rr_val) * NumVariables);

    /*
    * Initialize the arrays.
    */
    memset(LocalVars,0,sizeof(struct local_vars) * NumVariables);
    memset(Variables,0,sizeof(struct ha_rr_variable) * NumVariables);
    memset(Values,0,sizeof(struct ha_rr_val) * NumVariables);

    /*
    * Initialize the local variable structure for the non-instantiables.
    */
    for (i=0;i<NUM_NONINST_VARS;i++) {
        sprintf(LocalVars[i].var_name,"%s%d",NON_INST_VAR_NAME_PREFIX,i+1);
        LocalVars[i].var_resid[0] = '\0';
        LocalVars[i].value = 0;
    }

    /*
    * Register the non-instantiable variables with the RMAPI.
    */
    register_variables();

    /*
    * Make this Resource Monitor a server.
    */
    rc = SocketTable[RMAPI_SERVER_INDEX].socket_fd = ha_rr_makserv(HA_RR_NOTIFY_SIGIO, &ErrBlock);
    if (rc == HA_RR_FAIL) {
        display_rmapi_err(&ErrBlock);
        mon_exit(1);
    }
    fprintf(stdout,"Assigned rmapi server socket_fd(%d) to table index(%d).\n",
            rc,RMAPI_SERVER_INDEX);

    /*
    * Set up sampling timer.
    */
    IntervalTimerId = gettimerid(TIMERID_REAL, DELIVERY_SIGNALS);
    itimer.it_value.tv_sec = Interval;
    itimer.it_value.tv_nsec = 0;
    itimer.it_interval.tv_sec = Interval;
    itimer.it_interval.tv_nsec = 0;

    if (incinterval(IntervalTimerId, &itimer, (struct itimerstruct *)0) < 0) {
        fprintf(stderr,"incinterval() failed with errno=%d.\n",errno);
        mon_exit(1);
    }

    /*
    * Start the control loop - should never return.
    */
    control_loop();

    /*
    * Should never get here.
    */
    mon_exit(0);
}

/*
* Sets the signal handlers.
*/
void set_signals() {
    struct sigaction sigactn;

    sigemptyset(&SignalMask);
    sigaddset(&SignalMask, SIGALRM);
    sigaddset(&SignalMask, SIGTERM);
    sigaddset(&SignalMask, SIGIO);

    /*
    * SIGIO signal handler.
    */
    sigactn.sa_handler = catch_io_signal;
    sigactn.sa_mask = SignalMask;
    sigactn.sa_flags = 0;

    if (sigaction(SIGIO, &sigactn, (struct sigaction *)0) < 0) {
        fprintf(stderr,"sigaction() failed for SIGIO signal. errno=%d.\n",errno);
        mon_exit(1);
    }

    /*
    * SIGALRM caught on the interval boundary.
    * Handler sets a global flag to do sampling.
    */
    sigactn.sa_handler = catch_alrm_signal;
```

```

sigactn.sa_mask = SignalMask;
sigactn.sa_flags = 0;

if (sigaction(SIGALRM, &sigactn, (struct sigaction *)0) < 0) {
    fprintf(stderr, "sigaction() failed for ALRM signal. errno=%d.\n", errno);
    mon_exit(1);
}

/*
 * Need a terminate handler to cleanup RMAPI on exit.
 */
sigactn.sa_handler = catch_exit_signal;
sigactn.sa_mask = SignalMask;
sigactn.sa_flags = 0;

if (sigaction(SIGTERM, &sigactn, (struct sigaction *)0) < 0) {
    fprintf(stderr, "sigaction() failed for TERM signal. errno=%d.\n", errno);
    mon_exit(1);
}

/*
 * Set signal mask, i.e. block SIGIO, SIGALRM, SIGTERM
 */
if (sigprocmask(SIG_SETMASK, &SignalMask, (sigset_t *)0) < 0) {
    fprintf(stderr, "sigprocmask() failed. errno=%d.\n", errno);
    mon_exit(1);
}
}

/*
 * Sends new values to the RMAPI - called when SIGALRM is caught.
 */
void send_values() {
    int i, rc;
    int num_vals_to_send = 0;
    static int times_called = 0;

    times_called++;
    for (i=0; i<NumVariables; i++) {
        if (i < NUM_NONINST_VARS) {
            /*
             * Randomly increment the non-instantiable variable value (Counter).
             */
            LocalVars[i].value += (random() % 50);
        } else {
            /*
             * Assign a random value to the DynInstVar variable value (Quantity).
             */
            LocalVars[i].value = (random() % MAX_VALUE);
        }

        /*
         * Only need to send values for variables with non-NULL handles.
         */
        if (LocalVars[i].var_handle != (void *)0) {
            /*
             * Copy the handle and value to the next available RMAPI structure.
             */
            Values[num_vals_to_send].rr_var_hdl = LocalVars[i].var_handle;
            Values[num_vals_to_send].rr_value = &(LocalVars[i].value);

            /*
             * Increment the number to send count.
             */
            num_vals_to_send++;
        }
    }

    if (num_vals_to_send) {
        /*
         * Send the Value array to the RMAPI.
         */
        rc = ha_rr_send_val(Values, num_vals_to_send, 0, &ErrBlock);
        if (rc == HA_RR_FAIL) {
            display_rmapi_err(&ErrBlock);
            mon_exit(1);
        }

        /*
         * Reset the Timeout counter.
         */
        Timeout = MAX_INTERVAL;
    } else if ((Timeout -= Interval) <= 0) {
        /*
         * No values have been sent in the Timeout period. Need
         * to call ha_rr_touch to meet server requirements.
         */
        if (ha_rr_touch(&ErrBlock) == HA_RR_FAIL) {

```

## rmapi\_smplib.c

```
        display_rmapi_err(&ErrBlock);
        mon_exit(1);
    }

    /*
     * Reset the TimeOut counter.
     */
    TimeOut = MAX_INTERVAL;
}

/*
 * Catches SIGIO and sets global flag SigioCaught to send values.
 */
void catch_io_signal(int sig) {
    SigioCaught = 1;
}

/*
 * Catches SIGALRM and sets global flag SampleTime to send values.
 */
void catch_alrm_signal(int sig) {
    SampleTime = 1;
}

/*
 * Catches SIGTERM and sets global flag Terminate to exit.
 */
void catch_exit_signal(int sig) {
    Terminate = 1;
    SigCaught = sig;
}

/*
 * Handles a connection request on the RMAPI server
 * socket returned by ha_rr_makserv().
 */
void server_socket_handler(int table_index) {
    int i, mgr_sock, rc;

    mgr_sock = ha_rr_start_session(HA_RR_NOTIFY_SIGIO, &ErrBlock);
    if (mgr_sock < 0) {
        display_rmapi_err(&ErrBlock);
        if ((ErrBlock.em_errno != HA_RR_EMAXSESSIONS) && (ErrBlock.em_errno != HA_RR_EAGAIN)) {
            /*
             * Severe error in the RMAPI - exit.
             */
            mon_exit(1);
        }
        return;
    }
    /*
     * Find an available element in the session socket table.
     */
    for (i=0; (i < HA_RR_MAX_SESSIONS) && (SocketTable[i].socket_fd >= 0); i++);
    if (i < HA_RR_MAX_SESSIONS) {
        /*
         * Save the manager socket file descriptor returned by the RMAPI.
         * The socket function was initialized in main.
         */
        SocketTable[i].socket_fd = mgr_sock;
        NumMgrs++;
        fprintf(stdout, "New session accepted to index(%d) socket_fd(%d).\n", i, mgr_sock);
    } else {
        /*
         * Should never get here, RMAPI handles too many sessions.
         */
        rc = ha_rr_end_session(mgr_sock, &ErrBlock);
    }
}

/*
 * Handles a message being received from the manager session
 * specified by [table_index].
 */
void session_socket_handler(int table_index) {
    int len, session_sock;
    struct ha_rr_ctrl_msg *ctrl_msg_start, *ctrl_msg;

    /*
     * Get the socket file descriptor for this session.
     */
    if ((session_sock = SocketTable[table_index].socket_fd) < 0) {
        return;
    }

    /*
     * Call ha_rr_get_ctrl_msg() to read the command from the manager.
     */
    fprintf(stdout, "Calling ha_rr_get_ctrlmsg for session(%d) socket_fd(%d).\n",

```

```

        table_index,session_sock);
len = ha_rr_get_ctrlmsg(session_sock, &ctrl_msg, &ErrBlock);
if (len == 0) {
    /*
     * Message was for the RMAPI or was incomplete.
     */
    return;
}

if (len == HA_RR_FAIL) {
    display_rmap_i_err(&ErrBlock);
    if (ErrBlock.em_errno == HA_RR_EDISCONNECT) {
        end_session(table_index);
    } else if (ErrBlock.em_errno != HA_RR_EAGAIN) {
        /*
         * Severe RMAPI error.
         */
        mon_exit(1);
    }
    return;
}

/*
 * Since SIGIO is the notification protocol there could be
 * more than 1 message to read.
 */
ctrl_msg_start = ctrl_msg;
while ((char *)ctrl_msg < (((char *)ctrl_msg_start) + len)) {
    switch (ctrl_msg->rr_ctrl_cmd) {
        case HA_RR_CMD_ADDALL :
        case HA_RR_CMD_ADDV  :
            add_variables(table_index,ctrl_msg);
            break;
        case HA_RR_CMD_DELALL :
        case HA_RR_CMD_DELV   :
            del_variables(table_index,ctrl_msg);
            break;
        case HA_RR_CMD_INSTV  :
            inst_variables(ctrl_msg);
            break;
        default                :
            /*
             * Unknown or unsupported msg - ignore it.
             */
            fprintf(stderr,"Received an unexpected command(%d) from socket_fd(%d), ignoring...\n",
                ctrl_msg->rr_ctrl_cmd,session_sock);
            break;
    }
    ctrl_msg = (struct ha_rr_ctrl_msg *)(((char *)ctrl_msg) + ctrl_msg->rr_ctrl_msg_len);
}

/*
 * Free the message allocated by the RMAPI.
 */
free(ctrl_msg_start);
}

/*
 * Dynamically instantiates the variables requested in the control message. If the message
 * does not wildcard the resource ID, this routine creates instances of the requested
 * variables, otherwise it creates several variables to simulate wildcard matching.
 */
void inst_variables(struct ha_rr_ctrl_msg *ctrl_msg) {
    int num_var, num_to_create;
    struct ha_rr_ctrl_var *vp;
    char *inst_name;
    static int inst_num = 10;
    int i, j;

    fprintf(stdout,"Processing command HA_RR_CMD_INSTV to create %d variables.\n",
        ctrl_msg->rr_ctrl_num_vars);

    for (i=0;i < ctrl_msg->rr_ctrl_num_vars;i++) {
        num_to_create = 0;

        /*
         * Get a pointer to the variable structure in the control message.
         */
        vp = &(ctrl_msg->rr_ctrl_vars[i]);

        if (strcmp(DYN_INST_VAR_NAME, vp->rr_ctrl_name)) {
            /*
             * Variable name didn't match the dynamically instantiable variable name.
             */
            fprintf(stderr,"%s: Invalid variable (%s) name sent in INSTV ctrl msg.\n",
                PROGRAMNAME,vp->rr_ctrl_name);
            continue;
        }

        if (strcmp(vp->rr_ctrl_rsrc_ID,"") != 0) {

```

## rmapi\_smpsig.c

```

/*
 * A specific instance is requested. Accommodate the request by creating
 * the instance. If this were a real monitor, it would be created if
 * the request matched a valid variable.
 */
if (strcmp(vp->rr_ctrl_rsrc_ID,DYN_INST_VAR_RESID,strlen(DYN_INST_VAR_RESID))) {
    /*
     * The resource ID in the control message does not match the resource ID
     * for the variable.
     */
    fprintf(stderr,"%s: Invalid resource ID (%s) sent in INSTV ctrl msg.\n",
            PROGRAMNAME,vp->rr_ctrl_rsrc_ID);
    continue;
}
/*
 * Get a pointer passed the resource ID name - should be pointing to an '=' now.
 * The format of the resource ID is: ResourceID_Element_Name=Value
 */
inst_name = vp->rr_ctrl_rsrc_ID + strlen(DYN_INST_VAR_RESID);
if (*inst_name != '=' || *(inst_name + 1) == '\0') {
    /*
     * '=' not found or value missing in the ctrl message.
     */
    fprintf(stderr,"%s: Invalid resource ID (%s) sent in INSTV ctrl msg.\n",
            PROGRAMNAME,vp->rr_ctrl_rsrc_ID);
    continue;
}
/*
 * Adjust the instance name pointer to the value in the control message.
 */
inst_name++;
num_to_create = 1;
} else {
    /*
     * Resource ID in the message was NULL which means wildcard the value.
     * Since this is just a sample monitor, we'll make up some variable names
     * to simulate the wildcard match. The contrived instances will have
     * resource IDs: InstName=Inst<nn> where nn is the value of the local static
     * variable inst_num.
     */
    num_to_create = 5;
}

fprintf(stdout,"Instantiating %d new variables.\n",num_to_create);
/*
 * realloc() the local and RMAPI variable arrays to accommodate the new instances.
 */
num_var = NumVariables + num_to_create;
LocalVars = (struct local_vars *)realloc(LocalVars, sizeof(struct local_vars) * num_var);
Variables = (struct ha_rr_variable *)realloc(Variables, sizeof(struct ha_rr_variable) * num_var);
Values = (struct ha_rr_val *)realloc(Values, sizeof(struct ha_rr_val) * num_var);

/*
 * Initialize the new local variables.
 */
memset(LocalVars+NumVariables,0,sizeof(struct local_vars) * num_to_create);
for (j=NumVariables;j<num_var;j++) {
    /*
     * Copy the dynamic variable name to the local struct.
     */
    sprintf(LocalVars[j].var_name,DYN_INST_VAR_NAME);
    if (num_to_create == 1) {
        /*
         * Resource ID was specified in the control message - use the
         * value in the new variable's resource ID.
         */
        sprintf(LocalVars[j].var_resid,"%s=%s",DYN_INST_VAR_RESID,inst_name);
        fprintf(stdout,"Created dynamic instance: %s\n",LocalVars[j].var_resid);
    } else {
        /*
         * Resource ID not specified in the control message - contrive
         * a resource ID for the new instance.
         */
        sprintf(LocalVars[j].var_resid,"%s=Inst%d",DYN_INST_VAR_RESID,inst_num);
        /*
         * Increment the inst_num counter to keep new resource IDs unique.
         */
        inst_num++;
    }
    LocalVars[j].value = INITIAL_VALUE;
}
/*
 * Update the variable array size.
 */
NumVariables = num_var;
}
/*
 * Register the new variables with the RMAPI.
 */
register_variables();
}

```

```

/*
 * Adds the variables referenced in the ctrl_msg to the
 * manager session specified by [table index].
 */
void add_variables(int table_index, struct ha_rr_ctrl_msg *ctrl_msg) {
    int i;
    int sock_fd;
    int inst_id;
    int num_added, num_to_add = 0;

    sock_fd = SocketTable[table_index].socket_fd;
    switch (ctrl_msg->rr_ctrl_cmd) {
        case HA_RR_CMD_ADDALL:
            /*
             * Add all variables to this manager session.
             */
            fprintf(stdout,"Processing cmd HA_RR_CMD_ADDALL to add all variables to session_fd %d.\n",sock_fd);
            for (i=0;i<NumVariables;i++) {
                /*
                 * Copy the variable to the next RMAPI structure.
                 */
                Variables[num_to_add].rr_var_name = LocalVars[i].var_name;
                Variables[num_to_add].rr_var_rsrc_ID = LocalVars[i].var_resid;
                Variables[num_to_add].rr_varu.rr_var_hdl = &(LocalVars[i].var_handle);
                Variables[num_to_add].rr_value = &(LocalVars[i].value);

                /*
                 * Increment the count of variables to be added.
                 */
                num_to_add++;
            }
            break;
        case HA_RR_CMD_ADDV :
            /*
             * Add a vector of variables to the session.
             */
            fprintf(stdout,"Processing cmd HA_RR_CMD_ADDV to add %d vars for session_fd %d.\n",
                ctrl_msg->rr_ctrl_num_vars,sock_fd);
            for (i=0;i < ctrl_msg->rr_ctrl_num_vars;i++) {
                /*
                 * The instance id field in the ccontrol message is
                 * the index into the LocalVars array.
                 */
                inst_id = ctrl_msg->rr_ctrlv.rr_ctrl_vari[i];

                if (inst_id < NumVariables) {
                    /*
                     * Copy the variable to the next RMAPI structure.
                     */
                    Variables[num_to_add].rr_var_name = LocalVars[inst_id].var_name;
                    Variables[num_to_add].rr_var_rsrc_ID = LocalVars[inst_id].var_resid;
                    Variables[num_to_add].rr_varu.rr_var_hdl = &(LocalVars[inst_id].var_handle);
                    Variables[num_to_add].rr_value = &(LocalVars[inst_id].value);

                    num_to_add++;
                }
            }
            break;
        default :
            break;
    }
    if (num_to_add) {
        /*
         * Add the variables by calling ha_rr_add_var.
         */
        num_added = ha_rr_add_var(sock_fd, Variables, num_to_add, 1, &ErrBlock);

        if (num_added == HA_RR_FAIL) {
            display_rmap_err(&ErrBlock);

            if (ErrBlock.em_errno == HA_RR_EDISCONNECT) {
                /*
                 * Session closed by manager.
                 */
                end_session(table_index);
            } else {
                mon_exit(1);
            }
        }
    } else {
        fprintf(stdout,"ha_rr_add_var() added %d vars to session_fd %d.\n",num_added,sock_fd);
        if (num_to_add != num_added) {
            /*
             * Loop through the variables that were attempted
             * to be added and report any that had errors.
             */
            for (i=0;i<num_to_add;i++) {
                if (Variables[i].rr_var_errno != 0) {
                    fprintf(stderr,"Variable(%s) resource ID(%s) had bad errno(%d) from ha_rr_add_var().\n",

```

## rmapi\_smplib.c

```

        Variables[i].rr_var_name, Variables[i].rr_var_rsrc_ID, Variables[i].rr_var_errno);
    }
}

/*
 * Deletes the variables referenced in the control message from
 * the manager session specified by [table_index].
 */
void del_variables(int table_index, struct ha_rr_ctrl_msg *ctrl_msg) {
    int i;
    int sock_fd;
    int ctrl_cmd;
    int inst_id;
    int num_deleted, num_to_del = 0;

    sock_fd = SocketTable[table_index].socket_fd;
    if (ctrl_msg == (struct ha_rr_ctrl_msg *)0) {
        /*
         * Called with NULL message when a session is ending. Need to delete
         * all variables from the session before calling ha_rr_end_session.
         */
        ctrl_cmd = HA_RR_CMD_DELALL;
    } else {
        ctrl_cmd = ctrl_msg->rr_ctrl_cmd;
    }

    switch (ctrl_cmd) {
        case HA_RR_CMD_DELALL :
            fprintf(stdout, "Processing command HA_RR_CMD_DELALL to delete all variables for session_fd %d.\n",
                sock_fd);
            /*
             * Delete all variables from this manager session.
             */
            for (i=0; i<NumVariables; i++) {
                if (LocalVars[i].var_handle != (void *)0) {
                    /*
                     * Copy the variable handle to the next RMAPI structure.
                     */
                    Variables[num_to_del].rr_varu.rr_var_hndl = &(LocalVars[i].var_handle);

                    /*
                     * Increment the count of variables to be deleted.
                     */
                    num_to_del++;
                }
            }
            break;
        case HA_RR_CMD_DELV :
            fprintf(stdout, "Processing command HA_RR_CMD_DELV to delete %d variables for session_fd %d.\n",
                ctrl_msg->rr_ctrl_num_vars, sock_fd);
            /*
             * Delete a vector of variables from this manager session.
             */
            for (i = 0; i < ctrl_msg->rr_ctrl_num_vars; i++) {
                inst_id = ctrl_msg->rr_ctrlv.rr_ctrl_vari[i];
                if ((inst_id < NumVariables) &&
                    (LocalVars[i].var_handle != (void *)0)) {
                    /*
                     * Copy the variable handle to the next RMAPI structure.
                     */
                    Variables[num_to_del].rr_varu.rr_var_hndl = &(LocalVars[inst_id].var_handle);

                    /*
                     * Increment the count of variables to be deleted.
                     */
                    num_to_del++;
                }
            }
            break;
        default :
            break;
    }

    if (num_to_del) {
        /*
         * Delete the variables by calling ha_rr_del_var.
         */
        num_deleted = ha_rr_del_var(sock_fd, Variables, num_to_del, &ErrBlock);

        if (num_deleted == HA_RR_FAIL) {
            display_rmapi_err(&ErrBlock);

            if (ErrBlock.em_errno == HA_RR_EDISCONNECT) {
                /*
                 * Session closed by manager.
                 */
                end_session(table_index);
            } else {

```

```

        mon_exit(1);
    }
}

/*
 * rc>0 from ha_rr_del_var is the number of variables that no longer need
 * their values updated. The RMAPI will have set their handles to NULL.
 */
fprintf(stdout,"%d variables no longer need to be updated.\n",num_deleted);
}

/*
 * Ends the manager session indexed by the [table_index] parameter
 * in the socket table.
 */
void end_session(int table_index) {
int rc;

    fprintf(stdout,"Ending session index(%d) session_fd(%d).\n",
        table_index,SocketTable[table_index].socket_fd);

    /*
     * Make sure this is a valid session.
     */
    if ((table_index < 0) || (table_index > HA_RR_MAX_SESSIONS) ||
        (SocketTable[table_index].socket_fd < 0)){
        return;
    }

    /*
     * Delete all variables for this manager.
     */
    del_variables(table_index, (struct ha_rr_ctrl_msg *)0);

    /*
     * Call RMAPI to end the session.
     */
    rc = ha_rr_end_session(SocketTable[table_index].socket_fd, &ErrBlock);
    if (rc == HA_RR_FAIL) {
        display_rmap_i_err(&ErrBlock);
        mon_exit(1);
    }

    /*
     * Reset the socket file descriptor in the SocketTable and
     * decrement the manager count.
     */
    SocketTable[table_index].socket_fd = -1;
    NumMgrs--;

    if (NumMgrs <= 0) {
        /*
         * This was the last manager to end - exit with a good rc.
         */
        fprintf(stdout,"Last manager session closed, exiting...\n");
        mon_exit(0);
    }
}

/*
 * Registers variables with the RMAPI.
 */
void register_variables() {
int i, j, num_var;
int num_registered, num_to_reg = 0;

    for (i=0;i<NumVariables;i++) {
        if (LocalVars[i].registered == 0) {

            /*
             * Variable not registered yet. Copy name, resource ID and value
             * pointers. Use the index of the variable as the inst id.
             */
            LocalVars[i].registered = 1;
            Variables[num_to_reg].rr_var_name = LocalVars[i].var_name;
            Variables[num_to_reg].rr_var_rsrc_ID = LocalVars[i].var_resid;
            Variables[num_to_reg].rr_var_iid = i;
            num_to_reg++;
        }
    }

    if (num_to_reg) {

        /*
         * Call the RMAPI to register the variables.
         */
        num_registered = ha_rr_reg_var(Variables,num_to_reg,&ErrBlock);
        if (num_registered != num_to_reg) {
            if (num_registered == HA_RR_FAIL) {
                display_rmap_i_err(&ErrBlock);
            }
        }
    }
}

```

## rmapi\_smplib.c

```
    mon_exit(1);
}
printf(stdout,"%d variables registered with the RMAPI.\n",
       num_registered);
for (i=0; i < num_to_reg; i++) {
    if ((Variables[i].rr_var_errno) || (num_registered == HA_RR_FAIL)) {
        /*
         * Display a message for each variable with a registration error.
         */
        fprintf(stderr,"ha_rr_reg_var() error with variable(%) instid(%) RMAPI errno=(%).\n",
                Variables[i].rr_var_name,
                Variables[i].rr_var_inst_id,
                Variables[i].rr_var_errno);
    }
}
}
}

/*
 * Called once by main to establish the control loop.
 */
void control_loop() {
    int i, rc;
    int sock;

    for(;;) {

        /*
         * Unblock the signals.
         */
        if (sigprocmask(SIG_UNBLOCK, &SignalMask, (sigset_t *)0) < 0) {
            /*
             * sigprocmask sys call failed.
             */
            fprintf(stderr,"Cannot unblock signals. sigprocmask() errno=%d.\n",errno);
            mon_exit(1);
        }

        /*
         * Pause to receive a signal.
         */
        pause();

        /*
         * Block the signals.
         */
        if (sigprocmask(SIG_BLOCK, &SignalMask, (sigset_t *)0) < 0) {
            fprintf(stderr,"Cannot block signals. sigprocmask() errno=%d.\n",errno);
            mon_exit(1);
        }

        if (Terminate) {
            /*
             * SIGTERM caught - call mon_exit() to clean up.
             */
            fprintf(stdout,"Caught signal(%)...calling exit...\n",SigCaught);
            mon_exit(SigCaught);
        }

        if (SampleTime) {
            /*
             * SIGALRM caught - call send_values()
             * and reset the sample flag.
             */
            send_values();
            SampleTime = 0;
        }

        if (SigioCaught) {
            /*
             * Call the socket function for each valid socket.
             */
            SigioCaught = 0;
            for (i=0;i<SOCKET_TABLE_SIZE;i++) {
                if (SocketTable[i].socket_fd >=0) {
                    fprintf(stdout,"Calling socket function for session(%) socket_fd(%).\n",
                            i,SocketTable[i].socket_fd);
                    (SocketTable[i].sock_funcp)(i);
                }
            }
        }
    }
}

/*
 * Routine to terminate the monitor.
 */
void mon_exit(int s)
{
    static int recursively_called = 0;
    int i, rc;
```

```

/*
 * Check for recursive call - some routines mon_exit()
 * calls can likewise call mon_exit().
 */
if (recursively_called) return;
recursively_called = 1;

if (RmapiInit) {
    /*
     * Gracefully close all sessions.
     */
    for (i=0; i<HA_RR_MAX_SESSIONS; i++) {
        if (SocketTable[i].socket_fd >= 0) {
            end_session(i);
        }
    }
    /*
     * Terminate the RMAPI.
     */
    rc = ha_rr_terminate(&ErrBlock);
    if (rc) {
        display_rmap_err(&ErrBlock);
    }
}

fprintf(stdout, "%s: %s resource monitor exiting.\n", PROGNAME, RESOURCE_MONITOR_NAME);
exit(s);
}

/*
 * Display an RMAPI error.
 */
void display_rmap_err(struct ha_em_err_blk *errblk)
{
    fprintf(stderr,
        "RMAPI Error: File(%s) Version(%s) Line(%d) Errno(%d)\n\t%s",
        errblk->em_errfile,
        errblk->em_errlevel,
        errblk->em_errline,
        errblk->em_errno,
        errblk->em_errmsg);
}

```

---

## The rmapi\_smp.msg Message File

```

$ IBM_PROLOG_BEGIN_TAG
$ This is an automatically generated prolog.
$
$
$ Licensed Materials - Property of IBM
$
$ (C) COPYRIGHT International Business Machines Corp. 1996,1998
$ All Rights Reserved
$
$ US Government Users Restricted Rights - Use, duplication or
$ disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
$
$ IBM_PROLOG_END_TAG
$ =====*/
$ Module Name:  rmapi_smp.msg                               */
$ Description:  */
$   Script to load Event Manager configuration data into the SDR.  */
$   This data is for the Rmapi sample monitors.                  */
$   Use runcat to generate a catalog from this file:            */
$   runcat rmapi_smp rmapi_smp.msg                             */
$   Then copy the rmapi_smp.cat file to the nls directory       */
$   for the language set by your environment. */
$ =====*/
$ "@(#)45 1.6 src/rsct/pem/emtools/rmapi_samples/rmapi_smp.msg, emtools, rsct_rtro 2/20/98 14:23:17" */
$ The quote signifies that the " character will be used to delimit
$ messages in the catalog.
$
$quote "
$
$ Lines beginning with a dollar sign are comments. Lines
$ beginning with numbers are the messages.
$
$set 1
$
1 "IBM Data Suppliers"
$

```

## rmapi\_smp.msg

```
2 "RMAPI Sample Monitor"
$
3 "Example non-instantiable resource variable (Counter or Quantity)."
```

---

```
$
4 "Example dynamically instantiable resource variable."
$
5 "Example instantiable resource variable (Counter or Quantity)."
```

---

```
$
6 "Example instantiable resource variable (State)."
```

---

```
$
7 "Example non-instantiable variables."
$
8 "Example instantiable variables."
$
9 "Example resource identifier NodeNum. Used to instantiate variables as node \n\
specific and as the variable locator field."
$
10 "Example resource identifier for creating multiple instances of variables."
$
11 "Example SBS field: Action. Last action requested by the sample command."
$
12 "Example SBS field: Options. Last options parameter passed to the sample command."
$
13 "Example SBS field: StateChange. Success of last sample command call (!0==success,0==fail)."
```

---

```
$
14 "Example SBS field: State. State of the resource."
$
```

## The rmapi\_smp.loadsdr Shell Script

```
#!/bin/ksh
# IBM_PROLOG_BEGIN_TAG
# This is an automatically generated prolog.
#
#
# Licensed Materials - Property of IBM
#
# (C) COPYRIGHT International Business Machines Corp. 1996,1998
# All Rights Reserved
#
# US Government Users Restricted Rights - Use, duplication or
# disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
#
# IBM_PROLOG_END_TAG
#*=====*/
#*                                     */
#* Module Name:  rmapi_smp.loadsdr      */
#*                                     */
#* Description:                                     */
#*   Script to load Event Manager configuration data into the SDR. */
#*   This data is for the Rmapi sample monitors.                  */
#*                                     */
#* Usage: rmapi_smp.loadsdr <resource_monitor_path>              */
#*   Where <resource_monitor_path> is the path to the            */
#*   compiled monitor samples.                                     */
#*                                     */
#*=====*/
#@(#)40 1.8 src/rsct/pem/emtools/rmapi_samples/rmapi_smp.loadsdr, emtools, rsct_rtro 6/5/98 11:25:22"
#*=====*/

#-----
# Load resource monitor data for example resource monitor rmapi_smpdae
#-----

if [ "$#" -ne 1 ]
then
echo "Usage: $0 <resource_monitor_path>"
exit 1
fi
rmpath=$1

SDRCreateObjects EM_Resource_Monitor \
'rmName=IBM.PSSP.SampleDaeMon' \
'rmMessage_file=rmapi_smp.cat' \
'rmMessage_set=1' \
'rmConnect_type=server' \
'rmNum_instances=8' \
'rmPTX_prefix=IBM/PSSP.SampleDaeMon' \
'rmPTX_description=1,2' \
'rmPTX_asno=2'
if (( $? != 0 )) ; then exit 1 ; fi

SDRCreateObjects EM_Resource_Class \
'rcClass=IBM.PSSP.SampleDaeClass' \
'rcResource_monitor=IBM.PSSP.SampleDaeMon' \
'rcObservation_interval=10' \
'rcReporting_interval=10'
if (( $? != 0 )) ; then exit 1 ; fi
```

```

SDRCreatObjects EM_Resource_Variable \
'rvName=IBM.PSSP.SampleDaeMon.StaticVars.static_var1' \
'rvLocator=NodeNum' \
'rvDescription=3' \
'rvValue_type=Quantity' \
'rvInitial_value=0' \
'rvData_type=long' \
'rvPTX_name=StaticVars/static_var1' \
'rvPTX_description=7' \
'rvPTX_min=0' \
'rvPTX_max=500' \
'rvClass=IBM.PSSP.SampleDaeClass' \
'rvDynamic_instance=0'
if (( $? != 0 )) ; then exit 1 ; fi

```

```

SDRCreatObjects EM_Resource_Variable \
'rvName=IBM.PSSP.SampleDaeMon.StaticVars.static_var2' \
'rvLocator=NodeNum' \
'rvDescription=3' \
'rvValue_type=Quantity' \
'rvInitial_value=0' \
'rvData_type=long' \
'rvPTX_name=StaticVars/static_var2' \
'rvPTX_description=7' \
'rvPTX_min=0' \
'rvPTX_max=500' \
'rvClass=IBM.PSSP.SampleDaeClass' \
'rvDynamic_instance=0'
if (( $? != 0 )) ; then exit 1 ; fi

```

```

SDRCreatObjects EM_Resource_Variable \
'rvName=IBM.PSSP.SampleDaeMon.StaticVars.static_var3' \
'rvLocator=NodeNum' \
'rvDescription=3' \
'rvValue_type=Quantity' \
'rvInitial_value=0' \
'rvData_type=long' \
'rvPTX_name=StaticVars/static_var3' \
'rvPTX_description=7' \
'rvPTX_min=0' \
'rvPTX_max=500' \
'rvClass=IBM.PSSP.SampleDaeClass' \
'rvDynamic_instance=0'
if (( $? != 0 )) ; then exit 1 ; fi

```

```

SDRCreatObjects EM_Resource_Variable \
'rvName=IBM.PSSP.SampleDaeMon.InstVars.inst_var1' \
'rvLocator=NodeNum' \
'rvDescription=5' \
'rvValue_type=Quantity' \
'rvInitial_value=0' \
'rvData_type=long' \
'rvPTX_name=$InstName/inst_var1' \
'rvPTX_description=8' \
'rvPTX_min=0' \
'rvPTX_max=500' \
'rvClass=IBM.PSSP.SampleDaeClass' \
'rvDynamic_instance=0'
if (( $? != 0 )) ; then exit 1 ; fi

```

```

SDRCreatObjects EM_Resource_Variable \
'rvName=IBM.PSSP.SampleDaeMon.InstVars.inst_var2' \
'rvLocator=NodeNum' \
'rvDescription=5' \
'rvValue_type=Quantity' \
'rvInitial_value=0' \
'rvData_type=long' \
'rvPTX_name=$InstName/inst_var2' \
'rvPTX_description=8' \
'rvPTX_min=0' \
'rvPTX_max=500' \
'rvClass=IBM.PSSP.SampleDaeClass' \
'rvDynamic_instance=0'
if (( $? != 0 )) ; then exit 1 ; fi

```

```

SDRCreatObjects EM_Resource_ID \
'riResource_name=IBM.PSSP.SampleDaeMon.StaticVars' \
'riElement_name=NodeNum' \
'riElement_description=9'
if (( $? != 0 )) ; then exit 1 ; fi

```

```

SDRCreatObjects EM_Resource_ID \
'riResource_name=IBM.PSSP.SampleDaeMon.InstVars' \
'riElement_name=NodeNum' \
'riElement_description=9'
if (( $? != 0 )) ; then exit 1 ; fi

```

```

SDRCreatObjects EM_Resource_ID \
'riResource_name=IBM.PSSP.SampleDaeMon.InstVars' \
'riElement_name=InstName' \
'riElement_description=10'
if (( $? != 0 )) ; then exit 1 ; fi

```

## rmapi\_smp.loadsdr

```
#####  
# Load resource monitor data for example resource monitor rmapi_smpcmd  
#####  
  
SDRCreateObjects EM_Resource_Monitor \  
'rmName=IBM.PSSP.SampleCmdMon' \  
'rmMessage_file=rmapi_smp.cat' \  
'rmMessage_set=1' \  
'rmConnect_type=client' \  
'rmPTX_prefix=dummy' \  
'rmPTX_description=2' \  
'rmPTX_asno=1'  
if (( $? != 0 )) ; then exit 1 ; fi  
  
SDRCreateObjects EM_Resource_Class \  
'rcClass=IBM.PSSP.SampleCmdClass' \  
'rcResource_monitor=IBM.PSSP.SampleCmdMon' \  
'rcObservation_interval=0' \  
'rcReporting_interval=0'  
if (( $? != 0 )) ; then exit 1 ; fi  
  
SDRCreateObjects EM_Resource_Variable \  
'rvName=IBM.PSSP.SampleCmdMon.state' \  
'rvLocator=NodeNum' \  
'rvDescription=6' \  
'rvValue_type=State' \  
'rvInitial_value=0' \  
'rvData_type=long' \  
'rvClass=IBM.PSSP.SampleCmdClass' \  
'rvDynamic_instance=0'  
if (( $? != 0 )) ; then exit 1 ; fi  
  
SDRCreateObjects EM_Resource_Variable \  
'rvName=IBM.PSSP.SampleCmdMon.call' \  
'rvLocator=NodeNum' \  
'rvDescription=6' \  
'rvValue_type=State' \  
'rvInitial_value=0' \  
'rvData_type=SBS' \  
'rvClass=IBM.PSSP.SampleCmdClass' \  
'rvDynamic_instance=0'  
if (( $? != 0 )) ; then exit 1 ; fi  
  
SDRCreateObjects EM_Resource_ID \  
'riResource_name=IBM.PSSP.SampleCmdMon' \  
'riElement_name=NodeNum' \  
'riElement_description=9'  
if (( $? != 0 )) ; then exit 1 ; fi  
  
SDRCreateObjects EM_Resource_ID \  
'riResource_name=IBM.PSSP.SampleCmdMon' \  
'riElement_name=NAME' \  
'riElement_description=10'  
if (( $? != 0 )) ; then exit 1 ; fi  
  
SDRCreateObjects EM_Structured_Byte_String \  
'sbsVariable_name=IBM.PSSP.SampleCmdMon.call' \  
'sbsField_name=11' \  
'sbsField_type=long' \  
'sbsField_SN=0'  
if (( $? != 0 )) ; then exit 1 ; fi  
  
SDRCreateObjects EM_Structured_Byte_String \  
'sbsVariable_name=IBM.PSSP.SampleCmdMon.call' \  
'sbsField_name=12' \  
'sbsField_type=cstring' \  
'sbsField_SN=1'  
if (( $? != 0 )) ; then exit 1 ; fi  
  
SDRCreateObjects EM_Structured_Byte_String \  
'sbsVariable_name=IBM.PSSP.SampleCmdMon.call' \  
'sbsField_name=13' \  
'sbsField_type=long' \  
'sbsField_SN=2'  
if (( $? != 0 )) ; then exit 1 ; fi  
  
SDRCreateObjects EM_Structured_Byte_String \  
'sbsVariable_name=IBM.PSSP.SampleCmdMon.call' \  
'sbsField_name=14' \  
'sbsField_type=long' \  
'sbsField_SN=3'  
if (( $? != 0 )) ; then exit 1 ; fi  
  
#####  
# Load resource monitor data for example resource monitor rmapi_smpsig  
#####  
  
SDRCreateObjects EM_Resource_Monitor \  
'rmName=IBM.PSSP.SampleSigMon' \  
'rmPath=$rmpath/rmapi_smpsig' \  
'rmMessage_file=rmapi_smp.cat' \  

```

```

'rmMessage_set=1' \
'rmConnect_type=server' \
'rmPTX_prefix=IBM/PSSP.SampleSigMon' \
'rmPTX_description=1,2' \
'rmPTX_asno=2'
if (( $? != 0 )) ; then exit 1 ; fi

SDRCreatObjects EM_Resource_Class \
'rcClass=IBM.PSSP.SampleSigDynInstClass' \
'rcResource_monitor=IBM.PSSP.SampleSigMon' \
'rcObservation_interval=10' \
'rcReporting_interval=10'
if (( $? != 0 )) ; then exit 1 ; fi

SDRCreatObjects EM_Resource_Class \
'rcClass=IBM.PSSP.SampleSigNonInstClass' \
'rcResource_monitor=IBM.PSSP.SampleSigMon' \
'rcObservation_interval=10' \
'rcReporting_interval=10'
if (( $? != 0 )) ; then exit 1 ; fi

SDRCreatObjects EM_Resource_Variable \
'rvName=IBM.PSSP.SampleSigMon.NonInstVars.var1' \
'rvLocator=NodeNum' \
'rvDescription=3' \
'rvValue_type=Counter' \
'rvInitial_value=0' \
'rvData_type=long' \
'rvPTX_name=NonInstVars/var1' \
'rvPTX_description=7' \
'rvPTX_min=0' \
'rvPTX_max=500000' \
'rvClass=IBM.PSSP.SampleSigNonInstClass' \
'rvDynamic_instance=0'
if (( $? != 0 )) ; then exit 1 ; fi

SDRCreatObjects EM_Resource_Variable \
'rvName=IBM.PSSP.SampleSigMon.NonInstVars.var2' \
'rvLocator=NodeNum' \
'rvDescription=3' \
'rvValue_type=Counter' \
'rvInitial_value=0' \
'rvData_type=long' \
'rvPTX_name=NonInstVars/var2' \
'rvPTX_description=7' \
'rvPTX_min=0' \
'rvPTX_max=500000' \
'rvClass=IBM.PSSP.SampleSigNonInstClass' \
'rvDynamic_instance=0'
if (( $? != 0 )) ; then exit 1 ; fi

SDRCreatObjects EM_Resource_Variable \
'rvName=IBM.PSSP.SampleSigMon.NonInstVars.var3' \
'rvLocator=NodeNum' \
'rvDescription=3' \
'rvValue_type=Counter' \
'rvInitial_value=0' \
'rvData_type=long' \
'rvPTX_name=NonInstVars/var3' \
'rvPTX_description=7' \
'rvPTX_min=0' \
'rvPTX_max=500000' \
'rvClass=IBM.PSSP.SampleSigNonInstClass' \
'rvDynamic_instance=0'
if (( $? != 0 )) ; then exit 1 ; fi

SDRCreatObjects EM_Resource_ID \
'riResource_name=IBM.PSSP.SampleSigMon.NonInstVars' \
'riElement_name=NodeNum' \
'riElement_description=9'
if (( $? != 0 )) ; then exit 1 ; fi

SDRCreatObjects EM_Resource_Variable \
'rvName=IBM.PSSP.SampleSigMon.DynInstVar.var' \
'rvLocator=NodeNum' \
'rvDescription=4' \
'rvValue_type=Quantity' \
'rvInitial_value=0' \
'rvData_type=long' \
'rvPTX_name=DynInstVars/$InstName/var' \
'rvPTX_description=8,8' \
'rvPTX_min=0' \
'rvPTX_max=500000' \
'rvClass=IBM.PSSP.SampleSigDynInstClass' \
'rvDynamic_instance=1'
if (( $? != 0 )) ; then exit 1 ; fi

SDRCreatObjects EM_Resource_ID \
'riResource_name=IBM.PSSP.SampleSigMon.DynInstVar' \
'riElement_name=NodeNum' \
'riElement_description=9'
if (( $? != 0 )) ; then exit 1 ; fi

```

## rmapi\_smp.loadsdr

```
SDRCreatObjects EM_Resource_ID \
'riResource_name=IBM.PSSP.SampleSigMon.DynInstVar' \
'riElement_name=InstName' \
'riElement_description=I0'
if (( $? != 0 )) ; then exit 1 ; fi
```

---

## The rmapi\_smp.unloadsdr Shell Script

```
#!/bin/ksh
# IBM_PROLOG_BEGIN_TAG
# This is an automatically generated prolog.
#
#
# Licensed Materials - Property of IBM
#
# (C) COPYRIGHT International Business Machines Corp. 1996,1998
# All Rights Reserved
#
# US Government Users Restricted Rights - Use, duplication or
# disclosure restricted by GSA ADP Schedule Contract with IBM Corp.
#
# IBM_PROLOG_END_TAG
#-----*/
##                               */
## Module Name:  rmapi_smp.unloadsdr      */
##                               */
## Description:                               */
##   Script to remove Event Manager configuration data from the SDR.      */
##   This script is for sample resource monitors.          */
##                               */
#-----*/
# "@(#)41 1.7 src/rsct/pem/emtools/rmapi_samples/rmapi_smp.unloadsdr, emtools, rsct_rtro 2/20/98 14:23:22"
#-----*/
# Unload resource monitor data for example resource monitor rmapi_smpdae
#-----*/

SDRDeleteObjects EM_Resource_Monitor \
'rmName==IBM.PSSP.SampleDaeMon'

SDRDeleteObjects EM_Resource_Class \
'rcClass==IBM.PSSP.SampleDaeClass'

SDRDeleteObjects EM_Resource_Variable \
'rvName==IBM.PSSP.SampleDaeMon.StaticVars.static_var1'

SDRDeleteObjects EM_Resource_Variable \
'rvName==IBM.PSSP.SampleDaeMon.StaticVars.static_var2'

SDRDeleteObjects EM_Resource_Variable \
'rvName==IBM.PSSP.SampleDaeMon.StaticVars.static_var3'

SDRDeleteObjects EM_Resource_Variable \
'rvName==IBM.PSSP.SampleDaeMon.InstVars.inst_var1'

SDRDeleteObjects EM_Resource_Variable \
'rvName==IBM.PSSP.SampleDaeMon.InstVars.inst_var2'

SDRDeleteObjects EM_Resource_ID \
'riResource_name==IBM.PSSP.SampleDaeMon.StaticVars'

SDRDeleteObjects EM_Resource_ID \
'riResource_name==IBM.PSSP.SampleDaeMon.InstVars'

#-----*/
# Unload resource monitor data for example resource monitor rmapi_smpcmd
#-----*/

SDRDeleteObjects EM_Resource_Monitor \
'rmName==IBM.PSSP.SampleCmdMon'

SDRDeleteObjects EM_Resource_Class \
'rcClass==IBM.PSSP.SampleCmdClass'

SDRDeleteObjects EM_Resource_Variable \
'rvName==IBM.PSSP.SampleCmdMon.state'

SDRDeleteObjects EM_Resource_Variable \
'rvName==IBM.PSSP.SampleCmdMon.call'

SDRDeleteObjects EM_Resource_ID \
'riResource_name==IBM.PSSP.SampleCmdMon'

SDRDeleteObjects EM_Structured_Byte_String \
'sbsVariable_name==IBM.PSSP.SampleCmdMon.call'

#-----*/
```

```
# Unload resource monitor data for example resource monitor rmapi_smpsig
#-----
SDRDeleteObjects EM_Resource_Monitor \
'rmName==IBM.PSSP.SampleSigMon'

SDRDeleteObjects EM_Resource_Class \
'rcClass==IBM.PSSP.SampleSigDynInstClass'

SDRDeleteObjects EM_Resource_Class \
'rcClass==IBM.PSSP.SampleSigNonInstClass'

SDRDeleteObjects EM_Resource_Variable \
'rvName==IBM.PSSP.SampleSigMon.NonInstVars.var1'

SDRDeleteObjects EM_Resource_Variable \
'rvName==IBM.PSSP.SampleSigMon.NonInstVars.var2'

SDRDeleteObjects EM_Resource_Variable \
'rvName==IBM.PSSP.SampleSigMon.NonInstVars.var3'

SDRDeleteObjects EM_Resource_ID \
'riResource_name==IBM.PSSP.SampleSigMon.NonInstVars'

SDRDeleteObjects EM_Resource_Variable \
'rvName==IBM.PSSP.SampleSigMon.DynInstVar.var'

SDRDeleteObjects EM_Resource_ID \
'riResource_name==IBM.PSSP.SampleSigMon.DynInstVar'
```

`rmapi_smp.unloadsdr`

---

## Chapter 6. Using the EMAPI: Some Event Management Client Examples

This chapter contains the listings of four sample Event Management client programs:

- “The `emapi_v02_ex01.c` Sample Program” monitors four programs that are running on various nodes in two SP system partitions. Because two system partitions are involved, the program establishes two sessions with the EMAPI, one per system partition.

Because the program is not multi-threaded, it uses the **select** system call to multiplex responses from the two EMAPI sessions.

- “The `emapi_v02_ex02.c` Sample Program” on page 194 provides the same function as the **emapi\_v02\_ex01.c** program. It differs in that it uses callback routines to receive events.
- “The `emapi_v02_ex03.c` Sample Program” on page 206 provides the same function as the first two programs. It differs in that it is multi-threaded. Each session is managed by a separate thread.
- “The `emapi_v02_ex04.c` Sample Program” on page 219 illustrates the use of the EMAPI query function.

You can find the files for these programs online in the RSCT product samples directory, `/usr/sbin/rsct/samples/haem/emapi`.

---

### The `emapi_v02_ex01.c` Sample Program

```
/* IBM_PROLOG_BEGIN_TAG                               */
/* This is an automatically generated prolog.         */
/*                                                    */
/*                                                    */
/* Licensed Materials - Property of IBM               */
/*                                                    */
/* (C) COPYRIGHT International Business Machines Corp. 1996,1998 */
/* All Rights Reserved                               */
/*                                                    */
/* US Government Users Restricted Rights - Use, duplication or */
/* disclosure restricted by GSA ADP Schedule Contract with IBM Corp. */
/*                                                    */
/* IBM_PROLOG_END_TAG                               */

/* @(#)42 1.3 src/rsct/pem/emtools/emapi_test/emapi_ex/emapi_v02_ex01.c, emtools, rsct_rtr0 6/30/98 10:47:07 */

#include <unistd.h>
#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <errno.h>
#include <signal.h>
#include <time.h>

#include <sys/types.h>
#include <sys/select.h>
#include <sys/time.h>

#define HA_EM_VERSION 2
#include <ha_emapi.h>

/*
 * emapi_v02_ex01.c
 *
 * This program presents an example of using the Event Manager Application
 * Programming Interface (EMAPI). The program uses the EMAPI to monitor
 * four programs running on various nodes in two SP partitions. Since two
 * SP partitions are involved, the program establishes two sessions with the
 * EMAPI, one per SP partition. Since the program is not multi-threaded,
 * select() is used to multiplex responses from the two EMAPI sessions.
 *
 * This program can be compiled with the following command:
 *
 * cc -O emapi_v02_ex01.c -o emapi_v02_ex01 -lha_em
 */

/*
 * If a MAX macro isn't defined by one of the included header files, define
 * it.
 */

#ifndef MAX
#define MAX(x, y) ((x) > (y) ? (x) : (y))
#endif
```

## emapi\_v02\_ex01.c

```
/*
 * The following macros specify the programs this program will monitor through
 * the EMAPI.
 */

#define PART_1      "k21s"           /* 1st SP partition */
#define PROG_1A    "subsys_pgrm1"   /* 1st program in 1st SP partition, */
#define USER_1A   "root"           /* ... and user running 1st program, */
#define NODE_1A    5                /* ... and node running 1st program. */

#define PROG_1B    "subsys_pgrm2"   /* 2nd program in 1st SP partition, */
#define USER_1B   "root"           /* ... and user running 2nd program, */
#define NODE_1B    6                /* ... and node running 2nd program. */

#define PART_2      "k21sp2"        /* 2nd SP partition */
#define PROG_2A    "subsys_pgrm1"   /* 1st program in 2nd SP partition, */
#define USER_2A   "root"           /* ... and user running 1st program, */
#define NODE_2A    14               /* ... and node running 1st program. */

#define PROG_2B    "subsys_pgrm2"   /* 2nd program in 2nd SP partition, */
#define USER_2B   "root"           /* ... and user running 2nd program, */
#define NODE_2B    15               /* ... and node running 2nd program. */

/*
 * This program saves information about each EMAPI session it has established
 * in a session structure.
 */

struct session {
    char    *name;      /* Name of partition used by session */
    int     fd;         /* Session file descriptor */
    int     restart;    /* Boolean - Session's connection has been lost */
};

/*
 * This program saves information about each program it is monitoring through
 * the EMAPI in a program structure.
 */

struct program {
    char    *name;      /* Name of a program to be monitored */
    char    *user;      /* Name of user running program */
    int     node;       /* Node on which program is running */
    ha_em_eid_t eid;    /* Event identifier for program */
    int     unreged;    /* Event is unregistered */
};

/*
 * The terminate_requested global variable is set to a non-zero value when
 * the user of this program requests that the program be terminated. The
 * user requests program termination via the interrupt key (typically Ctrl-C).
 * The interrupt key causes a SIGINT signal to be delivered to this program.
 * This program installs a signal handler for the SIGINT signal, which sets
 * the terminate_requested global variable to a non-zero value when SIGINT
 * is delivered.
 */

static int terminate_requested = 0;

/*
 * The unregistrations_requested global variable indicates whether this
 * program has sent unregistration requests to the Event Manager.
 * Unregistration requests are sent when the user requests program
 * termination.
 */

static int unregistrations_requested = 0;

/*
 * Function prototypes for internal functions.
 */

static void interrupt_catch(int signo);
static void setup_signals(void);
static void start_session(struct session *sess_p);
static void register_for_events(struct session *sess_p,
    struct program *progs_p, int progs_elems, int *reg_cnt_p);
static void unregister_events(struct session *sess_p,
    struct program *progs_p, int progs_elems);
static void select_session(struct session *sess_p,
    fd_set *read_fds_p, int *fd_limit_p, int *timeout_p);
static int session_response_ready(struct session *sess_p, fd_set *read_fds_p);
static void process_response(struct session *sess_p,
    struct program *progs_p, int progs_elems, int *reg_cnt_p);
static void process_response_reg(struct session *sess_p,
    struct program *progs_p, int progs_elems,
    struct ha_em_rsp_blk *rsp_blk);
static void process_response_rerr(struct session *sess_p,
    struct program *progs_p, int progs_elems,
```

```

    struct ha_em_rsp_blk *rsp_blk, int *reg_cnt_p);
static void process_response_unreg(struct session *sess_p,
    struct program *progs_p, int progs_elems,
    struct ha_em_rsp_blk *rsp_blk, int *reg_cnt_p);
static void format_timestamp(struct timeval *timestamp_p,
    char *fmt_timestamp_p);
static void breakdown_rsrc_ID(char *rsrc_ID_p,
    char *prog_name_p, char *user_name_p, int *node_p);
static void find_rsrc_ID_value(char *rsrc_ID_p, char *name_p,
    char **value_pp, size_t *value_len_p);
static void end_session(struct session *sess_p);

/*
 * The main() function calls functions to initialize the program, establish
 * EMAPI sessions, register for events, wait for Event Manager responses,
 * process Event Manager responses, and respond to the user request for
 * program termination. When the program termination request is received,
 * main() calls functions to send unregistration requests through the EMAPI.
 * Once Event Manager responses have been received indicating that the
 * events have been unregistered, the program ends its EMAPI sessions,
 * and terminates.
 */
int main(int argc, char **argv)
{
    struct session sess1 = {PART_1, -1, 0};
    struct session sess2 = {PART_2, -1, 0};

    struct program progs1[] = {{PROG_1A, USER_1A, NODE_1A, 0, 0},
        {PROG_1B, USER_1B, NODE_1B, 0, 0}};
    struct program progs2[] = {{PROG_2A, USER_2A, NODE_2A, 0, 0},
        {PROG_2B, USER_2B, NODE_2B, 0, 0}};

    int    progs1_elems = sizeof progs1 / sizeof progs1[0];
    int    progs2_elems = sizeof progs2 / sizeof progs2[0];

    int    reg_cnt;           /* Count of registered events */
    fd_set read_fds;         /* select() file descriptor mask */
    int    fd_limit;         /* select() file descriptor limit */
    int    need_timeout;     /* Boolean - select() timeout needed */
    int    rc;               /* Return code */

    struct timeval timeout_value = {15, 0}; /* Fifteen second timeout */

    setup_signals();        /* Initialize signal dispositions */

    start_session(&sess1); /* Start EMAPI session for 1st SP partition */
    start_session(&sess2); /* Start EMAPI session for 2nd SP partition */

    /*
     * Register for events in both EMAPI sessions that will monitor the
     * programs of interest. Maintain count of registered events.
     */

    reg_cnt = 0;
    register_for_events(&sess1, progs1, progs1_elems, &reg_cnt);
    register_for_events(&sess2, progs2, progs2_elems, &reg_cnt);

    /*
     * Continue looking for Event Manager responses as long as there are
     * registered events.
     */

    while (reg_cnt > 0) {

        /*
         * If program termination has been requested, and unregistration
         * requests have not been sent, send them now.
         */

        if (terminate_requested && !unregistrations_requested) {
            printf("Termination requested.\n");
            unregister_events(&sess1, progs1, progs1_elems);
            unregister_events(&sess2, progs2, progs2_elems);
            unregistrations_requested = 1;
        }

        /*
         * Construct the parameters to be passed to select(), and then
         * call select() to wait for Event Manager responses.
         */

        FD_ZERO(&read_fds); fd_limit = 0; need_timeout = 0;
        select_session(&sess1, &read_fds, &fd_limit, &need_timeout);
        select_session(&sess2, &read_fds, &fd_limit, &need_timeout);

        rc = select(fd_limit + 1, &read_fds, NULL, NULL,
            need_timeout ? &timeout_value : NULL);

        if (rc == -1) { /* select() indicates error */

```

## emapi\_v02\_ex01.c

```
        if (errno == EINTR) {          /* select() interrupted by signal */
            continue;                 /* Return to top of loop          */
        }
        perror("select()");           /* select() really encountered error */
        exit(1);                       /* End program                    */
    }

    /*
     * If an Event Manager response is present for the 1st EAPI session,
     * process it. Also, maintain the count of registered events.
     */

    if (session_response_ready(&sess1, &read_fds)) {
        process_response(&sess1, progs1, progs1_elems, &reg_cnt);
    }

    /*
     * If an Event Manager response is present for the 2nd EAPI session,
     * process it. Also, maintain the count of registered events.
     */

    if (session_response_ready(&sess2, &read_fds)) {
        process_response(&sess2, progs2, progs2_elems, &reg_cnt);
    }

}

end_session(&sess1);                 /* End EAPI session for 1st SP partition */
end_session(&sess2);                 /* End EAPI session for 2nd SP partition */

return 0;                            /* Indicate program was successful      */
}

/*
 * The interrupt_catch() function is installed as the signal handler for the
 * SIGINT signal. The SIGINT signal is delivered to this process when the
 * user presses the interrupt key (usually Ctrl-C). When this routine is
 * invoked, it sets the global variable terminate_requested to a non-zero
 * value, indicating the program is to be terminated.
 */

static
void interrupt_catch(int signo)
{
    terminate_requested = 1;
    return;
}

/*
 * The setup_signals() function initializes the disposition of a couple of
 * signals.
 *
 * The disposition of the SIGPIPE signal is set such that the signal is
 * ignored. When a process writes to a pipe or connection-oriented socket for
 * which there is no reader, the SIGPIPE signal is delivered to the process.
 * The default disposition for SIGPIPE is to kill the receiving process. If
 * SIGPIPE is ignored, the SIGPIPE signal does not kill the process, and the
 * system call used to write to the pipe or connection-oriented socket sets
 * errno to EPIPE and returns an error indication. It is recommended that
 * clients of the EAPI ignore the SIGPIPE signal, since the EAPI uses
 * connection-oriented sockets to communicate with the Event Manager.
 *
 * The interrupt_catch() function is installed as the signal handler for the
 * SIGINT signal. See the comments pertaining to the interrupt_catch()
 * function to find out why this is done.
 */

static
void setup_signals(void)
{
    struct sigaction  sa;

    sa.sa_handler = SIG_IGN;          /* SIGPIPE to be ignored          */
    sigemptyset(&sa.sa_mask);        /* No signals to mask off         */
    sa.sa_flags = 0;                  /* No flags needed                */

    if (sigaction(SIGPIPE, &sa, NULL) == -1) { /* Change SIGPIPE disposition*/
        perror("sigaction(SIGPIPE)");
        exit(1);
    }

    sa.sa_handler = interrupt_catch; /* SIGINT signal handler          */
    sigemptyset(&sa.sa_mask);        /* No signals to mask off         */
    sa.sa_flags = SA_RESTART;        /* Restart interrupted syscalls   */

    if (sigaction(SIGINT, &sa, NULL) == -1) { /* Change SIGINT disposition */
        perror("sigaction(SIGINT)");
        exit(1);
    }
}
```

```

    return;
}

/*
 * The start_session() function establishes a session with the EAPI by
 * calling the EAPI routine ha_em_start_session(). If a session cannot be
 * established, the program is terminated.
 */

static
void start_session(struct session *sess_p)
{
    struct ha_em_err_blk errb;          /* EAPI error block */

    /*
     * The ha_em_start_session() routine receives the name of the partition
     * with which a session is to be established, and returns a file
     * descriptor with which the session can be used.
     */

    sess_p->fd = ha_em_start_session(HA_EM_DOMAIN_SP, sess_p->name, &errb);

    /*
     * If ha_em_start_session() indicates a session could not be established,
     * print the error information returned by the routine, and exit the
     * program.
     */

    if (sess_p->fd == -1) {
        fprintf(stderr, "%s:\tha_em_start_session() returned EM errno %d:\n%s",
            sess_p->name, errb.em_errno, errb.em_errmsg);
        exit(1);
    }

    /*
     * A session has been established with the EAPI, and the session's file
     * descriptor has been saved in the session structure passed to this
     * routine. Set the flag in the session structure that indicates the
     * session does not need to be restarted.
     */

    sess_p->restart = 0;

    return;
}

/*
 * The register_for_events() function uses the EAPI routine
 * ha_em_send_command() to request the registration of events pertaining to
 * programs of interest that are in the same SP partition.
 *
 * The resource variable name specified is "IBM.PSSP.Prog.xpcount". This is a
 * state variable whose value indicates whether or not processes are executing
 * a specific program. The resource ID for this resource variable
 * specifies the program name, the name of the user running the program, and
 * the node on which the program is running. The resource variable's value
 * is a structured byte string of three fields. Field 0 is of type long; its
 * value is the number of processes currently running the specified program
 * for the specified user on the specified node. Field 1 is also of type
 * long; its value is the previous number of processes running the program.
 * Field 2 is of type character string; its value is a comma separated list
 * of the PIDs of the processes running the program.
 *
 * When the events are registered, two expressions are specified. The primary
 * expression, "X00 == 0", specifies that an event is to be generated when
 * field 0 of the resource variable's structured byte string has a value of 0.
 * Taking into account the semantics of the resource variable, the expression
 * causes an event to be generated when no processes are running the specified
 * program for the specified user on the specified node. The re-arm
 * expression, "X00 > 0", specifies that after the primary expression has
 * generated an event, the primary expression is to be re-armed once a process
 * starts running the specified program for the specified user on the
 * specified node. Since the HA_EM_CMD_REG2 command is specified, the re-arm
 * expression will generate events.
 *
 * Notice that when events are registered, the HA_EM_SCMD_REVAL subcommand is
 * specified. As a result, the initial value of the programs being monitored
 * will be returned.
 */

static
void register_for_events(struct session *sess_p,
    struct program *progs_p, int progs_elems, int *reg_cnt_p)
{
    int          num_events;          /* Number of events to register */
    size_t       alloc_size;         /* Size of memory to allocate */
    struct ha_em_cmd_blk *cmd_blk;   /* Pointer to command block */
    struct ha_em_rb_reg *re;         /* Pointer to registered event */

```

## emapi\_v02\_ex01.c

```
int i; /* Index */
struct ha_em_err_blk errb; /* EMAPI error block */
int rc; /* Return code */
char rsrc_ID_buf[80]; /* Resource ID buffer */

num_events = progs_elems; /* Register one event per program */

/*
 * Allocate enough memory to hold the command block to be given to the
 * EMAPI. The ha_em_cmd_blk structure contains one ha_em_rb_reg
 * structure. Additional space must be allocated when more than one
 * event is to be registered.
 */
alloc_size = sizeof(struct ha_em_cmd_blk) +
             ((num_events - 1) * sizeof(struct ha_em_rb_reg));
if ((cmd_blk = malloc(alloc_size)) == NULL) {
    perror("malloc()");
    exit(1);
}

/*
 * Fill in the em_rb_reg array entries. There is one entry per event to
 * be registered.
 */
for (i = 0; i < num_events; i++) {
    re = &cmd_blk->em_res_blk.em_rb_reg[i];

    re->em_name = "IBM.PSSP.Prog.xpcount"; /* Resource variable name */

    sprintf(rsrc_ID_buf, "ProgName=%s;UserName=%s;NodeNum=%d",
            progs_p[i].name, progs_p[i].user, progs_p[i].node);
    re->em_rsrc_ID = strdup(rsrc_ID_buf); /* Resource ID */

    if (re->em_rsrc_ID == NULL) {
        perror("strdup()");
        exit(1);
    }

    re->em_expr = "X00 == 0"; /* Expression */
    re->em_raexpr = "X00 > 0"; /* Re-arm expression */

    re->em_cb = NULL; /* Callback routine not used*/
    re->em_cb_arg = NULL; /* Callback routine not used*/
}

/*
 * Fill in command block header, specifying number of elements in
 * em_rb_reg array, command, and subcommand.
 */
cmd_blk->em_cmd_num_elem = num_events;
cmd_blk->em_cmd = HA_EM_CMD_REG2;
cmd_blk->em_subcmd = HA_EM_SCMD_REVAL;

/*
 * Send the command. A more elaborate program might check for the
 * HA_EM_ECONNLOST Event Manager error number when an error is
 * returned, and attempt to restart the session. But, this program
 * just terminates if ha_em_send_command() detects a lost connection.
 */
rc = ha_em_send_command(sess_p->fd, cmd_blk, &errb);

if (rc == -1) {
    fprintf(stderr, "%s:\tha_em_send_command() returned EM errno %d:\n%s",
            sess_p->name, errb.em_errno, errb.em_errmsg);
    exit(1);
}

/*
 * Save the event identifiers assigned to the events just registered.
 */
for (i = 0; i < num_events; i++) {
    re = &cmd_blk->em_res_blk.em_rb_reg[i];
    progs_p[i].eid = re->em_event_id;
    progs_p[i].unreged = 0;
    free(re->em_rsrc_ID);
}

free(cmd_blk);

*reg_cnt_p += num_events; /* Update count of registered events */

return;
}
```

```

/*
 * The unregister_events() function uses the EMAPI routine
 * ha_em_send_command() to request the unregistration of events pertaining to
 * programs of interest that are in the same SP partition.
 */

static
void unregister_events(struct session *sess_p,
                      struct program *progs_p, int progs_elems)
{
    int          num_events;      /* Number of events to unreg. */
    size_t      alloc_size;      /* Size of memory to allocate */
    struct ha_em_cmd_blk *cmd_blk; /* Pointer to command block */
    ha_em_eid_t *re;             /* Pointer to unreg. event */
    int         i;               /* Index */
    struct ha_em_err_blk errb;    /* EMAPI error block */
    int         rc;              /* Return code */

    /*
     * Allocate enough memory to hold the command block to be given to the
     * EMAPI. The ha_em_cmd_blk structure contains one ha_em_eid_t
     * element. Additional space must be allocated when more than one
     * event is to be unregistered.
     */

    alloc_size = sizeof(struct ha_em_cmd_blk) +
                 ((progs_elems - 1) * sizeof(ha_em_eid_t));

    if ((cmd_blk = malloc(alloc_size)) == NULL) {
        perror("malloc()");
        exit(1);
    }

    /*
     * Fill in the event identifier array entries. There is one entry per
     * event to be unregistered. Only unregister events that are not already
     * unregistered.
     */

    num_events = 0;

    for (i = 0; i < progs_elems; i++) {
        if (!progs_p[i].unreged) {
            re = &cmd_blk->em_res_blk.em_rb_unreg[num_events];
            *re = progs_p[i].eid;
            num_events++;
        }
    }

    /*
     * Fill in command block header, specifying the number of elements in
     * the event identifier array, and the unregister command.
     */

    cmd_blk->em_cmd_num_elem = num_events;
    cmd_blk->em_cmd          = HA_EM_CMD_UNREG;

    /*
     * Send the unregister command. A more elaborate program might check for
     * the HA_EM_ECONNLOST Event Manager error number when an error is
     * returned, and attempt to restart the session. But, this program
     * just terminates if ha_em_send_command() detects a lost connection.
     */

    rc = ha_em_send_command(sess_p->fd, cmd_blk, &errb);

    if (rc == -1) {
        fprintf(stderr, "%s:\tha_em_send_command() returned EM errno %d:\n%s",
                sess_p->name, errb.em_errno, errb.em_errmsg);
        exit(1);
    }

    free(cmd_blk);

    return;
}

/*
 * The select_session() routine determines how the specified session affects
 * the next call to select(). If the session is marked as needing to be
 * restarted, ha_em_restart_session() is called to try to restart the session.
 * If the restart attempt fails, the next call to select() should specify
 * a timeout period - so the restart can be attempted again. If the restart
 * attempt succeeds, the new session file descriptor is saved, and the session
 * is marked as not needing to be restarted. If the session did not need
 * to be restarted, or had been successfully restarted, the next call to
 * select() should examine the session's file descriptor.
 */

```

## emapi\_v02\_ex01.c

```
static
void select_session(struct session *sess_p,
                   fd_set *read_fds_p, int *fd_limit_p, int *timeout_p)
{
    struct ha_em_err_blk  errb;          /* EMAPI error block */
    int                  new_fd;        /* New session file descriptor */

    /*
     * If the specified session needs to be restarted, try to do so.
     */
    if (sess_p->restart) {
        new_fd = ha_em_restart_session(sess_p->fd, &errb);

        if (new_fd == -1) {
            /*
             * The session could not be restarted. If the Event Manager
             * error number indicates connection refused, indicate the
             * next select() should specify a timeout value (so the restart
             * can be tried again later), and return without putting the
             * session's file descriptor in the select() file descriptor
             * mask.
             */

            if (errb.em_errno == HA_EM_ECONNREFUSED) {
                *timeout_p = 1;
                return;
            }

            /*
             * If execution reaches here, the session restart attempt failed
             * for an unexpected reason; terminate the program.
             */

            fprintf(stderr, "%s:\tha_em_restart_session() returned EM errno "
                    "%d:\n%s", sess_p->name, errb.em_errno, errb.em_errmsg);
            exit(1);
        }

        /*
         * If execution reaches here, a session restart was attempted and
         * was successful. Save the new session file descriptor, and indicate
         * that the session no longer needs to be restarted.
         */

        sess_p->fd = new_fd;
        sess_p->restart = 0;
    }

    /*
     * If execution reaches here, the specified session is not known to be
     * in need of a restart. Put the session's file descriptor in the
     * select() file descriptor mask. Also, adjust the maximum file
     * descriptor to select, if appropriate.
     */

    FD_SET(sess_p->fd, read_fds_p);
    *fd_limit_p = MAX(sess_p->fd, *fd_limit_p);

    return;
}

/*
 * The session_response_ready() function determines if a select() file
 * descriptor mask indicates that a session's file descriptor is ready for
 * reading.
 */

static
int session_response_ready(struct session *sess_p, fd_set *read_fds_p)
{
    return FD_ISSET(sess_p->fd, read_fds_p);
}

/*
 * The process_response() function processes an Event Manager response
 * in the specified EMAPI session. First, ha_em_receive_response() is called
 * to receive the response. If ha_em_receive_response() indicates that the
 * session's connection with the Event Manager has been lost, the session
 * is marked as being in need of a restart. Any other error will terminate
 * the program. The ha_em_receive_response() function may indicate that
 * there really isn't any response for the EMAPI client to process at this
 * time. In that case, this function just returns. If a response has
 * been received, it is processed depending on what type of response it is.
 */
```

```

static
void process_response(struct session *sess_p,
                    struct program *progs_p, int progs_elems, int *reg_cnt_p)
{
    struct ha_em_rsp_blk *rsp_blk; /* Pointer to the response block */
    int rc; /* Return code */
    struct ha_em_err_blk errb; /* Event Manager error block */
    int i; /* Index */

    /*
     * Receive response from session, if there is one to receive.
     */
    rc = ha_em_receive_response(sess_p->fd, &rsp_blk, &errb);

    if (rc == -1) {
        /*
         * If the connection with the Event Manager has been lost,
         * mark the session as needing to be restarted, and return.
         */
        if (errb.em_errno == HA_EM_ECONNLOST) {
            printf("%s:\tconnection to session lost.\n", sess_p->name);
            sess_p->restart = 1;
            return;
        }

        /*
         * Some error has occurred other than the loss of the connection.
         * Terminate the program.
         */
        fprintf(stderr, "%s:\tha_em_receive_response() returned EM errno %d:\n"
            "%s", sess_p->name, errb.em_errno, errb.em_errmsg);
        exit(1);
    }

    /*
     * If the ha_em_receive_response() routine returned zero,
     * there is no response for the EMAPI client (this program) to process
     * at this time. The response may have been for the EMAPI itself, or
     * a full response may not be available yet. Just return.
     */
    if (rc == 0) {
        return;
    }

    /*
     * A response for the EMAPI client (this program) has been returned.
     * The response buffer is pointed to by the rsp_blk variable. Appropriate
     * processing for the response depends on the command type.
     */
    switch (rsp_blk->em_cmd) {
        case HA_EM_CMD_REG: /* Event response */
        case HA_EM_CMD_REG2: /* Event response (possible re-arm) */
            process_response_reg(sess_p, progs_p, progs_elems, rsp_blk);
            break;

        case HA_EM_CMD_RERR: /* REG event registration error */
        case HA_EM_CMD_R2ERR: /* REG2 event registration error */
            process_response_rerr(sess_p, progs_p, progs_elems, rsp_blk,
                reg_cnt_p);
            break;

        case HA_EM_CMD_UNREG: /* Event unregistration response */
            process_response_unreg(sess_p, progs_p, progs_elems, rsp_blk,
                reg_cnt_p);
            break;

        default: /* Unexpected response */
            fprintf(stderr, "%s:\tProgram received unexpected command "
                "response: %d.\n", sess_p->name, rsp_blk->em_cmd);
            exit(1);
            break;
    }

    /*
     * The EMAPI client (this program) must free the memory associated with
     * the returned response block when it is no longer needed.
     */
    free(rsp_blk);

    return;
}

```

## emapi\_v02\_ex01.c

```
/*
 * The process_response_reg() function processes event responses from the
 * Event Manager. Several points should be kept in mind:
 *
 * - The response block may contain multiple event responses. The
 *   number of responses included in the response block is given in
 *   the response block header.
 *
 * - An event response may be an error response. An event error response
 *   does not indicate that a registered event has occurred. Instead, it
 *   indicates the Event Manager no longer knows the current value of
 *   the associated resource variable, and cannot generate events for
 *   the resource variable. This may be a temporary condition. If the
 *   Event Manager later obtains the current value of the associated
 *   resource variable, an event will be generated (possibly indicating
 *   that the expression is false).
 *
 * - The event for which an event error response is generated is still
 *   registered. There is no need to re-register the event.
 *
 * - Since events are registered by this program using the HA_EM_CMD_REG2
 *   command, the re-arm expressions will generate events. Therefore, the
 *   HA_EM_EVENT_RE_ARM flag must be tested to see if the expression or
 *   the re-arm expression generated the event.
 *
 * - It is possible for an event to be delivered indicating the expression
 *   is false. This can happen for two reasons. First, when events are
 *   registered by this program, the HA_EM_SCMD_REVAL subcommand is
 *   specified. That subcommand requests the initial value of the
 *   associated resource variable. Second, the current value of the
 *   associated resource variable is returned through an event once the
 *   Event Manager obtains the current value of a resource variable after
 *   an error has occurred.
 *
 * - If an event response does not have the HA_EM_EVENT_RE_ARM nor
 *   HA_EM_EVENT_EXPR_FALSE flags set, the event response indicates that
 *   the event's expression is true.
 */

static
void process_response_reg(struct session *sess_p,
                        struct program *progs_p, int progs_elems,
                        struct ha_em_rsp_blk *rsp_blk)
{
    struct ha_em_rpb_event *event_p;          /* Pointer to event response */
    struct ha_em_rpb_event *last_event_p;    /* Beyond last event response*/

    char time_stamp[80];                    /* Formatted event time stamp */
    char prog_name[80];                     /* Name of program to which event pertains */
    char user_name[80];                     /* Name of user to which event pertains */
    int node;                               /* Node to which event pertains */

    event_p = rsp_blk->em_resp_blk.em_rpb_event;
    last_event_p = event_p + rsp_blk->em_rsp_num_resp;

    for ( ; event_p < last_event_p; event_p++) {

        /*
         * Format the event's time stamp; extract the program name, user
         * name, and node number from the resource ID.
         */
        format_timestamp(&event_p->em_timestamp, time_stamp);
        breakdown_rsrc_ID(event_p->em_rsrc_ID, prog_name, user_name, &node);

        /*
         * If this is an event error response, print a message which includes
         * the general and specific error codes.
         */
        if (event_p->em_errnum != 0) {
            printf("%s:\t%s State of %s run by %s on node %d unknown (%d, %d)."\n",
                  sess_p->name, time_stamp, prog_name, user_name, node,
                  event_p->em_generr, event_p->em_specerr);
            continue;
        }

        /*
         * If the event response was generated for the re-arm expression, or
         * the event response was generated for the primary expression but
         * it is false, print a message indicating the program associated
         * with the event is being run.
         */
        if ((event_p->em_event_flags & HA_EM_EVENT_RE_ARM) ||
            (event_p->em_event_flags & HA_EM_EVENT_EXPR_FALSE)) {
            printf("%s:\t%s %s being run by %s on node %d.\n",
                  sess_p->name, time_stamp, prog_name, user_name, node);
            continue;
        }
    }
}
```

```

    }

    /*
     * If execution reaches this point, the event response was generated
     * for the primary expression, and it is true. Print a message
     * indicating the program associated with the event is not being run.
     */

    printf("%s:\t%s %s NOT being run by %s on node %d.\n",
           sess_p->name, time_stamp, prog_name, user_name, node);

    }

return;
}

/*
 * The process_response_rerr() function processes registration error
 * responses. When an event cannot be registered due to some detected error,
 * a registration error response is sent to the EMAPI client. This function
 * marks the event, as tracked in a program structure, as being unregistered.
 * This will prevent an attempt to unregister a non-registered event later
 * in the program.
 */

static
void process_response_rerr(struct session *sess_p,
                          struct program *progs_p, int progs_elems,
                          struct ha_em_rsp_blk *rsp_blk, int *reg_cnt_p)
{
    struct ha_em_rpb_rerr *error_p;          /* Pointer to error response */
    struct ha_em_rpb_rerr *last_error_p;    /* Beyond last error response*/
    struct program *prog_p;                 /* Pointer to program struct */

    char prog_name[80];                     /* Name of program to which error pertains */
    char user_name[80];                     /* Name of user to which error pertains */
    int node;                               /* Node to which error pertains */

    error_p = rsp_blk->em_resp_blk.em_rpb_rerr;
    last_error_p = error_p + rsp_blk->em_rsp_num_resp;

    for ( ; error_p < last_error_p; error_p++) {

        /*
         * Find the program structure which describes the program associated
         * with the event whose registration failed. The match is made
         * with event identifiers.
         */

        for (prog_p = progs_p; prog_p < progs_p + progs_elems; prog_p++) {
            if (prog_p->eid == error_p->em_event_id) {
                break;
            }
        }

        if (prog_p == progs_p + progs_elems) {
            fprintf(stderr, "%s:\tUnknown event identifier encountered (0x%x)."\n",
                   sess_p->name, error_p->em_event_id);
            exit(1);
        }

        /*
         * Extract the program name, user name, and node number from the
         * resource ID.
         */

        breakdown_rsrc_ID(error_p->em_rsrc_ID, prog_name, user_name, &node);

        /*
         * Print a message about the registration error.
         */

        printf("%s:\tRegistration error for %s run by %s on node %d (%d, %d)."\n",
               sess_p->name, prog_name, user_name, node,
               error_p->em_generr, error_p->em_specerr);

        /*
         * Mark the program as not having an associated event registered.
         */

        prog_p->unreged = 1;

    }

    /*
     * Update the number of events that are registered.
     */

    *reg_cnt_p -= rsp_blk->em_rsp_num_resp;

```

## emapi\_v02\_ex01.c

```
    return;
}

/*
 * The process_response_unreg() function processes unregistrations responses.
 * Note that the unregistration response uses the same type of structure
 * as the event response, but not all the fields in the structure have
 * meaningful values for an unregistration response.
 */

static
void process_response_unreg(struct session *sess_p,
    struct program *progs_p, int progs_elems,
    struct ha_em_rsp_blk *rsp_blk, int *reg_cnt_p)
{
    struct ha_em_rpb_event *event_p;          /* Pointer to unreg response */
    struct ha_em_rpb_event *last_event_p;    /* Beyond last unreg response */
    struct program *prog_p;                  /* Pointer to program struct */

    char time_stamp[80];                     /* Formatted unregister time stamp */

    event_p = rsp_blk->em_resp_blk.em_rpb_event;
    last_event_p = event_p + rsp_blk->em_rsp_num_resp;

    for ( ; event_p < last_event_p; event_p++) {
        /*
         * Find the program structure which describes the program associated
         * with the event just unregistered. The match is made
         * with event identifiers.
         */
        for (prog_p = progs_p; prog_p < progs_p + progs_elems; prog_p++) {
            if (prog_p->eid == event_p->em_event_id) {
                break;
            }
        }

        if (prog_p == progs_p + progs_elems) {
            fprintf(stderr, "%s:\tUnknown event identifier encountered (0x%x).",
                "\n", sess_p->name, event_p->em_event_id);
            exit(1);
        }

        /*
         * Format the unregistration time stamp.
         */

        format_timestamp(&event_p->em_timestamp, time_stamp);

        /*
         * Print a message indicating the program associated with the
         * unregistered event is no longer being monitored.
         */

        printf("%s:\t%s no longer monitoring %s run by %s on node %d.\n",
            sess_p->name, time_stamp, prog_p->name, prog_p->user,
            prog_p->node);

        /*
         * Note: em_errnum could indicate an error, but there is no point
         * looking at it here, since the program would not do anything
         * differently.
         */

        /*
         * Mark the program as not having an associated event registered.
         */

        prog_p->unreged = 1;
    }

    /*
     * Update the number of events that are still registered.
     */

    *reg_cnt_p -= rsp_blk->em_rsp_num_resp;
}

return;
}

/*
 * The format_timestamp() function takes a time stamp returned by the
 * Event Manager and converts it into 24-hour time.
 */

static
void format_timestamp(struct timeval *timestamp_p, char *fmt_timestamp_p)
```

```

{
    struct tm    *broken_down_time;

    broken_down_time = localtime((time_t *) &timestamp_p->tv_sec);
    (void) strftime(fmt_timestamp_p, 20, "(%X)", broken_down_time);

    return;
}

/*
 * The breakdown_rsrc_ID() function takes a resource ID
 * associated with the IBM.PSSP.Prog.xpcount resource variable and extracts
 * from it the program name, user name, and node number.
 */
static
void breakdown_rsrc_ID(char *rsrc_ID_p,
    char *prog_name_p, char *user_name_p, int *node_p)
{
    char    *value_p;
    size_t  value_len;

    find_rsrc_ID_value(rsrc_ID_p, "ProgName=", &value_p, &value_len);
    strncpy(prog_name_p, value_p, value_len);
    *(prog_name_p + value_len) = '\0';

    find_rsrc_ID_value(rsrc_ID_p, "UserName=", &value_p, &value_len);
    strncpy(user_name_p, value_p, value_len);
    *(user_name_p + value_len) = '\0';

    find_rsrc_ID_value(rsrc_ID_p, "NodeNum=", &value_p, &value_len);
    *node_p = atoi(value_p);

    return;
}

/*
 * The find_rsrc_ID_value() takes a resource ID and extracts
 * a value from it.
 */
static
void find_rsrc_ID_value(char *rsrc_ID_p, char *name_p,
    char **value_pp, size_t *value_len_p)
{
    char    *bp;
    char    *ep;
    size_t  len;

    if ((bp = strstr(rsrc_ID_p, name_p)) == NULL) {
        fprintf(stderr, "Unexpected resource ID encountered.\n");
        exit(1);
    }

    bp += strlen(name_p);

    if ((ep = strchr(bp, ';')) != NULL) {
        len = ep - bp;
    } else {
        len = strlen(bp);
    }

    *value_pp = bp;
    *value_len_p = len;

    return;
}

/*
 * The end_session() function terminates a session with the EMAPI by
 * calling the EMAPI routine ha_em_end_session().
 */
static
void end_session(struct session *sess_p)
{
    struct ha_em_err_blk errb;

    if (ha_em_end_session(sess_p->fd, &errb) == -1) {
        fprintf(stderr, "%s:\tha_em_end_session() returned EM errno %d:\n%s",
            sess_p->name, errb.em_errno, errb.em_errmsg);
        exit(1);
    }
}

```

## emapi\_v02\_ex01.c

```
    }  
    return;  
}
```

---

## The emapi\_v02\_ex02.c Sample Program

```
/* IBM_PROLOG_BEGIN_TAG                               */  
/* This is an automatically generated prolog.         */  
/*                                                    */  
/*                                                    */  
/* Licensed Materials - Property of IBM              */  
/*                                                    */  
/* (C) COPYRIGHT International Business Machines Corp. 1996,1998 */  
/* All Rights Reserved                               */  
/*                                                    */  
/* US Government Users Restricted Rights - Use, duplication or */  
/* disclosure restricted by GSA ADP Schedule Contract with IBM Corp. */  
/*                                                    */  
/* IBM_PROLOG_END_TAG                               */  
  
/* @(#)43 1.3 src/rsct/pem/emtools/emapi_test/emapi_ex/emapi_v02_ex02.c, emtools, rsct_rtro 6/30/98 10:47:26 */  
  
#include <unistd.h>  
#include <stdlib.h>  
#include <stdio.h>  
#include <string.h>  
#include <errno.h>  
#include <signal.h>  
#include <time.h>  
  
#include <sys/types.h>  
#include <sys/select.h>  
#include <sys/time.h>  
  
#define HA_EM_VERSION 2  
#include <ha_emapi.h>  
  
/*  
 * emapi_v02_ex02.c  
 *  
 * This program presents an example of using the Event Manager Application  
 * Programming Interface (EMAPI). The program uses the EMAPI to monitor  
 * four programs running on various nodes in two SP partitions. Since two  
 * SP partitions are involved, the program establishes two sessions with the  
 * EMAPI, one per SP partition. Since the program is not multi-threaded,  
 * select() is used to multiplex responses from the two EMAPI sessions.  
 *  
 * This program provides the same function as the emapi_v02_ex01.c program.  
 * This program differs in that it uses callback routines to receive events.  
 *  
 * This program can be compiled with the following command:  
 *  
 * cc -O emapi_v02_ex02.c -o emapi_v02_ex02 -lha_em  
 */  
  
/*  
 * If a MAX macro isn't defined by one of the included header files, define  
 * it.  
 */  
  
#ifndef MAX  
#define MAX(x, y) ((x) > (y) ? (x) : (y))  
#endif  
  
/*  
 * The following macros specify the programs this program will monitor through  
 * the EMAPI.  
 */  
  
#define PART_1 "k21s" /* 1st SP partition */  
  
#define PROG_1A "subsys_pgrm1" /* 1st program in 1st SP partition, */  
#define USER_1A "root" /* ... and user running 1st program, */  
#define NODE_1A 5 /* ... and node running 1st program. */  
  
#define PROG_1B "subsys_pgrm2" /* 2nd program in 1st SP partition, */  
#define USER_1B "root" /* ... and user running 2nd program, */  
#define NODE_1B 6 /* ... and node running 2nd program. */  
  
#define PART_2 "k21sp2" /* 2nd SP partition */  
  
#define PROG_2A "subsys_pgrm1" /* 1st program in 2nd SP partition, */  
#define USER_2A "root" /* ... and user running 1st program, */  
#define NODE_2A 14 /* ... and node running 1st program. */  
  
#define PROG_2B "subsys_pgrm2" /* 2nd program in 2nd SP partition, */  
#define USER_2B "root" /* ... and user running 2nd program, */  
#define NODE_2B 15 /* ... and node running 2nd program. */
```

```

/*
 * This program saves information about each EMAPI session it has established
 * in a session structure.
 */
struct session {
    char    *name;        /* Name of partition used by session */
    int     fd;          /* Session file descriptor */
    int     restart;     /* Boolean - Session's connection has been lost */
};

/*
 * This program saves information about each program it is monitoring through
 * the EMAPI in a program structure.
 */
struct program {
    char    *name;        /* Name of a program to be monitored */
    char    *user;        /* Name of user running program */
    int     node;         /* Node on which program is running */
    ha_em_eid_t eid;     /* Event identifier for program */
    int     unreged;     /* Event is unregistered */
    struct callback_data *cb_data; /* Callback data allocated for program */
};

/*
 * Data for a callback routine is stored in a callback_data structure.
 */
struct callback_data {
    struct session *sess_p; /* Session structure pointer */
    struct program *prog_p; /* Program structure pointer */
    int *reg_cnt_p; /* Registration counter pointer */
};

/*
 * The terminate_requested global variable is set to a non-zero value when
 * the user of this program requests that the program be terminated. The
 * user requests program termination via the interrupt key (typically Ctrl-C).
 * The interrupt key causes a SIGINT signal to be delivered to this program.
 * This program installs a signal handler for the SIGINT signal, which sets
 * the terminate_requested global variable to a non-zero value when SIGINT
 * is delivered.
 */
static int terminate_requested = 0;

/*
 * The unregistrations_requested global variable indicates whether this
 * program has sent unregistration requests to the Event Manager.
 * Unregistration requests are sent when the user requests program
 * termination.
 */
static int unregistrations_requested = 0;

/*
 * Function prototypes for internal functions.
 */
static void interrupt_catch(int signo);
static void setup_signals(void);
static void start_session(struct session *sess_p);
static void register_for_events(struct session *sess_p,
    struct program *progs_p, int progs_elems, int *reg_cnt_p);
static void unregister_events(struct session *sess_p,
    struct program *progs_p, int progs_elems);
static void select_session(struct session *sess_p,
    fd_set *read_fds_p, int *fd_limit_p, int *timeout_p);
static int session_response_ready(struct session *sess_p, fd_set *read_fds_p);
static void process_response(struct session *sess_p,
    struct program *progs_p, int progs_elems, int *reg_cnt_p);
static void process_response_rerr(struct session *sess_p,
    struct program *progs_p, int progs_elems,
    struct ha_em_rsp_blk *rsp_blk, int *reg_cnt_p);
static void event_callback(int sess_fd, struct ha_em_rpb_event *event_p,
    void *arg);
static void format_timestamp(struct timeval *timestamp_p,
    char *fmt_timestamp_p);
static void breakdown_rsrc_ID(char *rsrc_ID_p,
    char *prog_name_p, char *user_name_p, int *node_p);
static void find_rsrc_ID_value(char *rsrc_ID_p, char *name_p,
    char **value_pp, size_t *value_len_p);
static void end_session(struct session *sess_p);

/*
 * The main() function calls functions to initialize the program, establish
 * EMAPI sessions, register for events, wait for Event Manager responses,
 * process Event Manager responses, and respond to the user request for
 * program termination. When the program termination request is received,

```

## emapi\_v02\_ex02.c

```
* main() calls functions to send unregistration requests through the EMAPI.
* Once Event Manager responses have been received indicating that the
* events have been unregistered, the program ends its EMAPI sessions,
* and terminates.
*/

int main(int argc, char **argv)
{
    struct session sess1 = {PART_1, -1, 0};
    struct session sess2 = {PART_2, -1, 0};

    struct program progs1[] = {{PROG_1A, USER_1A, NODE_1A, 0, 0, NULL},
                               {PROG_1B, USER_1B, NODE_1B, 0, 0, NULL}};
    struct program progs2[] = {{PROG_2A, USER_2A, NODE_2A, 0, 0, NULL},
                               {PROG_2B, USER_2B, NODE_2B, 0, 0, NULL}};

    int    progs1_elems = sizeof progs1 / sizeof progs1[0];
    int    progs2_elems = sizeof progs2 / sizeof progs2[0];

    int    reg_cnt;           /* Count of registered events          */
    fd_set read_fds;         /* select() file descriptor mask      */
    int    fd_limit;         /* select() file descriptor limit     */
    int    need_timeout;     /* Boolean - select() timeout needed  */
    int    rc;               /* Return code                         */

    struct timeval timeout_value = {15, 0}; /* Fifteen second timeout */

    setup_signals();        /* Initialize signal dispositions     */

    start_session(&sess1);  /* Start EMAPI session for 1st SP partition */
    start_session(&sess2);  /* Start EMAPI session for 2nd SP partition */

    /*
     * Register for events in both EMAPI sessions that will monitor the
     * programs of interest. Maintain count of registered events.
     */

    reg_cnt = 0;
    register_for_events(&sess1, progs1, progs1_elems, &reg_cnt);
    register_for_events(&sess2, progs2, progs2_elems, &reg_cnt);

    /*
     * Continue looking for Event Manager responses as long as there are
     * registered events.
     */

    while (reg_cnt > 0) {
        /*
         * If program termination has been requested, and unregistration
         * requests have not been sent, send them now.
         */

        if (terminate_requested && !unregistrations_requested) {
            printf("Termination requested.\n");
            unregister_events(&sess1, progs1, progs1_elems);
            unregister_events(&sess2, progs2, progs2_elems);
            unregistrations_requested = 1;
        }

        /*
         * Construct the parameters to be passed to select(), and then
         * call select() to wait for Event Manager responses.
         */

        FD_ZERO(&read_fds); fd_limit = 0; need_timeout = 0;
        select_session(&sess1, &read_fds, &fd_limit, &need_timeout);
        select_session(&sess2, &read_fds, &fd_limit, &need_timeout);

        rc = select(fd_limit + 1, &read_fds, NULL, NULL,
                   need_timeout ? &timeout_value : NULL);

        if (rc == -1) {
            if (errno == EINTR) { /* select() indicates error          */
                continue;        /* select() interrupted by signal    */
            }                    /* Return to top of loop            */
            perror("select()");  /* select() really encountered error */
            exit(1);             /* End program                      */
        }

        /*
         * If an Event Manager response is present for the 1st EMAPI session,
         * process it. Also, maintain the count of registered events.
         */

        if (session_response_ready(&sess1, &read_fds)) {
            process_response(&sess1, progs1, progs1_elems, &reg_cnt);
        }

        /*
         * If an Event Manager response is present for the 2nd EMAPI session,
         */
    }
}
```

```

    * process it. Also, maintain the count of registered events.
    */
    if (session_response_ready(&sess2, &read_fds)) {
        process_response(&sess2, progs2, progs2_elems, &reg_cnt);
    }
}

end_session(&sess1);      /* End EMAPI session for 1st SP partition */
end_session(&sess2);     /* End EMAPI session for 2nd SP partition */

return 0;                /* Indicate program was successful */
}

/*
 * The interrupt_catch() function is installed as the signal handler for the
 * SIGINT signal. The SIGINT signal is delivered to this process when the
 * user presses the interrupt key (usually Ctrl-C). When this routine is
 * invoked, it sets the global variable terminate_requested to a non-zero
 * value, indicating the program is to be terminated.
 */
static
void interrupt_catch(int signo)
{
    terminate_requested = 1;
    return;
}

/*
 * The setup_signals() function initializes the disposition of a couple of
 * signals.
 *
 * The disposition of the SIGPIPE signal is set such that the signal is
 * ignored. When a process writes to a pipe or connection-oriented socket for
 * which there is no reader, the SIGPIPE signal is delivered to the process.
 * The default disposition for SIGPIPE is to kill the receiving process. If
 * SIGPIPE is ignored, the SIGPIPE signal does not kill the process, and the
 * system call used to write to the pipe or connection-oriented socket sets
 * errno to EPIPE and returns an error indication. It is recommended that
 * clients of the EMAPI ignore the SIGPIPE signal, since the EMAPI uses
 * connection-oriented sockets to communicate with the Event Manager.
 *
 * The interrupt_catch() function is installed as the signal handler for the
 * SIGINT signal. See the comments pertaining to the interrupt_catch()
 * function to find out why this is done.
 */
static
void setup_signals(void)
{
    struct sigaction sa;

    sa.sa_handler = SIG_IGN;          /* SIGPIPE to be ignored */
    sigemptyset(&sa.sa_mask);        /* No signals to mask off */
    sa.sa_flags = 0;                 /* No flags needed */

    if (sigaction(SIGPIPE, &sa, NULL) == -1) { /* Change SIGPIPE disposition*/
        perror("sigaction(SIGPIPE)");
        exit(1);
    }

    sa.sa_handler = interrupt_catch; /* SIGINT signal handler */
    sigemptyset(&sa.sa_mask);        /* No signals to mask off */
    sa.sa_flags = SA_RESTART;        /* Restart interrupted syscalls */

    if (sigaction(SIGINT, &sa, NULL) == -1) { /* Change SIGINT disposition */
        perror("sigaction(SIGINT)");
        exit(1);
    }

    return;
}

/*
 * The start_session() function establishes a session with the EMAPI by
 * calling the EMAPI routine ha_em_start_session(). If a session cannot be
 * established, the program is terminated.
 */
static
void start_session(struct session *sess_p)
{
    struct ha_em_err_blk errb;        /* EMAPI error block */

    /*

```

## emapi\_v02\_ex02.c

```
* The ha_em_start_session() routine receives the name of the partition
* with which a session is to be established, and returns a file
* descriptor with which the session can be used.
*/

sess_p->fd = ha_em_start_session(HA_EM_DOMAIN_SP, sess_p->name, &errb);

/*
 * If ha_em_start_session() indicates a session could not be established,
 * print the error information returned by the routine, and exit the
 * program.
 */

if (sess_p->fd == -1) {
    fprintf(stderr, "%s:\tha_em_start_session() returned EM errno %d:\n%s",
        sess_p->name, errb.em_errno, errb.em_errmsg);
    exit(1);
}

/*
 * A session has been established with the EAPI, and the session's file
 * descriptor has been saved in the session structure passed to this
 * routine. Set the flag in the session structure that indicates the
 * session does not need to be restarted.
 */

sess_p->restart = 0;

return;
}

/*
 * The register_for_events() function uses the EAPI routine
 * ha_em_send_command() to request the registration of events pertaining to
 * programs of interest that are in the same SP partition.
 *
 * The resource variable name specified is "IBM.PSSP.Prog.xpcount". This is a
 * state variable whose value indicates whether or not processes are executing
 * a specific program. The resource ID for this resource variable
 * specifies the program name, the name of the user running the program, and
 * the node on which the program is running. The resource variable's value
 * is a structured byte string of three fields. Field 0 is of type long; its
 * value is the number of processes currently running the specified program
 * for the specified user on the specified node. Field 1 is also of type
 * long; its value is the previous number of processes running the program.
 * Field 2 is of type character string; its value is a comma separated list
 * of the PIDs of the processes running the program.
 *
 * When the events are registered, two expressions are specified. The primary
 * expression, "X@0 == 0", specifies that an event is to be generated when
 * field 0 of the resource variable's structured byte string has a value of 0.
 * Taking into account the semantics of the resource variable, the expression
 * causes an event to be generated when no processes are running the specified
 * program for the specified user on the specified node. The re-arm
 * expression, "X@0 > 0", specifies that after the primary expression has
 * generated an event, the primary expression is to be re-armed once a process
 * starts running the specified program for the specified user on the
 * specified node. Since the HA_EM_CMD_REG2 command is specified, the re-arm
 * expression will generate events.
 *
 * Notice that when events are registered, the HA_EM_SCMD_REVAL subcommand is
 * specified. As a result, the initial value of the programs being monitored
 * will be returned.
 */

static
void register_for_events(struct session *sess_p,
    struct program *progs_p, int progs_elems, int *reg_cnt_p)
{
    int          num_events;      /* Number of events to register */
    size_t       alloc_size;     /* Size of memory to allocate */
    struct ha_em_cmd_blk *cmd_blk; /* Pointer to command block */
    struct ha_em_rb_reg *re;      /* Pointer to registered event */
    int          i;              /* Index */
    struct ha_em_err_blk errb;    /* EAPI error block */
    int          rc;             /* Return code */
    char         rsrc_ID_buf[80]; /* Resource ID buffer */
    struct callback_data *cd_p;   /* Pointer to callback data */

    num_events = progs_elems;     /* Register one event per program */

    /*
     * Allocate enough memory to hold the command block to be given to the
     * EAPI. The ha_em_cmd_blk structure contains one ha_em_rb_reg
     * structure. Additional space must be allocated when more than one
     * event is to be registered.
     */

    alloc_size = sizeof(struct ha_em_cmd_blk) +
        ((num_events - 1) * sizeof(struct ha_em_rb_reg));
}
```

```

if ((cmd_blk = malloc(alloc_size)) == NULL) {
    perror("malloc()");
    exit(1);
}

/*
 * Fill in the em_rb_reg array entries. There is one entry per event to
 * be registered.
 */
for (i = 0; i < num_events; i++) {
    re = &cmd_blk->em_res_blk.em_rb_reg[i];

    re->em_name = "IBM.PSSP.Prog.xpcount"; /* Resource variable name */

    sprintf(rsrc_ID_buf, "ProgName=%s;UserName=%s;NodeNum=%d",
            progs_p[i].name, progs_p[i].user, progs_p[i].node);
    re->em_rsrc_ID = strdup(rsrc_ID_buf); /* Resource ID */

    if (re->em_rsrc_ID == NULL) {
        perror("strdup()");
        exit(1);
    }

    re->em_expr = "X00 == 0"; /* Expression */
    re->em_raexpr = "X00 > 0"; /* Re-arm expression */

    if ((cd_p = malloc(sizeof(struct callback_data))) == NULL) {
        perror("malloc()");
        exit(1);
    }

    cd_p->sess_p = sess_p; /* Callback gets session pointer */
    cd_p->prog_p = &progs_p[i]; /* Callback gets program pointer */
    cd_p->reg_cnt_p = reg_cnt_p; /* Callback gets reg. count pointer */

    re->em_cb = event_callback; /* Callback routine */
    re->em_cb_arg = cd_p; /* Callback routine args. */

    progs_p[i].cb_data = cd_p; /* Save pointer to callback data */
}

/*
 * Fill in command block header, specifying number of elements in
 * em_rb_reg array, command, and subcommand.
 */
cmd_blk->em_cmd_num_elem = num_events;
cmd_blk->em_cmd = HA_EM_CMD_REG2;
cmd_blk->em_subcmd = HA_EM_SCMD_REVAL;

/*
 * Send the command. A more elaborate program might check for the
 * HA_EM_ECONNLOST Event Manager error number when an error is
 * returned, and attempt to restart the session. But, this program
 * just terminates if ha_em_send_command() detects a lost connection.
 */
rc = ha_em_send_command(sess_p->fd, cmd_blk, &errb);

if (rc == -1) {
    fprintf(stderr, "%s:\tha_em_send_command() returned EM errno %d:\n%s",
            sess_p->name, errb.em_errno, errb.em_errmsg);
    exit(1);
}

/*
 * Save the event identifiers assigned to the events just registered.
 */
for (i = 0; i < num_events; i++) {
    re = &cmd_blk->em_res_blk.em_rb_reg[i];
    progs_p[i].eid = re->em_event_id;
    progs_p[i].unreged = 0;
    free(re->em_rsrc_ID);
}

free(cmd_blk);

*reg_cnt_p += num_events; /* Update count of registered events */

return;
}

/*
 * The unregister_events() function uses the EMAPI routine
 * ha_em_send_command() to request the unregistration of events pertaining to
 * programs of interest that are in the same SP partition.
 */

```

## emapi\_v02\_ex02.c

```
static
void unregister_events(struct session *sess_p,
                      struct program *progs_p, int progs_elems)
{
    int                num_events;    /* Number of events to unreg. */
    size_t            alloc_size;    /* Size of memory to allocate */
    struct ha_em_cmd_blk *cmd_blk;    /* Pointer to command block */
    ha_em_eid_t       *re;           /* Pointer to unreg. event */
    int               i;             /* Index */
    struct ha_em_err_blk errb;       /* EMAPI error block */
    int               rc;           /* Return code */

    /*
     * Allocate enough memory to hold the command block to be given to the
     * EMAPI. The ha_em_cmd_blk structure contains one ha_em_eid_t
     * element. Additional space must be allocated when more than one
     * event is to be unregistered.
     */

    alloc_size = sizeof(struct ha_em_cmd_blk) +
                 ((progs_elems - 1) * sizeof(ha_em_eid_t));

    if ((cmd_blk = malloc(alloc_size)) == NULL) {
        perror("malloc()");
        exit(1);
    }

    /*
     * Fill in the event identifier array entries. There is one entry per
     * event to be unregistered. Only unregister events that are not already
     * unregistered.
     */

    num_events = 0;

    for (i = 0; i < progs_elems; i++) {
        if (!progs_p[i].unreged) {
            re = &cmd_blk->em_res_blk.em_rb_unreg[num_events];
            *re = progs_p[i].eid;
            num_events++;
        }
    }

    /*
     * Fill in command block header, specifying the number of elements in
     * the event identifier array, and the unregister command.
     */

    cmd_blk->em_cmd_num_elem = num_events;
    cmd_blk->em_cmd          = HA_EM_CMD_UNREG;

    /*
     * Send the unregister command. A more elaborate program might check for
     * the HA_EM_ECONNLOST Event Manager error number when an error is
     * returned, and attempt to restart the session. But, this program
     * just terminates if ha_em_send_command() detects a lost connection.
     */

    rc = ha_em_send_command(sess_p->fd, cmd_blk, &errb);

    if (rc == -1) {
        fprintf(stderr, "%s:\tha_em_send_command() returned EM errno %d:\n%s",
                sess_p->name, errb.em_errno, errb.em_errmsg);
        exit(1);
    }

    free(cmd_blk);

    return;
}

/*
 * The select_session() routine determines how the specified session affects
 * the next call to select(). If the session is marked as needing to be
 * restarted, ha_em_restart_session() is called to try to restart the session.
 * If the restart attempt fails, the next call to select() should specify
 * a timeout period - so the restart can be attempted again. If the restart
 * attempt succeeds, the new session file descriptor is saved, and the session
 * is marked as not needing to be restarted. If the session did not need
 * to be restarted, or had been successfully restarted, the next call to
 * select() should examine the session's file descriptor.
 */

static
void select_session(struct session *sess_p,
                   fd_set *read_fds_p, int *fd_limit_p, int *timeout_p)
{
    struct ha_em_err_blk errb;       /* EMAPI error block */
    int                 new_fd;     /* New session file descriptor */
}
```

```

/*
 * If the specified session needs to be restarted, try to do so.
 */
if (sess_p->restart) {
    new_fd = ha_em_restart_session(sess_p->fd, &errb);
    if (new_fd == -1) {
        /*
         * The session could not be restarted. If the Event Manager
         * error number indicates connection refused, indicate the
         * next select() should specify a timeout value (so the restart
         * can be tried again later), and return without putting the
         * session's file descriptor in the select() file descriptor
         * mask.
         */
        if (errb.em_errno == HA_EM_ECONNREFUSED) {
            *timeout_p = 1;
            return;
        }
        /*
         * If execution reaches here, the session restart attempt failed
         * for an unexpected reason; terminate the program.
         */
        fprintf(stderr, "%s:\tha_em_restart_session() returned EM errno "
            "%d:\n%s", sess_p->name, errb.em_errno, errb.em_errmsg);
        exit(1);
    }
    /*
     * If execution reaches here, a session restart was attempted and
     * was successful. Save the new session file descriptor, and indicate
     * that the session no longer needs to be restarted.
     */
    sess_p->fd = new_fd;
    sess_p->restart = 0;
}

/*
 * If execution reaches here, the specified session is not known to be
 * in need of a restart. Put the session's file descriptor in the
 * select() file descriptor mask. Also, adjust the maximum file
 * descriptor to select, if appropriate.
 */
FD_SET(sess_p->fd, read_fds_p);
*fd_limit_p = MAX(sess_p->fd, *fd_limit_p);
return;
}

/*
 * The session_response_ready() function determines if a select() file
 * descriptor mask indicates that a session's file descriptor is ready for
 * reading.
 */
static
int session_response_ready(struct session *sess_p, fd_set *read_fds_p)
{
    return FD_ISSET(sess_p->fd, read_fds_p);
}

/*
 * The process_response() function processes an Event Manager response
 * in the specified EAPI session. First, ha_em_receive_response() is called
 * to receive the response. If ha_em_receive_response() indicates that the
 * session's connection with the Event Manager has been lost, the session
 * is marked as being in need of a restart. Any other error will terminate
 * the program. The ha_em_receive_response() function may indicate that
 * there really isn't any response for the EAPI client to process at this
 * time. In that case, this function just returns. If a response has
 * been received, it is processed depending on what type of response it is.
 */
static
void process_response(struct session *sess_p,
    struct program *progs_p, int progs_elems, int *reg_cnt_p)
{
    struct ha_em_rsp_blk *rsp_blk; /* Pointer to the response block */
    int rc; /* Return code */
}

```

## emapi\_v02\_ex02.c

```
struct ha_em_err_blk  errb;    /* Event Manager error block */
int                  i;        /* Index */

/*
 * Receive response from session, if there is one to receive.
 */

rc = ha_em_receive_response(sess_p->fd, &rsp_blk, &errb);

if (rc == -1) {

    /*
     * If the connection with the Event Manager has been lost,
     * mark the session as needing to be restarted, and return.
     */

    if (errb.em_errno == HA_EM_ECONNLOST) {
        printf("\tconnection to session lost.\n", sess_p->name);
        sess_p->restart = 1;
        return;
    }

    /*
     * Some error has occurred other than the loss of the connection.
     * Terminate the program.
     */

    fprintf(stderr, "%s:\t ha_em_receive_response() returned EM errno %d:\n"
        "%s", sess_p->name, errb.em_errno, errb.em_errmsg);
    exit(1);

}

/*
 * If the ha_em_receive_response() routine returned zero,
 * there is no response for the EMAPI client (this program) to process
 * at this time. The response may have been for the EMAPI itself,
 * a full response may not be available yet, or the response may have
 * been handled with the callback routine. Just return.
 */

if (rc == 0) {
    return;
}

/*
 * A response for the EMAPI client (this program) has been returned.
 * The response buffer is pointed to by the rsp_blk variable. Appropriate
 * processing for the response depends on the command type. Note that
 * since all event registrations in this program specify the use of a
 * callback routine, event responses and event unregistration responses
 * are not expected by this function.
 */

switch (rsp_blk->em_cmd) {

    case HA_EM_CMD_RERR:          /* REG event registration error */
    case HA_EM_CMD_R2ERR:        /* REG2 event registration error */
        process_response_rerr(sess_p, progs_p, progs_elems, rsp_blk,
            reg_cnt_p);
        break;

    default:                      /* Unexpected response */
        fprintf(stderr, "%s:\t Program received unexpected command "
            "response: %d.\n", sess_p->name, rsp_blk->em_cmd);
        exit(1);
        break;

}

/*
 * The EMAPI client (this program) must free the memory associated with
 * the returned response block when it is no longer needed.
 */

free(rsp_blk);

return;
}

/*
 * The process_response_rerr() function processes registration error
 * responses. When an event cannot be registered due to some detected error,
 * a registration error response is sent to the EMAPI client. This function
 * marks the event, as tracked in a program structure, as being unregistered.
 * This will prevent an attempt to unregister a non-registered event later
 * in the program.
 */

static
```

```

void process_response_rerr(struct session *sess_p,
                          struct program *progs_p, int progs_elems,
                          struct ha_em_rsp_blk *rsp_blk, int *reg_cnt_p)
{
    struct ha_em_rpb_rerr *error_p;          /* Pointer to error response */
    struct ha_em_rpb_rerr *last_error_p;     /* Beyond last error response*/
    struct program *prog_p;                 /* Pointer to program struct */

    char prog_name[80];                     /* Name of program to which error pertains */
    char user_name[80];                     /* Name of user to which error pertains */
    int node;                               /* Node to which error pertains */

    error_p = rsp_blk->em_resp_blk.em_rpb_rerr;
    last_error_p = error_p + rsp_blk->em_rsp_num_resp;

    for ( ; error_p < last_error_p; error_p++) {
        /*
         * Find the program structure which describes the program associated
         * with the event whose registration failed. The match is made
         * with event identifiers.
         */
        for (prog_p = progs_p; prog_p < progs_p + progs_elems; prog_p++) {
            if (prog_p->eid == error_p->em_event_id) {
                break;
            }
        }
        if (prog_p == progs_p + progs_elems) {
            fprintf(stderr, "%s:\tUnknown event identifier encountered (0x%x)."\n",
                    sess_p->name, error_p->em_event_id);
            exit(1);
        }
        /*
         * Extract the program name, user name, and node number from the
         * resource ID.
         */
        breakdown_rsrc_ID(error_p->em_rsrc_ID, prog_name, user_name, &node);

        /*
         * Print a message about the registration error.
         */
        printf("%s:\tRegistration error for %s run by %s on node %d (%d, %d)."\n",
              sess_p->name, prog_name, user_name, node,
              error_p->em_generr, error_p->em_specerr);

        /*
         * Mark the program as not having an associated event registered.
         */
        prog_p->unreged = 1;

        /*
         * Free the callback data allocated for this program.
         */
        free(prog_p->cb_data);
        prog_p->cb_data = NULL;
    }

    /*
     * Update the number of events that are registered.
     */
    *reg_cnt_p -= rsp_blk->em_rsp_num_resp;
}

return;

/*
 * The event_callback() function is called whenever an event response has been
 * received from the Event Manager. Several points should be kept in mind:
 *
 * - The event response may indicate that an unregistration request has
 *   been completed. The HA_EM_EVENT_UNREG flag indicates the event
 *   response is actually an unregistration response.
 *
 * - An event response may be an error response. An event error response
 *   does not indicate that a registered event has occurred. Instead, it
 *   indicates the Event Manager no longer knows the current value of
 *   the associated resource variable, and cannot generate events for the
 *   resource variable. This may be a temporary condition. If the
 *   Event Manager later obtains the current value of the associated
 *   resource variable, an event will be generated (possibly indicating
 *   that the expression is false).
 */

```

## emapi\_v02\_ex02.c

```
*
* - The event for which an event error response is generated is still
*   registered. There is no need to re-register the event.
*
* - Since events are registered by this program using the HA_EM_CMD_REG2
*   command, the re-arm expressions will generate events. Therefore, the
*   HA_EM_EVENT_RE_ARM flag must be tested to see if the expression or
*   the re-arm expression generated the event.
*
* - It is possible for an event to be delivered indicating the expression
*   is false. This can happen for two reasons. First, when events are
*   registered by this program, the HA_EM_SCMD_REVAL subcommand is
*   specified. That subcommand requests the initial value of the
*   associated resource variable. Second, the current value of the
*   associated resource variable is returned through an event once the
*   Event Manager obtains the current value of a resource variable after
*   an error has occurred.
*
* - If an event response does not have the HA_EM_EVENT_RE_ARM nor
*   HA_EM_EVENT_EXPR_FALSE flags set, the event response indicates that
*   the event's expression is true.
*/

static
void event_callback(int sess_fd, struct ha_em_rpb_event *event_p, void *arg)
{
    struct callback_data *cd_p; /* Callback routine data pointer */
    struct session *sess_p; /* Session structure pointer */
    struct program *prog_p; /* Program structure pointer */
    int *reg_cnt_p; /* Registration counter pointer */

    char time_stamp[80]; /* Formatted event time stamp */

    /*
     * Get the callback routine's arguments.
     */

    cd_p = (struct callback_data *)arg;

    sess_p = cd_p->sess_p;
    prog_p = cd_p->prog_p;
    reg_cnt_p = cd_p->reg_cnt_p;

    /*
     * Format the time stamp.
     */

    format_timestamp(&event_p->em_timestamp, time_stamp);

    if (event_p->em_event_flags & HA_EM_EVENT_UNREG) {

        /*
         * If execution reaches this point, the event indicates an event
         * unregistration has completed.
         * Print a message indicating the program associated with the
         * unregistered event is no longer being monitored.
         */

        printf("%s:\t%s no longer monitoring %s run by %s on node %d.\n",
            sess_p->name, time_stamp, prog_p->name, prog_p->user,
            prog_p->node);

        /*
         * Note: em_errnum could indicate an error, but there is no point
         * looking at it here, since the program would not do anything
         * differently.
         */

        /*
         * Mark the program as not having an associated event registered.
         */

        prog_p->unreged = 1;

        /*
         * Indicate this program has one fewer registered event.
         */

        (*reg_cnt_p)--;

        /*
         * The callback routine will no longer be called about this event,
         * so free the data that was allocated specifically for this event.
         */

        free(prog_p->cb_data);
        prog_p->cb_data = NULL;

        return;
    }
}
```

```

/*
 * If execution reaches this point, the event response is not an
 * unregistration response.
 */

/*
 * If this is an event error response, print a message which includes
 * the general and specific error codes.
 */

if (event_p->em_errnum != 0) {
    printf("%s:\t%s State of %s run by %s on node %d unknown (%d, %d).\n",
        sess_p->name, time_stamp, prog_p->name, prog_p->user,
        prog_p->node, event_p->em_generr, event_p->em_specerr);
    return;
}

/*
 * If the event response was generated for the re-arm expression, or
 * the event response was generated for the primary expression but
 * it is false, print a message indicating the program associated
 * with the event is being run.
 */

if ((event_p->em_event_flags & HA_EM_EVENT_RE_ARM) ||
    (event_p->em_event_flags & HA_EM_EVENT_EXPR FALSE)) {
    printf("%s:\t%s %s being run by %s on node %d.\n",
        sess_p->name, time_stamp, prog_p->name, prog_p->user,
        prog_p->node);
    return;
}

/*
 * If execution reaches this point, the event response was generated
 * for the primary expression, and it is true. Print a message
 * indicating the program associated with the event is not being run.
 */

printf("%s:\t%s %s NOT being run by %s on node %d.\n",
    sess_p->name, time_stamp, prog_p->name, prog_p->user,
    prog_p->node);

return;
}

/*
 * The format_timestamp() function takes a time stamp returned by the
 * Event Manager and converts it into 24-hour time.
 */

static
void format_timestamp(struct timeval *timestamp_p, char *fmt_timestamp_p)
{
    struct tm    *broken_down_time;

    broken_down_time = localtime((time_t *) &timestamp_p->tv_sec);
    (void) strftime(fmt_timestamp_p, 20, "(%X)", broken_down_time);

    return;
}

/*
 * The breakdown_rsrc_ID() function takes a resource ID
 * associated with the IBM.PSSP.Prog.xpcount resource variable and extracts
 * from it the program name, user name, and node number.
 */

static
void breakdown_rsrc_ID(char *rsrc_ID_p,
    char *prog_name_p, char *user_name_p, int *node_p)
{
    char    *value_p;
    size_t  value_len;

    find_rsrc_ID_value(rsrc_ID_p, "ProgName=", &value_p, &value_len);
    strncpy(prog_name_p, value_p, value_len);
    *(prog_name_p + value_len) = '\0';

    find_rsrc_ID_value(rsrc_ID_p, "UserName=", &value_p, &value_len);
    strncpy(user_name_p, value_p, value_len);
    *(user_name_p + value_len) = '\0';

    find_rsrc_ID_value(rsrc_ID_p, "NodeNum=", &value_p, &value_len);
    *node_p = atoi(value_p);

    return;
}

```

## emapi\_v02\_ex02.c

```
/*
 * The find_rsrc_ID_value() takes a resource ID and extracts
 * a value from it.
 */
static
void find_rsrc_ID_value(char *rsrc_ID_p, char *name_p,
                       char **value_pp, size_t *value_len_p)
{
    char *bp;
    char *ep;
    size_t len;

    if ((bp = strstr(rsrc_ID_p, name_p)) == NULL) {
        fprintf(stderr, "Unexpected resource ID encountered.\n");
        exit(1);
    }

    bp += strlen(name_p);

    if ((ep = strchr(bp, ';')) != NULL) {
        len = ep - bp;
    } else {
        len = strlen(bp);
    }

    *value_pp = bp;
    *value_len_p = len;

    return;
}

/*
 * The end_session() function terminates a session with the EAPI by
 * calling the EAPI routine ha_em_end_session().
 */
static
void end_session(struct session *sess_p)
{
    struct ha_em_err_blk errb;

    if (ha_em_end_session(sess_p->fd, &errb) == -1) {
        fprintf(stderr, "%s:\tha_em_end_session() returned EM errno %d:\n%s",
                sess_p->name, errb.em_errno, errb.em_errmsg);
        exit(1);
    }

    return;
}
```

---

## The emapi\_v02\_ex03.c Sample Program

```
/* IBM PROLOG_BEGIN_TAG */
/* This is an automatically generated prolog. */
/* */
/* */
/* Licensed Materials - Property of IBM */
/* */
/* (C) COPYRIGHT International Business Machines Corp. 1996,1998 */
/* All Rights Reserved */
/* */
/* US Government Users Restricted Rights - Use, duplication or */
/* disclosure restricted by GSA ADP Schedule Contract with IBM Corp. */
/* */
/* IBM_PROLOG_END_TAG */

/* @(#)45 1.3 src/rsct/pem/emtools/emapi_test/emapi_ex_r/emapi_v02_ex03.c, emtools, rsct_rtrto 6/30/98 10:48:47 */

#include <pthread.h>
#include <unistd.h>
#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <errno.h>
#include <signal.h>
#include <time.h>

#define HA_EM_VERSION 2
#include <ha_emapi.h>

/*
 * emapi_v02_ex03.c
 *
 * This program presents an example of using the Event Manager Application
```

```

* Programming Interface (EMAPI). The program uses the EMAPI to monitor
* four programs running on various nodes in two SP partitions. Since two
* SP partitions are involved, the program establishes two sessions with the
* EMAPI, one per SP partition. Each session is managed by a separate thread.
*
* This program can be compiled with the following command:
*
*   cc_r -o emapi_v02_ex03.c -o emapi_v02_ex03 -lha_em_r -lpthreads -lc_r
*/

/*
* The following macros specify the programs this program will monitor through
* the EMAPI.
*/

#define PART_1      "k21s"           /* 1st SP partition           */
#define PROG_1A    "subsys_pgrm1"   /* 1st program in 1st SP partition, */
#define USER_1A    "root"           /* ... and user running 1st program, */
#define NODE_1A    5                /* ... and node running 1st program. */

#define PROG_1B    "subsys_pgrm2"   /* 2nd program in 1st SP partition, */
#define USER_1B    "root"           /* ... and user running 2nd program, */
#define NODE_1B    6                /* ... and node running 2nd program. */

#define PART_2      "k21sp2"        /* 2nd SP partition           */
#define PROG_2A    "subsys_pgrm1"   /* 1st program in 2nd SP partition, */
#define USER_2A    "root"           /* ... and user running 1st program, */
#define NODE_2A    14               /* ... and node running 1st program. */

#define PROG_2B    "subsys_pgrm2"   /* 2nd program in 2nd SP partition, */
#define USER_2B    "root"           /* ... and user running 2nd program, */
#define NODE_2B    15               /* ... and node running 2nd program. */

/*
* The STRERROR_BUF and STRERROR macros make it a little more convenient to
* use the thread safe version of strerror(), strerror_r().
*/

#define STRERROR_BUF  char strerror_buf[256]

#define STRERROR(en) ( \
    strerror_r(en, strerror_buf, sizeof strerror_buf) == 0 \
    ? strerror_buf : "Unknown error" \
)

/*
* This program saves information about each EMAPI session it has established
* in a session structure.
*/

struct session {
    char    *name;      /* Name of partition used by session */
    int     fd;        /* Session file descriptor           */
    int     restart;   /* Boolean - Session's connection has been lost */
};

/*
* This program saves information about each program it is monitoring through
* the EMAPI in a program structure.
*/

struct program {
    char    *name;      /* Name of a program to be monitored */
    char    *user;      /* Name of user running program      */
    int     node;       /* Node on which program is running   */
    ha_em_eid_t eid;    /* Event identifier for program       */
    int     unreged;    /* Event is unregistered              */
};

/*
* The thread_data structure is used by the initial thread to pass data to
* a secondary thread when it is created.
*/

struct thread_data {
    struct session *sess; /* EMAPI session the thread is to establish */
    struct program *progs; /* Programs the thread is to monitor */
    int     progs_cnt; /* Number of programs thread is to monitor */
};

/*
* Following are handles for threads created by the initial thread.
*/

pthread_t  thread1;    /* Handle to thread 1 - a program monitor */
pthread_t  thread2;    /* Handle to thread 2 - a program monitor */
pthread_t  thread3;    /* Handle to thread 3 - cancellation monitor */

/*
* Function prototypes for internal functions.

```

## emapi\_v02\_ex03.c

```
*/

static void setup_signals(void);
static void create_thread(pthread_t *thread_p, void *(* thread_rtn)(void *),
                          void *data_p);
static void cancel_thread(pthread_t thread);
static void wait_for_thread(pthread_t thread);
static void *program_thread_main(void *parm_p);
static void program_thread_main_cleanup(void *parm_p);
static void *cancel_thread_main(void *parm_p);
static void start_session(struct session *sess_p);
static void register_for_events(struct session *sess_p,
                               struct program *progs_p, int progs_elems, int *reg_cnt_p);
static void ensure_session_connected(struct session *sess_p);
static void process_response(struct session *sess_p,
                            struct program *progs_p, int progs_elems, int *reg_cnt_p);
static void process_response_reg(struct session *sess_p,
                                struct program *progs_p, int progs_elems,
                                struct ha_em_rsp_blk *rsp_blk);
static void process_response_rerr(struct session *sess_p,
                                 struct program *progs_p, int progs_elems,
                                 struct ha_em_rsp_blk *rsp_blk, int *reg_cnt_p);
static void format_timestamp(struct timeval *timestamp_p,
                             char *fmt_timestamp_p);
static void breakdown_rsrc_ID(char *rsrc_ID_p,
                              char *prog_name_p, char *user_name_p, int *node_p);
static void find_rsrc_ID_value(char *rsrc_ID_p, char *name_p,
                              char **value_pp, size_t *value_len_p);
static void end_session(struct session *sess_p);

/*
 * The main() function is executed by this program's initial thread.
 * The function creates three other threads. The first two threads
 * use the EMAPI to monitor other programs. One thread is used per
 * EMAPI session, and one EMAPI session is used per SP partition.
 * The third thread waits for the user to request the termination of this
 * program. When the third thread detects the user would like this program
 * to terminate, it cancels the two threads using the EMAPI. When those
 * two threads terminate, the initial thread cancels the third thread, waits
 * for the third thread to terminate, and terminates the program.
 */

int main(int argc, char **argv)
{
    struct session sess1 = {PART_1, -1, 0};
    struct session sess2 = {PART_2, -1, 0};

    struct program progs1[] = {{PROG_1A, USER_1A, NODE_1A, 0, 0},
                              {PROG_1B, USER_1B, NODE_1B, 0, 0}};
    struct program progs2[] = {{PROG_2A, USER_2A, NODE_2A, 0, 0},
                              {PROG_2B, USER_2B, NODE_2B, 0, 0}};

    int    progs1_elems = sizeof progs1 / sizeof progs1[0];
    int    progs2_elems = sizeof progs2 / sizeof progs2[0];

    struct thread_data td1 = {&sess1, progs1, progs1_elems};
    struct thread_data td2 = {&sess2, progs2, progs2_elems};

    setup_signals();                /* Initialize signal dispositions */

    /*
     * Create 2 threads that will monitor programs with the EMAPI.
     */
    create_thread(&thread1, program_thread_main, &td1); /* Create thread 1 */
    create_thread(&thread2, program_thread_main, &td2); /* Create thread 2 */

    /*
     * Create a thread that will wait for the user to request termination.
     */
    create_thread(&thread3, cancel_thread_main, NULL); /* Create thread 3 */

    /*
     * Wait for the 2 threads that are monitoring programs to end. They
     * will end due to errors or they will be canceled by thread 3 when
     * the user requests termination.
     */
    wait_for_thread(thread1);        /* Wait for thread 1 to terminate */
    wait_for_thread(thread2);        /* Wait for thread 2 to terminate */

    /*
     * Cancel the thread waiting for the user to request termination, and
     * then wait for that thread to terminate.
     */
    cancel_thread(thread3);          /* Cancel thread 3 */
    wait_for_thread(thread3);        /* Wait for thread 3 to terminate */
}
```

```

    return 0;                /* Indicate program was successful */
}

/*
 * The setup_signals() function initializes the disposition of a couple of
 * signals.
 *
 * The disposition of the SIGPIPE signal is set such that the signal is
 * ignored. When a process writes to a pipe or connection-oriented socket for
 * which there is no reader, the SIGPIPE signal is delivered to the process.
 * The default disposition for SIGPIPE is to kill the receiving process. If
 * SIGPIPE is ignored, the SIGPIPE signal does not kill the process, and the
 * system call used to write to the pipe or connection-oriented socket sets
 * errno to EPIPE and returns an error indication. It is recommended that
 * clients of the EMAPI ignore the SIGPIPE signal, since the EMAPI uses
 * connection-oriented sockets to communicate with the Event Manager.
 *
 * The SIGINT signal is blocked in the initial thread's signal mask.
 * Threads created by the initial thread will inherit the signal mask from
 * the initial thread. One thread will wait for the SIGINT signal
 * to be delivered by using sigwait().
 */
static
void setup_signals(void)
{
    struct sigaction sa;          /* Signal action description */
    int rc;                      /* Return code (error number) */
    STRError_BUF;               /* Buffer for strerror_r() */
    sigset_t sigset;            /* Signal mask to block SIGINT */

    sa.sa_handler = SIG_IGN;      /* SIGPIPE to be ignored */
    sigemptyset(&sa.sa_mask);    /* No signals to mask off */
    sa.sa_flags = 0;             /* No flags needed */

    if (sigaction(SIGPIPE, &sa, NULL) == -1) { /* Change SIGPIPE disposition*/
        perror("sigaction(SIGPIPE)");
        exit(1);
    }

    sigemptyset(&sigset);        /* Initialize signal set */
    sigaddset(&sigset, SIGINT);  /* Add SIGINT to signal set */

    rc = sigthreadmask(SIG_BLOCK, &sigset, NULL); /* Block SIGINT signal */
    if (rc != 0) {
        fprintf(stderr, "sigthreadmask(): %s\n", STRError(rc));
        exit(1);
    }

    return;
}

/*
 * The create_thread() function creates a thread, passing along the specified
 * main function address and parameter address. The function returns the
 * thread's handle. The thread is created such that it can be joined later
 * on.
 */
static
void create_thread(pthread_t *thread_p, void *(* thread_rtn)(void *),
                  void *data_p)
{
    pthread_attr_t thread_attr;  /* Thread attributes */
    int rc;                      /* Return code (error number) */
    STRError_BUF;               /* Buffer for strerror_r() */

    rc = pthread_attr_init(&thread_attr);
    if (rc != 0) {
        fprintf(stderr, "pthread_attr_init(): %s\n", STRError(rc));
        exit(1);
    }

    rc = pthread_attr_setdetachstate(&thread_attr, PTHREAD_CREATE_UNDETACHED);
    if (rc != 0) {
        fprintf(stderr, "pthread_attr_setdetachstate(): %s\n",
                STRError(rc));
        exit(1);
    }

    rc = pthread_create(thread_p, &thread_attr, thread_rtn, data_p);
    if (rc != 0) {
        fprintf(stderr, "pthread_create(): %s\n", STRError(rc));
        exit(1);
    }

    rc = pthread_attr_destroy(&thread_attr);
    if (rc != 0) {

```

## emapi\_v02\_ex03.c

```
        fprintf(stderr, "pthread_attr_destroy(): %s\n", STRERROR(rc));
        exit(1);
    }
    return;
}

/*
 * The cancel_thread() function requests the cancellation of the specified
 * thread.
 */
static
void cancel_thread(pthread_t thread)
{
    int rc;
    STRERROR_BUF; /* Return code (error number) */
                /* Buffer for strerror_r() */

    rc = pthread_cancel(thread);
    if (rc != 0) {
        fprintf(stderr, "pthread_cancel(): %s\n", STRERROR(rc));
        exit(1);
    }
    return;
}

/*
 * The wait_for_thread() function waits for the specified thread to
 * terminate.
 */
static
void wait_for_thread(pthread_t thread)
{
    int rc;
    STRERROR_BUF; /* Return code (error number) */
                /* Buffer for strerror_r() */

    rc = pthread_join(thread, NULL);
    if (rc != 0) {
        fprintf(stderr, "pthread_join(): %s\n", STRERROR(rc));
        exit(1);
    }
    return;
}

/*
 * The program_thread_main() function is the main function for threads that
 * are created to monitor programs in a SP partition. The function is told
 * what programs are to be monitored in what partition through the function
 * parameter, which is a pointer to a thread_data structure. The function
 * calls functions to establish an EMAPI session, register for events, wait
 * for process Event Manager responses, and cleanup when the thread is
 * canceled. The thread cleanup function that is installed ends the EMAPI
 * session that was established for the thread.
 */
static
void *program_thread_main(void *parm_p)
{
    struct thread_data *td; /* Thread data (parameters) */
    int reg_cnt; /* Count of registered events */
    int rc; /* Return code */

    td = (struct thread_data *) parm_p; /* Get parameters for the thread */
    start_session(td->sess); /* Start EMAPI session for SP partition */

    /*
     * Install a clean up handler that will end the session when the thread
     * is canceled.
     */
    pthread_cleanup_push(program_thread_main_cleanup, (void *)td->sess);

    /*
     * Register for events that will monitor the programs of interest.
     * Maintain count of registered events.
     */
    reg_cnt = 0;
    register_for_events(td->sess, td->progs, td->progs_cnt, &reg_cnt);

    /*
     * Continue looking for Event Manager responses as long as there are

```

```

    * registered events.
    */
while (reg_cnt > 0) {
    /*
     * Ensure the session is still connected to the Event Manager.
     */
    ensure_session_connected(td->sess);

    /*
     * Process the next response from the Event Manager.
     */
    process_response(td->sess, td->progs, td->progs_cnt, &reg_cnt);
}

/*
 * Pop and execute the thread cleanup handler installed by this routine.
 * Note that the cleanup handler will end the EMAPI session.
 */
pthread_cleanup_pop(1);
return NULL;          /* The thread is finished.          */
}

/*
 * The program_thread_main_cleanup() function is the thread cancellation
 * cleanup handler for the program_thread_main() function. When this
 * function is installed as a thread cancellation cleanup handler, the
 * thread has an EMAPI session. This function ends that session. The
 * function receives as input a pointer to the session structure describing
 * the session.
 */
static
void program_thread_main_cleanup(void *parm_p)
{
    struct session *sess_p;

    sess_p = (struct session *) parm_p;
    end_session(sess_p);
    printf("%s:\tno longer monitoring programs in this partition.\n",
        sess_p->name);
    return;
}

/*
 * The cancel_thread_main() function waits for the user to request
 * the termination of this program. The user requests program termination
 * by pressing the interrupt key (usually Ctrl-C). Pressing the interrupt
 * key will cause this program to receive the SIGINT signal. This function
 * waits for the SIGINT signal to be delivered. Once the SIGINT signal has
 * been delivered, this function cancels the two threads monitoring programs
 * through the EMAPI.
 */
static
void *cancel_thread_main(void *parm_p)
{
    int          rc;                /* Return code (error number) */
    sigset_t    sigset;            /* Signal mask to unblock SIGINT */
    int         sig;                /* Received signal */
    STRErrorBuf;                    /* Buffer for strerror_r() */

    sigemptyset(&sigset);          /* Initialize signal set */
    sigaddset(&sigset, SIGINT);    /* Add SIGINT to signal set */

    for ( ; ; ) {
        rc = sigwait(&sigset, &sig);
        if (rc != 0) {
            fprintf(stderr, "sigwait(): %s\n", STRERROR(rc));
            exit(1);
        }

        if (sig == SIGINT) {
            printf("Termination requested.\n");
            cancel_thread(thread1); /* Request cancellation of thread 1 */
            cancel_thread(thread2); /* Request cancellation of thread 2 */
            break;                  /* No longer wait for SIGINT */
        }
    }
}

```

## emapi\_v02\_ex03.c

```
        fprintf(stderr, "sigwait() reported unexpected signal: %d.\n", sig);
        exit(1);
    }
    return NULL;          /* The thread is finished      */
}

/*
 * The start_session() function establishes a session with the EAPI by
 * calling the EAPI routine ha_em_start_session(). If a session cannot be
 * established, the program is terminated.
 */
static
void start_session(struct session *sess_p)
{
    struct ha_em_err_blk errb;          /* EAPI error block      */

    /*
     * The ha_em_start_session() routine receives the name of the partition
     * with which a session is to be established, and returns a file
     * descriptor with which the session can be used.
     */
    sess_p->fd = ha_em_start_session(HA_EM_DOMAIN_SP, sess_p->name, &errb);

    /*
     * If ha_em_start_session() indicates a session could not be established,
     * print the error information returned by the routine, and exit the
     * program.
     */
    if (sess_p->fd == -1) {
        fprintf(stderr, "%s:\tha_em_start_session() returned EM errno %d:\n%s",
                sess_p->name, errb.em_errno, errb.em_errmsg);
        exit(1);
    }

    /*
     * A session has been established with the EAPI, and the session's file
     * descriptor has been saved in the session structure passed to this
     * routine. Set the flag in the session structure that indicates the
     * session does not need to be restarted.
     */
    sess_p->restart = 0;

    return;
}

/*
 * The register_for_events() function uses the EAPI routine
 * ha_em_send_command() to request the registration of events pertaining to
 * programs of interest that are in the same SP partition.
 *
 * The resource variable name specified is "IBM.PSSP.Prog.xpcount". This is a
 * state variable whose value indicates whether or not processes are executing
 * a specific program. The resource ID for this resource variable
 * specifies the program name, the name of the user running the program, and
 * the node on which the program is running. The resource variable's value
 * is a structured byte string of three fields. Field 0 is of type long; its
 * value is the number of processes currently running the specified program
 * for the specified user on the specified node. Field 1 is also of type
 * long; its value is the previous number of processes running the program.
 * Field 2 is of type character string; its value is a comma separated list
 * of the PIDs of the processes running the program.
 *
 * When the events are registered, two expressions are specified. The primary
 * expression, "X@0 == 0", specifies that an event is to be generated when
 * field 0 of the resource variable's structured byte string has a value of 0.
 * Taking into account the semantics of the resource variable, the expression
 * causes an event to be generated when no processes are running the specified
 * program for the specified user on the specified node. The re-arm
 * expression, "X@0 > 0", specifies that after the primary expression has
 * generated an event, the primary expression is to be re-armed once a process
 * starts running the specified program for the specified user on the
 * specified node. Since the HA_EM_CMD_REG2 command is specified, the re-arm
 * expression will generate events.
 *
 * Notice that when events are registered, the HA_EM_SCMD_REVAL subcommand is
 * specified. As a result, the initial value of the programs being monitored
 * will be returned.
 */
static
void register_for_events(struct session *sess_p,
                        struct program *progs_p, int progs_elems, int *reg_cnt_p)
```

```

{
    int                num_events;    /* Number of events to register */
    size_t            alloc_size;    /* Size of memory to allocate */
    struct ha_em_cmd_blk *cmd_blk;    /* Pointer to command block */
    struct ha_em_rb_reg *re;         /* Pointer to registered event */
    int               i;             /* Index */
    struct ha_em_err_blk errb;       /* EMAPI error block */
    int               rc;            /* Return code */
    char              rsrc_ID_buf[80]; /* Resource ID buffer */

    num_events = progs_elems;        /* Register one event per program */

    /*
     * Allocate enough memory to hold the command block to be given to the
     * EMAPI. The ha_em_cmd_blk structure contains one ha_em_rb_reg
     * structure. Additional space must be allocated when more than one
     * event is to be registered.
     */
    alloc_size = sizeof(struct ha_em_cmd_blk) +
        ((num_events - 1) * sizeof(struct ha_em_rb_reg));

    if ((cmd_blk = malloc(alloc_size)) == NULL) {
        perror("malloc()");
        exit(1);
    }

    /*
     * Fill in the em_rb_reg array entries. There is one entry per event to
     * be registered.
     */
    for (i = 0; i < num_events; i++) {
        re = &cmd_blk->em_res_blk.em_rb_reg[i];

        re->em_name = "IBM.PSSP.Prog.xpcount"; /* Resource variable name */

        sprintf(rsrc_ID_buf, "ProgName=%s;UserName=%s;NodeNum=%d",
            progs_p[i].name, progs_p[i].user, progs_p[i].node);
        re->em_rsrc_ID = strdup(rsrc_ID_buf); /* Resource ID */

        if (re->em_rsrc_ID == NULL) {
            perror("strdup()");
            exit(1);
        }

        re->em_expr = "X00 == 0"; /* Expression */
        re->em_raexpr = "X00 > 0"; /* Re-arm expression */

        re->em_cb = NULL; /* Callback routine not used*/
        re->em_cb_arg = NULL; /* Callback routine not used*/
    }

    /*
     * Fill in command block header, specifying number of elements in
     * em_rb_reg array, command, and subcommand.
     */
    cmd_blk->em_cmd_num_elem = num_events;
    cmd_blk->em_cmd = HA_EM_CMD_REG2;
    cmd_blk->em_subcmd = HA_EM_SCMD_REVAL;

    /*
     * Send the command. A more elaborate program might check for the
     * HA_EM_ECONNLOST Event Manager error number when an error is
     * returned, and attempt to restart the session. But, this program
     * just terminates if ha_em_send_command() detects a lost connection.
     */
    rc = ha_em_send_command(sess_p->fd, cmd_blk, &errb);

    if (rc == -1) {
        fprintf(stderr, "%s:\tha_em_send_command() returned EM errno %d:\n%s",
            sess_p->name, errb.em_errno, errb.em_errmsg);
        exit(1);
    }

    /*
     * Save the event identifiers assigned to the events just registered.
     */
    for (i = 0; i < num_events; i++) {
        re = &cmd_blk->em_res_blk.em_rb_reg[i];
        progs_p[i].eid = re->em_event_id;
        progs_p[i].unreged = 0;
        free(re->em_rsrc_ID);
    }

    free(cmd_blk);
}

```

## emapi\_v02\_ex03.c

```
    *reg_cnt_p += num_events;          /* Update count of registered events */
}

return;
}

/*
 * The ensure_session_connected() function ensures the specified session
 * is not known to be disconnected. If the session is marked as needing
 * to be restarted, the function does not return until the session has been
 * successfully restarted. Restart attempts are tried every 15 seconds.
 */

static
void ensure_session_connected(struct session *sess_p)
{
    struct ha_em_err_blk  errb;          /* EMAPI error block */
    int                   new_fd;       /* New session file descriptor */

    /*
     * If the specified session needs to be restarted, try to do so.
     */

    while (sess_p->restart) {
        new_fd = ha_em_restart_session(sess_p->fd, &errb);

        if (new_fd == -1) {
            /*
             * The session could not be restarted. If the Event Manager
             * error number indicates connection refused, wait 15 seconds
             * before trying the restart again.
             */

            if (errb.em_errno == HA_EM_ECONNREFUSED) {
                sleep(15);
                continue;
            }

            /*
             * If execution reaches here, the session restart attempt failed
             * for an unexpected reason; terminate the program.
             */

            fprintf(stderr, "%s:\tha_em_restart_session() returned EM errno "
                "%d:\n%s", sess_p->name, errb.em_errno, errb.em_errmsg);
            exit(1);
        }

        /*
         * If execution reaches here, a session restart was attempted and
         * was successful. Save the new session file descriptor, and indicate
         * that the session no longer needs to be restarted.
         */

        sess_p->fd = new_fd;
        sess_p->restart = 0;
    }

    /*
     * If execution reaches here, the specified session is not known to be
     * in need of a restart. Just return.
     */

    return;
}

/*
 * The process_response() function processes an Event Manager response
 * in the specified EMAPI session. First, ha_em_receive_response() is called
 * to receive the response. If ha_em_receive_response() indicates that the
 * session's connection with the Event Manager has been lost, the session
 * is marked as being in need of a restart. Any other error will terminate
 * the program. The ha_em_receive_response() function may indicate that
 * there really isn't any response for the EMAPI client to process at this
 * time. In that case, this function just returns. If a response has
 * been received, it is processed depending on what type of response it is.
 */

static
void process_response(struct session *sess_p,
                    struct program *progs_p, int progs_elems, int *reg_cnt_p)
{
    struct ha_em_rsp_blk  *rsp_blk;     /* Pointer to the response block */
    int                   rc;           /* Return code */
    struct ha_em_err_blk  errb;        /* Event Manager error block */
    int                   i;           /* Index */
}
```

```

/*
 * Receive response from session, if there is one to receive.
 */
rc = ha_em_receive_response(sess_p->fd, &rsp_blk, &errb);
if (rc == -1) {
    /*
     * If the connection with the Event Manager has been lost,
     * mark the session as needing to be restarted, and return.
     */
    if (errb.em_errno == HA_EM_ECONNLOST) {
        printf("%s:\tconnection to session lost.\n", sess_p->name);
        sess_p->restart = 1;
        return;
    }
    /*
     * Some error has occurred other than the loss of the connection.
     * Terminate the program.
     */
    fprintf(stderr, "%s:\tthe ha_em_receive_response() returned EM errno %d:\n"
        "%s", sess_p->name, errb.em_errno, errb.em_errmsg);
    exit(1);
}
/*
 * If the ha_em_receive_response() routine returned zero,
 * there is no response for the EMAPI client (this program) to process
 * at this time. The response may have been for the EMAPI itself, or
 * a full response may not be available yet. Just return.
 */
if (rc == 0) {
    return;
}
/*
 * A response for the EMAPI client (this program) has been returned.
 * The response buffer is pointed to by the rsp_blk variable. Appropriate
 * processing for the response depends on the command type.
 */
switch (rsp_blk->em_cmd) {
    case HA_EM_CMD_REG:          /* Event response */
    case HA_EM_CMD_REG2:        /* Event response (possible re-arm) */
        process_response_reg(sess_p, progs_p, progs_elems, rsp_blk);
        break;
    case HA_EM_CMD_RERR:        /* REG event registration error */
    case HA_EM_CMD_R2ERR:       /* REG2 event registration error */
        process_response_rerr(sess_p, progs_p, progs_elems, rsp_blk,
            reg_cnt_p);
        break;
    default:                    /* Unexpected response */
        fprintf(stderr, "%s:\tProgram received unexpected command "
            "response: %d.\n", sess_p->name, rsp_blk->em_cmd);
        exit(1);
        break;
}
/*
 * The EMAPI client (this program) must free the memory associated with
 * the returned response block when it is no longer needed.
 */
free(rsp_blk);
return;
}
/*
 * The process_response_reg() function processes event responses from the
 * Event Manager. Several points should be kept in mind:
 *
 * - The response block may contain multiple event responses. The
 *   number of responses included in the response block is given in
 *   the response block header.
 *
 * - An event response may be an error response. An event error response
 *   does not indicate that a registered event has occurred. Instead, it
 *   indicates the Event Manager no longer knows the current value of
 *   the associated resource variable, and cannot generate events for

```

## emapi\_v02\_ex03.c

```
* the resource variable. This may be a temporary condition. If the
* Event Manager later obtains the current value of the associated
* resource variable, an event will be generated (possibly indicating
* that the expression is false).
*
* - The event for which an event error response is generated is still
* registered. There is no need to re-register the event.
*
* - Since events are registered by this program using the HA_EM_CMD_REG2
* command, the re-arm expressions will generate events. Therefore, the
* HA_EM_EVENT_RE_ARM flag must be tested to see if the expression or
* the re-arm expression generated the event.
*
* - It is possible for an event to be delivered indicating the expression
* is false. This can happen for two reasons. First, when events are
* registered by this program, the HA_EM_SCMD_REVAL subcommand is
* specified. That subcommand requests the initial value of the
* associated resource variable. Second, the current value of the
* associated resource variable is returned through an event once the
* Event Manager obtains the current value of a resource variable after
* an error has occurred.
*
* - If an event response does not have the HA_EM_EVENT_RE_ARM nor
* HA_EM_EVENT_EXPR_FALSE flags set, the event response indicates that
* the event's expression is true.
*/

static
void process_response_reg(struct session *sess_p,
                        struct program *progs_p, int progs_elems,
                        struct ha_em_rsp_blk *rsp_blk)
{
    struct ha_em_rpb_event *event_p;          /* Pointer to event response */
    struct ha_em_rpb_event *last_event_p;     /* Beyond last event response*/

    char time_stamp[80];                      /* Formatted event time stamp */
    char prog_name[80];                       /* Name of program to which event pertains */
    char user_name[80];                       /* Name of user to which event pertains */
    int node;                                 /* Node to which event pertains */

    event_p = rsp_blk->em_resp_blk.em_rpb_event;
    last_event_p = event_p + rsp_blk->em_rsp_num_resp;

    for ( ; event_p < last_event_p; event_p++) {

        /*
         * Format the event's time stamp; extract the program name, user
         * name, and node number from the resource ID.
         */

        format_timestamp(&event_p->em_timestamp, time_stamp);
        breakdown_rsrc_ID(event_p->em_rsrc_ID, prog_name, user_name, &node);

        /*
         * If this is an event error response, print a message which includes
         * the general and specific error codes.
         */

        if (event_p->em_errnum != 0) {
            printf("%s:\t%s State of %s run by %s on node %d unknown (%d, %d)."\n",
                  sess_p->name, time_stamp, prog_name, user_name, node,
                  event_p->em_generr, event_p->em_specerr);
            continue;
        }

        /*
         * If the event response was generated for the re-arm expression, or
         * the event response was generated for the primary expression but
         * it is false, print a message indicating the program associated
         * with the event is being run.
         */

        if ((event_p->em_event_flags & HA_EM_EVENT_RE_ARM) ||
            (event_p->em_event_flags & HA_EM_EVENT_EXPR_FALSE)) {
            printf("%s:\t%s %s being run by %s on node %d.\n",
                  sess_p->name, time_stamp, prog_name, user_name, node);
            continue;
        }

        /*
         * If execution reaches this point, the event response was generated
         * for the primary expression, and it is true. Print a message
         * indicating the program associated with the event is not being run.
         */

        printf("%s:\t%s %s NOT being run by %s on node %d.\n",
              sess_p->name, time_stamp, prog_name, user_name, node);
    }

    return;
}
```

```

}

/*
 * The process_response_rerr() function processes registration error
 * responses. When an event cannot be registered due to some detected error,
 * a registration error response is sent to the EMAPI client. This function
 * marks the event, as tracked in a program structure, as being unregistered.
 * This will prevent an attempt to unregister a non-registered event later
 * in the program.
 */

static
void process_response_rerr(struct session *sess_p,
                          struct program *progs_p, int progs_elems,
                          struct ha_em_rsp_blk *rsp_blk, int *reg_cnt_p)
{
    struct ha_em_rpb_rerr *error_p;          /* Pointer to error response */
    struct ha_em_rpb_rerr *last_error_p;    /* Beyond last error response*/
    struct program *prog_p;                 /* Pointer to program struct */

    char prog_name[80];                     /* Name of program to which error pertains */
    char user_name[80];                     /* Name of user to which error pertains */
    int node;                               /* Node to which error pertains */

    error_p = rsp_blk->em_resp_blk.em_rpb_rerr;
    last_error_p = error_p + rsp_blk->em_rsp_num_resp;

    for ( ; error_p < last_error_p; error_p++) {

        /*
         * Find the program structure which describes the program associated
         * with the event whose registration failed. The match is made
         * with event identifiers.
         */

        for (prog_p = progs_p; prog_p < progs_p + progs_elems; prog_p++) {
            if (prog_p->eid == error_p->em_event_id) {
                break;
            }
        }

        if (prog_p == progs_p + progs_elems) {
            fprintf(stderr, "%s: Unknown event identifier encountered (0x%x).",
                   "\n", sess_p->name, error_p->em_event_id);
            exit(1);
        }

        /*
         * Extract the program name, user name, and node number from the
         * resource ID.
         */

        breakdown_rsrc_ID(error_p->em_rsrc_ID, prog_name, user_name, &node);

        /*
         * Print a message about the registration error.
         */

        printf("%s: Registration error for %s run by %s on node %d (%d, %d).",
              "\n", sess_p->name, prog_name, user_name, node,
              error_p->em_generr, error_p->em_specerr);

        /*
         * Mark the program as not having an associated event registered.
         */

        prog_p->unreged = 1;
    }

    /*
     * Update the number of events that are registered.
     */

    *reg_cnt_p -= rsp_blk->em_rsp_num_resp;

    return;
}

/*
 * The format_timestamp() function takes a time stamp returned by the
 * Event Manager and converts it into 24-hour time.
 */

static
void format_timestamp(struct timeval *timestamp_p, char *fmt_timestamp_p)
{
    struct tm broken_down_time;

```

## emapi\_v02\_ex03.c

```
    localtime_r((time_t *) &timestamp_p->tv_sec, &broken_down_time);
    (void) strftime(fmt_timestamp_p, 20, "%X", &broken_down_time);
}

return;
}

/*
 * The breakdown_rsrc_ID() function takes a resource ID
 * associated with the IBM.PSSP.Prog.xpcount resource variable and extracts
 * from it the program name, user name, and node number.
 */
static
void breakdown_rsrc_ID(char *rsrc_ID_p,
    char *prog_name_p, char *user_name_p, int *node_p)
{
    char *value_p;
    size_t value_len;

    find_rsrc_ID_value(rsrc_ID_p, "ProgName=", &value_p, &value_len);
    strncpy(prog_name_p, value_p, value_len);
    *(prog_name_p + value_len) = '\0';

    find_rsrc_ID_value(rsrc_ID_p, "UserName=", &value_p, &value_len);
    strncpy(user_name_p, value_p, value_len);
    *(user_name_p + value_len) = '\0';

    find_rsrc_ID_value(rsrc_ID_p, "NodeNum=", &value_p, &value_len);
    *node_p = atoi(value_p);
}

return;
}

/*
 * The find_rsrc_ID_value() takes a resource ID and extracts
 * a value from it.
 */
static
void find_rsrc_ID_value(char *rsrc_ID_p, char *name_p,
    char **value_pp, size_t *value_len_p)
{
    char *bp;
    char *ep;
    size_t len;

    if ((bp = strstr(rsrc_ID_p, name_p)) == NULL) {
        fprintf(stderr, "Unexpected resource ID encountered.\n");
        exit(1);
    }

    bp += strlen(name_p);

    if ((ep = strchr(bp, ';')) != NULL) {
        len = ep - bp;
    } else {
        len = strlen(bp);
    }

    *value_pp = bp;
    *value_len_p = len;
}

return;
}

/*
 * The end_session() function terminates a session with the EMAPI by
 * calling the EMAPI routine ha_em_end_session().
 */
static
void end_session(struct session *sess_p)
{
    struct ha_em_err_blk errb;

    if (ha_em_end_session(sess_p->fd, &errb) == -1) {
        fprintf(stderr, "%s: the ha_em_end_session() returned EM errno %d:\n%s",
            sess_p->name, errb.em_errno, errb.em_errmsg);
        exit(1);
    }

    return;
}
}
```

## The emapi\_v02\_ex04.c Sample Program

```

/* IBM_PROLOG_BEGIN_TAG                               */
/* This is an automatically generated prolog.         */
/*                                                    */
/*                                                    */
/* Licensed Materials - Property of IBM               */
/*                                                    */
/* (C) COPYRIGHT International Business Machines Corp. 1996,1998 */
/* All Rights Reserved                                */
/*                                                    */
/* US Government Users Restricted Rights - Use, duplication or */
/* disclosure restricted by GSA ADP Schedule Contract with IBM Corp. */
/*                                                    */
/* IBM_PROLOG_END_TAG                                 */

/* @(#)44 1.3 src/rsct/pem/emtools/emapi_test/emapi_ex/emapi_v02_ex04.c, emtools, rsct_rtro 6/30/98 10:47:45 */

#include <unistd.h>
#include <stdlib.h>
#include <stdio.h>
#include <string.h>
#include <errno.h>
#include <signal.h>

#define HA_EM_VERSION 2
#include <ha_emapi.h>

/*
 * emapi_v02_ex04.c
 *
 * This program presents an example of using the Event Manager Application
 * Programming Interface (EMAPI). The program uses the EMAPI to obtain
 * the definition of some resource variables.
 *
 * This program can be compiled with the following command:
 *
 * cc -O emapi_v02_ex04.c -o emapi_v02_ex04 -lha_em
 */

/*
 * The following array contains query requests that will be sent to the
 * EMAPI. Each array element is of type struct ha_em_rb_query, which is
 * defined by the EMAPI. Each element specifies a class, a resource
 * variable, and a resource ID.
 *
 * Request 1 requests the definitions of resource variables in the
 * IBM.PSSP.Prog class.
 *
 * Request 2 requests the definitions of resource variables in the
 * IBM.PSSP.SP_HW class whose resource variable names start with
 * IBM.PSSP.SP_HW.Node.
 *
 * Request 3 requests the definitions of resource variables in all
 * classes whose resource IDs include a resource ID element
 * named CPU.
 */

static struct ha_em_rb_query qry_array[] = {
    {"IBM.PSSP.Prog", "", "*"}, /* Request 1 */
    {"IBM.PSSP.SP_HW", "IBM.PSSP.SP_HW.Node", "*"}, /* Request 2 */
    {"", "", "CPU;*" } /* Request 3 */
};

static int qry_cnt = sizeof qry_array / sizeof qry_array[0];

/*
 * Function prototypes for internal functions.
 */

static void setup_signals(void);
static int start_session(void);
static ha_em_qid_t send_defn_request(int sess_fd,
                                   struct ha_em_rb_query *qry_array,
                                   int qry_cnt);
static int rcv_defn_reply(int sess_fd, ha_em_qid_t qid);
static int process_response_query(ha_em_qid_t qid,
                                  struct ha_em_rsp_blk *rsp_blk);
static int process_response_qerr(ha_em_qid_t qid,
                                 struct ha_em_rsp_blk *rsp_blk);
static void end_session(int sess_fd);

/*
 * The main() function calls functions to initialize the program, establish
 * an EMAPI session, send query requests, wait for and process Event Manager
 * responses, and terminate the EMAPI session.
 */

```

## emapi\_v02\_ex04.c

```
int main(int argc, char **argv)
{
    int      sess_fd;          /* Session file descriptor */
    ha_em_qid_t qid;          /* Query identifier */
    int      qend;             /* Query is finished */

    setup_signals();          /* Set signal dispositions */
    sess_fd = start_session(); /* Start EAPI session */
    qid = send_defn_request(sess_fd, qry_array, qry_cnt);
                                /* Send query request */

    /*
     * Receive query replies until all replies have been received.
     */
    while ((qend = rcv_defn_reply(sess_fd, qid)) == 0) {
        /* Nothing to do inside the loop */
    }

    end_session(sess_fd);     /* End EAPI session */
    return 0;
}

/*
 * The setup_signals() function initializes the disposition of a signal.
 *
 * The disposition of the SIGPIPE signal is set such that the signal is
 * ignored. When a process writes to a pipe or connection-oriented socket for
 * which there is no reader, the SIGPIPE signal is delivered to the process.
 * The default disposition for SIGPIPE is to kill the receiving process. If
 * SIGPIPE is ignored, the SIGPIPE signal does not kill the process, and the
 * system call used to write to the pipe or connection-oriented socket sets
 * errno to EPIPE and returns an error indication. It is recommended that
 * clients of the EAPI ignore the SIGPIPE signal, since the EAPI uses
 * connection-oriented sockets to communicate with the Event Manager.
 */

static
void setup_signals(void)
{
    struct sigaction sa;

    sa.sa_handler = SIG_IGN;          /* SIGPIPE to be ignored */
    sigemptyset(&sa.sa_mask);        /* No signals to mask off */
    sa.sa_flags = 0;                 /* No flags needed */

    if (sigaction(SIGPIPE, &sa, NULL) == -1) { /* Change SIGPIPE disposition*/
        perror("sigaction(SIGPIPE)");
        exit(1);
    }

    return;
}

/*
 * The start_session() function establishes a session with the EAPI by
 * calling the EAPI routine ha_em_start_session(). If a session cannot be
 * established, the program is terminated.
 */

static
int start_session(void)
{
    int sess_fd;                  /* Session file descriptor */
    struct ha_em_err_blk errb;     /* EAPI error block */

    /*
     * Since no specific partition name is passed to the
     * ha_em_start_session() routine, it attempts to establish a session
     * with the default partition. The routine returns a file
     * descriptor with which the session can be used.
     */
    sess_fd = ha_em_start_session(HA_EM_DOMAIN_SP, "", &errb);

    /*
     * If ha_em_start_session() indicates a session could not be established,
     * print the error information returned by the routine, and exit the
     * program.
     */
    if (sess_fd == -1) {
        fprintf(stderr, "ha_em_start_session() returned EM errno %d:\n%s",

```

```

        errb.em_errno, errb.em_errmsg);
    }
    exit(1);
}
return sess_fd;          /* Return session fd      */
}

/*
 * The send_defn_request() function uses the EMAPI routine
 * ha_em_send_command() to send the query defined resource variables requests
 * specified by the qry_array and qry_cnt parameters.
 */

static
ha_em_qid_t send_defn_request(int sess_fd,
                             struct ha_em_rb_query *qry_array, int qry_cnt)
{
    size_t          alloc_size;      /* Size of memory to allocate */
    struct ha_em_cmd_blk *cmd_blk;   /* Pointer to command block   */
    struct ha_em_err_blk errb;       /* EMAPI error block          */
    int             rc;              /* Return code                 */
    ha_em_qid_t    qid;             /* Query identifier            */

    /*
     * Allocate enough memory to hold the command block to be given to the
     * EMAPI. The ha_em_cmd_blk structure contains one ha_em_rb_query
     * structure. Additional space must be allocated when more than one
     * query is to be made.
     */

    alloc_size = sizeof(struct ha_em_cmd_blk) +
                 ((qry_cnt - 1) * sizeof(struct ha_em_rb_query));

    if ((cmd_blk = malloc(alloc_size)) == NULL) {
        perror("malloc()");
        exit(1);
    }

    /*
     * Fill in command block header, specifying number of elements in
     * em_rb_query array, command, and subcommand. Indicate a callback
     * routine will not be used to receive the responses.
     */

    cmd_blk->em_cmd_num_elem = qry_cnt;
    cmd_blk->em_cmd           = HA_EM_CMD_QUERY;
    cmd_blk->em_subcmd        = HA_EM_SCMD_QDEF;
    cmd_blk->em_qcb           = NULL;
    cmd_blk->em_qcb_arg       = NULL;

    /*
     * Copy the array of query requests into the command block.
     */

    memcpy(cmd_blk->em_res_blk.em_rb_query, qry_array,
           qry_cnt * sizeof(struct ha_em_rb_query));

    /*
     * Send the command. A more elaborate program might check for the
     * HA_EM_ECONNLOST Event Manager error number when an error is
     * returned, and attempt to restart the session. But, this program
     * just terminates if ha_em_send_command() detects a lost connection.
     */

    rc = ha_em_send_command(sess_fd, cmd_blk, &errb);

    if (rc == -1) {
        fprintf(stderr, "ha_em_send_command() returned EM errno %d:\n%s",
               errb.em_errno, errb.em_errmsg);
        exit(1);
    }

    /*
     * Save the query identifier assigned to the query request just sent.
     */

    qid = cmd_blk->em_qid;

    free(cmd_blk);          /* Free the allocated command block. */

    return qid;            /* Return the query identifier      */
}

/*
 * The recv_defn_reply() function processes an Event Manager response
 * in the specified EMAPI session. First, ha_em_receive_response() is called
 * to receive the response. If ha_em_receive_response() returns an error
 * the program is terminated. The ha_em_receive_response() function may
 * indicate that there really isn't any response for the EMAPI client to
 * process at this time. In that case, this function just returns. If a

```

## emapi\_v02\_ex04.c

```
* response has been received, it is processed depending on what type of
* response it is.
*
* This function returns 0 if all query responses have not been received yet.
* If all query responses have been received, 1 is returned.
*/

static
int recv_defn_reply(int sess_fd, ha_em_qid_t qid)
{
    struct ha_em_rsp_blk  *rsp_blk; /* Pointer to the response block */
    int rc;                /* Return code */
    struct ha_em_err_blk  errb;    /* Event Manager error block */
    int qend;             /* Query end flag */

    /*
     * Receive response from session, if there is one to receive.
     */

    rc = ha_em_receive_response(sess_fd, &rsp_blk, &errb);

    if (rc == -1) {
        /*
         * Some error has occurred receiving a response; terminate the
         * program. A more elaborate program might check for the
         * HA_EM_ECONNLOST Event Manager error number when an error is
         * returned, and attempt to restart the session. But, this program
         * just terminates if ha_em_receive_response() detects a lost
         * connection.
         */

        fprintf(stderr, "ha_em_receive_response() returned EM errno %d:\n%s",
            errb.em_errno, errb.em_errmsg);
        exit(1);
    }

    /*
     * If the ha_em_receive_response() routine returned zero,
     * there is no response for the EMAPI client (this program) to process
     * at this time. The response may have been for the EMAPI itself, or
     * a full response may not be available yet. Just return.
     */

    if (rc == 0) {
        return 0; /* Return indicating query is not ended */
    }

    /*
     * A response for the EMAPI client (this program) has been returned.
     * The response buffer is pointed to by the rsp_blk variable. Appropriate
     * processing for the response depends on the command type.
     */

    switch (rsp_blk->em_cmd) {
        case HA_EM_CMD_QUERY: /* Query response */
            qend = process_response_query(qid, rsp_blk);
            break;

        case HA_EM_CMD_QERR: /* Query error response */
            qend = process_response_qerr(qid, rsp_blk);
            break;

        default: /* Unexpected response */
            fprintf(stderr, "Program received unexpected command "
                "response: %d.\n", rsp_blk->em_cmd);
            exit(1);
            break;
    }

    /*
     * The EMAPI client (this program) must free the memory associated with
     * the returned response block when it is no longer needed.
     */

    free(rsp_blk);

    return qend;
}

/*
 * The process_response_query() function processes query responses from the
 * Event Manager. Several points should be kept in mind:
 *
 * - The response block may contain multiple query responses. The
 *   number of responses included in the response block is given in
 *   the response block header.
 */
```

```

*
* - An query response may be an error response. Such a response
*   may indicate an error of a temporary nature.
*
* - This routine compares the query identifier in the response block
*   with the query identifier assigned to the request sent by
*   this program to the EMAPI. This should not really be necessary.
*   If the program had sent multiple query requests to the EMAPI,
*   through multiple calls to ha_em_send_command(), the query identifiers
*   in the response blocks could be used to associate responses with
*   requests.
*
* - Multiple response blocks may be associated with one request block.
*   The response to a query request is not complete until a response
*   block is received with the em_gend field set to a non-zero value.
*/

```

```

static
int process_response_query(ha_em_qid_t qid, struct ha_em_rsp_blk *rsp_blk)
{
    struct ha_em_rpb_qdef *qp;          /* Pointer to query response */
    struct ha_em_rpb_qdef *last_qp;    /* Beyond last query response */

    /*
     * Check query identifier in response block.
     */
    if (rsp_blk->em_qid != qid) {
        fprintf(stderr, "Query response block contains unexpected "
            "query identifier: %d.\n", rsp_blk->em_qid);
        exit(1);
    }

    /*
     * Look at all the response elements in the response block.
     */
    qp = rsp_blk->em_resp_blk.em_rpb_qdef;
    last_qp = qp + rsp_blk->em_rsp_num_resp;

    for ( ; qp < last_qp; qp++) {
        printf("=====\n\n");

        /*
         * Check for error code.
         */
        if (qp->em_errnum != 0) {
            printf("Query response for class \"%s\", "
                "resource variable \"%s\", resource ID \"%s\" "
                "unexpectedly contains an error (%d, %d).\n",
                qp->em_class, qp->em_name, qp->em_rsrc_ID,
                qp->em_generr, qp->em_specerr);
            continue;
        }

        /*
         * Print definition of a resource variable.
         */
        printf("Resource Variable Name: \"%s\"\n"
            "Variable Value Type: %s\n"
            "Variable Data Type: %s\n"
            "Variable SBS Format: \"%s\"\n"
            "Variable Initial Value: \"%s\"\n"
            "Variable Class: \"%s\"\n"
            "Resource ID: \"%s\"\n"
            "PTX Name: \"%s\"\n"
            "Default Expression: \"%s\"\n"
            "Locator: \"%s\"\n"
            "\n"
            "Variable Description\n"
            "-----\n"
            "%s\n"
            "\n"
            "Resource ID Description\n"
            "-----\n"
            "%s\n"
            "\n"
            "Event Description\n"
            "-----\n"
            "%s\n"
            "\n",
            qp->em_name,
            qp->em_value_type == ha_emVTcounter ? "Counter" :
            qp->em_value_type == ha_emVTquantity ? "Quantity" :
            qp->em_value_type == ha_emVTstate ? "State" : "Unknown",
            qp->em_data_type == ha_emDTlong ? "long" :

```

## emapi\_v02\_ex04.c

```

        qp->em_data_type == ha_emDTfloat ? "float" :
        qp->em_data_type == ha_emDTsbs ? "SBS" : "Unknown",
        qp->em_sbs_format,
        qp->em_init_value,
        qp->em_class,
        qp->em_rsrc_ID,
        qp->em_ptx_name,
        qp->em_dflt_expr,
        qp->em_locator,
        qp->em_descr,
        qp->em_rsrc_ID_descr,
        qp->em_event_descr);
    }
}
return rsp_blk->em_qend;
}

/*
 * The process_response_qerr() function processes query error
 * responses.
 */
static
int process_response_qerr(ha_em_qid_t qid, struct ha_em_rsp_blk *rsp_blk)
{
    struct ha_em_rpb_qerr *error_p; /* Pointer to error response */
    struct ha_em_rpb_qerr *last_error_p; /* Beyond last error response*/

    /*
     * Check query identifier in response block.
     */
    if (rsp_blk->em_qid != qid) {
        fprintf(stderr, "Query response block contains unexpected "
            "query identifier: %d.\n", rsp_blk->em_qid);
        exit(1);
    }

    error_p = rsp_blk->em_resp_blk.em_rpb_qerr;
    last_error_p = error_p + rsp_blk->em_rsp_num_resp;

    for ( ; error_p < last_error_p; error_p++) {
        printf("=====\n\n");

        /*
         * Print error information for a resource variable.
         */
        printf("Resource Variable Name: %s\n"
            "Variable Class: %s\n"
            "Resource ID: %s\n"
            "General Error: %d\n"
            "Specific Error: %d\n"
            "Additional Error Info: %d\n",
            error_p->em_name,
            error_p->em_class,
            error_p->em_rsrc_ID,
            error_p->em_generr,
            error_p->em_specerr,
            error_p->em_errinfo);
    }

    return rsp_blk->em_qend;
}

/*
 * The end_session() function terminates a session with the EAPI by
 * calling the EAPI routine ha_em_end_session().
 */
static
void end_session(int sess_fd)
{
    struct ha_em_err_blk errb;

    if (ha_em_end_session(sess_fd, &errb) == -1) {
        fprintf(stderr, "ha_em_end_session() returned EM errno %d:\n%s",
            errb.em_errno, errb.em_errmsg);
        exit(1);
    }
}

```

```
    }  
    return;  
}
```

**emapi\_v02\_ex04.c**

---

## Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation  
Licensing  
2-31 Roppongi 3-chome, Minato-ku  
Tokyo 106, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation  
Department LJEB/P905  
2455 South Road

Poughkeepsie, NY 12601-5400  
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

#### COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

---

## Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States or other countries or both:

- AIX
- AIX/6000
- DATABASE 2
- DB2
- ES/9000
- ESCON
- HACMP/6000
- IBM
- IBMLink
- LoadLeveler
- Micro Channel
- Netfinity
- Netfinity Manager
- POWERparallel
- POWERserver
- RS/6000
- RS/6000 Scalable POWERparallel Systems
- Scalable POWERparallel Systems
- SP
- System/370
- System/390
- TURBOWAYS

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, BackOffice, MS-DOS, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Tivoli Enterprise Console is a trademark of Tivoli Systems Inc. in the United States, other countries, or both.

UNIX is a registered trademark in the United States and other countries licensed exclusively through The Open Group.

Other company, product, and service names may be the trademarks or service marks of others.

---

## Publicly Available Software

PSSP includes software that is publicly available:

**expect**

Programmed dialogue with interactive programs

**Perl** Practical Extraction and Report Language

**SUP** Software Update Protocol

**Tcl** Tool Command Language

**TclX** Tool Command Language Extended

**Tk** Tcl-based Tool Kit for X-windows

This book discusses the use of these products only as they apply specifically to the RS/6000 SP system. The distribution for these products includes the source code and associated documentation.

**/usr/lpp/ssp/public** contains the compressed **tar** files of the publicly available software. (IBM has made minor modifications to the versions of Tcl and Tk used in the SP system to improve their security characteristics. Therefore, the IBM-supplied versions do not match exactly the versions you may build from the compressed **tar** files.) All copyright notices in the documentation must be respected. You can find version and distribution information for each of these products that are part of your selected install options in the **/usr/lpp/ssp/README/ssp.public.README** file.

---

## Software Update Protocol

IBM has provided modifications to this software. The resulting software is provided to you on an "AS IS" basis and WITHOUT WARRANTY OF ANY KIND, WHETHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.



---

## Glossary of Terms and Abbreviations

This glossary includes terms and definitions from:

- The *IBM Dictionary of Computing*, New York: McGraw-Hill, 1994.
- The *American National Standard Dictionary for Information Systems*, ANSI X3.172-1990, copyright 1990 by the American National Standards Institute (ANSI). Copies can be purchased from the American National Standards Institute, 1430 Broadway, New York, New York 10018. Definitions are identified by the symbol (A) after the definition.
- The *ANSI/EIA Standard - 440A: Fiber Optic Terminology* copyright 1989 by the Electronics Industries Association (EIA). Copies can be purchased from the Electronic Industries Association, 2001 Pennsylvania Avenue N.W., Washington, D.C. 20006. Definitions are identified by the symbol (E) after the definition.
- The *Information Technology Vocabulary* developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions of published parts of this vocabulary are identified by the symbol (I) after the definition; definitions taken from draft international standards, committee drafts, and working papers being developed by ISO/IEC JTC1/SC1 are identified by the symbol (T) after the definition, indicating that final agreement has not yet been reached among the participating National Bodies of SC1.

The following cross-references are used in this glossary:

- Contrast with.** This refers to a term that has an opposed or substantively different meaning.
- See.** This refers the reader to multiple-word terms in which this term appears.
- See also.** This refers the reader to terms that have a related, but not synonymous, meaning.
- Synonym for.** This indicates that the term has the same meaning as a preferred term, which is defined in the glossary.

This section contains some of the terms that are commonly used in the SP publications.

IBM is grateful to the American National Standards Institute (ANSI) for permission to reprint its definitions from the *American National Standard Vocabulary for Information Processing* (Copyright

1970 by American National Standards Institute, Incorporated), which was prepared by Subcommittee X3K5 on Terminology and Glossary of the American National Standards Committee X3. ANSI definitions are preceded by an asterisk (\*).

Other definitions in this glossary are taken from *IBM Vocabulary for Data Processing, Telecommunications, and Office Systems* (SC20-1699) and *IBM DATABASE 2 Application Programming Guide for TSO Users* (SC26-4081).

### A

**adapter.** An adapter is a mechanism for attaching parts. For example, an adapter could be a part that electrically or physically connects a device to a computer or to another device. In the SP system, network connectivity is supplied by various adapters, some optional, that can provide connection to I/O devices, networks of workstations, and mainframe networks. Ethernet, FDDI, token-ring, HiPPI, SCSI, FCS, and ATM are examples of adapters that can be used as part of an SP system.

**address.** A character or group of characters that identifies a register, a device, a particular part of storage, or some other data source or destination.

**AFS.** A distributed file system that provides authentication services as part of its file system creation.

**AIX.** Abbreviation for Advanced Interactive Executive, IBM's licensed version of the UNIX operating system. AIX is particularly suited to support technical computing applications, including high function graphics and floating point computations.

**Amd.** Berkeley Software Distribution automount daemon.

**API.** Application Programming Interface. A set of programming functions and routines that provide access between the Application layer of the OSI seven-layer model and applications that want to use the network. It is a software interface.

**application.** The use to which a data processing system is put; for example, a payroll application, an airline reservation application.

**application data.** The data that is produced using an application program.

**ARP.** Address Resolution Protocol.

**ATM.** Asynchronous Transfer Mode. (See *TURBOWAYS 100 ATM Adapter*.)

**Authentication.** The process of validating the identity of a user or server.

**Authorization.** The process of obtaining permission to perform specific actions.

## B

**batch processing.** \* (1) The processing of data or the accomplishment of jobs accumulated in advance in such a manner that each accumulation thus formed is processed or accomplished in the same run. \* (2) The processing of data accumulating over a period of time. \* (3) Loosely, the execution of computer programs serially. (4) Computer programs executed in the background.

**BMCA.** Block Multiplexer Channel Adapter. The block multiplexer channel connection allows the RS/6000 to communicate directly with a host System/370 or System/390; the host operating system views the system unit as a control unit.

**BOS.** The AIX Base Operating System.

## C

**call home function.** The ability of a system to call the IBM support center and open a PMR to have a repair scheduled.

**CDE.** Common Desktop Environment. A graphical user interface for UNIX.

**charge feature.** An optional feature for either software or hardware for which there is a charge.

**CLI.** Command Line Interface.

**client.** \* (1) A function that requests services from a server and makes them available to the user. \* (2) A term used in an environment to identify a machine that uses the resources of the network.

**Client Input/Output Sockets (CLIO/S).** A software package that enables high-speed data and tape access between SP systems, AIX systems, and ES/9000 mainframes.

**CLIO/S.** Client Input/Output Sockets.

**CMI.** Centralized Management Interface provides a series of SMIT menus and dialogues used for defining and querying the SP system configuration.

**connectionless.** A communication process that takes place without first establishing a connection.

**connectionless network.** A network in which the sending logical node must have the address of the receiving logical node before information interchange can begin. The packet is routed through nodes in the network based on the destination address in the packet. The originating source does not receive an acknowledgment that the packet was received at the destination.

**control workstation.** A single point of control allowing the administrator or operator to monitor and manage the SP system using the IBM AIX Parallel System Support Programs.

**css.** Communication subsystem.

## D

**daemon.** A process, not associated with a particular user, that performs system-wide functions such as administration and control of networks, execution of time-dependent activities, line printer spooling and so forth.

**DASD.** Direct Access Storage Device. Storage for input/output data.

**DCE.** Distributed Computing Environment.

**DFS.** distributed file system. A subset of the IBM Distributed Computing Environment.

**DNS.** Domain Name Service. A hierarchical name service which maps high level machine names to IP addresses.

## E

**Error Notification Object.** An object in the SDR that is matched with an error log entry. When an error log entry occurs that matches the Notification Object, a user-specified action is taken.

**ESCON.** Enterprise Systems Connection. The ESCON channel connection allows the RS/6000 to communicate directly with a host System/390; the host operating system views the system unit as a control unit.

**Ethernet.** (1) Ethernet is the standard hardware for TCP/IP local area networks in the UNIX marketplace. It is a 10-megabit per second baseband type LAN that allows multiple stations to access the transmission medium at will without prior coordination, avoids contention by using carrier sense and deference, and resolves contention by collision detection (CSMA/CD). (2) A passive coaxial cable whose interconnections contain devices or components, or both, that are all active. It uses CSMA/CD technology to provide a best-effort delivery system.

**Ethernet network.** A baseband LAN with a bus topology in which messages are broadcast on a coaxial

cabling using the carrier sense multiple access/collision detection (CSMA/CD) transmission method.

**event.** In Event Management, the notification that an expression evaluated to true. This evaluation occurs each time an instance of a resource variable is observed.

**expect.** Programmed dialogue with interactive programs.

**expression.** In Event Management, the relational expression between a resource variable and other elements (such as constants or the previous value of an instance of the variable) that, when true, generates an event. An example of an expression is  $X < 10$  where  $X$  represents the resource variable `IBM.PSSP.aixos.PagSp.%totalfree` (the percentage of total free paging space). When the expression is true, that is, when the total free paging space is observed to be less than 10%, the Event Management subsystem generates an event to notify the appropriate application.

## F

**failover.** Also called failover, the sequence of events when a primary or server machine fails and a secondary or backup machine assumes the primary workload. This is a disruptive failure with a short recovery time.

**fall back.** Also called fallback, the sequence of events when a primary or server machine takes back control of its workload from a secondary or backup machine.

**FDDI.** Fiber Distributed Data Interface.

**FFDC.** First Failure Data Capture

**Fiber Distributed Data Interface (FDDI).** An American National Standards Institute (ANSI) standard for 100-megabit-per-second LAN using optical fiber cables. An FDDI local area network (LAN) can be up to 100 km (62 miles) and can include up to 500 system units. There can be up to 2 km (1.24 miles) between system units and/or concentrators.

**File Transfer Protocol (FTP).** The Internet protocol (and program) used to transfer files between hosts. It is an application layer protocol in TCP/IP that uses TELNET and TCP protocols to transfer bulk-data files between machines or hosts.

**file.** \* A set of related records treated as a unit, for example, in stock control, a file could consist of a set of invoices.

**file name.** A CMS file identifier in the form of 'filename filetype filemode' (like: TEXT DATA A).

**file server.** A centrally located computer that acts as a storehouse of data and applications for numerous users of a local area network.

**File Transfer Protocol (FTP).** The Internet protocol (and program) used to transfer files between hosts. It is an application layer protocol in TCP/IP that uses TELNET and TCP protocols to transfer bulk-data files between machines or hosts.

**First Failure Data Capture (FFDC).** A set of utilities used for recording persistent records of failures and significant software incidents. It provides a means of associating failures to one another, thus allowing software to link effects of a failure to their causes and thereby facilitating discovery of the root cause of a failure.

**foreign host.** Any host on the network other than the local host.

**FTP.** File transfer protocol.

## G

**gateway.** An intelligent electronic device interconnecting dissimilar networks and providing protocol conversion for network compatibility. A gateway provides transparent access to dissimilar networks for nodes on either network. It operates at the session presentation and application layers.

## H

**HACMP.** High Availability Cluster Multi-Processing for AIX.

**HACWS.** High Availability Control Workstation function, based on HACMP, provides for a backup control workstation for the SP system.

**Hashed Shared Disk (HSD).** The data striping device for the IBM Virtual Shared Disk. The device driver lets application programs stripe data across physical disks in multiple IBM Virtual Shared Disks, thus reducing I/O bottlenecks.

**help key.** In the SP graphical interface, the key that gives you access to the SP graphical interface help facility.

**High Availability Cluster Multi-Processing.** An IBM facility to cluster nodes or components to provide high availability by eliminating single points of failure.

**HiPPI.** High Performance Parallel Interface. RS/6000 units can attach to a HiPPI network as defined by the ANSI specifications. The HiPPI channel supports burst rates of 100 Mbps over dual simplex cables; connections can be up to 25 km in length as defined by the standard and can be extended using third-party HiPPI switches and fiber optic extenders.

**home directory.** The directory associated with an individual user.

**host.** A computer connected to a network, and providing an access method to that network. A host provides end-user services.

## I

**instance vector.** Obsolete term for resource ID.

**Intermediate Switch Board.** Switches mounted in the Sp Switch expansion frame.

**Internet.** A specific inter-network consisting of large national backbone networks such as APARANET, MILNET, and NSFnet, and a myriad of regional and campus networks all over the world. The network uses the TCP/IP protocol suite.

**Internet Protocol (IP).** (1) A protocol that routes data through a network or interconnected networks. IP acts as an interface between the higher logical layers and the physical network. This protocol, however, does not provide error recovery, flow control, or guarantee the reliability of the physical network. IP is a connectionless protocol. (2) A protocol used to route data from its source to its destination in an Internet environment.

**IP address.** A 32-bit address assigned to devices or hosts in an IP internet that maps to a physical address. The IP address is composed of a network and host portion.

**ISB.** Intermediate Switch Board.

## K

**Kerberos.** A service for authenticating users in a network environment.

**kernel.** The core portion of the UNIX operating system which controls the resources of the CPU and allocates them to the users. The kernel is memory-resident, is said to run in "kernel mode" and is protected from user tampering by the hardware.

## L

**LAN.** (1) Acronym for Local Area Network, a data network located on the user's premises in which serial transmission is used for direct data communication among data stations. (2) Physical network technology that transfers data at a high speed over short distances. (3) A network in which a set of devices is connected to another for communication and that can be connected to a larger network.

**local host.** The computer to which a user's terminal is directly connected.

**log database.** A persistent storage location for the logged information.

**log event.** The recording of an event.

**log event type.** A particular kind of log event that has a hierarchy associated with it.

**logging.** The writing of information to persistent storage for subsequent analysis by humans or programs.

## M

**mask.** To use a pattern of characters to control retention or elimination of portions of another pattern of characters.

**menu.** A display of a list of available functions for selection by the user.

**Motif.** The graphical user interface for OSF, incorporating the X Window System. Also called OSF/Motif.

**MTBF.** Mean time between failure. This is a measure of reliability.

**MTTR.** Mean time to repair. This is a measure of serviceability.

## N

**naive application.** An application with no knowledge of a server that fails over to another server. Client to server retry methods are used to reconnect.

**network.** An interconnected group of nodes, lines, and terminals. A network provides the ability to transmit data to and receive data from other systems and users.

**NFS.** Network File System. NFS allows different systems (UNIX or non-UNIX), different architectures, or vendors connected to the same network, to access remote files in a LAN environment as though they were local files.

**NIM.** Network Installation Management is provided with AIX to install AIX on the nodes.

**NIM client.** An AIX system installed and managed by a NIM master. NIM supports three types of clients:

- Standalone
- Diskless
- Dataless

**NIM master.** An AIX system that can install one or more NIM clients. An AIX system must be defined as a NIM master before defining any NIM clients on that system. A NIM master manages the configuration database containing the information for the NIM clients.

**NIM object.** A representation of information about the NIM environment. NIM stores this information as objects in the NIM database. The types of objects are:

- Network
- Machine
- Resource

**NIS.** Network Information System.

**node.** In a network, the point where one or more functional units interconnect transmission lines. A computer location defined in a network. The SP system can house several different types of nodes for both serial and parallel processing. These node types can include thin nodes, wide nodes, 604 high nodes, as well as other types of nodes both internal and external to the SP frame.

**Node Switch Board.** Switches mounted on frames that contain nodes.

**NSB.** Node Switch Board.

**NTP.** Network Time Protocol.

## O

**ODM.** Object Data Manager. In AIX, a hierarchical object-oriented database for configuration data.

## P

**parallel environment.** A system environment where message passing or SP resource manager services are used by the application.

**Parallel Environment.** A licensed IBM program used for message passing applications on the SP or RS/6000 platforms.

**parallel processing.** A multiprocessor architecture which allows processes to be allocated to tightly coupled multiple processors in a cooperative processing environment, allowing concurrent execution of tasks.

**parameter.** \* (1) A variable that is given a constant value for a specified application and that may denote the application. \* (2) An item in a menu for which the operator specifies a value or for which the system provides a value when the menu is interpreted. \* (3) A name in a procedure that is used to refer to an argument that is passed to the procedure. \* (4) A particular piece of information that a system or application program needs to process a request.

**partition.** See system partition.

**Perl.** Practical Extraction and Report Language.

**perspective.** The primary window for each SP Perspectives application, so called because it provides a unique view of an SP system.

**pipe.** A UNIX utility allowing the output of one command to be the input of another. Represented by the | symbol. It is also referred to as filtering output.

**PMR.** Problem Management Report.

**POE.** Formerly Parallel Operating Environment, now Parallel Environment for AIX.

**port.** (1) An end point for communication between devices, generally referring to physical connection. (2) A 16-bit number identifying a particular TCP or UDP resource within a given TCP/IP node.

**predicate.** Obsolete term for expression.

**Primary node or machine.** (1) A device that runs a workload and has a standby device ready to assume the primary workload if that primary node fails or is taken out of service. (2) A node on the SP Switch that initializes, provides diagnosis and recovery services, and performs other operations to the switch network. (3) In IBM Virtual Shared Disk function, when physical disks are connected to two nodes (twin-tailed), one node is designated as the primary node for each disk and the other is designated the secondary, or backup, node. The primary node is the server node for IBM Virtual Shared Disks defined on the physical disks under normal conditions. The secondary node can become the server node for the disks if the primary node is unavailable (off-line or down).

**Problem Management Report.** The number in the IBM support mechanism that represents a service incident with a customer.

**process.** \* (1) A unique, finite course of events defined by its purpose or by its effect, achieved under defined conditions. \* (2) Any operation or combination of operations on data. \* (3) A function being performed or waiting to be performed. \* (4) A program in operation. For example, a daemon is a system process that is always running on the system.

**protocol.** A set of semantic and syntactic rules that defines the behavior of functional units in achieving communication.

## R

**RAID.** Redundant array of independent disks.

**rearm expression.** In Event Management, an expression used to generate an event that alternates with an original event expression in the following way: the event expression is used until it is true, then the rearm expression is used until it is true, then the event expression is used, and so on. The rearm expression is commonly the inverse of the event expression (for example, a resource variable is on or off). It can also be used with the event expression to define an upper and lower boundary for a condition of interest.

**rearm predicate.** Obsolete term for rearm expression

**remote host.** See *foreign host*.

**resource.** In Event Management, an entity in the system that provides a set of services. Examples of resources include hardware entities such as processors, disk drives, memory, and adapters, and software entities such as database applications, processes, and file systems. Each resource in the system has one or more attributes that define the state of the resource.

**resource identifier.** In Event Management, a set of elements, where each element is a name/value pair of the form `name=value`, whose values uniquely identify the copy of the resource (and by extension, the copy of the resource variable) in the system.

**resource monitor.** A program that supplies information about resources in the system. It can be a command, a daemon, or part of an application or subsystem that manages any type of system resource.

**resource variable.** In Event Management, the representation of an attribute of a resource. An example of a resource variable is `IBM.AIX.PagSp.%totalfree`, which represents the percentage of total free paging space. `IBM.AIX.PagSp` specifies the resource name and `%totalfree` specifies the resource attribute.

**RISC.** Reduced Instruction Set Computing (RISC), the technology for today's high performance personal computers and workstations, was invented in 1975. Uses a small simplified set of frequently used instructions for rapid execution.

**rlogin (remote LOGIN).** A service offered by Berkeley UNIX systems that allows authorized users of one machine to connect to other UNIX systems across a network and interact as if their terminals were connected directly. The rlogin software passes information about the user's environment (for example, terminal type) to the remote machine.

**RPC.** Acronym for Remote Procedure Call, a facility that a client uses to have a server execute a procedure call. This facility is composed of a library of procedures plus an XDR.

**RSH.** A variant of RLOGIN command that invokes a command interpreter on a remote UNIX machine and passes the command line arguments to the command interpreter, skipping the LOGIN step completely. See also *rlogin*.

## S

**SCSI.** Small Computer System Interface.

**Secondary node.** In IBM Virtual Shared Disk function, when physical disks are connected to two nodes (twin-tailed), one node is designated as the primary

node for each disk and the other is designated as the secondary, or backup, node. The secondary node acts as the server node for the IBM Virtual Shared disks defined on the physical disks if the primary node is unavailable (off-line or down).

**server.** (1) A function that provides services for users. A machine may run client and server processes at the same time. (2) A machine that provides resources to the network. It provides a network service, such as disk storage and file transfer, or a program that uses such a service. (3) A device, program, or code module on a network dedicated to providing a specific service to a network. (4) On a LAN, a data station that provides facilities to other data stations. Examples are file server, print server, and mail server.

**shell.** The shell is the primary user interface for the UNIX operating system. It serves as command language interpreter, programming language, and allows foreground and background processing. There are three different implementations of the shell concept: Bourne, C and Korn.

**Small Computer System Interface (SCSI).** An input and output bus that provides a standard interface for the attachment of various direct access storage devices (DASD) and tape drives to the RS/6000.

**Small Computer Systems Interface Adapter (SCSI Adapter).** An adapter that supports the attachment of various direct-access storage devices (DASD) and tape drives to the RS/6000.

**SMIT.** The System Management Interface Toolkit is a set of menu driven utilities for AIX that provides functions such as transaction login, shell script creation, automatic updates of object database, and so forth.

**SNMP.** Simple Network Management Protocol. (1) An IP network management protocol that is used to monitor attached networks and routers. (2) A TCP/IP-based protocol for exchanging network management information and outlining the structure for communications among network devices.

**socket.** (1) An abstraction used by Berkeley UNIX that allows an application to access TCP/IP protocol functions. (2) An IP address and port number pairing. (3) In TCP/IP, the Internet address of the host computer on which the application runs, and the port number it uses. A TCP/IP application is identified by its socket.

**standby node or machine.** A device that waits for a failure of a primary node in order to assume the identity of the primary node. The standby machine then runs the primary's workload until the primary is back in service.

**subnet.** Shortened form of subnetwork.

**subnet mask.** A bit template that identifies to the TCP/IP protocol code the bits of the host address that are to be used for routing for specific subnetworks.

**subnetwork.** Any group of nodes that have a set of common characteristics, such as the same network ID.

**subsystem.** A software component that is not usually associated with a user command. It is usually a daemon process. A subsystem will perform work or provide services on behalf of a user request or operating system request.

**SUP.** Software Update Protocol.

**Sysctl.** Secure System Command Execution Tool. An authenticated client/server system for running commands remotely and in parallel.

**syslog.** A BSD logging system used to collect and manage other subsystem's logging data.

**System Administrator.** The user who is responsible for setting up, modifying, and maintaining the SP system.

**system partition.** A group of nonoverlapping nodes on a switch chip boundary that act as a logical SP system.

## T

**tar.** Tape ARchive, is a standard UNIX data archive utility for storing data on tape media.

**Tcl.** Tool Command Language.

**TclX.** Tool Command Language Extended.

**TCP.** Acronym for Transmission Control Protocol, a stream communication protocol that includes error recovery and flow control.

**TCP/IP.** Acronym for Transmission Control Protocol/Internet Protocol, a suite of protocols designed to allow communication between networks regardless of the technologies implemented in each network. TCP provides a reliable host-to-host protocol between hosts in packet-switched communications networks and in interconnected systems of such networks. It assumes that the underlying protocol is the Internet Protocol.

**Telnet.** Terminal Emulation Protocol, a TCP/IP application protocol that allows interactive access to foreign hosts.

**Tk.** Tcl-based Tool Kit for X Windows.

**TMPCP.** Tape Management Program Control Point.

**token-ring.** (1) Network technology that controls media access by passing a token (special packet or frame) between media-attached machines. (2) A network with a ring topology that passes tokens from one attaching device (node) to another. (3) The IBM Token-Ring LAN connection allows the RS/6000 system unit to participate in a LAN adhering to the IEEE 802.5

Token-Passing Ring standard or the ECMA standard 89 for Token-Ring, baseband LANs.

**transaction.** An exchange between the user and the system. Each activity the system performs for the user is considered a transaction.

**transceiver (transmitter-receiver).** A physical device that connects a host interface to a local area network, such as Ethernet. Ethernet transceivers contain electronics that apply signals to the cable and sense collisions.

**transfer.** To send data from one place and to receive the data at another place. Synonymous with move.

**transmission.** \* The sending of data from one place for reception elsewhere.

**TURBOWAYS 100 ATM Adapter.** An IBM high-performance, high-function intelligent adapter that provides dedicated 100 Mbps ATM (asynchronous transfer mode) connection for high-performance servers and workstations.

## U

**UDP.** User Datagram Protocol.

**UNIX operating system.** An operating system developed by Bell Laboratories that features multiprogramming in a multiuser environment. The UNIX operating system was originally developed for use on minicomputers, but has been adapted for mainframes and microcomputers. **Note:** The AIX operating system is IBM's implementation of the UNIX operating system.

**user.** Anyone who requires the services of a computing system.

**User Datagram Protocol (UDP).** (1) In TCP/IP, a packet-level protocol built directly on the Internet Protocol layer. UDP is used for application-to-application programs between TCP/IP host systems. (2) A transport protocol in the Internet suite of protocols that provides unreliable, connectionless datagram service. (3) The Internet Protocol that enables an application programmer on one machine or process to send a datagram to an application program on another machine or process.

**user ID.** A nonnegative integer, contained in an object of type *uid\_t*, that is used to uniquely identify a system user.

## V

**Virtual Shared Disk, IBM.** The function that allows application programs executing at different nodes of a system partition to access a raw logical volume as if it

were local at each of the nodes. In actuality, the logical volume is local at only one of the nodes (the server node).

## **W**

**workstation.** \* (1) A configuration of input/output equipment at which an operator works. \* (2) A terminal or microcomputer, usually one that is connected to a mainframe or to a network, at which a user can perform applications.

## **X**

**X Window System.** A graphical user interface product.

---

## Bibliography

This bibliography helps you find product documentation related to the RS/6000 SP hardware and software products.

You can find most of the IBM product information for RS/6000 SP products on the World Wide Web. Formats for both viewing and downloading are available.

PSSP documentation is shipped with the PSSP product in a variety of formats and can be installed on your system. The man pages for public code that PSSP includes are also available online.

You can order hard copies of the product documentation from IBM. This bibliography lists the titles that are available and their order numbers.

Finally, this bibliography contains a list of non-IBM publications that discuss parallel computing and other topics related to the RS/6000 SP.

---

### Information Formats

Documentation supporting RS/6000 SP software licensed programs is no longer available from IBM in hardcopy format. However, you can view, search, and print documentation in the following ways:

- On the World Wide Web
- Online (from the product media or the SP Resource Center)

---

### Finding Documentation on the World Wide Web

Most of the RS/6000 SP hardware and software books are available from the IBM RS/6000 Web site at:

<http://www.rs6000.ibm.com>

You can view a book or download a Portable Document Format (PDF) version of it. At the time this manual was published, the Web address of the "RS/6000 SP Product Documentation Library" page was:

[http://www.rs6000.ibm.com/resource/aix\\_resource/sp\\_books](http://www.rs6000.ibm.com/resource/aix_resource/sp_books)

However, the structure of the RS/6000 Web site can change over time.

---

### Accessing PSSP Documentation Online

On the same medium as the PSSP product code, IBM ships PSSP man pages, HTML files, and PDF files. In order to use these publications, you must first install the **ssp.docs** file set.

To view the PSSP HTML publications, you need access to an HTML document browser such as Netscape. The HTML files and an index that links to them are installed in the **/usr/lpp/ssp/html** directory. Once installed, you can also view the HTML files from the RS/6000 SP Resource Center.

If you have installed the SP Resource Center on your SP system, you can access it by entering the **/usr/lpp/ssp/bin/resource\_center** command. If you have the SP Resource Center on CD-ROM, see the **readme.txt** file for information about how to run it.

To view the PSSP PDF publications, you need access to the Adobe Acrobat Reader. The Acrobat Reader is shipped with the AIX Version 4.3 Bonus Pack and is also freely available for downloading from the Adobe Web site at:

| <http://www.adobe.com>

To successfully print a large PDF file (approximately 300 or more pages) from the Adobe Acrobat reader, you may need to select the "Download Fonts Once" button on the Print window.

---

## Manual Pages for Public Code

The following manual pages for public code are available in this product:

**SUP** /usr/lpp/ssp/man/man1/sup.1

### Perl (Version 4.036)

/usr/lpp/ssp/perl/man/perl.man

/usr/lpp/ssp/perl/man/h2ph.man

/usr/lpp/ssp/perl/man/s2p.man

/usr/lpp/ssp/perl/man/a2p.man

Manual pages and other documentation for **Tcl**, **TclX**, **Tk**, and **expect** can be found in the compressed **tar** files located in the **/usr/lpp/ssp/public** directory.

---

## RS/6000 SP Planning Publications

This section lists the IBM product documentation for planning for the IBM RS/6000 SP hardware and software.

*IBM RS/6000 SP:*

- *Planning, Volume 1, Hardware and Physical Environment, GA22-7280*
- *Planning, Volume 2, Control Workstation and Software Environment, GA22-7281*

---

## RS/6000 SP Hardware Publications

This section lists the IBM product documentation for the IBM RS/6000 SP hardware.

*IBM RS/6000 SP:*

- *Planning, Volume 1, Hardware and Physical Environment, GA22-7280*
- *Planning, Volume 2, Control Workstation and Software Environment, GA22-7281*
- *Installation and Relocation, GA22-7441*
- *System Service Guide, GA22-7442*
- *SP Switch Service Guide, GA22-7443*
- *SP Switch2 Service Guide, GA22-7444*
- *Uniprocessor Node Service Guide, GA22-7445*
- *604 and 604e SMP High Node Service Guide, GA22-7446*
- *SMP Thin and Wide Node Service Guide, GA22-7447*
- *POWER3 SMP High Node Service Guide, GA22-7448*

---

## RS/6000 SP Switch Router Publications

| The RS/6000 SP Switch Router is based on the Ascend GRF switched IP router product from Lucent Technologies. You can order the SP Switch Router as the IBM 9077.

The following publications are shipped with the SP Switch Router. You can also order these publications from IBM using the order numbers shown.

- | • *Ascend GRF GateD Manual, GA22-7327*
- | • *Ascend GRF 400/1600 Getting Started, GA22-7368*
- | • *Ascend GRF Configuration and Management, GA22-7366*
- *Ascend GRF Reference Guide, GA22-7367*
- | • *SP Switch Router Adapter Guide, GA22-7310*

---

## **RS/6000 SP Software Publications**

This section lists the IBM product documentation for software products related to the IBM RS/6000 SP. These products include:

- IBM Parallel System Support Programs for AIX (PSSP)
- IBM LoadLeveler for AIX (LoadLeveler)
- IBM Parallel Environment for AIX (Parallel Environment)
- IBM General Parallel File System for AIX (GPFS)
- IBM Engineering and Scientific Subroutine Library (ESSL) for AIX
- IBM Parallel ESSL for AIX
- IBM High Availability Cluster Multi-Processing for AIX (HACMP)
- IBM Client Input Output/Sockets (CLIO/S)
- IBM Network Tape Access and Control System for AIX (NetTAPE)

### **PSSP Publications**

*IBM RS/6000 SP:*

- *Planning, Volume 2, Control Workstation and Software Environment, GA22-7281*

*PSSP:*

- *Installation and Migration Guide, GA22-7347*
- *Administration Guide, SA22-7348*
- *Managing Shared Disks, SA22-7349*
- *Performance Monitoring Guide and Reference, SA22-7353*
- *Diagnosis Guide, GA22-7350*
- *Command and Technical Reference, SA22-7351*
- *Messages Reference, GA22-7352*

*RS/6000 Cluster Technology (RSCT):*

- *Event Management Programming Guide and Reference, SA22-7354*
- *Group Services Programming Guide and Reference, SA22-7355*
- | • *First Failure Data Capture Programming Guide and Reference, SA22-7454*

### **LoadLeveler Publications**

*LoadLeveler:*

- *Using and Administering, SA22-7311*
- *Diagnosis and Messages Guide, GA22-7277*

### **GPFS Publications**

*GPFS:*

- | • *Problem Determination Guide, GA22-7434*
- | • *Data Management API Guide, GA22-7435*
- | • *Guide and Reference, GA22-7452*

- *Installation and Tuning Guide*, GA22-7453

### **Parallel Environment Publications**

#### *Parallel Environment:*

- *Installation Guide*, GA22-7418
- *Messages*, GA22-7419
- *DPCL Programming Guide*, SA22-7420
- *DPCL Class Reference*, SA22-7421
- *MPI Programming Guide*, SA22-7422
- *MPI Subroutine Reference*, SA22-7423
- *Hitchhiker's Guide*, SA22-7424
- *Operation and Use, Volume 1*, SA22-7425
- *Operation and Use, Volume 2*, SA22-7426
- *MPL Programming and Subroutine Reference*, GC23-3893

### **Parallel ESSL and ESSL Publications**

- *ESSL Products: General Information*, GC23-0529
- *Parallel ESSL: Guide and Reference*, SA22-7273
- *ESSL: Guide and Reference*, SA22-7272

### **HACMP Publications**

#### *HACMP:*

- *Concepts and Facilities*, SC23-4276
- *Planning Guide*, SC23-4277
- *Installation Guide*, SC23-4278
- *Administration Guide*, SC23-4279
- *Troubleshooting Guide*, SC23-4280
- *Programming Locking Applications*, SC23-4281
- *Programming Client Applications*, SC23-4282
- *Master Index and Glossary*, SC23-4285
- *HANFS for AIX Installation and Administration Guide*, SC23-4283
- *Enhanced Scalability Installation and Administration Guide*, SC23-4284

### **CLIO/S Publications**

#### *CLIO/S:*

- *General Information*, GC23-3879
- *User's Guide and Reference*, GC28-1676

### **NetTAPE Publications**

#### *NetTAPE:*

- *General Information*, GC23-3990
- *User's Guide and Reference*, available from your IBM representative

---

## AIX and Related Product Publications

For the latest information on AIX and related products, including RS/6000 hardware products, see *AIX and Related Products Documentation Overview*, SC23-2456. You can order a hard copy of the book from IBM. You can also view it online from the “AIX Online Publications and Books” page of the RS/6000 Web site at:

[http://www.rs6000.ibm.com/resource/aix\\_resource/Pubs](http://www.rs6000.ibm.com/resource/aix_resource/Pubs)

---

## DCE Publications

The DCE library consists of the following books:

- *IBM DCE 3.1 for AIX: Administration Commands Reference*
- *IBM DCE 3.1 for AIX: Administration Guide—Introduction*
- *IBM DCE 3.1 for AIX: Administration Guide—Core Components*
- *IBM DCE 3.1 for AIX: DFS Administration Guide and Reference*
- *IBM DCE 3.1 for AIX: Application Development Guide—Introduction and Style Guide*
- *IBM DCE 3.1 for AIX: Application Development Guide—Core Components*
- *IBM DCE 3.1 for AIX: Application Development Guide—Directory Services*
- *IBM DCE 3.1 for AIX: Application Development Reference*
- *IBM DCE 3.1 for AIX: Problem Determination Guide*
- *IBM DCE 3.1 for AIX: Release Notes*

You can view a DCE book or download a Portable Document Format (PDF) version of it from the IBM DCE Web site at:

<http://www.ibm.com/software/network/dce/library>

---

## Red Books

IBM's International Technical Support Organization (ITSO) has published a number of redbooks related to the RS/6000 SP. For a current list, see the ITSO Web site at:

<http://www.redbooks.ibm.com>

---

## Non-IBM Publications

Here are some non-IBM publications that you may find helpful.

- Almasi, G., Gottlieb, A., *Highly Parallel Computing*, Benjamin-Cummings Publishing Company, Inc., 1989.
- Foster, I., *Designing and Building Parallel Programs*, Addison-Wesley, 1995.
- Gropp, W., Lusk, E., Skjellum, A., *Using MPI*, The MIT Press, 1994.
- Message Passing Interface Forum, *MPI: A Message-Passing Interface Standard, Version 1.1*, University of Tennessee, Knoxville, Tennessee, June 6, 1995.
- Message Passing Interface Forum, *MPI-2: Extensions to the Message-Passing Interface, Version 2.0*, University of Tennessee, Knoxville, Tennessee, July 18, 1997.
- Ousterhout, John K., *Tcl and the Tk Toolkit*, Addison-Wesley, Reading, MA, 1994, ISBN 0-201-63337-X.
- Pfister, Gregory, F., *In Search of Clusters*, Prentice Hall, 1998.



---

# Index

## A

- abnormal error condition
  - releasing PTX shared memory manually 26
- about this book vii
- adding registered variables to a resource monitor manager 39
- adding resource variables
  - resource monitor 13
- application programming interface
  - Event Management (EMAPI) 4
  - Resource Monitor (RMAPI) 4
- application types
  - command-based resource monitor 11
  - daemon-based resource monitor 10
  - EM client 2
  - resource monitor 2
    - considerations for choosing 11
  - resource monitor added to another program 11
- audience of this book viii

## C

- callback routine (Event Management)
  - and event errors 49
  - and query errors 50
  - and unregister errors 50
  - for handling events 24, 46
  - for query responses 24, 47, 49
  - for unregister responses 24, 46
- class, resource variable
  - definition 9
  - examples 9
  - naming convention 9
  - observation interval 9
- command-based resource monitor
  - characteristics 11
  - configuring 11
  - pseudocode example 14
- commands
  - haemcfg
    - configuring Event Management 12
  - haemqvar
    - finding out what resource data is available 18
  - ipcrm
    - using to release PTX shared memory manually 26
  - SDRCreateObjects
    - configuring Event Management 12
  - spevent
    - getting resource variable information 18
    - output example 19
- communication across nodes
  - Event Management 3
- components
  - Event Management 1
- concepts
  - Event Management 1
- configuration
  - command-based resource monitor 11

- configuration (*continued*)
  - daemon-based resource monitor 11
  - Event Management 12
    - haemcfg command 12
    - SDR classes 12
    - use of haemloadcfg command 12
    - use of SDRCreateObjects command 12
    - use of System Data Repository (SDR) 12
  - resource monitor logic 11
  - SDR data for EMAPI 28

- considerations
  - for choosing a resource monitor type 11
- control message commands
  - HA\_RR\_CMD\_ADDALL 15
  - HA\_RR\_CMD\_ADDV 15
  - HA\_RR\_CMD\_DELALL 15
  - HA\_RR\_CMD\_DELV 15
  - HA\_RR\_CMD\_INSTV 15

- Counter value type
  - and class, resource variable 9
  - and Event Management subsystem 6
  - and Performance Monitor subsystem 6
  - and PTX objects 6
  - and PTX shared memory 25
  - definition 6
  - rate 6
  - raw value 6

## D

- daemon-based resource monitor
  - characteristics 10
  - configuring 11
  - pseudocode example 16
- data types, Event Management
  - float 7
  - long 7
  - structured byte string (SBS) 7
- default expression
  - definition 22
- defining expressions
  - examples 21
  - rules 21
- defining resource data to be monitored 5
- deleting resource variable values
  - resource monitor 13
- deleting resource variables from a resource monitor manager session 40
- domains
  - support by Event Management 3
- dynamic data-supplier
  - definition (from PTX) 25
- dynamic instantiation
  - definition 9

## E

- EM client
  - application types 2
  - definition 2
  - diagram 2
  - ending a session 25
  - function 2
  - multiple sessions 4
  - receiving responses 24
  - registering for events 23
  - sending event management queries 24
  - SIGPIPE signal handling 64
  - starting a session 23
  - tasks for writing
    - coding and testing 22
    - defining expressions 20
    - finding out what resource data is available 17
    - overview 17
  - unregistering for events 24
- EM\_Resource\_Class SDR class 35
- EM\_Resource\_ID SDR class 35
- EM\_Resource\_Monitor SDR class 36
- EM\_Resource\_Variable SDR class 30
- EM\_Structured\_Byte\_String SDR class 34
- EMAPI (Event Management Application Programming Interface)
  - configuration data (SDR) 28
  - error reference 104
  - expression reference 123
  - resource variable reference 126
  - use by EM client 2
  - using 4
- EMAPI subroutines
  - ha\_em\_get\_ecgid subroutine 24
  - ha\_em\_receive\_response subroutine 24
  - ha\_em\_send\_command subroutine
    - HA\_EM\_CMD\_QUERY command 24
    - HA\_EM\_CMD\_REG command 23
    - HA\_EM\_CMD\_REG2 command 23
    - HA\_EM\_CMD\_UNREG command 24
  - ha\_em\_start\_session subroutine 23
  - ha\_rr\_end\_session subroutine 25
- ending a session
  - EM client 25
  - resource monitor 13, 15
- ending a session with a resource monitor manager 40
- ending a session with Event Management 39
- error condition
  - releasing PTX shared memory manually 26
- error responses (EM client)
  - response block 24
- errors
  - EMAPI 104
  - RMAPI 131
- establishing a session with Event Management 39
- establishing the resource monitor as a server 39
- event
  - definition 1, 21
- event command group
  - definition 24
- event expression
  - definition 22
- event generation
  - description 20
  - example 21
- Event Management
  - application programming interfaces 4
  - communication across nodes 3
  - component
    - EM client 2
    - Event Management subsystem 2
    - resource monitor 1
  - configuration 12
  - domains 3
  - event generation 20
  - graphical user interface 4
  - introduction 1
  - multiple sessions 4
  - name spaces 10
  - performance considerations 25
  - SDR classes
    - EM\_Resource\_Class 35
    - EM\_Resource\_ID 35
    - EM\_Resource\_Monitor 36
    - EM\_Resource\_Variable 30
    - EM\_Structured\_Byte\_String 34
  - software components 1
  - system partitioning 3
- Event Management components
  - diagram 2
- Event Management configuration
  - haemcfg command 12
  - SDR classes 12
  - use of haemloadcfg command 12
  - use of SDRCreateObjects command 12
  - use of System Data Repository (SDR) 12
- Event Management Configuration Database (EMCDB)
  - definition 12
  - refreshing 13
  - version number 13
- event management subroutines
  - ha\_em\_end\_session 41
  - ha\_em\_get\_ecgid 43
  - ha\_em\_receive\_response 44
  - ha\_em\_restart\_session 53
  - ha\_em\_send\_command 56
  - ha\_em\_start\_session 63
  - ha\_rr\_add\_var 66
  - ha\_rr\_del\_var 69
  - ha\_rr\_end\_session 72
  - ha\_rr\_get\_ctrlmsg 74
  - ha\_rr\_get\_interval 78
  - ha\_rr\_init 80
  - ha\_rr\_makserv 83
  - ha\_rr\_reg\_var 85
  - ha\_rr\_rm\_ctl 88
  - ha\_rr\_send\_val 91
  - ha\_rr\_start\_session 93
  - ha\_rr\_terminate 96
  - ha\_rr\_touch 98
  - ha\_rr\_unreg\_var 100

- Event Management subroutines
  - ending a session 39
  - establishing a session 39
  - get an event command group ID 39
  - receiving a response from the Event Management subsystem 39
  - restarting a session 39
  - send a command to the Event Management subsystem 39
- Event Management subsystem
  - concepts 1
  - definition 2
  - diagram 2
  - function 2
  - summary of subroutines 39
- Event Management terms
  - Counter value type 6
  - default expression 22
  - dynamic instantiation 9
  - EM client 2
  - event 1, 21
  - event command group 24
  - event expression 22
  - Event Management subsystem 2
  - expression 2, 20
  - float data type 7
  - instance 1, 7
  - location, resource variable instance 9
  - long data type 7
  - observation 2
  - Quantity value type 6
  - raw value 6, 124
  - rearm expression 22
  - resource 5
  - resource attribute 6
  - resource ID 8
  - resource ID element 8
  - resource ID element name 8
  - resource ID element value 8
  - resource monitor 1
  - resource name 6
  - resource variable 1, 5
  - resource variable class 9
  - resource variable description 6
  - resource variable name 6
  - serial number, structured byte string 7
  - State value type 6
  - structured byte string (SBS) data type 7
  - structured field 7
  - system resource 1
- Event Perspective 4
- event responses
  - callback routine 24
  - response block 24
- examples
  - class, resource variable 9
  - event generation 21
  - event management example
    - emapi\_v02\_ex01.c 181
    - emapi\_v02\_ex02.c 194
    - emapi\_v02\_ex03.c 206

- examples (*continued*)
  - event management example (*continued*)
    - emapi\_v02\_ex04.c 219
  - expressions 21
  - location, resource variable instance 9
  - rearm expression 22
  - resource data
    - using SP Perspectives 19
  - resource ID 8
  - resource monitor example
    - rmapi\_smp.loadsdr shell script 174
    - rmapi\_smp.msg message file 173
    - rmapi\_smp.unloadsdr shell script 178
    - rmapi\_smpcmd.c 141
    - rmapi\_smpdae.c 148
    - rmapi\_smpsig.c 161
  - resource monitor pseudocode
    - command-based resource monitor 14
    - daemon-based resource monitor 16
  - resource variable name 8
  - resource variables 5
  - structured bytes string (SBS) definition 7
  - system resources 5
- expression
  - creating 20
  - definition 2, 20
  - EMAPI 123
- expressions
  - examples 21
  - rules for defining 21

## F

- finding out what resource data is available
  - using the haemqvar command 18
- float data type
  - definition 7
- freeing resources and terminating use of the RMAPI 40

## G

- getting a control message
  - resource monitor 14
- getting a control message from a resource monitor manager 39
- getting an event command group ID 39
- getting resource variable information
  - from SP Perspectives 18
- getting the reporting interval for a class of resource variables 39
- graphical user interface
  - Event Management 4

## H

- ha\_em\_end\_session subroutine
  - ending a session 25
  - reference 41
- ha\_em\_get\_ecgid subroutine
  - getting the event command group 24

- ha\_em\_get\_ecgid subroutine (*continued*)
  - reference 43
- ha\_em\_receive\_response subroutine
  - receiving responses 24
  - reference 44
- ha\_em\_restart\_session subroutine
  - reference 53
- ha\_em\_send\_command subroutine
  - reference 56
  - registering for events
    - HA\_EM\_CMD\_REG command 23
    - HA\_EM\_CMD\_REG2 command 23
    - HA\_EM\_SCMD\_REVAL subcommand 24
  - sending event management queries
    - HA\_EM\_CMD\_QUERY command 24
  - unregistering for events
    - HA\_EM\_CMD\_UNREG command 24
- ha\_em\_start\_session subroutine
  - reference 63
  - starting a session 23
- ha\_emoji\_base.h header file
  - reference 114
- ha\_emoji.h header file
  - reference 110
- ha\_emcommon.h header file
  - reference 119
- ha\_rmapi.h header file
  - reference 136
- ha\_rr\_add\_var subroutine
  - adding resource variables 13
  - reference 66
- HA\_RR\_CMD\_ADDALL command 15
- HA\_RR\_CMD\_ADDV command 15
- HA\_RR\_CMD\_DELALL command 15
- HA\_RR\_CMD\_DELV command 15
- HA\_RR\_CMD\_INSTV command 15
- ha\_rr\_del\_var subroutine
  - deleting resource variable values 13
  - reference 69
- ha\_rr\_end\_session subroutine
  - ending a session 13, 15
  - reference 72
- ha\_rr\_get\_ctrlmsg subroutine
  - getting a control message 14
  - reference 74
- ha\_rr\_get\_interval subroutine
  - reference 78
- ha\_rr\_init subroutine
  - initializing 13, 14
  - reference 80
- ha\_rr\_makserv subroutine
  - making a server session 14
  - reference 83
- ha\_rr\_reg\_var subroutine
  - reference 85
  - registering resource variables and instances 13, 14
- ha\_rr\_rm\_ctl subroutine
  - reference 88
- ha\_rr\_send\_val subroutine
  - reference 91
  - sending resource variable values 13
- ha\_rr\_start\_session subroutine
  - reference 93
  - starting a session 13, 14
- ha\_rr\_terminate subroutine
  - and PTX shared memory 25
  - reference 96
  - terminating 14, 15
- ha\_rr\_touch subroutine
  - reference 98
- ha\_rr\_unreg\_var subroutine
  - reference 100
- haemcfg command
  - Event Management configuration 12
- haemloadcfg command
  - use for Event Management configuration 12
- haemloadlist file 12
- haemqvar command
  - finding out what resource data is available 18
- header files
  - ha\_emoji\_base.h 114
  - ha\_emoji.h 110
  - ha\_emcommon.h 119
  - ha\_rmapi.h 136

## I

- initializing
  - resource monitor 13, 14
- initializing RMAPI 39
- instance
  - definition 7
- instantiation, dynamic
  - definition 9
- ipcrm command
  - using to release PTX shared memory manually 26

## K

- killing a resource monitor
  - releasing PTX shared memory manually 26

## L

- loadlist file 12
- local data-consumer program
  - definition (from PTX) 25
- location, resource variable instance
  - definition 9
  - examples 9
- logic, resource monitor
  - characteristics 11
  - configuring 11
- long data type
  - definition 7

## M

- making a server session
  - resource monitor 14
- manual pages for public code 240

- maximum
  - number of resource ID elements 128
- meeting "send" frequency requirements 40
- multiple sessions
  - Event Management 4

## N

- name spaces
  - Event Management 10
- naming convention
  - class, resource variable 9
  - resource attribute 6
  - resource variable name 6

## O

- observation
  - definition (Event Management) 2
- observation, previous
  - specifying in an expression 21
- observation interval
  - and resource variable class 9
  - and resource variable value type 9

## P

- performance
  - of Event Management operation 25
- Performance Monitor subsystem
  - and Counter value type 6
  - and Quantity value type 6
  - and State value type 6
- prerequisite knowledge for this book viii
- processing a control message
  - resource monitor 14
- PTX object
  - and Counter value type 6
  - and Quantity value type 6
- PTX shared memory
  - and Quantity and Counter value types 25
  - and resource monitors 25
  - releasing manually 26

## Q

- Quantity value type
  - and class, resource variable 9
  - and Event Management subsystem 6
  - and Performance Monitor subsystem 6
  - and PTX objects 6
  - and PTX shared memory 25
  - definition 6
- query responses
  - callback routine 24
  - response block 24

## R

- raw value 124
  - definition 6
  - specifying in an expression 21
- rearm expression
  - definition 22
  - example 22
  - using for inverse conditions 22
  - using for upper and lower boundaries 22
- receiving a response from Event Management 39
- receiving responses
  - EM client 24
- registering a resource variable instance with the RMAPI 39
- registering for events
  - EM client 23
- registering resource variables
  - resource monitor 13, 14
- resource
  - definition 5
- resource attribute
  - definition 6
- resource data
  - defining what will be monitored 5
  - finding out what is available
    - using SP Perspectives 18
    - using the haemqvar command 18
- resource ID
  - definition 8
  - examples 8
- resource ID element
  - definition 8
  - wildcards 8
- resource ID element name
  - definition 8
- resource ID element value
  - definition 8
- resource monitor
  - abnormal termination
    - releasing PTX shared memory manually 26
  - added to another program 11
  - adding resource variables 13
  - and PTX shared memory 25
  - application types 2
  - command-based 11
  - control message commands
    - HA\_RR\_CMD\_ADDALL 15
    - HA\_RR\_CMD\_ADDV 15
    - HA\_RR\_CMD\_DELLALL 15
    - HA\_RR\_CMD\_DELV 15
    - HA\_RR\_CMD\_INSTV 15
  - daemon-based 10
  - definition 1
  - deleting resource variable values 13
  - diagram 2
  - ending a session 13, 15
  - examples
    - pseudocode for command-based resource monitor 14
    - pseudocode for daemon-based resource monitor 16

- resource monitor (*continued*)
  - function 1
  - getting a control message 14
  - initializing 13, 14
  - making a server session 14
  - processing a control message 14
  - registering resource variables and instances 13, 14
  - sending resource variable values 13
  - starting a session 13, 14
  - tasks for writing
    - choosing the resource monitor type 10
    - coding and testing 13
    - configuring 12
    - defining the resource data 5
    - overview 5
  - terminating 14, 15
  - testing 13
- resource name
  - definition 6
- resource variable
  - definition 1, 5
  - instance 1
  - naming convention 6
- resource variable description
  - definition 6
- resource variable information, displaying
  - using the `spevent` command 18
- resource variable name
  - definition 6
  - examples 8
- resource variables
  - dynamically instantiable 9
  - EMAPI 126
  - RMAPI 126
- response block (Event Management)
  - for error responses 24
  - for handling events 24
  - for query responses 24
  - for unregister responses 24
- restarting a session with Event Management 39
- RMAPI (Resource Monitor Application Programming Interface)
  - error reference 131
  - pseudocode for command-based resource monitor 14
  - pseudocode for daemon-based resource monitor 16
  - resource variable reference 126
  - summary of subroutines 39
  - use by resource monitor 2
  - using 4
- RMAPI control message commands
  - HA\_RR\_CMD\_ADDALL command 15
  - HA\_RR\_CMD\_ADDV command 15
  - HA\_RR\_CMD\_DELALL command 15
  - HA\_RR\_CMD\_DELV command 15
  - HA\_RR\_CMD\_INSTV command 15
- RMAPI subroutines
  - adding registered variables to a resource monitor manager 39
  - deleting resource variables from a resource monitor manager session 40

- RMAPI subroutines (*continued*)
  - ending a session with a resource monitor manager 40
  - establishing the resource monitor as a server 39
  - freeing resources and terminating use of the RMAPI 40
  - getting a control message from a resource monitor manager 39
  - getting the reporting interval for a class of resource variables 39
  - ha\_rr\_add\_var subroutine 13
  - ha\_rr\_del\_var subroutine 13
  - ha\_rr\_end\_session subroutine 13, 15
  - ha\_rr\_get\_ctrlmsg subroutine 14
  - ha\_rr\_init subroutine 13, 14
  - ha\_rr\_makserv subroutine 14
  - ha\_rr\_reg\_var subroutine 13, 14
  - ha\_rr\_send\_val subroutine 13
  - ha\_rr\_start\_session subroutine 13, 14
  - ha\_rr\_terminate subroutine 14, 15
  - initializing 39
  - meeting "send" frequency requirements 40
  - registering a resource variable instance with the RMAPI 39
  - sending variable values to the RMAPI 40
  - setting or getting RMAPI attributes 39
  - starting a session with a resource monitor manager 39
  - unregistering a resource variable instance 40

## S

- SDR (System Data Repository)
  - classes for Event Management configuration 12
  - EM configuration data 28
  - use for Event Management configuration 12
- SDR classes
  - Event Management
    - EM\_Resource\_Class 35
    - EM\_Resource\_ID 35
    - EM\_Resource\_Monitor 36
    - EM\_Resource\_Variable 30
    - EM\_Structured\_Byte\_String 34
- SDRCreateObjects command
  - configuring Event Management 12
- sending a command to Event Management 39
- sending event management queries
  - EM client 24
- sending resource variable values
  - resource monitor 13
- sending variable values to the RMAPI 40
- serial number, structured byte string
  - definition 7
- session with Event Management
  - ending 39
  - restarting 39
  - starting 39
- setting or getting RMAPI attributes 39
- signal handling
  - EM client 64

- SIGPIPE signal
  - handling for EM clients 64
- SP Perspectives
  - example of displaying resource data 19
  - getting resource variable information 18
  - using for Event Management 4
- spevent command
  - example of output 19
  - getting resource variable information 18
- starting a session
  - EM client 23
  - resource monitor 13, 14
- starting a session with a resource monitor manager 39
- State value type
  - and Event Management subsystem 6
  - and Performance Monitor subsystem 6
  - definition 6
- structured byte field
  - specifying in an expression 21
- structured byte string (SBS)
  - example 7
- structured byte string (SBS) data type
  - definition 7
- structured field
  - definition 7
- summary of subroutines
  - Event Management 39
  - Resource Monitor API 39
- system partition
  - support by Event Management 3
- system resource
  - definition 1

## T

- tasks
  - writing EM clients
    - coding and testing 22
    - defining expressions 20
    - finding out what resource data is available 17
    - overview 17
  - writing resource monitors
    - choosing the resource monitor type 10
    - coding and testing 13
    - configuring 12
    - defining the resource data 5
    - overview 5
- terminating
  - resource monitor 14, 15
  - use of the RMAPI 40
- termination, abnormal
  - resource monitor
    - releasing PTX shared memory manually 26
- terms, Event Management
  - Counter value type 6
  - default expression 22
  - dynamically instantiable resource variables 9
  - EM client 2
  - EMCDB (Event Management Configuration Database) 12
  - event 1, 21

- terms, Event Management (*continued*)
  - event command group 24
  - event expression 22
  - Event Management subsystem 2
  - expression 2, 20
  - float data type 7
  - instance 1, 7
  - location, resource variable instance 9
  - long data type 7
  - observation 2
  - Quantity value type 6
  - raw value 6, 124
  - rearm expression 22
  - resource 5
  - resource attribute 6
  - resource ID 8
  - resource ID element 8
  - resource ID element name 8
  - resource ID element value 8
  - resource monitor 1
  - resource name 6
  - resource variable 1, 5
  - resource variable class 9
  - resource variable description 6
  - resource variable name 6
  - serial number, structured byte string 7
  - State value type 6
  - structured byte string (SBS) data type 7
  - structured field 7
  - system resource 1
- testing
  - new resource monitor 13
- trademarks 228

## U

- unregister responses
  - callback routine 24
  - response block 24
- unregistering a resource variable instance 40
- unregistering for events
  - EM client 24
- using application programming interfaces
  - EMAPI 4
  - RMAPI 4
- using graphical user interfaces
  - SP Perspectives 4

## V

- value types, Event Management
  - Counter 6
  - Quantity 6
  - State 6
- version number
  - EMCDB 13

## W

- wildcard
  - resource ID element 8



---

# Readers' Comments — We'd Like to Hear from You

RS/6000 Cluster Technology  
Event Management Programming Guide and Reference

Publication No. SA22-7354-01

Overall, how satisfied are you with the information in this book?

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Overall satisfaction	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

How satisfied are you that the information in this book is:

	Very Satisfied	Satisfied	Neutral	Dissatisfied	Very Dissatisfied
Accurate	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Complete	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to find	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Easy to understand	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Well organized	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
Applicable to your tasks	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>

Please tell us how we can improve this book:

Thank you for your responses. May we contact you?  Yes  No

When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you.

---

Name

---

Address

---

Company or Organization

---

---

Phone No.

---



Fold and Tape

Please do not staple

Fold and Tape




NO POSTAGE  
NECESSARY  
IF MAILED IN THE  
UNITED STATES

# BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 40 ARMONK, NEW YORK

POSTAGE WILL BE PAID BY ADDRESSEE

IBM Corporation  
Department 55JA, Mail Station P384  
2455 South Road  
Poughkeepsie, NY 12601-5400



Fold and Tape

Please do not staple

Fold and Tape





Program Number: 5765-D51 (PSSP); 5765-D28 (HACMP)



Printed in the United States of America  
on recycled paper containing 10%  
recovered post-consumer fiber.

SA22-7354-01

