CICS/ESA XRF Guide

Version 3 Release 3

Document Number SC33-0661-02

Program Number
5685-083

Abstract

This book provides information to help systems programmers and designers to plan for and implement a CICS system running with XRF.

There is a high-level overview of CICS with XRF, and a description of the types of failures handled by XRF. There is a detailed description of the way XRF works. Some XRF configurations are suggested, as well as ways of managing them. There is guidance about the effect of XRF on different types of terminal and some network information is included. The way to define a CICS with XRF system is examined. There is information about the way CICS with XRF interacts with other products. There are also a short checklist and sample implementations.

Where appropriate, there are references to other CICS books which complement the XRF information in this book.

EDITION Notice

Third edition (March 1992)

This edition applies to Version 3 Release 3 of the IBM licensed program Customer Information Control System/ Enterprise Systems Architecture (CICS/ESA), program number 5685-083, and to all subsequent versions, releases, and modifications until otherwise indicated in new editions. Consult the latest edition of the applicable IBM system bibliography for current information on this product.

This book is based on the *XRF Guide* for CICS/ESA 3.2.1, SC33-0661-01. Changes from that edition are marked by vertical lines to the left of the changes.

This softcopy version is based on the printed version of the *XRF Guide* for CICS/ESA 3.3, SC33-0661-02, and includes the changes indicated in the printed version by vertical revision bars. Formatting amendments have been made to make this information more suitable for softcopy.

Additional changes made to this softcopy version of the manual since the hardcopy manual was published are indicated by the hash (#) symbol in the left-hand margin.

The CICS/ESA 3.2.1 edition remains applicable and current for users of CICS/ESA 3.2.1, and can now be ordered using the temporary order number ST00-5112-00.

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

Reader's comments on this publication should be addressed to:

International Business Machines Corporation, Attn: Dept ACV-H
1001 Wt Harris Blvd, Charlotte, NC 28257-0001, USA

or to:

IBM United Kingdom Laboratories Limited, Information Development,
Mail Point 095, Hursley Park, Winchester, Hampshire, England, SO21 2JN.

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

CONTENTS Table of Contents

[Summarize]

Figures

Tables

FRONT_1 Notices

copy or partial copy.

- Transfer the original unaltered copy of the documentation when you transfer the related IBM product (which may be either machines you own, or programs, if the program's license terms permit a transfer). You must, at the same time, destroy all other copies of the documentation.

You are responsible for payment of any taxes, including personal property taxes, resulting from this authorization.

THERE ARE NO WARRANTIES, EXPRESS OR IMPLIED, INCLUDING THE WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Some jurisdictions do not allow the exclusion of implied warranties, so the above exclusion may not apply to you.

Your failure to comply with the terms above terminates this authorization. Upon termination, you must destroy your machine readable documentation.

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates.

Any reference to an IBM licensed program or other IBM product in this publication is not intended to state or imply that only IBM's program or other product may be used. Any functionally equivalent program that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, is the user's responsibility.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Commercial Relations, IBM Corporation, Purchase, NY 10577.

Trademarks and service marks

The following terms, denoted by an asterisk(*), used in this publication, are trademarks or service marks of IBM Corporation in the United States or other countries:

ACF/VTAM, BookManager, CICS/ESA, CICS/MVS, CICS/VSE, DATABASE 2, DB2, ESA, IBM, IMS/ESA, MVS, MVS/ESA, NetView, Processor Resource/Systems Manager, PR/SM, SQL, VM/ESA, VM/XA, VTAM, 3090.

Preface

What this book is about

This book is intended to help you to understand the extended recovery facility (XRF) function of CICS/ESA (*) 3.3. It contains guidance about planning, setting up, and running a CICS with XRF configuration.

If you need to know where programming interface information is described, or about the definitions of the different types of information in the CICS library, you should read the CICS Library Guide.

Who this book is for

This book is for system designers and system programmers.

What you need to know to understand this book

You need a good understanding of CICS, and of the level of system availability that your users need. For planning and environmental information, and early guidance about XRF, see the *CICS/ESA Release Guide* and the *CICS/ESA Facilities and Planning Guide*.

How to use this book

Topics 1.x through 3.x introduce the XRF concept and explain how CICS with XRF works. Topic 4.x suggests possible configurations. Topic 5.x and 6.x give more detailed guidance to help you set up XRF. Topic 7.x relates XRF to other products.

The appendixes provide a checklist of what you do to create an XRF complex, and also a sample implementation with suitable definitions.

Additional task-specific information about XRF is given in other CICS books, and this book provides references to those books.

Notes on terminology

There is a glossary of terms of particular relevance to XRF in topic GLOSSARY. There is a general glossary of CICS/ESA terms in the *CICS/ESA Glossary*.

(*) IBM Trademark. For a list of trademarks, see topic FRONT_1.

Subtopics:

| PREFACE.1 Softcopy links

| This book is linked to the *CICS/ESA Glossary* and to the *CICS/ESA Messages and Codes* manual.

| If you are using BookManager (*) READ/MVS or BookManager READ/VM, you can view the definitions of terms and the messages directly from this book by selecting a term with your cursor and pressing the ENTER key.

PREFACE.2 CICS/ESA 3.3 library

Evaluation and planning

Brochure	GC33-0653
CICS General Information	GC33-0155
Facilities and Planning Guide	SC33-0654
Release Guide	GC33-0792
Front End Programming Interface Feature General Information	GC33-0803

General

CICS Library Guide	GC33-0356
Master Index	SC33-0671
User's Handbook	SX33-6076
Messages and Codes	SC33-0672
Processing Overview	SC33-0673

Administration

Installation Guide	SC33-0663
System Definition Guide	SC33-0664
Customization Guide	SC33-0665
Resource Definition (Online)	SC33-0666
Resource Definition (Macro)	SC33-0667
Operations Guide	SC33-0668
CICS-Supplied Transactions	SC33-0669
System Programming Reference	SC33-0670

Programming

CICS Application Programming Primer (VS COBOL II)	SC33-0674
Application Programming Guide	SC33-0675
Application Programming Reference	SC33-0676
Sample Applications Guide	SC33-0731
Distributed Transaction Programming Guide	SC33-0783

Service

Problem Determination Guide	SC33-0678
Diagnosis Handbook	LX33-6077
Diagnosis Reference	LY33-6072
Data Areas	LY33-6073

Special topics

Intercommunication Guide	SC33-0657
Recovery and Restart Guide	SC33-0658
Performance Guide	SC33-0659
CICS-IMS Database Control Guide	SC33-0660

XRF Guide	SC33-0661
CICS-RACF Security Guide	SC33-0749
CICS Communicating with CICS OS/2	SC33-0736
Communicating from CICS/ESA and CICS/VSE	SC33-0825
Front End Programming Interface Feature User's Guide	SC33-0804
IBM 3270 Data Stream Device Guide	SC33-0232

Related books

CICS OS/2 System and Application Guide	SC33-0616
CICS Application Migration Aid Guide	SC33-0768
Transaction Processing: Concepts and Products	GC33-0754

PREFACE.3 Books from related libraries

MVS/ESA

MVS/ESA General Information for MVS/SP Version 3, GC28-1359.
| *MVS/ESA General Information for MVS/SP Version 4*, GC28-1600.
MVS/ESA SP Version 4 Release 1 Implementation Guide, GG24-3628.
MVS/ESA Initialization and Tuning Reference, GC28-1635.
MVS/ESA Planning: Sysplex Management Guide, GC28-1620.
| *MVS/ESA JES2 Initialization and Tuning Reference*, SC23-0083.
| *MVS/ESA JES2 Initialization and Tuning Guide*, SC23-0082.
| *MVS/ESA JES3 Initialization and Tuning*, SC23-0059.

IMS/ESA

IMS/ESA General Information Manual, GC26-4275.
IMS/ESA Release Planning Guide, GC26-4386.

IMS/VS

IMS/VS Version 2 General Information Manual, GC26-4180.
IMS/VS Version 2.1 Release Guide, SC26-4185.
IMS/VS Version 2.2 Release Guide, GC26-4325.

NetView

Network Program Products Planning, SC30-3351.

NCP/Communication controllers

NCP, SSP and EP Resource Definition Guide, SC30-3349.
NCP, SSP and EP Resource Definition Reference, SC30-3254.

VTAM

VTAM Operation, SC23-0113.
VTAM Installation and Resource Definition, SC23-0111.
VTAM Programming, SC23-0115.
VTAM Diagnosis Guide, SC23-0116.
VTAM Diagnosis Reference, LY30-5582.

MSCM

Multisystem Configuration Manager Programming, SC23-0174.

ES/3090

ES/3090 Processor Complex - Processor Resource/Systems Manager Planning Guide, GC22-7123.

CHANGES Summary of changes

This section summarizes the major changes made to the *XRF Guide* during CICS/ESA Version 3.

Changes for this edition

The XRF function is unchanged for CICS/ESA 3.3.

The minor changes in this book (other than editorial ones) since the CICS/ESA 3.2.1 edition are shown by a vertical bar to the left of the changes.

Additional changes made to this softcopy version of the manual since the hardcopy manual was published are indicated by the hash (#) symbol in the left-hand margin.

Changes for the CICS/ESA 3.2.1 edition

For CICS/ESA 3.2.1, the book describes changes in takeover capability and strategy made possible using MVS/ESA (*) 4.1 and its cross-system coupling facility (XCF). These changes are reflected throughout the book.

There are changes in the way that RACF is defined. The effects on XRF users are described in this book.

The CEBT transaction may now be issued by TSO operators.

Changes for the CICS/ESA 3.1.1 edition

For CICS/ESA 3.1.1, the capability of retaining signon across a takeover was introduced. The capability is described in this book, in topic 5.3.3.

There are many possibilities for configuring a CICS with XRF system using the logical or physical partitioning capabilities of processing systems. Now that separate MVS images, each capable of supporting a CICS system, can run in a single CEC, the use of the term "MVS image" has often been substituted for the term "CEC" used in the previous edition of this book.

When using XRF, a DL/I database manager is supported using DBCTL. Apart

from a brief reference in topic 7.2, the XRF aspects of DBCTL are not described in this book. For the CICS user of DBCTL, guidance information is in the *CICS/ESA CICS-IMS Database Control Guide*.

(*) IBM Trademark. For a list of trademarks, see topic FRONT_1.

1.0 An overview of XRF

XRF stands for **extended recovery facility**, and it refers to a related set of programs that allows an installation to provide higher system availability to end users. CICS/ESA (*) works with a number of IBM products in an XRF environment. MVS (*) (including the cross-system coupling facility of MVS/ESA (*) System Product Version 4 Release 1), VTAM (*), and NCP all have XRF capability. You can combine them with CICS/ESA in any IBM processor that can operate in Enterprise Systems Architecture (ESA (*)) mode, and give your CICS systems improved recovery performance and thus provide improved availability to the end user.

The XRF approach to improved availability builds on two assumptions:

1. Many installations must minimize both planned and unplanned system outages. These installations are willing to devote extra resources to improve the service to their end users.

2. A defect that causes a failure in one environment does not necessarily cause a failure in a different environment.

By coding XRF=YES as a system initialization parameter, you obtain XRF support; by coding XRF=NO, you have a CICS/ESA system without XRF support. This book is for users who intend to run a system with XRF=YES.

XRF does not eliminate outages. It minimizes the duration of certain kinds of outage. Even if all unplanned failures, caused by both hardware and software failures, could be eliminated, there would still be planned downtime for maintenance, configuration changes, or migration. XRF reduces the impact of both unplanned and planned outages on the end user, and thus provides a higher level of availability than a non-XRF CICS system.

CICS with XRF is based on the use of an **active CICS system**, which supports the processing requests from the end user, in combination with an **alternate CICS system**, which can take over from the active if the active fails or if it is taken out of service.

The active and alternate systems must be at the same level. For example, you cannot match a CICS/ESA 3.3 active system with a CICS/MVS (*) 2.1 alternate, or a CICS/ESA 3.2.1 active with a CICS/ESA 3.3 alternate. Also, if the active and alternate CICS systems are running on separate MVS operating systems, it is advisable to use the same level of MVS for both.

(*) IBM Trademark. For a list of trademarks, see topic FRONT_1.

Subtopics:

1.1 XRF environments

An XRF complex is made up of:

- The active and alternate CICS systems.
- The associated software, including the MVS/ESA operating system (each copy of which may be called an MVS image) with JES2 shared spool or JES3, and VTAM.
- One or more IBM (*) 3745/3725/3720 Communication Controllers at the boundary network node (BNN).
- The network control program (NCP).
- The terminal network.

° Shared DASD.

° The processing systems. In this book, when referring to the whole of a physical machine, such as an IBM 3090 (*) Processor Complex, or a physical partition of that machine, the term "CEC" is used. CEC is short for "central electronic complex". The term is not used to refer to logical partitions of such a machine.

In addition, an XRF complex might include:

° The Processor Resource/Systems Manager (*) feature (PR/SM (*)), which provides flexible partitioning of a 3090 processing system into a number of logical partitions.

° The PR/SM Automatic Reconfiguration Facility (ARF) provided by function in both MVS/ESA SP 4.1 and PR/SM. ARF enables the alternate CICS's MVS image to reconfigure the storage of another logical partition on the same CEC (or the same side of a physically partitioned CEC) in the event of an MVS failure prior to the completion of a CICS XRF takeover.

CICS with XRF provides different levels of enhanced availability in different environments:

° Coverage against CICS failures is provided by active and alternate CICS systems running in the same MVS image.

° Improved availability when MVS, VTAM, and CICS outages occur is provided by:

- Placing the active and alternate CICS systems in separate logical partitions, made possible by the Processor Resource/Systems Manager (PR/SM) feature. Each of these partitions supports its own MVS image and VTAM, resulting in a multi-MVS environment. Use of MVS/ESA 4.1 and PR/SM provides further opportunities for enhanced availability. See "Multi-MVS configuration using PR/SM ARF" in topic 4.3.

- Placing the active and alternate CICS systems in separate physical partitions within the same processing system (each partition operating as a processing system in its own right). For example, an IBM 3090 Model 400 Processor Complex can be physically

partitioned into a multi-MVS environment, being the equivalent of two 3090 Model 200s, each running with its own MVS image and VTAM.

Such a configuration can also provide protection against partial processor failures, if one physical partition fails and the other continues to run.

° Enhanced availability against a complete processing system failure requires two completely separate processing systems. The active and alternate CICS systems must run on physically separate CECs (for example, two 3090s) as shown in Figure 1.

Figure 1. An XRF complex

(*) IBM Trademark. For a list of trademarks, see topic FRONT_1.

1.2 A brief description of XRF

Everything mentioned here is described more fully in the sections that follow.

CICS/ESA in XRF mode is a system approach to increased availability. It uses alternate resources to overcome hardware and software outages--both planned and unplanned.

When CICS is running with XRF, there is a pair of CICS systems:

1. The **active system** running the CICS workload

2. The partially initialized **alternate system**, standing by in case of failure.

This partially initialized alternate CICS system lets you provide greater availability to your end users. It can do this by reacting automatically to problems that cause interruptions in service. Through the **CICS availability manager (CAVM)**, the active constantly communicates with the alternate, so that the alternate can record changes in terminal usage--**tracking**--and monitor the well-being of the active system--**surveillance**. Surveillance and tracking information is passed through the CAVM data sets--the **message data set** and the **control data set**. These data sets are on shared DASD, accessible to both active and alternate CICS systems. When the alternate CICS system concludes that the active has failed, or when it is instructed to act, it has access to all the necessary information and resources to **take over** from the active system and reestablish service with the minimum of interruption.

XRF can help the operator by taking away some of the operator's decision-making. The alternate can react to a failure more quickly than the operator can. When XRF has identified a failure, it can help reduce operator reaction and decision time, because it can do most of the work for the operator. With certain configurations and types of failure, XRF can do all of the work to recover and restart from a failure. This capability is enhanced by the use of MVS/ESA SP 4.1 and PR/SM, which can further automate the takeover process.

There is an optional **overseer** function, in the form of a sample program, that provides status information to the operator about the active and alternate systems. The overseer is particularly useful when you are running many active and alternate systems, perhaps linked by multiregion operation (MRO), because it gives the operator an overview of the systems that are running. The overseer can also be used to automate some operator tasks.

When the alternate takes over the running of the CICS system, it performs an emergency restart similar to an emergency restart after the failure of a non-XRF CICS system. Resources are recovered in the same way as they are in an emergency restart. However, with XRF, the whole emergency restart process is faster. This is because:

° The alternate is already partially initialized.

° There are backup VTAM sessions from the BNN communication controller to the alternate already established for the terminals to be switched.

° The restart is initiated sooner because of the surveillance activity.

Most of your existing emergency restart procedures remain valid for XRF, because XRF builds on the existing CICS emergency restart facilities.

The alternate CICS is only partially initialized. It cannot complete its initialization until its active partner has terminated. It cannot do any normal processing until it has taken over and become the new active system. The alternate takes up very little resource, so, if you are using two MVS images, the second MVS is largely available for other work.

Subtopics:

1.2.1 Terminal capability

Although XRF is made up of active and alternate CICS systems, it presents a single-system image to the end user at a VTAM terminal. A terminal only has a working session with an active system.

In general, terminals can be divided into three classes. The terminals that benefit most are known as **XRF-capable** (class 1 terminals). XRF-capable terminals are SNA VTAM terminals connected to a boundary network node BNN communication controller and network control program (NCP) with XRF capability. Such terminals are switched to the alternate at takeover time, and retain their sessions. The operator does not have to intervene in this process. This is made possible by the alternate CICS system, which requests VTAM to create **backup sessions** between the alternate and the communication controller for all the XRF-capable terminals. The alternate system makes this request to VTAM as soon as it learns from the tracking mechanism that a new terminal has logged on to the active.

When other VTAM terminals (class 2 terminals, that are not eligible for backup sessions) log on, the alternate **tracks** them, and after a takeover it tries to reestablish their sessions.

Depending on the nature of the failure, non-VTAM terminals (normally class 3 terminals) may also benefit from the improved restart time after a failure. In a multi-MVS environment, terminals that do not have a path established to the alternate might need manual intervention to effect reconnection to the alternate system.

After a takeover, users of VTAM class 1 and class 2 terminals do not normally have to sign on to CICS, because signon security may be passed from the active to the alternate CICS system. If this facility is not implemented, end users have to follow their normal procedures for emergency restart. If there is a task in flight at the time of takeover, that task must be reentered.

More detailed information about different types of terminals and their XRF capabilities is given in "The terminal network" in topic 5.0.

1.2.2 The takeover

A takeover might occur because of:

- CEC outage
- MVS outage
- VTAM outage
- CICS outage.

"Outage" refers both to a failure, and to planned downtime for maintenance or upgrade.

In either case, XRF offers end users increased system availability. There is more information about the causes of a takeover in "Types of outage handled by CICS with XRF" in topic 2.0.

When a failure has occurred and the alternate has become the active system, you should initialize another alternate, and thus maintain the extended recovery facility. To make changes to your CICS system, you can initiate a takeover to an alternate CICS system that has already had software maintenance or its configuration changed. That alternate becomes the new active, which can then be backed up with a new alternate.

In this book, we describe the decisions you make about XRF. You decide under which conditions a takeover occurs, whether to restart failed active systems rather than have a takeover, whether the operator has to authorize a takeover, and how much involvement the operator has in the takeover.

1.2.3 Failures outside the scope of XRF

XRF cannot handle all failures at a CICS installation. It does not address those outages caused by the failure of system elements that are not duplicated. For example, XRF does not deal with:

- Failures in the telecommunication network, such as the communication controller, network control program (NCP), lines, and terminals
- Loss of, or damage to, the shared DASD for CICS system data sets such as the system log and similar resources, and also for user databases (however, note the write I/O error support provided for DL/I databases)
- Loss of, or damage to, essential system data sets, such as VSAM catalogs, or the JES checkpoint data set
- An environmental failure, such as a power or air-conditioning failure, that affects both active and alternate CICS systems
- Some software failures that recur after takeover
- Some operator errors, such as the corruption of a database because it was restored from the wrong backup tape.

Your installation might already have procedures for dealing with some of these other types of failure--an uninterruptible power supply, perhaps, or strict programming standards to avoid the risk of recurrent software failures.

2.0 Types of outage handled by CICS with XRF

In this topic, the types of outage that can be handled by CICS with XRF, already outlined in the last topic, are discussed in more detail.

The figures show a multi-MVS environment. This environment could be

provided by a single CEC or by two separate CECs. The single CEC may be partitioned, logically (using the PR/SM feature) or physically, into a multi-MVS environment. This environment provides cover against MVS, VTAM, and CICS failures, as described in "XRF environments" in topic 1.1. To guard against a CEC failure, you require two separate CECs, with a JES2 multiaccess spool or JES3 configuration. XRF running in one MVS image normally covers only against failures in the CICS address space, and against outages that would routinely be caused by CICS planned maintenance.

The way a takeover works is described in "How XRF works" in topic 3.0.

Subtopics:

2.1 CICS outage

Figure 2. CICS outage

XRF provides a rapid restart after the failure of the active CICS.

You do not need two MVS images to handle CICS outages. You can run XRF on a single MVS to give increased availability during outages in the CICS address space. For the benefits that can be gained from running XRF in this way, see "Single-MVS image, single-region XRF configuration" in topic 4.4.

If an application program causes CICS to fail, and there is a takeover, it is possible that the same application could cause another failure on the new active.

2.2 VTAM outage

Figure 3. VTAM outage

A VTAM failure may result in a takeover, or you may restart VTAM and leave the active running. If VTAM on the active's side fails, it drives the TPEND exit for the active CICS, which can then decide whether a takeover is the appropriate action. You may select beforehand the situations where a takeover is necessary, by coding an exit program or adding code to the overseer program to cause the takeover or other action.

If a takeover is not selected (the CICS default action), the active continues, in degraded mode. Meanwhile, if running in two MVS images, the alternate retains its backup terminal sessions, as it may be requested to take over. For example, it may be impossible to restart the active's VTAM within a certain time limit. If you do restart VTAM, ensure that the alternate has been taken down before the operator issues the CEMT SET OPEN command on the active. If this is not done, the existence of backup sessions prevents the active from establishing new sessions. Then a new alternate may be started, and backup sessions established.

In a multi-MVS environment, if the VTAM supporting the alternate fails, the active continues normally, without the cover of backup sessions. Here, the alternate terminates, a new alternate can be started, and backup sessions established, when VTAM has been restarted.

See "Multi-MVS, single-region XRF configuration" in topic 4.1 and "User exit for VTAM failure" in topic 6.3 for more information.

2.3 MVS outage

Note: XRF cannot guarantee recovery for any type of MVS outage.

Figure 4. MVS outage

If you have two MVS images, you can run the active CICS on one MVS, and have the alternate CICS partially initialized on the other MVS. VTAM terminals that you want to switch automatically from the active to the alternate, without having to log on to VTAM again, are connected to both CICS systems through a 3745/3725/3720 communication controller.

Without XRF, an MVS (or hardware) failure means that CICS could be unavailable for a long time. With XRF, when the active can no longer function properly, either because of an MVS or hardware failure, the alternate is notified, through the CAVM, of the active's failure and initiates a takeover.

When CICS with XRF is running under MVS/ESA SP 4.1 in a sysplex, you can
use the Cross-System Coupling Facility (XCF) PR/SM Automatic Reconfiguration Facility (ARF) to automate CICS takeovers if there is an MVS image failure. ARF can be used in either a single-CEC environment or a multiple-CEC environment, depending on your availability requirements.
| For more information about running CICS with XRF in a sysplex, see
| "Running CICS with XRF in a sysplex" in topic 3.1.6.

| If you are running CICS with XRF on a release of MVS before MVS/ESA SP
| 4.1, or if you are not running MVS/ESA SP 4.1 in sysplex mode, CICS cannot determine for itself the state of another MVS image and must ask the operator to confirm that an MVS image has failed. Here, the operator has to confirm to the alternate that its active counterpart has failed because of an MVS failure, and that a takeover can go ahead. See "Checking for termination of the active" in topic 3.1.4.2.

You are advised to run active and alternate CICS systems in MVS images that are at the same level, so that each can provide the same level of service to the current active CICS. For example, MVS/ESA Version 4 provides extended console support; a takeover to a prior level of MVS would take away that support.

2.4 CEC outage

To cope with the failure of a CEC, and the other failures detailed previously, the alternate CICS has to run in a separate CEC. The second CEC could be either in a physical partition in the same processing system as the active, or in a physically separate processing system. Running the active and alternate in different 3090s, for example, provides XRF cover against a failure of the active's 3090.

When running CICS with XRF under MVS/ESA SP 4.1, you can guard against CEC failure without operator intervention by running ARF in a multi-CEC environment. For more information, see "Running CICS with XRF in a sysplex" in topic 3.1.6.

If you are not using MVS/ESA SP 4.1, when a CEC fails (as when MVS fails) the alternate cannot always be certain of what has happened to its active counterpart. The operator has to confirm to the alternate that its active counterpart has failed because of a CEC failure and that a takeover can go ahead. See "Checking for termination of the active" in topic 3.1.4.2.

2.5 Planned takeover

CICS with XRF gives you improved availability if a failure occurs. It also allows you to shut down the active system and instruct the alternate to take over to do CICS software maintenance, or to introduce changes into your CICS system more easily. In a multi-MVS or two-CEC environment, XRF also helps you to take care of the maintenance of the CECs or of other software.

There are some maintenance activities that must be performed concurrently to both the active and the alternate systems, and so upgrading through a takeover is impossible. Operation in a single MVS image is also more restrictive, because some changes cannot be made without an IPL of MVS. This applies, for example, to maintenance of any CICS software that must reside in the LPA (link pack area).

For guidance information about the use of XRF takeovers as a maintenance aid, see the *CICS/ESA Installation Guide*.

XRF gives you the flexibility, through a planned takeover, to choose when you do maintenance. You probably would not want to do a takeover during a peak period, while there are many end users on the system, unless there is a good reason for it. But you might choose to make changes more frequently, to tables for which RDO is not available, or to parameters, or to apply PTFs, for example.

To initiate a takeover, your operator can use the CEBT transaction, or an extension to the CEMT transaction, both described in "Supplied transactions for controlling the alternate" in topic 6.5.

3.0 How XRF works

Before CICS/MVS Version 2, a CICS failure meant that you needed to restart your system, probably using an emergency restart. An XRF takeover, which is simply an enhanced emergency restart, provides the same **integrity** as an emergency restart in a non-XRF system. To the end user, the takeover has a similar appearance to an emergency restart. Most of your existing emergency restart procedures will remain valid for XRF. However, an XRF takeover does not allow you to delay the restart to allow (for example) postprocessing or preprocessing job steps.

Subtopics:

3.1 An XRF sequence

Figure 5 shows a possible XRF sequence. The five stages in the sequence are described in the next five sections:

1. Initialization

2. Synchronization

3. Surveillance and tracking

4. Takeover

5. After takeover.

The description of an XRF sequence does not completely apply if you have a CICS system running under MVS/ESA SP 4.1 in a sysplex. In that case, you should still read the following sections, and then read "Running CICS with XRF in a sysplex" in topic 3.1.6. Such a configuration can reduce operator involvement and improve takeover times.

Figure 5. An XRF sequence

Subtopics:

- 3.1.1 1. Initialization
- 3.1.2 2. Synchronization
- 3.1.3 3. Surveillance and tracking
- 3.1.4 4. Takeover
- 3.1.5 5. After takeover
- 3.1.6 Running CICS with XRF in a sysplex

•

3.1.1 1. Initialization

Figure 6. Initialization of alternate after active has started processing

Figure 6 shows that you need a pair of CICS systems called the active and the alternate to use XRF. You start the active and the alternate separately, and you can start them concurrently, or in either order. The startup job streams for active and alternate must be very similar except for some of the system initialization parameters (probably overrides), and certain data-set definitions. (For examples of startup job streams, see the *CICS/ESA Operations Guide* .)

The actives and alternates have their own local catalog, dump, and auxiliary trace data sets. They either share or have their own extrapartition transient data data sets. The alternate has its own transient data destination, CXRF, which is dynamically defined and is available to the alternate before takeover. For guidance information about how to use CXRF, see the description of the DFHCXRF data set in the *CICS/ESA System Definition Guide*. Apart from such minor differences, the active and alternate must be compatible, with the same recoverable resource definitions. This ensures that, after a takeover, the new active provides the same service as before.

The active and alternate **sign on** to the CICS availability manager (CAVM) at the start of initialization. The CAVM is the mechanism that allows actives and alternates to coordinate their processing. The CAVM uses a shared pair of data sets: a control data set and a message data set. Each active and each alternate has its own CAVM (in the CICS region), and the active and alternate pair share the CAVM data sets.

This pair of data sets is logically a single entity which contains:

° State data whose main purpose is to ensure that one of the CICS jobs

sharing that particular pair of data sets is allowed to perform the active role at any time

° Primary and secondary surveillance signals of actives and alternates, so that each system can tell whether its partner is working correctly

° Messages about the state of some resources in use on the active, which are written by the active, and read and processed by the alternate.

CAVM rejects a request from a CICS job to sign on as the active if the control data set shows that an active is already present, or that a takeover is in progress. This ensures that the integrity of files and databases cannot be lost because of uncontrolled concurrent updating by two or more actives. When an active or alternate signs on, it starts to write its own surveillance signals, and to look for its partner's surveillance signals.

The control data set is used:

° To record the presence or absence, identities, and current state of active and alternate CICS jobs

° For the primary surveillance signals of the active and alternate.

The message data set is used:

° Principally to pass messages about the current state of certain resources from the active to the alternate

° For the secondary surveillance signals of the active and alternate systems, when the control data set is unavailable for this purpose, either because the last write has not completed or because of I/O errors.

Once a pair of CAVM data sets has been used by the active and alternate systems that share a generic applid, those data sets may not subsequently be used by another active or alternate with a different generic applid.

For more guidance information about the CAVM data sets, see the *CICS/ESA System Definition Guide*.

The active completes its initialization normally. It then begins to provide a service to its end users.

The alternate cannot be fully initialized because, until it takes over from its active counterpart, it does not own the resources that can be used by only one system at a time, such as the system log and user data sets. The alternate is initialized only to the point at which it can monitor the active. VTAM must be running before the alternate can complete its initialization. Only one alternate at a time is allowed to sign on to the CAVM. If the alternate is started first, it waits, watching for its active partner's surveillance signals to start when it signs on to the CAVM.

The alternate cannot perform any active CICS function (users cannot log on to it, for example), and it takes up very little resource. The only means of external communication with the alternate is through the MVS MODIFY command or the overseer. The MODIFY command is limited to a small set of CEBT commands, described in "Supplied transactions for controlling the alternate" in topic 6.5. The overseer is described in "The overseer" in topic 6.4. The alternate carries out **surveillance** and **tracking**, writing its own surveillance signals, reading the active's surveillance signals, and reading messages describing the status of terminals in the active.

Subtopics:

3.1.1.1 Running the active by itself

The active can run by itself without a matching alternate. This is shown in Figure 5 in topic 3.1. You may start an active and not start a matching alternate, or you might choose to take down the alternate at periods of low activity.

3.1.2 2. Synchronization

When the active is initialized, and it detects that the alternate has signed on to CAVM, they are both at the synchronization stage. The active uses CAVM message services to send a stream of messages describing the

current state of all its VTAM terminals via the message data set to the alternate. This is called the **catch-up** process, which allows the alternate to build a complete picture of the active's terminal resources and the status of those terminals. In this way, the alternate is aware of the existing terminal network, can request backup sessions for XRF-capable terminals, and can track any remaining VTAM terminals.

If the alternate stops for any reason, and the active runs by itself for some time before another alternate is started, the same catch-up process is used for the new alternate.

Then the active and alternate enter the surveillance and tracking stage.

3.1.3 3. Surveillance and tracking

Most of the time, CICS with XRF is in the third stage: surveillance and tracking, as shown in Figure 7.

The active sends out surveillance signals to the alternate, and the alternate monitors them, checking for any sign of failure in the active. If the active itself detects a failure that prevents it from continuing to provide a service, it signs off abnormally from the CAVM to inform the alternate of its failure. A CEC, MVS, or serious CICS failure causes the active's surveillance signals to stop.

While running normally, the active uses CAVM message services to inform the alternate about changes made to the terminals installed in the system. The active also informs the alternate of changes to the installed, logged-on, and logged-off state of all VTAM terminals and sessions as they are acquired or released. In this way, the alternate tracks the installed, logged-on and logged-off state of all VTAM terminals. If the tracked logged-on state suggests XRF capability, the alternate requests VTAM to set up a backup session for that terminal so that it can be switched at takeover. This session is set up between the alternate and the NCP in the BNN communication controller.

The emphasis in surveillance is that the alternate monitors the state of the active. But, at the same time, the active continually checks the status of the alternate and its surveillance signals, to ensure that an alternate exists to receive the messages it is sending. If the alternate's surveillance signal disappears, or it signs off abnormally from the CAVM, the active warns the system operator. Loss of the alternate does not affect the running of the active. When another alternate is started, synchronization begins again.

Figure 7. Use of the CAVM data sets for surveillance and tracking

---

3.1.4 4. Takeover

A takeover can be started by several events:

- The alternate detects that the active has signed off abnormally from the CAVM.
- The alternate detects the disappearance of the active's surveillance signal.
- The operator or an MRO-connected region that is taking over sends the alternate a CEBT PERFORM TAKEOVER command.
- The operator issues a CEMT PERFORM SHUTDOWN TAKEOVER or a CEMT PERFORM SHUTDOWN IMMEDIATE command to the active.

The type of event and the TAKEOVR system initialization parameter determine whether a takeover occurs and also the level of operator involvement in that takeover. The system initialization TAKEOVR parameters--AUTO, MANUAL, and COMMAND--are described in "Starting the alternate" in topic 6.1.2.

Active signs off abnormally from the CAVM: If the active signs off abnormally from the CAVM, for whatever reason, and TAKEOVR=COMMAND is **not** specified, the alternate starts a takeover.

Alternate detects the disappearance of the surveillance signal: If the alternate detects that the active's surveillance signals have disappeared, the action taken by the alternate is dependent on its current takeover operand, as follows:

TAKEOVR=AUTO
    The alternate initiates a takeover automatically, when the alternate delay interval (ADI) has elapsed.

TAKEOVR=COMMAND
    The alternate does **not** initiate a takeover.

TAKEOVR=MANUAL
    After the ADI interval has elapsed, the alternate sends a message asking the operator whether it should try to takeover, or ignore the apparent failure of the active. If the operator can repair the active, the alternate can be told to ignore the loss of the surveillance signal. If the active recovers, the alternate detects the reappearance of its surveillance signal, cancels the message to the operator, and continues with its standby role. If the operator cannot repair the active, the alternate should be told to begin takeover.

CEBT PERFORM TAKEOVER command: This command may be issued to the alternate by the operator, by another alternate taking over in a multi-MVS MRO configuration, or by the overseer. On receipt of this command, the alternate starts taking over, without reference to the operator, regardless of the takeover operand.

A CEMT PERFORM SHUTDOWN TAKEOVER (IMMEDIATE) command is issued: The CEMT PERFORM SHUTDOWN IMMEDIATE or the CEMT PERFORM SHUTDOWN TAKEOVER command can be used to start a takeover by telling the active to shut down and sign off abnormally from the CAVM. However, a takeover only occurs if TAKEOVR=AUTO or TAKEOVR=MANUAL has been defined in the system initialization parameter for the alternate.

Figure 8. Takeover

Subtopics:

3.1.4.1 Takeover begins

Once it has been decided that the alternate will try to take over from the active, a takeover request is passed to the CAVM, as shown in Figure 8 in topic 3.1.4. In most cases this request will be accepted, but may be rejected for any of the following reasons:

- The active has already signed off normally.
- The active is not the same active as the one that the alternate had been tracking. The CAVM detects that it is a new active, probably because of a restart in place. Here, the alternate cannot continue its role, and a new alternate should be started.
- The active and alternate are on different MVS images, and the alternate has not been monitoring the active's surveillance signals long enough to assess the difference between the time-of-day clocks on the two MVS images.

When the CAVM has accepted the takeover request from the alternate, an attempt by another CICS to sign on to the CAVM as an active will be rejected. The alternate next issues the command:

```
MVS MODIFY netname,USERVAR
```

to redefine the CICS application name, and begins to switch the sessions of the XRF-capable terminals.

During takeover, the alternate uses two different mechanisms to try to force the termination of the active CICS job, as follows:

1. If the active is still signed on to the CAVM, the alternate uses the surveillance mechanism to try to pass a "takeover-requested" message to the active, including a "dump" or "no-dump" indicator. If the active receives the message, it responds by issuing abend U206, and eventually signs off abnormally from the CAVM.

2. If the active job is still executing, the alternate also issues a CANCEL command (prefixed by a JES routing command in a multi-MVS configuration). The CANCEL command is issued if the active is unable to respond to the alternate's request to take over.

Next, the alternate starts to process the command list table (CLT). You build your CLT to describe what will happen at takeover. It provides the authorization to cancel the active system, and can also contain routing information, MVS system commands, and messages to the operator. For more information, see "Command list table (CLT)" in topic 6.2.

3.1.4.2 Checking for termination of the active

The alternate asks JES periodically about the status of the active. Job termination ensures that all I/O activity has been completed (or will subsequently be backed out), and thus ensures data integrity. If JES replies that the job has terminated, the next phase, "Completing the takeover" in topic 3.1.4.3, can start immediately.

If JES replies that the job is still executing, the alternate continues to check the status until the interval defined by the JESDI system initialization parameter expires. After that interval, the alternate prompts the operator (with message DFH6561 or DFH6581) to investigate why the job has not stopped. There might be a JES problem, or an authorization problem in the CLT. The alternate also offers this prompt if JES is not running, or does not respond.

When active and alternate are running in different MVS images, JES might continue to tell the alternate that the active job is still running even though the active's MVS or CEC has failed. Here, the alternate cannot complete its takeover without operator intervention. Another possibility is that the active job is still running, and either never received the CANCEL command, or received it but could not terminate because a system error necessitating a FORCE command has occurred.

If the active's MVS has not failed, the operator must ensure that the active job really has terminated before informing the alternate that the active job has ended.

If the active's MVS has failed, and the operator decides that an IPL is required, the operator should stop the processors of the failed MVS and IPL the system, after which the operator can reply to the alternate's question, notifying it that the CEC has failed.

Here, an internal record is kept that the MVS image, identified by its SMF system identification (SID), has failed. Other alternates examine this record while they are taking over, to try to avoid operator intervention.

The alternate cannot complete takeover until the operator replies to its question, unless either of the following occurs:

- The alternate receives a late reply from JES that the active job has terminated

- A previous reply to another alternate's message has already confirmed CEC or MVS failure.

In either case, the operator does not have to reply, and takeover continues.

3.1.4.3 Completing the takeover

When CAVM has received confirmation that the active CICS job has terminated, it notifies the alternate that it may now assume the fully active role, and updates the CAVM control data set to this effect.

Takeover resumes. In a multi-MVS environment, if the time-of-day clock of the new active's MVS is slow compared with the time-of-day clock of the old active's, the takeover is delayed until the new active's time-of-day clock has reached the value of the old active's clock at the time of job termination. This is because recovery processing depends on time-of-day clock readings to establish the correct sequence of events. Then the alternate completes its takeover, and becomes the active.

The restart time is improved by the existence of backup sessions for XRF-capable terminals. The new active does not have to establish new sessions for these terminals. It cleans up the sessions for them, and

reestablishes the sessions for other VTAM terminals. Terminal users still use the same generic applid (the USERVAR) to log on to CICS.

3.1.4.4 Logging and archiving

Because the aim is to provide a rapid recovery from a failure, your system log must be on two disk data sets. To avoid any archiving delay, and consequent unnecessary takeover delay, you are advised to use automatic archiving, specified by the JOUROPT=AUTOARCH operand of the DFHJCT macro. For further guidance information about automatic archiving, see the *CICS/ESA Recovery and Restart Guide*.

If you submit the archiving job for execution on the active's MVS, and that MVS fails while an archiving job is running, the job has to be resubmitted, and takeover might be delayed until it finishes. This problem could be avoided by making a practice of submitting the archiving job for execution on the other MVS.

3.1.4.5 Failure analysis

Diagnostic information about the failure of the active is provided by the usual termination dumps. Taking a dump is a part of the CICS job, and the alternate cannot complete its takeover until the active job has taken its dump and terminated.

You are recommended to specify SDUMP as the termination dump, to provide adequate diagnostics, and to ensure that the active closes down as quickly as possible. For more information, see the ADI parameter description in topic 6.1.2. You can format and analyze the dump using the interactive problem control system (IPCS). For more guidance information about how to do this, refer to the *CICS/ESA Operations Guide*.

If the active is running normally and it is being taken over because of a command from the operator or from another CICS region, no dump is taken.

3.1.5 5. After takeover

In a multi-MVS environment, after the takeover, the operator manually switches any devices that need to be physically connected to the new

active: perhaps local VTAM, or TCAM terminals, or other software outside the control of CICS.

Depending on the options you set, end users of VTAM class 1 and class 2 terminals do not normally have to sign on again after their terminals have been switched to the new active.

As in an emergency restart, an end user might have to reenter the last transaction, if that transaction was in flight when the active failed. This applies to all classes of terminals. If you prefer, you can send your own message, telling the end user what to do. You should consider your methods of establishing what was the last successful logical unit of work (LUW) before the takeover occurred, so that you can provide users with a significant recovery message.

Subtopics:

3.1.5.1 Initiating network changes

To allow additional end users to log on after a takeover, VTAM must change the application name (specific applid) in its USERVAR table. The alternate issues the command:

```
MVS MODIFY vtamname,USERVAR
```

to change the entry of the local USERVAR. If you are running XRF with ACF/VTAM (*) 3.2 (with the appropriate PTF), USERVAR values in remote VTAMs communicating with the local VTAM are changed by VTAM. See "The terminal network" in topic 5.0 for more information about USERVARs and applids.

(*) IBM Trademark. For a list of trademarks, see topic FRONT_1.

3.1.5.2 Reestablishing the system

When the old alternate has become the new active, there is a period when it runs without an alternate as its partner. You should plan to start an alternate as quickly as possible to restore the protection of XRF to your users. You can use the old active job's JCL for the new alternate job, ensuring that the correct START override is coded, or you can use different JCL. The job to start a new alternate may begin execution when you know that the old alternate has become the new active. This will probably be before the new active has finished the takeover.

When the new alternate has been initialized, new backup sessions are established for end users with XRF-capable terminals.

3.1.6 Running CICS with XRF in a sysplex

If you are running CICS with XRF in a multi-image MVS/ESA SP 4.1 sysplex, there are additional considerations to those described in the sequence above. You should read through the sequence, and then read this section.

Subtopics:

3.1.6.1 Single-CEC sysplex

When CICS with XRF is running under MVS/ESA SP 4.1, you can use the Automatic Reconfiguration Facility (ARF) in a single-CEC environment to provide a improved availability platform, at minimal hardware cost,

against software outages. Specifically, this protects you against MVS system failures.

Figure 9. Single-CEC improved availability using XCF

If the primary and backup MVS images (MVS1 and MVS2 respectively in Figure 9) are defined to be in the same sysplex, the CICS takeover completes without any operator intervention if:

1. Both the active and the alternate CICS are running under different MVS/ESA SP 4.1 images.

2. Both MVS systems are running in logical partitions (LPs) on the same CEC, or the same side of a physically partitioned CEC.

3. Both MVS systems are members of the same sysplex.

4. A Cross-system Coupling Facility (XCF) PR/SM policy is defined specifying the actions that the backup MVS (MVS2) is to take if the primary MVS fails.

In the example shown in Figure 9, the primary LP (MVS1) is defined to have most of the CEC resources while the backup LP (MVS2) is defined to have the minimal resources required to monitor the active CICS. The backup system, when it detects that the primary system appears to be nonfunctional, processes the XCF PR/SM policy. XCF PR/SM policy allows you to specify that the backup system is to reset the primary LP or deactivate the primary LP and acquire the storage resources defined to the primary LP. ARF enables the backup LP to acquire those storage resources.
| In this case, because the backup LP has only minimal resources, the policy
| is to deactivate the primary LP and take over its resources. If the
| policy was to reset the primary LP, the backup LP would need equivalent
| resources to those of the active LP. When the backup LP has taken over from the failed LP, CICS can continue the takeover.

XCF PR/SM policy allows you to define the delay intervals before a failing LP should be reset to allow MVS to be re-IPLed in it or deactivated to

allow its storage to be acquired by another MVS system running in another LP. For more details, see "Defining your XCF PR/SM policy" in topic 6.6.

In a single-CEC environment, the backup MVS executes the XCF PR/SM policy when either:

- A 'status update missing' condition occurs (a system apparently failed because its status has not been updated within the XCF INTERVAL), or

- A 'system gone' condition occurs (a system has been removed from the sysplex).

The effect is to perform a system reset of the primary-MVS LP, or to deactivate the primary-MVS LP and acquire the storage resources owned by the primary-MVS LP.

3.1.6.2 Multiple-CEC sysplex

The multiple-CEC environment, in which the active and alternate CICS systems are executing on different CECs, provides improved availability against software and hardware outages. Specifically, this protects you against hardware failures and MVS image failures.

Figure 10. Multiple-CEC improved availability using XCF

The primary system (MVS1), executing in basic mode or in a large LP, runs the active CICS on one CEC, and the alternate CICS, running in a small LP (MVS2) on another CEC, monitors the active CICS. The backup system processes the XCF PR/SM policy when the primary system appears to be nonfunctional. XCF PR/SM policy allows you to specify that the backup system can deactivate expendable logical partitions (MVSX) and acquire the storage resources defined to them. ARF enables the backup logical

partition to acquire the storage resources previously owned by the expendable logical partitions.

When running CICS with XRF in a multi-CEC sysplex, conditions 1, 3, and 4 stated previously still apply. If appropriate values are assigned to the CICS ADI and JESDI system initialization parameters (see "Defining your XCF PR/SM policy" in topic 6.6), CICS itself does not require the operator to respond to any prompts if an MVS or CEC failure occurs. However, XCF produces the IXC402D message to ask the operator to confirm the failure of the primary MVS or primary CEC, after doing a system reset of the primary system. However, the system reset and the reply to the message could be made by an automation tool such as NetView (*) with Target System Control Facility (TSCF). System reset clears any allegiance between the primary system and any shared I/O devices (for example, any reserves on shared DASD are cleared). Data integrity of shared data is maintained only if the automation tool successfully resets the primary **before** replying to IXC402D. As soon as a reply to this message has been received to confirm the failure, the XCF PR/SM policy is executed to deactivate the specified expendable LPs on the same CEC (or the same side of a physically partitioned CEC) to acquire their resources. These expendable LPs may be members of the same sysplex but need not be. The LPs that have been deactivated depend on the specified XCF PR/SM policy:

- If DEACTIVATE (sysname) is specified (implying that the system is a member of the sysplex and it may or may not have failed), the "sysname" system's LP is deactivated and its storage resources acquired by the MVS image processing the policy, if STORE (YES) and ESTORE(YES) are specified.

- If DEACTIVATE (ALL) is specified (implying that the LPs are not executing MVS images that are in the sysplex), one or more LPs with storage within the addressing range of the MVS image processing the policy are deactivated and their storage resources acquired.

3.1.6.3 Further XCF considerations

The meaning of the CICS system initialization parameter TAKEOVR=COMMAND is changed if an MVS image failure occurs. If XCF informs CICS surveillance that an active CICS system's MVS has failed, it triggers an XRF takeover regardless of whether or not you used TAKEOVR=COMMAND to request operator

intervention. (The other TAKEOVR= parameters automatically cause a takeover after an MVS failure, and the meaning of TAKEOVR=COMMAND remains the same for failures of components other than MVS.)

| In an XCF sysplex, use the AUTOESYS and RESTART options of the MASDEF
| initialization statement to:

| ° Reset the software checkpoint lock if only the JES2 subsystem fails
| while holding that lock

| ° Make jobs eligible for restart if another JES2 member (MVS and JES2)
| in a sysplex has failed.

| Note that the AUTOESYS option has no effect on the hardware lock. If it
| is held by a failed member, the operator must reset the lock. Through the
| MASDEF initialization statement, indicate whether:

| ° A JES2 member can have other members make its jobs eligible for
| restart if it fails (by specifying AUTOESYS=ON)

| ° A JES2 member can make jobs from any other failed member eligible for
| restart (by specifying AUTOESYS=ON and RESTART=YES)

| ° A JES2 member can reset the software checkpoint lock held by a failed
| member (by specifying AUTOESYS=ON and RESTART=YES).

| For more information, see the *MVS/ESA JES2 Initialization and Tuning
| Reference*.

In an XCF sysplex, there are no time-of-day (TOD) clock considerations. All systems in the same XCF sysplex must have synchronized TOD clocks. This is guaranteed in a PR/SM environment, where all MVS images share the same hardware clock. In a non-PR/SM environment, clocks are synchronized by a sysplex timer. This means that the operator no longer has to synchronize TOD clocks in a multi-MVS/ESA SP 4.1 environment.

A further advantage of the multi-MVS/ESA SP 4.1 XCF sysplex environment under PR/SM is that down-level CICS systems can benefit from fully-automated takeover if an MVS image failure occurs, if at least one CICS system running in the backup MVS image is a CICS ESA 3.3 system. A down-level CICS alternate, running on the same MVS image, can detect the failure of its active CICS's MVS by interrogation of the internal record of known MVS image failures. This relieves the operator of any involvement in the initiation of a takeover, and also means that it is not

necessary to migrate all CICS regions to CICS/ESA 3.3 before some of the benefits of automated takeover can be achieved.

For example, consider the configuration shown in Figure 12 in topic 4.2. Takeover of all 3 regions completes without operator intervention if MVS1 fails, if

- MVS1 and MVS2 are running in LPs on the same CEC, or on the same side if the CEC is physically partitioned, and

- MVS1 and MVS2: are members of the same sysplex, and

- The terminal-owning regions (active and alternate) are CICS/ESA 3.3 systems.

For more information about these environments, see the *MVS/ESA SP Version 4 Release 1 Implementation Guide*.

3.2 Operations and management

For operations staff, the XRF environment brings new tasks. For example, there is the CEBT transaction for controlling the alternate, described on topic 6.5. The overseer, described on topic 4.2.3, also has an operator command interface. To play their part in a rapid takeover, operators must understand what they have to do during a takeover, and this in turn depends on the sort of takeover. Use of a sysplex, described in "Running CICS with XRF in a sysplex" in topic 3.1.6, can reduce operator involvement in takeovers following MVS image failures.

In a large installation, it might be worthwhile to rearrange system consoles, so that operators can easily communicate, or to simplify operator control of an XRF complex after a takeover. A second master terminal for each active, permanently available, is a useful addition.

Your existing CICS application programs and user exits should execute unchanged in an XRF environment. You might have to make changes to programs running in an ISC environment; see "LUTYPE6 ISC application-to-application sessions" in topic 5.4.1.

In a multi-MVS environment, you must ensure that databases and other shared information, like the system log, are placed on shared DASD. (Some shared information, such as user journals, may be on tape.) Data specific

to the active or to the alternate does not have to be on shared DASD. If you want to collect data across a takeover, you might have to modify utilities to read unique data from the old active and from the new active.

Clearly, XRF involves new and changed procedures for your installation. By careful planning and organization, you can minimize this overhead.

3.3 Performance

The *CICS/ESA Performance Guide* contains further guidance information about XRF performance. This section contains some general points.

Subtopics:

3.3.1 Takeover performance

Takeover performance may be considered as the time it takes to close down the active, establish the alternate as the running system, and switch the terminal network. This performance depends on many factors, including the:

- Number of CECs
- Model and characteristics of the CECs
- Use of logical or physical partitioning
- Availability of XCF services and whether all MVS images are running in PR/SM LPs with the PR/SM Automatic Reconfiguration Facility activated

- Number of related regions to be taken over
- Number of open databases or files
- Number of recoverable inflight transactions
- Number of active terminals, lines, and NCPs
- Recovery mode chosen for terminals
- Frequency of activity keypointing
- Type of dump, if any, taken by the active
- Setting of the alternate delay interval (ADI) parameter
- Communication management configuration in use
- Time difference between the two time-of-day clocks in a multi-MVS environment (but not applicable in a sysplex).

XRF improves recovery times by detecting the failures automatically, and by automating the recovery and restart process (fully or partially, depending on your configuration and the work, if any, that you want to leave to an operator). The benefits are particularly evident in a multi-MVS, large network environment.

3.3.2 Performance during normal running

During normal running, the working of the CAVM is the main difference, in performance terms, between an active XRF system and a non-XRF system. The additional overhead of the surveillance mechanism of the CAVM on the active and alternate operations is small, as it normally involves only the reading and writing of the surveillance signals in the CAVM data sets. Greatest use is most likely to occur during synchronization, when the active is sending the catchup messages to the alternate. If a system performs adequately in non-XRF mode, moving to XRF should not introduce a performance problem.

The alternate is potentially the active, so you should normally assign to it the same priority and performance group that you assign to the active. You should also consider the real storage isolation of the CICS system.

3.3.3 Workload on a second MVS image

This section is concerned with multi-MVS configurations. The second MVS (where the alternate is running) does not have to be used entirely for processing the alternate, which incurs only a small overhead. It can be used for other processing, perhaps batch work, or as "the active's MVS" for another pair of CICS systems, or for a non-XRF CICS system used to test and debug application programs.

Both the single and the multiple-CEC environments described above do not guard you against CICS failures. If CICS fails in either environment the CICS XRF takeover might also fail. (The backup MVS image may not be of sufficient size.) Restart in place of failing CICS regions should be performed (TAKEOVR=COMMAND), but this can be automated using the overseer. See "The overseer" in topic 6.4.

At takeover, the alternate issues an MVS fast workload accept request, to ask MVS to adapt more rapidly to the sudden increase in its resource needs. After a takeover, the new active provides the same service as the old. In a two-CEC environment, if the new active is in a CEC that is already running near capacity, you should make arrangements to suspend some of the work. This could be a particular concern if the alternate's CEC is smaller than the active's CEC. You might, for example, have to suspend some batch jobs temporarily. You could use the MVS system resources manager to automate this.

If there are other subsystems running in the alternate's CEC, such as IMS/VS or TSO, and they continue to run after takeover, performance will be degraded because the new active takes up more of the CEC's resources.
A lot depends on how the MVS tuning parameters have been set. Refer to
the *MVS/ESA Initialization and Tuning Reference* manual.

Subtopics:

- 3.3.3.1 Workload considerations when using ARF
-

3.3.3.1 Workload considerations when using ARF

If you are running your CICS XRF systems under MVS/ESA SP 4.1, you can take advantage of the PR/SM Automatic Reconfiguration Facility (ARF). It allows you to define the second MVS image such that it has minimal storage resource, just enough for it to support an alternate CICS. If the first MVS fails, the backup MVS image can acquire the necessary extra storage resource from another LP on the same CEC (or on the same side of physically partitioned CEC).

ARF applies to both single- and multiple-CEC environments. In a single-CEC environment, ARF allows efficient running of your active CICS system in an LP of sufficient size. At the same time your alternate CICS system is running under another MVS image in an LP with just enough storage to monitor the status of the active. If the active image fails, the XCF PR/SM policy allows you to specify that the backup LP is to deactivate the primary LP and acquire its storage resource. Such a single-CEC environment provides a guard against MVS failures only. If you need to guard against hardware failure, a multi-CEC environment should be considered.

In a multiple-CEC environment the XCF PR/SM policy allows the backup LP to acquire the storage resource of expendable LPs on the same CEC as the backup LP (or on the same side of CEC if physically partitioned). The backup LP requires minimal resources, so the rest of the CEC could be used to run other work until the active CICS's MVS image or CEC fails. To
| allow a CICS XRF takeover to commence, the work in the expendable
| partitions is destroyed and the resource reallocated to the backup LP.
| The work may be restarted on the backup LP. Such a configuration provides an improved availability platform while allowing you to make best use of the spare capacity on the second CEC until it is required to take over the CICS workload.

4.0 XRF configurations

There are many ways in which you can set up CICS systems for XRF. This topic describes some example configurations:

- "Multi-MVS, single-region XRF configuration" in topic 4.1
- "Multi-MVS, MRO XRF configuration" in topic 4.2
- "Multi-MVS configuration using PR/SM ARF" in topic 4.3
- "Single-MVS image, single-region XRF configuration" in topic 4.4
- "Single-MVS image, MRO XRF configuration" in topic 4.5
- "Further configurations" in topic 4.6.

With each configuration diagram, there is a short explanation of the

availability enhancements that each configuration provides. This topic does not tell you how to set up CICS with XRF, nor how to control the takeover, restart in place, or hierarchy of regions. That information is in "Defining CICS for XRF" in topic 6.0.

A single 3090, logically or physically partitioned, can run multi-MVS images, making possible a CICS with XRF system providing cover against MVS, VTAM, and CICS outages.

A single-MVS configuration provides protection against outages in the CICS address space only. But, if you want to reduce the downtime caused by CICS failures, or if you are interested in applying CICS maintenance with less impact on your system, such a configuration might be a suitable choice.

You need a two-CEC configuration if you want to provide protection against outages of the CEC, MVS, VTAM, and CICS.

The examples that follow begin with multi-MVS configurations. Even if you are not concerned with multi-MVS configurations, it is best to read them first, because the information builds up through the examples.

Subtopics:

4.1 Multi-MVS, single-region XRF configuration

The multi-MVS, single-region XRF configuration, shown in Figure 11, offers

increased protection against outages of MVS, VTAM, and CICS, and, in a two-CEC configuration, of the CEC. The active and alternate must be in the same equipment complex, so that they can share DASD, and they must be coupled by JES2 multiaccess spool or by JES3.

Figure 11. Multi-MVS, single-region XRF configuration

If a CEC, MVS, or CICS failure occurs, the CAVM surveillance mechanism of the alternate recognizes the failing state of the active. The alternate can take over, and resume the workload of the failed system.

The alternate may be attached to a different RACF with a different RACF inventory. You must ensure that the RACF definitions are the same in the active and alternate systems.

VTAM is a special case. When VTAM fails, you can initiate a takeover, but if, for example, you have a system where many terminals are not XRF-capable, you might prefer not to do so. You might gain better availability by allowing other, unaffected users to continue to work without the interruption of a takeover. There are two ways that you can select your course of action:

1. The XXRSTAT user exit allows you to decide what to do if VTAM fails. The exit allows you to abend CICS, which could lead to a takeover, or you could do nothing and wait for VTAM to restart.

2. The overseer program, introduced more fully in topic 4.2.3, can be customized to allow you to initiate a takeover, or to wait for VTAM to recover and then act appropriately.

More information about the exit and the overseer is given in "Defining CICS for XRF" in topic 6.0. In this configuration, a simple exit program is probably a more suitable tool for deciding whether to take over, rather than the more complex overseer program.

If you are using XRF primarily to protect against non-CICS failures, for a CICS failure you might prefer to try to restart the failing CICS region (restart in place) before taking over, to try to minimize the disruption

to the end user. You might choose to restart in place if many terminals need manual switching, or if (in a two-CEC configuration) the alternate CEC is heavily loaded at the time of the CICS failure, or if the time taken by a restart in place compares well with the time taken by a takeover. There is a further discussion of restarting in place, in an MRO environment, in topic 4.2.2.

At takeover time, the logged-on terminals that are XRF-capable are automatically switched from the active sessions to the backup sessions on the alternate CICS. The end users of class 1 and class 2 terminals do not have to log on to VTAM again, and, depending on the options set, they do not have to sign on to CICS again, because signon security may be passed from the active to the alternate. A user who is in the middle of a transaction when the system goes down will have to go through the same procedures as in a non-XRF emergency restart. You can provide your own message to tell end users what to do. XRF will certainly shorten the length of the interruption.

Terminals that are not XRF-capable--such as local VTAM terminals, TCAM terminals, or remote non-SNA VTAM terminals--could also have faster recovery because of the quicker restart that XRF provides.

4.2 Multi-MVS, MRO XRF configuration

The multi-MVS MRO configuration offers increased availability against outages of MVS, VTAM, and CICS, and (in a two-CEC environment) of the CEC. The CICS system is divided into several MRO-connected active regions, each with its own alternate. Figure 12 shows active and alternate CICS regions for terminals, applications, and databases. There could be, for example, several active terminal regions, each backed up by an alternate region, or several application regions or database regions. However, the division into regions could also be along different functional lines from the ones suggested here. Note that there is no communication between the alternate regions before takeover.

In this multiregion configuration, there are more things to consider about a takeover than in a single-region configuration. The takeover is across MVS images. If one alternate region takes over, all the related alternate regions must take over, because interregion communication does not operate across MVS images. A CEC or MVS failure clearly should result in a takeover of all the regions.

Figure 12. Multi-MVS, MRO XRF configuration

VTAM failures are a special case, as discussed in the previous section.

If a region owning XRF-capable terminals were to experience a CICS failure, you might want a takeover by the alternate, because some or all of your most important end users would have lost their sessions with CICS. The takeover of that region would require the takeover of all the other regions in that MRO complex. However, if you have several CICS regions owning XRF-capable terminals, and only one of them fails, you might decide not to have a takeover, but to retain maximum availability for all other users at the expense of users of the failed region.

In an MRO configuration, you decide how important each region is, and whether there should be a takeover if a region fails. The alternative to a takeover is to restart a region in place, rather than involving all the related regions in a takeover.

Subtopics:

4.2.1 Hierarchy of regions

To help understand a takeover strategy that handles regions of varying importance, you might find it useful to think of your regions as forming a hierarchy. A typical arrangement is shown in Figure 13.

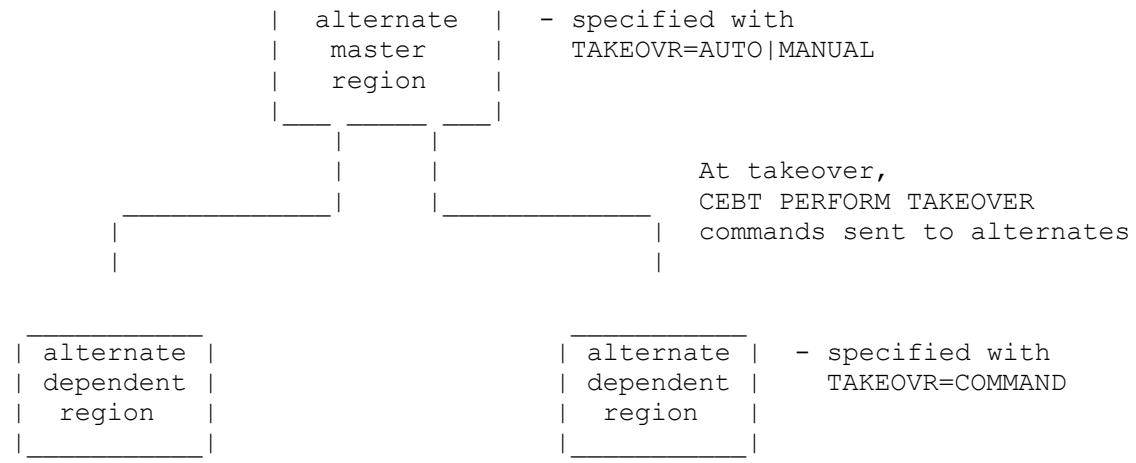

Figure 13. Hierarchy of one master and two dependent regions

Each region in an MRO complex may be considered as a **master**, **dependent**, or **coordinator** region. A region that instructs other connected regions to take over, in the event of its own takeover, may be regarded as a master or coordinator region. A region that does not initiate the takeover of other connected regions in the event of its own failure may be regarded as a dependent region.

A dependent region differs from a master or coordinator region in that its takeover operand is TAKEOVR=COMMAND. This means that the failure of a dependent region does not result in its own takeover, nor does it force a takeover of the entire complex of regions. Instead, the system operator (or perhaps the XRF overseer) tries a restart in place using existing emergency restart procedures.

A region running XRF-capable terminals would probably be a master region. The failure of an active master region results in its takeover by its alternate region. That alternate master region initiates its own takeover, and issues:

```
CEBT PERFORM TAKEOVER
```

commands in its CLT to all the other alternate regions, instructing them to take over from their active counterparts. These other regions are the dependent regions, probably application-owning or database-owning regions.

If there is more than one master region, one of them may be made the coordinator region. If a master region or the coordinator region fails, then only the alternate coordinator region issues commands to all the other alternate regions, instructing each alternate to take over from its active counterpart. By using a coordinator, you avoid having several

master alternate regions all instructing the other alternate regions to take over. Any region may be nominated as dependent, master, or coordinator.

In this way, the coordinator is responsible for the takeover of all its MRO-connected regions. If an alternate coordinator region is called on to start a general takeover, and that alternate coordinator is not running for some reason, an automatic takeover is impossible, and the operator must intervene.

There is no specific definition of a region as dependent, master, or coordinator. A region is related to its connected regions by the contents of the CLT, one for each alternate, and by the TAKEOVR system initialization parameter. You code your own CLTs to suit the structure of your system. If you prefer, you can write one CLT for a set of MRO-connected alternate regions, with a separate section for each region. The CLT, the system initialization parameter, and the CEBT transaction are described in "Defining CICS for XRF" in topic 6.0.

4.2.2 Restarting regions in place

Usually, a master region may be regarded as one that causes a takeover if it fails. That takeover involves all the related regions in their own takeovers. A dependent region is one that does not cause a takeover (its own or that of any other region) if it fails. When a dependent region fails, it is normal to try to restart that region in place. A restart in place of one region could cause less disruption to your end users than a takeover of all the related regions.

A restart in place might be particularly appropriate for an application-owning region. These regions are usually quick to restart. Then, if you could not restart that region in the active CICS complex, and the region was necessary to your operation, you could force a takeover of all the related regions to the other MVS image. When the active is restarted in place, the alternate closes down automatically, because the old alternate cannot provide support to the new active. To continue XRF support, you start up the alternate again.

An important consideration is the restart time for a particular region. An application-owning region is usually quick to restart. Terminal-owning regions usually take longer to restart, because of the overhead of establishing the VTAM sessions. Even for a vital region, you might try a restart in place before calling for a takeover of all regions, because the restart even of a vital region might still cause less disruption than a takeover.

You must work out in advance the strategy for each situation. For a speedy restart, your operations should be automated wherever possible. Operators must understand clearly what to do when any type of failure occurs. They must also know what is happening automatically, so that they can take the speediest path to recovery.

4.2.3 Using the overseer

The overseer program can help you to restart regions in place. You can use it to do some of the work that would otherwise have to be done by the operator.

The XRF overseer is supplied as a sample program and associated CICS functions, including an operator interface and macros for identifying CICS systems to it. The overseer runs in its own address space, and can operate only on CICS systems defined with XRF, because it obtains its status information from the CAVM data sets. You can extend the sample program if it does not meet your needs.

The sample overseer can:

° Monitor the status of active and alternate XRF regions, to help the operator to keep track of your systems. You determine how often the overseer checks the status of each system, and the operator can request a display of the information that the overseer collects.

° Restart a failed active region in place. This region would probably be a dependent region, or a single CICS system. Compared with an operator-controlled restart, using the overseer has the advantage that you can automate, and so accelerate, the restart process.

° Restart an alternate region in place, after it has failed, or after the restart in place of its active partner. When an active restarts, it is necessary to start a new alternate to reestablish XRF protection.

In a multi-MVS environment, where you want to restart actives and alternates in place, there must be two overseers, one for each MVS.

The overseer can be particularly useful in a large installation, where you might have many XRF regions that are connected by MRO, with a hierarchy of coordinator, master, and dependent regions.

There is further discussion of the overseer in topic 6.4.

4.3 Multi-MVS configuration using PR/SM ARF

CICS with XRF can take advantage of enhanced function in MVS/ESA SP 4.1 and PR/SM, offering improved availability, response time, ease of operation, and data integrity because:

- The time taken to perform an XRF takeover after an MVS failure can be significantly reduced because there is no longer a need to issue a message to the operator, thus delaying the takeover until the operator responds.

- XRF takeover is further automated and the process is made less susceptible to potential human error.

- Operations are simplified in a multi-MVS/ESA SP 4.1 environment, as there is no longer a need for the operator to synchronize time-of-day (TOD) clocks.

The multi-MVS, single-region XRF configuration, shown in Figure 14 offers increased protection against outages of MVS while allowing you to run both MVS images on the same CEC with the consequential savings in hardware cost.

The active CICS runs in the primary LP (MVS1) and is defined to have most of the CEC resources. The alternate CICS runs in the backup LP (MVS2) on the same CEC (or on the same side of a physically partitioned CEC) and is defined to have the minimum resources required to run an alternate CICS.

The backup system (MVS2), when it detects that the primary system (MVS1) appears to be nonfunctional, processes the XCF PR/SM policy. XCF PR/SM policy allows you to specify that the backup system is to reset the primary LP or deactivate the primary LP and acquire the storage resources defined to the primary LP. ARF enables the backup LP to acquire the storage resources previously owned by the primary LP.

CICS uses MVS services to detect when all this is happening and completes an automatic takeover regardless of the value specified for the TAKEOVR system initialization parameter when the storage reconfiguration is complete.

Figure 14. Multi-MVS configuration using PR/SM ARF

If CICS fails in such an environment, takeover is unlikely to be successful because the backup MVS image is defined to have only the minimum storage resources it needs to run an alternate CICS. As such, TAKEOVR=COMMAND should be used to inhibit takeover in the event of a CICS failure and restart in place (perhaps automated using the overseer) performed instead.

For further details, refer to "Running CICS with XRF in a sysplex" in topic 3.1.6.

4.4 Single-MVS image, single-region XRF configuration

Figure 15 shows that, for CICS outages only, you can increase availability by using XRF in a single-MVS environment. Even if you have more than one MVS image available, you might choose a single-MVS configuration. This might be because of terminal-switching considerations, lack of capacity on the second MVS, or shared DASD limitations.

If you usually run XRF on two MVS images, but one is temporarily unavailable because of maintenance or because it has other work to do, you might choose a single-MVS configuration to provide cover against CICS failures during that period.

With this configuration, you are able to cover yourself against CICS outages, whether they are scheduled, for service or maintenance, or unscheduled, perhaps because of a program error. There is no protection against outages of the CEC, MVS, or VTAM, because these parts of the system are not duplicated. But there are two paths from the network control program through VTAM: one to the active CICS system, and one to the alternate. If the active fails, or if you require a planned takeover, the alternate takes over.

Figure 15. Single-MVS image, single-region XRF configuration

Figure 16. Single-MVS image, multiregion operation XRF configuration

---

4.5 Single-MVS image, MRO XRF configuration

Like the single-MVS image, single-region XRF system, the single-MVS, MRO XRF configuration also improves availability for CICS failures.

For each active region shown in Figure 16 in topic 4.4--terminal, application, and database--there is a corresponding alternate region. Each active-alternate pair has its own CAVM and associated data sets.

Whichever active region fails, its alternate takes over and becomes the new active. The other active regions are unchanged, and the new active reestablishes MRO links with them. The effect observed by the end user depends on which region fails. In this example, failure of the terminal-owning region would result in the effects already described in "Multi-MVS, MRO XRF configuration" in topic 4.2 (and more fully in "The terminal network" in topic 5.0). Failure of other regions is observable at the terminal only if the user is running a transaction that uses the failing region. Such an end user would suffer a transaction failure, but would not lose the session to CICS, nor have to sign on again.

In this sort of configuration, there is no need for the restart in place suggested for multi-MVS configurations.

4.6 Further configurations

"XRF configurations" in topic 4.0 has examined some XRF configurations. Clearly, there are other ways to configure a system. When you are running many systems with XRF, the overseer, described in topic 4.2.3, can give the operator an overview of the active and alternate CICS systems in the XRF complex.

The examples are divided into single- and multi-MVS configurations, but even if you are able to run XRF on two MVS images, there might be some systems that you would prefer to run with the active and alternate in the same MVS.

If you have three MVS images available, you could use the third MVS for a new alternate CICS, if the failure of the first MVS meant that it would be unavailable for an unacceptably long time.

The examples also make a division into MRO and single regions, but you might find that you want to use a combination of MRO and non-MRO XRF regions. You can also have non-XRF regions running with XRF regions in the same MVS image.

In multi-MVS operation, you can place actives and alternates from different CICS systems in the same MVS image.

You could place terminals with backup sessions in one region, and other terminals in a non-XRF region. If you have applications or databases that are rarely used, or applications that rarely fail, they too can be placed in non-XRF regions. This non-XRF region could be a CICS/MVS 2.1 system defined with XRF=NO as a system initialization parameter, or a CICS/OS/VS 1.7 system. A failure in a non-XRF region would then be handled by an emergency restart.

Multiregion operation links can be maintained between the non-XRF region and the active XRF regions. In a single-MVS operation, if a takeover occurs in one of the XRF regions, the MRO link between the new active and the non-XRF region is reestablished. To that non-XRF region, the takeover looks like an emergency restart.

A configuration that you might consider is a communication management configuration (CMC), involving three MVS images: two for the active and alternate CICS systems, and one for a VTAM network owner. A CMC and its possible advantages are described in topic 5.5.

5.0 The terminal network

When you implement XRF, there are implications for your existing terminal network. The information that follows is to help you organize your terminals in an XRF environment.

Any terminal that you currently use with CICS can be used in an XRF environment. Remote SNA VTAM terminals, connected through a boundary network node (BNN) with the network control program (NCP) in a 3745, 3725, or 3720 communication controller (with the appropriate VTAM and NCP levels), are XRF-capable (class 1) terminals. These terminals can show the most improvement in restart time. These terminals have backup sessions to the alternate CICS. But XRF offers benefits to all terminals. Terminals that cannot have backup sessions may also experience a faster restart. This is because the alternate can recognize failure earlier, and because it tracks the installed, logged-on, or logged-off state of other VTAM terminals and attempts to reestablish sessions after takeover.

For XRF-capable terminals, the fundamental networking feature of XRF is **session switching**. If the active CICS system in a multi-MVS environment suffers a failure in MVS, VTAM, CICS, or (in a two-CEC environment) in the CEC, so that the CICS system is no longer available, the alternate CICS is notified of the failure. It requests that the affected terminals are switched to the alternate system where they continue the application-to-terminal sessions. The active and alternate systems are connected by the common switch point at the NCP in the 3745/3725/3720. As long as there is an appropriate communication controller at the BNN, terminal sessions may pass through an outboard IBM 3705 communication controller and still have backup capability.

Each terminal can have a working session with only one CICS system. However, the active CICS system notifies its alternate of all its sessions (except those defined with RECOVOPTION(NONE)). If the session is with an XRF-capable terminal, the alternate CICS asks VTAM to establish a backup session for the terminal to the BNN communication controller. When the alternate takes over, it issues a:

```
VTAM SESSIONC CONTROL=SWITCH
```

command to switch the terminal-to-CICS session to the alternate. This is much faster than the alternate reacquiring the terminals. The backup session to the terminal means that the end user does not have to log on again to VTAM. Transactions that are in flight at the point of takeover are backed out by CICS and must be reentered by the end user (or by your normal restart practices). However, depending on the signon options set, end users do not normally have to sign on to CICS again.

Figure 21 in topic 5.6 shows a representative sequence of SNA flows for an XRF-capable terminal.

Before specific terminal types and levels of service are discussed, note that there are many factors that can affect the performance of a terminal at takeover, as follows:

- The type of terminal and its access method
- The total number of terminals connected
- What the end user is doing at the time of takeover
- Whether the terminal has signon security
- The signon options set
- The type of failure of the active CICS system
- Whether the terminal has to be physically switched to a second MVS image
- How the terminal is defined by the systems programmer.

Subtopics:

5.1 VTAM and NCP considerations for active and alternate

Users are unaware of being attached to the active side of an XRF pair. They have an image of a single system processing the workload. So it should be irrelevant to them which system is the active. Similarly, the XRF-capable terminal is not aware that a second session has been created from the alternate to the BNN communication controller.

The active and alternate share a common **generic applid**. In addition, each active and alternate has a unique **specific applid** to identify itself to VTAM. The end user is only aware of the generic applid used at logon. For existing systems that you convert to XRF, you could retain the applid that is familiar to the end user as the generic applid, and have two new names, probably based on the generic applid, as the specific applids.

For more VTAM information, you should consult the *VTAM Installation and Resource Definition* manual and the *VTAM Operation* manual. This is particularly important if you are not accustomed to multi-MVS network environments.

The generic applid is known in VTAM terms as the USERVAR; the specific applid is the VTAM application id. The generic applid is used by CICS for many purposes: for example, it indicates the active-alternate pairing to the CAVM; it is also used for IRC, for DBRC, and for IRLM.

Subtopics:

5.1.1 Defining the applids

The active and alternate are defined as specific applids to VTAM by VTAM APPL definition statements; for example:

```
CICS1 APPL AUTH=(ACQ),HAVAIL=YES

CICS2 APPL AUTH=(ACQ),HAVAIL=YES
```

The first part of the APPL statement defines to VTAM the specific applids

(known to VTAM as the application ids). The HAVAIL option enables the defined CICS system to support alternate sessions and session switching.

The generic and specific applids have to be defined to CICS using the APPLID system initialization parameter. See topic 6.1.

5.1.2 Controlling the use of the applids by USERVAR

Note: The description under this heading, "USERVAR propagation to remote VTAMs" in topic 5.1.2.1, and "Transferring a terminal session to the active" in topic 5.1.2.2 does not apply to VTAM releases prior to Version 3.2 (with the appropriate PTF).

To control these generic and specific applids, XRF makes use of the VTAM USERVAR facility. VTAM maintains a USERVAR table which records the relationship between the generic and specific applids. The entries in the USERVAR table are built dynamically by VTAM. The generic and specific applids are added to the table by VTAM when the first MVS MODIFY USERVAR command is issued from the first active CICS. The specific applid may subsequently be changed dynamically at a takeover.

When a terminal logs on, the "logon message", which refers to the generic applid, is interpreted as a logon request to the application whose specific applid is contained in the USERVAR. In this way, the USERVAR table relates the generic applid (which does not change) to the specific applid of the current active, and VTAM can identify the CICS system to which the terminal's active session should be connected.

Figure 17 shows a set of definitions, with CICS1 as the active system and VTAM1 as the network owner. At startup, the active uses the:

```
MVS MODIFY vtamname,USERVAR
```

command to set its specific applid (CICS1 in the figure) in the VTAM USERVAR table. The USERVAR table contains an entry like this:

```
CICS, CICS1
```

which ensures that logons are directed to the current active. The TYPE=DYNAMIC parameter (the default) specifies that this USERVAR entry is for an XRF system that is likely to change its specific applid periodically.

The user's logon message "PRODCICS" is associated with the correct specific applid by VTAM's USERVAR processing.

At the start of a takeover, the alternate changes the setting of the USERVAR to its own specific applid, so that logons to a failing active are stopped as soon as possible. It issues a second:

```
MVS MODIFY vtamname,USERVAR
```

command when it issues the:

```
SET LOGON START
```

command, which tells VTAM that the new CICS system is ready to accept logons.

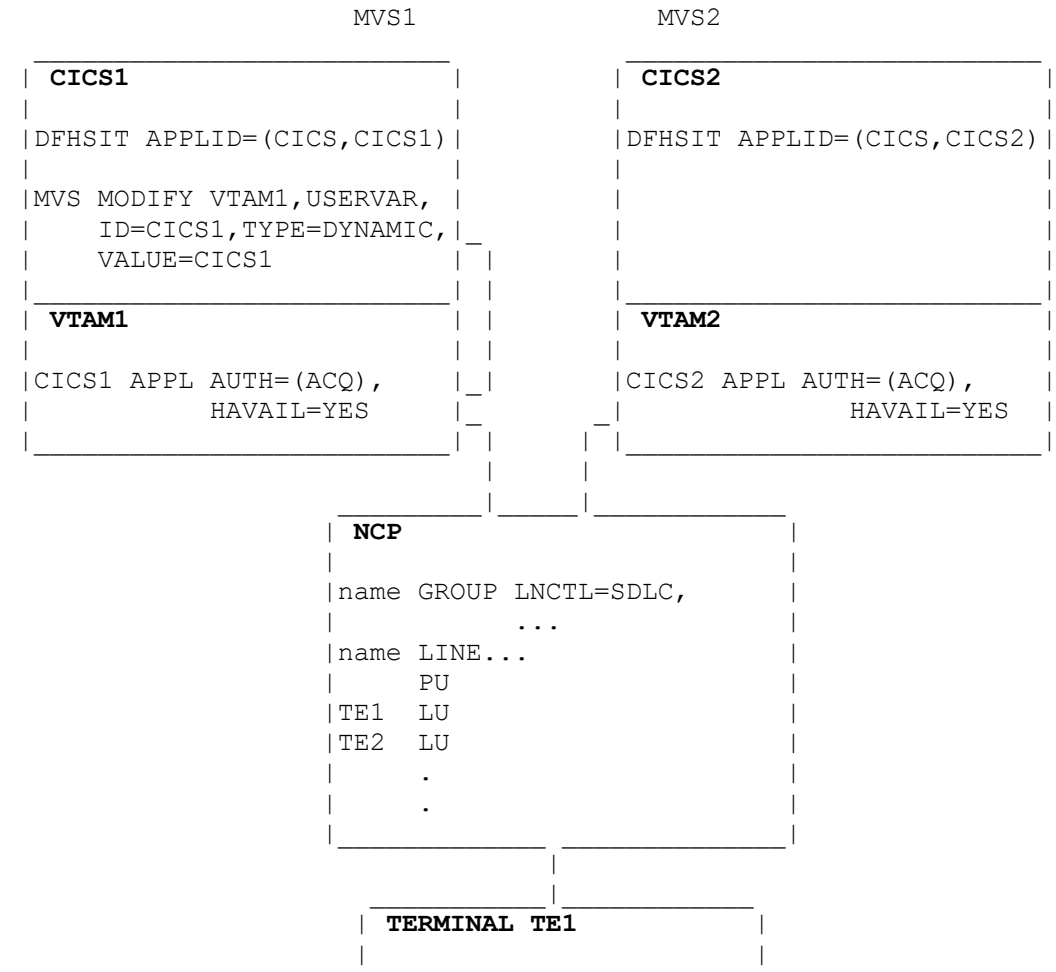

Figure 17. Logging on to the active

Subtopics:

- 5.1.2.1 USERVAR propagation to remote VTAMs
- 5.1.2.2 Transferring a terminal session to the active

5.1.2.1 USERVAR propagation to remote VTAMs

The USERVAR modified by an MVS MODIFY USERVAR command issued by an active is known as a user-managed USERVAR. USERVARs in remote VTAMs that communicate with the VTAM that is local to the XRF system can be modified by VTAM with no involvement by CICS. VTAM does this in response to a change in the user-managed USERVAR. These remote USERVARs are known as automatic USERVARs.

Unless you have other, non-XRF, uses for USERVARs that conflict with such USERVAR processing, you are recommended to allow VTAM to manage this propagation of USERVARs. If you leave the operator to propagate the USERVAR, and there is a delay before the operator issues the command, some new users cannot log on to CICS during that delay.

There are no XRF-specific changes for the SNA unformatted system services (USS) tables. You may use SNA USS tables with XRF-capable terminal definitions.

5.1.2.2 Transferring a terminal session to the active

With VTAM Version 3.2 (with the appropriate PTF) in all the relevant regions, you can transfer a terminal session to an active using the generic applid in the VTAM CLSDST PASS command, as follows:

```
EXEC CICS ISSUE PASS LUNAME(generic applid) .........
```

You do not need code to establish the specific applid of the active. An application that already contains such code continues to work unchanged.

5.1.3 Ownership of the network

In an XRF environment, terminals may be owned by a VTAM in a different MVS image from that of the active CICS system. Because of this, terminals must be defined to be **cross-domain**, which means that:

- Terminals may log on to the active (CICS1 in Figure 18)
- CICS1 may acquire terminals
- The alternate (CICS2 in the figure) may have backup sessions established
- After takeover, CICS2 may acquire terminals
- New terminals may log on to CICS2.

In the example in Figure 18, there are the following considerations:

- The ownership of the network by the VTAM in MVS1
- The cross-domain definitions of the network to MVS2
- The local definition of application CICS1 in MVS1
- The cross-domain definition of application CICS1 in MVS2
- The local definition of application CICS2 in MVS2
- The cross-domain definition of application CICS2 in MVS1.

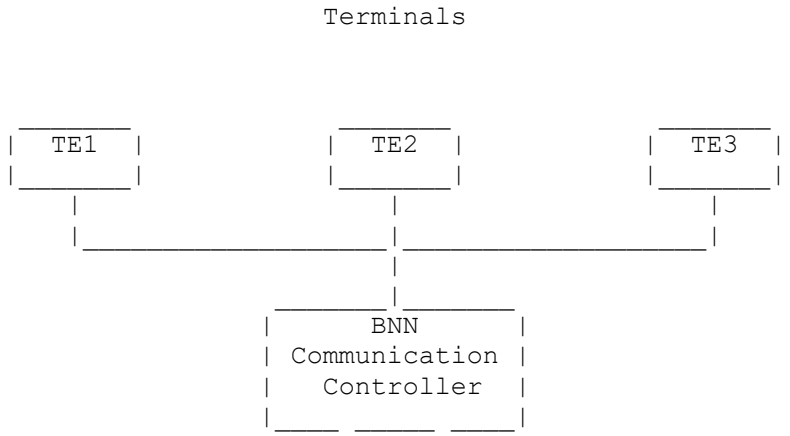

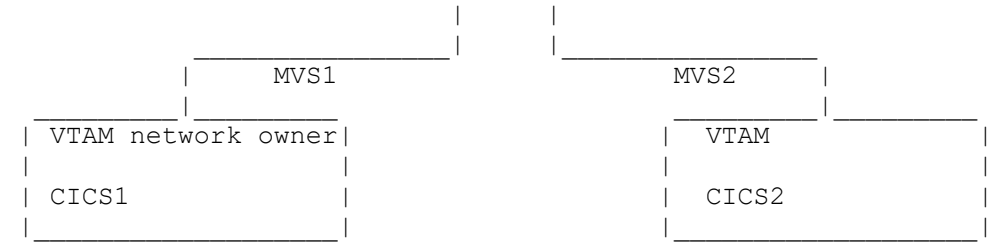

Figure 18. VTAM network ownership

The following partial NCP definition defines MVS1 as the network owner, and the terminals in that network:

```
             BUILD.......,BACKUP=350
             GROUP....,LNCTL=SDLC,.....,OWNER=MVS1
             LINE...
             PU...
  TE1        LU...
  TE2        LU...
  TE3        LU...
```

The following partial definition defines CICS1 on MVS1, with a cross-domain definition for CICS2:

```
  CICS1      APPL....HAVAIL=YES
             VBUILD TYPE=CDRSC


  CICS2      CDRSC CDRM=MVS2
```

("CDRSC" is the cross-domain resource, and "CDRM" is the cross-domain resource manager.)

Here is a cross-domain definition in MVS2 for the terminals:

```
             VBUILD TYPE=CDRSC
  TE1        CDRSC CDRM=MVS1
  TE2        CDRSC CDRM=MVS1
  TE3        CDRSC CDRM=MVS1
```

The following partial definition defines CICS2 to run on MVS2, with a cross-domain definition for CICS1:

```
CICS2       APPL....HAVAIL=YES
            VBUILD TYPE=CDRSC

CICS1       CDRSC CDRM=MVS1
```

The advantages of assigning ownership of the network to a third MVS image are described in "Advantages of a CMC configuration" in topic 5.5. The way that ownership of the network affects terminals at takeover is shown in Table 2 in topic 5.5.

For terminals owned by VTAMs other than the VTAM for the active, the use of the automatic USERVAR for USERVAR propagation is described in "USERVAR propagation to remote VTAMs" in topic 5.1.2.1.

5.1.4 Preparing NCP for XRF

To define NCP support for class 1 terminals, code the XRF-related operands on the BUILD definition statement in the NCP generation deck. You should also review other operands on NCP definition statements.

The BACKUP operand on the BUILD definition statement specifies how many class 1 terminals you expect the NCP to support. This information allows NCP to reserve the right amount of storage for control blocks for primary and backup sessions. For example, if you expect to have 350 class 1 terminals, code:

```
BUILD ...,BACKUP=350
```

You should consider carefully the number of backup sessions you define. If you fail to reserve enough storage for your class 1 terminals, VTAM cannot open either active or backup sessions for any additional terminals above the limit you specified. Each time that happens, an error message is issued to the CSMT log, and CICS then tries to open a class 2 session for the terminal. However, each additional backup session takes up storage in the NCP. Do not define more backup sessions than you need. You should consult your IBM representative for advice about the storage requirements of backup sessions.

The PAUSE operand on the LINE definition statement sets the polling values

for the NCP. The switching of sessions from primary to backup at takeover places a large demand on the NCP. To decrease the time it takes the NCP to complete session switching, increase the polling values on the PAUSE operand. For information about changing the polling values, see the *NCP, SSP and EP Resource Definition Reference* manual, or the equivalent manual for your communication controller.

5.2 Levels of terminal support

A typical CICS installation may have a wide range of terminal connections in its network, including VTAM and non-VTAM, local, and remote devices. The full list of IBM terminals and devices that can be used with CICS is in the *CICS/ESA Facilities and Planning Guide*.

Table 1 describes the three classes of terminals in an XRF environment, how XRF supports them, and what the user can expect at a takeover.

Table 1. Terminal support

Terminal class	**How XRF supports terminals at logon**	**How XRF supports terminals at takeover**	**How takeover affects terminal user**
Class 1	CICS opens active session from active, and backup session from alternate.	Alternate issues SWITCH command to take over session.	Service is almost continuous.
Class 2	No change to normal CICS support. (No backup session established.)	Alternate tries to reestablish session.	Brief delay in service while alternate acquires session.
Class 3	No change to normal CICS support. (No backup session established.)	No change to normal CICS emergency restart procedures.	User loses service. Operator must reestablish session.

In this table, the word "terminal" does not just describe a simple terminal device, but also describes a component of a terminal system,

including a programmable controller and its attached operator terminals, printers, and remote subsystems.

The RECOVOPTION and RECOVNOTIFY terminal definition keywords and the signon options modify the service that CICS gives to each terminal, but initially the default values of these keywords are assumed. The defaults give each terminal the best service that its characteristics allow. The effects of using alternative settings of the terminal definition keywords, and of signon security, are discussed under "Defining the recovery process" in topic 5.3.

Subtopics:

5.2.1 Class 1 terminals

To take full advantage of the support available in XRF, the terminals in your complex that need high availability should be class 1. These are the terminals that are described as being "XRF-capable".

A class 1 terminal:

- Is a VTAM terminal that uses SNA protocols
- Is controlled by release levels of VTAM and NCP that support XRF
- Has an IBM 3745/3725/3720 Communication Controller as its BNN.

CICS does not determine the XRF capability of a terminal. VTAM tells CICS of a terminal's capability when the terminal logs on. In the VTAM APPL definition for the CICS system, you must code HAVAIL=YES to obtain the capability indicator.

For class 1 terminals, the active session is matched by a corresponding backup session to the alternate. At takeover, the alternate tells VTAM to switch the terminal sessions from the old active to the new active. Under

certain circumstances, this switch could be transparent to the end user. However, if the takeover occurs during transaction input, the end user might have to reenter the last transaction.

A terminal that is executing a transaction at takeover (a "busy" terminal) must have its session state recovered by CICS before it continues operation. If the terminal has only sent data through the NCP, and that data has not yet been processed, the terminal is still classed as busy. CICS recovers the session state as quickly as it can, depending on the exact state of the session at takeover, the session's characteristics, and the setting of the RECOVOPTION keyword.

CICS has the choice of:

- Sending an end-bracket indicator to close the current bracket.
- Issuing a CLEAR command to reset the conversation.
- Issuing an UNBIND (followed by a simlogon) to reset the session. In this last case, the recovery is similar to that for a class 2 terminal. See "Class 2 terminals" in topic 5.2.2.

The user might be aware of loss of service during a takeover, but the terminal will remain in session at all times, unless the session is reset by issuing an UNBIND command.

If the takeover occurs during the catchup process, it is possible that the end user of an autoinstalled terminal might have to log on again. For more information on how to control this, see the AIRDELAY parameter described in topic 6.1.1.

Depending on the options you choose, end users of terminals with signon security do not have to sign on after a takeover, because signon information is passed from the active to the alternate. If you choose to tell end users of the takeover, by using the RECOVNOTIFY option, the system can send a message to the terminal. You could use this message to explain to end users that they might have to reenter the last transaction. If the terminal is unbound and then rebound at takeover, perhaps because of the setting of RECOVOPTION, the "good morning" message (if defined) is sent instead of the recovery message. For more information, see "Defining the recovery process" in topic 5.3.

5.2.2 Class 2 terminals

A Class 2 terminal is:

- A remote VTAM terminal that is not connected through a BNN communication controller and its NCP, or through the appropriate level of VTAM.

- A locally attached VTAM terminal or a VTAM non-SNA terminal, including a BSC 3270 terminal. In a multi-MVS environment, locally attached VTAM terminals qualify as class 2 if they are definable as cross-domain resources to both active and alternate, and able to connect to the alternate after takeover. For local terminals, see note 1 at the end of this section.

- A BSC 3270 terminal attached to a BNN communication controller.

- A terminal supported by the network terminal option (NTO) or network routing facility (NRF).

- A VTAM terminal using session-level cryptography.

- An LU6.1 or APPC ISC system.

Terminals in this group are not XRF-capable because they are not routed through a BNN communication controller and an NCP. They do not have backup sessions to the alternate. However, class 2 local and remote VTAM terminals can still benefit from an XRF environment, through the tracking procedure. The alternate **tracks** the installed, logged-on, or logged-off status of all VTAM terminals and sessions as they are acquired or released. Terminals that are already logged on and active on the active CICS when the alternate is started are catered for by the catchup process. If RECOVOPTION(NONE) has been specified for a terminal, that terminal is not tracked, and it becomes a class 3 terminal.

After a takeover, the new active tries to establish new sessions for terminals that it tracked when they were in session with the old active. This reconnection may not succeed immediately because you may need to transfer the connection of some of these terminals manually from one MVS to the other. So CICS tries again at intervals of 1, 2, 4, and 8 minutes after the first attempt. The timing of the first attempt depends on the number of backup sessions requiring switching, and on the value set by the AUTCONN system initialization parameter. After the reconnection transaction has finished, you either use operator intervention to reacquire remaining sessions, or the users themselves log on again. This situation could arise if the VTAM that owns the network has failed, and it takes more than 8 minutes to restart it. In that case, all terminals that are normally reconnected will require some sort of intervention.

If the network owner has not failed, end users might experience a short interruption in service, and the takeover has the appearance of an emergency restart. If the session is successfully reestablished, end users of such terminals do not have to log on again, nor, depending on the options set, do they have to sign on to CICS again. However, unlike class 1 terminals, class 2 terminals do not receive specific messages after takeover. Instead, the "good morning" message is displayed. The end user must be aware that logon or signon might not be necessary. For more information about the options that control signon, see "Signon after takeover" in topic 5.3.3.

You must consider how your operations staff will transfer class 2 terminals from one MVS to another in a multi-MVS environment. In a single-MVS system, this is not a problem, but you might still need procedures for connecting class 2 terminals to a new active after a takeover.

Notes:

1. There is a technique that allows local terminals to be reconnected to the new active, but it involves you in additional programming. If local terminals are attached to an IBM 3814 communication controller and a multisystem configuration manager (MSCM), you can write a program to provide the physical transfer from the active to the alternate. If you add to the program an operator interface that could be driven from the CLT, the operator is not involved in the physical switching. If you already have terminals attached through a 3814 and MSCM, you might be interested in this form of switching. For more information about MSCM, see the *Multisystem Configuration Manager Programming* manual.

2. It is possible that class 2 terminals will not be reacquired after a takeover if you have the combination of (1) long-running tasks updating recoverable resources without syncpointing, and (2) a high value in the AKPFREQ system initialization parameter. With this combination, a terminal or session that is installed, subsequently reinstalled, and then acquired, might not be reacquired after a takeover. If this happens, you should ensure that long-running tasks take regular syncpoints, and you should set a lower AKPFREQ value.

3. A takeover initiated by CEMT PERFORM SHUTDOWN TAKEOVER is different from other forms of takeover, and might affect the recovery of class 2 terminals on subsequent takeovers. For more information, see topic 6.5.2.

5.2.3 Class 3 terminals

A class 3 terminal is a terminal that is not tracked, because it is a:

- VTAM terminal with the recovery option suppressed; for this class of terminal, the installed, logged-on or logged-off state is not tracked. The end user has to log on again when service is reestablished.

- Non-VTAM terminal, such as a TCAM terminal.

In a multi-MVS environment, after a takeover, end users of class 3 terminals can communicate with the new active only after the operator has created a physical path to it.

There is no specific XRF support for TCAM terminals. They are supported as they are in a non-XRF environment. Users of these terminals log on to the specific applid. You might consider isolating all your TCAM terminals into a separate non-XRF CICS region that is connected by MRO to the other CICS regions. In a single-MVS environment, the failure of an XRF region does not affect a region that isolates TCAM terminals. After the takeover of the failed region, the MRO connection to the TCAM region is reestablished.

To the end user of a class 3 terminal, a takeover has the appearance of an emergency restart.

5.3 Defining the recovery process

You can use RDO to define the recovery process for each terminal, by the RECOVOPTION and RECOVNOTIFY keywords, or their macro equivalents. For reference information about these keywords, see the *CICS/ESA Resource Definition (Online)* manual. The options that control whether or not an end user has to sign on again after a takeover are described in "Signon after takeover" in topic 5.3.3.

Subtopics:

- 5.3.3 Signon after takeover
-

5.3.1 Using the RECOVOPTION keyword

The RECOVOPTION keyword gives you control over the way the alternate system tracks and recovers the session state of a terminal. The default action is to allow CICS to determine the most efficient way of recovering the session after takeover, based on the particular type of terminal and its activity at takeover.

By specifying either CLEARCONV or RELEASESESS for the RECOVOPTION keyword, you can force CICS to use a more drastic way of recovering sessions that are busy at takeover. This could be desirable if you have specialist knowledge of a terminal, and believe that it will not respond correctly to receiving an unpredictable flow that the alternate CICS might send to recover it. However, if the option is not suitable for a particular terminal, CICS will override it.

Coding RECOVOPTION(CLEARCONV) prevents CICS from sending just an end-bracket indicator to terminate the current bracket for a terminal that is active at takeover. For terminals with session characteristics that support the VTAM SESSIONC CONTROL=CLEAR command, the alternate system will issue the CLEAR command under these circumstances. If the session characteristics show that the terminal cannot support a clear command, then CICS will unbind and simlogon the session.

RECOVOPTION(RELEASESESS) restricts the alternate to using the unbind and simlogon option to recover active sessions at takeover.

RECOVOPTION(UNCONDREL) is a very drastic form of recovery at takeover. It forces the alternate to unbind and simlogon the terminal after takeover regardless of the state of the session. It differs from the RELEASESESS option, because that option is invoked only if the terminal is found to be active at takeover. It would be useful in cases where the terminal needs to know which CICS system it is connected to, so that a transparent takeover would be unacceptable.

Notes:

1. For both UNCONDREL (which means that any session is unbound) and RELEASESESS (which means that only active sessions are unbound) the RECOVNOTIFY message or transaction is not run. The "good morning" message (if defined) is sent instead.

2. If the VTAM network owner fails, any session that is to be unbound and then rebound will only be unbound. It cannot be rebound until VTAM network ownership is reestablished.

RECOVOPTION(NONE) may be used to prevent the alternate system from tracking the installed, logged-on or logged-off status of the terminal in the active system. It may be used for any class of terminal. After takeover, the end user or the operator will have to initiate the session.

5.3.2 Using the RECOVNOTIFY keyword

The RECOVNOTIFY keyword applies only to class 1 terminals. The takeover may be completely transparent to the end user. In many cases this is what you want. In other cases it may be desirable to inform the end user after takeover that there has been a system failure, but that service has now been restored.

The default operand for RECOVNOTIFY is NONE, which means that the end user will not receive any notification that a takeover has occurred. The other operands available are MESSAGE and TRANSACTION. These are used to inform end users that a takeover is complete, so that they can check that the last transaction terminated normally. Also, they are used to tell end users whether or not they should sign on again after the takeover.

The MESSAGE option results in a simple message to inform the end user that the system has recovered. It also reminds the end user to initiate any recovery procedures before continuing processing. The message that is sent depends on whether or not the end user has to sign on again:

```
CICS/ESA has recovered after a system failure
Execute recovery procedures
Please sign on
```

or:

```
CICS/ESA has recovered after a system failure
Execute recovery procedures
Already signed on
```

You can modify the text of either message by replacing the IBM-supplied map set. The map set name is DFHXMSG, and it contains two messages in the form required by BMS, under the names DFHXRC1 and DFHXRC2. You can edit those messages. Alternatively, you can specify a different map set in the

node error program. As you increase the size of the message, you might increase the takeover time. For guidance information and definitive product-sensitive programming interface information about node error programs, see the *CICS/ESA Customization Guide*.

The TRANSACTION option initiates a transaction in a similar way to the "good morning" message transaction. This is more versatile than the MESSAGE option, but it takes more processing and it can slow down takeover. The default transaction ID is CSGM. However, you can use the system initialization RMTRAN parameter to request that a transaction of your choice be run at this time. In this way, specifying that a transaction should run allows wider control compared to the MESSAGE option. Like the MESSAGE option, the TRANSACTION option allows you to use two transactions: one for end users who do not have to sign on again, and one for those who do have to sign on.

Note that this keyword does not take effect for:

- Terminals that are unbound and rebound
- Terminals that are not cleaned up by the system and are left unbound.

These terminals receive the "good morning" message (if defined).

5.3.3 Signon after takeover

Users of class 1 and class 2 terminals do not normally have to sign on after a takeover has switched the terminal session to a new active. This is made possible by the transfer of signon security information from the active to the alternate through the message data set.

There is a hierarchy to control whether or not particular terminals, or sets of terminals, or all terminals, have to be signed on again. It is also possible to sign off terminals if the takeover takes more than a specified time.

The three ways are:

XRFSOFF=FORCE|**NOFORCE** system initialization parameter
: If you specify *FORCE*, **all** end users have to sign on again after a takeover. *FORCE* always takes precedence over the same operand in both the CEDA transaction and the DFHSNT macro. If you specify *NOFORCE*, the CEDA transaction and the DFHSNT macro can be used to make smaller groups of terminals sign on again. The system initialization parameters are further described in topic 6.1.

CEDA DEFINE TYPETERM XRFSIGNOFF(FORCE|**NOFORCE**)
: You use this transaction to define the signon characteristics of a set of terminals. You might choose to force the sign off of a set of terminals if they are located in a security-sensitive area. An SNTTE entry set to *NOFORCE* for an individual terminal has no effect if the *TYPETERM* definition for the terminal is set to *FORCE*, but if you opt for a *TYPETERM* definition of *NOFORCE*, you can then use the SNTTE entry to force a terminal or group of terminals to be signed off.

DFHSNT TYPE=ENTRY, XRFSOFF=FORCE|**NOFORCE**
: The lowest level at which you can force a terminal to be signed off is in its SNTTE entry. One SNTTE entry could apply to a number of terminals.

With RACF 1.9, the operator information may be held in the RACF CICS segment, in the form XRFSOFF(FORCE|NOFORCE). For more information, see the *CICS/ESA CICS-RACF Security Guide*.

So, to summarize, there are three levels at which terminals may be forced to sign off at takeover and end users have to sign on again. This is shown in Figure 19.

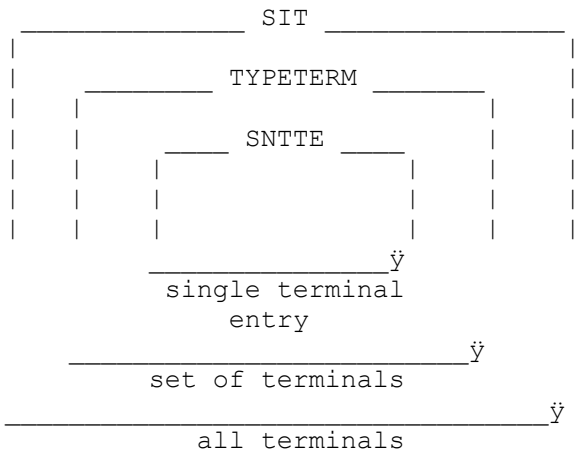

Figure 19. Signoff levels

In addition to these signon options, there is also the
XRFSTME=decimal-value|5 system initialization time-out parameter, which enables you to sign off users if the takeover takes more than the specified time in minutes: For this parameter, takeover time is defined as the time between the initiation of the takeover to the time a user is able to input data again. So, if takeover takes 4 minutes, and the

default is set, end users are still signed on. If the takeover takes 6 minutes, end users are signed off. Note that this option applies only to those terminals that have the SNT *TIMEOUT* option set. Without that option, the end user may still be signed on after a takeover that takes longer than the period set by the XRFSTME option.

For non-XRF-capable terminals, you must consider the effect of the system initialization *AUTCONN* parameter. *AUTCONN* delays the reconnection of non-XRF-capable terminals (see "Starting the alternate" in topic 6.1.2), so you might choose to extend the XRFSTME value to allow these terminals to be reconnected and remain signed on.

Note: When a CEMT PERFORM SECURITY (REBUILD) command is issued to the
active CICS, it uses the message data set to tell the alternate that the
RACF resource profiles have been rebuilt. RACF definitions must be the same for the active and alternate. If the active fails at the time of the rebuild, a message warns the operator if the rebuild has not been successful.

5.4 Specific session types

Generally, the way in which sessions are acquired and taken over in an XRF environment is transparent to the terminal. However, you might find the information in the following sections helpful when considering the settings of system parameters.

Subtopics:

- 5.4.1 LUTYPE6 ISC application-to-application sessions
- 5.4.2 Programmable terminals
- 5.4.3 Pipeline logical units

5.4.1 LUTYPE6 ISC application-to-application sessions

If you run an XRF system against VTAM 3.2 (with the appropriate PTF), VTAM USERVAR support extends to subsystems that communicate with an active

through LUTYPE6.1 or APPC ISC links. Application programs can initiate the session to the active using the generic applid. The INQUIRE USERVAR command, if used, returns the name given as input.

If you have an earlier level of VTAM, the subsystems must first determine which of the two CICS systems is the active by issuing the INQUIRE USERVAR command to VTAM. This returns the specific applid that has been set in that user variable.

Subtopics:

5.4.1.1 CICS-to-CICS and CICS-to-IMS communication

An active can communicate, using ISC, with:

- Another CICS/ESA 3.3 active
- A CICS/ESA 3.3 system specified with XRF=NO
- A CICS/OS/VS 1.7 system with the PTF that enables ISC to an active
- A CICS/MVS 2.1 system
- A CICS/ESA 3.1.1 system
- A CICS/ESA 3.2.1 system
- A CICS/DOS/VS 1.7 system
- A CICS/VSE (*) 2.1 system

- IMS/VS Version 2

- IMS/ESA (*) Version 3 or Version 4.

If you are using a level of VTAM earlier than 3.2, CICS and IMS take care of the USERVAR processing.

(*) IBM Trademark. For a list of trademarks, see topic FRONT_1.

5.4.1.2 Bind format

The format of the bind that the active sends to the terminal or secondary logical unit (SLU) contains the normal primary logical unit (PLU) name field. The contents of this name field depend on whether the PLU or the SLU initiated the session; that is, whether the terminal user logged on to CICS, or CICS acquired the terminal.

- If the PLU initiated the session, the field contains the PLU name. This will be the specific applid of the CICS system.

- If the SLU issued the INITSELF, the name field contains the uninterpreted name as carried in that RU. This is the generic applid of the CICS system.

This is no different from what happens in the normal SNA environment, but in an XRF environment it may become significant if the SLU examines this name field. If the SLU relies on the host to initiate the session (using the RDO attribute AUTOCONNECT(YES), for example), the contents of this name field vary according to which system is the active.

APPC architecture has defined the structure of the bind user data fields. One of these user data fields is reserved for the PLU name, and CICS uses this field to pass its generic name. The APPC terminal should examine this user data PLU name field to determine the name of the LU requesting the session. Thus APPC terminals will find a common PLU name regardless of which CICS is the active system, and so these terminals can connect directly to CICS.

5.4.2 Programmable terminals

You may have programmable, or "intelligent", LU0 terminals that examine the bind parameters they receive from CICS. As discussed above, if such terminals examine the PLU name in the bind, their programs might need modification to accept a bind from both the active and the alternate.

5.4.3 Pipeline logical units

Pipeline logical units cannot be supported in the same manner as other class 1 XRF-capable terminals. Pipeline LUs use a restricted set of communication protocols and cannot use the normal CICS restart facilities.

If the network sees a pipeline LU as a class 1 terminal, the alternate still establishes the backup session and issues the SWITCH command to acquire the session. However, after the SWITCH, the alternate issues CLSDST to send unbind to the terminal, so that the terminal can reset its state.

In effect, therefore, pipeline LUs are handled as class 1 terminals with RECOVOPTION=UNCONDREL defined.

5.5 Advantages of a CMC configuration

You might consider assigning ownership of class 1 terminals to a separate VTAM communication management configuration (CMC). This involves a third MVS image that has ownership of the VTAM terminals, as shown in Figure 20.

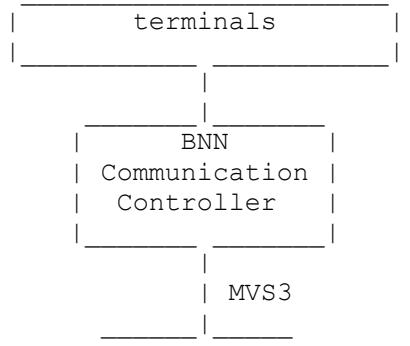

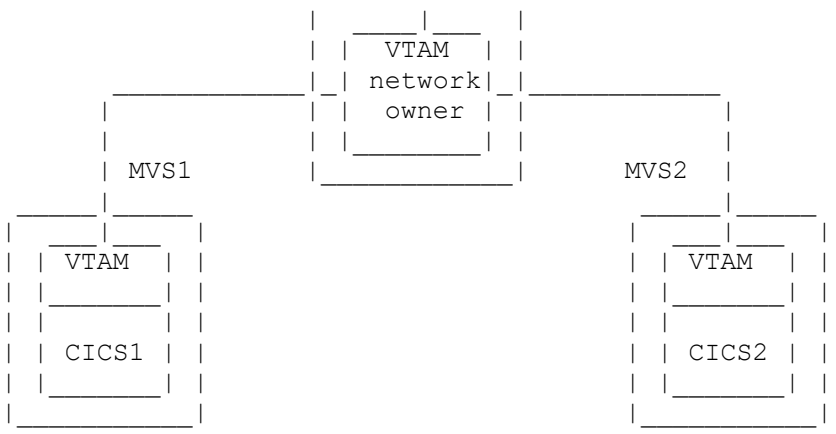

Figure 20. A CMC configuration

The advantage of this arrangement is that, if the MVS of the active or alternate fails, VTAM ownership is not lost. If the VTAM owner does fail, new end users cannot log on until network ownership has been reestablished, probably by transferring to the other MVS. Existing end users are unaffected. If the active's MVS fails, and VTAM ownership is lost, tracked terminals cannot be reconnected after takeover until ownership has been reestablished.

If you do not have a CMC, assign the ownership of the class 1 terminals to either of the VTAMs in the complex, and do not change the ownership after a takeover unless you are restarting VTAM. Where you place the ownership of VTAM might have an effect on the takeover time following a VTAM failure.

For more information about reestablishing or changing ownership of terminals, see the *VTAM Operation* manual.

Table 2 summarizes the ways in which VTAM ownership may affect class 1 and class 2 terminals when there is a failure.

Table 2. VTAM ownership and terminal failure

Owner	**Cause of failure**	**Take-over?**	**Effect on class 1 terminals**	**Effect on class 2 terminals**
Active	Active VTAM fails	Yes	Leased lines switched; dial lines disrupted. See notes 1 and 2.	Sessions terminate; automatic session restart after owner reestablished.

				See notes 1 and 2.
Alternate	Active VTAM fails	Yes	All terminals switched.	Sessions terminate; automatic session restart.
Alternate	Alternate VTAM fails	No	Active sessions maintained; alternate CICS terminates. See note 1.	Session maintained; alternate CICS terminates. See note 1.
CMC	Active or alternate VTAM fails	No	All sessions maintained or switched.	Automatic session restart after takeover.

Notes:

1. No new logons are possible until the VTAM network owner has been reestablished.

2. Terminals that are unbound will not be rebound until the VTAM network owner has been reestablished.

5.6 XRF SNA flows

Figure 21 shows a representative sequence of SNA flows for:

- A class 1 terminal logging on to the active
- The backup session being established
- The alternate taking over after a failure of the active.

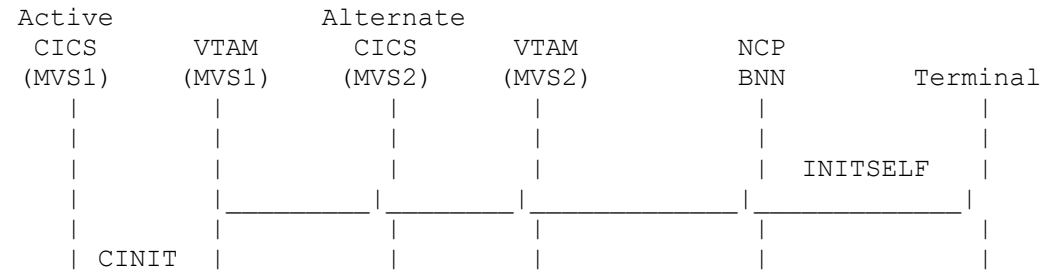

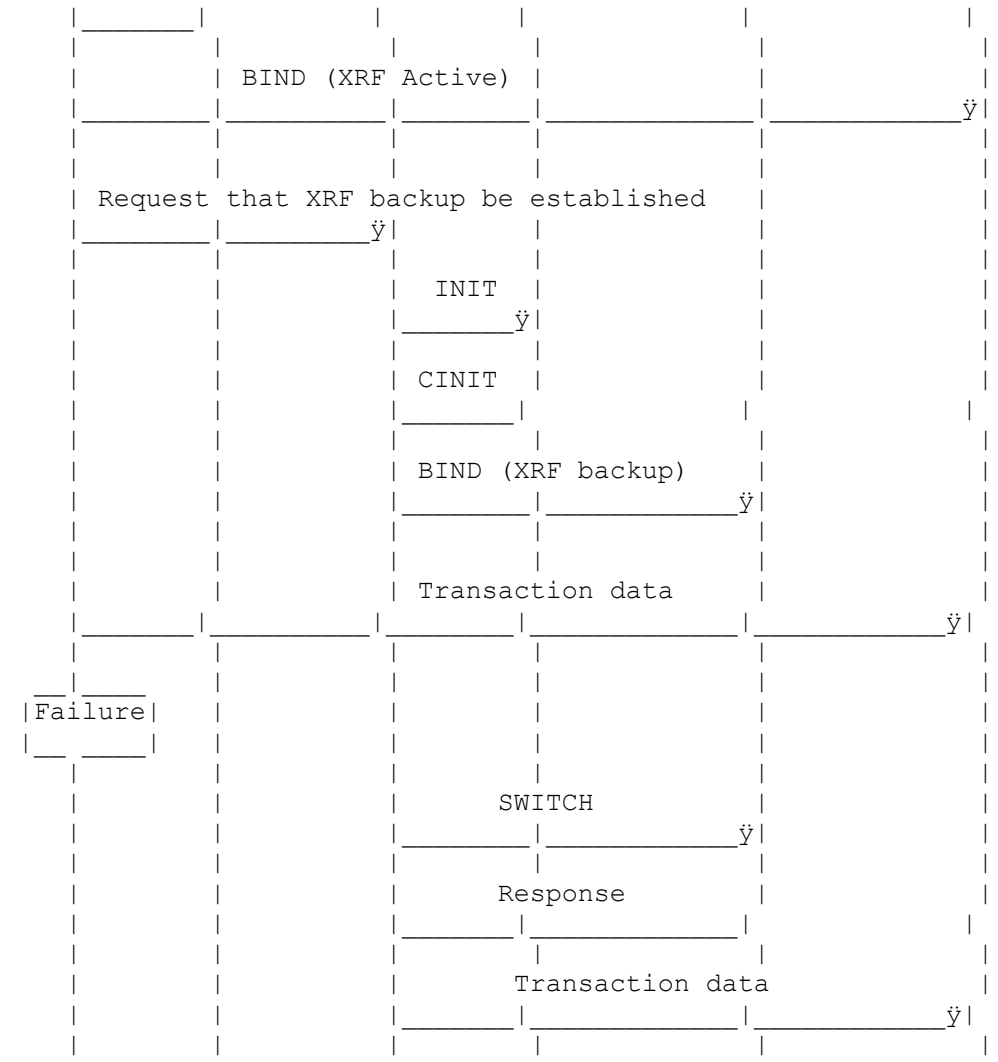

Figure 21. Abbreviated XRF SNA flows

6.0 Defining CICS for XRF

This topic gives you the information you need to define an active and alternate pair and the takeover appropriate for them. To create a system (which could be made up of MRO-connected regions), you combine the functions described in the following sections:

- "System initialization parameters" in topic 6.1
- "Command list table (CLT)" in topic 6.2
- "User exit for VTAM failure" in topic 6.3

- "The overseer" in topic 6.4
- "Supplied transactions for controlling the alternate" in topic 6.5
- "Defining your XCF PR/SM policy" in topic 6.6
- "Sharing data sets" in topic 6.7
- "Storage protection considerations" in topic 6.8.

For reference information for tables, see the *CICS/ESA Resource Definition (Macro)* manual. For system initialization, see the *CICS/ESA System Definition Guide*. For a sample startup job stream, see the *CICS/ESA Operations Guide*. Two specific sample implementations are given in Appendix B, "Sample XRF implementations" in topic B.0.

Advice about terminal operands that can influence the takeover characteristics for individual terminals is given in "The terminal network" in topic 5.0.

Subtopics:

6.1 System initialization parameters

You start your active and alternate CICS systems in the same way as you start a non-XRF CICS system. You are recommended to use the same SIT for active and alternate, and define the system you are starting as either the active or the alternate by overrides. However, you can have separate SITs for active and alternate.

Most of the system initialization parameters operands are the same as for a system specified with XRF=NO. When an active is started, operands that are only for an alternate do not take effect. If that system is subsequently started as an alternate, those operands then apply. Similarly, when an alternate is started, operands for actives only take effect if it takes over and becomes the new active. Only operands affecting XRF are described in this section.

Subtopics:

6.1.1 Starting the active

The following parameters apply to actives:

```
START=AUTO
XRF=YES
APPLID=(generic-applid,specific-applid)
PDI=30|decimal-value
AIRDELAY=700|hhmmss
XRFSOFF=FORCE|NOFORCE
...
```

START=AUTO
This gives you a normal cold, warm, or emergency restart.

XRF=YES
The system signs on to CAVM because XRF support is required.

APPLID=(generic-applid,specific-applid)
The *generic applid* is the applid of this matching active and alternate pair. It is the applid by which the system is known to the end user. It is also used in interregion communication.

The *specific applid* is the applid for the active. It is used by CICS

when CICS opens the VTAM ACB. See "VTAM and NCP considerations for active and alternate" in topic 5.1 for more information.

PDI=30|decimal-value
decimal-value is the interval (in seconds) before the active tells the operator that it cannot detect the alternate's surveillance signal. This value is not critical. The default value is 30 seconds. No other action is taken; the active continues to operate as if the alternate were still present.

AIRDELAY=700|hhmmss
hhmmss is the restart delay (in hours, minutes, and seconds) that will elapse after a takeover before autoinstalled terminal entries are deleted if they are not in session. The default value is 700, that is, 7 minutes. A zero value means that the TCTTE of an autoinstalled terminal is not written to the catalog. You might choose a zero value to improve normal emergency restart times or your autoinstall performance. (For further guidance about the use of this operand, see the *CICS/ESA Performance Guide* .) For XRF systems, a zero value means that you might lose some autoinstalled terminal entries if there is a takeover during the catchup process. This is because the information about an autoinstalled terminal might not have been passed to the alternate through the message data set, and the alternate cannot learn about that terminal from the catalog. The end user of that terminal has to log on again. You should set the same restart delay value for both the active and the alternate, to maintain the takeover characteristics for autoinstalled terminals over several takeovers.

XRFSOFF=FORCE|NOFORCE
This operand is used by the active to determine whether it should send signon information to the alternate.

FORCE specifies that the active ensures that the alternate does not have any terminals signed on after a takeover.

NOFORCE (the default) allows you to be more selective about the terminals that are signed off, by using the TYPETERM definition or the SNTTE entry.

For more information, see "Signon after takeover" in topic 5.3.3.

6.1.2 Starting the alternate

You use the following parameters to start the alternate:

```
START=STANDBY
APPLID=(generic-applid,specific-applid)
XRF=YES
CLT=01
TAKEOVR=AUTO|MANUAL|COMMAND
ADI=30|decimal-value
JESDI=30|decimal-value
AUTCONN=0|hhmmss
RMTRAN=(transaction1,transaction2)
XRFSTME=nn|5
...
```

START=STANDBY
: Specifies that the system you are starting is an alternate.

APPLID=(generic-applid,specific-applid)
: *generic-applid* must be the same as that in the SIT of its matching active, but the alternate has a different *specific applid*.

CLT=xx
: Specifies the command list table to be used if a takeover occurs. *xx* specifies that table DFHCLTxx is to be used. The CLT applies only to the alternate. The CLT is described in "Command list table (CLT)" in topic 6.2.

TAKEOVR=AUTO|MANUAL|COMMAND
: *AUTO* specifies that the takeover is to be automatic, requiring no intervention by the operator. The alternate requests help from the operator only if it needs confirmation that the takeover can proceed safely. Possible causes of a request to the operator are described in "Checking for termination of the active" in topic 3.1.4.2. The operator can always issue a takeover command to an alternate, whatever takeover system initialization parameter is specified. So, if you define a system with TAKEOVR=AUTO, you retain the right to order a takeover. You can also change the takeover operand dynamically. "Supplied transactions for controlling the alternate" in topic 6.5 tells you about issuing operator commands to the alternate.

  COMMAND is the most restrictive type of takeover, whereby the alternate sends a message to the operator and takes over only when it receives a command to do so. This command could come from the operator (or the overseer), or, if the region is a dependent region in an MRO complex, from a master or coordinator region. If the alternate has noted the failure of the active, but has not received a command, it continues to run as an alternate.

  MANUAL ensures that the operator must approve a takeover if the alternate cannot determine that the active has failed. This could

occur if the active has stopped sending surveillance signals, but has not signaled a definite failure by signing off abnormally from the CAVM. The *MANUAL* operand is useful if you particularly want to avoid unnecessary takeovers. In a multi-MVS environment, it could also be useful if activity on the active CICS MVS (perhaps only for brief periods) prevents the active from sending a regular surveillance signal. With the *MANUAL* operand, operators can make decisions based on their knowledge of the other activity in the system. If the alternate receives a specific takeover command, or the active signs off abnormally from the CAVM, the takeover is automatic.

Table 3 summarizes the *TAKEOVR* operands and the types of takeover associated with each operand. An unconditional takeover involves no request to the operator for permission to take over. In a conditional takeover, a message to the operator asks for permission to start the takeover. For more detailed guidance about operator involvement, see the *CICS/ESA Operations Guide*.

Table 3. Types of takeover

Event	**TAKEOVR= AUTO**	**TAKEOVR= MANUAL**	**TAKEOVR= COMMAND**
Operator or program issues CEBT transaction	Unconditional takeover	Unconditional takeover	Unconditional takeover
Signoff abnormal	Unconditional takeover	Unconditional takeover	No takeover
Missing surveillance signal (see note 1)	Unconditional takeover (see note 2)	Conditional takeover	No takeover
Operator issues a CEMT transaction	Unconditional takeover	Unconditional takeover	No takeover

Notes:

1. If there is an MVS image failure when the active and alternate are in the same PR/SM CEC, with an XCF PR/SM policy in effect for the failing MVS image, the TAKEOVR option is not applicable, and there is an unconditional takeover. If an XCF PR/SM policy is not in effect, the operator is prompted by XCF to confirm the failure of another MVS image in the sysplex.

2. If the active CICS MVS image fails, the operator must confirm to the alternate that takeover may proceed.

ADI=30|decimal-value

Defines the delay (in seconds) before the alternate takes action after it has noted the disappearance of the active's surveillance signal. If you have coded TAKEOVR=AUTO, the alternate initiates a takeover. The ADI value here has to be a compromise, as follows:

° A low *ADI* value means that the alternate does not wait long before it starts its takeover process. So, a low value could mean a more rapid takeover after the active fails.

° A high *ADI* value reduces the risk of unnecessary takeovers, which might otherwise happen, when the active system has not failed, but has been temporarily prevented from transmitting its surveillance signals.

For TAKEOVR=COMMAND and TAKEOVR=MANUAL, the *ADI* value can be smaller, because the takeover is subject to intervention anyway.

An unnecessary takeover is not a serious error. It is more of an inconvenience; you have to try to determine the level of inconvenience when you set the *ADI* value. But you can prevent unnecessary takeovers in some predictable situations. The CEBT SET SURVEILLANCE command, described in topic 6.5.1, can prevent the alternate from reacting to the disappearance of the active's surveillance signal while, for example, the MVS image of the active CICS is stopped. However, there may be unpredictable, temporary stoppages of the active CICS (for example, when an unrelated address space in its MVS image issues an SDUMP). You should take such occurrences into account when choosing your *ADI* value.

You should also consider how to avoid some of the causes of unnecessary takeovers. You should carefully consider the implications of the MVS QUIESCE=YES|NO operand, as defined in the dump definition member of SYS1.PARMLIB. The choice in this parameter affects the dispatchability of SDUMP processing. This should influence your choice of the *ADI* value. For example, you can set the QUIESCE=NO operand for SDUMP (using the MVS CHNGDUMP command) to allow CICS to continue running during an SDUMP for another address space. Use of QUIESCE=NO keeps common storage unfrozen during the dump. Here, the *ADI* value can be kept low, and unnecessary takeovers are still avoided. (However, this has the disadvantage that there could be inconsistent dumped contents of the CSA, and similar effects.)

If you are using MVS/ESA SP 4.1, see also "Defining your XCF PR/SM policy" in topic 6.6.

JESDI=30|decimal-value

Defines the interval (in seconds) between takeover initiation and the point at which the alternate first prompts the system operator to investigate why the alternate cannot proceed. The alternate asks for this help if JES is unable to inform the alternate that the active has stopped. The *JESDI* value might have to be a compromise, as follows:

- A low *JESDI* value might avoid delaying the completion of a takeover, because the alternate system does not wait a long time before requesting operator assistance.

- A high value might avoid some unnecessary operator involvement. By waiting, the alternate allows the active more time to terminate, and then the alternate can continue the takeover by itself.

MVS or CEC failure is a typical case in which operator action is necessary. This is because neither JES nor the alternate CICS is able to determine that the other MVS or CEC has failed. A high *JESDI* value would delay the completion of the takeover here.

A CICS failure, on the other hand, can usually be handled automatically if the JES systems can access the shared spool. A low *JESDI* value would result in requests for operator action even though MVS is probably about to terminate the active CICS, and thus start a takeover sequence.

Even after the alternate requests the operator to confirm that the active job has terminated, the alternate continues to ask JES for the status of the active job. If it discovers from JES that the active has terminated, it cancels the request for an operator reply.

The operator can reply either that a CICS region has failed, or that the MVS or CEC has failed. If the operator replies "CEC" to the first alternate system that takes over, any other alternates taking over from actives that have failed on that MVS image do not have to ask for operator intervention, and their takeovers proceed without interruption.

If you are using MVS/ESA SP 4.1, see also "Defining your XCF PR/SM policy" in topic 6.6.

AUTCONN=0|hhmmss
Delays the reconnection, after a takeover, of class 2 terminals in session at the time of failure. The default is zero. You might set a delay for either of these reasons:

- You want your XRF-capable terminals working again as quickly as possible after a takeover. Delaying the reconnection of other terminals means that there is no interference with the switching of XRF-capable terminals.

- The operator needs time to do some manual switching of lines, for class 2 or class 3 terminals.

This delay also applies to any class 1 terminals that have failed to switch sessions. A terminal might not switch because it had logged on just before the failure and there was no time to build a backup session.

AUTCONN also applies to an active start. If you specify a long delay, terminals at normal start will be affected, unless you specify *AUTCONN* as an override.

RMTRAN=(transaction1,transaction2)
Defines the transaction that is initiated when logged-on class 1 terminals, with RECOVNOTIFY(TRANSACTION) specified, are switched. The default name is that specified in the system initialization GMTRAN parameter. If you provide your own transactions, transaction1 is run if the terminal has been signed off, and transaction2 if the terminal is still signed on. If you specify only one transaction, the second defaults to the value of GMTRAN. For more information, see "Using the RECOVNOTIFY keyword" in topic 5.3.2.

XRFSTME=nn|5
This operand has already been described in topic 5.3.3. It gives a time limit for signed-on terminals. When a takeover has not completed by the expiry of the time limit, terminals that would normally be in a signed-on state after a takeover are signed off.

6.2 Command list table (CLT)

Before you start to look at how the CLT works, you need to consider the role of the CAVM and its relationship to the CLT.

Subtopics:

6.2.1 CAVM and CLT

When the alternate takes over from the active, it cannot safely start to use resources such as files, databases, and the system log until it is certain that the old active has stopped using them. The CAVM ensures this integrity by making the alternate wait until the active job has terminated before allowing the use of those resources. The CAVM tries to minimize the wait time by issuing an MVS CANCEL command to remove the active CICS job. If the active and alternate are running in different MVS images, the CAVM uses JES facilities to send the CANCEL command to the destination MVS.

If an alternate in one MVS takes over from an active that is one of a set of MRO-connected regions running in a second MVS, the remaining alternates must be forced to take over, so that the MRO communication can continue. The CAVM can achieve this by issuing MVS system commands, which are coded in the CLT, causing each of the related alternates to take over.

The CAVM functions just described require MVS services that are restricted to authorized programs. Since CICS normally runs without MVS authorization, the CICS SVC has been extended to allow these services to be used in a strictly controlled way.

6.2.2 The CLT--background information

The CLT applies only to XRF. It is used only by the alternate; every

alternate must have a CLT. The CLT contains sensitive information that is needed to control the use of MVS-authorized services during takeover. The authenticity of that information must be guaranteed because the integrity and security of the entire MVS system might be compromised if an alternate could be made to use data supplied by an unauthorized person. Resource definition online (RDO) cannot provide the global protection needed for this sensitive data, because RDO security only operates at the level of individual CICS systems.

This information is therefore placed in the CLT. After assembly, you link-edit the CLT into an APF-authorized library. Unlike other CICS tables, the CLT is not loaded permanently when the alternate is initialized. It is loaded temporarily during initialization of the alternate, and when the alternate detects that an active job has signed on to the CAVM. This temporary loading is only for validity checking, after which it is discarded until takeover. (The validity check gives an opportunity to correct any problems, before the CLT is needed at takeover.) Loading only at takeover time means that you do not have to stop and subsequently restart an alternate to provide it with a changed CLT. During takeover, the CICS SVC loads the CLT into write-protected storage, and deletes it again after the CAVM has processed the information.

A CLT can contain the following information:

- Authorizations to cancel named jobs. Every CLT must contain the name of the active job that is to be canceled.

- Routing information needed to send CANCEL commands to the appropriate target MVS system (in a multi-MVS environment). You do not need this information in a single-MVS environment.

- MVS system commands and messages to the operator, to be issued during takeover. Typically, the function of these commands might be to tell other alternates to take over from actives in the same MRO-connected configuration. There could also be commands to handle non-XRF subsystems, such as DB2 (*). A master region would have such system commands. Messages to the operator might be instructions to perform some operator tasks to help the takeover.

Usually, each alternate needs a different CLT, but you may combine several of these CLTs in a single CLT load module. The specific applid of the alternate is used to select the relevant part of the single CLT when that alternate takes over. Using a single CLT might make it easier for you to manage your CLTs, especially in a large installation with many interconnected CICS systems.

There are examples of CLTs in Appendix B, "Sample XRF implementations" in

topic B.0. For reference information about the CLT, see the *CICS/ESA Resource Definition (Macro)* manual.

(*) IBM Trademark. For a list of trademarks, see topic FRONT_1.

6.2.3 The CLT in a single CICS configuration

Figure 22 shows you the relationship between the system initialization parameters and the way the CLT uses them.

If CICS2 is running as the alternate and it is told of a failure in the active (CICS1), or the operator instructs CICS2 to take over, DFHCLT02 is used. The FORALT operand of the DFHCLT macro allows CICS2 to cancel JOB1. Putting together the macros described in the *CICS/ESA Resource Definition (Macro)* manual, the sample CLT following the figure defines the CICS2 system illustrated in the figure.

Figure 22. System initialization parameters and CLT working together

```
DFHCLT02 DFHCLT TYPE=INITIAL,                                          *
               SUFFIX=02

         DFHCLT TYPE=LISTSTART,                                        *
               FORALT=(CICS2,JOB1)

         DFHCLT TYPE=WTO,                                              *
```

```
                WTOL=MESSAGE
MESSAGE         WTO 'CICS2 IS TAKING OVER, PERFORM MANUAL OPS',    *
                ROUTCDE=(1),                                       *
                DESC=number,                                       *
                MF=L


          DFHCLT TYPE=LISTEND


          DFHCLT TYPE=FINAL


          END
```

In a multi-MVS environment, the DFHCLT TYPE=INITIAL macro also contains information to route the CANCEL command to the active job, as follows:

```
DFHCLT02 DFHCLT TYPE=INITIAL,                                      *
                JES=JES2,                                          *
                JESCHAR=$,                                         *
                SUFFIX=02,                                         *
                JESID=(SMF2,JESID=JES2,2)
```

For reference information about the JESID operands, see the *CICS/ESA Resource Definition (Macro)* manual.

6.2.4 The CLT in a multi-MVS, MRO configuration

In an MRO configuration, each alternate needs a CLT, which can be loaded at takeover. As with the single CICS configuration, the CLTs are used only by the alternates.

In a multi-MVS, MRO configuration, when there is a takeover of one region to the second MVS, all the alternates must take over from their active counterparts to retain communication between the regions. This is because MRO does not operate across MVS images.

The system initialization parameters and the CLT determine the takeover policy for each active-alternate pair, and for groups where the actives are connected by MRO. In a hierarchy of communicating XRF regions, you use the CLT and the TAKEOVR operand of the SIT to structure the regions into dependent, master, and coordinator regions. The effect of a takeover of each type of region is as follows:

- The failure of an active dependent region does not automatically cause a takeover. Such a takeover is always initiated by a command from the operator or from another region. An alternate dependent region does not command other alternate regions to take over.

- The takeover of a failing master region forces the takeover of all communicating regions to the alternates in the second MVS image.

- If there is more than one master region, one of them may be used as a coordinator to organize the takeovers.

There is no need for such a hierarchy in a single-MVS MRO environment, because regions can be taken over from active to alternate (which becomes the new active region), and reestablish MRO links to all the regions with which the previous active communicated.

In the next example, shown in Figure 23, there are two active regions, connected by MRO, in a multi-MVS configuration. The master region has TAKEOVR=AUTO as its system initialization parameter. Its dependent region has the TAKEOVR=COMMAND system initialization parameter. The alternate master region's CLT authorizes the cancelation of the active master job, and the alternate dependent region's CLT authorizes the cancelation of the active dependent job.

In this hierarchy, if the alternate master region takes over from its failing active counterpart, it sends a command to the alternate dependent region telling it to take over from the active dependent region; the

```
MODIFY JOBD2,CEBT PERFORM TAKEOVER
```

command for the dependent region is coded in the CLT of the master region, and is shown in the figure. On receipt of this command, the dependent alternate region initiates a takeover. The CEBT transaction is described in "Supplied transactions for controlling the alternate" in topic 6.5.

If the dependent region fails, its alternate does not take over because of the TAKEOVR=COMMAND system initialization parameter. It takes over only on receipt of a command, and not automatically. Instead, the alternate sends a message to the operator stating that the active's surveillance signal is missing or that the active has signed off abnormally. The operator, or the overseer, might decide to try to restart the failed region in MVS1. This would avoid the disruption in the service provided by the master region that would occur on a takeover to MVS2. If the restart failed, it might be necessary to effect a takeover of both regions by issuing a CEBT PERFORM TAKEOVER command to the master alternate region. For restart in place, see "Restarting regions in place" in topic 4.2.2.

This is relevant to individual CICS failures. If the CEC or MVS failed, all regions would have to be taken over to the other MVS. A VTAM failure is a special case, and you use the XXRSTAT exit or the overseer to determine appropriate action.

With an MRO configuration, you can code a single CLT for all the regions involved. So, in the configuration discussed here, it could be for both master regions and both dependents. The FORALT operand indicates the section for a particular region. In the example CLT following the figure, only the entries for the current alternates (M2 and D2) are shown, for clarity.

Figure 23. System initialization parameters and CLT in an MRO configuration

```
DFHCLT01 DFHCLT TYPE=INITIAL,                                  *
               JES=JES2,                                       *
               JESCHAR=$,                                      *
               SUFFIX=01,                                      *
               JESID=(SMF2,JES2,2)

MAS2     DFHCLT TYPE=LISTSTART,                                *
               FORALT=(M2,JOBM1)

         DFHCLT TYPE=COMMAND,                                  *
               COMMAND='MODIFY JOBD2, CEBT PERFORM TAKEOVER'

         DFHCLT TYPE=WTO,                                      *
               WTOL=MESSAGE
MESSAGE        WTO 'TAKEOVER TO NUMBER 2 REGIONS',             *
               ROUTCDE=(1),                                    *
               DESC=number,                                    *
               MF=L

         DFHCLT TYPE=LISTEND
```

```
DEP2     DFHCLT TYPE=LISTSTART,                                  *
               FORALT=(D2,JOBD1)

         DFHCLT TYPE=LISTEND

         DFHCLT TYPE=FINAL

         END
```

You can extend the usefulness of the CLT by adding other commands to the CEBT commands shown here. The CLT can be used to issue any MVS commands that are needed to complete the takeover, for example, VTAM VARY NET commands. In this way, you can reduce the need for the operator to be involved.

6.2.5 Use of the coordinator

In a large multi-MVS, MRO configuration, you might have more than one master region and any number of dependent regions. Figure 24 shows that you might find it convenient to nominate one master region as the coordinator. You do not have to do this, but you might find that it reduces the number of redundant commands that would otherwise be issued during a takeover of many regions (if, for example, three master regions all give takeover commands to several dependent regions).

Figure 24. Flow of control and the coordinator region

The following notes apply to the figure:

Notes:

1. When the active master region fails, it triggers the alternate master region.

2. The alternate master region issues a CLT command to the alternate coordinator region to initiate a takeover.

3. The alternate coordinator region issues CLT commands to alternate dependent regions to initiate takeovers.

4. The alternate coordinator region sends a redundant command back to the alternate master region to initiate a takeover. If the coordinator active region had failed, rather than the master, this command would not be redundant.

If a coordinator region fails, its alternate uses the CLT to issue CEBT PERFORM TAKEOVER commands to all other alternate regions, master and dependent. If a master region fails, its alternate will initiate a takeover, and issue a command to the alternate coordinator region to take over. Then the coordinator will issue its own commands to all regions, in the way that a single master region would.

There is an example of a CLT with a coordinator region in Appendix B, "Sample XRF implementations" in topic B.0.

6.3 User exit for VTAM failure

For XRF, the global user exit, XXRSTAT, allows you to code a decision after a VTAM failure. It runs in the active system only. For definitive product-sensitive programming interface information about exits, see the *CICS/ESA Customization Guide*.

User exit XXRSTAT is called after CICS has been told of a VTAM failure by the TPEND exit. This occurs just before the update of status information that will become available to the alternate through the CAVM data sets. In the exit you can choose what to do following a VTAM failure. You can tell CICS to take any of the following actions:

° Abend CICS and thus force a takeover, or whatever action you have specified if that region abends. You may specify a dump with the abend. The status information is not written to the control data set. If you do require a takeover, you need the TAKEOVR=AUTO or TAKEOVR=MANUAL system initialization parameter.

° Allow the CICS region to continue, after updating the status information to tell the overseer that VTAM has failed. The overseer

then performs the action that you have specified for this particular combination of circumstances, as described in the next section.

° Suppress the update of the status information, and allow the CICS region to continue, on the assumption that the VTAM region will be restarted. In this way, the overseer, if present in the system, is not made aware of the VTAM failure and does not go through its VTAM failure procedure. The alternate must be taken down by the time the operator issues the CEMT SET OPEN command on the active, to remove its unusable backup sessions. When it is restarted, new backup sessions are built to match the new sessions to the active.

The alternate terminates by itself if its VTAM fails. In a multi-MVS environment, if the active's VTAM fails, and you choose to restart VTAM, you must manually take down the alternate.

In some configurations, you might prefer to handle VTAM failures in the exit program (by initiating a takeover or tolerating the VTAM failure) instead of in the overseer. The exit program is probably quicker and relatively simple to implement. The overseer is more complex, and could be slower. However, the overseer allows you to use more complicated logic to deal with the situation.

6.4 The overseer

The overseer was introduced in topic 4.2.3. The IBM-supplied sample overseer can perform two functions. It can display the status of XRF regions, and it can restart a failed region in place. The overseer sample is named DFH$AXRO, and is in the CICS sample library, CICS330.SDFHSAMP. There is also a pregenerated version ready to use. See the *CICS/ESA Customization Guide* for guidance information about using the overseer, and for definitive product-sensitive programming information about the interface for defining actives and alternates to the overseer. See the *CICS/ESA Operations Guide* for guidance information about its operator interface.

You can write your own overseer program to extend its capabilities. The overseer can perform non-CICS functions. Here are some examples of what the overseer can do:

° Display its status information in a suitable format at regular intervals.

° Examine information about VTAM failure passed by the user exit, and act accordingly. Information is available to the overseer about the

last eight failures detected by the active CICS. Make sure that the overseer and user exit actions are consistent. The overseer could make its own enquiries into the state of VTAM. Its action could depend on many things: the length of the VTAM outage, the number of times VTAM has failed, the number of end users affected, or the time of day. Its most likely action would be to initiate a takeover by issuing a CEBT PERFORM TAKEOVER command.

- Make decisions beyond the capability of the CLT, if the system initialization parameters and CLT definitions do not provide the required flexibility. The overseer can provide additional control, and thus take actions that would otherwise have to be taken by the operator. For example, you could put logic in the overseer so that it could make decisions based on the time of day. If a region failed during a period when you knew it was lightly used, you might prefer not to initiate a takeover, involving many regions, but to restart the failed region in place. At other times, the overseer could initiate a takeover, by issuing a CEBT PERFORM TAKEOVER command.

- Issue commands during takeover, not only to CICS regions. You might choose to put a command in the overseer rather than in the CLT, because the overseer can handle variables in the commands, and the CLT cannot.

- Detect the possibility of a looping or waiting active. The sample overseer can do this after minor changes and reassembly.

| - Operate on CICS Version 2 and Version 3.1.1 and 3.2.1 systems running in XRF mode in the same MVS images as a Version 3.3 CICS with XRF system.

The sample overseer carries out basic functions, which will be adequate for some installations. Other installations will accept the added complexity and significant programmer effort involved, and extend the scope of the overseer.

6.5 Supplied transactions for controlling the alternate

Because the alternate is only partially initialized, the usual transactions for a CICS system do not apply to it. There is a system console transaction specifically for the alternate--the CEBT transaction. The CEMT transaction may be used to initiate a takeover. For reference information about CEBT and CEMT, see the *CICS/ESA CICS-Supplied Transactions* manual.

Subtopics:

6.5.1 The CEBT transaction

The CEBT transaction can be issued from a master or coordinator region to a dependent region, when it is normally used to start a takeover. The operator, too, can issue CEBT transactions, from the system console or from a TSO console if extended console support is in use.

The CEBT transaction is usable from the time when the alternate is initialized to the time after takeover when CEMT becomes usable. The operator can use CEBT to do the following:

° Request the alternate to take over.

This is relevant for a failed dependent region, which is taken over only when its alternate receives specific instructions. The failure of a dependent region results in a message to the operator, and the operator can then decide what to do. The first thing to do would probably be to try to restart the failed region; you can use the overseer to automate that process. If it is impossible to restart the region, the operator might initiate a general takeover to the other MVS image, by issuing a CEBT PERFORM TAKEOVER command to a master or coordinator region.

The operator can use a CEBT PERFORM TAKEOVER command to cause a takeover when the alternate has not recognized that the active is not working properly.

For planned maintenance, you use this command to request a takeover. You also use it to return the CICS workload to the preferred MVS image, when it has been recovered after a failure. If you want to move a set of MRO regions from one MVS image to another, you need to issue this command only to the alternate coordinator region, which then issues its own commands to the other regions.

A CEBT PERFORM TAKEOVER command is not governed by the takeover type

specified at system initialization. If TAKEOVR=AUTO is specified, the operator is still able to initiate a takeover.

° Change the takeover type specified at system initialization.

In this way, you can change the takeover operand without shutting down the alternate. (The takeover types are described under "System initialization parameters" in topic 6.1.) Using CEBT you could, for example, change the automatic takeover operand to the manual takeover operand.

You might find this command useful for altering the takeover characteristics of a region during a particular working period, at the end of the working day, or if the level of operator coverage is changing, for example.

° Shut down the alternate.

° Make the alternate ignore the active surveillance signals, thereby removing its capability to take over. CEBT can also restore surveillance of the active's signals.

For example, by switching off surveillance, you are able to stop the active's MVS image, and not cause a takeover. When the MVS is restarted, the active starts work again. Then surveillance can be switched on again. However, tracking continues normally while surveillance is switched off.

° Manage dump data sets, and request a dump.

° Manage auxiliary trace data sets, and switch trace on and off.

6.5.2 The CEMT transaction

Another way to control the alternate is to issue a CEMT PERFORM SHUTDOWN TAKEOVER or CEMT PERFORM SHUTDOWN IMMEDIATE command to the active, which causes a takeover by the alternate. If you specify TAKEOVER rather than IMMEDIATE, normal shutdown processing is carried out before the takeover starts. This is unlike takeovers initiated in any other way. In particular, a warm keypoint, which includes the current TCT state, is written to the catalog. When the catchup process uses the catalog, it will use the information written at the warm keypoint. If IMMEDIATE is specified, a warm keypoint is not written and therefore the catalog

information is unchanged. If either IMMEDIATE or TAKEOVER is specified, sessions which are not XRF-capable are terminated immediately.

6.6 Defining your XCF PR/SM policy

XCF PR/SM policy is defined in the SYS1.PARMLIB member XCFPOLxx. For details, refer to the *MVS/ESA Planning: Sysplex Management Guide*. When defining your XCF PR/SM policy there are several points you must consider. These are:

- To avoid the DFH6582 and DFH6583 messages, ensure that the interval expressed by the sum of the CICS delay intervals ADI and JESDI is greater than the sum of the values specified for XCF INTERVAL and the XCF PR/SM policy interval RESETTIME or DEACTTIME. This means that XCF notes the MVS image failure before the alternate CICS does, and therefore the operator does not receive prompts enquiring about the state of the failing system.

- In a single-CEC environment, to prevent message IXC402D from being issued, set the interval expressed by the sum of XCF INTERVAL plus XCF PR/SM policy RESETTIME or DEACTTIME to be greater than XCF OPNOTIFY.

  In a two-CEC configuration, the message IXC402D is produced by the backup system if a failure of the primary MVS image occurs, but the necessary system reset and reply of DOWN to the IXC402D message could be automated using NetView with TSCF. In such an environment, setting the ADI plus JESDI interval greater than the XCF OPNOTIFY interval prevents the DFH6582 and DFH6583 messages.

- To prevent the backup system from performing unnecessary takeovers, the interval expressed by RESETTIME or DEACTTIME should be greater than the time for which the primary system would be temporarily nonfunctional, if, for example, the system is in a loop, or taking an SVC dump.

The system administrator should define to MVS the preferred policy for handling MVS failures in a PR/SM environment, and define to PR/SM each logical partition's authorization to cause reset or deactivation of another logical partition. For further guidance, see the *ES/3090 Processor Complex - Processor Resource/Systems Manager Planning Guide* (GC22-7123).

PR/SM ARF (automatic reconfiguration facility) also allows you to define to MVS the actions that PR/SM takes to adjust the relative sizes of the failing logical partition, and also to define which logical partition is

to take over its resources. For further information, see the *MVS/ESA Initialization and Tuning Reference* manual.

6.7 Sharing data sets

There are three ways data sets can be shared between the active and the alternate, as follows:

1. Actively shared, like the CAVM data sets.

2. Passively shared, meaning that only one system at a time accesses a data set, normally the active, or the alternate when it begins its takeover processing. The system log and user data sets are examples.

   The CSD (CICS system definition file) is passively shared. An active and alternate pair running in two MVS images cannot share the CSD with another CICS system unless some form of global enqueuing, such as GRS (global resource serialization), is in use. The same applies to the DFHCSDUP utility.

3. Unique to active or alternate. For example, the active and alternate each has its own auxiliary trace data sets and dump data sets.

Because the system log is shared (although not used at the same time) by the active and alternate, it must be defined DISP=SHR in the JCL. For user VSAM data sets and DL/I data sets, you cannot allocate the data sets using DISP=OLD through JCL, because this prevents the alternate from starting. Instead:

- Use dynamic allocation, thus avoiding the JCL.

- Specify DISP=SHR in the JCL, and take the risk of other, unrelated jobs using your data sets.

- Specify DISP=SHR, and use RACF to reduce the possibility of any integrity exposure.

- For DL/I, specify DISP=SHR and use the data sharing functions of DBRC to maintain integrity. (For more information, see "IMS--local DL/I" in topic 7.3.)

- For VSAM, specify DISP=SHR and set appropriate share options.

- Ensure that RACF definitions are the same for the active and alternate.

For further guidance about the way to use specific data sets, see the *CICS/ESA System Definition Guide*.

| 6.8 Storage protection considerations

| CICS with XRF is fully supported by the storage protection facility. Either the active or the alternate system can operate without storage protection even though its partner does. This is necessary, for example, in circumstances where the alternate is running on a processing system, or under a level of MVS, that does not support the storage override facility. In this situation you should specify one system initialization table for use on both the active and alternate CICS regions, and modify it as appropriate for either the active or alternate by providing system initialization override parameters at run-time.

| CICS does not save any of the storage-related system initialization parameters in the global catalog, including the DSA and cushion sizes.

7.0 XRF and other products

This topic provides information about XRF and its relation to other products.

Subtopics:

7.1 DB2

IBM DATABASE 2 (*) (DB2) is supported in its own address space through the CICS task-related user exit interface, also known as the CICS resource manager interface (RMI).

(*) IBM Trademark. For a list of trademarks, see topic FRONT_1.

Subtopics:

7.1.1 Single-MVS environment

In a single-MVS environment, after a CICS failure has caused a takeover, the alternate attempts to reestablish communication with DB2 in the same way as a non-XRF CICS system after an emergency restart. You can use PLT entries to try to initiate the reconnection between DB2 and the new active CICS, to minimize the time that the CICS/DB2 interface is unavailable.

7.1.2 Multi-MVS environment

In a multi-MVS environment, you must be able to run DB2 in both MVS images. A takeover by an alternate in a second MVS means that DB2 in that MVS has to become the operative DB2 for the new CICS active systems. To ensure data integrity, the two DB2s, in the two MVS images, must not both be allowed to use the same resources at the same time. This means that the STOP DB2 command must complete on the old active system before the START DB2 command is issued on the new system. Note that in the case of CICS failures you must consider non-CICS use of DB2. Interactive or batch

SQL (*) and DB2 utilities must be completed or terminated before DB2 service may be switched to the second MVS. You can use the CLT or the overseer to issue MVS commands to stop DB2 on the MVS of the failed active CICS, and to issue a START DB2 command to a second DB2, which uses the same logs and databases as the first, on the MVS of the new active.

You can use PLT entries to initiate the reconnection to the second DB2. If DB2 restart has not been completed when you initiate the reconnection, DB2 data will not be available to application programs. You might consider temporarily blocking transactions that access DB2. One way to do this is to group all those transactions in an XLT and enable them all together when DB2 data becomes available. You should also consider the STRTWT=YES option on the INIT macro in the DB2 resource control table (RCT). This option allows the reconnection process between CICS and DB2 to wait until DB2 is restarted. Other PLT and further CICS processing is not affected. See the appropriate DB2 manuals for further information about this option.

There are two possible techniques using global resource serialization (GRS) for reducing the risk of data integrity problems caused by concurrent execution of DB2 on both the active and alternate MVS images:

1. Control all DB2 data set activity by GRS, or equivalent operational procedures, under the major name SYSVSAM. You should evaluate carefully the overhead of this technique in your environment.

2. Use GRS only on the DB2 bootstrap data set, preventing multiple allocation of the data set by using GRS services. This is a viable alternative because you need to protect only the first instance of that data set. Using GRS in this way incurs a much smaller overhead.

CICS systems and the DB2 system with which they communicate must all be in the same MVS image. If two active CICS systems (whether or not they are MRO-connected) access the same DB2 system, and one of them is taken over by its alternate in a second MVS, the other CICS and the DB2 must be taken over to that MVS if both CICS systems are to continue using the same DB2 system. To handle this, the relevant CLTs could contain CEBT PERFORM TAKEOVER commands to the other alternates, even if they are not MRO-connected.

7.2 DBCTL

CICS with XRF supports the use of DBCTL as the DL/I database manager. This support is not described in this book. For guidance information about DBCTL, see the *CICS/ESA CICS-IMS Database Control Guide*. DBCTL support is available with IMS/ESA Version 3.

7.3 IMS--local DL/I

CICS/ESA 3.3 offers local DL/I support in CICS for IMS/VS Version 2 and IMS/ESA Version 3. With IMS/VS Version 2 and later levels, you can run CICS with XRF=YES or with XRF=NO. The XRF=NO cases are not discussed in this book, which only deals with aspects of local DL/I data sharing that are different for CICS running with XRF.

The following list shows the operating scenarios that are discussed. The scenarios apply to single- or multi-MVS configurations.

- CICS with IMS (without DBRC)--no data sharing
- CICS with IMS with DBRC (recovery control only)--no data sharing
- CICS with IMS with database-level data sharing
- CICS with IMS with block-level data sharing.

These topics are discussed in the following sections, with the emphasis on data sharing. Even if you're only interested in block-level data sharing, you are advised to read the earlier sections, because they contain information relevant to all data-sharing scenarios.

Subtopics:

7.3.1 No DBRC and no data sharing

There are no XRF-specific considerations for CICS systems that use DL/I databases without DBRC and that do not share the databases with other subsystems, unless you use IRLM as a single lock manager. In that case, the remarks about the generic applid in the following section apply.

7.3.2 DBRC recovery control and no data sharing

In this and the following scenarios, you should note that only the generic applid applies to IRLM and DBRC. It may be regarded as a consistent subsystem name. It is the name that appears in a RECON display. Also, the alternate CICS is not signed on to DBRC, or identified to IRLM, until takeover.

7.3.3 Database-level data sharing

When running CICS with XRF in a data-sharing configuration, the normal considerations of data sharing still apply. The figures in topic 7.3.3.2 show examples of data sharing.

This is particularly so when you run XRF in a single-MVS image, because other systems see a takeover as an emergency restart. Therefore, data-sharing recovery should be based on procedures you have for existing data-sharing configurations. In general, for single- and multi-MVS configurations, the guidance information about data sharing in the *CICS/ESA Operations Guide* and the *CICS/ESA Recovery and Restart Guide* remains valid for your XRF systems.

In the sections that follow there are a few points to note about the various types of failure.

Subtopics:

7.3.3.1 MVS or CEC failure with database-level sharing

Single-MVS image: In this situation, there are no XRF-specific considerations. Everything has failed and must be restarted in the same way as after a non-XRF failure. Your existing non-XRF procedures will continue to work in this situation.

Multi-MVS images: When there is a failure of MVS or a CEC (in a two-CEC environment), all sharing subsystems move to the second MVS image. The failed batch jobs cannot run until restart procedures are carried out. In this case, the CICS systems all share on the new MVS.

Batch backout of any failed batch jobs should be accomplished as quickly as possible because, if a batch job has update authority, it prevents
other subsystems from obtaining update authority. This means that CICS
transactions cannot start if they have read-with-integrity specified.

Like the single-MVS scenario, there should not be many changes to your existing procedures.

7.3.3.2 CICS failure with database-level sharing

Single-MVS image: The situation is the same as for an emergency restart of a non-XRF system. The XRF takeover means that the duration of the outage is reduced. The other participating subsystems react to the takeover as if it were the restart of a failed CICS region.

If CICS had update authority when it failed, it retains that authority during takeover, as a non-XRF CICS system does during emergency restart. After the takeover, when the emergency restart processing has taken place, including backout, CICS gives up authorizations on all databases. It reacquires authorizations as the databases are first accessed by applications. However, this leaves an interval in which other subsystems might obtain authorization.

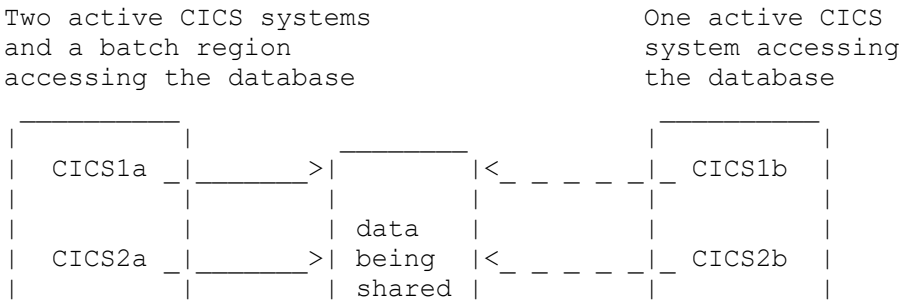

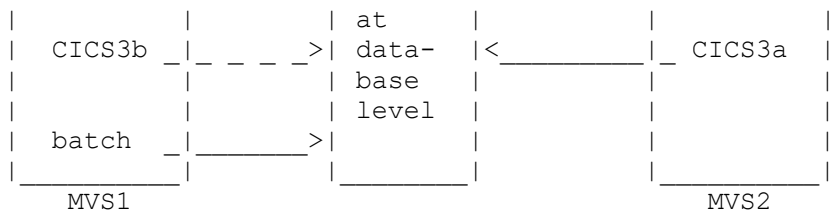

Figure 25. Multi-MVS database-level data sharing

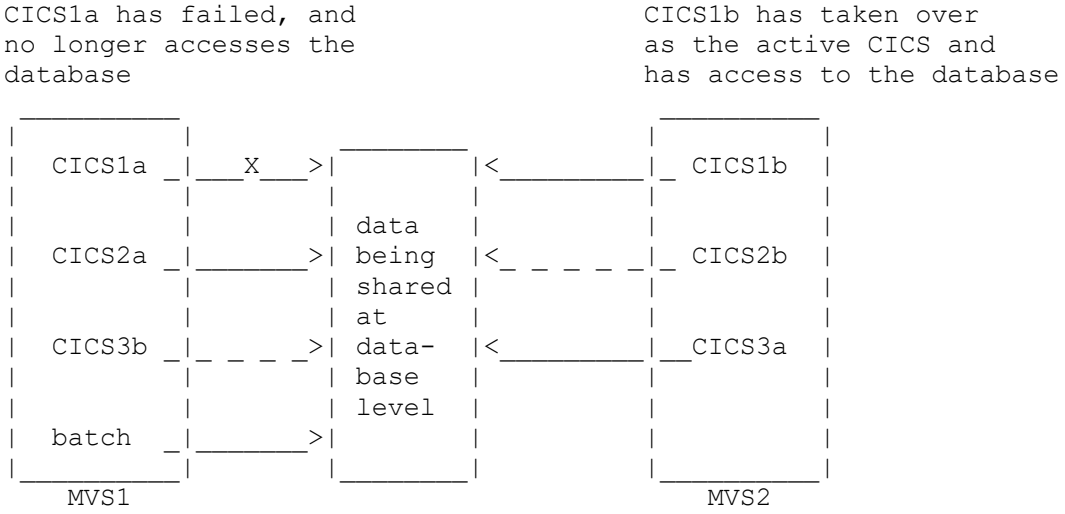

Figure 26. Multi-MVS database-level data sharing after failure of one active CICS

Multi-MVS images: This configuration is shown in Figure 25 and Figure 26. There are active sharers on MVS1 and MVS2. When an active CICS fails, its alternate on MVS2 takes over, and database-level data sharing continues with the new active when backout has been completed.

7.3.3.3 Database-level data sharing with IRLM

IMS/VS Version 2.2 and later levels of IMS support improved integrity for subsystems reading with ACCESS=RO, if they are using IRLM. To obtain this improvement in read integrity, you should specify SCOPE=GLOBAL for the IRLM. The usual data-sharing performance considerations apply, but in an XRF environment you should consider the effects of your availability

policy. You might have all the active sharers on one MVS, and all the alternate sharers on a second MVS. If a CICS system is taken over to the second MVS, while other sharers remain on the first MVS, the after-takeover phase will incur the normal overhead for NOTIFY traffic in an inter-MVS data-sharing environment. See Figure 27.

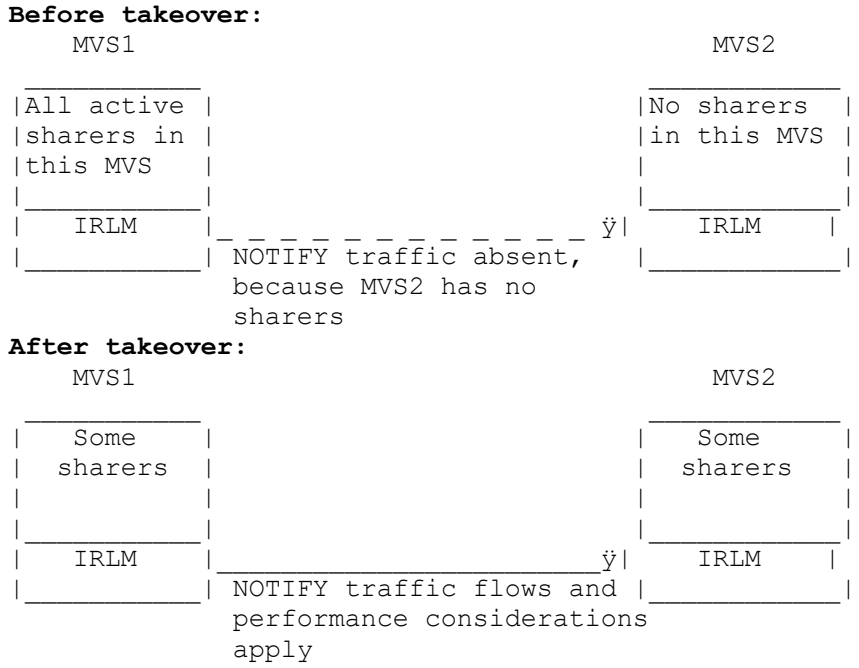

Figure 27. Database level sharing with SCOPE=GLOBAL defined for the IRLM

7.3.4 Block-level data sharing

CICS running with XRF can share at block level with:

- Other CICS systems, XRF or non-XRF
- IMS/DB subsystems (including batch)
- IMS/DC, XRF or non-XRF.

For block-level sharing in multi-MVS configurations, you must use IMS/VS 2.2 or later, because of their support for the ready region.

Subtopics:

- 7.3.4.1 Single-MVS block-level data sharing
- 7.3.4.2 Multi-MVS block-level data sharing
- 7.3.4.3 Performance considerations

7.3.4.1 Single-MVS block-level data sharing

In one MVS, block-level data sharing involves no XRF-specific considerations, apart from the generic applid considerations outlined in "DBRC recovery control and no data sharing" in topic 7.3.2. If there is a CICS failure, the takeover acts in the same way as an emergency restart. IRLM deals with a takeover as though it was an emergency restart, and any batch jobs continue to run. Your existing non-XRF procedures should continue to work in this situation.

7.3.4.2 Multi-MVS block-level data sharing

CICS with XRF gives you the capability of running your data-sharing CICS systems on one MVS with supporting alternates in another MVS. If you have a failure of the environment supporting the active CICS systems, you can initiate takeovers to the second MVS to maintain data availability.

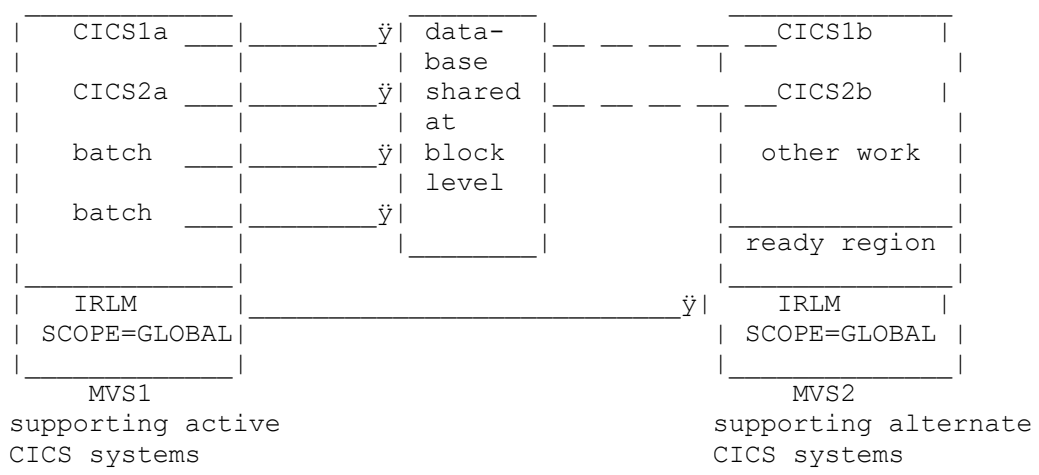

Figure 28. Block-level data sharing in two MVS images

To obtain this level of support, there are the following requirements:

- Reregister, if necessary, your databases as SHARELEVEL(3).

- Configure the IRLM on each MVS as SCOPE=GLOBAL and RULES=AVAIL.

  The GLOBAL option enables the passage of information (including lock information) from the IRLM in the MVS of the active CICS to the IRLM in the second MVS. This information makes possible data availability after a takeover caused by either a CICS failure, an MVS failure, or (in a two-CEC system) a total CEC failure, because a new active CICS system on the second MVS is provided with the information it needs to take part in data sharing with systems (CICS or batch) in either MVS.

  The lock information allows full block-level data sharing on the second MVS, even if the active's IRLM or MVS fails, or (in a two-CEC system) if the CEC fails.

  A ready region (see the next item in this list) is required to obtain advantages in terms of data-sharing availability. The loss of availability is limited to those database blocks containing uncommitted updates issued by CICS or IMS/DB or IMS/DC subsystems that have failed. These blocks are locked until the corresponding updates have been backed out by CICS or IMS restart processing or by batch backout.

  As in other XRF configurations, it is essential that operators understand set procedures, such as manual batch backout operations, and follow them carefully and as quickly as possible to restore availability. This includes issuing the SETSTATE SYSTEM command to the IRLM in response to IRLM message DXR027A to obtain the benefits of the retained locks.

- Install a special type of DL/I batch job, called the **ready region**, in MVS images that support alternate CICS systems. The *CICS/ESA Operations Guide* contains further guidance information, including a sample job to start up a ready region.

  A ready region is included in Figure 28. The ready region must be left running on the second MVS all the time, to obtain the availability improvements mentioned above.

If an IRLM fails, normal considerations apply. You should restart the IRLM as quickly as possible. There is no need for a takeover.

7.3.4.3 Performance considerations

The implementation of data sharing at block level on two MVS images involves decisions about performance and capacity planning. This section introduces the topic, and for further guidance information you should refer to the *CICS/ESA Performance Guide*. The use of data sharing introduces performance considerations that you should take into account when you are designing your system. The effect on performance varies according to the characteristics of the system and the workload on it. When CICS is running with XRF, you should consider carefully configurations in which the sharing takes place (perhaps only temporarily) between active CICS systems on one MVS and active CICS systems or batch jobs on another MVS. To obtain the protection of XRF, you accept the overhead of the traffic between the IRLMs. While all the active data sharers are in one MVS, and the second MVS contains only alternates, there is an optimization to reduce this traffic to a minimum.

If MVS or (in a two-CEC configuration) the CEC of the active sharers fails, all the sharing subsystems may be moved to another MVS. Until XRF cover is reestablished, they share on the new MVS, and there is no sharing overhead between MVS images.

If there is an active CICS or VTAM failure, you have to decide the best way to minimize disruption in your configuration. As in a non-data-sharing configuration, the type of region that fails is a significant factor. Unless it is vital to move the failed CICS region to the second MVS, you might choose to restart the failed region in place and keep all the active sharers on one MVS.

However, it might be your policy, for some regions, to take over to another MVS if they fail. Your system is now running full block-level data sharing between MVS images. This involves an increased overhead of lock passing between the IRLMs, and response-time considerations, because the optimization that was possible when the CICS systems on the second MVS were only alternates, and were not involved in the data sharing, is no longer possible. You have to decide whether to accept the increased overhead. It might be only temporary, until you can take over the region back to its original MVS image.

When one active CICS system fails, and there has to be a takeover to the second MVS, you might consider taking over all the data-sharing CICS regions to the second MVS, even though they might not be MRO-connected. You can put CEBT commands in the CLT to act on regions that are not connected by MRO, to effect the takeover. You could then either allow short-running, sharing batch jobs to complete on the first MVS, or stop the batch jobs and restart them on the second MVS, thus avoiding the overhead of sharing between MVS images.

7.4 NetView

You can use the network management product NetView to add function to XRF. One possible use of NetView is to propagate changes in the USERVAR value to remote VTAMs that are in communication with the local VTAM of the XRF complex. However, you are recommended to leave this propagation to the VTAM automatic USERVAR facility, described in "VTAM and NCP considerations for active and alternate" in topic 5.1.

Subtopics:

7.4.1 Restarting a 37xx or the NCP

You can use NetView to obtain rapid notification of a failed 3705, 3720, 3725, or 3745 Communication Controller, and its network control program (NCP). You may also use NetView to restart them. This adds to the restart capability of XRF. Figure 29 shows the way NetView can do this.

In this section, we give you an overview. For further reading, see the *Network Program Products Planning* manual.

When a 37xx or its NCP fails, VTAM issues an error message. You can pass this message to NetView, which compares the message with its message table. If there is a match, NetView initiates a CLIST that corresponds to that message.

You code CLISTs yourself, and you can choose the sequence of recovery actions. You can refresh the message table, thereby changing your recovery procedure, without stopping NetView.

If you prefer not to automate such a procedure, you can send messages to the operator, requesting intervention. Alternatively, the CLIST can attempt to reload the 37xx communication controller. If the 37xx communication controller cannot be reloaded, you can use a further CLIST to prompt the operator to switch to another, if one is available. You can then use a CLIST to acquire resources from the failed 37xx and activate them for the new one.

Figure 29. Automating 37xx recovery with NetView

| 7.5 VM

| CICS with XRF will work under VM/XA (*) and VM/ESA (*). Such usage is not
| recommended for production purposes, because there is no cover against VM
| failures. Running CICS with XRF under VM is suitable for test
| environments.

| (*) IBM Trademark. For a list of trademarks, see topic FRONT_1.

A.0 Appendix A. Checklist

To help you organize your work for XRF, this alphabetic checklist contains XRF-related activities for the systems programmer. Much of the information summarized here is in the appropriate CICS books, whose titles are given. Long-term planning items, such as setting up the correct XRF environment, and selecting the configurations you need, are not included here. For guidance information about the early stages of planning, see the *CICS/ESA Facilities and Planning Guide*.

Application programs
: Ensure that your existing application programs run in an XRF environment. You should look at those programs that depend on the specific applid, or that have unsupported interfaces into CICS code.

Automatic reconfiguration facility (ARF)

Define to XCF your policy for MVS failures in a PR/SM environment. See "Defining your XCF PR/SM policy" in topic 6.6.

Backup while open
: For files defined as eligible for backup while open (BWO), that is, BACKUPTYPE=DYNAMIC has been specified in the FCT and the files are currently open for output, tie-up records are written during activity keypointing. To reduce the amount of tie-up record log activity if keypoints occur at intervals of less than thirty minutes, tie-up record log writes will be limited to one set per thirty minutes. Therefore, in a large XRF system with a high keypoint rate, there will be many keypoints without tie-up record sets. It is likely that your XRF system will have a high keypoint rate. You should note possible performance effects.

CLISTs
: Provide CLISTs to propagate the USERVAR, if required. See topic 7.4.

Collection of data
: Do you have any utilities that will have to collect data across a takeover? This data could be SMF data, in two separate data sets if you are running in a multi-MVS environment. There must be unique SMF IDs on all participating MVS images in an XRF complex, so that the ID identifies the data from each CICS system. For guidance information, see the *CICS/ESA Operations Guide*.

Dump
: Determine the amount of information you want dumped by the failing active. For further guidance information, see the *CICS/ESA Operations Guide*.

  You should use the MVS SDUMP facility, with an appropriate system initialization ADI value to avoid unnecessary takeovers. See topic 6.1.2.

  Modify AMDPRECT for CICS PRDMP formatting of SDUMPs. For further guidance information, see the *CICS/ESA Operations Guide*.

JES (in multi-MVS environment)
: Set up a shared spool environment, with JES2 multiaccess spool, or
| JES3. In an XCF sysplex, use the AUTOESYS and RESTART options of the
| MASDEF initialization statement.

MVS considerations
: Define CICS as a subsystem to MVS. The *CICS/ESA Installation Guide* tells you how to do this.

In a multi-MVS environment, there are advantages if both MVS images are at the same level. For example, MVS/ESA SP 4.1 introduces extended console support; a takeover to an image using a prior MVS level would mean a change in console support.

If you are running CICS using MVS/ESA SP 4.1, in a PR/SM environment, define to XCF your policy for handling MVS failures within your sysplex. See "Defining your XCF PR/SM policy" in topic 6.6.

If you use cross-memory services for MRO, consider making CICS systems non-swappable, to eliminate swap-out and later swap-in when the alternate takes over. For guidance about how to make CICS non-swappable, see the *CICS/ESA Performance Guide*.

NCP

Define NCP for XRF. See topic 5.1.4.

Node error program

The *CICS/ESA Customization Guide* contains the definitive product-sensitive programming interface information about the node error program.

Operator instructions

Prepare operator instructions, so that the operators understand the CEBT transaction, the way an XRF takeover works, and any extra tasks they might have to perform. There is information about operating XRF throughout this book. For further guidance information, see the *CICS/ESA Operations Guide*.

Overseer (if required)

Define the active and alternate CICS systems to the overseer. Create your own overseer program, if required. The *CICS/ESA Customization Guide* contains the definitive product-sensitive programming interface information, and further guidance, about the overseer.

Programmable terminals

Ensure that your terminals have any extra code they need to enable them to connect to whichever system is the active.

Programs run at shutdown

Review programs run in the PLTSD phase and post-execution batch runs. Evaluate the need for the data they extract, and whether the data is needed by the alternate, because these programs only run when a takeover occurs after an orderly shutdown of the active, initiated by a CEMT PERFORM SHUT TAKEOVER command. For definitive

product-sensitive programming interface information about PLTSD programs, see the *CICS/ESA Customization Guide*.

Recoverable resources (in a multi-MVS environment)
: Ensure that all recoverable resources and their dependencies are accessible from both MVS images.

Shared DASD
: Many data sets for XRF must be on shared DASD, in particular the CAVM data sets. The *CICS/ESA Operations Guide* gives advice about the characteristics of each data set.

Signon options
: Ensure that each terminal has the correct characteristics for signon after a takeover. See topic 5.3.3 and following topics.

System initialization programs
: Check that any user programs that run at initialization perform as expected in an XRF environment.

System logging
: System logging must be on two disk extents.

  Consider using automatic archiving for journal archiving. The *CICS/ESA Operations Guide* describes automatic archiving.

System naming conventions
: Review the need for changes or additions to system naming conventions.

Table definitions
: You need to consider the definitions for the:

  - SIT and overrides
  - CLT
  - RDO TYPETERM options or TCT entries.

  There is some guidance about definitions in this book. For more details, see the *CICS/ESA Resource Definition (Online)* and the *CICS/ESA Resource Definition (Macro)* manuals.

Takeover message
: Code a message, or write a transaction, to provide information to terminal users after takeover, if required. See topic 5.3.2.

Time-of-day clock
: The setting of the clocks in a multi-MVS environment must be as close as possible at IPL. If the alternate clock is running later than the active clock there is a delay at takeover. This is not a consideration if running CICS in a multiple MVS/ESA SP 4.1 sysplex.

User exits
: Ensure that the current user exit programs run in an XRF environment. You should check the function, timing, and use of data of each exit program.

VTAM
: You must define one generic and two specific applids for each active-alternate pair.

  In multi-MVS operations, you need XDOMAIN definitions for CICS systems and logical units. These enable LUs owned by either MVS, or by a third MVS in a CMC configuration, to log on to CICS. They also enable CICS to acquire logical units after takeover.

  For VTAM information, see "VTAM and NCP considerations for active and alternate" in topic 5.1.

Workload on second MVS image
: Consider the effects of the workload on the second MVS after a takeover. See topic 3.3.3.

XCF PR/SM Policy
: Consider the takeover policy you are going to adopt. The information given in this book may need to be augmented by referring to the MVS books listed under "Books from related libraries" in topic PREFACE.3.

XXRSTAT exit
: Create a user exit program for the XXRSTAT exit, if required. See topic 6.3. The *CICS/ESA Customization Guide* contains the definitive product-sensitive programming interface information about global user exits.

B.0 Appendix B. Sample XRF implementations

In this appendix there are two sample implementations:

1. A single CICS region with an alternate in a second MVS image

2. An MRO configuration, with a dependent region, a master region, and a coordinator region, with actives and alternates in separate MVS images.

This appendix gives an overview of the SIT and SIT overrides, and CLT definitions. If you need more information about the SIT and CLT, see
"Defining CICS for XRF" in topic 6.0. The *CICS/ESA System Definition*
Guide contains a sample startup job stream.

In the following examples it is assumed that SIT overrides are entered using SYSIN and not the CONSOLE.

Subtopics:

B.1 Single CICS implementation

In this example, the operator is requested to confirm takeover when the surveillance signal is lost. If a takeover occurs because the active CICS issues "signoff abnormal", or if a CEBT PERFORM TAKEOVER command is issued, the alternate tries to take over automatically. This is done by specifying TAKEOVR=MANUAL in the SIT.

In this example, CICS1 is started as the active and CICS2 as the alternate.

Subtopics:

B.1.1 SIT and SIT overrides for a single CICS system

The SIT (DFHSITAA) and SIT overrides (CICS jobs JOB1 and JOB2) are as follows:

DFHSITAA

```
DFHSIT .....
    ,SUFFIX=AA
    ,XRF=YES
    ,START=STANDBY              /* (May be altered by override)
    ,APPLID=(CICS,CICS1)        /* (May be altered by override)
    ,ADI=40                     /* (Alternate only)
    ,PDI=40                     /* (Active only)
    ,TAKEOVR=MANUAL             /* (Alternate only)
    ,CLT=01                     /* (Alternate only)
    ,JESDI=35                   /* (Alternate only)
    ,AUTCONN=0
    ,AIRDELAY=700               /* (Active only)
    ,RMTRAN=(trans1,trans2)     /* (Alternate only)
    ,XRFSOFF=NOFORCE            /* (Active only)
    ,XRFSTME=5                  /* (Alternate only)
    ,.....
```

CICS job JOB1: The SIT overrides in JOB1 required to initialize CICS1 as the active on MVS1 are as follows. SIT parameters for an alternate are ignored during an active startup. If you want to start CICS1 as an alternate, remove the START=AUTO override from the SYSIN data.

```
//JOB1  JOB
 ...
//SYSIN  DD *
              ....
             ,SIT=AA
             ,START=AUTO               /* (Could be COLD or EMER)
             ,APPLID=(CICS,CICS1)      /* (Not strictly necessary, but
```

```
                ,.....                  /* (compatible with the job for
                ,.....                  /* (specific applid CICS2)
```

CICS job JOB2: This job initializes CICS2 as an alternate. When the alternate starts up, it ignores SIT operands for an active until it takes over and becomes an active itself. Then the SIT parameters for an active apply to it.

```
//JOB2  JOB
 ...
//SYSIN  DD *
                ....
               ,SIT=AA
               ,APPLID=(CICS,CICS2)
               ,.....
```

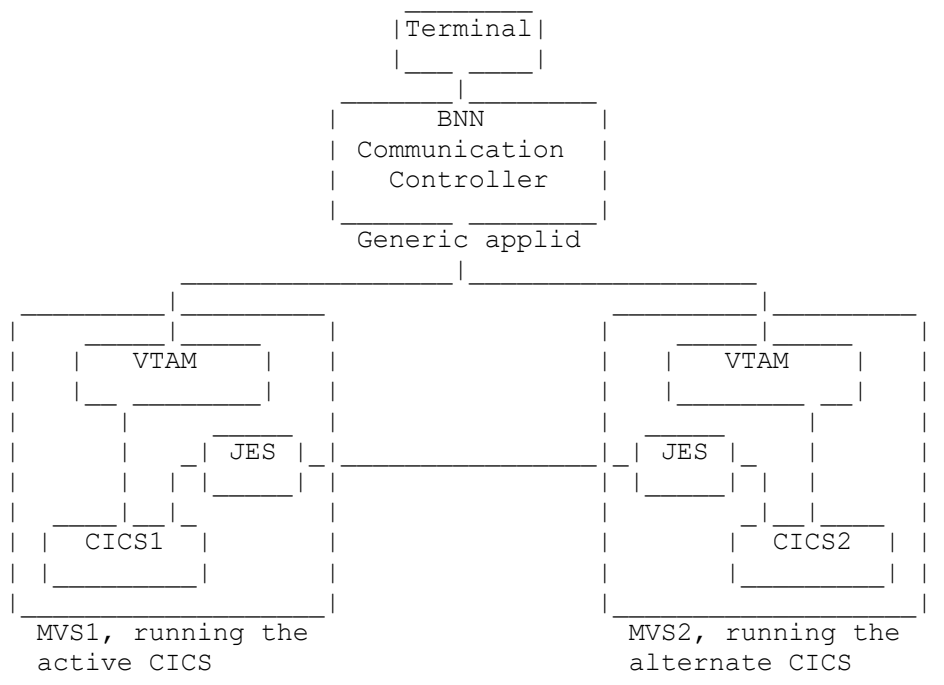

Figure 30. Sample single CICS implementation

B.1.2 CLT for a single CICS system

The sample CLT shown below is intended for use by either JOB1 or JOB2 running as an alternate. The CLT is processed by an alternate only at takeover time.

Each alternate uses the CLT entries that apply to its specific applid. The FORALT option indicates that the entries that follow it are for the systems with the specific applids shown in the FORALT option. Each system using this CLT will have been initialized with the START=STANDBY and CLT=01 parameters.

The sample CLT demonstrates that a single CLT, with one sequence of commands and messages, can be used for both CICS jobs. This is possible here because both jobs execute the same set of commands and messages. If you wanted to issue different commands or send messages that depend on which job is taking over, you could still use a single CLT, but you would have a separate LISTSTART and LISTEND for each of the specific applids.

The sample CLT for a single CICS system is as follows:

```
DFHCLT01 DFHCLT TYPE=INITIAL,                                          *
               JES=JES2,                JES variant                    *
               JESCHAR=$,               Prefix char for JES cmds       *
               SUFFIX=01,               CLT suffix                     *
               JESID=((MVS1,JES2,1),    SMFID JES-SSNM SPOOL-NUM MVS1  *
               (MVS2,JES2,2))           SMFID JES-SSNM SPOOL-NUM MVS2
*
label    DFHCLT TYPE=LISTSTART,                                        *
               FORALT=((CICS1,JOB2),     Alternate system applid      *
               (CICS2,JOB1))             Name of job it is allowed
*                                        to cancel
         DFHCLT TYPE=WTO,                Put out a console message     *
               WTOL=MSG1
MSG1     WTO 'CICS TAKEOVER IN PROGRESS,PLEASE SWITCH LOCALS',         *
               ROUTCDE=(1),                                            *
               DESC=(number),                                          *
               MF=L


*
         DFHCLT TYPE=LISTEND
*
         DFHCLT TYPE=FINAL
         END
```

B.2 MRO CICS implementation

In this example, shown in Figure 31, there are three MRO-connected regions: dependent, master, and coordinator. If either the master or coordinator region fails, there is an automatic takeover. If the dependent region fails by itself, it is restarted in place by an operator or by the overseer.

The operator can initiate a takeover of all the regions by issuing a CEBT PERFORM TAKEOVER command to the coordinator region. By doing this, all regions are taken over by their alternates. A CEBT PERFORM TAKEOVER command issued to a dependent region does not cause a takeover of all the regions. To allow this would require additional entries for the dependent portions of the CLT. There would be no benefit in having extra entries, because the advantage of issuing the CEBT command to the coordinating region is that doing so minimizes the flow of commands from the CLTs.

Note: For this example, only three regions are shown, one of each kind. Adding more dependent regions to the example would not illustrate anything new, because the entries for each of them would be basically the same. However, in a real system with only three regions, you probably would not want the added complexity of a coordinator because it saves very few CLT commands.

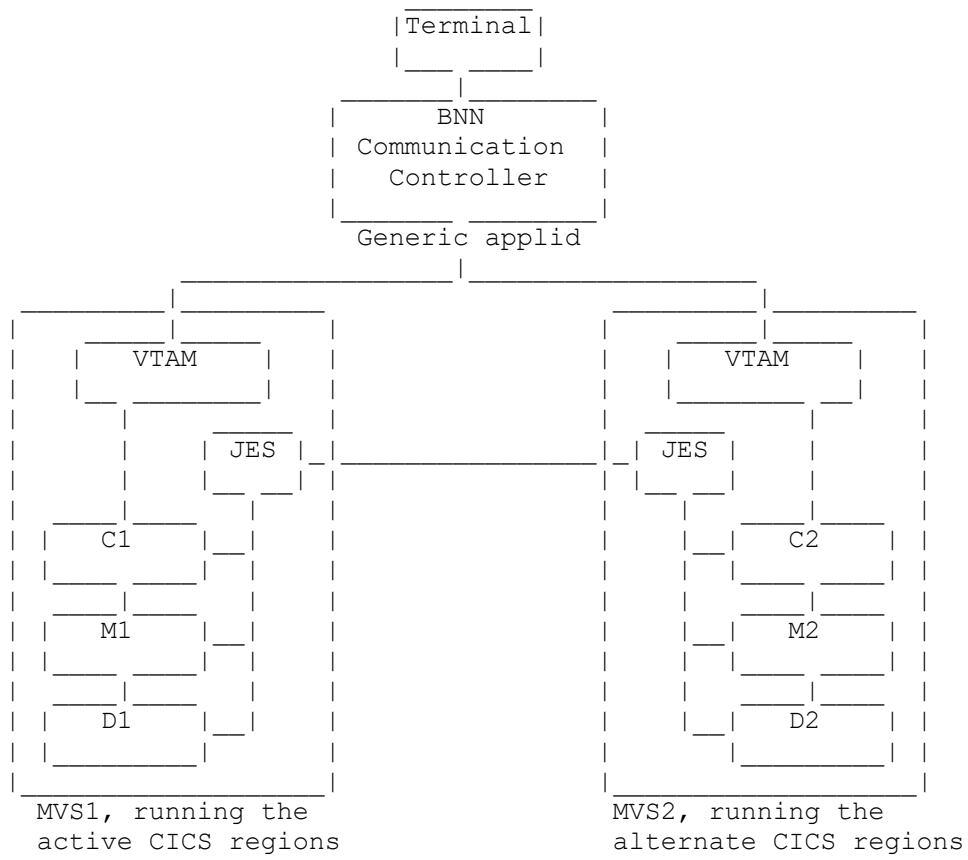

Figure 31. Sample MRO CICS implementation

Subtopics:

- B.2.1 SIT and SIT overrides for MRO-connected regions
- B.2.2 CLT for MRO-connected regions

B.2.1 SIT and SIT overrides for MRO-connected regions

Each active-alternate pair has its own SIT. As with the SIT for the single-region CICS system, SYSIN overrides are used to tailor the SIT.

Subtopics:

- B.2.1.1 CICS region C--the coordinator region
- B.2.1.2 CICS region M--the master region
- B.2.1.3 CICS region D--the dependent region

B.2.1.1 CICS region C--the coordinator region

DFHSITCO

```
  DFHSIT .....
      ,SUFFIX=CO
      ,XRF=YES
      ,START=STANDBY
      ,APPLID=(C,C1)
      ,ADI=20
      ,PDI=20
      ,TAKEOVR=AUTO
```

```
      ,CLT=02
      ,JESDI=25
      ,AUTCONN=0
      ,AIRDELAY=700
      ,RMTRAN=(trans1,trans2)
      ,XRFSOFF=NOFORCE
      ,XRFSTME=5
      ,.....
```

CICS job JOBC1: The following SIT overrides are required to initialize the active coordinator region on MVS1.

If you want to start JOBC1 as an alternate, you should remove the START=AUTO override. This applies to all of the jobs that follow that are initially started with START=AUTO.

```
  //JOBC1  JOB
   ...
  //SYSIN  DD *
                ....
               ,SIT=CO
               ,START=AUTO
               ,APPLID=(C,C1)
               ,.....
```

CICS job JOBC2

```
  //JOBC2  JOB
   ...
  //SYSIN  DD *
                ....
               ,SIT=CO
               ,APPLID=(C,C2)
               ,.....
```

B.2.1.2 CICS region M--the master region

DFHSITMA

```
DFHSIT .....
    ,SUFFIX=MA
    ,XRF=YES
    ,START=STANDBY
    ,APPLID=(M,M1)
    ,ADI=20
    ,PDI=20
    ,TAKEOVR=AUTO
    ,CLT=02
    ,JESDI=25
    ,AUTCONN=0
    ,AIRDELAY=700
    ,RMTRAN=(trans1,trans2)
    ,XRFSOFF=NOFORCE
    ,XRFSTME=5
    ,.....
```

CICS job JOBM1

```
//JOBM1  JOB
 ...
//SYSIN  DD *
                ....
               ,SIT=MA
               ,START=AUTO
               ,APPLID=(M,M1)
               ,.....
```

CICS job JOBM2

```
//JOBM2  JOB
 ...
//SYSIN  DD *
                ....
               ,SIT=MA
               ,APPLID=(M,M2)
               ,.....
```

B.2.1.3 CICS region D--the dependent region

DFHSITDE

```
DFHSIT .....
    ,SUFFIX=DE
    ,XRF=YES
    ,START=STANDBY
    ,APPLID=(D,D1)
    ,ADI=20
    ,PDI=20
    ,TAKEOVR=COMMAND
    ,CLT=02
    ,JESDI=25
    ,AUTCONN=0
    ,AIRDELAY=700
    ,RMTRAN=(trans1,trans2)
    ,XRFSOFF=NOFORCE
    ,XRFSTME=5
    ,.....
```

CICS job JOBD1

```
//JOBD1  JOB
 ...
//SYSIN  DD *
              ....
             ,SIT=DE
             ,START=AUTO
             ,APPLID=(D,D1)
             ,.....
```

CICS job JOBD2

```
//JOBD2  JOB
 ...
//SYSIN  DD *
              ....
             ,SIT=DE
             ,APPLID=(D,D2)
             ,.....
```

B.2.2 CLT for MRO-connected regions

This sample CLT is for use by all six jobs in the MRO group when they run as alternates.

If the alternate coordinator region is taking over, it uses CEBT to force the other regions to take over. If the master region fails and is being taken over by its alternate, that alternate forces the alternate coordinator to take over, and the coordinator instructs the other regions to take over. In this example, the command to the alternate master region is redundant, because it has already begun its takeover processing. But in a larger MRO complex, where the addition of a coordinator is more worthwhile, the number of redundant commands would not increase with the extra regions.

However, you might not want the added complexity of a coordinator. If there were no coordinator, each master region would contain two CEBT commands to the other regions in the complex.

```
*-------------------------------------------------------------------*
*
*Composite CLT for use with all six regions in this MRO-connected group
*
*-------------------------------------------------------------------*
*
DFHCLT02 DFHCLT TYPE=INITIAL,                                       *
               JES=JES2,              JES variant                   *
               JESCHAR=$,             Prefix char for JES cmds      *
               SUFFIX=02,             CLT suffix (CLT02 for both MVSs *
               JESID=((MVS1,JES2,1),  SMFID JES-SSNM SPOOL-NUM MVS1  *
               (MVS2,JES2,2))         SMFID JES-SSNM SPOOL-NUM MVS2
*
*-------------------------------------------------------------------*
*
*The following CLT entries govern a takeover of the MRO group from C1,
*M1, D1 running on one MVS to C2, M2, D2 running on the other MVS
*
*-------------------------------------------------------------------*
*
COORD1   DFHCLT TYPE=LISTSTART,                                     *
               FORALT=((C2,JOBC1))    Alternate system applid
*                                     Name of job it is allowed
*                                     to cancel
         DFHCLT TYPE=COMMAND,         M2 takeover from M1           *
               COMMAND='MODIFY JOBM2,CEBT PERFORM TAKEOVER'
         DFHCLT TYPE=COMMAND,         D2 takeover from D1           *
               COMMAND='MODIFY JOBD2,CEBT PERFORM TAKEOVER'
*
         DFHCLT TYPE=COMMAND,         Insert a user command
*                                     for any job running under MVS  *
               COMMAND='MODIFY USERJOB,USER COMMAND'
*
         DFHCLT TYPE=WTO,             Put out a console message     *
               WTOL=MSG1
```

```
MSG1     WTO 'NOTE TAKEOVER TO NUMBER 2 REGIONS',                      *
               ROUTCDE=(1),                                            *
               DESC=(number),                                          *
               MF=L
*
         DFHCLT TYPE=LISTEND
*
MASTER1  DFHCLT TYPE=LISTSTART,                                        *
               FORALT=((M2,JOBM1))     Alternate system applid
*                                      Name of job it is allowed
*                                      to cancel
         DFHCLT TYPE=COMMAND,          C2 take over the complex        *
               COMMAND='MODIFY JOBC2,CEBT PERFORM TAKEOVER'
*
         DFHCLT TYPE=LISTEND
*
*
DEPEND1  DFHCLT TYPE=LISTSTART,                                        *
               FORALT=((D2,JOBD1))     Alternate system applid         *
*                                      Name of job it is allowed
*                                      to cancel
*
         DFHCLT TYPE=LISTEND
*
*-----------------------------------------------------------------------*
*
*The following CLT entries govern a takeover of the MRO group from C2,
*M2, D2 running on one MVS to C1, M1, D1 running on the other MVS
*
*-----------------------------------------------------------------------*
*
COORD2   DFHCLT TYPE=LISTSTART,                                        *
               FORALT=((C1,JOBC2)      Alternate system applid
*                                      Name of job it is allowed
*                                      to cancel
*
         DFHCLT TYPE=COMMAND,          M1 takeover from M2             *
               COMMAND='MODIFY JOBM1,CEBT PERFORM TAKEOVER'
         DFHCLT TYPE=COMMAND,          D1 takeover from D2             *
               COMMAND='MODIFY JOBD1,CEBT PERFORM TAKEOVER'
*
         DFHCLT TYPE=COMMAND,          Insert a user command
*                                      for any job running under MVS   *
               COMMAND='MODIFY USERJOB,USER COMMAND'
*
         DFHCLT TYPE=WTO,              Put out a console message       *
               WTOL=MSG2


MSG2     WTO 'NOTE TAKEOVER TO NUMBER 1 REGIONS',                      *
               ROUTCDE=(1),                                            *
               DESC=(number),                                          *
               MF=L
*
         DFHCLT TYPE=LISTEND
```

```
*
MASTER2  DFHCLT TYPE=LISTSTART,                                        *
               FORALT=((M1,JOBM2))     Alternate system applid
*                                      Name of job it is allowed
*                                      to cancel
*
         DFHCLT TYPE=COMMAND,          C1 take over the complex        *
               COMMAND='MODIFY JOBC1,CEBT PERFORM TAKEOVER'
*
         DFHCLT TYPE=LISTEND
*
*
DEPEND2  DFHCLT TYPE=LISTSTART,                                        *
               FORALT=((D1,JOBD2))     Alternate system applid         *
*                                      Name of job it is allowed
*                                      to cancel
*
         DFHCLT TYPE=LISTEND
*
         DFHCLT TYPE=FINAL
         END
```

Glossary

This is a selective glossary of XRF terms.

active-alternate pair. The active and alternate CICS systems that work together to provide an XRF service to the end user.

active CICS system (active). The CICS system that currently supports user processing requests.

active session. The session that connects the active CICS to the end user of a class 1 terminal.

alternate CICS system (alternate). The CICS system that stands by, waiting to take over the user workload when the active CICS system fails or a takeover is initiated.

automatic reconfiguration facility (ARF). Also known as XCF PR/SM Policy. In a multisystem sysplex on PR/SM, the actions that XCF takes when one MVS system in the sysplex fails. This policy provides high availability for multisystem applications in the sysplex.

backup session. The session built by VTAM to the alternate CICS system for XRF-capable terminals, used after a takeover to reestablish service to the terminals.

boundary network node (BNN). The point at which terminal sessions are switched from the failing active to the new active. The communication controller at the BNN must be able to operate in an XRF configuration.

CAVM. See CICS availability manager

CEBT. A CICS-supplied transaction. that the operator can issue from the MVS console to control an alternate CICS system.

central electronic complex (CEC). A processing system, such as a 3090, with one or more central processors, running under the control of a single MVS/ESA operating system. Or, a physical partition of such a processing system, with its own copy of the operating system. The processor can be either a uniprocessor or a multiprocessor (including a dyadic processor).

CICS availability manager (CAVM). The mechanism that provides integrity for a CICS system with XRF. The CAVM uses the control file and the message file to handle communication between the active and alternate systems.

class 1 terminal. A remote SNA VTAM terminal connected through a boundary network node IBM 3745/3725/3720 Communication Controller with an NCP that supports XRF. Such a terminal has a backup session to the alternate CICS system.

class 2 terminal. A terminal belonging to a class mainly comprised of VTAM terminals that are not eligible for class 1. For these terminals, the alternate tracks the session, and attempts reestablishment after takeover.

class 3 terminal. A terminal belonging to a class mainly comprised of TCAM (DCB) terminals. These terminals lose their sessions at takeover.

command list table (CLT). A table that contains information required by the alternate CICS system during takeover.

communication management configuration (CMC). A configuration in which the VTAM subsystem that owns the terminals is in a different MVS image from the active or the alternate CICS system.

control data set. A data set that ensures XRF system integrity by allowing only one active CICS system to access a particular set of resources. It is used by the active and the alternate CICS systems to monitor each other's state.

coordinator. A region, in a multi-MVS MRO XRF configuration, that receives requests from master regions to initiate a takeover. It then instructs all the alternate regions to take over.

dependent. A region, in a multi-MVS MRO XRF configuration, that receives commands from a master or coordinator region at takeover time. It cannot initiate a takeover.

extended recovery facility (XRF). A software function that minimizes the effect of various failures on the end users of the system.

generic applid. The name by which the active-alternate pair of CICS systems is known to the end user. In VTAM terms, this is the USERVAR. The generic name is also used in intersystem communication.

initialization. The stage of the XRF process when the active or the alternate CICS system is started, signs on to the control data set, and begins to issue its surveillance signal.

JES (job entry subsystem). The subsystem used in CICS with XRF to route commands and queries from the alternate to the active system.

logical partition. A partition, in a CEC, capable of running its own MVS image.

master. A region, in a multi-MVS MRO XRF configuration, that issues commands to dependent regions at takeover time. See also **coordinator.**

message data set. A data set used by the active CICS system to transmit messages to the alternate CICS system.

MVS image. A single copy of the MVS operating system.

multi-MVS environment. An environment, in one or more CECs, that supports more than one MVS image.

NCP (network control program). A control program that resides in a BNN communication controller. It builds the backup sessions to the alternate CICS system for XRF-capable terminals.

NetView. A network management product that can provide rapid notification of events and automated operations.

network control program (NCP). See NCP

overseer. A program running in its own address space that provides status information about active and alternate CICS systems. The overseer can be used to automate the restart of failed active regions in place.

physical partition. Part of a CEC that operates as a CEC in its own right, with its own copy of the operating system.

planned takeover. A planned shutdown of the active CICS system, and takeover by the alternate, perhaps for maintenance or operational reasons.

PR/SM. The PR/SM feature on a 3090 Processor Complex offers flexible partitioning into a number of logical partitions, each of which can run its own MVS image.

restart in place. The restart of a failed active CICS system, instead of initiating a takeover to its alternate.

specific applid. The name used by the active CICS system when it opens the VTAM ACB.

surveillance. The active and alternate CICS systems use the CAVM surveillance mechanism to monitor each other's state.

surveillance signal. The signal continuously written to the CAVM data sets by the active and alternate CICS systems to inform the other of its state.

synchronization. The stage of the XRF process when the active and the alternate are both initialized, are aware of each other's presence, and the alternate is ready to begin tracking.

sysplex. A sysplex (SYStem comPLEX) is a set of one or more MVS systems

that is given an XCF name and in which programs in the systems can then use XCF services.

takeover. The shift of the workload from the active to the alternate CICS system, and the switching of resources needed for this to happen.

tracking. The monitoring of the terminals in the active CICS system by the alternate, through the message data set.

USERVAR. A VTAM function that enables VTAM to connect XRF-capable terminals to the active CICS system by recording the relationship between the generic and specific applids.

XCF (cross-system coupling facility). A set of services necessary to support a multisystem application (for example, CICS).

XRF. See extended recovery facility.

XRF-capable terminal. See class 1 terminal

XRF complex. The CEC, or CECs, and licensed programs that provide an XRF service.

3745/3725/3720 Communication Controller. The IBM 3745, 3725, or 3720 Communication Controller that may be used at the boundary network node in an XRF configuration. The 37xx and its resident NCP can switch terminals automatically from the failing active CICS system to the alternate.

Index

Numerics

A

B

C

D

E

F

G

H

I

J

L

M

N

O

P

Q

R

S

T

U

V

W

X