E-SRF

EKC Security
Event Reporting Facility

Release 2.1
User Guide

E-SRF V2R1 – **GENERAL AVAILABILITY, Revised February 13, 2005**

EKC Inc.

E-SRF™ is a proprietary product
developed and maintained by

EKC Inc.
10400 West Higgins Road
Rosemont, Illinois 60018
USA

(847) 296-8010

Technical Support:
(847) 296-8035

EKC, Inc. provides only software program products, which fully comply with, and maintain MVS integrity.

The vendor hereby warrants that:

1) E-SRF™ ("Software") performs only those functions which are described in the published specifications;

2) there are no methods for gaining access to the Software or other computer resources or data of Licensee (such as a master access key, ID, password, or trap door) other than set forth in the published specifications;

3) the Software does not introduce any MVS integrity exposures. The program code, with the exception of one utility, runs totally in non-authorized, problem state. The one utility, EKCRXCAT, requires APF-authorization to read the MVS System Catalogs. A non-APF authorized utility, EKCRGCAT, is supplied to perform the same function, but at a considerably slower speed.

4) the software shall be year 2000 compliant, and shall function correctly in the next century according to published specifications as long as regular software maintenance is applied.

Version 2, Release 1 February 13, 2005, (Revised for: LE00450)

Contents

E-SRF Publications

Name	Contents
Installation Guide	E-SRF installation including: installation and maintenance steps, startup and shutdown considerations, and backup and recovery procedures.
Change Summary Guide	Contains all new features and system function changes.
General Overview	An overview of E-SRF and its components.
Resource Grouping Facility Guide	Brief overview of the Resource Grouping Facility, its relationship to E-SRF, language command syntax, TSO commands and JCL.
Access Analysis Reports Guide for ACF2 *Access Analysis Reports Guide for RACF*	Brief overview of Access Analysis reports, explanation of the DataOwner and Userid/LogonidOwner reports, command syntax, utilities necessary for creating input to reports, and JCL.
Event Reporting User Guide	A “How To” guide for users of E-SRF Event Reporting.
Event Reporting Facility - Command Reference	Explains the Event Reporting Facility command processor, command syntax, and JCL.
Event Reporting Facility - Masterfile and Data Dictionary Reference	Explains the structure of the E-SRF Masterfile and describes all Masterfile fields.
Event Reporting Facility - Messages and Codes	Lists Event Reporting Facility messages and codes.
Event Reporting Facility - Report Overlays Guide	An overview of the report overlays provided with the Event Reporting Facility.

This page intentionally left blank

Chapter 1: Introduction

The Reporting "Wish List"

There are many resources within your security environment that require protection from misuse or destruction. Security administrators spend a great deal of time evaluating this protection. For example, access must be controlled when a dataset is initially created and again if a change is required. There are many other resources besides datasets contained in a typical system, including CICS or IMS transactions, terminals, DASD volumes, tape volumes, TSO commands, etc. Who should have access to these resources and to what extent? What reporting mechanism identifies activity on these resources? These are the questions to which security administrators need answers.

Have you ever wanted the ability to assign "owners" to resources?

The definition of an "owner" is an individual who is *responsible* for the resource. Owners know who needs access to their resource(s) and when it is required, as well as an understanding of how the resource is used or referenced. Owners can then assist the security administrator in determining what access is allowed.

Has management ever needed to know what resources a specific user or groups of users have access to?

Within the structure of the current resident security systems (CA-ACF2, CA-Top Secret and IBM RACF), group accesses and extensive resource rules make it difficult to determine an individual's resource access. An automated tool that quickly produces this information is a valuable time-saver.

Have you ever wanted a better reporting facility to alert security administrators, the "owners" of resources, and any other interested parties of security violations or other security loggings?

EKC offers a reporting facility that can interpret security event data, be customized to make these reports meaningful for an installation, and alert owners to an inappropriate access or an attempted access violation to their resource. In addition, these reports can be automatically distributed.

These items are on the *"wish list*" of many security professionals. The security systems today have little, if any, means by which you can associate a "*true*" owner with a resource. They may provide the ability to assign an owner to a resource profile or within a security rule, but not to the resource itself. In addition, significant and succinct event reporting is not easily obtained. The E-SRF product provides all of this and more.

Presenting: E-SRF

Resident security systems (RSS) currently on the market secure your system but are lacking in administration tools, accountability of resources, reporting, and report distribution. The **EKC Security Reporting Facility** focuses on these areas.

Major Functions

E-SRF security reports provide the comprehensive information needed by the security administrator to respond to system events. They are organized in a manner that allows for quick and easy interpretation. Once it is determined which reports best meet an installation's needs, they can be scheduled to run at regular intervals (daily, weekly, etc.). This eases administrative burden while still providing timely access to the reports.

Rather than only providing for formatted reports, E-SRF produces "intelligent" reporting; that is, E-SRF reporting decisions are based on the relevance of the data contained within the report.

The E-SRF product is an essential tool for the security administrator, auditor, and manager in order to assess the effectiveness of their security systems. In addition, E-SRF enables a proactive approach to adjusting access as changes in that system occur in order to maintain its integrity.

EKC Security Reporting Facility

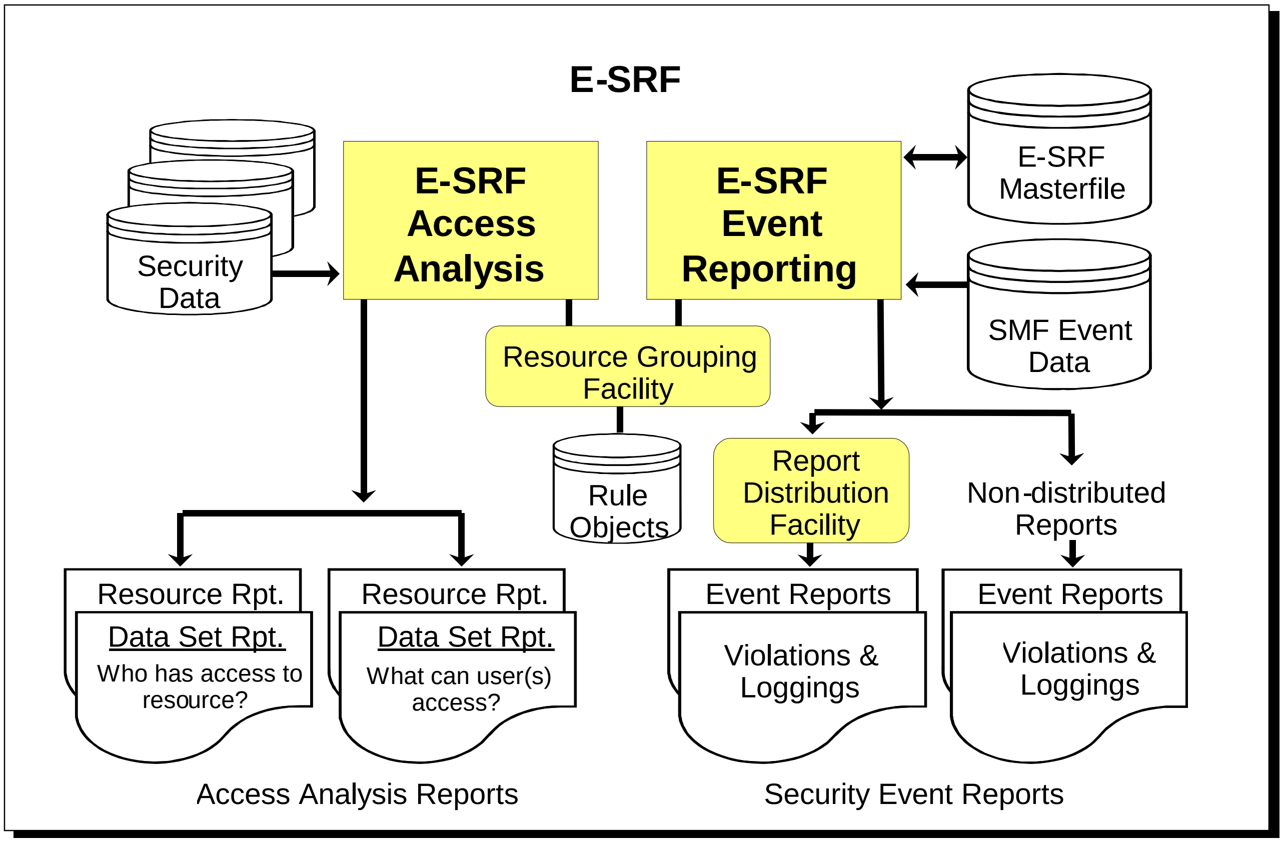

The EKC Security Reporting Facility (E-SRF) major components are:

The **Access Analysis Reporting Facility** is a robust reporting facility that can report against your security system's access information from several perspectives. These reports answer the questions: "Who can access a resource/dataset?" and "What resources/datasets can a specific user or users access?" In addition to answering the questions of who has access, the comprehensive Access Analysis Reports include such information as who can gain access due to privileges instead of a rule or profile. This is an invaluable tool for resource "owners."

The **Event Reporting Facility** has three major functions. First, it is a powerful mechanism that can report on security events and loggings in a comprehensive yet concise manner. Second, it allows you to define owners to resources. Finally, it automates distribution of these reports to the owners of resources, security personnel, or any other interested party.

The **Resource Grouping Facility** provides the ability to associate a group name to a single resource or group of resources. Resource types include datasets, CICS and IMS transactions, terminals, tapes, DASD volumes, and many more. Users may also be classified as a resource type.

A summary of E-SRF's components and their major functions follows in Chapter 2.

Chapter 2: E-SRF Components

This chapter highlights the components of E-SRF.

- EKC Integrated Grouping Facility provides you with the option to group resources and datasets for reporting.
- Access Analysis Reports answer key questions about who has access to critical business data and resources.
- Event Reports summarize actual security events.

What is a "resource"?

By definition, a "resource" is a name of any item for which your resident security system is providing protection and accountability.

Normally, a typical data processing center has two main types of resources: datasets (files that contain data) and non-dataset resources (such as CICS and IMS transactions and other "labels" placed on computer resources). Most security systems, as well as E-SRF Access Analysis, make a distinction between dataset and non-dataset resources.

However, E-SRF Event System does not make a distinction between the two. All resources will be assigned to their respective classes (e.g., dataset resources will be in the DATASET class). As far as grouping is concerned, even a USERID is considered to be a "resource;" the "classname" might be designated as USER, and the "resource name" would be the User's identification.

Resource Grouping Facility

The EKC Integrated Resource Grouping Facility is a common utility component that is bundled into E-SRF along with the other two separate components, Access Analysis and Event Reporting. It is used to *dynamically* associate a group name to one or more resources on demand. This is a key component of the system and allows for automatically distributing security event reports to interested parties within your organization. Access Analysis and Event Reporting share this facility. Each of these components may use this facility for processing requests by the user in whatever manner is required.

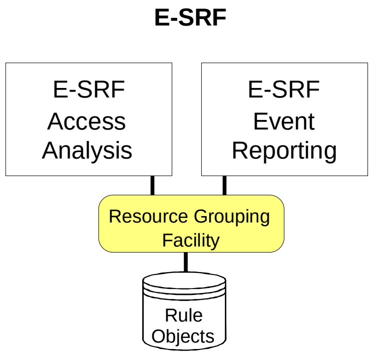

Utilizing this grouping facility is optional by one or both reporting components, especially in the early stages of product implementation. *This means you do not have to have any type of grouping established to initially use E-SRF until you so desire to incorporate this facility.* New users to Event Reporting are recommended NOT to attempt grouping until after they have run reports and are familiar with the basic reporting components of the E-SRF system. As product usage matures, the use of grouping will especially enhance Event Reporting. It can provide the means for Automated Report Distribution to various data Owners within your organization.

The Resource Grouping Facility uses "*Grouping Rules*" to associate a group name to a resource(s) for reporting purposes. These grouping rules are not to be confused with any type of *rules* that you may already have in place for your Resident Security System.

"Rules" are written and stored on a separate file in order to maintain multiple grouping schemes within a single organization. The EKC Integrated Resource Grouping Facility provides a mechanism for grouping both dataset and non-dataset resources using separate schemes. Resources of differing types can be grouped together under a single group name.

These group names are retained in the Rule Objects file for subsequent analysis and reporting. The group name associated with a resource(s) is retrieved from the Resource Grouping Facility and is supplied to the requesting E-SRF component.

A "*resource*" is a dataset or other entity (such as a transaction, command, terminal, or userid) that you have protected. To help determine what resources may need to be grouped, Access Analysis utilities can be used to identify your dataset and resource names. It accomplishes this by either reporting the dataset name information obtained from your system catalogs or by reporting on the security definitions, which exist in your Resident Security System.

For example, if you want to create a report for the Payroll Manager to review, you can group all payroll datasets and other resources into a group called PAYROLL. By specifying to E-SRF that you want a report for the PAYROLL group, it will evaluate only those dataset and other resource security definitions or events that fall within the PAYROLL group.

The Resource Grouping Facility allows "grouping" of various resource types within a single group name. As an example, datasets PROD.TECH1, PROD.TECH2 and CICS transaction resources TR01, TR02, and TR03 can all be grouped under group name "MVSTECH." In addition, grouping is dynamic and can be changed whenever the installation makes the determination that the associated group names are no longer appropriate.

More information on why to group, how to group, and the benefits of grouping are covered throughout this publication. The *Resource Grouping Facility Guide* provides detailed information on how to set up E-SRF grouping in your installation.

E-SRF Access Analysis

The Access Analysis Reporting Facility provides several reports that will analyze and summarize access to resources. Security Professionals need these reports to identify the current levels of security controls; i.e., who has access and to what extent.

The first set of reports identifies resource or dataset access and identifies who has access to a specific resource/dataset. These reports answer the question, "Who is accessing and changing my data?" The second set of reports identifies user access, specifically presenting the resources a particular user or groups of users can access. These reports answer the question, "What critical business information are my employees able to gain access to and at what level of access (i.e. read/update)?" They are very helpful to both security administrators and resource owners, who will also be alerted to any access changes made.

The Access Analysis Facility of E-SRF evaluates the security controls in place for a set of users or a set of resources. This is accomplished by reviewing the actual security definitions, evaluating and comparing information, and presenting a matrix of the data, users, and types of accesses allowed.

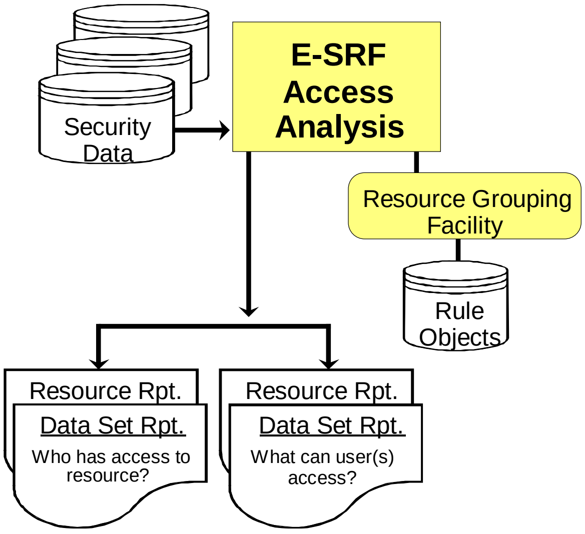

Access Analysis Reports

Important Note:

Access analysis reporting was designed for specific security systems and reporting is done on demand. No data is stored beyond the execution of the desired Access Analysis report.

The E-SRF Event Reporting Masterfile is **NOT** used for E-SRF Access Analysis and therefore does not have to be available (or even exist) in order to produce Access Analysis reports.

The information about E-SRF Access Analysis reporting provided in this publication is just an introduction to its features. To learn more about E-SRF Access Analysis, please refer to the specific E-SRF Access Analysis publications that relate to your resident security system.

DataOwner Report

DataOwner Reports answer the question, "Who has the ability to read or update my data or resources?"

E-SRF evaluates the security access definitions, compares them to the users defined to the security system, and determines who has access.

The DataOwner Report can be run for datasets or other defined secured resources. Data owners benefit from this report because it aids in evaluating any exposures to “their” data by users with the currently defined access controls.

Group, High Level Qualifier, or a list of datasets can be specified when requesting a DataOwner Report. This report is an especially useful tool to assist managers and auditors in investigating security breaches. Which users have the capability to see, update, or delete the data and resources anywhere on your system may be quickly and easily determined.

LogonidOwner (ACF2) or UseridOwner (RACF) Reports

LogonidOwner and UseridOwner Reports answer the question, “What can my users do?”

These reports provide a different perspective. For ACF2 reports, a group of users is specified based on their Logonid attributes. For RACF reports, a group of users is specified based on their User Profile attributes or RACF Groups to which they are connected.

E-SRF evaluates all of the security rules either for datasets or other resources and determines what accesses those users have. The LogonidOwner and UseridOwner Reports allow you to make quick evaluations (e.g., a user with too much access or users who are able to change information beyond their normal day-to-day requirements).

In cases where auditors or others are investigating a security breach, the LogonidOwner Reports may conversely assist in eliminating possible suspects since it becomes readily apparent what access users have (or don’t have).

ACF2 Specific Reports

For ACF2, E-SRF also provides the Proposed Rule Processor and the System Differences Reports. The combination of these two reports allows the Security Administrator to determine what different access users will have after a set of *proposed* ACF2 Access and Resource Rules are implemented. This eliminates much guesswork and manual analysis.

Additionally for ACF2, E-SRF provides a Database Differences Report. This report shows the changes made to the ACF2 Security Database between two points in time. This might be between yesterday and today, or between last week and now, or between last month and now. The Database Differences Report provides Security Administrators a needed tool to simply audit changes in Logonid definitions or access control rules and definitions.

These reports answer the questions:

- “What are the differences in access for two different databases?”
- “If I make a change to the rules, how can I be sure that the different access allowed will be what I expected?”
- “What portions of the database changed between yesterday (or last week) and today?”
- “What logonid record fields, rule lines, and information storage records are different between two separate databases?”

RACF Specific Reports

E-SRF also provides additional reports for RACF. The Userid Differences Report analyzes the different accesses users have and groups them. The differences between the access for each group are detailed if they are within a certain set of limits specified by the Security Administrator. This provides the Security Administrator the ability to combine RACF groups or to highlight a user who has a slightly different access than his or her colleagues.

Additionally for RACF, E-SRF provides a DataOwner Open Edition Report supporting the OS/390 UNIX System Services portion of the Operating System.

Note: As of this release, the Access Analysis reports *do not* reference the E-SRF Masterfile definitions for Data Owners.

For more detailed information on Access Analysis, see the *ACF2 Access Analysis Reports Guide and the RACF Access Analysis Reports Guide.*

E-SRF Event Reporting

Event Reporting assists security administrators, auditors, managers, and owners of resources in evaluating the events that have occurred in their system. These events are recorded by the Resident Security System (RSS) as a result of access activity and access logging specifications.

While Resident Security Systems provide security event reporting capabilities, the information is presented in a manner that can be very difficult to understand, evaluate, and distribute to appropriate personnel. A need was identified for reports that are clear, concise, and summarize the important information for easy evaluation.

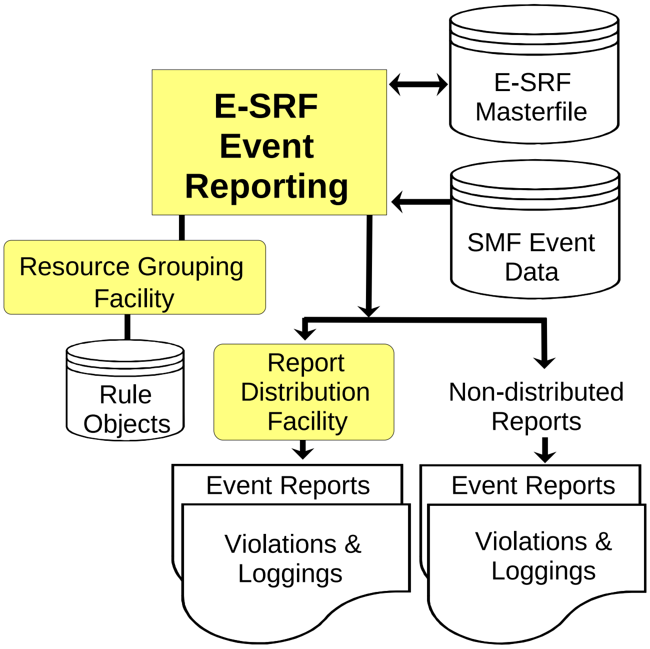

Security Event Reports

The E-SRF Event Reporting Facility was designed with that objective in mind. This Facility evaluates data provided by the RSS, updates its Masterfile, and can produce a wide variety of security reports based upon this data. Event reports may contain either detailed event information and/or summary data and include the necessary information to determine who is accessing key resources and if the level of security is appropriate.

Some of these Event Reports are customizable, while others cannot be modified. Fixed-format reports provide a consistent and reliable form for reporting event data. Customizable reports provide flexibility in both the information and the format in which it will be presented to best suit the needs of the installation.

In order to obtain the *group names* associated with a resource or user, Event Reporting optionally utilizes the Resource Grouping Facility. (Defining group names through Grouping Rules is explained in the E-SRF *Resource Grouping Facility Guide*.) Group names may represent one or more resource and will be reported along with all other pertinent access information.

The Event Reporting Facility additionally allows you to define *group headers* to specify a group's characteristics to the Masterfile, including the group *owner* (identifies who is *responsible* for the group), as well as up to 16 possible *interested parties* (for example: managers, resource owners, security personnel, and auditors).

Owners and defined interested parties will receive reports regarding resources that have been accessed. E-SRF may be set up to provide the option of automatically distributing copies of these reports.

Security event journalized information, as well as additional information about users and data in your environment, are consolidated and normalized into a relational-style database called the E-SRF Event System Masterfile (normally referred to as "the Masterfile").

Event reporting contains an “engine” that is used to maintain the Masterfile data from any *supported* Resident Security System (RSS). As of release 1.4, the supported RSS systems are CA-ACF2 and IBM RACF. The journalized data is read by “update overlays” and normalized into common event related data. This data is stored on the Masterfile. Reporting is accomplished by executing the Command Processor and requesting one or more “Report Overlays.”

Reporting is done in a normalized fashion, providing reports in as close to “*English*” as possible. A single report may contain security events from multiple systems and multiple RSS systems.

The goal of Event Reporting is an enterprise-wide reporting facility, where the audience does not have to know anything about the security system being executed or the computer systems being reported.

Presenting Information

The primary function of E-SRF is to present information in a clear, concise format. E-SRF reports are unique because the information is simplified and summarized as much as possible during the report creation process.

E-SRF does the evaluation so you can focus on the critical information you need. Both summary and detail reports are available.

People, Places & Things

The Masterfile is divided up into logical *segments* that contain specific types of data. Three of these segments are “*people, places,* and *things.*”

You relate the **people** in your organization to the **places** (terminals and workstations) they work and the **things** (resources) they access during the course of their workday. For example, E-SRF consolidates, summarizes, and presents information if you need a report to:

- show who accessed what datasets from a particular terminal, or
- describe everything a particular user did during a day, or
- show everyone who attempted to update a particular dataset.

Report Overlays

Report overlays are programs that produce reports from the Event Reporting engine. The engine, providing a common look and feel across all reports, performs most of the processing.

All report overlays allow the user to select the data required for reporting. This includes date ranges, resource classes, types of events and so forth.

There are three classes of report overlays:

- Control reports that allow you to view information about the actual Event System
- Utility reports that allow you to customize reports the way you want them or download data for other systems to use
- Fixed, non-modifiable reports that produce data in a format pre-determined by the design of the system. There are currently many fixed reports available, and there are more to come.

Because of the abundance of report overlays and ways to customize and control them, a separate publication, the *E-SRF Event System Report Overlays Guide* is provided. The ability to select data is best explained in the *E-SRF Event System Command Reference*. Other chapters in this publication also explore Report Overlays.

MVS Security System Journal Data

As security events take place, the Resident Security System optionally creates a journal record based on the criticality of the event, or whether or not the security administrator deemed it necessary to "*log*" the event. In the case of MVS, the journal facility is normally the System Management Facility (SMF).

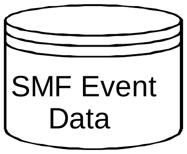

These journal records contain information relevant to the security event. These records may include the user and resource involved, the date and time the event occurred, the identification of the system, the type of access attempted, etc. E-SRF uses this journal information to update its Masterfile.

All Event System reporting comes from the Masterfile. Your Masterfile may be considered a "window" to security events that occur during a pre-determined period of time (which is controlled by the installation).

The E-SRF Event System Masterfile

The E-SRF Masterfile is the heart of the E-SRF Event Reporting Facility (*yet is not used or required for E-SRF Access Analysis reporting!)*. It contains data collected from your Resident Security System(s) (RSS) and the security events that have occurred, as well as global system activity deemed important to the installation.

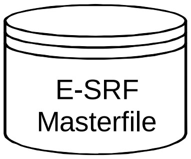

The RSS records this data and events on a journal file such as the MVS operating system's System Management Facility (SMF). At a scheduled point in time, the installation copies this journal data to a secondary file. The system journals are normally unloaded into a flat file for subsequent processing. There is, of course, more on the system journal than just security journal data. Many types of data used for operating system event reporting and possible resource recovery are recorded by the system in journals. Most installations "strip" the security related data from the unloaded flat file and create a separate file that can be used as input to the Masterfile "*update*" function.

The E-SRF Event Facility has a function that "updates" the Masterfile with data from this secondary file. The Update Function will analyze the unloaded flat file or "stripped" file and use this data to update its Masterfile. Updates must only be executed using the unloaded flat file, as Event Reporting will not use the MVS recording files as input, and if attempted, will be detected and terminated.

The E-SRF Masterfile is only as current as its last update. Once established, it must be maintained and should be monitored for space and backed up. The size of your Masterfile along with the time it takes to update is based upon how many security events you choose to log as well as how long you wish to retain this information. EKC recommends you schedule the E-SRF Update Function after the journal "strip" or unload prior to producing reports. Most installations routinely do this early in the morning. Additionally, make sure the correct records are stripped off. Refer to the information provided regarding the Update Function to make sure you have all that is required. An example of this is RACF, where the SMF type 80 records are used, but you may need "some" of the TYPE 30 records for TSO and batch submission signon journals, well as the TYPE 230 in the event your installation is running EKC's product ETF/R.

Please note that if the Masterfile is updated with the CACHE on during the UPDATE function, it will never require any type of reorganization processing (such as a IDCAMS REPRO) to reclaim the empty space generated by the CI (control interval) and CA (control area) splits.

During the process of updating the Masterfile, RSS specific UPDATE Overlays are used to "*normalize*" the data from the various RSS formats to a common consistent E-SRF terminology prior to being stored on the Masterfile for further processing. With few exceptions, an event from one particular RSS (such as ACF2) will look identical to the equivalent event from another RSS (such as RACF). A multi-RSS environment subsequently benefits from these common security reports.

RSS data, such as the description of each USER defined to the RSS, is also stored on the Masterfile and can be included in specified reports. *Any* information contained in the Masterfile can be used in reports. This gives you the ability to include user information (name, phone, etc.) on non-user reports.

The E-SRF Masterfile is a VSAM KSDS cluster, containing records that make up a warehouse of maintained security data. The access method is VSAM, but the logic to relate the data is a three-dimensional relational database scheme made up of row, column and array elements. To gain access to the data, a Data Dictionary has been provided.

For a comprehensive discussion of the Masterfile, please reference the *E-SRF Event System Masterfile and Data Dictionary Reference*, as well as other pertinent topics including *Setting Up Your Masterfile.*

Normalizing and Relating Data

As previously stated, the E-SRF Masterfile data is "normalized". This is important because E-SRF is designed to maintain security event data and provide security event reporting for multiple resident security systems, such as ACF2 or RACF. The goal is to provide a "common face" to the data and the resulting reports

For example, when viewing a report, a violation is simply a violation, regardless of the RSS responsible for posting it. The explanation of the violation that occurred is *normalized* down to common violation reasons, such as access was attempted outside the normal time period, access was attempted from an unauthorized terminal, or the user simply was not authorized to gain access to the resource. It may be desirable to have a more detailed explanation of the violation at the RSS level. E-SRF provides the RSS type and the actual reason codes that are presented in easy-to-understand format.

An example of RSS-specific information would be an ACF2 system denying access to a particular online system transaction because an I/O error occurred when attempting to reference the Generalized Resource Ruleset containing the required rule from its database. The access was denied as a violation, but the *reason* was an I/O error rather than unauthorized access.

For more information, see the *Security Event Reports* chapter later in this guide.

Ownership & Report Distribution

Again, the Access Analysis and the Event Reporting components optionally use the Resource Grouping Facility to associate the group(s) with resources. Additionally, the Event System provides a means to define the GROUP names to the Masterfile and assign them to actual data OWNERS. The owners, as well as all interested parties along with their location and other related data, are defined to the Event Reporting Facility for the purpose of automated report distribution.

For complete information on Grouping Rules, see both the *Administrative Procedures* chapter in this guide and the *Resource Grouping Facility Guide.*

Event Reporting Automatic Report Distribution

The Event Reporting system has an optional feature that will allow you to execute reports in the distributed mode. By using the Automatic Report Distribution Facility, owners that have been assigned to the E-SRF Masterfile will receive reports in a timely manner on resources for which they are responsible. The reports will provide information on what security related events have affected their resources.

A single report execution actually produces a separate report for each data OWNER defined to the E-SRF Masterfile and places the report on the OWNER's output reporting file. This file can then be routed to the particular owner after E-SRF completes its processing. In order to implement this feature, the Masterfile must be provisioned with a definition of each data OWNER and definitions of the GROUPS that the OWNER "owns." Additionally, the Grouping Facility must be provisioned with the desired grouping scheme.

This feature does not have to be implemented immediately. When you are ready to set this up, resources may be grouped a few at a time or all at once. Any *ungrouped* resource will be associated to a default group, which is "owned" by the default, OWNER. This allows you to implement report distribution gradually.

There are Event System reports than can assist Report Distribution set up and analysis. Please refer to the *Event System Report Overlays Guide* for information on the following report overlays: ESRFGRPS, ESRFGRPT, ESRFGRPV, ESRFGRPX, ESRFOGL and ESRFOWNX.

Chapter 3: Event Reporting: Be Productive *Today*

This chapter will instruct you on how to define your Masterfile in order to lead you to your initial production of sample useable reports.

This journey will consist of the following steps:

The Implementation Terms section defines the Domain and Image terms as they relate to E-SRF.

There are five steps normally required to get Event Reporting "up and running."

Step 1: *Creating the E-SRF Masterfile VSAM Cluster* - explains how to define and initialize the Masterfile.

Step 2: *Gathering Resident Security System (RSS) Dependent Data* - provides the necessary job to acquire information from an RSS. This information allows E-SRF to "learn" about the RSS and capture certain control parameters on the Masterfile.

Step 3: *Image Configure, Assigning Domains, and RSS Synchronize* - will take you step by step through these important processes.

Step 4: *Updating the Masterfile* - discusses keeping the Masterfile up to date by applying RSS event updates.

Step 5: *The Quick Start* - identifies a job to produce meaningful Event Reports the very first day.

Resident Security Systems

MVS systems utilize Resident Security Systems coupled with the proper MVS interfaces to provide the level of security that is available under the MVS Operating System.

The MVS Operating System provides a very high level of "integrity" but minimal security. It was designed to interface to a *Resident Security System* (RSS), also known as a *Security Manager* (SM), to extend to MVS the type and level of security desired.

There are currently three approaches available to address security issues:

- The first is to only implement whatever security exists within the MVS Operating System itself.
- A second approach is to design and write your own security system that could precisely address your organization's security needs.
- The third approach is considered to be the most desirable option: selecting one of the three commercially available Resident Security Systems and implementing this choice on your system. E-SRF is designed to deal with reporting issues associated with this option.

This chapter explains general issues and takes you through Step 1. The next two chapters cover Steps 2 through 4 for ACF2 (Chapter 4) and RACF (Chapter 5) specifically (CA-Top Secret is not currently supported in this release of E-SRF Event Reporting). Chapter 6 continues with how to run the sample reports found in the Sample Library (Step 5). *Other Important Matters* (Chapter 7) discusses record retention, rebuilding, reorganizing, and backing up the Masterfile.

Implementation Terms

Domains and Images

These topics will be discussed in more detail later in this publication but need to be touched upon at this point in order to understand the implementation process. These are *E-SRF terms*.

A ***DOMAIN*** (or sysid) is a logical MVS system that is IPL'ed and maintained. Another name for this is an LPAR (Logical PARtition). As far as MVS users are concerned, a single LPAR *is* the MVS system, and two LPARS are two MVS systems, even if both LPARS run on the same physical computer hardware.

An ***IMAGE*** is one or more DOMAINS that share the same security database. For example, two physical computer systems could be at one location. One has a single LPAR (domain), and the other has two LPARs (two more domains). All three run as separate MVS systems and share the same ACF2 LOGONID cluster or RACF database. For E-SRF Event Reporting, the three LPARs are considered a single IMAGE and should be configured as such.

Implementation Considerations

The next sections will explain how to modify JCL (Job Control Language) and submit batch jobs. If you are unfamiliar with JCL or batch job submission, it is recommended that you partner with someone in your organization that understands these concepts.

Some items from this chapter may also be contained in the Installation Guide. Work with your E-SRF product installer to determine what parts, if any, have already been implemented.

Some initial set up is required In order to begin using the *E-SRF Event Reporting Facility*. This includes allocating and initializing the VSAM Masterfile, configuring an IMAGE, and populating the VSAM Masterfile with event data from the Resident Security Systems (RSS).

'SYS1.ESRF.ESRFSMPL' is the default name of the E-SRF *SAMPLIB* dataset, which was created as part of the installation process. *Your SAMPLIB may have been assigned a different name by your installation*. All job names listed in the following sections are included in the E-SRF-supplied *SAMPLIB.*

The next two chapters will explain how the following processing is performed for each RSS (first for ACF2 and then for RACF). The basic procedure is the same, but the RSS differences in JCL, as well as some of the actions, may be confusing if explained together in a single chapter.

The following members in the E-SRF *SAMPLIB will be discussed and used:*

RUNBUILD Defines the VSAM Masterfile cluster and identifies a preliminary set of control information that customizes your E-SRF system.

RUN*xxx*P Looks at a Resident Security System (RSS) and uses its current established parameters to build a file that allows E-SRF to "learn" about the security system that will be running in your Image.

RUN*xxx*I Builds an E-SRF **IMAGE**, assigns an RSS to that **IMAGE**, and **SYNCHRONIZES** the userids or logonids from your security system to the E-SRF Masterfile.

RUN*xxx*U Updates the E-SRF Masterfile with Journal Event Information captured from the RSS being reported on.

RUNQUICK Provides an example of E-SRF Event Reporting. The information contained in these reports is the live data contained in the Masterfile supplied by the **RUN*xxx*P**, **RUN*xxx*I**, and **RUN*xxx*U** jobs above.

Remember: Be sure to make ***your own copy*** of any members before modifying them with information relevant to your operating environment.

Additionally, please first read all instructions provided in the appropriate chapter for the above job. This will insure having a thorough understanding of how these jobs are intended before attempting to execute them. This includes the possibility of having to prepare and execute other processing steps for their execution.

As you proceed through the JCL for jobs identified in the chapter, you may need to add or remove a DD statement from the JCL prior to running the job.

The **//RULES** DD is used to indicate which dataset contains the E-SRF grouping rules. Please note that these "*rules*" are a component of the EKC Integrated Grouping Facility and do not relate to "rules" in any Resident Security System. *It is important not to confuse E-SRF rule syntax with ACF2 rules, although they are similar in appearance.*

If you are a new user of E-SRF, you probably have not yet defined grouping rules. There is no need to define them at this time since doing so may complicate implementation. The E-SRF Event System makes reference to grouping rules for two reasons:

- They may be used for selecting and displaying data (based on GROUP or OWNER ID).
- They are also used in the *Report Distribution Facility*. Grouping Rules are not referenced until Report Distribution is discussed later in this guide.

If you have not created your E-SRF grouping "rules" or are not prepared to implement Report Distribution at this time, simply remove the **//RULES** DD statement from the JCL stream. You will see E-SRF message **E305** with a return code '24' from the Grouping initialization processing. This is *normal* and indicates the grouping facility could not be activated. The "error" will ignore requests for group information, including the overhead required to build the grouping tables. This is the recommended way to run Event System processing until there is a need for grouping.

Once you have set up the grouping rules and want to use them in a particular E-SRF job, you can put the **//RULES** DD statement back in the JCL for any execution that requires grouping.

Step 1: Creating the E-SRF Masterfile VSAM Cluster

Discussion

The E-SRF Masterfile is a VSAM KSDS cluster. It likely will be your responsibility to create the Masterfile VSAM dataset. The process for creating the VSAM Masterfile is described in the following topics.

For a comprehensive discussion of the Masterfile, including its physical characteristics, refer to the *E-SRF Event System Masterfile and Data Dictionary Reference.*

Definition: INITIALIZE

The INITIALIZE command is used to set up the initial records in the appropriate format for the E-SRF Masterfile. This command, issued with a RESET parameter (INITIALIZE RESET), will allow you to replace the System Dictionary. This command should only be used when you are ready to start over with a fresh Masterfile database.

Masterfile VSAM Cluster Creation JCL - RUNBUILD

The following is a copy of the JCL stream (**RUNBUILD**) found in the SAMPLIB containing the necessary IDCAMS commands to define your cluster.

There is an execution of the ESRF Command Processor immediately following the IDCAMS DEFINE of your Masterfile's VSAM cluster. The specifications here will work for you. As you get more familiar with the product, you may want to change some of these specifications, but for now, this will achieve the goal of a functional Masterfile that may be used to produce reports.

Each line in the JCL is identified with a number. Following the JCL is an explanation of each line. Refer to the number of the line to find the appropriate definition.

Remember: make ***your own copy*** of any members before modifying them with information relevant to your operating environment.

```
1)     //JOBNAME  JOB (XXXX),'XXXXXXXXX',MSGCLASS=X,MSGLEVEL=(1,1)
2)     //*
3)     //********************************************************************
4)     //*                                                                  *
5)     //*       JOB TO CREATE AND INITIALIZE ESRF MASTERFILE CLUSTER.      *
6)     //*                                                                  *
7)     //*       SAMPLIB MEMBER: RUNBUILD                                   *
8)     //*------------------------------------------------------------------*
9)     //*                                                                  *
10)    //*       NOTE: --- PRIOR TO RUNNING THIS JOB:                       *
11)    //*                                                                  *
12)    //*    1) MAKE SURE DD STATEMENTS ARE VALID FOR YOUR INSTALLATION.   *
13)    //*                                                                  *
14)    //*    2) REVIEW ALL OPTIONS SPECIFIED IN BOTH STEPS OF THIS JOB     *
15)    //*       TO MAKE SURE THEY ARE APPLICABLE TO YOUR INSTALLATION.     *
16)    //*                                                                  *
17)    //*------------------------------------------------------------------*
18)    //*                                                                  *
19)    //*   --> THIS JOB IS INTENDED TO BE USED WHEN INITIALLY             *
20)    //*       CREATING AN E-SRF MASTERFILE FROM SCRATCH.                 *
21)    //*                                                                  *
22)    //*   --> THIS JOB WILL CREATE THE 'SHELL' OF A MASTERFILE           *
23)    //*       THAT WILL BE MADE USEFUL BY SUBSEQUENT JOBS.               *
24)    //*                                                                  *
25)    //*   --> PLEASE IGNORE THE GROUPING FACILITY ERROR.                 *
26)    //*                                                                  *
27)    //********************************************************************
28)    //*
```

```
//*       CREATE MASTERFILE VSAM CLUSTER.
//*
//*
//DEFINE  EXEC PGM=IDCAMS
//SYSPRINT DD  SYSOUT=*
//INPUT    DD  *
A CONTROL
//SYSIN    DD  *
      DELETE ESRF.MASTER PURGE
      DEFINE CLUSTER -
             (  NAME(ESRF.MASTER) -
                OWNER(ESRF) -
                VOLUME(XXXXXX) -
                SHAREOPTIONS(2 3) -
                SPEED -
                REUSE ) -
             DATA -
             (  NAME(ESRF.MASTER.DATA) -
                KEYS(62 0) -
                SPANNED -
                RECORDS(200000 10000) -
                RECORDSIZE(2048,8192) -
                FREESPACE(0,0) ) -
             INDEX -
             (  NAME(ESRF.MASTER.INDEX) )
      REPRO  IFILE(INPUT) ODS(ESRF.MASTER)
//*
//*       INITIALIZE MASTERFILE AND ESTABLISH SYSTEM OPTIONS.
//*
//INIT    EXEC PGM=ESRFCMD,REGION=0M
//*
//STEPLIB  DD  DSN=ESRF.LOAD,DISP=SHR        ESRF SYSTEM.
//MASTER   DD  DSN=ESRF.MASTER,DISP=SHR      ESRF MASTERFILE.
//SYSPRINT DD  SYSOUT=*                      CONTROL REPORT.
//*
//SYSIN    DD  *

   /*      SAMPLE COMMANDS TO INITIALIZE E-SRF MASTERFILE SHELL.

           NOTE: 'COMPRESS' IN SET COMMAND IS COMMENTED OUT.
                 (CONSULT COMMAND GUIDE FOR MORE INFORMATION).

                 RETENTION VALUES ARE FOR EXAMPLE ONLY, DAYS FOR SOURCE
                 RECAP AND USER TRACE ARE SET TO ZERO.  IF YOU DECIDE
                 THE CHOSEN VALUES ARE NOT DESIRABLE, THEY MAY BE
                 MODIFIED NOW OR LATER WHEN REQUIRED.

   */

           INITIALIZE  RESET

           SET         NOBANNERS      -
                       MAXLINES(58)   -
                       ELEMENTS(6144) -
                       EXCLUDE (NONE)  -
                       GDG(TRUNCATE)  -
                       WIDTH(133)     -
                   *   COMPRESS(A0000000,B0000000,G0000V00) -
                       TITLE(YOUR COMPANY NAME UP TO 64 CHARACTERS)

           RETAIN      RESET

           RETAIN      OBJECT(FC) DAYS(30)

           RETAIN      OBJECT(MC) DAYS(30)

           RETAIN      OBJECT(RC) DAYS(5)
           RETAIN      OBJECT(RM) DAYS(60)
           RETAIN      OBJECT(RR) DAYS(30)
           RETAIN      OBJECT(RS) DAYS(10)
```

```
89)              RETAIN     OBJECT(SR) DAYS(0)
89)              RETAIN     OBJECT(SU) DAYS(0)
91)              RETAIN     OBJECT(UA) DAYS(30)
92)              RETAIN     OBJECT(UB) DAYS(30)
93)              RETAIN     OBJECT(UC) DAYS(5)
94)              RETAIN     OBJECT(UF) DAYS(60)
95)              RETAIN     OBJECT(UM) DAYS(60)
96)              RETAIN     OBJECT(UR) DAYS(30)
97)              RETAIN     OBJECT(US) DAYS(10)
98)              RETAIN     OBJECT(UT) DAYS(0)

99)              RUN        REPORT(ESRFSHOW)
```

JCL Line Definitions for RUNBUILD:

1) Standard MVS “JOB” statement. Code this to your installation’s specifications.

2 - 27) Standard MVS “comment statements.”

28 - 31) Standard MVS “comment statements” documenting the execution of IDCAMS.

32) Standard MVS “EXECUTE” statement. It indicates the name of the program you are executing. In this case, it is the MVS utility IDCAMS.

33) Standard MVS “DD” statement. It relates to a file called SYSPRINT where the IDCAMS utility writes its printed report output.

34) Standard MVS “DD” statement. It relates to a file called INPUT specified in the IDCAMS commands as an input file.

35) An initial data record that is going to be copied into the VSAM cluster by statement number 54.

36) Standard MVS “DD” statement. It relates to a file called SYSIN where the IDCAMS utility reads its command input. In this case, the input data is in line with the JCL.

37) IDCAMS control statement. This command deletes any existing E-SRF Masterfile VSAM cluster.

Initially, the Masterfile will not exist so this command will fail with a return code of eight. This is acceptable as long as it is the only reason you received the return code of eight. The operation ensures you will have the latest version of the Masterfile format in cases where you are re-installing the product.

38) IDCAMS control statement. This command defines the cluster.

39) IDCAMS control statement. This command identifies the dataset name of the cluster. Choose a name using the dataset naming conventions established at your site.

40) IDCAMS control statement. This command identifies the “owner” of the cluster. Specify an owner in accordance with your installation standards for VSAM ownership.

41) IDCAMS control statement. This command identifies the target volume(s) that will contain the cluster.

42) IDCAMS control statement. This command specifies the VSAM sharing options. Typically, code your options to allow a single update (2) and multiple read-only (3).

43) IDCAMS control statement. This command specifies the “speed” attribute. This attribute identifies the data component’s space as not pre-formatted.

44) IDCAMS control statement. This command identifies the REUSE option. This is important and **must** be specified. During CACHE upgrades, E-SRF may decide to completely rebuild your Masterfile. If this option is not in effect, the CACHE rebuild will fail and the cluster will not be updated.

45) IDCAMS control statement. This command defines the data component of your KSDS VSAM cluster.

46) IDCAMS control statement. This command identifies the name of the data component of your VSAM cluster. It should be the same name as the base cluster name you chose in statement #39 with the ".DATA" qualifier appended to the end. For example, if the Masterfile cluster name is ESRF.MASTER, the data component name should be ESRF.MASTER.DATA.

47) IDCAMS control statement. This command identifies the length and relative position of the file KEY within a VSAM record. In this case, the KEY is sixty-two positions long and starts at relative position zero which is position one of the record.

48) IDCAMS control statement. This command informs VSAM that control intervals may be spanned. Logical records may cross Control Interval boundaries.

49) IDCAMS control statement. This command defines how large to make the data component of the VSAM Masterfile. Specify the size of your data component, in any denomination you choose (bytes, records, cylinders, etc.). We recommend records. Suggestions on how large to make this file are discussed in the *Event System Masterfile and Data Dictionary Reference*.

50) IDCAMS control statement. This command defines the average and maximum record lengths contained on the VSAM Cluster. The average record length is followed by the maximum record length. E-SRF considers the cluster to be variable length with most records being 8192 bytes in length. All record segments for any given object will be the length required to contain them. The settings in the example are appropriate.

51) IDCAMS control statement. This command specifies how much free space (as a percentage) is to be set aside in the data component of the cluster for adding and/or lengthening records. This specification is dependent on whether or not you use the CACHE during update. If you do not use CACHE, leave a large amount of freespace. If you use the CACHE, **which we recommend**, specify zero freespace.

52) IDCAMS control statement. This command defines the index component of your KSDS VSAM cluster. In most cases, IDCAMS and VSAM will do a better job of determining space and other related values for the index component than people can. Let VSAM and IDCAMS do as much as possible in defining the index component.

53) IDCAMS control statement. This command identifies the name of the index component of your VSAM cluster. It should be the same name as the base cluster name in statement #39 you chose with the ".INDEX" qualifier appended to the end. For example, if the Masterfile cluster name is ESRF.MASTER, the index component name should be ESRF.MASTER.INDEX.

54) IDCAMS control statement. This command copies a single record (from statement #35) to the newly created VSAM Cluster. This record is placed in the cluster to prevent E-SRF from having problems relating to KSDS structures with an "unloaded" KSDS VSAM Cluster. The single record is carried in-stream with the JCL. E-SRF knows about this record and deals with it when encountered. E-SRF has no dependencies on this record and it is eventually removed.

55 - 57) Standard MVS "comment statements."

58) Standard MVS "EXECUTE" statement. It indicates the name of the program you are executing. In this case, it is the E-SRF Command Processor ESRFCMD.

59) Standard MVS "comment statement."

60) Standard MVS "DD" statement. It relates to a file called STEPLIB where the MVS program loader will look for E-SRF. This includes the Command Processor and any other E-SRF programs required to execute the E-SRF session. The Load Library containing E-SRF must be specified here. This DD is not required if the E-SRF programs are contained in the MVS LINK LIST. Consult your E-SRF installer for more information.

61) Standard MVS "DD" statement. It relates to a file called MASTER. This is the E-SRF Masterfile cluster previously created. At the beginning of this step, there is a single record in the cluster.

62) Standard MVS "DD" statement. It relates to a file called SYSPRINT where program ESRFCMD writes its control report output.

63) Standard MVS "comment statement."

64) Standard MVS "DD" statement. It relates to a file called SYSIN where program ESRFCMD reads its command input. In this case, the input data is in line with the JCL.

65 - 72) E-SRF Command Processor “comment statements.”

73) E-SRF major command to initialize the Masterfile. The “reset” minor command will reset the data dictionary if there is an existing one. At the initial Masterfile VSAM cluster creation, a data dictionary should not exist.

74) Start of E-SRF major command to SET system options.

Minor command NOBANNERS tells E-SRF you do not want two-page separator pages printed for each report.

75) Minor command MAXLINES tells E-SRF the maximum number of lines a report overlay will print on a single page.

76) Minor command ELEMENTS tells E-SRF the maximum number of array elements, which a single E-SRF object may contain before an “unexpired rolloff” occurs. For more information on unexpired rolloffs, refer to *Setting Up Your Masterfile* in this guide.

77) Minor command to determine how E-SRF should process EXCLUDE Grouping Rules.

78) Minor command GDG(TRUNCATE) tells the E-SRF update control program to drop the G0000V00 qualifier from any dataset name that has one. The result will be to associate all GDG activity to the “base” dataset name. For information about other parameters for the GDG command, refer to the *Command Reference Guide*.

79) Minor command WIDTH tells E-SRF the global width standard for executing Report Overlays.

80) Minor command COMPRESS specifies to E-SRF a list of dataset qualifier compression masks.

81) Minor command TITLE assigns a one to sixty-four character constant title that will appear on all report output produced by E-SRF. We recommend you include your company name here.

82 - 98) RETAIN commands. These commands establish the number of days that Masterfile event data will stay on the file before being purged. They refer to the available objects within the various segments. These numbers represent the “window” of available event activity that may be used to produce reports. Set these values according to your reporting needs. The RETAIN RESET resets the retention days to the system defaults. Refer to the *Command Reference Guide* for additional information.

99) E-SRF major command to RUN the ESRFSHOW report overlay. It lists the current E-SRF processing options in effect.

The next two chapters provide you with the necessary information to continue “getting productive” and are based on which Resident Security System you use. The chapters are identical in what topics are covered, differing only as they apply to the respective Resident Security System. Chapter 6 then continues with how to run the “Quick Job” reports.

Chapter 4: Be Productive *Today* (with ACF2)

This chapter is intended to take you from defining your Masterfile to being ready to produce sample useable reports when you are using ACF2 as your Resident Security System. For a discussion on RACF, please review the next chapter.

Step 2: Gathering RSS Dependent Data for ACF2

Discussion

The E-SRF Event System requires some basic information about the Resident Security System (RSS) that will be reported on. The program **ESRFACFP** is provided to build a parameter file that may be used in the remaining jobs in this chapter. This job is run on the system that E-SRF is to report on. Its purpose is for E-SRF to *learn* about the system.

This step may be optional on ACF2 systems in certain situations as described below.

Gathering RSS Dependent Data JCL - RUNACFP

On ACF2 systems, if you are going to run the IMAGE configuration on the system you intend to report on, *this step may not be necessary*.

E-SRF has the ability to examine the executing ACF2's system and gather all this information during its configuration process. If you are in a position to do this, you can skip this step. Remember to leave out the DDNAME parameter on the CONFIGURE command later in this process.

By omitting the DDNAME parameter from an ACF2 IMAGE's "*CONFIGURATION,*" E-SRF will "learn" about the system that is actually hosting the configuration.

If the above option is not possible, or not desired, use member: **RUNACFP** to create the configuration parameters for the ACF2 IMAGE.

This step is **required** if you are configuring an IMAGE on your Masterfile that resides on a different system and has different characteristics that the one hosting your execution. This would include the ACF2 UID string or if the system hosting this configuration is running a Resident Security System other than ACF2.

This step 'PORTS" information from the system you want to configure to the system that is hosting the actual configuration processing. This is very common because unless you only have a single system, your Masterfile was designed to maintain security event data for multiple systems across your enterprise. Some of these may not be attached and therefore cannot directly access the Masterfile.

This step will take you through the porting procedure.

The characteristics of this file are indicated in the sample JCL.

RUNACFP is the sample job for you to run to extract the information from a running ACF2 system.

The following is a copy of **RUNACFP**. Each line in the JCL is identified with a number. Following the JCL is an explanation of each line. Refer to the number of the line to find the appropriate definition.

Remember: Make your own copy of any members before you modify them with information relevant to your operating environment.

RUNACFP:

```
1)     //JOBNAME  JOB (XXXX),'XXXXXXXXX',MSGCLASS=X,MSGLEVEL=(1,1)
2)     //*
3)     //*********************************************************************
4)     //*                                                                   *
5)     //*       RUN ESRFACFP -- ACF2 PARAMETERIZATION PROGRAM               *
6)     //*                                                                   *
7)     //*-------------------------------------------------------------------*
8)     //*                                                                   *
9)     //*       THIS IS AN OFFLINE BATCH UTILITY THAT MAY BE USED           *
10)    //*       TO EXTRACT INFORMATION FROM THE ACF2 SYSTEM BEING           *
11)    //*       CONFIGURED AS AN IMAGE TO E-SRF.                            *
12)    //*                                                                   *
13)    //*-------------------------------------------------------------------*
14)    //*                                                                   *
15)    //*       THIS PROGRAM LOOKS AT AN ACF2 SYSTEM AND USES THE CURRENT   *
16)    //*       ACF2 ESTABLISHED PARAMETERS TO BUILD A PARAMETER FILE       *
17)    //*       THAT MAY BE TRANSPORTED TO THE SYSTEM WHERE THE IMAGE IS    *
18)    //*       BEING CONFIGURED.                                           *
19)    //*                                                                   *
20)    //*-------------------------------------------------------------------*
21)    //*                                                                   *
22)    //*       PARMOUT:  MUST BE A 255 BYTE FIXED FILE LARGE ENOUGH TO     *
23)    //*                 HOLD TWO HUNDRED RECORDS.                         *
24)    //*                                                                   *
25)    //*********************************************************************
26)    //*
27)    //RUNACFP EXEC PGM=ESRFACFP
28)    //STEPLIB  DD  DSN=ESRF.LOAD,DISP=SHR        ESRF SYSTEM.
29)    //PARMOUT  DD  DSN=ESRF.ACF2PARM,DISP=SHR    PARAMETER FILE.
30)    //SYSPRINT DD  SYSOUT=*                      CONTROL REPORT.
```

JCL Line Definitions for RUNACFP

1) Standard MVS “JOB” statement. Code this to your installation’s specifications.

2 - 26) Standard MVS “comment statements.”

27) Standard MVS “EXECUTE” statement. It indicates the name of the program you are executing. In this case, it is the E-SRF program **ESRFACFP**.

28) Standard MVS “DD” statement. It relates to a file called STEPLIB where the MVS program loader will look for E-SRF. This includes the Command Processor and any other E-SRF programs required to execute the E-SRF session. The Load Library containing E-SRF must be specified here. (This DD is not required if the E-SRF programs are contained in the MVS LINKLIST). Consult your E-SRF installer for more information.

29) Standard MVS “DD” statement. It relates to the file produced by the program **ESRFACFP** that builds parameters from ACF2 for input into E-SRF. This is the input that will “teach” E-SRF about the particular ACF2 IMAGE.

30) Standard MVS “DD” statement. It relates to a file called SYSPRINT where **ESRFACFP** writes its printed output.

Output file characteristics

ESRFACFP creates an output file that becomes input to the IMAGE Configuration process. The characteristics of this file are as follows:

The file should be large enough to contain several hundred 255-character data records. Any suitable blocksize may be specified.

The records may be edited with ISPF to alter the descriptions of the fields contained on the User Header object, if desired.

The following JCL may be incorporated in lieu of PARMOUT (found on line 29 of the example).

```
29)   //PARMOUT    DD    DSN=ESRF.ACF2PARM,DISP=(NEW,CATLG),
      //                 UNIT=SYSDA,VOL=SRR=VOLID,
      //                 DCB=(LRECL=255,BLKSIZE=7650,RECFM=FB),
      //                 SPACE=(7650,(10,5))
```

Porting the data to the hosting system

ESRFACFP creates an output file that must be "ported" to the system that hosts the E-SRF Masterfile. This file is very small. Any suitable means to transport the data may be used. Most 3270 terminal emulators have a file download/upload facility. FTP (File Transfer protocol) may also be used. The only requirement is the data must be present for the Masterfile IMAGE configuration. After that, the file is no longer referenced.

Step 3: ACF2 Image Configure, Assigning Domains, and Synchronize

Discussion

Once an E-SRF Masterfile has been defined and initialized, it is time to tie together the associations of Domains, Images, and RSS specific data. The **RUNACFI** job found in the sample library will assist you with this process for an ACF2 IMAGE.

CONFIGURE

The CONFIGURE command builds an E-SRF IMAGE and assigns a Resident Security System (RSS) to the IMAGE. For example, if you are creating an image called CHICAGO, which uses ACF2 as its RSS, you would use a CONFIGURE command which should look similar to the example shown below.

```
CONFIGURE    IMAGE(CHICAGO)             -
             RSS(ACF2)                  -
             NAME(CHICAGO CENTRAL)
```

This command will allow E-SRF to "learn" about the ACF2 RSS that the IMAGE will represent.

BE PRODUCTIVE TODAY (with ACF2)

In this example, E-SRF will use the RSS ACF2 data from the system on which this command was issued. If this is not possible (perhaps the system you are running this job is not the IMAGE you are working with), you will have to provide the following:

```
     CONFIGURE    IMAGE(CHICAGO)            -
                  RSS(ACF2)                 -
                  NAME(CHICAGO CENTRAL)     -
                  DDNAME(ACF2PARM)
```

The ACF@PARM dataset (if required) was created by step two described in this chapter. The DDNAME could have been almost anything you wanted it to be, as long as it matched your definition in the JCL.

You may or may not have needed to run ESRFACFP. If you omit the DDNAME configuration specification, the characteristics of the hosting ACF2 system will be applied to the IMAGE you are configuring. If the system hosting this process is the IMAGE you are configuring, or the IMAGE you are configuring has the exact same ACF2 UID string and LOGONID FDR characteristics, the execution of ESRFACFP is not required.

SYNCHRONIZE

The SYNCHRONIZE command loads the Userids present in the security database that exists for the IMAGE being SYNCHRONIZED to the E-SRF Masterfile, making this information available for userid (UA header) reporting (related to the target IMAGE contained on the Masterfile).

When configuring an IMAGE, the process should normally be followed up with a SYNCHRONIZE to bring the Masterfile up to date with the latest RSS userid data. This is vital in the event that the ACF2 Logonid database was altered externally or there have been days skipped during the update process.

You can print the user's name or any other ACF2 Logonid fields on a report (or use Logonid fields as selection criteria) because that information will be contained on the E-SRF Masterfile.

Normally, the updating of the Masterfile will keep this information in sync with routine ACF2 updating. Certain processing may not be able to provide this type of updating. In these cases, you will have to re-SYNCHRONIZE your Masterfile to make sure the userid data reported in E-SRF reflects what is contained on the ACF2 system's database. If you are in doubt, re-SYNCHRONIZE. It cannot hurt anything even if executed when not required.

If you make "offline" changes (such as an ACF2 UID string conversion), you will have to re-CONFIGURE and re-SYNCHRONIZE your Masterfile.

Caution: Always make sure you use the correct backup ACF2 Logonid file as input. E-SRF cannot insure this is properly done. The use of an improper Logonid backup file will cause unmatched users to be added to the Masterfile, and data contained on matched users may be altered causing unwanted changes.

An example of the SYNCHRONIZE command is indicated below:

```
             SYNCHRONIZE         DDNAME(backup Logonid flat file DDNAME) -
                                 IMAGE(CHICAGO)
```

You must obtain a very recent Logonid database backup flat file. If you do not have one, execute the appropriate ACF2 console command to create one.

ASSIGN

The ASSIGN statement is how E-SRF knows which IMAGE to place data from a particular domain.

You must provide the necessary ASSIGN statement to associate your DOMAIN(S) to your IMAGE. In the sample job, we assume there is only one IMAGE, and we provided you with an ASSIGN statement that associates all domains to an image called MVSACF2. This will only work if there is one IMAGE. If you omit this step, all the data will go to whatever IMAGE assignment matches the domain id when presented during an Update Function.

Assignments are *maskable*. A domain ID is the identification of the actual system (or LPAR) that created a journal. You can have journal files that contain journals from multiple domains, but they must all be from the same RSS (*in this case, ACF2*). The Update Function will update multiple domains, but only to a specific RSS (*in this case, ACF2*). If you have multiple RSS systems, such as ACF2 and RACF, you would have to run an Update Function for ACF2 images and another for RACF images.

You may have existing assignments that may have to be changed. Look over all of your assignments when adding or changing IMAGES to make sure you are directing your journals to their proper IMAGES.

```
ASSIGN DOMAIN(********)    -
             IMAGE(MVSACF2)       -
             NAME(ACF2 SECURITY IMAGE)
```

Configuring, Synchronizing, and Assigning JCL - RUNACFI

RUNACFI contains the necessary JCL commands to CONFIGURE your ACF2 IMAGE, SYNCHRONIZE the E-SRF Masterfile to the existing users defined on the ACF2 Logonid database, and ASSIGN the applicable DOMAINS (individual LPARS that use the Logonid database) to the IMAGE. Each line in the JCL is identified with a number. Following the JCL is an explanation of each line. Refer to the number of the line to find the appropriate definition.

The following is a copy of job RUNACFI. Each line in the JCL is identified with a number. Following the JCL is an explanation of each line. Refer to the number of the line to find the appropriate definition.

If you require the output from **ESRFACFP** (created earlier), then you must supply and edit the “DD” card in the JCL, and make sure the CONFIGURE parameter points to its ddname in the DDNAME parameter.

Consult the *Command Reference* if you need assistance with the syntax editing.

Remember: Make ***your own copy*** of any members before you modify them with information relevant to your operating environment.

RUNACFI:

```
1)     //JOBNAME  JOB (XXXX),'XXXXXXXXX',MSGCLASS=X,MSGLEVEL=(1,1)
2)     //*
3)     //********************************************************************
4)     //*                                                                  *
5)     //*       JOB TO CREATE AN ACF2 IMAGE ON THE ESRF MASTERFILE.        *
6)     //*                                                                  *
7)     //*       SAMPLIB MEMBER: RUNACFCI                                   *
8)     //*------------------------------------------------------------------*
9)     //*                                                                  *
10)    //*       NOTE: --- PRIOR TO RUNNING THIS JOB:                       *
11)    //*                                                                  *
12)    //*    1) YOU MUST HAVE AN EXISTING MASTERFILE.                      *
13)    //*                                                                  *
14)    //*    2) MAKE SURE DD STATEMENTS ARE VALID FOR YOUR INSTALLATION.   *
15)    //*                                                                  *
16)    //*    3) REVIEW ALL OPTIONS ARE APPLICABLE TO YOUR INSTALLATION.    *
17)    //*                                                                  *
18)    //*    4) IF REQUIRED, YOU MUST HAVE THE DATASET CONTAINING THE      *
19)    //*       OUTOUT CREATED BY 'RUNRACP'.                               *
20)    //*                                                                  *
21)    //*    5) YOU MUST HAVE THE VSAM CLUSTER CONTAINING INFORMATION      *
22)    //*       FROM YOUR RACF DATABASE, CREATED BY MEMBER: EKCRRCDB.      *
23)    //*                                                                  *
24)    //*------------------------------------------------------------------*
25)    //*                                                                  *
26)    //*       IGNORE THE GROUPING FACILITY ERROR.                        *
27)    //*                                                                  *
28)    //********************************************************************
29)    //*
30)    //INIT    EXEC PGM=ESRFCMD,REGION=0M
31)    //*
32)    //STEPLIB  DD  DSN=ESRF.LOAD,DISP=SHR        ESRF SYSTEM.
33)    //MASTER   DD  DSN=ESRF.MASTER,DISP=SHR      ESRF MASTERFILE.
34)    //*
35)    //*ACF2PARM DD  DSN=ESRF.ACF2PARM,DISP=SHR    ESRFACFP PARAMETERS.
36)    //ACFLIDS  DD  DSN=ACF2.BKLIDS,DISP=SHR      ACF2 LID BACKUP FILE.
37)    //SYSPRINT DD  SYSOUT=*                      CONTROL REPORT.
38)    //*
39)    //SYSIN    DD  *

40)       /*      THE FOLOWING COMMANDS WILL CONFIGURE AND SYNCHRONIZE A
41)               NEW RACF IMAGE ON YOUR EXISTING MASTERFILE.

42)               IT IS ASSUMED THAT THIS IS THE ONLY IMAGE THAT WILL BE
43)               CONTAINED ON YOUR MASTERFILE.  IF THIS IS NOT THE CASE,
44)               YOU WILL HAVE TO ADJUST YOUR ASSIGN STATEMENTS TO SELECT
44)               THE DOMAINS (MVS SYSIDS) YOU WANT TO BE INCLUDED IN THIS
45)               IMAGE.  (CONSULT COMMAND GUIDE FOR MORE INFORMATION).
46)      */

47)               CACHE       ON

48)               INITIALIZE  RESET

49)               CONFIGURE   IMAGE(MVSACF2)        -
50)               *           DDNAME(ACF2PARM)      -
51)                           PARAMETER()           -
52)                           NAME(ACF2 SYSTEM)     -
53)                           RSS(ACF2)             -
54)                           RESET

55)               SYNCHRONIZE IMAGE(MVSACF2) DDNAME(ACFLIDS)

56)               ASSIGN      DOMAIN(********) IMAGE(MVSACF2) -
57)                           NAME(ACF2 SECURITY IMAGE)

58)               RUN         REPORT(ESRFSHOW)
59)               RUN         REPORT(ESRFDICT)
60)               RUN         REPORT(ESRFSTAT)
```

JCL Line Definitions for RUNACFI

1) Standard MVS “JOB” statement. Code this to your installation’s specifications.

2 - 29) Standard MVS “comment statements.”

30) Standard MVS “EXECUTE” statement. It indicates the name of the program you are executing. In this case, it is the E-SRF Command Processor ESRFCMD.

31) Standard MVS “comment statement.”

32) Standard MVS “DD” statement. It relates to a file called STEPLIB where the MVS program loader will look for E-SRF. This includes the Command Processor and any other E-SRF programs required to execute the E-SRF session. The Load Library containing E-SRF must be specified here. This DD is not required if the E-SRF programs are contained in the MVS LINK LIST. Consult your E-SRF installer for more information.

33) Standard MVS “DD” statement. It relates to a file called MASTER. This is the E-SRF Masterfile created previously.

34) Standard MVS “comment statement.”

35) Standard MVS “DD” statement. It relates to a file called ACF2PARM. This file was created in job **RUNACFP**. *If the ACF2 RSS resides on the hosting image, the job may not have been required and “DD” can be commented out. If you are importing an ACF2 RSS from another image, RUNACFP must be run and the “DD” statement must be active. Please see the RUNACFP job in your SAMPLIB to note these differences.*

36) Standard MVS “DD” statement. It relates to a file called “*ACF2LIDS*” but could have been any DDNAME that was specified in the SYNCHRONIZE command. This is the input used to perform the SYNCHRONIZE function for a particular security *IMAGE*. It relates to the flat file representation of the ACF2 Logonid Database, which the E-SRF IMAGE is being configured to run with. Our example calls this file *ACFLIDS*.

37) Standard MVS “DD” statement. It relates to a file called SYSPRINT where program ESRFCMD writes its control report output.

38) Standard MVS “comment statement.”

39) Standard MVS “DD” statement. It relates to a file called SYSIN where program ESRFCMD reads its command input. In this case, the input data is in line with the JCL.

40 - 46) E-SRF Command Processor “comment statements.”

47) E-SRF command to turn the Masterfile Caching option ON. CACHE OFF may also be specified for those sites not using the Caching option. It is highly recommended that you run with CACHE ON for this process.

48) E-SRF major command to initialize the Masterfile. The “reset” minor command will reset the data dictionary if there is an existing one.

49) E-SRF major command to CONFIGURE the Masterfile for a specific security IMAGE. This statement contains the major command and identifies the security Image to be configured. “*MVSACF2*” is the IMAGE chosen in this example. It can be called any UNIQUE NAME desired.

50) This statement *OPTIONALLY* identifies the DDNAME of a dataset containing the configuration parameters from the Resident Security System being configured to the Image. *The “DD” name ACF2PARM matches the “DD” name ACF2PARM in the RUNACFP job previously submitted*. *If RUNACFP was not run, this statement must be commented out*.

51) Optional configuration parameter. The value specified is dependent on the particular RSS being configured to the target security Image. Refer to the *Command Reference Guide* for more information. Currently, this parameter should be set as shown in this example.

52) This statement identifies the NAME of the security IMAGE being configured. In this example, we have chosen “ACF2 SYSTEM” as the name of the security IMAGE.

53) This statement identifies the particular RSS being configured to the target security IMAGE. In this example, the RSS is RACF.

54) The RESET attribute indicates that if a previous IMAGE (MVSACF2) already exists, it will be reset to the specifications declared here.

55) This major command will SYNCHRONIZE the IMAGE (MVSACF2) to the specified Resident Security System. Each RSS has its own requirements. This example is for ACF2, which uses the ACF2 Logonid database backup flat file (*mentioned in statement #36). It is a flat file representation of the ACF2 Logonid Database*.

56) Major command to ASSIGN all DOMAINS, designated by the "*xxxxxxxx,*" to the target Image (MVSACF2).

57) Name of the assignment. "ACF2 SECURITY IMAGE" was chosen for this example. All domains in this IMAGE will match this assignment. All activity updated against the Masterfile will end up in the IMAGE MVSACF2.

58) E-SRF major command to RUN the ESRFSHOW report overlay. It lists the current E-SRF processing options.

59) E-SRF major command to RUN the ESRFSTAT report overlay. It provides a byte count of the data contained on the Masterfile by E-SRF segment and object. This report can be useful when determining where the majority of your security event records are being stored and where the majority of the security violations and/or loggings are taking place.

60) E-SRF major command to RUN the ESRFDICT report overlay. It provides an alphabetical list of all E-SRF data names and object names contained on the E-SRF data dictionary. For more information about the E-SRF data dictionary, see *E-SRF Data Dictionary Structure* in this guide and the *Masterfile and Data Dictionary Reference Guide.*

UID String Conversion Issues

If you run a UID string conversion, you will probably have to re-configure the IMAGE(s) that you converted. Additionally, you may have to re-SYNCHRONIZE the IMAGE(s) using your converted ACF2 Logonid database. This depends on how you did your conversion.

Important Note: If you are not sure what to do, *please contact EKC Technical Support* prior to the conversion for advise on what you may have to do to adjust ESRF Event Reporting to reflect your new structures.

Step 4: Updating the Masterfile with Your Data

Discussion

At this point, reports could actually be run. If you run the "**RUNQUICK**" SAMPLIB member, you will end up with userid reports but empty event reports, which would serve no purpose. Yes, you have a configured Masterfile full of users, but as of yet, no security events.

In order for E-SRF to report on events, you must provide Journal Event Information from the particular RSS from which you are reporting. The act of placing the journals on E-SRF is referred to as "updating the Masterfile," or more technically the "UPDATE FUNCTION." There are entire chapters of E-SRF publications dedicated to this topic. For now, let us review just enough to allow us to get it up and running.

After the E-SRF Masterfile has been **defined, configured, and synchronized**, it must be **updated** to include records on security events that occurred on your system. This is accomplished by applying your daily security journals (SMF data) to your Masterfile. The job RUNACFU in the SAMPLIB contains the JCL to do this. There is also another job for RACF. You are interested in ACF2, so use RUNACFU. If you have ACF2 and RACF, they can live on the same Masterfile provide they are in separate IMAGES. Process each update separately using the same Masterfile. When you publish reports, all security events may be treated as a single set of data.

There is not much external difference between one security system and another when setting up the UPDATE Function. Each security system has its own Update Overlay that "*normalizes*" the journalized data into something common for E-SRF's Masterfile Update Processor to use when placing events on the Masterfile.

However, you **must** declare which security system you want to process during a single UPDATE execution. You may have as many UPDATE commands in a single run as desired. The SMF data may be supplied in any sequence. You can first update today's, then last week's and then yesterday's. E-SRF will sort out and collate the journals from various UPDATE runs. The end result will be a seamless recording of security events across whatever span of time that was declared when E-SRF was configured. Old data will be automatically rolled off the Masterfile.

It is possible that you may present too much data for E-SRF to store. Limits are established when E-SRF is configured, and in the event a particular limit was exceeded, the data will be *'rolled off*", oldest first. This may seem severe, but experience has shown that the vast majority of the data rolled off this way consists of excessive logging that is really not desired. This was set up this way to limit overly large reports, huge DISK and STORAGE requirements, and unusable reports.

All of this can be examined and controlled after operational experience has been gained using this product. Ultimately, you will have a reporting system, tailored to your auditing and reporting needs, and containing only the information you desire to have reported, delivered to the data owners you deem responsible for reviewing.

Updating an ACF2 IMAGE

The UPDATE command applies SMF journal data regarding security events to the E-SRF Masterfile. This command should be run periodically to update the Masterfile with recent security events. We recommend you schedule the UPDATE job to run nightly after the SMF extract job.

```
UPDATE     RSS(ACF2)    SMFTYPE(230)
```

This command is explained in more detail in the *E-SRF Event Reporting System Command Reference*.

Update JCL: RUNACFU for ACF2

These jobs contain the necessary JCL commands to UPDATE your E-SRF Masterfile with the requested security events from your RSS.

The following is a copy of the job, **RUNACFU.** Each line in the JCL is identified with a number. Following the JCL is an explanation of each line. Refer to the number of the line to find the appropriate definition.

The sample jobs are set up to run four update commands. A "DD" statement defining the SMF file and an UPDATE command to invoke the update function for each desired update is required. Updates may be run in any order.

<u>If you are applying journals from an ACF2 system</u>, copy SAMPLIB job **RUNACFU** into your working library.

The SMFTYPE for ACF2 defaults to 230 (if it is not specified). If your journals are anything other than 230, you MUST supply the SMFTYPE(*nnn*) specification. The SMF record ID for your particular system is "*nnn."*

If an attempt is made to run the UPDATE more than once for a particular SMF file, E-SRF will reject the update. If an update run is executed with no valid input data, return code 20 is presented to the Command Processor.

<u>**Remember**</u>: Make ***your own copy*** of any members before you modify them with information relevant to your operating environment.

RUNACFU:

```
1)     //JOBNAME  JOB (XXXX),'XXXXXXXXX',MSGCLASS=X,MSGLEVEL=(1,1)
2)     //*
3)     //********************************************************************
4)     //*                                                                  *
5)     //*       JCL TO APPLY JOURNAL UPDATE TO ACF2 SECURITY IMAGES.       *
6)     //*                                                                  *
7)     //*       SAMPLIB MEMBER: RUNACFU                                    *
8)     //*------------------------------------------------------------------*
9)     //*                                                                  *
10)    //*       MUST BE CUSTOMIZED FOR YOUR INSTALLATION.                  *
11)    //*                                                                  *
12)    //*       SAMPLE SET UP FOR FOUR SMF INPUT DATASETS.                 *
13)    //*       (YOURS MAY BE ONE OR MORE, CHANGE AS REQUIRED).            *
14)    //*                                                                  *
15)    //*       THIS JOB MUST BE SETUP TO HAVE NO CPU OR STORAGE           *
16)    //*       LIMITATIONS.                                               *
17)    //*                                                                  *
18)    //*       GROUPING RULES NOT NEEDED UNLESS UPDATE EXCLUDE IS         *
19)    //*       REQUIRED FOR THIS JOB.  IGNORE THE GROUPING ERROR.         *
20)    //*                                                                  *
21)    //********************************************************************
22)    //*
23)    //UPDATE  EXEC PGM=ESRFCMD,REGION=0M,TIME=1440
24)    //*
25)    //STEPLIB  DD  DSN=ESRF.LOAD,DISP=SHR         ESRF SYSTEM.
26)    //*
27)    //MASTER   DD  DSN=ESRF.MASTER,DISP=SHR       ESRF MASTERFILE.
28)    //*
29)    //SMF1     DD  DSN=SMF.INPUT.FILE.1,DISP=SHR  SMF JOURNAL FILE 1.
30)    //SMF2     DD  DSN=SMF.INPUT.FILE.2,DISP=SHR  SMF JOURNAL FILE 2.
31)    //SMF3     DD  DSN=SMF.INPUT.FILE.3,DISP=SHR  SMF JOURNAL FILE 3.
32)    //SMF4     DD  DSN=SMF.INPUT.FILE.4,DISP=SHR  SMF JOURNAL FILE 4.
33)    //*
34)    //SYSPRINT DD  SYSOUT=*                       CONTROL REPORT.
35)    //*
36)    //SYSIN    DD  *
37)              CACHE       ON
38)              UPDATE      RSS(ACF2) DDNAME(SMF1) SMFTYPE(230)
39)              UPDATE      RSS(ACF2) DDNAME(SMF2) SMFTYPE(230)
40)              UPDATE      RSS(ACF2) DDNAME(SMF3) SMFTYPE(230)
41)              UPDATE      RSS(ACF2) DDNAME(SMF4) SMFTYPE(230)
```

JCL Line Definitions

1) Standard MVS “JOB” statement. Code this to your installation’s specifications.

2 - 22) Standard MVS “comment statements.”

23) Standard MVS “EXECUTE” statement. It indicates the name of the program you are executing. In this case, it is the E-SRF Command Processor ESRFCMD.

24) Standard MVS “comment statement.”

25) Standard MVS “DD” statement. It relates to a file called STEPLIB where the MVS program loader will look for the E-SRF Command Processor and any other E-SRF programs required to execute the E-SRF session. The Load Library containing E-SRF must be specified here. This DD is not required if the E-SRF programs are contained in the MVS LINK LIST. Consult your E-SRF installer for more information.

26) Standard MVS “comment statement.”

27) Standard MVS “DD” statement. It relates to a file called MASTER. This is the E-SRF Masterfile you intend to update.

28) Standard MVS “comment statement.”

29 - 32) Four SMF input datasets. Either remove or add to if necessary. These file(s) contain the SMF journal data produced by the Resident Security System that you are going to use to update your Masterfile.

33) Standard MVS "comment statement."

34) Standard MVS "DD" statement. It relates to a file called SYSPRINT where ESRFCMD writes its control report output.

35) Standard MVS "comment statement."

36) Standard MVS "DD" statement. It relates to a file called SYSIN where ESRFCMD reads its command input. In this case, the input data is in line with the JCL.

37) CACHE specification. **Always** set the CACHE **ON** when running an UPDATE. This will be explained later in this publication. Refer to the *Command Reference* for more information about the CACHE command.

38 - 41) UPDATE Commands. One is required for each SMF input file. You can run multiple updates for one or more IMAGES in a single execution.

This page intentionally left blank

Chapter 5: Be Productive *Today* (with RACF)

This chapter is intended to take you from defining your Masterfile to being ready to produce sample useable reports when you are using RACF as your Resident Security System. For a discussion on ACF2, please review the previous chapter.

Step 2: *Gathering RSS Dependent Data for RACF*

Discussion

The E-SRF Event System requires information about the Resident Security System (RSS) that will be reported on. The program **ESRFRACP** is provided to build a parameter file that is required for the remaining jobs in this chapter. This job is run on the system that E-SRF is to report on. Its purpose is for E-SRF to *learn* about the system.

Gathering RSS Dependent Data JCL - RUNRACP

On RACF systems, this job is required. Use member: **RUNRACP**

E-SRF Event Reporting uses a single Masterfile to report on multiple systems, which may or may not be secured by the same security system. E-SRF must “learn” about each system individually. In this case, your intention is to configure a RACF IMAGE. You may be doing this for the system on which you are currently running E-SRF or a on a different system. The system you are running on may not even be a RACF system.

RACF may be configured in different ways on different systems. In addition, current implementations of E-SRF cannot *learn* directly how the current RACF IMAGE (the one you are currently running on) is to be configured. This information must be extracted via running ESRFRACP. This includes the local system that is hosting this execution.

The characteristics of this file are indicated in the sample JCL.

RUNRACP is the sample job to run to extract the information from a running RACF system.

The following is a copy of the job **RUNRACP.** Each line in the JCL is identified with a number. Following the JCL is an explanation of each line. Refer to the number of the line to find the appropriate definition.

Remember: Make your own copy of any members before you modify them with information relevant to your operating environment.

RUNRACP:

```
1)     //JOBNAME  JOB (XXXX),'XXXXXXXXX',MSGCLASS=X,MSGLEVEL=(1,1)
2)     //*
3)     //********************************************************************
4)     //*                                                                  *
5)     //*       RUN ESRFRACP -- RACF PARAMETERIZATION PROGRAM              *
6)     //*                                                                  *
7)     //*------------------------------------------------------------------*
8)     //*                                                                  *
9)     //*       THIS IS AN OFFLINE BATCH UTILITY THAT MAY BE USED          *
10)    //*       TO EXTRACT INFORMATION FROM THE RACF SYSTEM BEING          *
11)    //*       CONFIGURED AS AN IMAGE TO E-SRF.                           *
12)    //*                                                                  *
13)    //*------------------------------------------------------------------*
14)    //*                                                                  *
15)    //*       THIS PROGRAM LOOKS AT A RACF SYSTEM AND USES THE CURRENT   *
16)    //*       RACF ESTABLISHED PARAMETERS TO BUILD A PARAMETER FILE      *
17)    //*       THAT MAY BE TRANSPORTED TO THE SYSTEM WHERE THE IMAGE IS   *
18)    //*       BEING CONFIGURED.                                          *
19)    //*                                                                  *
20)    //*------------------------------------------------------------------*
21)    //*                                                                  *
22)    //*       PARMOUT:  MUST BE A 255 BYTE FIXED FILE LARGE ENOUGH TO    *
23)    //*                 HOLD TWO HUNDRED RECORDS.                        *
24)    //*                                                                  *
25)    //********************************************************************
26)    //*
27)    //RUNACFP EXEC PGM=ESRFACFP
28)    //STEPLIB  DD  DSN=ESRF.LOAD,DISP=SHR        ESRF SYSTEM.
29)    //PARMOUT  DD  DSN=ESRF.ACF2PARM,DISP=SHR    PARAMETER FILE.
30)    //SYSPRINT DD  SYSOUT=*                      CONTROL REPORT.
```

JCL Line Definitions for RUNRACP

1) Standard MVS “JOB” statement. Code this to your installation’s specifications.

2 - 26) Standard MVS “comment statements.”

27) Standard MVS “EXECUTE” statement. It indicates the name of the program you are executing. In this case, it is the E-SRF program **ESRFRACP**.

28) Standard MVS “DD” statement. It relates to a file called STEPLIB where the MVS program loader will look for E-SRF programs. This includes the Command Processor and any other E-SRF programs required to execute the E-SRF session. The Load Library containing E-SRF must be specified here. (This DD is not required if the E-SRF programs are contained in the MVS LINKLIST.) Consult your E-SRF installer for more information.

29) Standard MVS “DD” statement. It relates to the file produced by the program **ESRFRACP** that builds parameters from RACF for input into E-SRF. This is the input that will “teach” E-SRF about the current RACF IMAGE being configured. *For RACF, this file is required for E-SRF IMAGE configuration.*

30) Standard MVS “DD” statement. It relates to a file called SYSPRINT where **ESRFRACP** writes its printed output.

Output file characteristics

ESRFRACP creates an output file that becomes input to the IMAGE Configuration process. The characteristics of this file are as follows:

The file should be large enough to contain several hundred 255-character data records. Any suitable blocksize may be specified.

The records may be edited with ISPF to alter the standard descriptions of the fields contained on the User Header object, if needed.

The following JCL may be incorporated in lieu of PARMOUT (found on line 29 of the example).

```
29)    //PARMOUT     DD     DSN=ESRF.RACFPARM,DISP=(NEW,CATLG),
       //                   UNIT=SYSDA,VOL=SRR=VOLID,
       //                   DCB=(LRECL=255,BLKSIZE=7650,RECFM=FB),
       //                   SPACE=(7650,(10,5))
```

Porting the data to the hosting system

ESRFRACP creates an output file that must be “ported” to the system that hosts the E-SRF Masterfile. This file is very small. Any suitable means to transport the data may be used. Most 3270 terminal emulators have a file download/upload facility. FTP (File Transfer protocol) may also be used. The only requirement is the data must be present for the Masterfile IMAGE configuration. After that, the file is no longer referenced.

Step 3: RACF Image Configure, Assigning Domains, and Synchronize

Discussion

Once an E-SRF Masterfile has been defined and initialized, it is time to tie together the associations of Domains, Images, and RSS specific data. The **RUNRACI** job found in the sample library will assist you with this process.

CONFIGURE

The CONFIGURE command builds an E-SRF IMAGE and assigns a Resident Security System (RSS) to the IMAGE. For example, if you are creating an IMAGE called CHICAGO, which uses RACF as its RSS, and the IMAGE contains multiple LPARs (domains), you would use the CONFIGURE command below:

```
      CONFIGURE     IMAGE(MVSRACF)      -
                    DDNAME(RACFPARM)    -
                    PARAMETER()         -
                    NAME(RACF SYSTEM)   -
                    RSS(RACF)    -
                    RESET
```

This command will allow E-SRF to “learn” about the RACF RSS the IMAGE represents. This is accomplished by processing the file created by the **ESRFRACP** batch utility. This utility was executed on the system using the RACF security database that you are creating the IMAGE for. The job is executed and the output data is ported to the system running ESRF.

A “*DD*” JCL statement must be provided naming the dataset name to a unique *DDNAME*. This DDNAME is also specified in the DDNAME(*ddname*) CONFIGURE specification as shown above.

SYNCHRONIZE

The SYNCHRONIZE command takes a copy of the userids present on the RACF security database and applies them to the E-SRF Masterfile so this information can be used for reporting. When configuring an IMAGE, it should be followed up with a SYNCHRONIZE to bring the Masterfile up to date with the userid data contained on the security database.

You can print the user’s name or any other fields on a report or use fields as selection criteria since that information is contained in the E-SRF Masterfile.

An example of the SYNCHRONIZE command is shown below:

```
SYNCHRONIZE         DDNAME(backup Logonid flat file DDNAME) -
                    IMAGE(CHICAGO)
```

Where to Obtain the RACF Data for SYNCHRONIZE

You must run an additional step that unloads the RACF database belonging to the system you are configuring into a flat file using a RACF supplied utility, 'IRRDBU00'. This utility is a standard component of RACF and is documented in the RACF product documentation.

This file will be used as input to the SYNCHRONIZATION command for the target RACF IMAGE. When running this program, it is recommended that the following parameter be placed on the execute card: PARM=NOLOCKINPUT.

The *SAMPLIB* member 'RACFUNLD' contains a sample job to execute this utility. The output dataset is a variable blocked (RECFM=VB), and a 4092 byte record length (LRECL=4092) works.

Note About RACF Synchronization

RACF does not create composite journals of its security database update. This means E-SRF must individually, field-by-field, track all database updates (*which are maintained on the E-SRF Masterfile*). Most fields are accounted for at this product release level, but there are some that may not be, or the journalizing is not adequate for proper maintenance of the User Header Masterfile object. Product development goals will insure all fields are properly accounted for in future enhancements to E-SRF.

Until such time, it is recommended that a SYNCHRONIZE be executed some time after the Masterfile Update Function completed and the execution of any report that references fields that begin with RACF.*xxxxxxx*. These fields are contained on the User Header (UA) object.

If in doubt, assume you can run a SYNCHRONIZE any time after an Update Function, provided you supply it the correct data (*a proper RACF database unload data file as described above*) and do it for the correct IMAGE.

ASSIGN

The ASSIGN statement is how E-SRF knows what IMAGE to place data from a particular domain.

You must provide the necessary ASSIGN statement to associate your DOMAIN(S) to your IMAGE. In the sample job, we assume there is only one IMAGE, and we provide an ASSIGN statement that associates all domains to an image called MVSRACF. This will only work if there is one IMAGE. If you omit this step, all the data will go to whatever IMAGE assignment matches the domain id when presented during an Update Function.

Assignments are *maskable*. A domain ID is the identification of the actual system (or LPAR) that created a journal. Journal files may contain journals from multiple domains, but they must all be from the same RSS (in this case, RACF). The Update Function will update multiple domains, but only to a specific RSS (*in this case, RACF*). If you have multiple RSS systems, such as RACF and ACF2, you would have to run an Update Function for RACF IMAGES and another for ACF2 IMAGES.

You may have existing assignments that have to be changed. Look over all assignments when adding or changing IMAGES to make sure you are directing your journals to their proper IMAGES.

Configuring, Synchronizing, and Assigning JCL - RUNRACI

The **RUNRACI** job (found in SAMPLIB) contains the necessary JCL commands to CONFIGURE your IMAGE, SYNCHRONIZE the E-SRF Masterfile to your RACF system, and ASSIGN the applicable DOMAINS to the newly created IMAGE. Each line in the JCL is identified with a number. Following the JCL is an explanation of each line. Refer to the number of the line to find the appropriate definition.

Consult the *Command Reference* for assistance with the syntax editing.

Remember: Make ***your own copy*** of any members before you modify them with information relevant to your operating environment.

RUNRACI:

```
1)     //JOBNAME  JOB (XXXX),'XXXXXXXXX',MSGCLASS=X,MSGLEVEL=(1,1)
2)     //*
3)     //********************************************************************
4)     //*                                                                  *
5)     //*       JOB TO CREATE AN RACF IMAGE ON THE ESRF MASTERFILE.        *
6)     //*                                                                  *
7)     //*       SAMPLIB MEMBER: RUNRACI                                    *
8)     //*------------------------------------------------------------------*
9)     //*                                                                  *
10)    //*       NOTE: --- PRIOR TO RUNNING THIS JOB:                       *
11)    //*                                                                  *
12)    //*    1) YOU MUST HAVE AN EXISTING MASTERFILE.                      *
13)    //*                                                                  *
14)    //*    2) MAKE SURE DD STATEMENTS ARE VALID FOR YOUR INSTALLATION.   *
15)    //*                                                                  *
16)    //*    3) REVIEW ALL OPTIONS ARE APPLICABLE TO YOUR INSTALLATION.    *
17)    //*                                                                  *
18)    //*    4) YOU MUST HAVE THE DATASET CONTAINING THE OUTPUT            *
19)    //*       CREATED BY 'RUNRACP'.                                      *
20)    //*                                                                  *
21)    //*    5) YOU MUST HAVE THE FLAT FILE CONTAINING THE OUTPUT FROM     *
22)    //*       THE RACF DATABASE UNLOAD UTILITY.                          *
23)    //*                                                                  *
24)    //*------------------------------------------------------------------*
25)    //*                                                                  *
26)    //*       IGNORE THE GROUPING FACILITY ERROR.                        *
27)    //*                                                                  *
28)    //********************************************************************
29)    //*
30)    //INIT    EXEC PGM=ESRFCMD,REGION=0M
31)    //*
32)    //STEPLIB  DD  DSN=ESRF.LOAD,DISP=SHR        ESRF SYSTEM.
33)    //MASTER   DD  DSN=ESRF.MASTER,DISP=SHR      ESRF MASTERFILE.
34)    //*
35)    //RACFPARM DD  DSN=ESRF.RACFPARM,DISP=SHR    ESRFRACP PARAMETERS.
36)    //EXTRACT  DD  DSN=SYS1.ESRF.UNLOAD,DISP=SHR RACF EXTRACT FILE.
37)    //SYSPRINT DD  SYSOUT=*                      CONTROL REPORT.
38)    //*
39)    //SYSIN    DD  *

40)       /*      THE FOLOWING COMMANDS WILL CONFIGURE AND SYNCHRONIZE A
41)               NEW RACF IMAGE ON YOUR EXISTING MASTERFILE.

42)               IT IS ASSUMED THAT THIS IS THE ONLY IMAGE THAT WILL BE
43)               CONTAINED ON YOUR MASTERFILE.  IF THIS IS NOT THE CASE,
44)               YOU WILL HAVE TO ADJUST YOUR ASSIGN STATEMENTS TO SELECT
44)               THE DOMAINS (MVS SYSIDS) YOU WANT TO BE INCLUDED IN THIS
45)               IMAGE.  (CONSULT COMMAND GUIDE FOR MORE INFORMATION).
46)      */

47)               CACHE       ON
```

```
48)            INITIALIZE  RESET

49)            CONFIGURE   IMAGE(MVSRACF)        -
50)                        DDNAME(RACFPARM)      -
51)                        PARAMETER()           -
52)                        NAME(RACF SYSTEM)     -
53)                        RSS(RACF)             -
54)                        RESET

55)            SYNCHRONIZE IMAGE(MVSRACF) DDNAME(EXTRACT)

56)            ASSIGN      DOMAIN(********) IMAGE(MVSRACF) -
57)                        NAME(RACF SECURITY IMAGE)

58)            RUN         REPORT(ESRFSHOW)
59)            RUN         REPORT(ESRFDICT)
60)            RUN         REPORT(ESRFSTAT)
```

JCL Line Definitions for RUNRACI

1) Standard MVS "JOB" statement. Code this to your installation's specifications.

2 - 29) Standard MVS "comment statements."

30) Standard MVS "EXECUTE" statement. It indicates the name of the program you are executing. In this case, it is the E-SRF Command Processor ESRFCMD.

31) Standard MVS "comment statement."

32) Standard MVS "DD" statement. It relates to a file called STEPLIB where the MVS program loader will look for E-SRF. This includes the Command Processor and any other E-SRF programs required to execute the E-SRF session. The Load Library containing E-SRF must be specified here. This DD is not required if the E-SRF programs are contained in the MVS LINK LIST. Consult your E-SRF installer for more information..

33) Standard MVS "DD" statement. It relates to a file called MASTER. This is the E-SRF Masterfile created previously.

34) Standard MVS "comment statement."

35) Standard MVS "DD" statement. It relates to a file called RACFPARM. This file was created in job **RUNRACP**.

36) Standard MVS "DD" statement. It relates to a file called "*EXTRACT*" but could have been any DDNAME that was specified in the SYNCHRONIZE command. This is the input used to perform the SYNCHRONIZE function for a particular security *IMAGE*. ***For RACF***, it is the flat file output from a prior execution of the RACF offline database unload utility (IRRDBU00) and it contains information from your RACF database. You create this file using the utility supplied by RACF. Please refer to the RACF product documentation for information on the execution of this utility.

37) Standard MVS "DD" statement. It relates to a file called SYSPRINT where program ESRFCMD writes its control report output.

38) Standard MVS "comment statement."

39) Standard MVS "DD" statement. It relates to a file called SYSIN where program ESRFCMD reads its command input. In this case, the input data is in line with the JCL.

40 - 46) E-SRF Command Processor "comment statements."

47) E-SRF command to turn the Masterfile Caching option ON. CACHE OFF may also be specified for those sites not using the Caching option. It is highly recommended that you run with CACHE ON for this process.

48) E-SRF major command to initialize the Masterfile. The "reset" minor command will reset the data dictionary if there is an existing one.

49) E-SRF major command to CONFIGURE the Masterfile for a specific security IMAGE. This statement contains the major command and identifies the security IMAGE to be configured. "*MVSRACF*" is the IMAGE chosen in this example. You can call yours anything you desire.

50) This statement identifies the DDNAME of a dataset containing the configuration parameters from the Resident Security System being configured to the image. The file is created by the ESRFRACP stand alone batch program. In this case, "DD" name RACFPARM matches the "DD" name RACFPARM in the RUNRACP job previously submitted.

51) Optional configuration parameter. The value specified is dependent on the particular RSS being configured to the target security IMAGE. Refer to the *Command Reference Guide* for more information. Currently, this parameter should be set as shown in this example.

52) This statement identifies the NAME of the security IMAGE being configured. In this example we have chosen "RACF SYSTEM" as the name of the security IMAGE.

53) This statement identifies the particular RSS being configured to the target security IMAGE. In this example, the RSS is RACF.

54) The RESET attribute indicates that if a previous IMAGE (MVSRACF) already exists, it will be reset to the specifications declared here.

55) This major command will SYNCHRONIZE the IMAGE (MVSRACF) to the specified Resident Security System. Each RSS has its own requirements. This example is for RACF, which uses the RACF VSAM extract created in EKCRRCDB as input. This is reflected by the DDNAME specified on this statement.

56) Major command to ASSIGN all DOMAINS, designated by the "xxxxxxxx", to the target IMAGE (MVSRACF).

57) Name of the assignment. We chose "RACF SECURITY IMAGE" for this example. All domains in this IMAGE will match this assignment. All activity updated against the Masterfile will end up in the IMAGE MVSRACF.

58) E-SRF major command to RUN the ESRFSHOW report overlay. It lists the current E-SRF processing options.

59) E-SRF major command to RUN the ESRFSTAT report overlay. It provides a byte count of the data contained on the Masterfile by E-SRF segment and object. This report can be useful when determining where the majority of your security event records are being stored and where the majority of the security violations and/or loggings are taking place.

60) E-SRF major command to RUN the ESRFDICT report overlay. It provides an alphabetical list of all E-SRF data names and object names contained on the E-SRF data dictionary. For more information about the E-SRF data dictionary, see *E-SRF Data Dictionary Structure* in this guide and the *Masterfile and Data Dictionary Reference Guide.*

Step 4: Updating the Masterfile with Your Data

Discussion

At this point, reports can actually be run. If you run the "**RUNQUICK**" SAMPLIB member, you will end up with userid reports but empty event reports. This is not what you really had in mind. Yes, you have a configured Masterfile full of users, but no security events.

In order for E-SRF to report on events, you must provide Journal Event Information from the particular RSS you are reporting on. The act of placing the journals on E-SRF is referred to as "updating the Masterfile," or more technically the "UPDATE FUNCTION." There are entire chapters of E-SRF publications dedicated to this topic. For now, just enough will be reviewed to allow us to get it up and running.

After the E-SRF Masterfile has been **defined, configured, and synchronized**, it must be **updated** to include records on security events that occurred on your system. This is accomplished by applying your daily security journals (SMF data) to your Masterfile. The job RUNRACU in the SAMPLIB contains the JCL to do this. There is also another job for ACF2. You are interested in RACF, so use RUNRACU. If you have RACF and ACF2, they can live on the same Masterfile provide they are in separate IMAGES. Process each update separately using the same Masterfile. When you publish reports, all security events may be treated as a single set of data.

There is little difference (externally) between one security system and another when setting up the UPDATE Function. Each security system has its own Update Overlay that "*normalizes*" the journalized data into something common for E-SRF's Masterfile Update Processor to use when placing events on the Masterfile.

The only hard and fast rule is to declare which security system to process during a single UPDATE execution. You may have as many UPDATE commands in a single run as desired. The SMF data may be supplied in any sequence. You can update today's first, last week's and then yesterday's. E-SRF will sort out and collate the journals from various UPDATE runs. The end result will be a seamless recording of security events across whatever span of time that was declared when E-SRF was configured. Old data will be automatically rolled off the Masterfile.

It is possible to present too much data for E-SRF to store. Limits are established when E-SRF is configured, and in the event a particular limit was exceeded, the data will be '*rolled off*", oldest first. This may seem brutal at first, but experience has shown that the vast majority of the data rolled off this way consists of excessive logging that are really not desired. The limits were set to keep you from "shooting yourself in the foot" with large reports, huge DISK and STORAGE requirements and unusable reports.

All of this can be examined and controlled after operational experience has been gained using this product. Ultimately, when things come together, you will have a reporting system tailored to your auditing and reporting needs, containing only the information you desire to have reported, and delivered to the data owners deemed responsible for reviewing.

Updating a RACF IMAGE

UPDATE

The UPDATE command applies SMF journal data about security events to the E-SRF Masterfile. This command should be run periodically to update the Masterfile with recent security events. We recommend you schedule the UPDATE job to run nightly after the SMF extract job.

```
UPDATE      RSS(RACF)   SMFTYPE(80)
```

This command is explained in more detail in the *E-SRF Event Reporting System Command Reference*.

Update JCL: RUNRACU

RUNRACU contains the necessary JCL commands to UPDATE your E-SRF Masterfile with the requested security events from your RSS.

The following is a copy of the job, **RUNRACU.** Each line in the JCL is identified with a number. Following the JCL is an explanation of each line. Refer to the number of the line to find the appropriate definition.

The sample jobs are set up to run four update commands. A "DD" statement defining the SMF file, and an UPDATE command to invoke the update function for each desired update is required. Updates may be run in any order.

The SMFTYPE for RACF defaults to 80 (if it is not specified). If your journals are anything other than 80, you MUST supply the SMFTYPE(*nnn*) specification. "*nnn"* is the SMF record ID for your particular RACF system.

If an attempt is made to run the UPDATE more than once for a particular SMF file, E-SRF will reject the update. If an update run is executed without valid input data, return code 20 is presented to the Command Processor.

Important Note About RACF Journals

RACF does not provide system entry validation (signon) journals for TSO logons or batch job submission. This information must come from other SMF journals. Additionally, if you have the EKC ETF/R product installed, its FIRECALL records are written to User Journal 230.

Despite the fact that RACF only creates SMF type 80 journals, the Update Function it will also process type 230 (looking only for specific FIRECALL records).

Additionally, either TYPE 20 or TYPE 30 records are used to provide signon loggings. You have a choice whether to use type 20 or type 30 records. The E-SRF will default to type 30 because of the quality of the data.

If you want the above-mentioned data, make sure the SMF input file supplied in the Update Function contains the record types needed to provide the data.

If TYPE 20 records are processed (NOT the default), all type 20 records are required.

If TYPE 30 records are processed (recommended mode of processing), most are skipped and do not have to be present on the SMF input file. E-SRF needs the STEP TOTALS (SMF30STP=x’0004’) only for the first step for each job (SMF30STN=x’0001’). All others may be skipped.

Failure to supply this data simply means E-SRF will not report on the missing data.

For more information on this topic, please refer to the *E-SRF Event Reporting System: Command Reference* under the UPDATE function topic.

Remember: Make ***your own copy*** of any members before you modify them with information relevant to your operating environment.

RUNRACU:

```
//JOBNAME  JOB (XXXX),'XXXXXXXXX',MSGCLASS=X,MSGLEVEL=(1,1)
//*
//********************************************************************
//*                                                                  *
//*       JCL TO APPLY JOURNAL UPDATE TO RACF SECURITY IMAGES.       *
//*                                                                  *
//*       SAMPLIB MEMBER: RUNRACF                                    *
//*------------------------------------------------------------------*
//*                                                                  *
//*       MUST BE CUSTOMIZED FOR YOUR INSTALLATION.                  *
//*                                                                  *
//*       SAMPLE SET UP FOR FOUR SMF INPUT DATASETS.                 *
//*       (YOURS MAY BE ONE OR MORE, CHANGE AS REQUIRED).            *
//*                                                                  *
//*       THIS JOB MUST BE SETUP TO HAVE NO CPU OR STORAGE           *
//*       LIMITATIONS.                                               *
//*                                                                  *
//*       GROUPING RULES NOT NEEDED UNLESS UPDATE EXCLUDE IS         *
//*       REQUIRED FOR THIS JOB.  IGNORE THE GROUPING ERROR.         *
//*                                                                  *
//********************************************************************
//*
//UPDATE  EXEC PGM=ESRFCMD,REGION=0M,TIME=1440
//*
//STEPLIB  DD  DSN=ESRF.LOAD,DISP=SHR         ESRF SYSTEM.
//*
//MASTER   DD  DSN=ESRF.MASTER,DISP=SHR       ESRF MASTERFILE.
//*
//SMF1     DD  DSN=SMF.INPUT.FILE.1,DISP=SHR  SMF JOURNAL FILE 1.
//SMF2     DD  DSN=SMF.INPUT.FILE.2,DISP=SHR  SMF JOURNAL FILE 2.
//SMF3     DD  DSN=SMF.INPUT.FILE.3,DISP=SHR  SMF JOURNAL FILE 3.
//SMF4     DD  DSN=SMF.INPUT.FILE.4,DISP=SHR  SMF JOURNAL FILE 4.
```

```
33)    //*

34)    //SYSPRINT DD  SYSOUT=*                       CONTROL REPORT.
35)    //*
36)    //SYSIN    DD  *
37)              CACHE       ON
38)              UPDATE      RSS(RACF) DDNAME(SMF1) SMFTYPE(80)
39)              UPDATE      RSS(RACF) DDNAME(SMF2) SMFTYPE(80)
40)              UPDATE      RSS(RACF) DDNAME(SMF3) SMFTYPE(80)
41)              UPDATE      RSS(RACF) DDNAME(SMF4) SMFTYPE(80)
```

JCL Line Definitions

1) Standard MVS “JOB” statement. Code this to your installation’s specifications.

2 - 22) Standard MVS “comment statements.”

23) Standard MVS “EXECUTE” statement. It indicates the name of the program you are executing. In this case, it is the E-SRF Command Processor ESRFCMD.

24) Standard MVS “comment statement.”

25) Standard MVS “DD” statement. It relates to a file called STEPLIB where the MVS program loader will look for the E-SRF Command Processor and any other E-SRF programs required to execute the E-SRF session. The Load Library containing E-SRF must be specified here. This DD is not required if the E-SRF programs are contained in the MVS LINK LIST. Consult your E-SRF installer for more information.

26) Standard MVS “comment statement.”

27) Standard MVS “DD” statement. It relates to a file called MASTER. This is the E-SRF Masterfile you intend to update.

28) Standard MVS “comment statement.”

29 - 32) Four SMF input datasets. Either remove or add to if necessary. These file(s) contain the SMF journal data produced by the Resident Security System that you are going to use to update your Masterfile.

33) Standard MVS “comment statement.”

34) Standard MVS “DD” statement. It relates to a file called SYSPRINT where ESRFCMD writes its control report output.

35) Standard MVS “comment statement.”

36) Standard MVS “DD” statement. It relates to a file called SYSIN where ESRFCMD reads its command input. In this case, the input data is in line with the JCL.

37) CACHE specification. **Always** set the CACHE **ON** when running an UPDATE. This will be explained later in this publication. Refer to the *Command Reference* for more information about the CACHE command.

38 - 41) UPDATE Commands. One is required for each SMF input file. You can run multiple updates for one or more IMAGES in a single execution. Either remove or add if necessary.

Chapter 6: Sample Reports With Your Data

Step 5: The Quick Start

Discussion

In order to run the Quick Start Reports, you must have initialized, configured, and synchronized your Masterfile and applied some Resident Security System Event data through the UPDATE function. Attempting to run the Quick Start job without doing this first will yield empty reports.

The Quick Start Report execution job output provides you with an example of E-SRF Event reporting. The job will make date selections and assumes you have updated your Masterfile with *yesterday's* event data. It also assumes that you may update one or more days of data and may have stopped updating the Masterfile so you can get accustomed to the reports. This process may continue for several weeks, until you have become familiar with the product. This may not seem like it would be a problem, but when one considers that date selection is one of the most common selection criteria, it could get confusing when you report on the same data day after day.

It is also assumed that you may want to keep some of the command input from the sample job to use for "real" reporting when you are finished experimenting with the system. When the Quick Start job was initially put together, it was set up to run as a real production job (i.e., running the reports early in the morning using the previous day's event data). All "date required" selections make this assumption.

Using the RELATIVEDATE Option to Override the System Date

To enable you to experiment, an optional command is placed in front of the command input that "freezes" the base date that is used for report data selection:

```
OPTION      RELATIVEDATE(LASTRANS)
```

This sets the relative base date to "LASTRANS" (*the most recent transaction date*). This allows you to run reports day after day with date specifications and still get the same reports while you are experimenting (*without having to run subsequent Update Functions and dealing with changing data*).

When you want to run these reports on a daily basis with actual daily updated events, simply remove the RELATIVEDATE command statement from the command input. This will make the relative base data the same as the job's execution date.

Refer to the *Command Reference* for more information on how dates are used for selections and how they may be overridden.

Quick Start Report JCL: RUNQUICK

The E-SRF Event System sample library contains the job: RUNQUICK. The following is an abbreviated copy of this job for you to review. It contains the necessary JCL commands to run a sampling of some basic E-SRF Event System Reports.

The JCL to run the Quick Start reports is contained on SAMPLIB in a member called "RUNQUICK." Each report will be produced in its own print DD so you can easily locate the printed output.

A list of the reports contained in RUNQUICK, their DDNAMES, and the names of the report overlays that created them can be found on the following page.

Once you have run the Quick Start reports several times, you may want to add additional reports or remove reports from the job. This often becomes the production job that is run.

SAMPLE REPORTS With Your Data

Some of the reports contained in the Quick Start job use ESRFLIST, which also may be used to produce numerous other reports (see *Customizing Reports and Selecting Data* later in this chapter). Some of the reports are “canned” reports that cannot be changed other than the data being selected to report on.

Reports Available Through the RUNQUICK Job

The following is a list of the reports that are produced by the RUNQUICK job. The “Specified DD Name” is the report output file for the particular report.

Report Description	Report Overlay	Specified DD Name
Count of Violations/Loggings By Userid	ESRFUVLC	USERVL
Top 20 User Violation Report	ESRFRVU	T20UV
Top 20 User Log Report	ESRFRLU	T20UL
Top 20 Signon Errors by User	ESRFRUSE	T20SEU
Violation/Logging Summary By Resource	ESRFLIST	VLSUMM
Daily Ranked Log Recap By Resource	ESRFRDRE	RSALLOW
Daily Ranked Log Recap By User	ESRFRDUE	USRALLOW
Daily Violation Recap By Resource	ESRFRDRV	BADBOYR
Daily Violation Recap By User	ESRFRDUV	BADBOYU
Violation Logging Detail By Resource	ESRFLIST	LOGDET
Userids That Will Expire Within 30 Days	ESRFLIST	EXP30
System Console Activity Report	ESRFLIST	CONSOLE
Security Changes By Target Resource	ESRFLIST	SECMAINT
Security System Administration Change Summary	ESRFLIST	SECADM
Security Rules Maintenance History	ESRFLIST	RMAINT
User Maintenance Detail History	ESRFLIST	UMAINT
Summary of Loggings For Users	ESRFLIST	LOGSUMM

Remember: Make ***your own copy*** of any members before you modify them with information relevant to your operating environment.

RUNQUICK: Partial JCL Listing of the SAMPLIB member used to produce the sample reports:

```
//JOBNAME  JOB (XXXX),'XXXXXXXXX',MSGCLASS=X,MSGLEVEL=(1,1)

//*
//********************************************************************
//*                                                                  *
//*       SAMPLE REPORT EXECUTION JOB FOR ACF2 AND RACF.             *
//*                                                                  *
//*       SAMPLIB MEMBER: RUNQUICK                                   *
//*------------------------------------------------------------------*
//*                                                                  *
//*   1)  JCL MUST BE CUSTOMIZED FOR YOUR INSTALLATION.              *
//*       ...OTHER CUSTOMIZATION MAY BE DESIRABLE.                   *
//*                                                                  *
//*   2)  EACH REPORT WILL GO TO ITS OWN OUTPUT DATASET.             *
//*       THESE DATASETS WILL BE DYNAMICALLY ALLOCATED TO            *
//*       'SYSOUT=*'.  IF THIS IS DESIRED, LEAVE THINGS ALONE.       *
//*                                                                  *
//*   3)  YOU CAN SUPPLY ANY DD NAME YOU LIKE, CAUSING ONE OR        *
//*       MORE REPORTS TO BE PLACED ON THE SAME OUTPUT FILE.         *
//*       THIS IS DONE IN THE REPORT'S RUN STATEMENT.                *
//*                                                                  *
//*   4)  IF YOU WANT REAL OUTPUT DATASETS, PLACE THE APPROPORATE    *
//*       'DD' JCL STATEMENT(S) REFERENCING THE DDNAME AND THE       *
//*       DESIRED DATASET NAME.                                      *
//*                                                                  *
//*   5)  BECAUSE THIS JOB IS GENERIC BETWEEN ACF2 AND RACF,         *
//*       YOU WILL GET DATANAME ERRORS.  PLEASE IGNORE THEM.         *
//*                                                                  *
//*   6)  GROUPING IS NOT REQUIRED FOR THIS EXECUTION.               *
//*       IF YOU ARE NOT USING GROUPING, IGNORE THE GROUPING ERROR.  *
//*                                                                  *
//*       IF YOU WANT TO USE GROUPING, YOU MUST HAVE A GROUPING      *
//*       RULES OBJECT FILE AND A 'DD' JCL STATEMENT THAT POINTS     *
//*       'RULES' TO YOUR RULES OBJECT FILE.  ALSO, REMOVE           *
//*       THE 'OPTION GROUPING(NONE)' FROM THE COMMAND INPUT.        *
//*                                                                  *
//*   7)  REVIEW THE COMMAND INPUT 'OPTIONS' TO SEE IF THEY          *
//*       APPLY TO THE WAY YOU ARE RUNNING YOUR JOB.                 *
//*                                                                  *
//********************************************************************
//*
//QUICK   EXEC PGM=ESRFCMD,REGION=0M
//*
//STEPLIB  DD  DSN=ESRF.LOAD,DISP=SHR        ESRF SYSTEM
//MASTER   DD  DSN=ESRF.MASTER,DISP=SHR      ESRF MASTERFILE
//SYSPRINT DD  SYSOUT=*                      CONTROL REPORT
//*
//SYSIN    DD  *                             COMMAND STATEMENTS.

**********************************************************************
*                                                                    *
*        ESRF SAMPLE REPORT EXECUTION.                               *
*                                                                    *
*--------------------------------------------------------------------*
*                                                                    *
*        THE FOLLOWING  COMMANDS WILL PRINT REPORT SAMPLES.          *
*                                                                    *
*        THESE REPORTS REPRESENT A SAMPLING OF EVENT REPORTING.      *
*                                                                    *
*        THESE REPORTS ARE SET UP TO PRINT DATA FOR THE TWENTY-FOUR  *
*        HOUR PERIOD OF TIME AFTER THE MOST RECENT EVENT UPDATED     *
*        ON THE MASTERFILE.                                          *
*                                                                    *
**********************************************************************
```

```
          /* -----> THE FOLLOWING OPTION COMMAND TURNS OFF E-SRF GROUPING
                    REMOVE IT WHEN YOU START SETTING UP YOUR ESRF
                    GROUP STRUCTURES AND HAVE WRITTEN EXTERNAL GROUPING
                    FACILITY RULES                                  */

                    OPTION  GROUPING(NONE)

          /* -----> THE FOLLOWING OPTION COMMAND SETS THE CURRENT RELATIVE DATE
                    TO THE DATE OF THE MOSE RECENT EVENT TRANSACTION POSTED ON
                    YOUR MASTERFILE.  THIS WILL ALLOW YOU TO RUN A SINGLE
                    UPDATE AND PROCESS REPORTS FOR MULTIPLE DAYS AFTERWARDS
                    WITHOUT HAVING TO CHANGE ALL THE RELATIVE DATES  */

                    OPTION  RELATIVEDATE(LASTRANS)

          /*        COUNT OF VIOLATIONS/LOGGINGS BY USERID          */

                    RUN     REPORT(ESRFUVLC)     -
                            DD(USERVL)           -
                            WHEN(US.DATE EQ *-1)
```

To view the remaining commands in the job, refer to the RUNQUICK member in the SAMPLIB provided.

Customizing Reports and Selecting Data

ESRFLIST and ESRFDXD are specialized E-SRF Report Overlays that may be customized by simply editing the input statements to create hundreds of new reports to meet specific needs. All E-SRF reports can be reduced by using selection criteria (IF and WHEN/OR/AND) or by grouping reports to specific areas of responsibility. For more information about record selection and grouping, see the *Report Overlays Guide*.

The *Event System Command Reference* explains all the commands needed for both the report execution and data selection.

The *Event System Report Overlays Guide* shows examples of reports that may or may not appear in the Quick Start job. Please refer to this publication to find out about the many reports that are available to you and how to set them up for execution.

Chapter 7: The Masterfile - Other Important Matters

Masterfile Record Retention

Now that you have created a VSAM cluster for the Masterfile, as well as populating it with your security event SMF data, how long do you want to retain the information on the Masterfile?

Remember, SMF backups are kept according to your installation's procedural requirements. The E-SRF Masterfile information is stored for the purpose of creating security event reports based on your needs coupled with the expense of storing and maintaining the data.

Information is expired and removed automatically based on your controls. Additionally, unexpired security events may roll off due to the threshold set for an object's capacity, which is also something you control. This latter situation is known as an "unexpired rolloff."

Once you are more familiar with E-SRF and the Masterfile, you may also want to exclude some security event data from your Masterfile. For example, you may have no need to see anything at the "*source*" level. If this is the case, set the source segment's recap object retention to zero, which will conserve disk space and reduce execution time.

Rebuilding the Masterfile

Rebuilding the Masterfile may be required for one or more of the following reasons:

1) The Masterfile is either too large or too small to contain the data you require.

2) The Masterfile has become fragmented with wasted space due to numerous Control Interval and Control Area splits. This is a normal VSAM issue that consumes disk space and impacts overall performance.

3) You want to physically move the Masterfile to another DASD volume.

To rebuild or reorganize your VSAM Masterfile cluster, use the same procedures used for any other VSAM KSDS cluster. Use the IDCAMS utility to "REPRO" the cluster to a flat backup file. Redefine the VSAM cluster (if necessary), and "REPRO" the flat backup file back into the redefined VSAM cluster.

Reorganizing Your Masterfile VSAM Cluster

As with any VSAM cluster, reorganization may become necessary. With E-SRF, if you use the CACHE option when altering the Masterfile (such as running the UPDATE command), you should **never** have to reorganize your cluster.

When the Masterfile Control Program applies the CACHE to the cluster, a decision is made whether to update the current VSAM record structure in place or totally rebuild the cluster.

When the cluster is <u>rebuilt</u> by the CACHE upgrade function, the data contained on the file is automatically reorganized as if an IDCAMS "REPRO" occurred.

If the CACHE is in effect, and ten percent or more of the cluster requires adding or updating, the cluster is automatically rebuilt.

Backing up your E-SRF Masterfile

Your E-SRF Masterfile is a critical resource and should be backed up like any other important VSAM cluster.

This page intentionally left blank

Chapter 8: Event Reporting Masterfile Organization

Introduction

This chapter describes the organization of the E-SRF Event Reporting Masterfile. The information provided is a basic description intended to give you a general overview of how your security event data is stored on the Masterfile. For a much more detailed explanation, please refer to the *E-SRF Masterfile and Data Dictionary Reference*, which is to be considered “the final authority” on this topic.

The topics in this chapter include information relating to:

- Events grouped by their relationship to the resources affected.
- Masterfile organization into Segments, Objects, and Datanames.
- Objects identified by their names.
- Processing options that determine how the Masterfile maintains events and processes reports.
- How selecting event information affects the selection of Masterfile and Reporting options.

The E-SRF Masterfile is a central repository that contains the data used for E-SRF Event Reporting. This repository is stored on a VSAM cluster in a relational manner.

The Masterfile includes security events recorded by the Resident Security System (RSS). This data may have been extracted from the MVS System Management Facility (SMF) journals. During the Masterfile UPDATE process, all security events are separated into individual objects for **people**, **places** and **things**. The events are stored on the Masterfile based on how the data is classified by E-SRF as described in this chapter.

Organization

Any information stored on the Masterfile may be used to create reports. You simply “tell” E-SRF what you want to see by giving E-SRF the *name* of the place where that information is stored. E-SRF divides up the information into individual items of data and stores them in groupings with other related information. Those groupings are called IMAGES, DOMAINS, SEGMENTS and OBJECTS.

The following sections detail the structure of the Masterfile and how that structure is used to identify which information is included in a report.

IMAGES

A Security Image is the largest Masterfile *grouping* in the Event Reporting system. IMAGE refers to the users and resources that share a common set of security databases.

Everyone should initially start out with one IMAGE to keep it simple.

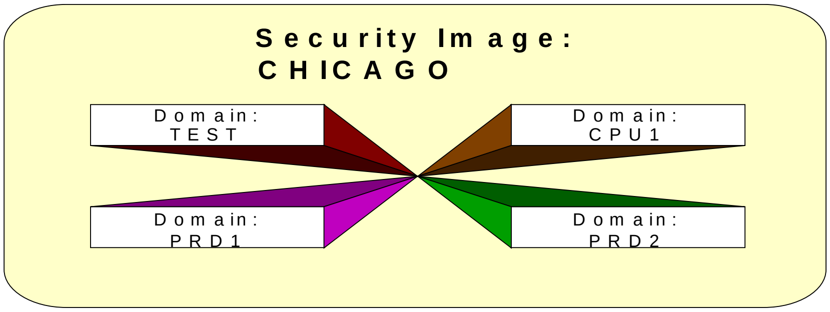

For example, ABC Company has three MVS machines in Chicago: *TEST*, *CPU1,* and *PRODUCTION*. PRODUCTION is divided into two LPARs (Logical PARtitions) - PRD1 and PRD2. There are four different MVS systems referred to as SYSIDs, LPARs, or DOMAINS. All four DOMAINS use the same set of security databases. Therefore, in E-SRF terms, there is one Chicago Security Image.

Each E-SRF security IMAGE includes information about the people, places, and things on the DOMAINS assigned to that IMAGE. All events that take place on those DOMAINS are captured by the RSS and applied to the E-SRF Masterfile for that Security Image.

IMAGES are CONFIGURED on the Masterfile and represent a specific set of Resident Security System (RSS) databases used to administer security.

An IMAGE is always associated with a specific RSS. You may have multiple IMAGES on the same Masterfile that deal with different implementations of the same or different RSS systems.

The Resource and Source Masterfile segments do not physically separate their data by IMAGE, but IMAGE may be used to refer to data related to a specific IMAGE.

The User Masterfile segment does separate its data based on IMAGE. The IMAGE ID is part of the User Segment object key. *This is because the same USERID may not be the same user across IMAGES.*

If the USERID "JOHN" is in two IMAGES, the ID may be representing John Smith in one IMAGE and John Jones in another. For this reason, the User Segment data must be separated.

The DOMAIN (*explained in the next section*) is used to separate events from one system to another. DOMAINS are assigned to IMAGES. This tells the system which DOMAINS belong to which IMAGES.

For example, CICS transaction TR01 can be executed on the *TEST* and *PROD* MVS systems. The resource segment for TR01 will contain any event recorded for TR01 on *TEST* and *PROD*. On the User Segment, the situation is the same, except if the two systems (DOMAINS) are *assigned* to different IMAGES, they will then reside on the objects belonging to the proper IMAGE.

Even though User information is the only data dependent on the IMAGE, IMAGE ID may be used to reference resource *or* user events. If IMAGE is used with a RESOURCE Segment object, it is determined by the assignment of the DOMAIN on which the resource event occurred. If IMAGE is used with a USER Segment object, it is processed at the object key level.

Domains

As previously discussed, an E-SRF DOMAIN refers to a specific MVS system that is commonly referred to as an LPAR or SYSID.

In the world of MVS, the operating system executes on a mainframe either all by itself or in a portion of the mainframe, referred to as an LPAR (L*ogical* PAR*tition*). Other hardware vendors may have other names for this, but LPAR appears to be the universal term.

When MVS runs in an LPAR, it executes as if it were the only operating system on the physical mainframe computer. The mainframe computer may host one or more LPARs, and some of these may be running operating systems other than MVS.

MVS has a term that represents the individual execution of the MVS operating system. The usual term is the four-character MVS "SYSID." It is also sometimes referred to as the SMF ID, or even a JES ID. No matter what it may be referred to in your installation, it means the same thing to E-SRF.

Do not confuse the MVS term SYSID with the *CICS "SYSID."* They have completely different meanings. The CICS SYSID has no relevance to E-SRF.

Other operating systems may refer to their existence on a specific piece of hardware as a "DOMAIN."

Because of the confusion with this terminology, E-SRF has adopted the term "**DOMAIN**" to describe a specific copy of an operating system (such as MVS) running on a specific computer system complex. In the Data Dictionary, the name "SYSTEM" may be used interchangeably with "DOMAIN." However, DOMAIN is the preferred term.

Event Reporting supports an eight-character DOMAIN ID associated with each event. Events are stored on a Masterfile object in chronological order from oldest to most recent. Events from all DOMAINS are stored together. The DOMAIN ID is treated as another data field captured along with an event. This is how E-SRF determines on which system the event occurred. It also tells the Update Function how to store the event on the Masterfile.

IMAGES consist of one or more specific *DOMAINS* that share a common set of security databases.

Domains are NOT actually defined to the Masterfile. You *assign* DOMAINS to existing IMAGES using the ASSIGN command. The only Masterfile objects directly affected by a Domain assignment are those contained in the USER Segment. The IMAGE is part of that object's key. The key is USERID and IMAGE. These two datanames, however, can be treated separately in reports. Refer to the *Event Reporting Command Reference* and the *User Guide* for more information on how to *assign* DOMAINS to IMAGES.

In Event Reporting, you cannot have two systems with the same DOMAIN name. Event Reporting will attempt to treat both as the same DOMAIN. If this situation exists, you must pre-process your journalized event data by altering the supplied *domain ID* to something unique, and separate out unlike Resident Security System (RSS) data to individual journal datasets. An individual Update Function execution can only process a single RSS.

In summary, security events originate in DOMAINS. DOMAINS live in IMAGES. A DOMAIN is an eight - character Masterfile field that is captured and used by Event Reporting to track on which system a specific event occurred. In the case of MVS, a DOMAIN is the four-character SYSID, left justified, and padded with blanks.

IMAGE Structure Summary

Each Security IMAGE has user information specific to the RSS. Information such as the user's name, telephone number, social security number and RSS dependent statistics is stored for each user on each IMAGE. It is possible to have a single userid "JSMITH" representing JOHN on one system and JILL on another. The Security Image ID separates user information across multiple security databases. Each IMAGE could have its own complete set of User Objects relating to a specific user.

MASTERFILE ORGANIZATION

Security IMAGES are individually defined to E-SRF during CONFIGURATION. Each IMAGE is CONFIGURED and the RSS dependent data SYNCHRONIZED to the Masterfile. When these operations are completed, an E-SRF Masterfile exists that contains objects for each user in each defined security IMAGE. If JSMITH is defined in all of your security databases, there will be a separate set of user objects including a User Header (UA) describing JSMITH in each IMAGE.

The *ASSIGN* command is used to associate DOMAINS to IMAGES.

When the IMAGE is defined in E-SRF, you *ASSIGN* DOMAINS to the IMAGE they belong to. If you are running with a single security database structure, you should ASSIGN all DOMAINS (********) to your IMAGE. It would be better to specify each DOMAIN separately, but for now, a full mask (********) meaning all DOMAINS will suffice.

When the E-SRF Masterfile was built, a job was run that related to the security system you initially wanted to report on. An IMAGE ID was chosen. That IMAGE is what you will report on.

When adding additional IMAGES, remember to review your ASSIGNments and make corrections if necessary. Refer to the *Command Reference* for more information on how this is done.

When you run the UPDATE, the update function is performed at the Resident Security System (RSS) level. This means the update may contain data from one or more DOMAINS and there may be multiple IMAGES involved. The only restriction is that if you run the Update Function for RACF, all journal data must be from a RACF RSS.

The only objects directly affected by DOMAIN assignments are those contained in the USER Segment. A user segment key is: USERID and IMAGE. These two datanames, however, can be treated separately in reports.

Resource Segment objects contain all DOMAIN activity across all IMAGES.

User Segment objects contain only the DOMAINS belonging to the current IMAGE.

Even though user information is the only data dependent upon the IMAGE, IMAGE ID may be used to reference both *user* and *resource* events.

If IMAGE is used with a Resource Segment object, the data for the proper DOMAINS is provided by cross-checking the DOMAIN ID against the DOMAIN assignment table to determine if the selected IMAGE matches.

Segments

The Masterfile is divided into SEGMENTS. A "segment" *is a grouping of information that is related in some way*. For example, the USER segment contains information about the users on your computer system, and the RESOURCE segment contains information about the datasets and resources in your computing environment. A segment may also be viewed as a "file within a file."

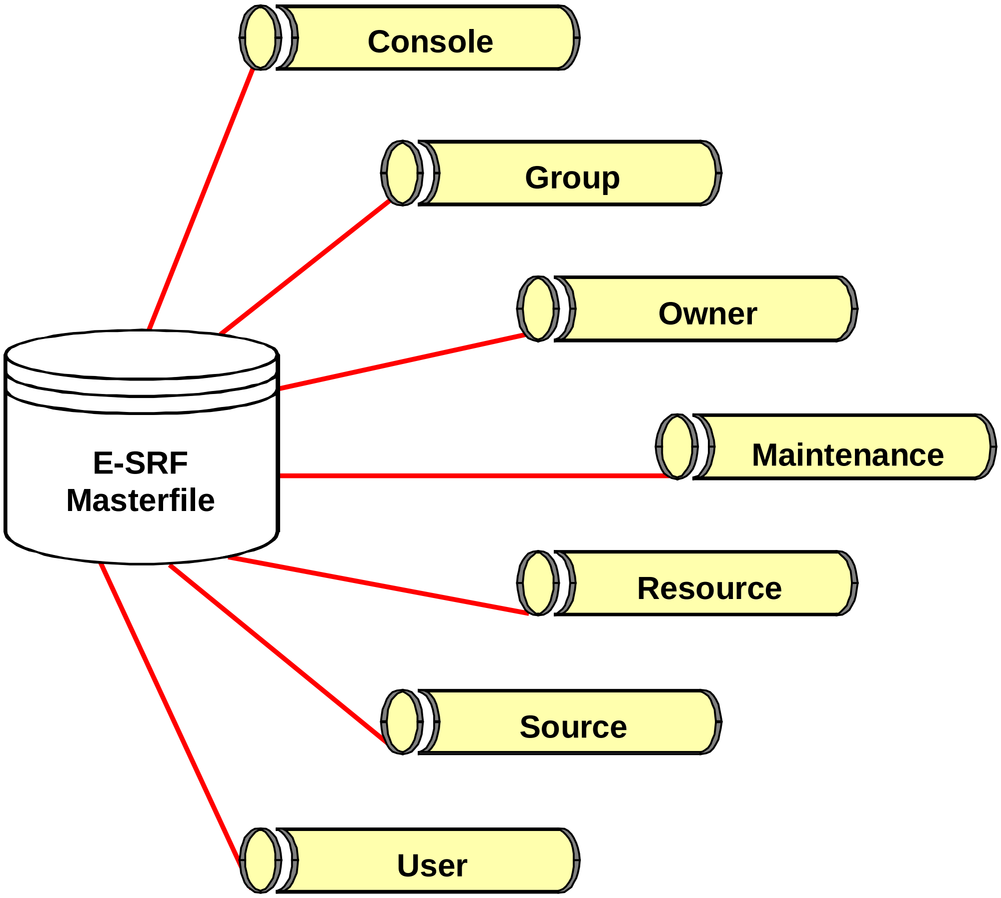

There are seven segments in the current release of E-SRF, not including the "Control Segment" that is used to administer the Event System product itself. This segment contains the customization "parameters" you provide, and warehouses the SYSTEM and IMAGE data dictionaries. In most discussions regarding the segments of the Masterfile, the Control Segment is omitted.

The segments that you use for reporting and control are: Console, Group, Owner, Maintenance, Resource, Source, and User.

Data from one segment may interact with data from another, such as tying information from one segment to another in order to produce the desired report.

Following is a list of each segment, a brief description, and the object types used in that segment.

Masterfile Segments and Object Descriptions

Control Segment

The control segment is used to internally maintain control over the Masterfile and all E-SRF processing options. This segment is not available to the user through the data dictionary. There is no security event data contained here.

OBJECT TYPES: Control

Console Segment

All logged console activity carried out by the Resident Security System is stored in this segment.

OBJECT TYPES: Chronological

Group Segment

This segment contains a single object type that identifies GROUPS, as viewed by E-SRF. To use E-SRF automatic report distribution, the target groups must be defined in this segment.

OBJECT TYPES: Header

Owner Segment

This segment contains a single object type that identifies OWNERS, as viewed by E-SRF. E-SRF automatic report distribution routes the reports to the identified owners and interested parties. The target owners must be defined in this segment.

OBJECT TYPES: Header

Maintenance Segment

All Resident Security System maintenance information is stored in this segment.

OBJECT TYPES: Chronological

Resource Segment

This segment contains objects that describe “things,” i.e. the various resources the Resident Security System is attempting to protect. Datasets are considered a type of resource.

OBJECT TYPES: Chronological
Maintenance
Recap
Statistical Summary

Source Segment

This segment contains objects that describe “places” where access is being attempted.

OBJECT TYPES: Recap
Userid (Invalid Userid signon attempts)

User Segment

This segment contains objects that describe “people” who are making access requests.

OBJECT TYPES: Header
Administration
Chronological
ETF/* Firecall
Maintenance
Profile
Recap
Statistical Summary
Trace Chronological

Objects

Segments are further divided into smaller groupings called OBJECTS. An Object is ***“****a collection of data fields that are associated to a particular item.”* For example, the resource “ACCT.LEDGER.MASTER” may be a dataset that contains general ledger information for your accounting system. In E-SRF, Objects would be created for this resource if any activity were reported against it. This Object would contain all E-SRF data relative to it. Physically, the objects individually could contain one or more VSAM RECORDS on the Masterfile, depending on how much room was needed to store the data. An Object in E-SRF could be as large as *sixteen million* characters.

The *Object Types* represent similar types of data about that segment. The following table is a list of possible Object Types.

Object Type	Description
HEADER	Information about a user, group, or owner.
ADMINISTRATION	Summary of changes made by a user with security privileges.
CHRONOLOGICAL	Detailed list of events that have happened in date and time sequence.
FIRECALL	Events that have occurred as a result of the ETF/* Firecall Facility.
MAINTENANCE	Information about changes made to a user or resource.
PROFILE	RACF specific data. A list of all CONNECT groups and their attributes for the user.
RECAP	Summary of daily access event information summarized for resources, users, or sources.
STATISTICAL	Summary of logged events in a single day by user or by resource.
TRACE	Detailed list of events that were logged as a result of an RSS TRACE facility.

As demonstrated by the illustration below, some segments contain a single object, while others contain multiple objects.

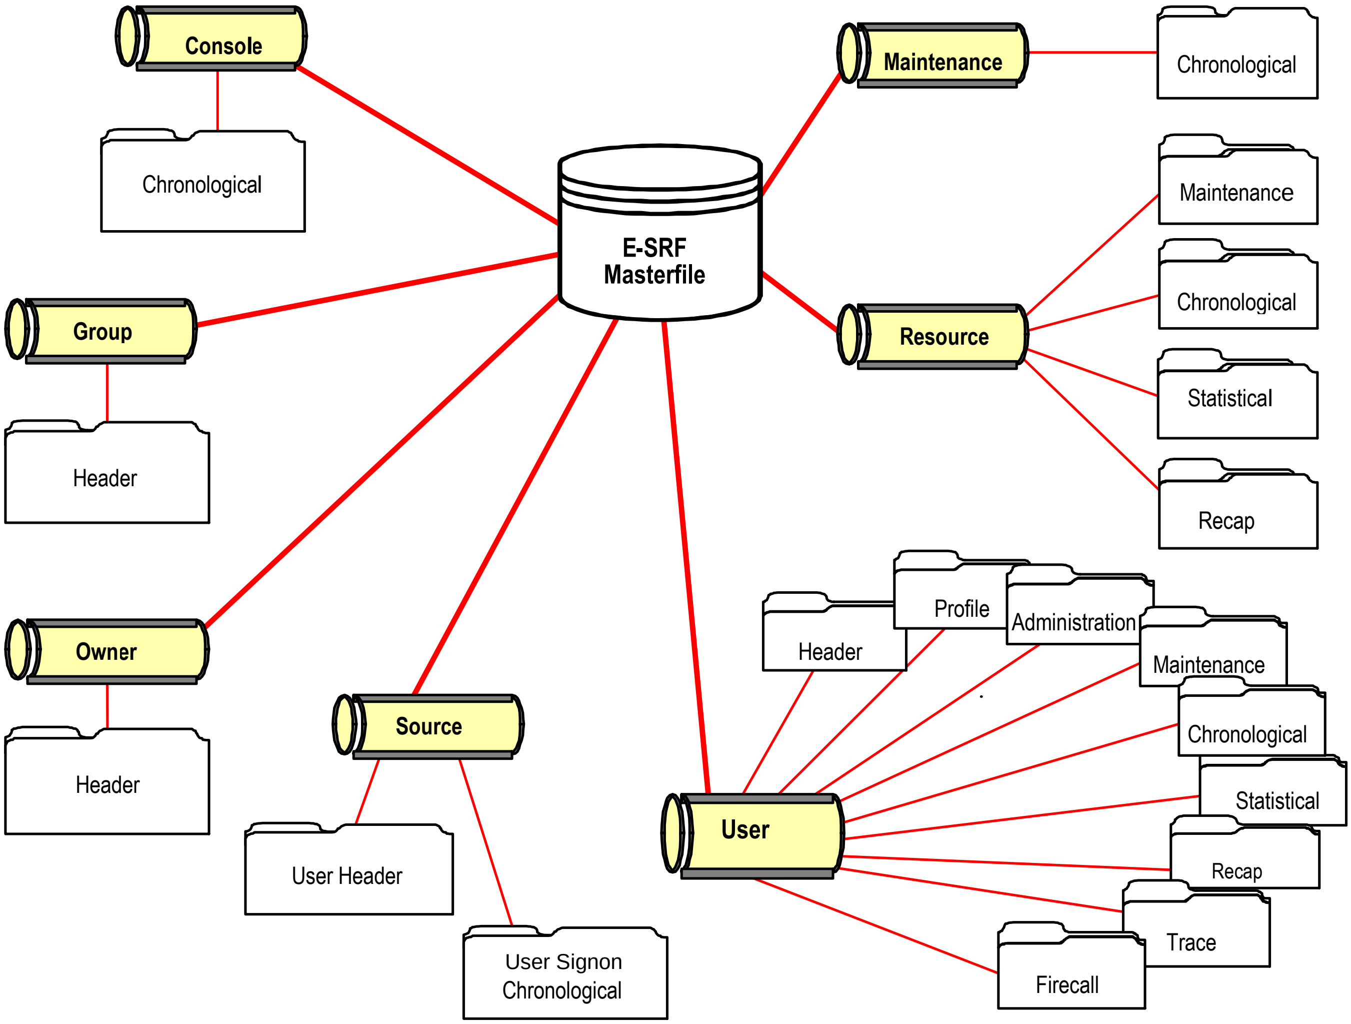

The Masterfile is comprised of many forms of data records stored on a single VSAM cluster.

The Masterfile contains multiple segments (User, Source, Resource, etc.). The Segments contain multiple objects (header, chronological, maintenance, etc.). It is important that this concept is understood at this time.

The following chart provides a brief description of each Masterfile segment and its associated types.

Segment	Type(s)	Description
CONSOLE	**Chronological**	Detailed list of console interactions with RSS.
GROUP	**Header**	Detailed information about a group: name, contacts, interested parties, etc.
OWNER	**Header**	Detailed information about owners for report distribution when using the grouping facility: address, JES class, JES destination, name, phone, routing, etc.
MAINTENANCE	**Chronological**	Chronological list of updates to RSS databases. This is a virtual segment made up of the Resource Maintenance and the User Maintenance objects.
RESOURCE	**Chronological**	Detailed information about access events for resources: type of access requested, action by RSS, date, jobname, program executing, source, resource group, user, user group and other information.
	Maintenance	RSS database maintenance events.
	Recap	Summary of all logging and violations for each resource for each day.
	Statistical	Summary of any logged access attempt events for a resource by userid.
SOURCE	**Recap**	Information about sources where events occurred: date, number of dataset access loggings and violations in a day, times of events, allowed access due to privileges, signon activity, program access and others.
	Userid Invalid Chronological	Detailed information about invalid userid signon attempts at sources
USER	**Header**	Information about users who caused events to occur. RSS userid fields: name, UID, etc.
	Administration	Summary of security administration events the user performed.
	Chronological	Detailed access event activity for users.
	Firecall	Detailed list of events that occurred as a result of the ETF/* Firecall Facility
	Maintenance	RSS database maintenance events.
	Recap	Summary of all events that took place for a particular user.
	Profile	RACF specific data. A list of all CONNECT groups and their attributes.
	Statistical	Summary of loggings and violations: user, class, resource, date and time.
	Trace	Detail list of events that were logged because of an RSS TRACE facility

If you have more than one IMAGE configured for your E-SRF Masterfile, you will have a set of user objects for each IMAGE that has activity reported against it and a user header for every user in every IMAGE.

Object Identifiers

Each Object Type within a segment has an associated *Object Identifier*. Segments that contain multiple Object Types require specific identifiers to differentiate between segment types (Chronological, Recap, etc.). Therefore, an object identifier consists of a SEGMENT and an OBJECT TYPE code for the kind of data stored.

SEGMENT TYPE

+ 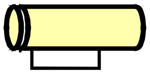 = **Object Identifier**

Resource **R**ecap ➔ **RR**

In most cases, the Object Identifier will be one character for Segment and one character for Object Type. Some object identifiers do not follow this convention. The Segments and Object Types making up the Object Identifiers are shown below:

Object Identifier	Segment	Type
FC	**F** - Console	**C** – Chronological
GA	**G** - Group	**A** - Header
MC	**M** - Maintenance	**C** - Chronological
OA	**O** - Owner	**A** - Header
RC	**R** - Resource	**C** - Chronological
RM		**M -** Maintenance
RR		**R** - Recap
RS		**S** - Statistical
SR	**S** - Source	**R** - Recap
SU		**U** – Userid Inv Signon
UA	**U** - User	**A -** Header
UB		**B -** Administration
UC		**C -** Chronological
UF		**F -** Firecall
UM		**M -** Maintenance
UP		**P -** Profile
UR		**R -** Recap
US		**S -** Statistical
UT		**T** - Trace

For more information about naming conventions of Segments and Object Types, refer to *The Masterfile and Data Dictionary Reference.*

Datanames

Each individual piece of information kept in the Object and Segment groupings is referred to as a *Dataname*. The dataname represents a piece of information kept in the Masterfile. For example, the dataname for the eight-character userid is USERID. The dataname for the twenty-character string that identifies a user is NAME.

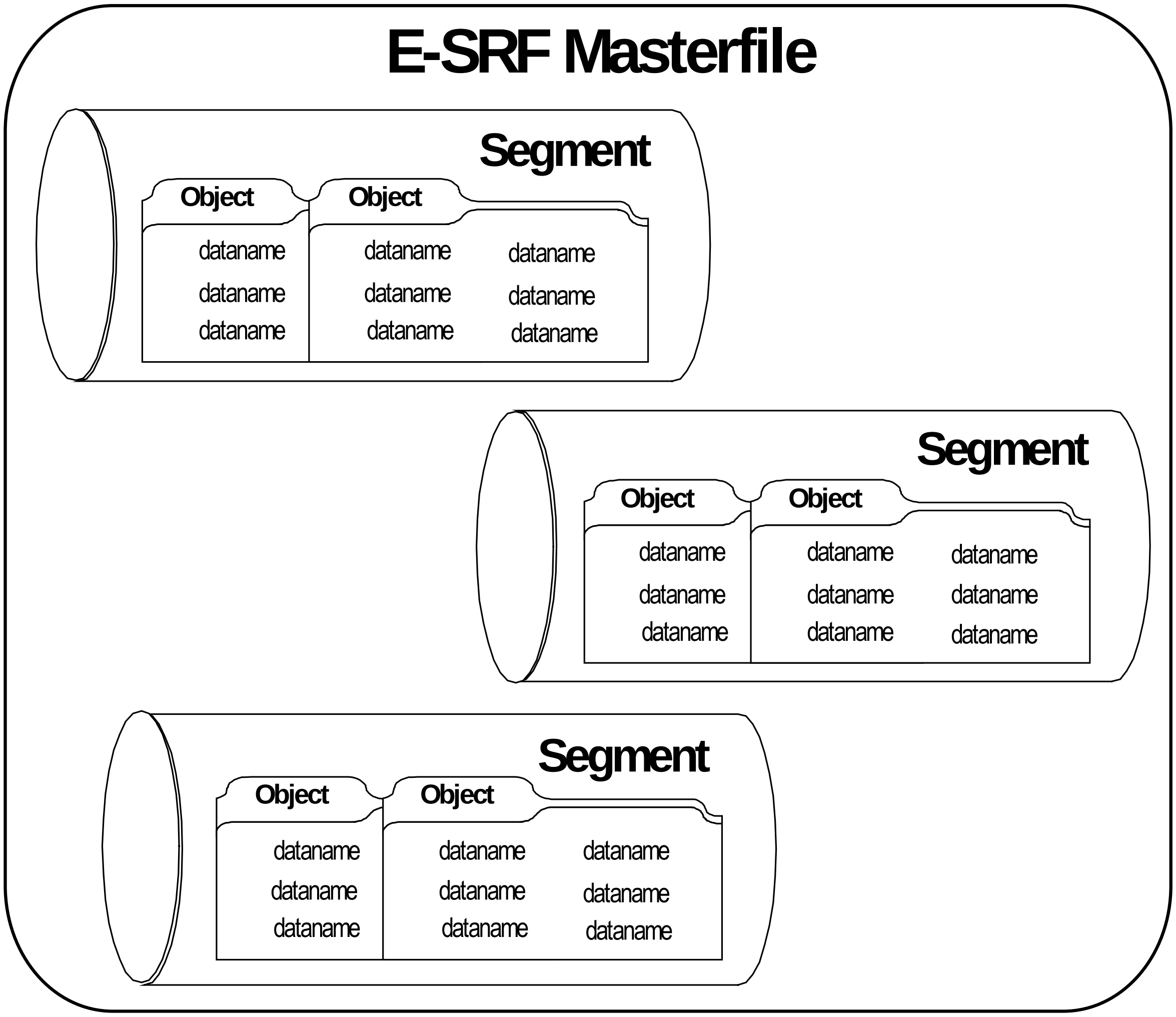

The E-SRF Masterfile is organized so that all information is readily available.

Reports are produced describing events by identifying which Datanames and Objects from which Segments you want to see. This flexibility enables a view of security events from many perspectives.

Dataname Naming Conventions

As previously discussed, the Masterfile is divided up into segments that contain objects that are used to maintain your data.

The Masterfile is a random access VSAM file. It contains a KEY and data associated with the key.

The overall Masterfile key has datanames that are associated with it. These names may be used to reference any object within any segment.

The mapping of the Masterfile key is altered based on the type of segment being processed. The RESOURCE dataname contained on the key may be considered a superset of other datanames when relating to a different segment being processed. In the case of the User Segment, the RESOURCE dataname really contains the USERID and IMAGE of the user.

The naming conventions employed in the Data Dictionary should help you understand the usage of the dataname by the actual name given to the data field.

KEY FIELDS

These fields stand by themselves, meaning the entire dataname is the name of the data field. These names may be referenced any time for any object within any segment. The names of these data fields are: SEGMENT, CLASS, RESOURCE, VOLUME and OBJTYPE.

The following fields are not actually stored on the Masterfile, but are *related* to the key.

COMMENT	provides grouping “comment” text for the object.
GROUP	provides the group ID of the group associated with this object.
OWNER	provides the owner ID of the owner associated with this object.

OBJECT KEY FIELDS

These fields redefine the RESOURCE KEY field mentioned above and also stand by themselves. These datanames may be referenced for any object within a specific segment. The names of these data fields are:

Source Segment:

SOURCE	The name of the source the objects in the segment represent.

User Segment:

LOGONID	The userid of the user the objects in the segment represent.
IMAGE	The IMAGE a particular userid belongs to.
USER	The userid of the user the objects in the segment represent.
USERID	The userid of the user the objects in the segment represent.

OBJECT DATA FIELDS

These fields contain your data. The datanames allow you the ability to access them in reporting. The dataname naming conventions allow you to know at a glance the segment and the object the data represents and the dictionary name of the data itself.

The name is divided up into two character strings separated by a period, similar to the convention used in naming datasets.

The first string consists of two characters. The first character is the SEGMENT, followed by the second character that identifies the OBJECT within the segment.

The second string consists of the dictionary name for a specific piece of data.

Using Datanames to Identify Information

To identify which particular item of data you want to E-SRF, identify the Segment and Object Type, followed by the Dataname. For example, if you want to include the number of loggings for a particular user, use the following:

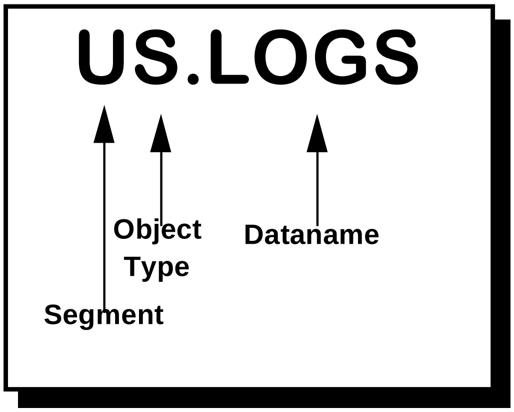

Where “U” = the User Segment, “S” = the Statistical Object Type.
“LOGS” is the Dataname contained in the Object Type.

To include the type of access that was attempted in a security violation for a resource, use the following:

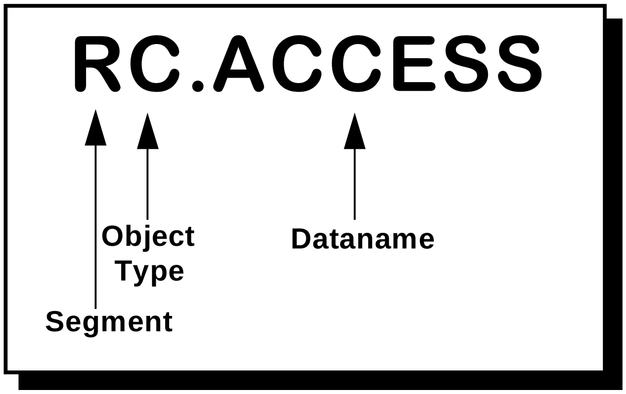

Where “R” = the Resource Segment and “C” = the Chronological Object Type

Each individual piece of information within an Object Type and Segment must be identified uniquely. Datanames are the labels used to identify those individual data items.

For example, the name RC.ACCESS is the type of access requested for an event. This is carried on the Chronological Object of the Resource Segment. In other words the complete name of this data is Resource Segment, Chronological Object, ACCESS.

Since there are too many datanames to list in this chapter, please refer to the *Event System Masterfile and Data Dictionary Reference* for the entire list.

Once you know the datanames within a particular Segment and Object Type, you can create customized reports that include those datanames, such as in the example that follows:

```
RUN  REPORT(ESRFLIST)  PARM(USER)  -
     TITLE(USER ACTIVITY REPORT)   -
     WHEN(USERID EQ MANAG01)       -
     WHEN(UC.DATE EQ *-1)          -
     FIELDS(                       -
               USERID              -
               UC.DATE             -
               UC.RESOURCE         -
               UC.ACCESS           -
               UC.REASON           -
                        )
```

Keys

Each Object Type also needs a *key* to identify how E-SRF relates to the data in each segment. For example, the key in the User segment is the *userid*. The key is referenced only by the dataname; no Object identifier is used. Note in the following example that USERID only is used, not UC.USERID. This only applies to Datanames contained in the Key of the Segment that is being referenced. If the Dataname being referenced were "Reason," UC.REASON would be coded.

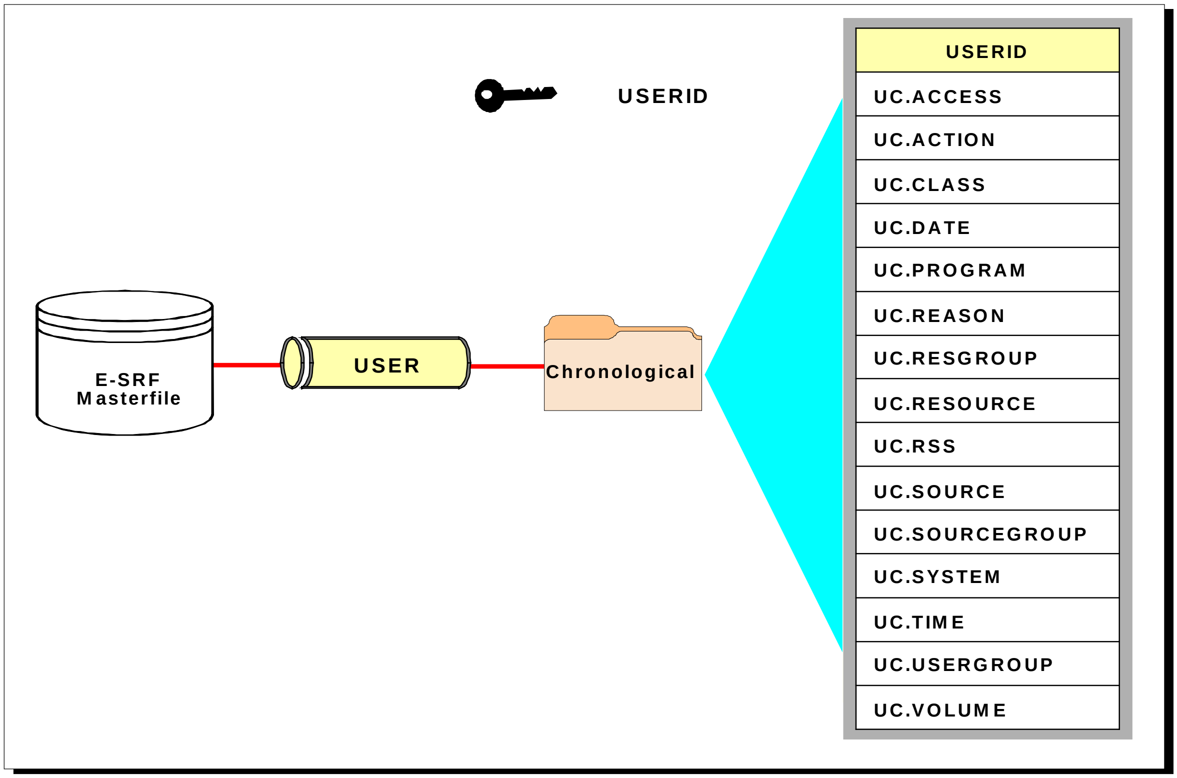

Data Item Types

In addition to knowing the dataname of an object, it is helpful to know how the information is stored on the Masterfile. The actual information can be stored as a single data item or an array data item.

SINGLE Data Items

Data fields stored in a "linear representation" on a particular E-SRF object are referred to as single data items.

An example of this is a data processing record consisting of several fields (*field1 field2 field3*... *field20*). These twenty fields make up a single logical record, each field having its own name and purpose. Below is an example of a data processing record and a Masterfile object for comparison:

Record	key and other fields	*fieldxxxx*1, *fieldxxxx*2, *fieldxxxx*3, *fieldxxxx*4, ….field*xxxx*20
Object	key and other fields	*dataname*1, *dataname*2, *dataname*3, *dataname*4, …*dataname*20

ARRAY Data Items

A ***list*** of several data fields stored in a "linear representation" is called an array. Each occurrence on the list is considered to be an array element. Individual fields within the array element are called array items. The sequence of fields referred to as an *ARRAY ELEMENT* may be repeated over and over. Array data items "live" in array elements; array elements live in objects.

E-SRF relates the individual array items with their respective datanames within the array element structure.

An example of this concept is a single group of several fields (*field1 field2 field3...)* repeated one or more times in a single data processing data record. E-SRF utilizes this concept when maintaining array element objects on its Masterfile, as shown below:

***Object* key and other fields**

Array Element	Array Items
Event 1	*dataname1, dataname2, dataname3, …dataname20*
Event 2	*dataname1, dataname2, dataname3, …dataname20*
Event 3	*dataname1, dataname2, dataname3, …dataname20*
	. . .
Event *n*	*dataname1, dataname2, dataname3, …dataname20*

E-SRF Data Dictionary Structure

System Dictionary

E-SRF contains a *SYSTEM DICTIONARY* that is placed on the Masterfile when the system is initialized. This dictionary contains datanames common to all Security Images on the Masterfile.

Image Dictionaries

When a specific Security Image is configured to E-SRF, a separate *IMAGE DICTIONARY* is placed on the Masterfile. This dictionary contains datanames that relate to the *IMAGE* assigned to it. The datanames normally map the USER HEADER object because this data is completely related to the specific RSS database represented to E-SRF. Remember that a single IMAGE is a set of all DOMAINS (SYSIDS) being serviced by a common set of Resident Security System databases.

More than one IMAGE may be contained on a single Masterfile. Each contains its own Image Dictionary, and these Image Dictionaries may or may not contain the same datanames as they relate to the specific RSS databases. This structure allows you to CONFIGURE multiple Security Images to a single Masterfile, even across Resident Security Systems. A single userid may be represented across multiple Security Images and may or may not be the same individual.

E-SRF’s reference to a particular dataname in an Image Dictionary is transparent to you. You specify the dataname and E-SRF relates it to the proper IMAGE and location of the data.

NOTE: **In this release, all data referenced by the same dataname across IMAGES must be in the same format and location (this will change in a future release).**

This Reference Manual describes the information related to security events, but it does not attempt to describe the information contained in the Resident Security System datanames that are located in the USER segment of the Masterfile. The ESRFDICT Report Overlay prints a listing of all datanames in the current E-SRF Masterfile. This includes datanames from the Resident Security System for all IMAGES. Refer to the *Report Overlays Guide* for instructions to produce that listing.

E-SRF Data Dictionary Summary

When you INITIALIZE an E-SRF Masterfile, E-SRF creates a System Dictionary. When you CONFIGURE one or more Security Images, E-SRF creates an Image Dictionary for each IMAGE that you CONFIGURE. The System Dictionary and the additional Image Dictionaries make up the *E-SRF Data Dictionary*,

Run the ESRFDICT Report Overlay to produce a listing of the System Dictionary and all subsequent Image Dictionaries. For more information about running the ESRFDICT Report Overlay, reference the *Report Overlays Guide.*

Chapter 9: Controlling the Event System

Masking and Grouping Facility Rules

The masking conventions in the E-SRF system had to be equal or greater than all Resident Security Systems (RSS) this product supports.

A careful analysis was performed to determine which masking was the most powerful and easiest to use.

The masking conventions were modeled after ACF2 and further enhanced. Any masking requirements for RACF or other systems can be supported with the approach developed in E-SRF.

The EKC Integrated Grouping Facility underwent the same analysis, with the same results. The "*rules,*" however, are not ACF2 oriented. They are a cross between ACF2 and RACF. The power of the E-SRF implementation is far greater than what is in either ACF2, RACF, or any other security system available.

You will find that very little "rule writing" is required to provide the grouping of your organization's secured resources. This is because the rules are truly *algorithmic,* providing the most efficient approach to dealing with this issue.

To learn more about the EKC Integrated Grouping Facility, please refer to the *EKC Integrated Facilities Resource Grouping Facility Guide*. Additionally, the *E-SRF Event System Command Reference* discusses masking conventions used in various commands and features.

Working With Date and Time

Date Formatting Options

There are two date formats to choose from. The date format option may be set at the SYSTEM level with a SET command and may be overridden at the report level using the RUN command or at the OWNER level, as a specification in the owner header.

The date formatting option has control over how dates are formatted everywhere in the system regardless of which type of formatting is being used to represent the date.

If you want your dates to be formatted in the "international" format (i.e. day, month and year), specify the following command:

```
SET/RUN      DATEFORMAT(INTERNATIONAL)
```

If you want your dates to be formatted in the "United States" format (i.e. month, day and year), specify the following command:

```
SET/RUN      DATEFORMAT(USA)
```

Entry Dates

Dates may be entered into the Event System in either INTERNATIONAL or USA format regardless how the system was configured.

Dates must be entered in the eight-character format, which is using only a two-character year specification. The turnover is as follows: Years 00 to 49 are considered to be in the twenty-first century. Years 50 to 99 are considered to be in the twentieth century.

As far as the Resident Security System Journal entries are concerned, the dates are processed in the format they are presented. No matter how dates are entered, they are converted to a standard E-SRF internal format, which is common throughout the Event System.

Internal Date Formats

E-SRF maintains its dates internally in a format that is not affected by day, week, month, year, decade, century or millennium changes. This was necessary to provide users the ability to float across any date span based on relationships instead of actual dates. This type of processing insulates the product from any irregularities resulting from crossing any of the above mentioned date boundaries.

This means there are no Year 2000 or any other date issues in this product.

Century Formatting

Normally the Event System displays all control information dates in the format, which includes the full century. In all cases, no matter in which format control information dates are presented, they will contain a full century.

On report detail when dates are presented in either INTERNATIONAL or USA format, they will be formatted in the classic eight-character date format. This is due to the following reasons:

- The century will add two characters to the width of a report line. In most cases, this could push the report width over the maximum report data width.
- The data is redundant. It is unnecessary to display the same “19” or “20” for a hundred years. The only time this may be a matter of concern is when data for one century is mixed with the data from the other. This may occur for the month before and the month after the century changes.

If you want full century formatting on detailed report lines, you may specify the **CENTURY** command in the RUN statement. This will cause the date field to be expanded from eight characters to ten characters. There is no system option to force this into effect. It must be done at the report level.

Time Formatting Options

There are two time formats to choose from. The time format option may be set at the SYSTEM level with the SET command and may be overridden at the report level using the RUN command. It may also be SET at the OWNER level, as a specification in the owner header.

If times are to be formatted in the “standard” format (i.e. hours, minutes, and AM/PM designation using a twelve-hour clock), specify the following command:

SET/RUN **TIMEFORMAT(STANDARD)**

In this instance, twelve noon would appear as 12:00 PM.

If times are to be formatted in the “military” format (i.e. hours, minutes and seconds using a twenty-four hour clock), specify the following command:

SET/RUN **TIMEFORMAT(MILITARY)**

In this instance, six thirty in the evening would appear as 18:30:00.

System Control Settings

Set Control Settings

SET commands allow you to tailor the event system to your needs. It is possible to run E-SRF with the Default OPTIONS. However, some defaults may not be appropriate for your requirements.

SET commands establish your specifications to E-SRF and will remain until either re-specified by another SET command or overridden by other commands such as the RUN command.

SET specifications are stored in the CONTROL segment of the Masterfile. This segment is not accessible for reporting.

The current SET specifications may be reviewed by issuing the following command:

```
RUN     REPORT(ESRFSHOW)
```

Most of these options are established automatically for you when the Masterfile is initially built. These default options may not meet your needs. E-SRF provides a means to control these options.

The *E-SRF Command Reference* will explain how to code the various commands. While reading this User Guide, you may want to have a copy of the *Command Reference* readily available so if the need arises, you may reference the actual command for more information.

Control over “permanent” processing options in E-SRF is maintained by the “SET” major command. Once a SET command is issued, the named option is in effect immediately. Additionally, the setting is stored in the Masterfile’s Control Segment for subsequent executions.

These “permanent” specifications for parameters are set to control how E-SRF *normally* would work. In most cases, this is the way the parameters will remain.

There are exceptions. A large set of these parameters may be overridden either when *specific* reports are being run or at the data OWNER level for reports going to a specific owner. Normally, options established by the owner header override all others and are specific to the data owner only. Consider certain modes of operation, such as the format of the date on information reports. An owner in Europe may want dates represented in the “international” format, but owners in the United States would prefer the “USA” date format.

We have not yet discussed OWNERS. This discussion will be covered under a separate topic later in this Guide.

Some of the parameters established by the SET command may be overridden by the RUN command for a specific report execution. In this case, the settings are overridden ONLY FOR THE DURATION OF THE SPECIFIC REPORT EXECUTION. At the conclusion of the RUN command, the parameters are reset back to their established settings.

There are additional options, which normally only apply to a single execution of the product. These options are controlled by the major command: OPTION.

To fully understand any E-SRF command, refer to the *E-SRF Event System Command Guide* for complete information about commands and their usage**.**

Options to Control Grouping

Controlling Grouping for a Particular Execution

A significant strength of the E-SRF Event Reporting system is its ability to group items and distribute reports to data owners automatically. For this to occur, GROUPING MUST BE AVAILABLE and by default, it is.

It is **strongly recommended** that grouping be set to NONE, unless you need grouping for a specific execution. Use the following command to accomplish this:

```
OPTION      GROUPING(NONE).
```

There is no system parameter to control whether or not grouping will be used in E-SRF. It is assumed to be required, and group structures will be built the first time any function is performed that has the potential to need them.

If you need grouping, make sure the OPTION GROUPING(NONE) command is not present in the job you are running. Also, make sure there is a JCL DD statement pointing to your Integrated Grouping RULES dataset.

When the Command Processor begins, the grouping option is set to EXTERNAL. This means all required resources are grouped according to an EXTERNAL Grouping Facility. E-SRF currently uses the EKC Resource Grouping Facility for EXTERNAL grouping. You may override this setting by specifying a GROUPING OPTION other than EXTERNAL, such as discussed above. The following statement will reinstate grouping for the duration of the current job, or until another GROUPING option statement is detected:

```
OPTION      GROUPING(EXTERNAL)
```

Temporarily Turning On or Off SOURCE Grouping

Using the OPTION GROUPING command, you can temporarily turn off the grouping of the Source segment for an individual execution. The use of this option has no affect on the grouping status of the Resource or Source segments.

To turn off SOURCE grouping, specify:

```
OPTION      GROUPING(NOSOURCE)
```

To turn SOURCE grouping back on during the same run, specify:

```
OPTION      GROUPING(SOURCE)
```

As with any OPTION command, this setting is only in effect for the duration of the run, or until altered by another OPTION command.

Temporarily Turning On or Off USER Grouping

Using the OPTION GROUPING command, you can temporarily turn off the grouping of the User segment for an individual execution. The use of this option has no affect on the grouping status of the Resource or User segments.

To turn off USER grouping, specify:

```
OPTION      GROUPING(NOUSER)
```

To turn USER grouping back on during the same run, specify:

```
OPTION      GROUPING(USER)
```

As with any OPTION command, this setting is only in effect for the duration of the run, or until altered by another OPTION command.

Performance Impact of Grouping

The grouping OPTION controls whether or not the E-SRF Grouping Structures are to be built and how it occurs. As mentioned in other topics in this publication, grouping in E-SRF is *dynamic.* This means the grouping of resources is performed on demand when grouping information is required. When an E-SRF job is initiated, the EKC Resource Grouping Facility and its interface are initialized. No other grouping provision is made. During this time, E-SRF does not contain any grouping structures and would not be able to associate a resource to a group name.

The first time the potential exists where a group name must be associated to a resource, a user, or a source, the Group Control Program will build the group structures.

The purpose of this OPTION is to control how the group structures are built.

The EKC Resource Grouping Facility will associate a group name to a specific resource. The Masterfile is read sequentially to build grouping structure tables containing lists of all grouped entities, and calls are made to the EKC Resource Grouping Facility to supply the group name. The group names are based on the RULES supplied to it from the **RULES DD** statement in the JCL.

Please note that if you omit the RULES DD statement, a return code of 24 is presented to you. This is OK if you are not intending to use grouping facilities. If this is the case, it is more efficient to also supply the OPTION GROUPING(NONE) command before grouping is referenced.

If you are an ACF2 user, **DO NOT POINT TO THE ACF2 RULES** data. If you do this, you will end up with an abend from the EKC Resource Grouping Facility. The EKC Resource Grouping "*rules*" are unique to E-SRF and are not related at all to the "*rules*" used in ACF2 or any other security product.

The group structures must be in place when an UPDATE or a REPORT using grouping is requested. If they do not already exist, they will be built on demand at this time (*unless told not to with the OPTION GROUPING(NONE) command set*).

It takes time, main storage and CPU resources to build these tables. Although it does not hurt anything to build them, they may not always be necessary.

It would be wise to specify **OPTION GROUPING(NONE)** if:

- You are not EXCLUDING specific resource activity from being applied to the Masterfile during your UPDATE run.
- The reports you are running do not involve grouping (that is you are not displaying GROUP NAMES on any report).
- You are not executing the report under Report Distribution.
- You are not using anything relating to groups and owners during report selection

If **OPTION GROUPING(NONE)** is specified, E-SRF will NOT build group structures. In the event that a group name is required, its intended function is ignored. The OPTION GROUPING choices are:

NONE indicates do not build structures of any kind.

CLASS means build structures assigning the resource's CLASSNAME as the group name.

EXTERNAL means use the EKC External Grouping Facility to relate group names to resources.

As mentioned before, exclusion of this option defaults the system to EXTERNAL. There is no command to alter this default action.

Controlling Resource Grouping

Resources are <u>always</u> grouped. You <u>cannot</u> turn off RESOURCE GROUPING. The only way not to group resources is to turn grouping off for a particular execution as described above.

Controlling the Way <u>Datasets</u> are Grouped

The EKC Integrated Grouping Facility makes a distinction between DATASET resources and all other resources, just as most security systems do. *The E-SRF Event System does not*. For the most part, it treats them the same. In order for the Event System to properly utilize the Integrated Grouping Facility, the grouping logic examines the CLASSNAME of the resource being processed. If it is DATASET, it uses the DATASET grouping routine. If it is anything else, the *RESOURCE grouping routine* is used.

If you want to use *RESOURCE grouping schemes* for datasets, you can accomplish this by specifying the following command:

```
SET    DATASET(classname)
```

You can supply any desired classname, including DATASET. This command supplies a grouping classname for DATASETS, which is normally set to blanks. The presence of a grouping classname for DATASETS causes the Event System to use *RESOURCE grouping routines* for DATASETS.

To nullify the command, you must specify the following command:

```
SET    DATASET()
```

The absence of a DATASET grouping classname will cause the Event System to use *DATASET grouping routines*.

This specification is different than its counterparts for SOURCE and USERID. It <u>cannot</u> be used to turn grouping off.

If classname is used, whatever you specified as a classname is used by E-SRF when determining dataset group names. You will have to treat datasets as any other resource when providing the necessary grouping rules.

The SET DATASET(classname) command set only affects grouping. On reports, datasets will always appear as DATASET as far as the resource CLASS is concerned, no matter what you specify here.

This specification is normally not made unless you want to administer DATASET resource grouping rules the same way you do all other resources.

Controlling the Way Sources are Grouped

The following SET command establishes the SOURCE Masterfile segment group association provision. It is only relevant if you are updating the SOURCE Segment of the E-SRF Masterfile.

Most users do not retain the SOURCE . If you retain the SOURCE segment, you can decide whether or not you want to group your sources.

> If you do not want to group your sources, set the source grouping classname to NONE. The E-SRF group structure will not contain groupings for the SOURCE Segment of the Masterfile.
>
> If you do not want to group your sources and you are using grouping, the source grouping routine will not build empty grouping structures for sources.

To control source grouping, specify the following command:

SET SOURCE(NONE/*classname*)

If a classname is specified, SOURCE objects in the SOURCE segment of the Masterfile will be included in the group structures. The specified classname and the eight-character SOURCE ID will be presented to the source grouping routine for grouping association.

The recommendation is **NONE**, unless you have a reason to associate a group name to the SOURCE segment of the Masterfile.

Controlling the Way USERS are Grouped

This specification establishes the USER Masterfile segment group association provision.

You must decide if you want to group your users or not. Normally users are not grouped, but many installations have legitimate reasons to group them.

If you do not want users grouped, make sure you set the user grouping classname to NONE. By doing this, the E-SRF group structure will not contain groupings for the USER Segment of the Masterfile. If you do not want users grouped, and you are using grouping, the USER grouping routine will not build grouping structures for users.

To control user grouping, specify the following command:

USERID(NONE/*classname*)

If a classname is specified, USER objects in the USER Segment of the Masterfile will be included in the group structures. The specified classname, as well as the user's E-SRF Universal Identification data will be presented to the user grouping routine.

The recommendation is **NONE**, unless you have a reason to associate a group name to the USERID segment of the Masterfile. This recommendation is based on not needing to group your users in an effort to conserve resources. If you want to group your users, the recommendation is to supply a classname that matches your grouping rules.

The Event System User's Universal Identifier

When establishing Userid grouping, the "resource name" is the user's IMAGE followed by a twenty-four character E-SRF Universal Identification.

The E-SRF Universal Identifier will vary according to the RSS being represented by the Security Image.

- In the case of ACF2, it is the twenty-four ACF2 UID assigned to the user.
- In the case of RACF, it is the OWNER, Default GROUP and the USERID.

Please note that the userid itself is not specifically involved in the grouping process. If it were, you would have to group each user one by one.

You must also consider a single userid that may be in more than one security IMAGE. It may or may not be the same user. They would have to be considered separate userids.

The group association is made using the chosen classname as the grouping resource class, IMAGE as the first qualifier of the resource name, followed by the entire E-SRF Universal Identifier as the second qualifier.

The following example is a key for a user called JOHN, with an ACF2 UID string of ABCDEFGHIJKLMNOPJOHN, which exists in IMAGE MVSACF2:

MVSACF2.ABCDEFGHIJKLMNOPJOHN

Because the two fields are separated by a period, you can mask either one.

In RACF, the E-SRF Event System Universal Identification is a composite of three RACF fields:

Characters 1 to 8:	the RACF OWNER
Characters 9 through 16:	the RACF DEFAULT GROUP
Characters 17 through 24:	the RACF USERID

Consider the following example. You have multiple IMAGES and all userids are the same in each IMAGE (i.e. John is really John in all IMAGES). You are able to group by RACF OWNER, and the default group has no relevance, nor does the userid.

The following rule key may be used:

-.PERSONEL–

Notice the dash in place of the IMAGE. This means any IMAGE matches the rule. The group is the personnel department and its owner is PERSONEL, and it happens to be eight characters long. Because we do not care about the RACF default group or the userid, the dash masks it out.

This technique could also be used in ACF2.

Grouping Granularity

When grouping resources, you can control the granularity of grouping using the SYSIDGROUPING command.

The EKC Integrated Grouping Facility provides a means to specify a SYSID to use in addition to the resource name. The Event System will control how this is used.

A resource normally consists of a CLASS, a RESOURCE NAME, a DOMAIN id (where the event took place), and if it is a DATASET resource, it may have a VOLUME associated with it.

You can group on all of these particulars. However, the more you use, the harder the grouping, the slower the processing, and the more complex the reporting.

A command is provided that allows you to maintain control of how granular you want to group - the less granular, the better. First analyze your needs, then specify as required.

SET SYSIDGROUPING(NONE/IMAGE/DOMAIN)

This specification instructs E-SRF as to the “scope” of group association using the EKC External Grouping Facility’s SYSID keyword.

If you specify NONE, SYSID is not considered during group association.

If IMAGE is specified, the Security Image associated with the request is presented to External Grouping as the SYSID.

If DOMAIN is specified, the individual DOMAIN ID (SYSID) is presented to External Grouping as the SYSID.

You should specify **NONE**, if possible. If you must group across systems, use IMAGE. *DOMAIN should be used only if absolutely necessary, as this is CPU resource intense.*

Note: USERID grouping includes IMAGE as part of the resource name and this specification has no affect on this.

Option to Control the E-SRF Base Dates

Altering Processing Dates

E-SRF maintains two "Base Dates" that are used for all date related references. The "Base Dates" are the SYSTEM and RELATIVE date.

The SYSTEM Date is used for global date references such as posting dates relating to when a report was produced, determining what to retain on the Masterfile, and a base for calculating *relative* date. The SYSTEM date is normally set to the Operating System's (MVS) current date.

The RELATIVE Date is used by E-SRF to determine offsets from the SYSTEM Date when specifying reporting options such as the following:

WHEN(RC.DATE EQ *-1).

In this case "*" = RELATIVE Base Date (or *Current Date*)
"-1" = minus one day.

If the relative base date were today, this selection would target events that occurred yesterday.

The RELATIVE Date is normally set based on the current SYSTEM Date. Consider the RELATIVE Date as the date used for report selection processing and the SYSTEM Date as the date used for everything else.

There are times when one or both of these dates may not be appropriate for what you are doing and may need to be overridden.

You may use the date controlling OPTION commands to alter the date anywhere in your command input data.

Use the following OPTION to modify the SYSTEM Date. Remember, if you modify the SYSTEM Date, it will also override the RELATIVE Date, unless a RELATIVE DATE override option is in effect.

OPTION SYSTEMDATE(RESET/*actual date*)

Use the following OPTION to modify the RELATIVE Date. Remember, if you modify the RELATIVE Date, there is no affect on the SYSTEM Date.

OPTION RELATIVEDATE(SYSTEM/LASTRANS/DAYAFTER/*actual date*)

Normally, you would never alter either of these dates unless you want to re-run something or are experimenting.

Consider running a packet of reports you normally run, but for a point in time thirty days ago. You probably have to restore an old backup of your Masterfile to get the data. You would now have to hard code all selection dates to the exact day you want. Instead, use the exact same command input you normally use to run the reports, *except* include the RELATIVEDATE command before the first RUN command.

An example of this is:

```
OPTION        RELATIVEDATE(11/30/98)
```

The Midnight Hour and Date Control

It is possible that E-SRF could be executing when the clock ticks past midnight. This could be a problem because many functions within E-SRF are date dependent (data retention on UPDATE and report selection criteria). E-SRF has several options available to deal with this situation.

Starting Before and After Midnight

The system date and time information used by E-SRF is normally refreshed each time a new command is issued. This provides exact date and time information relating to a specific processing function.

There are situations where this mode of operation may create problems:

If an execution is scheduled to run ***before*** midnight and the job actually runs after midnight. If this occurs, the current date used by E-SRF will reflect the next day, rendering your selections one day too late.

If an execution is scheduled to run ***after*** midnight and the job actually runs before midnight. If this should occur, the current date used by E-SRF will reflect the previous day, rendering your selections one day too early.

Proper scheduling is imperative for date-sensitive selections that start before or after midnight.

Straddling the Midnight Hour

Consider the effects of running an E-SRF execution that starts before midnight and executes commands after midnight. If any of the commands are date sensitive, there will be inconsistencies that may result in incorrect report production.

There are two ways to resolve this problem:

The first is to ensure that E-SRF executions are properly scheduled.

The second is to require E-SRF to **"*freeze*"** its current system date and time for the duration of the job's execution. The default is to "*upgrade*" this information each time a command is encountered.

Use the following OPTION to cause E-SRF to obtain the current system date and time information, and HOLD it for the duration of the current execution:

```
OPTION    SYSTEMDATE(FREEZE)
```

Use the following OPTION to cause E-SRF to obtain the current date and time information and REFRESH the date and time at every command. This allows E-SRF to revert to its normal default of UPGRADING this information for the duration of the current execution.

```
OPTION    SYSTEMDATE(UPGRADE)
```

Please note that the setting of this option is *only* referenced when E-SRF has to refer to the Operating System to get the date and time information. Obtaining the date and time from the Operating System is only performed if there were no overrides in effect for this information.

Consolidating Masterfile Resource Names

E-SRF provides two options that help address issues relating to consolidating like-named dataset resources into a single-named resource. For example, consider Generation Data Groups (GDGs). There may be hundreds of generations of a single dataset. The access rights could be identical, but the reports are cluttered with these accesses, and all you want to know is if someone gained access to *any* of the GDG generations.

There are two dataset consolidation options available. The ability to consolidate the GDG dataset names mentioned above and a "compress" facility allows you to supply a mask of a specific numeric character string within a dataset qualifier.

Both of these options alter the resource name to a common name that may be more appropriate for reporting. The GDG option is very desirable and highly recommended. The "compress" facility is more specialized and would need some analysis before attempting to use.

To implement the GDG consolidation, you would specify the following command:

```
SET     GDG(TRUNCATE)
```

When a GDG's last qualifier is detected, it is simply truncated. The resulting dataset name becomes the original name minus the G0000V00 qualifier. The default is to RETAIN this information. The sample problem shipped was set up with SET GDG(TRUNCATE). This is the recommended setting.

To implement COMPRESSION, you supply a list of qualifier masks using the "SET COMPRESS" command with zeros where the compression should take place. All non-numeric characters of the specified qualifier mask are matched against the dataset name qualifier. If the qualifier matches, and the numeric characters in the mask are actually numeric (in the qualifier), then the qualifier is replaced by the mask. For example, if you have a mask "CICS0000," and a qualifier "CICS1242" is detected, the "CICS1242" will be replaced with "CICS0000." The qualifier "CICS000A" would not match and would be left unchanged.

If you like the idea of consolidating GDG resources, but still desire them to be shown as a GDG, then consider using the following command:

```
SET COMPRESS(G00V0000)
```

It uses more CPU, but all GDGs will be identified with G00V0000 as their last qualifier.

The COMPRESS consumes CPU time and should be used only if necessary.

Keeping Data On or Off The Masterfile

How Many Elements Can an Object Contain?

The Masterfile attempts to store all events inputted through the Update Function. There must be a sensible limit as to how many events may be stored for any single object. Consider the affect of granting a userid the ability to access all resources whether or not the security definitions allow it. In ACF2, this would be NON-CNCL; in RACF, the OPERATIONS attribute provides this. You could get hundreds of thousands of loggings.

To prevent this, a limit is SET on how many elements an array object can contain. The event will be contained in the Recap object, but may not be in the Chronological object.

To set the limit, use the following command:

```
SET     ELEMENTS(1024/number)
```

This specification controls how many events may be stored on a particular E-SRF Masterfile object. This setting is a capacity-oriented specification. You logically decide how large a particular object's *window* is by establishing data retention using the RETAIN command. This setting tells E-SRF that no matter what the RETAIN setting for a particular object is, more events will not be stored than the value declared in this setting.

When the number of events to be retained is greater than this setting, ten percent of the oldest events will be removed to make room for new elements. The removed elements are termed UNEXPIRED ROLLOFFs and are reported in the UPDATE Function's control report.

CONTROLLING the Event System

Experience has taught us that you will never be able to set up a Masterfile that can hold every event that the RSS generates and still be useful for reporting unless your RSS specifications are *very well tuned*. Consider the logging of a dataset such as 'SYS1.BRODCAST' or the logging of a very popular CICS transaction. Unexpired rolloffs are not always bad, especially in the beginning when you are attempting to "fine tune" your RSS journalizing specifications.

We have found **a specification of 6144 is best suited for most implementations**. Your needs may be different. Start with 6144 and then evaluate your UPDATE control report information.

Keeping Unwanted Events Off The Masterfile

The obvious answer is to fix the definitions in the Resident Security System, but this is not always possible. For instance, RACF and an "all or nothing" approach to profile access logging can cause huge amounts of unwanted journalizing that will clutter up your reports.

E-SRF provides you with the ability, using GROUPING RULES, to EXCLUDE resource data from your Masterfile.

You must first set up an Exclusion grouping scheme for this purpose. *Do not* confuse this with your reporting grouping schemes; they are separate. The Exclusion grouping scheme may even consist of its own RULES dataset containing grouping rules you have written solely for the purpose of excluding data from your Masterfile.

Once you have a RULES dataset established for exclude processing, alter the UPDATE job's JCL to contain a //RULES DD pointing to the exclude rules.

If you choose to use the EXCLUDE feature during the UPDATE Function and then run reports in the same job as the Update Function, grouping will then be performed using the rules you set up for EXCLUDE processing. It is possible to set up a single grouping structure for both, but it is not recommended.

Unless you are able to use the same rules for both functions, it is recommended that separate jobs be set up - one for the Update Function using the EXCLUDE grouping rules and another for the reporting that requires grouping rules, using the grouping rules which you have set up for this purpose.

Run a job that activates EXCLUDE processing in the Update Function. For auditing purposes, this is an option with a default set to NO. If you activate EXCLUDE processing, the Update Function will attempt to exclude based on EXCLUDE rules. If there are no EXCLUDE specifications in your rules, no exclusion will take place. Use the following command:

SET EXCLUDE (NONE/YES/RECTYPE)

This command manages the Update Function's exclude provision. It is possible to exclude Resident Security System loggings from ever being updated to the Masterfile by enabling this option and using the EXCLUDE keyword on the External Grouping Facility's rules being used during the update.

NONE Do not exclude regardless of what is in the rules.

YES Exclude at the resource level based on rules. The event's RECTYPE is ignored. If the rule excludes, then the activity is excluded regardless of the event's RECTYPE.

RECTYPE Use the event's RECTYPE as part of the exclude criteria. Each event includes a RECTYPE determined as part of the Resident Security System's E-SRF Update normalization process. As an example, this can give you the ability to exclude loggings, but not violations or special privilege access grants.

Set this specification to NONE, unless you have a good reason to exclude events from being updated on the Masterfile. If you use the EXCLUDE facility, make sure you have procedures in place to track its usage. Review the update control report to ensure that the volume of excludes makes sense according to your requirements.

Eliminating VOLUME From The Masterfile

Dataset resources are normally presented to the Event System associated with a VOLUME SERIAL identification. This may have been important previous to SMS (System Managed Storage). However, SMS routinely moves datasets from one volume to another, as well as manually moving datasets from one volume to another.

If you retain VOLUME information, you may end up with many Masterfile objects containing data for the same resources. **We recommend NOVOLUME**.

The Event System assumes you **do not** want volume information captured and sets the VOLUME information to blanks.

If volume information is desired, be sure it's really necessary. Ask yourself: "How important is it?" "Is it worth the entire overhead (larger reports and more processing time)?"

If you want volume information, specify the following command:

SET VOLUME

If you want to eliminate volume processing, specify the following command:

SET NOVOLUME

Note: Changing it back to NOVOLUME **does not clean up the Masterfile**. The objects with VOLUME will be retained until they expire off.

Masterfile Data Retention Control

The RETAIN command is used to control how long to keep elements on the Masterfile before expiring them. This command controls the size of the reporting "*window*" available to you. There is a separate command for each Object Type contained on the Masterfile. If you set zero days, the data for a particular object will not be stored on the Masterfile.

You specify the name of the object and the number of days for which you want to keep events. If you specify RETAIN RESET, all retention days are reset to system defaults.

RETAIN OBJECT(*xx*) DAYS(*nn*)

xx identifies the object to be controlled (Segment and Type)
nn indicates the number of days to retain events.

User Header (UA) Object Special Considerations

The UA (User Header) object is **special**. When the Resident Security System reports that a user has been deleted, E-SRF will mark the user as being deleted by storing the delete date in UA.DELETED.

If the UA retain days is set to zero, the UA object is deleted immediately.

If it is not zero, the UA object remains on the Masterfile, marked with the delete date for the number of days, which is equal to the highest retain value set for any USER object. When the highest retain value has been reached, the user header object will be purged from the Masterfile. This ensures that references from other user objects will be satisfied. The role of the RETAIN command for the UA object is to lengthen the retention, perhaps to make it as long as other non-user objects on the Masterfile.

Masterfile Object Retention Recommendations

There are no recommendations. Retention depends on your reporting needs and available disk space. For instance, you may not need the SOURCE Recap object. If this is the case, you can specify:

RETAIN OBJECT(SR) DAYS(0)

This will cause all activity going to the SOURCE Recap object to be ignored. The update will fun faster, the Masterfile will be smaller, and this data will not be missed.

The following may be used as a guide for setting up RETAIN values:

FC	Console Chronological	14 days	
RC	Resource Chronological	4 days	High volume (should cover a weekend)
RM	Resource Maintenance	30 days	Security maintenance (cover one month)
RR	Resource Recap	30 days	Daily recap (cover one month)
RS	Resource Statistical	7 days	**Summarized statistics** (cover one week)
SR	Source Recap	0-3 days	If used (should cover a weekend)
UA	User Header	60 days	Keep for historical purposes
UB	User Administration	14 days	Administrator's history
UC	User Chronological	4 days	**High volume** (should cover a weekend)
UF	User Firecall Log	7 days	Access to FIRECALL (cover one week)
UM	User Maintenance	30 days	Security maintenance (cover one month)
UR	User Recap	30 days	Daily recap (cover one month)
US	User Statistical	7 days	Summarized statistics (cover one week)
UT	User Trace Log	0 days	RSS trace information (rarely used)

CACHE Control Settings

Turning on the CACHE

The CACHE command is used to control the presence of the Caching Facility and its use when accessing the E-SRF Masterfile.

The Cache Facility was incorporated into E-SRF due to the large number of VSAM I/O operations required to UPDATE the Masterfile, as well as the I/O required to produce distributed reports, especially at the GROUP level.

When the CACHE is turned ON, the Masterfile is sequentially read and placed into an internally managed Virtual Storage Caching Facility.

When a Masterfile Object is needed, the requester calls the Masterfile Control Program for the data. The requester does not know (or care) if the CACHE is in use. The object will be returned to the requester regardless of where it came from. The purpose of the CACHE is to reduce CPU overhead and shorten the run times required to execute E-SRF by eliminating most of the disk I/O required to manage the Masterfile.

CACHE ON

Activates Caching Facility. CACHE is built. Requests to read or write to E-SRF objects are processed through the CACHE. CACHE stays active until it is turned off or the current E-SRF session ends.

CACHE OFF

Deactivates Caching Facility. Requests to read and write to E-SRF are processed through the VSAM Masterfile. The Masterfile's VSAM cluster is upgraded. Storage required to contain the CACHE is returned to the Operating System.

CACHE UPGRADE

The Masterfile's VSAM Cluster is upgraded with data contained in the current CACHE. The CACHE is still active and continues to be used to provide Masterfile objects when requested.

It should be noted that when the Masterfile is upgraded, it is normally rebuilt. This removes the need to reorganize your Masterfile. Additionally, if the CACHE is left on when E-SRF terminates successfully, the Masterfile's VSAM Cluster will automatically be upgraded (if necessary).

If E-SRF terminates, the CACHE is NOT AUTOMATICALLY upgraded to the Masterfile's VSAM Cluster. This is good thing. If an Update Function terminates, it normally can be started over without having to restore the Masterfile's VSAM Cluster. This is because all the changed VSAM data is in the CACHE (and not on the cluster).

If the ***CACHE REBUILD*** function was running during the termination, you will have to restore your Masterfile's VSAM Cluster.

The CACHE should **ALWAYS** be ON while updating the Masterfile. Once it is on, it should be left on for the duration of the E-SRF execution. Any reporting that is run in the same job with the Update Function with the CACHE ON will benefit.

Storage Issues with the CACHE

Most Masterfiles contain large amounts of data. This data placed in the CACHE greatly improves processing overhead, but at the expense of Virtual Storage.

Please note, the storage used by the CACHE is managed by E-SRF's Pooled Storage Facility. As of Release 1.6, pooled storage is provided using a data-only address space and is done so by default. In Release 2.1, multiple data only address spaces are used to maintain the cache. This was implemented because there are Masterfiles that became very large; and it was impossible to contain the processing programs, various control areas, grouping structures and the CACHE in a single address space.

To learn more about this topic, please refer to the *E-SRF Event Reporting Masterfile and Data Dictionary Reference Manual* for information on the CACHE and how it operates in E-SRF.

Reporting and Data Retention

There is a close relationship between which reports you use and how long to keep data on the Masterfile. Information may be kept on the Masterfile indefinitely. However, most environments will not allow for an infinitely large Masterfile VSAM cluster. This section will outline a method to determine how long to keep data on the Masterfile based on the type and frequency of the reports you create.

The following steps outline the process to determine which reports you want and what Masterfile data is used to create them.

- Which reports are you using on a regular basis?
  How frequently is each one created?
  Which Masterfile Object Types are used for those reports?

- Are there any ad-hoc reports you run?
  Which Masterfile Object Types are used for those reports?

Categorizing Reports

When selecting Event Reports to be created on a regular basis in scheduled jobs, it is beneficial to categorize the selected reports. For example, select reports that will be reviewed on a daily basis, then those that will be viewed on a weekly basis, monthly, quarterly, and perhaps annually. An example of such a list follows:

Report Description	Report Overlay	Frequency
Ranked Security Violations by User	ESRFRVU	Daily
Ranked User Signon Errors	ERSFRUSE	Daily
Ranked Security Violations by Resource	ESRFRVR	Daily & Annually
Environment Change	ESRFLIST	Weekly

Masterfile RETAIN Option

Before deciding how long to keep data, it is important to know how much and for how long data will be stored on the Masterfile. The intent is to ensure that required event data is available for desired reporting.

To review the current Masterfile options in effect, run the ESRFSHOW report. This report will display the RETAIN DAYS specification for each Object Type.

E-SRF event data is stored on the Masterfile during the UPDATE process. Each data item has a limit of how many events can be stored.

Data will be dropped from the Masterfile due to either excessive activity (too many events, so the oldest ones drop off) or reaching the date when the data “expires” (the number of days specified in the RETAIN setting is exceeded). If data is lost because the number of events is excessive, it is considered an *unexpired rolloff*.

It is important to analyze the reports you select and the Masterfile RETAIN "days" settings. Information should be kept long enough to meet your needs and objectives without an unexpired rolloff.

The chart below shows the Masterfile Object Type used and the RETAIN days setting adjusted to the report frequency.

Report Description	Report Overlay	Frequency	Masterfile Segment Used	Retain Days
Ranked Security Violations by User	ESRFRVU	Daily	RR	4
Ranked User Signon Errors	ESRFRUSE	Daily	UR	4
Ranked Security Violations by Resource	ESRFRVR	Daily and Annually	SR	4
Environment Change	ESRFLIST	Weekly	CONSOLE	9

Using the above chart, RETAIN days settings required for each segment used in the selected reports are easily analyzed. It is recommended to add a day or two to the RETAIN setting for a report. For example, a daily report would at most include Friday, Saturday, and Sunday data, so the setting is 4. A weekly report could potentially include both the weekend before and the weekend after, so the setting is 9. It might also be prudent to schedule weekly jobs on Wednesdays versus Mondays. By deferring until Wednesday, the SMF data will accurately span Monday-to-Monday, allowing time for Monday's data to be available.

A RETAIN DAYS setting of 4 will trigger E-SRF to keep an object in the Masterfile for four days. Once that data has been in the Masterfile for five days, it expires and is removed. This means that the oldest event data for that object type will be four days old.

To determine the RETAIN days setting needed, consider:

- How often is the copied SMF file available?
- What day of the week will reports run?
- Will an entire day's worth of data always reside on one journal (SMF) backup?
- Can some objects be established with a zero RETAIN DAYS?

It is important to note that some Object Types may *not* be required for reporting. Establishing zero RETAIN DAYS for some objects can reduce the overall amount of storage required for the Masterfile.

Chapter 10: Event System Reporting

This chapter describes specific highlights of the Security Event Reporting Facility of E-SRF including:

- Running report overlays designed for specific needs.
- Applying titles and headings to add meaning to reports.
- Output reporting to each specific report overlay.
- Using selection criteria to narrow the scope of the report.
- Automatically distributing to OWNERS and INTERESTED PARTIES.
- Creating custom reports using The Data Dictionary.

Report Overlays

The Event Reporting Facility produces many reports (called report overlays). A report overlay is a program designed to collect the appropriate security event information from the E-SRF Masterfile and report it back to you in a readable, sorted format. Each report overlay has been designed to fill a very specific need.

The design concept of report overlays provides a means to produce almost any type of desired report by providing a guide to the E-SRF Event Reporting System's reporting *engine* during the production of reports. This gives tremendous power to any report design and provides a consistent "look and feel" to all Event System reports.

The EKC E-SRF product development team currently creates report overlays. The evolution is dictated primarily by design and analysis of customer requests and suggestions.

If you have a need for a specific report that is not provided by the Event System, you are encouraged to contact EKC to discuss your requirements. We welcome the suggestions and requirements to assist us in enhancing this product.

Normally, report overlays can be developed quickly and can be added to your E-SRF system as routine maintenance.

Non-Modifiable Reports

Some reports cannot be modified by your installation. This type of report overlay was created to assure managers and auditors that the information they receive has not been tampered with along the way. You do, however, maintain control of what is selected for processing and other non-content report characteristics. Some of these reports even allow you to add additional data to the report.

The formats of these non-modifiable reports should meet the needs of most E-SRF users. The list below shows the variety of format and content currently available as of the writing of this publication

ESRFRDRE	Ranked Daily Resource Events
ESRFRDRV	Ranked Daily Resource Violations
ESRFRDUE	Ranked Daily User Events
ESRFRDUV	Ranked Daily User Violations
ESRFRLR	Ranked Loggings by Class and Resource
ESRFRLU	Ranked Loggings by User
ESRFRVR	Ranked Security Violations by Class and Resource
ESRFRVU	Ranked Security Violations by Userid
ESRFRSSE	Ranked Source Signon Errors
ESRFRUSE	Ranked User Signon Errors
ESRFUVLC	Statistical Count of Violations and Loggings by User Within Resource Class
ESRFUVLR	Violations/Loggings for Each Resource by Userid
ESRFVLCS	Statistical Count of Violations and Loggings by Resource Class

.

The following report is an example of the ESRFVLCS report showing a summary count of how many violations or loggings occurred for each resource class (ACF2 type code). Remember, with Event reports, DATASET is just another resource class, so there is not a separate report for dataset related events. You can select them for inclusion in a separate report if you want only that information.

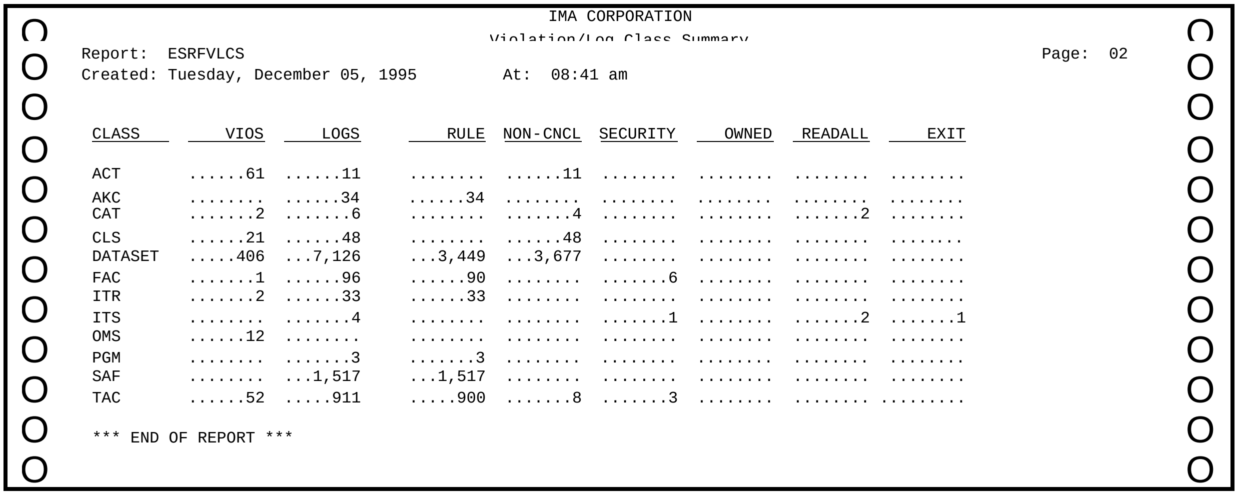

```
                                                IMA CORPORATION
                                          Violation/Log Class Summary
Report:  ESRFVLCS                                                                                     Page:  02
Created: Tuesday, December 05, 1995        At:  08:41 am


 CLASS         VIOS      LOGS         RULE  NON-CNCL  SECURITY     OWNED   READALL      EXIT

 ACT       ......61  ......11     ........  ......11  ........  ........  ........  ........
 AKC       ........  ......34     ......34  ........  ........  ........  ........  ........
 CAT       .......2  .......6     ........  .......4  ........  ........  .......2  ........
 CLS       ......21  ......48     ........  ......48  ........  ........  ........  ........
 DATASET   .....406  ...7,126     ...3,449  ...3,677  ........  ........  ........  ........
 FAC       .......1  ......96     ......90  ........  .......6  ........  ........  ........
 ITR       .......2  ......33     ......33  ........  ........  ........  ........  ........
 ITS       ........  .......4     ........  ........  .......1  ........  .......2  .......1
 OMS       ......12  ........     ........  ........  ........  ........  ........  ........
 PGM       ........  .......3     .......3  ........  ........  ........  ........  ........
 SAF       ........  ...1,517     ...1,517  ........  ........  ........  ........  ........
 TAC       ......52  .....911     .....900  .......8  .......3  ........  ........ .........

 *** END OF REPORT ***
```

E-SRF performs some evaluation of data for you. In addition to the summary information provided by the report, the following sample evaluated the events and ranked them from most to least occurrences for a particular resource. This enables security administrators to concentrate on the most critical issues without having to calculate events themselves.

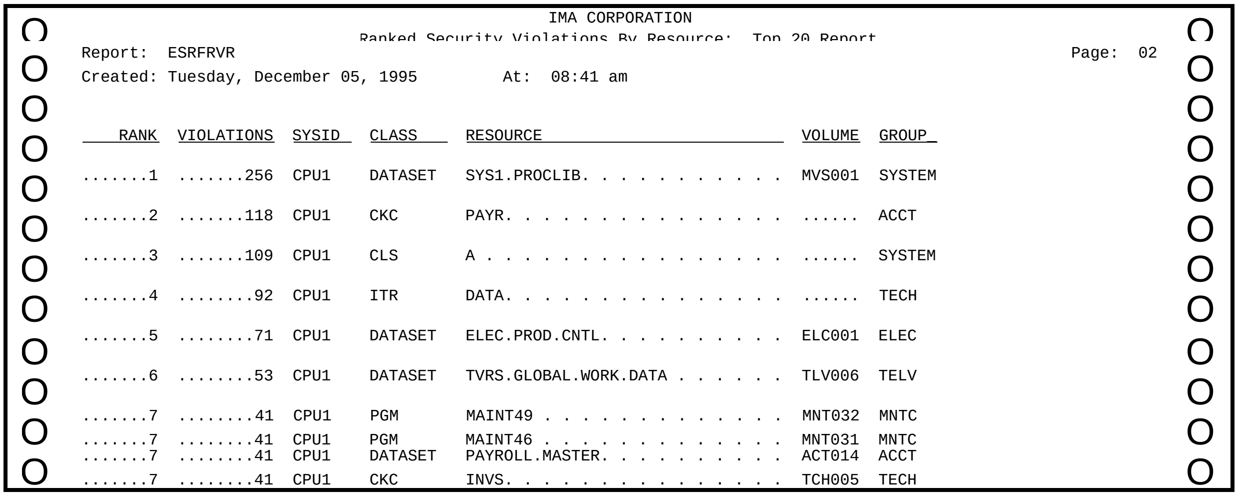

```
                                   IMA CORPORATION
                    Ranked Security Violations By Resource:  Top 20 Report
Report:  ESRFRVR                                                                    Page:  02
Created: Tuesday, December 05, 1995        At:  08:41 am


    RANK  VIOLATIONS  SYSID   CLASS     RESOURCE                          VOLUME  GROUP

.......1  .......256  CPU1    DATASET   SYS1.PROCLIB. . . . . . . . . . . MVS001  SYSTEM

.......2  .......118  CPU1    CKC       PAYR. . . . . . . . . . . . . . . ......  ACCT

.......3  .......109  CPU1    CLS       A . . . . . . . . . . . . . . . . ......  SYSTEM

.......4  ........92  CPU1    ITR       DATA. . . . . . . . . . . . . . . ......  TECH

.......5  ........71  CPU1    DATASET   ELEC.PROD.CNTL. . . . . . . . . . ELC001  ELEC

.......6  ........53  CPU1    DATASET   TVRS.GLOBAL.WORK.DATA . . . . . . TLV006  TELV

.......7  ........41  CPU1    PGM       MAINT49 . . . . . . . . . . . . . MNT032  MNTC
.......7  ........41  CPU1    PGM       MAINT46 . . . . . . . . . . . . . MNT031  MNTC
.......7  ........41  CPU1    DATASET   PAYROLL.MASTER. . . . . . . . . . ACT014  ACCT
.......7  ........41  CPU1    CKC       INVS. . . . . . . . . . . . . . . TCH005  TECH
```

EKC is continually developing additional report overlays and will make them available upon completion. If your company has a need for non-modifiable report overlays not already provided, please contact EKC for information on whether the desired report has already been or is being developed.

Consult the *Event System Report Overlays Guide* for examples, a complete description, and instructions on how to obtain these reports.

Modifiable Reports

These are reports that are completely user modifiable. In fact, a report will not be produced unless you specify what is to appear. The flexibility of these reports should provide you with most of what is necessary for day-to-day security event reporting even if the previously mentioned reports were not included with the product.

The list below shows the modifiable report overlays available as of the writing of this publication.

ESRFDXD	E-SRF Data Download Utility
ESRFLIST	E-SRF Utility List
ESRFUVAR	E-SRF User Variance Reporting Utility

ESRFLIST Report Overlay

ESRFLIST is a utility that enables you to create custom reports. To make this report overlay work, you must supply a list of fields you want included on the report. Any other selection parameters are also considered. ESRFLIST will render a report of the data fields you specified with the selection parameters you requested.

The following sample report output shows violations sorted by resource with very detailed information. The fields were chosen to identify who caused an event to occur, what they were attempting, and what action the RSS took. This is only one example of the limitless number of combinations you can create.

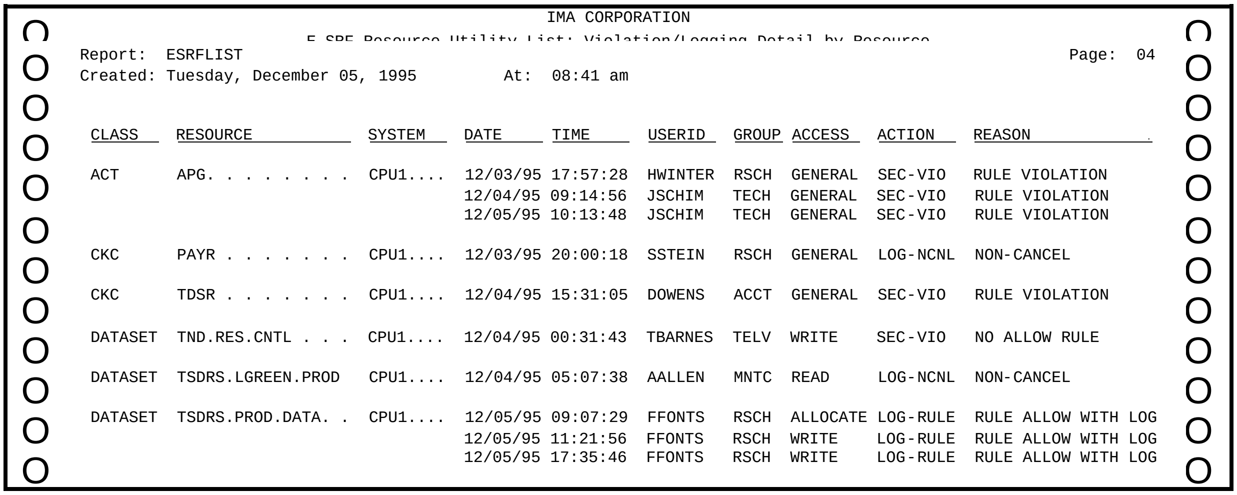

```
                                         IMA CORPORATION
                   E-SRF Resource Utility List: Violation/Logging Detail by Resource
Report:  ESRFLIST                                                                 Page:  04
Created: Tuesday, December 05, 1995        At:  08:41 am


 CLASS    RESOURCE            SYSTEM    DATE      TIME      USERID   GROUP ACCESS   ACTION    REASON

 ACT      APG. . . . . . . .  CPU1....  12/03/95 17:57:28  HWINTER  RSCH  GENERAL  SEC-VIO   RULE VIOLATION
                                        12/04/95 09:14:56  JSCHIM   TECH  GENERAL  SEC-VIO   RULE VIOLATION
                                        12/05/95 10:13:48  JSCHIM   TECH  GENERAL  SEC-VIO   RULE VIOLATION

 CKC      PAYR . . . . . . .  CPU1....  12/03/95 20:00:18  SSTEIN   RSCH  GENERAL  LOG-NCNL  NON-CANCEL

 CKC      TDSR . . . . . . .  CPU1....  12/04/95 15:31:05  DOWENS   ACCT  GENERAL  SEC-VIO   RULE VIOLATION

 DATASET  TND.RES.CNTL . . .  CPU1....  12/04/95 00:31:43  TBARNES  TELV  WRITE    SEC-VIO   NO ALLOW RULE

 DATASET  TSDRS.LGREEN.PROD   CPU1....  12/04/95 05:07:38  AALLEN   MNTC  READ     LOG-NCNL  NON-CANCEL

 DATASET  TSDRS.PROD.DATA. .  CPU1....  12/05/95 09:07:29  FFONTS   RSCH  ALLOCATE LOG-RULE  RULE ALLOW WITH LOG
                                        12/05/95 11:21:56  FFONTS   RSCH  WRITE    LOG-RULE  RULE ALLOW WITH LOG
                                        12/05/95 17:35:46  FFONTS   RSCH  WRITE    LOG-RULE  RULE ALLOW WITH LOG
```

Sorting ESRF Reports

A SORT command has been provided that will allow you to alter the sequence of the ESRFLIST output to any desired sequence. This command may also be used to cause page breaks in the report output.

To learn how to control the use of the ESRFLIST report overlay, please reference the *Event System Report Overlays Guide*.

ESRFDXD Report Overlay

ESRFDXD is a another utility identical to ESRFLIST, except instead of producing a report, it creates a Variable Length file that is suitable for downloading into a PC application such as Microsoft's EXCEL. The following sample report was created to indicate how many changes each security administrator made during the day. This was created with Microsoft Excel, version 4.0.

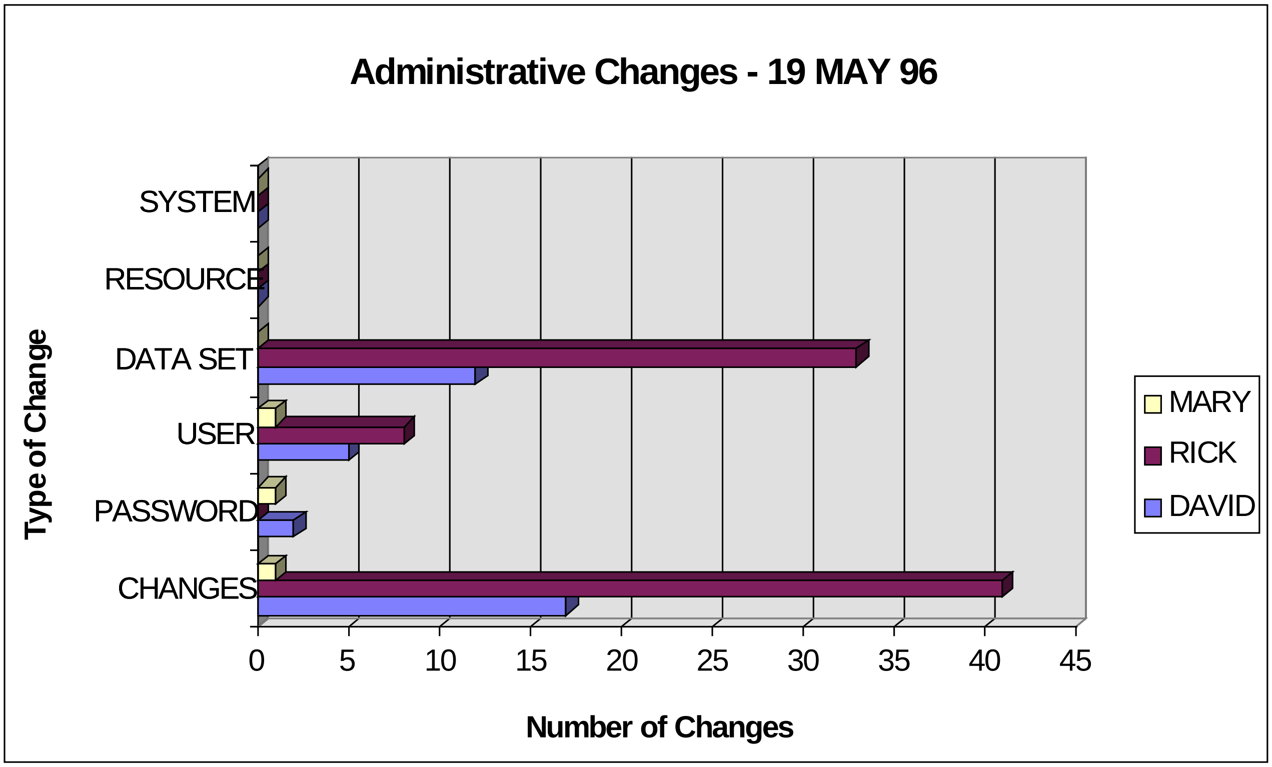

The next example was created with Microsoft Word 7.0 to generate a memo for the manager of a user who caused a violation. This memo can be used as a proactive approach to security administration by asking the manager if the attempted access is required for that user. Once created, you can use your current Email system to distribute these memos as attachments.

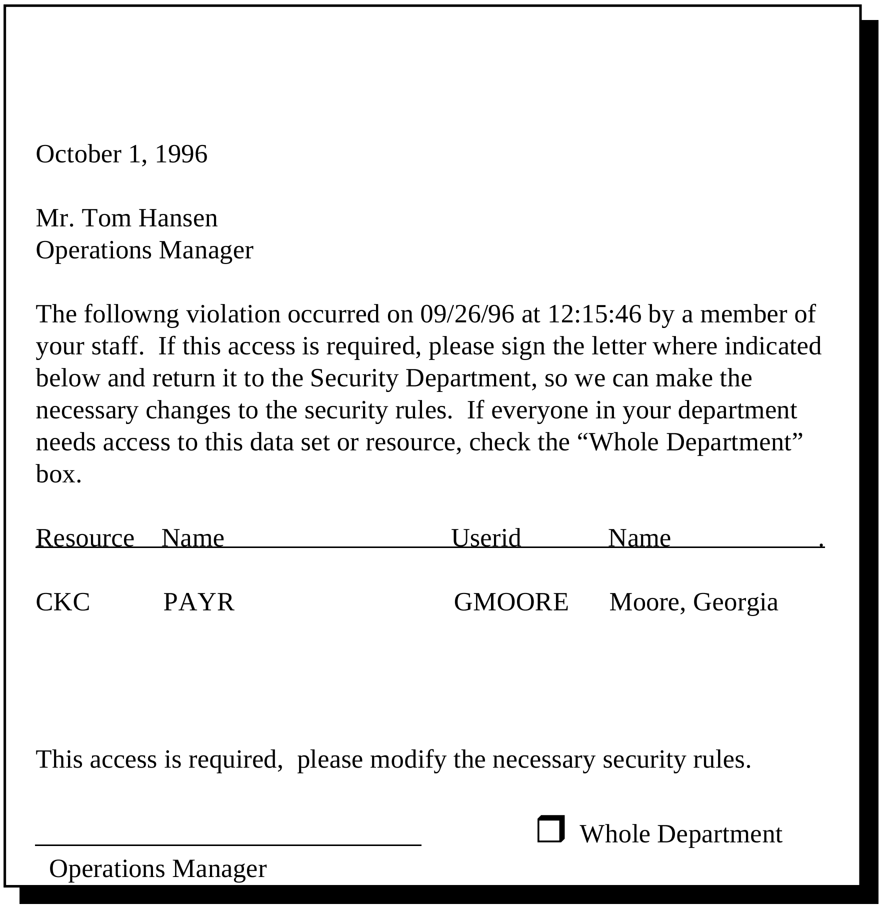

October 1, 1996

Mr. Tom Hansen
Operations Manager

The followng violation occurred on 09/26/96 at 12:15:46 by a member of your staff. If this access is required, please sign the letter where indicated below and return it to the Security Department, so we can make the necessary changes to the security rules. If everyone in your department needs access to this data set or resource, check the "Whole Department" box.

Resource	Name	Userid	Name .
CKC	PAYR	GMOORE	Moore, Georgia

This access is required, please modify the necessary security rules.

______________________________ □ Whole Department
Operations Manager

ESRFUVAR Report Overlay

The ESRFUVAR report overlay will provide a "*delta*" report of userid definition differences for userids across IMAGES.

As in ESRFLIST, a list of User Header (UA) fields is supplied. Instead of just printing these listed fields, all users with the same userid across all IMAGES will have these fields compared. If they are all the same, nothing is reported. If there is any variance in the compared fields for each userid, all listed fields are printed on the report.

The end result is a report that looks exactly like ESRFLIST, but shows all appearances of any userid that had a difference detected in the listed fields.

The following example will produce a list of all users across all images that have variances in the below referenced fields. Please note that there is no field specification for UA.USERID. The userid is automatically established by the report overlay and it will appear as the first field on the report. The references to ACF2.*datanames* are obtained from the Image Dictionary built during Configuration. The RSS happened to be ACF2.

```
RUN  REPORT(ESRFUVAR)                   -
     TITLE(TEST USER VARIANCE REPORT)   -
     FIELDS(                            -
               ACF2.NAME                -
               ACF2.UID                 -
               ACF2.NON-CNCL            -
               ACF2.SECURITY            -
               ACF2.MUSASS              -
                              )
```

Note: Like ESRFLIST, data selection criteria may be used to "scope" down this report. The report overlay will look at the data and make comparisons AFTER the selection process is complete.

Report Control

Report output varies depending upon what information has been requested. The *Report Overlays Guide* provides samples for each of the current E-SRF reports showing what information is presented.

All information in E-SRF reports is pre-sorted into an order that makes sense for the specific report being created. If information is being presented about resources, the report is sorted by resource name. If the report shows signon errors by user, it is sorted by userid.

No meaningless abbreviations, technical jargon or codes are used in E-SRF reports. Every attempt is made to use "English" words and phrases to describe the security events. Security terminology is used only where it is needed to explain the reason for the event logging. E-SRF reports are designed to be understood by non-technical, non-security personnel.

Several E-SRF reports summarize and evaluate event information before presenting it in a report. For example, the "ranked" reports evaluate the events and rank them from the dataset, resource, or user with the most to the least violations or loggings. The most critical issues are always at the beginning of the report so that security administrators can focus on those issues immediately. By using this information to tune the security system, administrators can reduce future security event loggings.

Summary information takes up very little space in the Masterfile and on reports. You can create summary reports over longer periods of time, from monthly to yearly. This type of report enables trend analysis of security activity in your system.

There are occasions when there may not be any events that match your selection criteria. An empty report will be produced to assure you that the report did in fact process correctly, but there were no events to report. The control page will show what was requested and the end of report summary page will show if anything was selected. Auditors and managers can be confident that the report completed and nothing is missing.

Report Overlay General Format

Reports are produced by the E-SRF Event Reporting "engine" (using report overlays to direct and customize the processing). This provides a common look and feel to all reports.

Each report has a "*report wrapper.*" This wrapper consists of a title page, showing the name of the report overlay in block letters, followed by any user specified parameters and control options that were used to produce the report. An end-of-report wrapper page is also provided, indicating the completion of the report and some statistics on how much data was processed and how many lines were produced. "*Flower boxes*" may also appear on the report wrapper.

Sandwiched between the report wrappers is the actual report, or the report "*body.*" This is the data that was requested. Each page has a set of heading lines, consisting of a major heading (common to all report overlays) and minor detail headings (common to the report being produced).

The report detail is then presented. You have complete control over what data is selected for reporting. Page Breaks are at the discretion of the report being produced.

Additional information is available depending upon the report overlay.

A summary listing titled "End of Activity Summary Totals" is presented at the end of every report. Summary data includes information about any processing errors that may have occurred, recap counts, totals, and error codes detected. The final message on the summary page indicates the completion of the report.

You may specify report footers at both the SYSTEM level (footers on every report) or at the individual report request.

Report Overlays make use of various optional RUN minor commands. Consult the *Event System Report Overlays Guide* to determine which RUN commands are available for a specific report. Also consult the *E-SRF Event System Command Guide* to learn more about the RUN command and its various minor commands.

Report "Look and Feel" Control

When the Event System was developed, an attempt was made to "*normalize*" as much security terminology as possible. ACF2 was the first Resident Security System implemented, with RACF coming into existence afterwards. During the product's Beta and Early ship testing of the RACF component, it was found that despite the effort, some RACF users indicated that the reports looked too much like ACF2.

To address this concern, a new system specification was added. The Look And Feel (LAF) specification was added at the system, report, and owner level.

If the LAF(*x*) specification is omitted, the E-SRF reports will look the same way they have always looked.

- If you want your reports to have a look and feel of ACF2, specify: LAF(A).
- If you want your reports to have a look and feel of RACF, specify: LAF(R).

This specification will be enhanced over the next few releases of the product and will be shipped as regular product maintenance. Changes and enhancements will be guided by user requests. If there are no more requests, the assumption will be made that the look and feel is appropriate.

Consult the *Event System Command Reference* for additional information on this feature.

Maximum Lines On a Page

You may specify how many lines will appear on any report output. This includes the control report produced on SYSPRINT, as well as any report directed to a report output DD.

The lines specification may be established as a system wide option provided in a SET command, or you may override the system setting when producing a report by providing it in a RUN command.

To set the maximum lines per page, supply the following command:

SET/RUN LINES(60/lines)

E-SRF manages the print control function by using standard CCW machine control characters (same type as RECFM=FBM) and manages page overflow by line counting. It does not matter whether the report output DD is pointing to a real printer, a JES SYSOUT, or a sequential dataset. Reports are reports. All DDs that are considered reports follow this rule.

There is **no** recommendation for this setting. You have to find out how your printing environment is set up. If you set lines too low, you will waste paper. If you set lines too high, there may be a conflict between what the printer detects as end of page and what E-SRF through line counting detects as end of page. Therefore, a little investigation will eliminate the possibility of wasted paper.

Providing Report Flower Boxes

It is possible to provide "flower boxes" on the front report wrapper, the back report wrapper, or both. You may have as many lines in your flower box as desired. The lines will be framed in a box made of asterisks. These flower boxes will appear despite any WRAPPER options, including the one that turns the wrapper off. The flower box will be constructed by lines supplied by you, in the order supplied.

To create a flower box on the front wrapper, supply the following command:

RUN FRONT(zero to sixty four characters of text)

To create a flower box on the back wrapper, supply the following command:

RUN BACK(zero to sixty four characters of text)

If you want your text centered in the flower box, place a % (percent sign) in the first position of your text line.

Providing Report Major Title

The first line of every report, as well as the control report, contains the system TITLE information.

This setting allows you to specify from one to sixty-four characters of title information. The title text will be centered to the middle of the report. The name of your organization is usually specified.

SET TITLE(*title information*)

Do not confuse this with the "title" specification in the RUN command. This has a completely different meaning.

Augmenting a Report Title

The system wide title is printed at the top of the page on reports. The second line contains the report overlay title information. This title is sixty-four characters long and is centered under the main title.

You have the ability to augment this title information by specifying an additional title line, up to sixty-four characters long. This provision is useful in identifying exactly why the report is being produced. If this is specified, it will be formatted as part of the report overlay title explained below.

The report overlay title will be first separated by a dash, followed by the text you have supplied. The complete line may be up to 131 characters long. This line is then centered under the main title. The report overlay title will be right justified to the dash and your added title will be left justified to the dash.

Use the following specification in the RUN command to augment the report overlay title line:

RUN TITLE(*title information*)

Do not confuse this with the “title” specification in the SET command. This has a completely different meaning.

Controlling Report Widths

In the Event system, the maximum width of any report is 255 characters. However, most printers do not provide this width.

The actual report width may be controlled as a system option using the SET command or overridden by a specification in the RUN command.

WIDTH(133/*report width value*)

Normally, reports are produced with the intent to print on a printer. The standard printer width is 133 characters; the first character is the print control character, followed by 132 characters of text to be printed.

Print control is ultimately up to the print file receiving the data. If it were opened with a record length of 150, then that is how wide all reports going to that file will be. Reports that do not conform to the “OPEN” width are processed with warning messages.

If the actual report being printed is wider, it will be right truncated.

If the actual report is narrower, it will be left justified, filled with blanks.

This value is only looked at when a report output dataset is being opened. The open length overrides everything. **The rules of MVS record length DCB merge DO NOT APPLY here**.

The established FILE record length is looked at first. If set to zero, then the WIDTH value is supplied to determine the length at file open time.

Once the report file is opened, the open length is retained for all reports written to that file regardless of whether or not additional WIDTH values are specified.

This value may also be specified in a RUN command. If it is, and the particular RUN command is causing the report output file to be opened, the width value of the RUN specification will be used.

The recommendation for this value is to set it to 133. If you desire reports to be greater than 133 characters, override the value at the RUN level. Once the output file is opened, all other RUN requests use the original open record length.

LINEMODE Controls

In the Event system, the maximum width of any report is under the control of the WIDTH specification. Such matters as the physical characteristics of the printer or the final output device receiving the report may influence this.

To overcome this restriction, LINEMODE was provided. If the specification is omitted, STANDARD is assumed. If you specify LINEMODE with a valid linemode specification, the report output format will be altered in an attempt to provide data widths greater than what is normally possible. To learn more about LINEMODE and its operational characteristics, please refer to the *Event Reporting Command Reference* and the *Event Reporting Report Overlays Guide*.

Printer Control Characters

In the Event system, print position “zero” contains a standard “Machine” print control character. This character instructs the printer as to the function the printer is to perform on the data being presented. E-SRF provides a means to alter this processing, such as to remove the print control character. Please reference the E-SRF *Event Reporting Command Reference* publication for more information.

Generating Event Reports

E-SRF Event reports are generated in a batch environment using JCL to submit jobs. The JCL identifies all required datasets, including the E-SRF Masterfile, the execution library, and the Grouping Rules dataset if you are using the Resource Grouping Facility.

The E-SRF Command Processor

When creating an E-SRF Event report, the program that is executed will always be ESRFCMD. The Command Processor will use the commands described below to format and create the Event report you specify. Many subcommands are available to generate event reports and maintain the E-SRF Masterfile database. All commands are executed through the ESRFCMD command processor.

NOTE: Do not attempt to run report overlays as stand alone programs.

RUN Commands

The command that creates an E-SRF Event report is RUN. The RUN command is issued with the name of the report overlay you want to execute. Some report overlays have additional subcommands and parameters to specify; some can simply be executed without any parameters. For example, to generate a report showing all items in the E-SRF Data Dictionary, use:

```
     RUN  REPORT(ESRFDICT)
```

No parameters are necessary because this report does not need any formatting flexibility. In addition, most of the non-modifiable report formats can also be run without any parameters or subcommands. However, the ranked reports, for example, do have a parameter for how many ranking levels you want to see, and all Event reports can produce a user-defined title after the report specific one is generated. The Event reports can also use IF and WHEN criteria to select a certain range of information to include.

```
     RUN  REPORT(ESRFUVLC)         -
          WHEN(UL.DATE EQ *-1)
```

“Today” is represented by “*”. Therefore, “*-1” means create a report including only information from yesterday. This variable allows you to create a job, schedule it to run daily, and only include information from the previous day. The “-“ at the end of the line is a continuation character, identifying to E-SRF that there is more information about this report run.

```
     RUN  REPORT(ESRFRVU)          -
          TITLE(TOP 20 REPORT)     -
          LIMIT(20)                -
          WHEN(UR.DATE EQ *-1)
```

This report has several parameters, including: LIMIT (for how many ranking levels are requested), an additional user-defined title, and the date specification for yesterday.

Some parameters exist in both the SET and RUN commands. If identical, the RUN parameter supersedes the Set parameters. There is one exception: the “TITLE” parameter has a completely different meaning in the RUN command. For more information, consult the *Event System Command Reference*.

```
     RUN  REPORT(ESRFLIST) PARM(USER)        -
          TITLE(VIOLATION DETAIL BY USERID)  -
          FIELDS(                            -
                    USERID                   -
                    UA.NAME                  -
                    UA.UID                   -
                    UC.DATE                  -
                    UC.CLASS                 -
                    UC.RESOURCE              -
                    UC.VIOS                  -
                                   )         -
          WHEN(UC.DATE = *-1)
```

This RUN command is executing the ESRFLIST utility to create a custom report. PARM identifies from which segment of the Masterfile to get the information. PARM(USER) indicates that the information will be from the USER segment, and the report will be sorted by userid. The title is specified, followed by each field of information in the order the columns should appear in the report.

```
     RUN  REPORT(ESRFLIST) PARM(RESOURCE)    -
          TITLE(VIOLATION DETAIL BY RESOURCE)     -
          FIELDS(                                 -
                    CLASS                         -
                    RESOURCE                      -
                    RC.DATE                       -
                    RC.TIME                       -
                    RC.USERID                     -
                    UA.NAME                       -
                    RC.ACCESS                     -
                    RC.ACTION                     -
                    RC.REASON                     -
                                   )              -
          WHEN(RC.DATE = *-1)
```

Notice in this example that the segment to retrieve information from is RESOURCE. However, when printing the columns on the report, we want to include the user’s name in addition to the userid. The user’s name is stored in the USER segment, not the RESOURCE segment, so that field is called UA.NAME. You can include information from other segments as long as it is from the HEADER (A) object type. In other words, information about users, groups, and owners (UA, GA, OA) can be included in any report, regardless of from what segment the report is getting the event information.

Selecting Dictionary Fields for Display on Reports

Some report overlays allow you the ability to specify the name(s) of items found on the Data Dictionary that are to be included on the report.

The desired fields may be specified using the listed FIELDS RUN command specification. Some reports do not make use of the FIELDS specification, while others do. The *Event System Report Overlays Guide* indicates which specifications may apply.

To find out more about the FIELDS specification, refer to the *Event System Command Reference*. To learn more about the data items in the dictionary, refer to the *Event System Masterfile and Data Dictionary Reference*.

Using the Grouping Rule Comment to Augment Reports

The EKC Integrated Grouping Facility provides you with the ability to specify comment text. This information is retained as part of the grouping definition and is made available to the Event System for both report data selection and report display.

A useful example of what may be done with the comment would be to associate a plain “English” description of your individual resources during the grouping process. This information could then be displayed on reports, making the reports user friendly.

To learn more about how to specify comments to grouping rules, please refer to the *EKC Integrated Facilities Resource Grouping Facility Guide*. The actual reference to the comment is made using the dictionary name “COMMENT”. For more information on the Data Dictionary, please refer to the *Event System Masterfile and Data Dictionary Reference*

User Information Reports

ESRFLIST will create custom reports including any information in the E-SRF Masterfile. Therefore, you can create reports to show user-related data, without event-related data. For example, you can create a report to show all ACF2 userids with special high-access privileges.

```
RUN  REPORT(ESRFLIST)     PARM(USER)         -
     TITLE(USERIDS WITH SPECIAL PRIVILEGES)  -

     FIELDS(                                      -
               USERID                             -
               UA.NAME                            -
               ACF2.SECURITY                      -
               ACF2.AUDIT                         -
               ACF2.LEADER                        -
               ACF2.NON-CNCL                      -
               ACF2.READALL                       -
                                        )         -

     IF(  ACF2.SECURITY                           -
          ACF2.AUDIT                              -
          ACF2.LEADER                             -
          ACF2.NON-CNCL                           -
          ACF2.READALL   )     ANY
```

This RUN command will generate a report showing a column for each privilege with a YES indicator for the privileges that a userid has. If the userid does not have any of the listed privileges, it will not be included in the report. The ANY keyword means that a userid will be included if it has any of the previously listed fields in the IF specification. If the ANY were left off, a userid would have to possess all of the IF fields.

Userids that are no longer usable can also be reported with ESRFLIST. The following RUN command shows how to create a report of all userids that are EXPIRED, SUSPENDED, or CANCELLED according to the ACF2 RSS. In this example, we chose to use “WHEN” instead of “IF”.

```
     RUN  REPORT(ESRFLIST)    PARM(USER)     -
          TITLE(USERIDS NO LONGER USABLE)    -

          FIELDS(                       -
                    USERID              -
                    IMAGE               -
                    UA.NAME             -
                    UA.UID              -
                    UA.DELETED          -
                    ACF2.CANCEL         -
                    ACF2.SUSPEND        -
                    ACF2.CSDATE         -
                    ACF2.EXPIRE         -
                                   )    -

          WHEN(ACF2.CANCEL  EQ YES)     -
          OR(ACF2.SUSPEND  EQ  YES)     -
          OR(ACF2.EXPIRE  EQ  YES)      -
          OR(UA.DELETED  EQ  YES)
```

For more information about the fields specified in the above RUN command examples, see the *Data Dictionary Reference* or your RSS documentation.

Down-Loadable Output - ESRFDXD

ESRFDXD is a report utility similar to ESRFLIST. The difference is that ESRFLIST generates formatted report output, whereas ESRFDXD generates a comma-delimited sequential output file. The comma-delimited output can be downloaded into any PC application to create graphic reports, formatted letter/memo reports, etc. (Samples were shown earlier in this chapter.)

ESRFDXD uses the RUN command formatted identical to ESRFLIST, except that you run report utility ESRFDXD. There are additional parameters that can be used. For instance, DATADD specifies a DDNAME that contains the dataset you want the comma-delimited output to be sent to. ID( ) may be used to identify the beginning of each report RUN if you are appending multiple reports into the same output dataset.

The comma-delimited output is sent to a sequential file that can be used as input to any subsequent processing. If you have existing MVS applications that make use of comma-delimited files, ESRFDXD output can also be used.

TRIM/NOTRIM is an option to specify whether or not to have each field in the report generated at a fixed length. TRIM takes out all trailing blanks for each field in the comma-delimited list. NOTRIM can be specified to force each field to start at a specific location on the output record. Therefore, each field of information will always start in a particular column. Depending upon how you are using the output, you may want to specify NOTRIM on the RUN command. TRIM is the default.

There is no limit in ESRFDXD to the number of fields you can include. Several shortcuts have been provided, such as *xxxx*.ALL, where x*xxx* is the first qualifier of the dataname. For all ACF2 fields, you would specify "ACF2.ALL". For all RACF fields, you would specify "RACF.ALL".

The output dataset must be a Variable Blocked file with a Logical Record Length (LRECL) greater than sixteen characters. If the data you select requires a file with a LRECL larger than what you have provided, the data will be right truncated with warning messages posted. The maximum record length that will be formatted by ESRFDXD is 8,192 characters, including the record's four character Record Descriptor Word (RDW).

Using Selection Criteria: IF and WHEN

The information presented here is intended to acquaint you with this facility. To really understand and use the comprehensive package of data selections, refer to the RUN command description in the *Event System Command Reference.*

IF and WHEN are included as part of the selection criteria to narrow the scope of data contained on your reports. These specifications may be used with any dataname on the Masterfile relative to the report you are producing.

For example, if you want to see only those users with the ACF2 Audit attribute, you would specify: IF(ACF2.AUDIT). ACF2.AUDIT is the dataname for the field containing the ACF2 Audit privilege.

This selection criterion enables you to focus your reports on the specific information you need. The WHEN specification may also be used with all information contained on the Masterfile, including fields representing yes/no conditions. An example of this was previously shown.

Thresholds can also be created using a range of dates or counts. Full ACF2-style masking is available when using userids, resources, and dates in WHEN logic.

Relative date WHEN processing gives you the flexibility to design and schedule report runs to execute daily, weekly, or any frequency schedule you need. E-SRF establishes a *Relative base date* from the SYSTEM date. If "* - 1" is specified, then the statement would be interpreted as: today minus one day = yesterday.

There may be some cases where you want to temporarily modify that processing without having to change already scheduled report jobs. There are system options called SYSTEMDATE and RELATIVEDATE that allow you to alter the System or Relative date. Consult the *E-SRF Event System Command Reference* to obtain an understanding of these system options.

For example, if today is August 1, *-1 would refer to July 31. But, if you change the RELATIVEDATE to July 1, then *-1 will refer to June 30. This may be necessary during the implementation of E-SRF or the debugging of a particular issue. However, we **do not** recommend using this feature unless it is necessary.

If you want to determine if a date field exists, such as an Access Date, you can use YES and NO in a WHEN or IF specification. This would be used when you don't care what the value is, as long as there *is* one. For example, if you only want to see ACF2 userids on a report if they have been used, you can use:

WHEN(ACF2.ACC-DATE = YES)

This parameter will select userids for the report only if there is some value in the ACF2.ACC-DATE field.

To report on the userids that do not have a password associated with them, you can use:

WHEN(ACF2.PASSWORD EQ NO)

Userids would only be selected if the password field were blank.

The *Report Overlays Guide* has additional information about using IF and WHEN logic for reporting purposes. The *Command Reference* provides the best reference on how to use these and other criteria selection commands.

Using Selection Criteria: GROUP List

If you are grouping your resources, you can specify a list of groups to temporarily add to the group list associated with the report you are running. The group list is emptied out at the beginning of each RUN command.

Please note that if the report is running under Automatic Report Distribution, this specification will be ignored.

GROUP(*group1,group2,…groupn*)

Using Grouping Rule Comments for Data Selection

The comment information may also be used for data selection. Any information that is specified in the grouping rule's COMMENT is available for IF and WHEN selection.

Using Selection Criteria: OWNER List

If you are grouping, and you have your GROUPS and OWNERS defined to the Masterfile, you may specify a list of owners. The owner list is emptied out at the beginning of each RUN command.

If this is a non-distributed report, all groups targeted to owners named on the list will be *appended* to the current group list associated to the current report. The rules of group lists apply.

If this is a distributed report, the specified owners will be appended to the owner list. Owners named in the owner list will limit the scope of report creation to only those owners on the current owner list. If the owner list is empty, all owners will participate in the report creation and distribution.

OWNER(*owner1,owner2,…,ownern*)

Century Dates in Reports

Many customers use the two-character date format and it will always be supported in E-SRF.

Dates will appear on reports only in the standard format: MM/DD/YY (06/21/98) or the International format: MMDDD/YY (21JUL/99)

If you want the century added, specify CENTURY in the RUN command. This will cause the dates displayed on reports to be extended by two characters in length, thus requiring ten characters instead of eight characters. The two-character year will be replaced with the four-character full year, including the millennium and century.

Please note that all control dates posted anywhere in the Event System will appear in full four-character year formats.

Report Output Dataset(s)

When creating E-SRF reports, there are several possible destinations for the report output. By default, E-SRF Event Reports go to an output DDNAME called REPORTS. If REPORTS is not specified in the JCL, the Event System will automatically allocate the output dataset for you as needed for report production. The equivalent JCL is shown in the DD statement below:

```
//REPORTS   DD    SYSOUT=*
```

If the output goes directly to a dataset, or another location such as an online viewing mechanism, you can specify the REPORTS DD in the JCL with any type of characteristics appropriate for a print output file.

If you want a specific report to go to a report output file other than REPORTS, you may specify the DDNAME of the desired output file in the RUN command using the DDNAME specification.

A report output file is opened based on the characteristics of the current report. If the WIDTH of the current report is 133 characters, the dataset will be opened that way. If the dataset were pre-allocated with a larger LRECL, then the report will still be formatted 133 characters, but the remaining characters will be blanked out.

The largest valid LRECL for a report output file is 255 characters.

If subsequent reports reference the same report output file, the data will be placed behind the previous report data contained on the file.

The only limit to how many report datasets may be present in an execution is the limit set by the operating system for the maximum number of files that may be opened for a single job's execution.

Automatic Report Distribution

It is possible to let the Event System produce individual reports for each USER defined to the Event system. This would include data belonging to each GROUP the USER participates in or is an interested party to.

To run a comprehensive report distribution, you must have a group structure in place and a GROUP and OWNER structure defined to the E-SRF Event System Masterfile.

The only requirements of report distribution is setting up your OWNER and GROUP definitions on the Masterfile and providing the necessary grouping of your resources. Once this is ready, the actual production of distributed reports is quite easy.

You do not have to complete either of the above in order to start distributing reports.

To run a report in distributed mode, simply add the DIST(OWNER) specification to your RUN statement. Almost every report (except several control reports) may be distributed. Owner distribution will produce a single report containing data from all groups that are owned by the owner. This means that if ten groups are owned (by a particular owner), the owner will get a single report containing data from all ten groups initiated by the single RUN command.

You can distribute by group by specifying DIST(GROUP), but it is not recommended. You still distribute reports to their owners, but it will build a separate report for each group. So if an owner has ten groups, the owner will receive ten reports, one for each group initiated by a single RUN command.

Owners can also be considered interested parties to a group. This means data associated to the group will be presented to the owner (or interested party), just as if he owned the group. The owner of the group would also get the same data, which means it is possible to report the same data to more than one owner.

If you have an empty set of grouping rules and have not defined any OWNERS or GROUPS to the E-SRF Masterfile, report distribution can still be run. All of the data will be associated to the DEFAULT GROUP, which is owned by the DEFAULT OWNER. You will end up with a single distributed report for the DEFAULT OWNER containing all the report data.

You can phase this process in gradually as you build your grouping structure and define the GROUP and OWNERS to E-SRF.

To understand more about report distribution, please review the chapter *Grouping and Ownership* in this publication and also reference the *Event System Command Guide* for additional information.

This page intentionally left blank

Chapter 11: Event System Grouping

This chapter introduces E-SRF grouping concepts. The following topics are discussed:

- What is *grouping* and why do we group?
- How are grouping rules set up for users and resources?

E-SRF Grouping

Grouping is an integral part of E-SRF processing. It is used for associating datasets, non-datasets, and user resources into meaningful collections, selecting records to update, selecting information to report on, and automatically distributing reports. Grouping may affect updating the Masterfile as well as Event Report processing.

The association of a group name with resources is *not* stored on the E-SRF Masterfile. The associations are dynamically determined when requested because of update or report processing. When applying security event data during a Masterfile update, it is possible to exclude information based on grouping associations.

The purpose of grouping users and resources is to make it easier to create useable reports containing only information that a particular department head or owner wants to see. Grouping Rules can provide selection criteria and additional information in reports. Once grouped, the group names are used to identify the specific resources or users to include in a report. They can also be used to add a column on a report that identifies which group a user or resource belongs to. The group name can identify more information about the ownership of a resource than the actual resource name or class.

If the Report Distribution Facility is used, E-SRF will create and distribute reports based on the defined groupings. In order to take advantage of this facility, groups and report jobs need to be set up. The jobs can then be scheduled and the defined owners will periodically receive information about *their* resources and users.

With this type of automatic functional reporting, true ownership of resources is assured. The designated owners receive the information they need about what is happening to the resources they are responsible for. They can follow-up and change security definitions so that no unauthorized access is allowed, and unauthorized attempts are dealt with quickly.

What is a Grouping Rule?

The Resource Grouping Facility associates datasets and non-datasets with a group name in **Grouping Rules**. This grouping is used to put resources together into a functional group. For example, all resources that are used by the Payroll Department can be grouped together regardless of their actual naming conventions.

Once grouped, E-SRF uses that association to create reports showing events or access definitions for only the resources that are part of that group. The group names themselves can be printed as a column on E-SRF Event Reports.

If you wish to include users and sources (terminal IDs) in groups, you can create special Grouping Rules for that purpose. This is not recommended unless you have a specific need for them. The process to create them is extensive and E-SRF provides other mechanisms for reporting on users and sources.

What is an E-SRF Group?

An E-SRF group is a one- to sixteen-character name identified in a Grouping Rule that relates to one or more specific resources contained on the Masterfile.

If a group is going to participate in Report Distribution, or if information specified on the group header is required for report selection, the group will also have a definition on the E-SRF Masterfile.

The Masterfile definitions provide information to the Event System about the group, who owns it, what owners (if any) may be interested in the group, and other various data that may or may not be relative to your reporting requirements. Please refer to the chapter on *Report Distribution* in this publication for more information. Defining group headers is discussed in the next chapter.

In summary, an E-SRF group is a one- to sixteen-character name given to a logical or functional set of users, datasets, and non-dataset resources. For example, Group: PAYROLL may contain certain datasets that only apply to the Payroll Department.

Working with Grouping Rules

To utilize grouping in E-SRF Access Analysis or Event Reporting, Grouping Rules are required. The rule writing process does not have to take place all at once. Grouping is an advanced feature that can be implemented over a period of time.

Any E-SRF report can be created and used without Grouping Rules. You can slowly phase in groupings for several resources or groups at a time. The datasets and non-dataset resources that have not yet been grouped can be assigned a default group until you have the time to create rules for them.

Because grouping is done dynamically on demand, you can maintain more than one grouping rule file. At execution time, the entire grouping scheme (consisting of one or more grouping rule files) is evaluated. This provides a way to group resources in more than one grouping scheme.

For more information about the grouping process and rule writing, refer to the *EKC Integrated Facilities Resource Grouping Facility Guide*.

Writing and Changing Grouping Rules

Once you have determined the grouping structure by identifying what datasets and non-dataset resources belong in which Groups, as well as who “owns” the group(s) for reporting purposes, these groups are then related to their objects contained on the Masterfile.

EKC provides the Integrated Resource Grouping Facility as a means to associate group names to Masterfile objects.

There are three levels of grouping within the Event System:

Resources (which include both datasets and non-dataset resources). Non-dataset resources include resources such as CICS transactions. If grouping is active, resource grouping is mandatory.

Sources (which indicate where a particular event took place). A VTAM nodename is an example of a source. You have the option to either group sources or not. The SET SOURCE(…) command controls the capability to group sources. If you specify SET SOURCE(NONE), there will be no source grouping, and no source rules will be referenced. If you specify SET SOURCE(*classname*), sources will be grouped (as a resource) using the specified classname.

Users (which are the individual users that cause events to take place). You have the option to either group users or not. The SET USERID(…) command controls the capability to group users. If you specify SET USERID(NONE), there will be no user grouping, and no userid rules will be referenced. If you specify SET USERID(*classname*), users will be grouped (as a resource) using the specified classname.

The Integrated Resource Grouping Facility uses *"rules"* to define the relationships between *resources* (objects) contained on the Masterfile with the proposed *group names*. Grouping rules may look a little like ACF2 rules, but are not interchangeable and have a completely different function. Grouping rules are NOT ACF2 rules. E-SRF Grouping Rules are only used to associate resources to group names.

Each Grouping Ruleset, representing a high level qualifier of a dataset or a resource class is written in its own member of a rules ***source*** partitioned dataset (PDS).

These "*rulesets*" are maintained in the source PDS using a text editor such as ISPF. This PDS of *"rulesets"* provides a base that is *compiled* and stored as *object code* in a non-PDS **Rules Object File**. This output dataset is considered to be the *"RULES"* file that is discussed throughout this publication and other publications relating to the use of the EKC Integrated Grouping Facility.

Do not attempt to use the base PDS containing non-compiled rules as a rule input file for the E-SRF Event Reporting System. If you do, you will receive errors from the grouping routines and possibly a program abnormal termination (ABEND). When Grouping Rules are required, *ALWAYS* use the compiled **Rules Object File**.

In order to provide some familiarity with the Grouping Rule concept, a high level discussion follows. When you are ready to build your rules, please reference the *EKC Integrated Facilities Resource Grouping Facility Guide* for assistance.

There are two different types of Grouping Rules:

- Dataset Grouping Rules - with a $INDEX keyword
- Resource Grouping Rules - with a $CLASS keyword

Grouping RESOURCE Objects

There are two types of resources stored in the Resource Segment of the Event Reporting System Masterfile, dataset and non-dataset resources. The EKC Integrated Resource Grouping Facility maintains two types of Grouping Rules, dataset and non-dataset Grouping Rules. The Event Reporting System makes use of non-dataset resource Grouping Rules for grouping all non-dataset resource objects. For dataset resource objects, the default is to use dataset Grouping Rules. However, you can use non-dataset Grouping Rules by providing the SET DATASET(*classname*) command. When this command is specified with any classname (including '*DATASET*'), resource Grouping Rules will be used to group dataset resources instead of dataset Grouping Rules. Please note that the classname used for reporting will always be DATASET regardless of what classname is present in this parameter.

Grouping Non-Dataset Resources

Grouping Rules for non-dataset resources are written based on their resource class or type code.

For example, all CICS transactions are grouped for RACF systems with a $CLASS(TCICSTRN).
In ACF2, they are placed in a resource type, such as CKC using resource rule $CLASS(CKC).

```
$CLASS(TCICSTRN)    DEFGROUP(UNKNOWN)
ACFM                GROUP(SECURITY)
ADF*                GROUP(ADMSYSTEMS)
A-                  GROUP(ACCOUNTING)
B-                  GROUP(BUDGETING)
CEMT                GROUP(SYSTEMS)
E-                  GROUP(DEVELOPMENT)
F-                  GROUP(DEVELOPMENT)
```

The above rule example shows the basic syntax for writing Grouping Rules for resources such as CICS transactions secured by a RACF system. Each rule line represents a transaction or group of transactions.

The DEFGROUP keyword defines the default group name (UNKNOWN in this case) if the resource does not match any of the lines specified in the rules.

CICS transaction ACFM is associated to the SECURITY group. Any transaction id beginning with ADF and followed with any one character is associated with the ADMSYSTEMS group. Transactions beginning with A belong to the ACCOUNTING group. All transactions beginning with B belong to the BUDGETING group. The CEMT transaction is associated to SYSTEMS. Any transaction beginning with D is associated with the DEVELOPMENT group. All others are associated with the Default Group UNKNOWN.

If the example were for ACF2, the $CLASS(TCICSTRN) would have been specified as the three- character ACF2 Resource Rule type code for CICS transactions.

```
$CLASS(CKC)        DEFGROUP(UNKNOWN)
```

For an expanded discussion of masking, refer to the *EKC Integrated Facilities Resource Grouping Facility Guide*.

Non-dataset ACF2 Grouping Considerations

ACF2 uses generalized Resource Rules to provide non-dataset access control. An ACF2 Resource Rule consists of an ACF2 Information Storage "CLASS" (normally "*R*") followed by the three-character ACF2 "*TYPE*" code, creating a four-character string. The *key* is normally a one- to 40-character masked resource name, which follows the previously described four-character string, creating a full 44- character ACF2 resource name. Later releases of ACF2 augmented the resource name with additional resource name data appended to the resource key.

The E-SRF Event Reporting System (E-SRF) uses the eight-character industry standard classname, followed with a one- to 44-character resource name.

When the E-SRF classname is constructed from an ACF2 event, the E-SRF classname is represented by the three-character ACF2 *TYPE* code. *The ACF2 Information Storage "CLASS" is NOT included*. Thus the CICS transaction "CEMT" would have an E-SRF classname of "CKC" (assuming CKC was the ACF2 *TYPE* code), with a resource name of CEMT.

Grouping Rules for ACF2 non-dataset access are written to this specification (classname of CKC, followed by the resource CEMT). The ACF2 Information Storage class "*R*" is not specified in class associated with the resource access.

Maintenance functions for ACF2, such as the ones contained on the Resource Maintenance (RM) Masterfile object are slightly different. All maintenance for ACF2 Information Storage data contains an E-SRF classname consisting of four characters. The first character is the ACF2 Information Storage "*CLASS*", followed by the three-character ACF2 "*TYPE*" code. This is because there may be multiple uses of the same ACF2 "*TYPE*" code for various ACF2 Information Storage "*CLASSES.*" Additionally, it provides for better E-SRF Event Reporting data selection.

Therefore, in the current E-SRF product offering, different grouping rules are required for resource reference (*three-character classname*) and the maintenance of resources (*full four-character classname*).

Grouping Datasets - Using Dataset Grouping Rules

To create groups for datasets, the keyword $INDEX is used to represent a high-level qualifier or high-level index. This will allow all datasets that start with the specified high-level index to be included in the Grouping Rule. For example, to group all datasets that start with a high level qualifier of SYS1, the rule would look similar to this:

```
$INDEX(SYS1)   DEFGROUP(UNKNOWN)
BRODCAST       GROUP(TELECOM)
MAN*           GROUP(SECURITY)
PARMLIB        GROUP(SYSTEMS)
PROCLIB        GROUP(SYSTEMS)
UADS           GROUP(SECURITY)
```

In this simplified example, each dataset that starts with SYS1 is associated with the appropriate group name. SYS1.BRODCAST is associated with the TELECOM group. SYS1.PROCLIB and SYS1.PARMLIB are associated with a group called SYSTEMS. SYS1.UADS is associated with a group called SECURITY. All others not defined are associated with the DEFGROUP (default group) UNKNOWN.

"SYS1.MAN*" is a *mask*. The asterisk (*) represents any **one** character. So all SYS1 datasets that contain four characters in the second qualifier starting with MAN will be grouped in the SECURITY Group. This masking allows the rules to be written more generically if needed. In addition to the asterisk as a masking character, the dash (-) is also a masking character representing multiple characters. A rule line such as the one below could have been be supplied in the example above:

```
P-                  GROUP(SYSTEMS)
```

This rule would include all SYS1 datasets with a P in the beginning of the second qualifier. In the above example, both SYS1.PARMLIB and SYS1.PROCLIB would be represented by P-. This would replace the need for two statements to cover PARMLIB and PROCLIB.

The masking characters can be thought of as a shorthand method to writing Grouping Rules in a quick and concise manner.

There are many ways to utilize masking in Grouping Rules. For more information and examples, consult the *EKC Integrated Facilities Resource Grouping Facility Guide*.

Grouping Datasets - Using Resource Grouping Rules

To create groups for datasets using resource rules, specify the SET DATASET(*classname*) command. The recommended classname is DATASET. This will cause resource rules to be used to associate groups to datasets. To reverse this, specify SET DATASET(NONE). This will cause the Event Reporting System to revert to its system default of using dataset Grouping Rules for dataset resources.

Grouping datasets in this way allows datasets to be treated like any other resource.

The *classname* must be a class that is unique and not currently being used. Assume you chose DATASET as the *classname*. You would specify:

```
SET    DATASET(DATASET)
```

For example, all dataset resources will be grouped using resource Grouping Rules with a class of DATASET. This was afforded by the SET DATASET(DATASET) command. The dataset Grouping Rule concept example is illustrated below:

A special $CLASS rule can be created to group your datasets similar to grouping resources. The $CLASS(DATASET) rule allows you to associate actual dataset names to specific groups.

```
$CLASS(DATASET)     DEFGROUP(UNKNOWN)
SYS*.-              GROUP(TECHSUPPORT)
SYS1.BRODCAST       GROUP(TELECOM)
SYS1.MAN*           GROUP(SECURITY)
SYS1.PARMLIB        GROUP(SYSTEMS)
SYS1.PROCLIB        GROUP(SYSTEMS)
SYS1.UADS           GROUP(SECURITY)
SYS4*               GROUP(APPSUPPORT)
ACCTG.-             GROUP(ACCOUNTING)
INV.-               GROUP(INVENTORY)
SALES               GROUP(SALES)
SALES.MKTG.-        GROUP(MARKETING)
MKTG.-              GROUP(MARKETING)
```

The above rule example shows the basic syntax for writing a Resource Grouping Ruleset for datasets. Each rule line represents a particular masked representation of dataset name(s).

Grouping SOURCE Objects

The Event System allows you to group Source Objects that exist on the Masterfile. The Objects will contain source names, such as the network addresses (VTAM LUs) of the physical sources (terminals) you wish to group. For example, it could be used to track events occurring within a specific set of terminals. Before this commitment is made, however, make sure it is required. It takes a considerable amount of virtual storage and CPU resources to maintain it. Grouping the Source segment will increase E-SRF overhead. Because of the explosion of dial-in capabilities and VTAM LU pooling, Grouping Rules of this type will have limited use.

Grouping the SOURCE Masterfile segment is not recommended unless you have a specific need. The default is SET SOURCE(NONE).

Please note it is possible to turn off Source grouping for specific executions where user grouping is not needed. This will greatly reduce the associated overhead.

The following command may be presented to temporarily turn off user grouping:

OPTION GROUPING(NOSOURCE)

To set up the E-SRF Event System to group the Source Segment specify the following command:

SET SOURCE(*classname*)

The *classname* must be a class that is unique and not currently being used. Assume you chose TERMINAL as the *classname*. You would specify the following:

SET SOURCE(TERMINAL)

A special $CLASS rule can be created to group sources similar to grouping resources. The $CLASS(TERMINAL) rule allows you to associate terminal ids (VTAM LUs) to a specific group. The resource name is the eight-character source name, as shown in the example below:

```
$CLASS(TERMINAL)    DEFGROUP(UNKNOWN)
DIAL-               GROUP(TELECOM)
POOL-               GROUP(TELECOM)
LUTS-               GROUP(TECHSUPPORT)
LUA-           GROUP(ACCOUNTING
LUS-           GROUP(SALES)
LUM-           GROUP(MARKETING)
```

Grouping USER Objects

The Event System allows you to group User Objects that exist on the Masterfile. It may be very desirable to group your users, but it takes a considerable amount of virtual storage and CPU resources to maintain the user grouping structures. Grouping the User Segment will increase E-SRF overhead and should be considered a major factor in determining whether or not to group User Objects.

Please note it is possible to turn off user grouping for specific executions where user grouping is not needed. This will greatly reduce the associated overhead.

The following command may be presented to temporarily turn off user grouping:

OPTION GROUPING(NOUSER)

If you need to group your users, set up the E-SRF Event System to group the User Segment by specifying the following command:

SET USERID(*classname*)

The *classname* must be a class that is unique and not currently being used. Assume you chose USER as the *classname*. You would specify:

SET USERID(USER)

A special $CLASS rule can be created to group users similar to grouping resources. The $CLASS(USER) rule allows you to identify the user with a sixteen-character resource name consisting of the IMAGE and E-SRF Event System *Universal Identifier* (UID).

The resource name consists of two fields separated by a period (.). The E-SRF IMAGE is the first field, and because of the period (.), the IMAGE becomes a separate, maskable field from one to eight characters in length. The second field is the E-SRF Event Reporting System's Universal User Identifier, which is Resident Security System independent.

The E-SRF Event Reporting System's UID is formatted as follows:

For ACF2: The E-SRF UID field is constructed from the ACF2 LOGONID record identical to the way ACF2 builds it for general access validation. The fields are concatenated together one behind the other with nothing separating the fields. The fields are the exact length contained on the LOGONID record. Normal UID fields match up against the ESRF UID and work the same way as when writing ACF2 rules.

For RACF: The E-SRF UID field is constructed from three eight-character RACF fields: the RACF *OWNER*, the *DEFAULT GROUP*, and the *USERID*. The fields are concatenated together one behind the other with nothing separating the fields. The fields are each eight characters long, totaling twenty-four characters.

The following example shows you how to code the character string needed to group users. The resource name for the sample RACF user: **CAROLYN**, owned by: **TECH**, in the default group **MVS** belonging to the **CHICAGO** E-SRF Image would be:

```
0        1         2         3         4
1...5....0....5....0....5....0....5....0
CHIGAGO.TECH    MVS     CAROLYN
```

We would provide a resource name that consists of two individual fields. These fields are the IMAGE and an ESRF Universal Identifier (UID) that will represent that particular user within a particular IMAGE. These two fields are treated as separate maskable components of the resource name by the grouping facility.

In our example, positions one through seven contain the E-SRF *IMAGE.* Even though the IMAGE may be from one to eight characters long, you only have to specify the characters you need to represent the desired IMAGE. All other characters up to the period are considered masked. If necessary, you may specify spaces before the period (to identify 'CHGO' 'CHGO1' and 'CHGO1A').

Position eight in our example consists of the field separator (a period), which makes the two fields separate maskable entities.

Position nine through thirty-three consists of the twenty-four character E-SRF UID. Notice the blanks. These blanks must be present, since all three fields collectively participate as a single maskable entity (since the match up is on an entire twenty-four character field). Trailing blanks are treated as described in the IMAGE discussion above.

The resource name is constructed and presented to the user grouping facility for evaluation.

As you can see, rules are written against an IMAGE and the universal user identifier, not the userid itself. This provides great flexibility in rule writing.

An example of a User grouping rule for a RACF Image may look like:

```
$    CLASS(USERS)                 DEFGROUP(UNKNOWN)
     CHICAGO.TECH****MVS-         GROUP(TECHSUPT)
     CHICAGO.PAYROLL-             GROUP(PAYROLL)
     CHICAGO.-               GROUP(CHICAGO)
```

Note how the two fields participate in logical grouping. The first two lines direct specific sets of users that are in the CHICAGO E-SRF image and match the UID masks to their respective groups. The third line denotes anything else in the CHICAGO image should be in the CHICAGO group. If an image other than CHICAGO is encountered, there are no rules that will match it. The resource will be grouped in the UNKNOWN group as directed by the DEFGROUP specification.

Grouping Rule Syntax Reminder:

Remember that when you are creating Grouping Rules for these resources, and the resource name contains any blanks, you must enclose the resource name in single quotation marks.

```
     $CLASS(USERS)                DEFGROUP(UNKNOWN)
     'CHICAGO.TECH    MVS-'       GROUP(TECHSUPT)
```

Notice the four blanks between TECH and MVS. The single quotation marks actually establish the resource name boundary and allow the use of embedded blanks in the actual resource name text.

Chapter 12: Defining Owners and Groups to the Masterfile

Defining OWNERS and GROUPS to the Masterfile is mandatory only if the Report Distribution Facility is used and/or if information contained in either of these definitions is required for reporting or report data selection.

What is an E-SRF OWNER?

The concept of *owners* is unique to the Event Reporting System. This is *not* to be confused with the terminology used by the Access Analysis product component. In the Event System, resources are grouped together dynamically using the EKC Integrated Grouping Facility. This determines the resource's group name. In Access Analysis, the group name is the owner. In Event Reporting, this is NOT the case.

By defining your individual groups to the Masterfile, grouping may be further extended. A group definition on the Masterfile allows you to relate specific groups to *OWNERS*, as well as relating as many as 16 other OWNERS as *interested parties* to the same groups.

An *OWNER* in Event Reporting is a definition on the Masterfile. In report distribution, the OWNER ID is the output file used to contain all the owner's reports. Owners are not specified in grouping rules. Grouping Rules target GROUPS; group headers on the Masterfile target OWNERS.

Once an OWNER is defined on the Masterfile, it can be targeted by group header definitions. A GROUP *must be owned* by an OWNER and optionally have other owners targeted as *interested parties*. There may be as many as 16 interested parties associated to a group. This means data relating to a GROUP could be distributed to as many as 17 (one primary and 16 secondary) owners in your organization (when the Report Distribution Facility is utilized).

In summary:

- OWNERS *own* GROUPS. GROUPS are also defined on the Masterfile. When defining a GROUP on the Masterfile, you target the defined OWNER and any defined OWNERS that are considered *interested parties* to the information associated to the group.

- Grouping Rules dynamically determine the group's name for a resource. The group header controls how the group will be processed (by supplying information such as the group's OWNER). Owner definitions identify the various owners and contain information on how to process reports.

What Exactly is an E-SRF GROUP?

Grouping objects contained on the Masterfile has been discussed in previous chapters of this publication. Relating a resource to a particular group is provided on demand (by the EKC Integrated Grouping Facility). The group name ascertained by this grouping process may be further defined to enable the Event Reporting System the ability to distribute reports, provide extended reporting data selection, augment report sort sequences and list group-related information on reports.

When you defined your grouping rules, you determined a grouping structure. If you were to publish a report with the dictionary name for the group name, you would see the group name along with the information being reported. The EKC Integrated Grouping Facility only provides group names and comment information relative to the object being grouped. The Event Reporting System may require additional information about a particular group such as: *who owns it?*

To address the extended needs of the Event Reporting System, a *group header* object can be established for some or all of your groups. Consider the group header as an extension of the group name.

The GROUP Masterfile definitions provide information to the Event System about the characteristics of the GROUP. It contains information such as: who owns it, what owners may be interested in the group, and other pieces of data that can be relevant to your reporting requirements. The group's resources are not specified on the Masterfile. They were defined earlier in a RULES source PDS discussed in the previous chapter.

Do I Need To Do All of This?

There is a certain amount of effort required to provide the OWNER and GROUP definitions on the Masterfile.

This process is required if you intend to do one of more of the following:

- Distribute reports to individual OWNERS (the Report Distribution Facility).
- Use information on either the owner header or the group header for one or more of the following reasons:
    - Selecting data to report on.
    - Controlling the sequence of reports using information contained on either of these headers.
    - Displaying information contained on either of these headers.

The OWNER and GROUP definitions can be phased in gradually. There is no need to implement this powerful option in its entirety all at once.

Relationship of Owners, Groups, and Grouping Rules

The Owners, Groups, and Grouping Rules are interrelated in E-SRF Event System processing. E-SRF asks the Resource Grouping Facility for a group name to associate to a specific resource.

Once the Event System knows the group name for a particular resource, that group name may be used to select the resource for reporting.

To run an undistributed report showing the event journals for all resources associated to the owner NELSON, the following selection specification in the run statement could be used:

OWNER(NELSON)

The OWNER selection specification will cause the Event System to build a *grouplist* of all groups that target the owner: NELSON.

E-SRF knows all the groups that NELSON is involved in. This is made possible because the group headers were set up with OWNERS and interested parties. In the above example, any group that is owned by NELSON or NELSON is specified as an interested party will be added to the grouplist. When the report is produced, the group names will be matched to groups appearing on the grouplist. In our example, a report would be produced with all events associated to groups that are related to the owner: NELSON.

The same type of processing could be used for report data selection. For example, the owner header contains a field called OA.ADDRESS5 that may be used to contain information about the owner’s address. Let’s say you supplied a location code of “CHICAGO” in this field for certain owners who work in the Chicago office. If you want to produce a report containing events that occurred for any groups whose owners have “CHICAGO” in OA.ADDRESS5, specify the following WHEN statement:

WHEN(OA.ADDRESS5 EQ CHICAGO)

The report production processing would look at an event, obtain the event’s group name (from your grouping rules), use the group name to relate a group header, use the OWNER specification in the group header to relate to the owner header, and finally compare the OA.ADDRESS5 data contained on the owner header object to the specification contained on the WHEN statement. The result would be a report containing event data that is related to owners who have CHICAGO specified in OA.ADDRES5.

Consider sorting. You may want to sort a report by the criticality code contained on the group header: (GA.CRITICAL). Like the example above, the criticality code could be used in the sort control statement:

SORT(GA.CRITICAL)

The end result will be a report sorted in the GA.CRITICAL sequence.

There is a wealth of information about the Masterfile and how things are related in the *E-SRF Event System Masterfile and Data Dictionary Reference*. This manual deals with how the Masterfile is constructed and a brief description of each dataname contained on every segment and object of the Masterfile.

If you know the Masterfile, you know this product. The programs and commands are there to accommodate the Masterfile. The Masterfile is there to accommodate your reporting needs. Groups and owners are necessary to provide the flexibility needed to select and report the way *you* want and most of all, to provide a means to automatically distribute reports to the data owners within your organization.

Defining/Changing GROUPS and OWNERS

Commands to Define Groups and Owners

OWNERS and GROUPS are defined on the Masterfile using the Command Processor.

The syntax is identical for both OWNER and GROUP definitions. The only difference is the actual datanames associated to the object being defined.

There is no crosschecking between OWNERS and GROUPS. If you target an owner that does not exist, the DEFAULT owner is used it its place.

You can INSERT, ALTER, and DELETE any OWNER or GROUP definition. You cannot delete the DEFAULT definitions.

The information associated with OWNERS and GROUPS is contained in the *E-SRF Event System Masterfile and Data Dictionary Reference.* This publication should be the ultimate source of what fields are available and the characteristics of these fields.

The following is an example of how this is done. Please refer to the *E-SRF Event System Masterfile and Data Dictionary Reference* for more information on the many fields contained on these headers and what can be placed in them.

```
INSERT    GROUP(PAYROLL)      -
          DESC(PAYROLL TRANS) -
          OWNER(PAYMGR)

INSERT    OWNER(PAYMGR)       -
          NAME(PAYROLL MGR)   -
          JESCLASS(A)         -
          JESDEST(USER01)
```

For more information about the E-SRF commands to insert, change, or remove Groups and Owners, refer to the *Command Reference* and the *E-SRF Event System Masterfile and Data Dictionary Reference*

DEFAULT Owner and Group Definitions

When the Masterfile was first established, default owner and group headers were inserted on the Masterfile for you. These headers must always be there and cannot be removed.

These headers are present in the event that a specific header was required, but the targeted header is not defined on the Masterfile. In report distribution, the *reporting group* can only be a group previously defined to the Masterfile. The same is true for the *reporting owner*.

<u>Defining OWNERS and GROUPS to the Masterfile</u>

The Owner default ID is DEFAULT. The Group default ID is DEFAULT.

The information supplied on these headers is the bare minimum for the product to function. You can alter any field contained on these headers. The normal mode of operation is for the DEFAULT GROUP to be owned by the DEFAULT OWNER. You can change this if desired.

Sample Scenario

To implement grouping in E-SRF, several clients have successfully followed the plan below:

- Create your group structure and define your rules as discussed in previous chapters of this publication.

- You may now take advantage of the group designations that associate groups to objects in the Event Reporting Facility. To use Automatic Report Distribution, continue through the following steps.

- Define the owners in the OWNER segment of the Masterfile. In order to use Automatic Report Distribution, any targeted OWNER or INTERESTED PARTY specified in a group header must be defined in the OWNER segment of the Masterfile. If they are not, the data will go to the default OWNER. This is what allows you to phase in Automatic Report Distribution a little at a time.

- Supply any other information about your owners you deem necessary. See the *E-SRF Event System Masterfile and Dictionary Reference* for additional information that you may include on the Owner definition.

- Define the groups in the GROUP segment of the Masterfile. Each group name you specify in the Grouping Rules (including all default group names) must be defined as a GROUP in the GROUP segment. If the groups are not defined to the Masterfile, the *reporting group* will be the DEFAULT group. E-SRF will route the supplied default group header to the specified OWNER. Again, this allows you to phase in Automatic Report Distribution.

- Target an OWNER for each group and provide any other information you deem necessary. See the *E-SRF Event System Masterfile and Dictionary Reference* for additional information that you may include on the group definition. Each group requires an owner. However, you can additionally include other owner(s) who should receive reports about the group by supplying those owner Ids in the *interested party* specifications.

Run the ESRFOWNX (and other diagnostic reports) to make sure you are satisfied with your grouping structures and your OWNER and GROUP definitions.

Special Processing of Owner and Group Headers

Mass Deleting Owner or Group Headers

It may be desirable to maintain all owner and group header definitions in a single job. In this case, all of your headers are redefined each time you want to make changes. For this type of maintenance to be effective, a means to delete all existing headers must be provided. This will enable the owner and group header definitions to be re-installed on the Masterfile all together.

```
DELETE      GROUP(!ALL)
DELETE      OWNER(!ALL)
```

If the above commands are entered, the result will be a Masterfile containing only the default headers.

Defining Groups and Owners With a Spreadsheet

A program is supplied with the Event System that enables you to create an external spreadsheet and upload the spreadsheet into a file that may be used as an input to a program that will generate the necessary ESRFCMD input.

Owner and Group Masterfile Definitions

As previously discussed, data owner definitions (**OWNERS**) should be established on the E-SRF Masterfile if Report Distribution is utilized or if data selection, sequence controlling, or display of any owner header fields are required. GROUPS are related to these OWNERS when distributing or creating reports that are owner based.

Additionally, **GROUPS** should also be defined on the E-SRF Masterfile. These definitions provide the means to relate a GROUP to an OWNER and possibly other owners who may be *interested parties*.

As explained in this guide, the actual relating of *resources* to groups is performed on demand by the EKC *Integrated Grouping Facility,* which tells E-SRF the *group name* associated to the resource.

If the group name is not defined on the Masterfile, it will display correctly on reports but when distributing reports, the distribution *group name* will be DEFAULT. The group (DEFAULT) is associated to the OWNER established for it. When the product is initially installed, the DEFAULT Owner for the DEFAULT Group is DEFAULT. This allows you the ability to phase in a report distribution scheme over time.

In summary, OWNER and GROUP definitions on the Masterfile are key players in the way you use E-SRF, how the reports are displayed, and how they are distributed.

Optional Owner/Group Command Generator

Up until now, the only way to maintain these definitions was to use the E-SRF Command Processor to process the individual definitions. The current release of E-SRF is a batch processing system and therefore has imposed restrictions on how this type of administration can be performed. The following enhancement was developed to provide some relaxing of these constraints until a more suitable approach evolves with the growth of the product.

ESRFCGEN Offline Batch Utility

The offline batch utility (ESRFCGEN) was developed to allow the use of a "*comma delimited*" text file to be generated (for example, from Microsoft's EXCEL product). This uploaded text file is then used as input to the ESRFCGEN program that will build the standard E-SRF commands normally created by you. These commands can be reviewed and possibly edited before being presented to the E-SRF Event System Command Processor (ESRFCMD).

With this tool, you have the option to maintain your OWNER and GROUP definitions in an external spreadsheet, instead of entering maintenance directly into the batch job with control statements supplied by you.

You can retain your definitions in the spreadsheet and use this data to maintain your OWNER and GROUP definitions outside of E-SRF.

This facility is OPTIONAL.

This feature is a processing layer above the normal E-SRF Command Processor and its use is strictly at the discretion of the E-SRF administrator.

The Spreadsheet

The spreadsheet can be used as the controlling vehicle to maintain your Group and Owner definitions. It can also be used to initially define and/or provide subsequent maintenance when required.

The only validation performed when the data is pre-processed is what is required to convert the spreadsheet's uploaded file into E-SRF Event Reporting command input.

An example of a spreadsheet follows.

	A	B	C	D	E	F
1	**OA**	**ID**				
2	DELETE	KAREN				
3						
4	**OA**	**ID**	**NAME**	**ADDRESS1**	**PHONE**	
5		TOM	TOM SMITH	ROSEMONT, IL	847 296-8010	
6		SUE	SUE SMITH	ROSEMONT, IL	847 296-8010	
7		GEORGIE	GEORGIE SMITH	ROSEMONT, IL	847 296-8010	
8		BARRY	BARRY SMITH	ROSEMONT, IL	847 296-8010	
9		DIANE	DIANE SMITH	ROSEMONT, IL	847 296-8010	
10		SARAH	SARAH SMITH	ROSEMONT, IL	847 296-8010	
11						
12	**GA**	**ID**	**NAME**	**OWNER**	**PARTY1**	
13		TEST.ACCT	ACCOUNTING TEST	TOM	DIANE	
14		TEST.OE	ORDER ENTRY TEST	TOM	GEORGIE	
15		TEST.SAM	SALES AND MKTG TEST	TOM	SARAH	
16		PAYROLL	PAYROLL SYSTEM	DIANE		
17		ACCOUNT	ACCOUNTING SYSTEM	DIANE		
18		ORDENT	ORDER ENTRY SYSTEM	BARRY		
19		SHIPPING	SHIPPING SYSTEM	SUE	BARRY	
20		ACCTPAY	ACCOUNTS PAYABLE	DIANE		
21		ACCTREC	ACCOUNTS RECEIVABLE	DIANE	BARRY	
22		SALES	SAVES REPORTING	SARAH		
23		MKTG	MARKETING SYSTEM	GEORGIE		

Establishing a header (processing mode) row (rows 1, 4 and 12):

The first column is the operation.

First, establish a processing mode by specifying GA (for a group header) or OA (for an owner header).

The remaining columns are individual datanames.

Specify the desired datanames, which relate to the desired object (processing mode of GA or OA).

They may be in any order, and their segment/object prefix may be omitted.

You only need to specify the fields you are using. The command generator will create a complete shell of all supported datanames for the specific object, with your specifications "sprinkled in." This means if you only specify a subset of the required fields, the fields unspecified will be blanked out when the maintenance is applied to the Masterfile.

In this example, we chose to define all the owners first, with the groups following. You are not restricted to this approach. You could define a header row for a single owner, followed by the single owner, and then supply a group header row with the owner's groups defined. This would visually make relating the owners and their groups easier.

Specifying objects to be maintained (rows 2, 5 to 10, and 13 to 23):

The first column is the operation.

If left blank, the command will DELETE the existing definition, INSERT a shell, and REPLACE all fields relating to the object with a completed shell that contains all datanames belonging to the object, along with your specifications "sprinkled in." **This is the recommendation**.

YOU may specify DELETE as in row 2 of the example. The object named by the ID will be deleted. No other fields need to be specified when deleting an object.

The remaining columns are individual datanames.

Specify the fields that relate to the columns established by the processing mode (GA or OA).

Organization

You may flip back and forth between processing modes. The example specified owner, followed by group definitions. You may do it in any order. One approach is to define an owner, followed by the owner's groups.

The File Transfer

E-SRF is a mainframe product and relies on its data residing there. Transferring the data from your spreadsheet to the mainframe requires an UPLOAD from your PC to the Mainframe.

In our discussion, the assumption is that data is coming from an EXCEL spreadsheet. We will also assume a single sheet of data exists in a workbook.

The first step is to create or maintain the data in a spreadsheet described above. Before attempting an upload, **SAVE YOUR WORKBOOK. This is important.**

The next step is to save the single sheet in a "*comma delimited*" text file. This is accomplished by clicking on *FILE* and then choosing *SAVE AS*.

You will then be presented with the SAVE AS dialog box. Normally, the file name to be saved is the name of the WORKBOOK, and the type of save is "MICROSOFT EXCEL WORKBOOK." **Do not save with either of those settings**.

Instead…………..

Supply a different file name other than your workbook name (unless you want to overlay your workbook with the comma delimited text file) and change the file type to **COMMA DELIMITED TEXT**.

With the save characteristics established as described above, click *SAVE*. Remember that EXCEL now refers to this entire workbook as the name you chose to save it as. When you exit**, DO NOT SAVE AGAIN**.

Please note that a suffix CSV is appended to your filename. You must remember this when you upload or the PC file will not be found.

Now you must upload it to your mainframe file system. The command generator will process your data in either FIXED or VARIABLE record length formats and there is no preference. However, ISPF may not let you edit a record that is longer than 255 characters. Remember this when you define the file to receive your data.

When running the upload, make sure to specify ASCII and CRLF.

Switch into your mainframe editor (ISPF) and look at your file. It should consist of a single record for each row in your spreadsheet. Cells that did not contain anything are represented only by the comma separating them. Do not be alarmed if you find strings of commas (or the absence of commas) if the ending fields were omitted. The command generator will interpret this for you.

The Mainframe Upload Output File Characteristics

The uploaded output from the spreadsheet is used as input to ESRFCGEN.

This file may be either FIXED or VARIABLE length and its logical record length may be anything that is appropriate to contain the data.

You can populate this file a number of ways. In addition to the discussion above, you could code the data yourself, use another data manipulation product, or anything else that will produce the same end result.

Running the Command Generator

Output From the Upload

The following is a sample of what the data looks like after you upload the spreadsheet’s comma delimited text output using the sample spreadsheet shown previously.

```
0         1         2         3         4         5         6         7
1....5....0....5....0....5....0....5....0....5....0....5....0....5....0

OA,ID,,
DELETE,KAREN

OA,ID,NAME,ADDRESS1,PHONE
,TOM,TOM SMITH,”ROSEMONT,IL”,847 296-8010
,SUE,SUE SMITH,”ROSEMONT,IL”,847 296-8010
,GEORGIE,GEORGIE SMITH,”ROSEMONT,IL”,847 296-8010
,BARRY,BARRY SMITH,”ROSEMONT,IL”,847 296-8010
,DIANE,DIANE SMITH,”ROSEMONT,IL”,847 296-8010
,SARAH,SARAH SMITH,”ROSEMONT,IL”,847 296-8010

GA,ID,NAME,OWNER,PARTY1
,TEST.ACCT,ACCOUNTING TEST,TOM,DIANE
,TEST.OE,ORDER ENTRY TEST,TOM,GEORGIE
,TEST.SAM,SALES AND MKTG TEST,TOM,SARAH
,PAYROLL,PAYROLL SYSTEM,DIANE,
,ACCOUNT,ACCOUNTING SYSTEM,DIANE,
,ORDENT,ORGER ENTRY SYSTEM,BARRY,
,SHIPPING,SHIPPING SYSTEM,SUE,BARRY
,ACCTPAY,ACCOUNTS PAYABLE,DIANE,
,ACCTREC,ACCOUNTS RECEIVABLE,DIANE,BARRY
,SALES,SALES REPORTING,SARAH,
,MKTG,MARKETING SYSTEM,GEORGIE,
```

Again, it should be noted that any programmatic means to create the above input could be used instead of the spreadsheet approach.

The Job to Create E-SRF Command Input

Once the uploaded data is available, it must be converted into a format acceptable to the Command Processor for execution.

If you review the uploaded data, you will notice that only minimal information is present. There are no commands to DELETE, INSERT or CHANGE objects. Additionally, only the fields declared by the object headers exist.

The data is presented to an offline command generation program (ESRFCGEN). This program uses the uploaded data as input (whether fixed or variable length records) and creates the E-SRF Event System command statements suitable for DELETING, INSERTING and ALTERING the targeted owner and group headers. The following JCL may be modified and used to run the command generator.

```
//ESRF     JOB  ('Your Accounting Information'),'Yourname',
//              CLASS=A,MSGCLASS=X,MSGLEVEL=(0,0)
//*
//       RUN E-SRF COMMAND GENERATOR
//*
//CGEN     EXEC PGM=ESRFCGEN
//*
//STEPLIB  DD  DSN=ESRF.LOAD,DISP=SHR
//SYSPRINT DD  SYSOUT=*
//SYSPUNCH DD  DSN=ESRF.COMMAND.OUTPUT,DISP=SHR
//SYSIN    DD  DSN=ESRF.INPUT,DISP=SHR
```

Description of DD statements in above sample JCL:

JOBCARD Must be edited to your installation’s specifications.

CGEN E-SRF Command Generator program: ESRFCGEN.

STEPLIB The load library which contains the E-SRF system.

SYSPRINT A Control Report will be produced by ESRFCGEN and written to this file.

SYSPUNCH This is where your ESRF commands will be written to (LRECL=80).

SYSIN This is the input file that was uploaded from your PC.

This job is then run and the output is then ready to present to the E-SRF Event System’s Command Processor (ESRFCMD) as input (SYSIN). This is accomplished the same way as you would any other E-SRF command request.

What Will Be Generated For You

The following will be generated from the job shown above. For the sake of brevity, only the OWNER delete and the first OWNER definition is shown.

```
0         1         2         3         4         5         6
0....5....0....5....0....5....0....5....0....5....0....5....0

**************************************************
*                                                *
*    ESRFCGEN - GENERATED ESRF COMMANDS          *
*               START OF GENERATED COMMANDS      *
*                                                *
**************************************************

DELETE OWNER(KAREN)

DELETE OWNER(TOM)
INSERT OWNER(TOM)
REP    OWNER(TOM)                      -
        NAME(TOM SMITH)                -
       PHONE(847 296-8010)             -
    ADDRESS1(ROSEMONT, IL)             -
    ADDRESS2()                         -
    ADDRESS3()                         -
    ADDRESS4()                         -
    ADDRESS5()                         -
     ROUTING()                         -
    JESCLASS()                         -
     JESDEST()                         -
        DATE()                         -
        TIME()                         -
         LAF()                         -

* ESRFCGEN... WARNINGS WERE DETECTED, REVIEW 'SYSPRINT' OUTPUT

**************************************************
*                                                *
*    ESRFCGEN - GENERATED ESRF COMMANDS          *
*               END OF GENERATED COMMANDS        *
*                                                *
**************************************************
```

Please note that the entire field compliment will be produced despite the number of fields you specify. Omitted fields result in null field definitions, as seen in the output shown above.

These commands will render a DELETE of Owner Karen, a DELETE of Owner Tom, an INSERT of Owner Tom with no sub-fields presented, followed by a REP of Owner Tom that replaces the empty shell provided by the INSERT.

The sequence of these commands may seem illogical. They can cause errors during execution (specifically, a return code of 4 from the DELETE if the target object does not exist), but they afford the most flexibility.

You may edit the command data prior to presenting it to the Event Reporting System's Command Processor.

If there are processing errors, a warning is posted at the end of the commands as an alert to you. A sample of this warning is shown in the previous report sample.

You should review the SYSPRINT output from ESRFCGEN to insure no errors occurred.

How to Delete Unwanted Definitions

If you delete a definition row from your spreadsheet, the original definition on the Masterfile will still exist after you apply your commands. This is because we only process what was uploaded from the spreadsheet to the Masterfile. In this case the row was removed, so there is nothing to upload.

There are several approaches to deal with this situation:

- You can add the DELETE command to the command input (SYSIN file) before presenting it to the Command Processor.
- You can delete it yourself in a separate job.
- You can code the word DELETE in the first cell of the affected row on the spreadsheet. This will cause ESRFCGEN to generate a DELETE for the target object.
- You can code a DELETE for the ID of **!ALL** in the beginning of your spreadsheet, or add it to the beginning of the command input (SYSIN file) prior to presenting it to the Command Processor. This is a special ID that instructs the Command Processor to physically DELETE **all** objects within the current segment (except the default headers). This will give you an empty segment before applying the maintenance. **Remember, if you do this**, you may need one for OWNERS and GROUPS.

This page intentionally left blank

Chapter 13: Report Distribution

This chapter describes the Automatic Report Distribution Facility in E-SRF. This facility is currently available for use only with E-SRF Event Reports.

Distributed Reporting

This facility provides for an automated distribution of reports to data owners and any interested parties. The output is controlled by the OWNER's Masterfile definition. The OWNER Masterfile header contains information that can help route the reports through an already existing print report routing and control system if one exists.

The owner's report file may be defined in the JCL or dynamically allocated on demand. The dynamic allocation uses definitions found on the owner's Masterfile definition to complete the allocation.

If you choose to provide JCL, the DDNAME is the owner ID, followed by whatever E-SRF Report Output file DD parameters you desire.

Like any report output file, the owner's target report file is opened once on demand and stays open until the current Command Processor execution is completed. This means the same file is receiving all reports for the owner, for all RUN commands.

Distributed and Non-Distributed Event Reporting

E-SRF processes reports two different ways:

- **Non-distributed** reports contain all security events in a single report depending on the information requested. This is the way the Event System normally publishes reports.
- **Distributed** reports are created based on groupings. The report is republished for each data owner, containing only those events that relate to the owner. This is the type of reporting provided by the Automatic Report Distribution Facility.

Each style has its own advantages:

- **Non-distributed** reports present *all* reporting items and are appropriate for security administrators and others who want a global view of the entire environment.
- **Distributed** reports are designed to include only the information needed for a specific owner. This style enables mangers, data owners, decentralized security administrators, auditors, and others to view only the information they need.

Default Owner and Group

When the E-SRF Masterfile was initially created, the system automatically inserted a default group (DEFAULT) and a default owner (DEFAULT). These were placed on the Masterfile as "catch all" definitions in the event that all GROUP and OWNER definitions were not defined.

In report distribution, the first choice is to use the associated group name as the reporting group name, as specified in your grouping rules. If the associated group name is not properly defined on the Masterfile (*with a GA header*), the reporting group name is substituted with the default group name defined on the Masterfile, which is DEFAULT.

Commands That Involve Report Distribution

There are certain commands that affect the way Report Distribution works. Below is a brief explanation of those commands. Consult the *E-SRF Event System Command Reference* for more detailed information regarding these commands.

Option Grouping

Option Grouping determines whether or not grouping is involved in the subsequent RUN command specifications.

If you are running Report Distribution, the mode must be:

OPTION GROUPING(EXTERNAL)

EXTERNAL is the default.

- If distribution is attempted with GROUPING(NONE), the request will fail.
- If distribution is attempted with GROUPING(CLASS), the request will be honored but the distribution may not be particularly useful. Regardless, the option is available.
- If you disable portions of external grouping, the request will be honored but the distribution may or may not be useful, depending on what Masterfile segment is being reported on relative to what part of grouping is turned off.
- Portions of external grouping may be disabled by issuing one or more of the following grouping options:

  OPTION GROUPING(NOSOURCE) or OPTION GROUPING(NOUSER)

- By turning off part of grouping, you may save considerable time. Consider publishing a report based on the RESOURCE segment and its groupings. The grouping structures may include USERID grouping, which may take a considerable time to establish. By specifying OPTION GROUPING(NOUSER), user grouping is turned off. This saves time and resources, but the desired RESOURCE reporting is still received.

GROUP(…) RUN Selection Command

The use of a *user specified* group selection list specified in the RUN command is ignored when running report distribution. The command: GROUP(*group1,group2,…groupn*) is dynamically built by the Automatic Report Distribution facility based upon the current owner being reported on. All groups that target the actual owner or one of the interested parties are selected.

OWNER(…) RUN Selection Command

The use of the OWNER(*owner1,owner2,…,ownern*) RUN selection specification will limit report production to only the owners named in the list. This is usually specified during a rerun situation where reports are produced for a subset of the owners defined on the Masterfile. When owner selection is made, a grouplist is dynamically created for each owner.

This is different than its use in a non-distributed report, where the owner list creates a group list and all data is placed on the same report.

DISTRIBUTE(…) RUN Processing Option

This command initiates Report Distribution.

There are three choices to distribute reports: NONE, GROUP, or OWNER.

NONE: Do not distribute reports. The report produced is a single undistributed report produced in report sequence and placed on the DDNAME specified or on the REPORTS output. The report is produced under the full control of the report overlay and augmented by any additional RUN minor commands. **This is the default**:

DIST(NONE)

OWNER: Activate Report Distribution. Run a single report for each owner containing all groups targeted to the owner.

Each owner will receive a separate report consisting of **all** the report items associated to **all** groups, which the owner owns or is an *interested party* to. If fifteen groups targeted a particular owner, then the owner will receive a single report consisting of all reporting items related to the fifteen groups.

The single report will be bundled into the owner's report package, along with all other reports for that owner from other RUN commands.

To distribute using the OWNER option, the following distribution would be specified on the RUN command.

This is the recommended distribution method to use:

DIST(OWNER)

GROUP: Activate Report Distribution. Run a separate report for each GROUP and place the reports on the OWNER's output report file. ***This is not recommended***.

A Group associates resources via Resource Grouping Rules. There may be hundreds of Groups created through Grouping Rules.

Each owner will receive a separate report **for EACH GROUP** that the owner either owns or is an *interested party* to. If an owner is targeted for fifteen groups, the owner will receive fifteen reports, each one containing items that relate to the group being processed for the owner.

All reports will be bundled into the owner's report package, along with all other reports for that owner from other RUN commands.

This will consume a great deal of paper and CPU resources to create.

If you are running in this mode, it may be prudent to *turn on the Masterfile CACHE Facility*.

To distribute using the GROUP option, the following distribution would be specified on the RUN command:

DIST(GROUP)

Each RUN command is independent. You may submit a job that has multiple RUN commands. Each individual RUN command can specify Report Distribution independent of the others. This means you must restate DIST(OWNERS) or DIST(GROUP) on *each* RUN command that you want distributed. Remember, NONE is the default.

Masterfile CACHING for Report Distribution

The use of the Masterfile Caching Facility is recommended when using Automatic Report Distribution.

Report distribution causes the reporting engine to produce a specific report for each report output. If you are running DIST(OWNER) and there are 25 owners defined on your Masterfile, a single report will be produced for each of the 25 owners for each RUN statement that is set to do report distribution in your job.

This results in the generation of 25 separate reports: one for each owner, for each RUN command. Each report is unique and treated as if the data reported for that owner is the only data that exists on the system.

If you are running DIST(GROUP), the overhead is much greater. It would still be 25 owners, but a separate pass is required for each group that each owner is targeted to. If five groups targeted the average owner, this would result in the report being generated 125 times. In our example, this would be five times the number of reports than with DIST(OWNER).

The system was optimized for this type of processing, but the Masterfile is still VSAM. The use of the CACHE Facility is **strongly advised** unless you have a DASD environment that internally caches the disk drive's current data work set.

Because there are so many variables, EKC would not normally dictate the best approach. When you begin running distributed reports, it is recommended that you benchmark both approaches and compare the throughput, CPU requirements, central storage usage and actual execution wall times. Use this information to help you make the decision on whether or not to CACHE the Masterfile during Report Distribution.

For more information on Masterfile CACHING, please refer to both the *Event System Command Guide* and the *Event System Masterfile and Data Dictionary Reference*.

Report Distribution Example

The following is an example of how to set up Automatic Report Distribution. This example assumes that you have not done anything to either define the required control information on the Masterfile or write the required Grouping Rules to associate the reporting events to your desired groups.

This example is purely hypothetical and does not necessarily represent the actual definitions you will need. It does, however, represent the type of definitions required.

Step One:

Ten **groups** will be defined to the Masterfile GROUP segment for each of the six functional areas:

Functional Area	**Groups**
SECURITY	**DFTSEC, SECADM, SECDATA**
SYSTEMS	**DFTSYS, SYSPROC, SYSTEST**
MARKETING	**MKTG**
FINANCE	**FINANCE**
PERSONNEL	**HRIS**
General CICS	**CICSPROD**

Step Two:

Seven **owner**s will be defined to the Masterfile. One owner is an *Interested Party*. Owners and interested parties are all considered Owners by E-SRF.

TOM, GEORGIE, SARAH, BARRY, PAT, and MARK (Owners)
DIANE (Interested Party)

Step Three:

Resource Grouping Rules need to be verified if they already exist. If they do not exist, they must be written and reviewed.

Step Four:

Verify the grouping schemes for Automatic Report Distribution. This is accomplished by running report overlays designed to help insure that your resource grouping scheme (Grouping Rules) and your Masterfile definitions (owner and group headers) are appropriate for your report distribution requirements.

Step Five:

Run your distributed reports. Decide upon your distribution method and generate your reports.

Masterfile GROUP Setup (Step One):

Each ***grou**p* must be defined to the Masterfile using the E-SRF Command Processor.

Group	**E-SRF Command**
SYSTEMS area:	
DFTSYS	**INSERT GROUP(DFTSYS) DESC(SYS* DSNS) OWNER(TOM)**
SYSPROC	**INSERT GROUP(SYSPROC) DESC(SYS* DSNS) OWNER(TOM)**
SYSTEST	**INSERT GROUP(SYSTEST) DESC(SYS* DSNS) OWNER(TOM)**
SECURITY area:	

DFTSEC	**INSERT GROUP(DFTSEC) DESC(SECURITY DATASET) -**
	OWNER(GEORGIE) PARTY1(DIANE)

SECADM	**INSERT GROUP(SECADM) DESC(SECURITY DATASET) -**
	OWNER(GEORGIE) PARTY1(DIANE)

SECDATA	**INSERT GROUP(SECDATA) DESC(SECURITY DATASET) -**
	OWNER(GEORGIE) PARTY1(DIANE)

MARKETING area:	
MKTG	**INSERT GROUP(MKTG) DESC(CICS TXNS) OWNER(SARAH)**
FINANCE area:	
FINANCE	**INSERT GROUP(FINANCE) DESC(CICS TXNS) OWNER(BARRY)**
PERSONNEL area:	
HRIS	**INSERT GROUP(HRIS) DESC(CICS TXNS) OWNER(PAT)**
PRODUCTION CICS area:	
DFTPKC	**INSERT GROUP(DFTPKC) DESC(CICS TXNS) OWNER(MARK)**

Masterfile OWNER Setup (Step Two):

Each **Owner** must be defined to the Masterfile using the E-SRF Command Processor. The following table presents a sample of the syntax necessary to add owners and interested parties to the Masterfile:

Owner Name	Masterfile Command
TOM	**INSERT OWNER(TOM) - ADDRESS1(EKC) - ADDRESS2(ROSEMONT,IL) - NAME(TOM SMITH) - JESCLASS(H)**
GEORGIE	**INSERT OWNER(GEORGIE) - ADDRESS1(EKC) - ADDRESS2(ROSEMONT,IL) - NAME(GEORGIE TAYLOR) - JESCLASS(H)**
DIANE	**INSERT OWNER(DIANE) - ADDRESS1(EKC) - ADDRESS2(ROSEMONT,IL) - NAME(DIANE DURBIN) - JESCLASS(H)**
SARAH	**INSERT OWNER(SARAH) - ADDRESS1(EKC) - ADDRESS2(ROSEMONT,IL) - NAME(SARAH NELSON) - JESCLASS(H)**
BARRY	**INSERT OWNER(BARRY) - ADDRESS1(EKC) - ADDRESS2(ROSEMONT,IL) - NAME(BARRY SAINTIAGOSTINO)- JESCLASS(H)**
PAT	**INSERT OWNER(PAT) - ADDRESS1(EKC) - ADDRESS2(ROSEMONT,IL) - NAME(PAT ROGERS) - JESCLASS(H)**
MARK	**INSERT OWNER(MARK) - ADDRESS1(EKC) - ADDRESS2(ROSEMONT,IL) - NAME(MARK SHULTZ) - JESCLASS(H)**

Owners and Interested Parties

Owners and Interested Parties are treated exactly the same by the E-SRF Report Distribution Facility. Owners are considered the primary contact for Groups. Interested Parties are optional and are simply additional stations for report distribution. Owners or interested parties are “targets” for grouped reporting items.

All owners and interested parties must be added to the Masterfile. Every group must have an owner.

Resource Grouping Rules Setup (Step Three):

The following is a sample of how to set up ***Grouping Rules*** for each group:

Functional Area	Resource Grouping Facility Rule Example			Description
SYSTEMS	**$INDEX(SYS*)**	**DEFGROUP(DFTSYS)**	**NAME(SYSALL)**	**System**
	PROC-.-	**GROUP(SYSPROC)**		**Datasets**
	TEST-.-	**GROUP(SYSTEST)**		
SECURITY	**$INDEX(ACF*, SEC-)**	**DEFGROUP(DFTSEC)**	**NAME(SECURITY)**	**Security**
	PROD.-	**GROUP(SECADM)**		**Datasets**
	DATA.-	**GROUP(SECDATA)**		
CICS	**$CLASS(PKC)**	**DEFGROUP(DFTPKC)**	**NAME(PRODPKC)**	**CICS**
MARKETING	**MKT-**	**GROUP(MKTG)**		**Production**
FINANCE	**FIN-**	**GROUP(FINANCE)**		**Transactions**
PERSONNEL	**HRS-**	**GROUP(HRIS)**		

Verify Your Report Distribution Scheme (Step Four):

The analysis report overlays (ESRFOWNX, ESRFGRPX, ESRFGRPS, ESRFGRPT, ESRFGRPV) may be very useful in this phase of the project. The *E-SRF Event System Report Overlays Guide* contains information on how to run these reports. You may not need to run all of these at once, but you will run each one of them sometime during the distribution process. These reports evolved as tools from actual attempts to implement report distribution for E-SRF clients.

Review the reports and run ESRFGRPT in the distributed mode as a final test. This report will produce a simple list of Masterfile Objects without any detail. This will give you a close look at what data will go to which OWNERS and may be used as a final distribution checkout.

Sample Using: DIST(OWNER) (Step Five - A):

The following **RUN** command will be submitted to the E-SRF Command Processor to produce reports at the **Owner** level. This RUN execution will generate a separate report for each Owner defined to the Masterfile.

```
RUN REPORT(ESRFLIST) PARM(RESOURCE)              -
     TITLE(DATASET & RESOURCE ACTIVITY REPORT)    -

     DIST(OWNER)                        -

     FIELDS(                            -
          RESOURCE                      -
          GROUP                         -
          RC.SYSTEM                     -
          RC.DATE                       -
          RC.USERID                     -
          RC.ACCESS                     -
          RC.ACTION                     -
               )                        -

     WHEN(RC.DATE EQ *-1)
```

REPORT DISTRIBUTION

Eight reports will be produced from the above RUN command. One report is generated for each owner plus one report for the DEFAULT group's owner if applicable.

Distribution by Owner results in one or more Groups combined into one report.

For example:

Information relating to **Groups**:

DFTSYS
SYSPRO
SYSTEST

Will produce: A **single** report

That is distributed to **Owner** Tom.

This table presents the report distribution results for the above report request:

Report Contents: **Groups in Each Report**	**Reports Produced**	**Owner Destination**	**Interested Parties Destination**
DFTSYS **SYSPROC** **SYSTEST**	1	**Tom's Station**	
DFTSEC **SECADM** **SECDAT**	1	**Georgie's Station**	**Diane's Station**
MKTG	1	**Sarah's Station**	
FINANCE	1	**Barry's Station**	
HRIS	1	**Pat's Station**	
DFTPKC	1	**Mark's Station**	
Non-grouped resources via the external Resource Grouping Facility will be deposited in a single report.	1	**Default's Station**	

The seven owners defined in the Masterfile each receive one report. The Default owner will also receive a separate report for all "left over" resources that have not been grouped.

Sample Using: DIST(GROUP) (Step Five - B):

The following **RUN** command will be submitted to the E-SRF Command Processor to produce reports at the **Group** level. This RUN execution will generate a minimum of one report for each Group defined to the Masterfile.

```
RUN  REPORT(ESRFLIST) PARM(RESOURCE)              -
     TITLE(DATASET & RESOURCE ACTIVITY REPORT)    -

     DIST(GROUP)              -

     WHEN(RC.DATE EQ *-1)     -

     FIELDS(                  -
          RESOURCE            -
          GROUP               -
          RC.DOMAIN           -
          RC.DATE             -
          RC.USERID           -
          RC.ACCESS           -
          RC.ACTION           -
               )
```

Ten reports will be produced from the above RUN command. There will be one report for each Group, plus one for the DEFAULT Group if applicable.

DIST(GROUP) option with Report Distribution produces the following:

Owners to Receive Report	Interested Parties to Receive Report	Groups in Each Report	Total Reports Produced
TOM		DFTSYS SYSPROC SYSTEST	3
GEORGIE	DIANE	DFTSEC SECADM SECDATA	6
SARAH		MKTG	1
BARRY		FINANCE	1
PAT		HRIS	1
MARK		DFTPKC	1
DEFAULT		DEFAULT	1

The ten groups defined in the Masterfile generated a total of fourteen reports, including the default group.

This page intentionally left blank

Chapter 14: Administrative Procedures

There are many required and/or optional procedures to administer in E-SRF. Some procedures require a one-time-only set up, while others are ongoing. This chapter presents an encapsulation of those procedures. This is by no means a complete list. It does, however, identify the important procedures needed to ensure the successful administration of the E-SRF Event System Facility.

Brief explanations of the procedures along with references to other publications with more specific details are provided.

Configuring IMAGES

This is a required procedure. Configuring an Image is initially an implementation consideration. As time goes on, additional Images may be added and must be CONFIGURED. A Security Image is the largest Masterfile *grouping* in the Event Reporting system. An IMAGE refers to the users and resources that share a common set of security databases. Chapters Four and Five of this publication explain Images and how to CONFIGURE them. For additional information, please consult the *E-SRF Event Reporting Facility Command Reference*.

Assigning DOMAINS

This is a required procedure. A DOMAIN is a logical MVS system that is IPL'ed and maintained. Another name for this is an LPAR (Logical PARtition). You must provide the necessary ASSIGN statement to associate the DOMAINS to the IMAGE. Again, this may be both an implementation procedure and a modification procedure. Chapters Four and Five of this publication explain not only what a Domain is but how to Assign Domains to Images. The *E-SRF Event Reporting Facility Command Reference* can also assist you.

Defining and Maintaining the MASTERFILE

This is a required procedure. The Masterfile is ordinarily defined once at the time of implementation. An in-depth discussion of the Masterfile and its components is contained in the *E-SRF Event Reporting Facility Masterfile and Data Dictionary Reference*. Information regarding the creation of the Masterfile is also included in Chapters Four and Five of this publication. Once created, the Masterfile does require some maintenance such as fine-tuning record retention, rebuilding, reorganizing, and backup. In addition to the above referenced sources, Chapter Six in this publication discusses Masterfile maintenance considerations.

Gathering and SYNCHRONIZING RSS Data

This is a required procedure. The E-SRF Event System requires information about the Resident Security System (RSS) that it will report on. The SYNCHRONIZE command takes a copy of the userids or logonids present in the RSS and applies them to the E-SRF Masterfile so information in the userids or logonids can be used for reporting. Please refer to Chapters Four and Five of this publication for an explanation and the *E-SRF Event Report Facility Command Reference* for the syntax required to perform this function.

RESOURCE Grouping

This is an OPTIONAL procedure. Grouping is an integral part of E-SRF processing. It is used for associating datasets, non-datasets, and user resources into meaningful collections; selecting records to update, selecting information to report on, and automatically distributing reports. Grouping may affect updating the Masterfile as well as Event Report processing.

The association of a group name with resources is *not* stored on the E-SRF Masterfile. The associations are dynamically determined when requested because of update or report processing. When applying security event data during a Masterfile update, it is possible to exclude information based on grouping associations.

The purpose of grouping users and resources is to make it easier to create useable reports containing only information that a particular department head or owner wants to see. Grouping Rules can provide selection criteria and additional information in reports.

Chapter Eight of this publication is devoted to Event Reporting Grouping. Another publication, the E-SRF *Resource Grouping Facility Guide* explains grouping and should also be referenced to fully understand this concept and how to use it. In addition, the *E-SRF Event Reporting Facility Command Reference* will guide you through the syntax of the definitions.

GROUP and OWNER Header Maintenance

This is an OPTIONAL procedure. This function need only be performed if the Report Distribution Facility is used. For more information on this procedure, please refer to Chapter Nine of this publication and the *E-SRF Event Reporting Facility Command Reference*.

Controlling the Event System

This is both a required procedure and an OPTIONAL procedure. There are a myriad of commands that control the E-SRF Event System and how it presents data. Some are defaults, others are optional, and still others are required to define to E-SRF what processing options to use. Among these are commands to control date and time, CACHE control, grouping options, resource masking, and others. The best sources for this information are in Chapter Six of this guide and the *E-SRF Event Reporting Facility Command Reference*.

Automatic Report Distribution

This is an OPTIONAL procedure. The Report Distribution Facility is a powerful function that allows the Event System to produce and distribute individual reports for each Owner and/or interested party defined to the Event system.

The only requirements of report distribution are setting up your OWNER and GROUP definitions on the Masterfile and providing the necessary grouping of your resources.

Information available in Chapters Seven and Nine of this guide and the *E-SRF Event Reporting Facility Command Reference* will assist you in taking advantage of this function.

IMAGE Maintenance

This is an OPTIONAL procedure. You may have the need to either completely remove or rename a particular Masterfile Image. These functions may be accomplished using the following procedures.

When executing these functions, numerous Masterfile data record *update* operations are performed. It is recommended that Masterfile CACHING be turned on. The following illustrates this command:

```
CACHE ON
```

All information relating to these functions are displayed on the Control Report for your evaluation. These functions should be considered a major event and should be carefully planned and evaluated.

RENAMING an Existing IMAGE

By using the RENAME command, you may rename an Image that already exists to another IMAGE ID that does not already exist on the Masterfile. As of this release, there is no function to *merge* one Image into another.

```
REMOVE       IMAGE(NEWYORK)            NEWNAME(DALLAS)
```

When this command is executed, all User Segment data stored on objects for the current Image ID will be RENAMED to the target IMAGE ID. No event data is altered or removed.

All Domain assignments that target the current IMAGE ID, whether masked or specific, are adjusted to target the new IMAGE ID.

All IMAGE configuration data will be renamed to the new target IMAGE ID.

When the operation is finished, it will appear as if the current IMAGE never existed, and the new target IMAGE is associated with all Domains presently existing on the Masterfile.

The ESRFSHOW report should be run and the output examined. This is especially true for the IMAGE definitions and DOMAIN assignments. You may have to further adjust your assignments if you are relying on a generic scheme, or if you want the transactions belonging to Domains assigned to the Image to be related to another IMAGE.

REMOVING an Existing IMAGE

By using the REMOVE command, you may completely delete an IMAGE that presently exists on the Masterfile.

```
REMOVE       IMAGE(NEWYORK)
```

All User Segment data stored on objects for the current Image ID will be REMOVED from the Masterfile.

Event data for DOMAINS that target the removed IMAGE will be deleted from the RESOURCE Segment.

All Domain assignments that target the current IMAGE ID, whether masked or specific, are deleted from the Masterfile.

All IMAGE configuration data will be deleted from the Masterfile.

When the operation has been completed, it will appear as if the current IMAGE never existed. If an Update is presented for Domains that were removed, the events will be mapped into an IMAGE based on the remaining assignments. Normally, this would result in the transactions being rejected, unless you have a domain assignment targeting another image that generically or specifically matches the Domain ID.

ADMINISTRATIVE Procedures

The ESRFSHOW report should be run and the output examined. This is especially true for the IMAGE definitions and DOMAIN assignments. You may have to further adjust your assignments if you were relying on a generic scheme or want the transactions belonging to the removed Domains to be related to another IMAGE.

If you are running this function to "*clean up*" your Masterfile, you will have to replace the deleted assignments and re-configure your IMAGE. If this is what you are attempting to do, you may want to consider removing the Domains individually.

DOMAIN Maintenance

This is an OPTIONAL procedure. You may have the need to completely remove a particular Masterfile Domain. This function may be accomplished using the following procedures.

When executing this function, numerous Masterfile data record *update* operations are performed. It is recommended that Masterfile CACHING be turned on. The following illustrates this command:

```
CACHE ON
```

All information relating to this function is displayed on the Control Report for your evaluation. This function should be considered a "*big hammer*" and its use should be carefully planned and evaluated.

REMOVING an Existing Domain

By using the REMOVE command, you may completely delete a DOMAIN that presently exists on the Masterfile.

```
REMOVE     DOMAIN(TST1)
```

All event data stored on objects for the target Domain ID will be REMOVED from the Masterfile.

All Domain assignments that target the current IMAGE ID, that **specifically** match the target Domain, are deleted from the Masterfile. Assignments that generically match the target domain will REMAIN on the Masterfile. This is because you may have other Domains that either NOW or IN THE FUTURE will generically match the assignment. This is a serious consideration if you are using generic (*masked*) assignments.

Domain assignments that were not deleted due to masking are indicated on the Control Report.

When the operation has been completed, it will appear as if the current Domain never existed. If Update is presented for domains that were removed, the events will be mapped into IMAGES based on the remaining assignments. Normally, this would result in the transactions being rejected unless you have a generic domain assignment that matches the removed Domain, either generically or one added after the REMOVE that redirects the assignment.

The ESRFSHOW report should be run and the output examined. This is especially true for the DOMAIN assignments. You may have to further adjust your Domain assignments if you were relying on a generic scheme or want the transactions belonging to the removed Domains to be related to another IMAGE.

Masterfile Maintenance

This is a REQUIRED procedure. You must consider your E-SRF Event System Masterfile a critical application data file and provision it with suitable backup *and recovery* procedures. Reorganization procedures may also be required.

These procedures should be tested. If they are ever required, they may be performed with minimal impact on normal day-to-day operations.

Backing Up Your Masterfile

Your Masterfile normally gets updated with Resident Security System (RSS) data on a daily basis. Besides applying new transaction updates, older data are automatically purged. This data is lost.

Please… *BACK UP YOUR DATA*

Producing Reports From an OLD Masterfile

You may be required to construct a report consisting of data that has been previously purged off the current production E-SRF Masterfile. This may be accomplished by creating a temporary Masterfile from a previous backup, then running the desired report production job, using the RELATIVEDATE or SYSTEMDATE command(s) to alter the base operating dates.

It should be noted that if a new release of E-SRF Event Reporting were placed into production, the old Masterfile images would be automatically converted. The Event Reporting System is able to function with any previous Masterfile format. This is accomplished automatically for you via the automatic Masterfile Upgrade Facility.

Masterfile Reorganization

As mentioned before, the Masterfile is contained on a VSAM KSDS Cluster. The amount of update activity on this cluster is very high. Normally you would have to monitor and adjust FREESPACE and reorganize the Cluster on a daily basis.

The bulk of update activity occurs during Resident Security System (RSS) update. When the Update Function is executed, the recommended approach is to activate the Masterfile CACHE facility. When the CACHE facility is used and more than ten percent of the objects were updated, the current High Used RBA (Relative Byte Address) is set to zero, and the file is sequentially rebuilt.

If the CACHE is used during your update, you should NEVER have to reorganize your Cluster.

If the CACHE is not used, you must reorganize your cluster and monitor it similar to other VSAM clusters with high update activity.

This page intentionally left blank

Chapter 15: Troubleshooting

This chapter is presented to assist you if you are having difficulties with the product. Several topics are discussed including:

- EKC Technical Support information.
- Reading the Control Report for problem determination.
- Event System Messages and Return Codes.
- Using the ESRFSHOW Report to answer "why?"
- Common JCL issues.

EKC Technical Support

If you are experiencing difficulties with the product, the first step is to determine what the problem is to the best of your ability. This can be accomplished by:

- Reviewing the Control Report of the Event Report that was run.
- Obtaining a thorough understanding of the problem by looking up the messages and referencing other related parts of the E-SRF documentation.
- To the best of your ability, determine if the error is user or product related.
- Try to identify "what changed?" if applicable.
- Removing any user written report overlays or other exits and resubmit the failed request.

You may however, required additional assistance. EKC provides technical support, which may help you with abnormal situations as well as assistance with the actual usage of the product.

To contact EKC Technical Support, please call:

(847) 296-8035

Support is provided twenty-four hours a day, seven days a week.

During regular business hours, Technical Support is available from 8:00 AM to 5:00 PM CST. Normally the response during the day is from immediate to one hour.

During off-hour support, technicians are either on the premise or must be located. EKC attempts to provide the fastest off-hour Technical Support possible. There may be a delay in locating a qualified technician for a specific product during off hours. If the problem is an emergency or is critical to your processing, please call any time. If it can wait, a faster response is provided during the normal Technical Support hours.

When calling Technical Support, please have materials relevant to the problem available. This will afford the technician a better understanding into what is needed to help you resolve the issue you may have. Be in a position to discuss and answer questions about your environment.

Event System Reporting

Control Report (SYSPRINT)

The E-SRF Command Processor, and other components of the Event System record their activities on the E-SRF Event System Control Report. This report is written to an output print file called SYSPRINT.

This report is the most valuable tool for finding out what actually occurred during a particular run.

It should be reviewed after each execution and MUST be reviewed if the E-SRF Command Processor's execution ended with anything other than a return code of ZERO.

Many issues can be resolved quickly by reviewing the Control Report. For example, the following are two classic situations that usually prompt a call to the EKC Technical Group:

My report distribution does not work. Most of the time the grouping rules "RULES" DD was not specified, or an OPTION GROUPING(NONE) was specified. Sometimes the DIST(OWNER) specification was omitted. All of this information is displayed on the control report.

Some of my users no longer get reported on, and the ones that do are incorrect. This can occur if you run out of VSAM space while the CACHE attempts to rebuild or upgrade your Masterfile VSAM cluster. The cluster ends up being truncated to the VSAM space available. Since the User Segment is at the end of the cluster, the users at the end along with their event data will no longer exist. The situation may go unnoticed. When the next update is run, more VSAM space is available either because the disk volume has more space to contain more of the Masterfile cluster, or enough unexpired events were purged off, or both. Some of the users may be reinstated because they had activity in the second update run. However, those users lost from the first run with no activity in the second run are still lost until some activity reinstates them.

At this point the Masterfile no longer has integrity and is not useful for reporting.

Both of these, and many other problems can be caught. If the Control Report is examined when the problem first occurs, the problem can be detected. Corrective action can be taken immediately with little or no disruption of the reporting flow in your organization.

The Event System is currently a batch system. As with any batch job, the errors can only be discovered by examining the output it produces. **Please, review your Control Report**.

Event System Messages

As you review the Control Report, you will notice messages are posted on it. Messages can also appear on report output. All messages are documented in the *Event System Messages and Codes* publication.

Everything presented on the Control Report is posted in message format. The message itself contains a message ID, followed by message text.

The message ID consists of the letter E (*Event Reporting*), followed by a three-character unique message number. A dash separates the message ID from the actual four-character Event System component that posted the message. For example, the message E050-CMD came from the Event System, the message number was 50, and it was issued by the CMD (Command Processor) component.

The message text contains text explaining the message. There may be data relative to the message imbedded in the text. For example the message E050-CMD NON-ZERO RETURN: 8, FUNCTION: VSAM READ was presented indicating there was an error reading a VSAM record from the Masterfile. The return code and function were inserted into the message text. In a case like this, there will be more information indicating what was occurring at the time, such as the VSAM object key, or why the VSAM read was being carried out.

Most messages are informative and indicate what is occurring within the E-SRF System. The example above is NOT of that nature. An error has occurred somewhere because this message should never be received. If the message remained unnoticed or not acted upon, a potential problem could escalate into a **real** problem. You may ultimately contact EKC Technical Support to assist you with this error.

To learn more about messages, please refer to the *Event System Message and Codes* publication.

Event System Codes

In the data processing world, the only good code is a ZERO return code.

If a *non*-ZERO code is detected, consult the Control Report to determine what is wrong so that it may be corrected in a timely fashion.

Please, refer to the *Event System Messages and Codes* publication to find out about the *non*-ZERO code(s) you may have received. Normally there are message(s) associated with the *non* ZERO code(s) that will describe what was being attempted when the problem occurred.

ESRFSHOW

The ESRFSHOW report publishes all system options you currently have in place on your Masterfile. This report can be used to answer the “why?” questions when things appear to be working improperly. For instance, on 5/1 you ran a report selecting Resource Chronological events for a particular resource class. The selection dates were 4/15 through 4/25. You received an empty report, but you are absolutely sure the events took place. What is wrong?

Review the ESRFSHOW report and look at the RETAIN specification for the Resource Chronological object. You find it is set for four days. This is why you received an empty report. The data was purged off the Masterfile on 4/26.

The ESRFSHOW report publishes a list of the last 256 updates you have run. This may be useful if you think data is missing. Perhaps an UPDATE was not run when it should have.

JCL - Items To Watch For

Once you have your Masterfile established, the JCL to run the Event System is very simple. Sometimes, however, changes are made to the JCL that were never intended to be permanent. Because this occurs, a list of some common JCL errors that have prompted Technical Support calls are indicated below:

> Review your JCL. Make sure it is pointed to the proper Masterfile. If you are using Grouping Rules, check to see that you are specifying the proper RULE OBJECT dataset in the RULES DD and not the RULE SOURCE PDS.
>
> Verify that all DD references you made in the Command Processor input are included in the JCL. If you do not, some DDs will be dynamically allocated for you. Normally this will be desirable, but sometimes it may not.
>
> Make sure the only program you specify in the EXEC PGM=*xxxxxxxx* JCL statement is ESRFCMD (the Event System Command Processor). If you specify a report overlay, your job will abend with a user ABEND code.
>
> If you are running an UPDATE, you must specify TIME=1440 and run the job in an execution class that will permit this TIME parameter.
>
> If you are running an execution that uses the Masterfile CACHE, make sure REGION=0M is specified and run the job in an execution class that will permit this REGION size.

CPU Time and Main Storage Usage

The UPDATE function runs more efficiently when the Masterfile CACHE is active. After the CACHE is built, the only I/O processing that occurs is reading the Resident Security System journal data. It may appear that the job is looping. When the CACHE is on, the entire contents of the Masterfile VSAM Cluster is loaded into main storage.

Remember, the data being processed is YOUR data. E-SRF will attempt to process whatever you give it. If it takes a long time, or uses a lot of main storage, it may be because of the volume of journal data it is being presented with.

There are procedures that may be performed to analyze the logged journal data and provide corrective measures to scale it down to something more reasonable. A call to EKC Technical Support to discuss this issue should yield proposed solutions to this problem.

EKC Product Development is constantly improving the processing routines used to manage the Masterfile. Each new release out performs the previous release. However, there may be a time when it is no longer possible to speed up the processing routines. When that time is realized, it will be up to you to take a closer look at how the journals are managed.

Storage DUMPS

As with any data processing application, you may encounter a situation where a storage dump may be produced.

If the CACHE is on, the dump will be too large to be useful.

Do not print storage dumps for EKC unless they are specifically requested. E-SRF storage dumps are "enormous" when running with the Masterfile CACHE option enabled. Attempting to update the Masterfile with the CACHE disabled could take hours if not days. If a storage dump becomes necessary, and you are not updating the Masterfile, run the request with the CACHE disabled.

Access Analysis Reporting

Section 0 of all Access Analysis Reports has information about report processing, the parameters specified, and the records used to create the report. If you are receiving an empty report, it may be because no records matched the specified parameters. Section 0 statistics will verify that is what occurred.

Chapter 16: Index