E-SRF

EKC Security
Event Reporting Facility

Release 2.1
Change Summary

E-SRF V2R1
EKC, Inc.

GENERAL AVAILABILITY

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 1, Release 6 LE00450 February, 2005

Printed in USA

Contents

E-SRF Event Reporting Release 2.1 Changes....................1

SCOPE....................1

E-SRF Event Reporting System....................2

PERFORMANCE IMPROVEMENTS....................2

Release 2.1 was primarily developed as a performance enhancement release. *2*

DOCUMENTATION 3

GENERAL 3

This document only contains changes made by enhancement level LE00450. ....*3*

Corrective service applied to this release....................*3*

LE00450 places the E-SRF Event Reporting System at Release 2.1 ....*3*

MASTERFILE CHANGES....................4

Resource Name length increased from 44 to 1000 characters. *4*

Resource Maintenance (RM), User Maintenance (UM) and Maintenance Chronological (MC) datanames are now consistent.....................*4*

Maintenance Chronological (MC) is now a Virtual Object *4*

Masterfile Level 2 CACHE storage. *4*

THE DEFAULT IS TO USE ONE OR MORE DATA ONLY ADDRESS SPACES....................*4*

E-SRF MASTERFILE LEVEL UPGRADED TO 02.01.01 5

Masterfile level has been upgraded due to the following changes: *5*

Masterfile data has been tokenized....................*5*

Tokenized data and what this means....................*5*

Masterfile performance objectives for the Update Function *5*

Masterfile performance objectives for Report Production....................*6*

Future Masterfile performance objectives.....................*6*

E-SRF EVENT REPORTING GROUPING ENHANCEMENTS 6

New approach to grouping the Masterfile....................*6*

"COLD" Group Structure Build:....................*7*

"HOT" Group Structure build:....................*7*

REPORT OUTPUT ENHANCEMENTS 7

Fixed or variable length report output files....................*7*

Report output can be a Partitioned Dataset (PDS)....................*7*

Report output can be separated by the IEBUPDTE "./ NAME=name" statement....................*7*

Report output can formatted in HTML format....................*8*

Report output can formatted in ASCII format....................*8*

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.

E-SRF Event Reporting Release 2.1 Changes

This guide identifies the major enhancements incorporated into E-SRF Version 2, Release 1. The intent is to describe the changes that were incorporated into the product and their possible impact on your E-SRF system currently in place. The enhancements are discussed assuming you have a basic understanding of E-SRF.

If you are new E-SRF customer, this information serves as a point of interest. New customers should utilize the *User Guide* for information on the use of this product.

For E-SRF Access Analysis, please refer to the documentation relating to those components.

These publications provide the best overall explanation of E-SRF Event Reporting functionality. The *User Guide* is organized to "bring together" E-SRF Event Reporting functionality and concepts and direct you to other manuals when information that is more detailed is required.

Release 2.1 is a critical release of this product, and contains enhancements that were requested by customers using the product, as well as planned enhancements needed to develop the product into an Enterprise-wide facility for reporting on journalized security events. This release is critical to future offerings of this product.

Please review this information to determine what, if any, impact this release may have on your operating environment. Review the product's documentation. If you have questions or concerns about anything mentioned in this document, additional questions or comments, please contact EKC Technical Support.

The E-SRF Event Reporting System, with the exception of the enhancements stated in this document is functionally identical to previous releases from a user perspective. Many other changes were made to accommodate new enhancements targeted for future releases. Some enhancements may not have any external appearance and therefore not mentioned in this publication.

Scope

Information provided in this document represents all changes that occurred from the *most recent* release 1.6 offering (LE00432) to the current product offering (release 2.1 at LE00450). If you are on a release prior to 1.6 at LE00432 and are interested in what changed before LE00432, please consult the change summaries published for all maintenance levels between the maintenance level you are currently using and this level.

At the time of this writing, LE00432 is an APAR only release and has not been provided as a separate release. The various APARS included LE00432 have been made available in LE00431.

Because Release 2.1 represents such a major milestone in this product's evolution, EKC plans to functionally stabilize E-SRF Event Reporting version 1 at release 6 levels. There will be offerings, such as LE00432 in the future, but only for problem resolution purposes. Version 1 will remain at release 6.

We encourage all of our customers to migrate to Release 2.1 (or higher when available).

E-SRF Event Reporting System

This offering will place your E-SRF Event Reporting System at release 2.1 at maintenance level LE00450.

Performance Improvements

An ongoing effort is being made to improve the operational performance of the E-SRF Event Reporting System.

As stated in previous releases, the main objective of E-SRF performance improvements are to decrease elapsed time (*wall time*), reduce CPU consumption, and reduce the virtual storage and disk space required to contain the E-SRF Masterfile VSAM cluster.

One of our primary development goals is to provide performance enhancements with each new release of this product.

This release contains performance enhancements that have been planned for the past several releases. Release 2.1 represents the most performance enhancements that were ever incorporated in a single release. Most of what was done to improve the performance of this product will have no impact on what your reports look like. Many functions, such as the Update Function will simply run faster and use less storage.

In a typical operating environment, normally, you can expect to see the following:

A fifty percent (or better) reduction of the amount of CPU resources required for the execution of the Update Function.

Report production CPU time was not directly enhanced in this release.

From ten to fifty percent reduction of the VSAM space required to contain the Masterfile. There wil be a reduction of space. The amount of reduction depends on how your Masterfile was configured, length of your resource names, number of userids related to resources and the way your Masterfile compressed.

From ten to fifty percent (or better) reduction of the amount of Virtual Storage required for general program execution. As with the previous point, much of this is dependent on your data.

Resources required for grouping the Masterfile were greatly reduced.

Please note that additional functionality was added to the product. Additional data is also being stored on the Masterfile. All additional resource requirements for additional functionality have been considered in the above estimates.

Overall, Release 2.1 represents the most aggressive effort to improve the performance of this product.

Release 2.1 was primarily developed as a performance enhancement release.

Release 2.1 was developed as a performance enhancement release. The design of this release started during Release 1.5. Enhancements contained in all releases after 5.1 were positioning the product for this release.

Additional performance enhancements are scheduled for subsequent releases during the life of Version 2.

Please note there are additional enhancements in this release. Great effort was expended to not change the appearance of reports. If such is noted, it may be due to enhancements required by the security systems that produce the loggings, or problem resolution changes made to the product.

Documentation

The Event Reporting System documentation was upgraded and should be considered a total replacement. In addition to this summary, please review the entire product documentation for additional information as needed.

General

Provide corrective service maintenance throughout the entire E-SRF Event Reporting System.

Some of the notable changes are contained on this document.

This document only contains changes made by enhancement level LE00450.

Other enhancements that were made in previous offerings of this product are contained in their respective documents. During the life of release 1.5 and 1.6, the product was extensively enhanced during maintenance releases.

If you are migrating from anything older than LE00431, and have an interest in what changed at each respective level, *please refer to documents distributed with those offerings*.

Corrective service applied to this release

All corrective service that was made available to customers in any release of E-SRF Event Reporting up to February 13, 2005 has been incorporated into this release.

Release 1.6 will be continue to be supported for one year. After this time, support will be provided on a best effort bases. Any problem reported and resolved in a previous release will be retro fitted to release 2.1 and made available as demand APARS or in subsequent release 2.1 product level sets.

Any maintenance applied to release 2.1 and not reported in previous releases may not be retro fitted into previous releases.

LE00450 places the E-SRF Event Reporting System at Release 2.1

This release is downward compatible with prior releases. This means if you have an existing Masterfile created by any older release it will be upgraded to the current formatting requirements.

Caution: *LE004500 alters the format of the Masterfile* and upgrades the Masterfile's release level. A conversion will occur (*provided the UPGRADE parameter was specified*) on any Masterfile that is older than the current Masterfile release (02.01.01).

Masterfile Changes

Resource Name length increased from 44 to 1000 characters.

The maximum length of a resource name has been extended from 44 characters to 1000 characters in length.

E-SRFLIST report overlay formatting of resource names has been enhanced to use as much space as is available to format resource names. If resource names are found to be too long, the lane is formatted left justified with right character truncation. A footnote is appended to names that were truncated. At the end of the report, all truncated names are formatted along side of the footnote that references it.

Other report overlays were changed in similar manners with or without the footnoting.

Resource Maintenance (RM), User Maintenance (UM) and Maintenance Chronological (MC) datanames are now consistent.

All three of these objects represent maintenance performed on security definition data. These objects now all contain the same data elements. Please note that not all data will be provided for all maintenance events. The level of data is the same as it was in release 1.6. What was changed is the organization.

The MC object now has more data items. It is the same as what is contained in the RM and UM objects.

You can get a full maintenance report by simply replacing the two reports you may have reporting on the RM and UM objects with a single report that reports on the same datanames with the MC identifier.

You can logically consider all three of these objects to be the same. The only difference is the segment scope. RM's scope is the Resource Segment, UM's scope is the User Segment while MC's scope is both segments.

Maintenance Chronological (MC) is now a Virtual Object

The MC object no longer physically exists on the Masterfile. The Data Dictionary still carries the datanames for this object, but in reality, they reference the Resource Maintenance (RM) and User Maintenance (UM) objects in a seamless fashion.

If you run the exact same MC report you have in the past, you will get more lines of data. To run the exact same MC report showing the exact same data, you will have to add the following to your RUN command:

```
WHEN(MC.DATANAME EQ ' ')     -
  OR(MC.DATANAME MATCH +-)  -
  OR(MC.DATANAME EQ ACCESS) -
  OR(MC.DATANAME EQ ID)     -
```

This will filter out the added data to this Virtual Object and make the output look like what the old MC object would have produced. You can do the same with RM or UM if desired.

Masterfile Level 2 CACHE storage.

Introduced in Release 6, you have a choice of where to allow E-SRF to establish and maintain its Masterfile Level 2 CACHE. Prior to that release, the CACHE (when specified) resided on the same address space that hosts the E-SRF execution.

The default is to place the Level 2 cache in a data-only address space. This release extends the process to use multiple data-only address spaces as required.

THE DEFAULT IS TO USE ONE OR MORE DATA ONLY ADDRESS SPACES.

E-SRF does not use more storage using this scheme. E-SRF simply divides the storage requirement between the job's executing address space and one or more data-only address spaces.

Please note that E-SRF runs as a normal "key 8" application program. This means it is not APF authorized. APF authorized programs may normally utilize any size data only address spaces without any conditioning. Normal application programs may have to filter through installation limits on address space memory allocations. Please insure that the execution environment is suitable for maximum data only address space memory allocations. This may require a discussion with your system's programmer.

E-SRF Masterfile level upgraded to 02.01.01

Masterfile level has been upgraded due to the following changes:

Changes to data formats contained on the Masterfile.

Additional data now carried on the Masterfile.

Additions and changes to the data dictionary.

Maintenance Chronological object is now a virtual object.

Resource name size limit has been increased from 44 characters to 1000 characters.

Masterfile data has been tokenized

Many of the E-SRF data structures contained on the Masterfile are now tokenized in Release 2.1.

Tokenized data and what this means

Tokenizing E-SRF Event data is a major step forward.

The majority of storage and processor requirements required by the E-SRF Event Reporting System are due to the extensive amount of data that must be applied to and maintained on the E-SRF Masterfile, and subsequently reported on. To speed up E-SRF, all of this must be reduced. At the same time, customers have required the resource name length to be extended from the current 44 characters to something considerably longer. The current industry resource name specification is 1000 characters. Other data has also been requested to be maintained on the Masterfile.

Data Compression was used (*and still is*) as a means of reducing the physical size of the data. The E-SRF Masterfile is so large that in most installations that you could not simply decompress the data, perform processing against the data, and recompress the data back into the file. Additionally, during the Update Function, the data structures are constantly being lengthened as more events are associated to objects and stored. The compressed data is cached in the Level 2 cache in its compressed form. When a particular object is required for processing, it is decompressed into the Level 1 cache for processing. A decompressed E-SRF object could be up to sixteen million characters long. This means the Level 1 cache is only able to maintain a set number of objects (*currently fifteen objects may be maintained in the level 1 cache*). When the Level 1 cache is full, the oldest referenced object is compressed (*if required*) and stored back in the Level 2 cache. This is expensive. Not compressing the data would mean most people could not run this product with the amount of security logging that exists and the current hardware available.

To solve this problem, you must make the data smaller and at the same time reduce the need to compress the data.

When you tokenize data, what you really are doing is creating a unique label (token) consisting of 3 to 4 bytes that uniquely represents a long character string. This means that if you have a 1000 character resource name, you can represent it with the token that may be three or four characters long.

Token processing can be very complex, but it totally removes previous limits and the overall performance is greatly enhanced.

In this release, we established a token processing model and used it for such things as resource names, UID data, and profile names.

Token processing requires data structures to tokenize and resolve existing tokens. These structures consume resources. The performance yield is the difference between executing a function under the previous release compared to executing the same function under Release 2.1. All overhead is accounted for when performance information is referenced.

Masterfile performance objectives for the Update Function

Release 2.1 development objectives for the Update Function was to reduce storage and CPU usage by fifty percent or better. This objective was met and in most cases exceeded in this release.

Masterfile performance objectives for Report Production

Release 2.1 development objectives for actual report production was to try to keep it the same as it was in release 1.6. Report production CPU performance appears to be stable from 1.6 to 2.1, with a major storage usage improvement. Event Reporting was always optimized for report production. The more something is optimized, the more difficult it becomes improve on it.

Any additional processing introduced to report production could be measurable and noticed. It takes additional processor resources to process data that has been tokenized. This additional processing appears to be offset by what is saved by the smaller Masterfile and other storage worksets required by this release.

Future Masterfile performance objectives

Performance enhancements for this product will not end with release 2.1. This release must be made available to address current performance concerns. What is provided in this release will go a long way toward helping our clients manage the huge amounts of security logging data and produce meaningful reports for management.

There are still more ways this product will be enhanced to improve its overall operational performance and this effort will continue.

We will continue to address more of these issues and provide future offerings that will build on our tokenization model, zOS architecture and other design elements that are planned for future releases of Event Reporting.

Other enhancements to the core value of the product will continue to be enhanced. Included in this but not limited are more reports and Update Functions for more security system Platforms.

E-SRF Event Reporting Grouping Enhancements

E-SRF Event Reporting optionally makes extensive use of the EKC Integrated Grouping Facility. Grouping is used for report data selection, object identification, (for example, displaying the name in plain English of a four character CICS Transaction ID on a report), and the creation of groups that are associated to owners for Report Distribution.

The Masterfile has always been grouped “on demand”, which allows you to change your grouping schemes and use them immediately. This is a wonderful feature, except grouping appears to be a very expensive process.

To address the performance issues grouping objects on the Masterfile, the grouping process was redesigned providing two facilities to group your Masterfile:

New approach to grouping the Masterfile

Three primary objectives had to be met relating to the Event Reporting grouping process. The first was the ability to group very long resource names. Second was to reduce the overall storage requirement. Third was to reduce the time and overhead of grouping Masterfile entities.

To meet these three objectives, an entirely new grouping architecture, based on the current EKC External Grouping facility was implemented in Release 2.1. This new architecture met all of the above goals.

The new architecture addressed the storage usage and long resource name issue by tokenizing the entire grouping structure. This means the Event Reporting can now maintain grouping structures for resources up to 1000 characters long. Please note, in Release 2.1, the EKC External Grouping Facility still maintains a 256- character limit on the maximum length of an entity to be grouped. This means only the first 256 characters will be actually presented for grouping. Currently, this is the only restriction that remains, and will be addressed in a future release. For most installations, this temporary restriction will not impose any problems.

The last issue, (the time it takes to group the entities on the Masterfile) was also addressed. Due to the nature of grouping and its use of grouping Rules, this process will always be a time consuming operation.

The new grouping architecture uses a dual approach to grouping the entities contained on the Masterfile. This completely addressed the requirement. Unless you changed your grouping rules, this architecture will make the time, it takes to establish the grouping structures appear to be almost non-existent.

When Event Reporting detects the need for grouping, one of the two grouping structure startup facilities will be used:

"COLD" Group Structure Build:

Initially, group structures are built "COLD". A COLD build means every entity is grouped and its resulting group name is made available for processing.

Grouping structures are normally built on demand whenever the potential for grouping is detected. It has been seen that a report that takes ten seconds to run can actually take fifteen minutes due to the time it takes to group all the entities contained in the Masterfile.

The COLD start process is similar to the way the Masterfile was grouped in previous releases. Each Masterfile entity is assigned a group name based on the EKC Grouping Rules you provide from your Rule Object file.

This type of process runs every time grouping is involved in releases prior to release 2.1

"HOT" Group Structure build:

In release 2.1, grouping structures can now be built "HOT".

The entire grouping structure is tokenized and stored on the Masterfile for future use during a COLD start.

Each time E-SRF is started with a grouping HOT start, the grouping structure will be built from the existing data built from a previous COLD start.

A COLD start will only be run if it is determined that things relating to grouping have changed and the current HOT start would not properly reflect what the group structures should look like. These changes include grouping rule changes, the name of the rule object file being presented for processing, and other items relating to grouping that E-SRF now keeps track of.

A group structure build that takes fifteen minutes to do in a COLD start may now be established in seconds HOT. Previous releases always COLD started grouping structures and this was the reason OPTION GROUPING(NONE) was recommended on executions that did not require grouping structures to be present.

The decision use of either a COLD or HOT start is normally an automatic function and does not require any intervention.

REPORT output enhancements

E-SRF Event Reporting will direct report output to the DDNAME 'REPORTS", unless overridden by a particular report's RUN command: DDNAME(*ddname*). This is still true in release 2.1, except you now have many additional options relating to the report output file.

Fixed or variable length report output files

Your output file can now be FIXED or VARIABLE length.

Report output can be a Partitioned Dataset (PDS)

The report output file DDNAME(*ddname*) from a RUN command can be a Partitioned dataset (PDS). Each RUN command can create its report output as a member of a particular PDS as specified by the DDNAME RUN parameter. The member name may either be declared in the particular report RUN command using the MEMBER(*membername*) parameter or have E-SRF assign a sequential member name relative to the particular PDS being written to.

The output may be fixed or variable length.

Report output can be separated by the IEBUPDTE "./ NAME=*name*" statement

You can create output in a standard physical sequential file with the IEBUPDTE control statement './ name=xxxxxxxx' added to the beginning of each report RUN that directs its output to the particular file. This may be useful if you have a need to divide a report file for post processing. Keep in mind that the IEBUPDTE cannot process this data unless the data is 80 character fixed data.

The output may be fixed or variable length.

Report output can formatted in HTML format

Using any format described above, you can also format your data in HTML (Hyper Text Markup Language). The output may then be viewed by any WEB browser.

This output must be written as a variable length file. The output may be directed to a PDS or IEBUPDTE file type.

Report output can formatted in ASCII format

You can create report output in ASCII format, complete with basic print control characters. The output may be downloaded (binary transfer) and directed to a PCL printer.

This is a special type of file that may be downloaded to a machine capable of processing ASCII data.

This output may also be processed using any text editor or word processing program.

This output must be written as a variable length file.