CICS® Transaction Server for VSE/ESA™

XRF Guide

Release 1

SC33-1671-01

CICS® Transaction Server for VSE/ESA™

XRF Guide

Release 1

SC33-1671-01

Note!

Before using this information and the product it supports, be sure to read the general information under "Notices" on page 93.

First Edition (June 1999)

This edition applies to Release 1 of CICS Transaction Server for VSE/ESA, program number 5648-054, and to all subsequent versions, releases, and modifications until otherwise indicated in new editions. Make sure you are using the correct edition for the level of the product.

The CICS for VSE/ESA Version 2.3 edition remains applicable and current for users of CICS for VSE/ESA Version 2.3.

Order publications through your IBM representative or the IBM branch office serving your locality.

At the back of this publication is a page entitled "Sending your comments to IBM". If you want to make any comments, please use one of the methods described there.

Contents

Preface

What this book is about

This book is intended to help you to understand the extended recovery facility (XRF) function. It contains guidance about planning, setting up, and running a CICS system 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 Resource Definition 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.

How to use this book

Chapters 1 through 3 introduce the XRF concept and explain how CICS with XRF works. Chapter 4 suggests possible configurations. Chapters 5 and 6 give more detailed guidance to help you set up XRF. How XRF relates to other products is discussed in Chapter 7.

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 on page /GLOSSY/. There is a general glossary of CICS terms in the *CICS Glossary* GC33-1649.

New terms are explained when they first occur.

Notes on terminology

The terms listed in Table 1 are commonly used in the CICS Transaction Server for VSE/ESA Release 1 library. See the *CICS Glossary* for a comprehensive definition of terminology.

Table 1 (Page 1 of 2). Commonly used words and abbreviations

Term	**Definition (and abbreviation if appropriate)**
$(the dollar symbol)	In the character sets and programming examples given in this book, the dollar symbol ($) is used as a national currency symbol and is assumed to be assigned the EBCDIC code point X'5B'. In some countries a different currency symbol, for example the pound symbol (£), or the yen symbol (¥), is assigned the same EBCDIC code point. In these countries, the appropriate currency symbol should be used instead of the dollar symbol.
BSM	BSM is used to indicate the basic security management supplied as part of the VSE/ESA product. It is RACROUTE-compliant, and provides the following functions: • Signon security • Transaction attach security
C	The C programming language
CICSplex	A CICSplex consists of two or more regions that are linked using CICS intercommunication facilities. Typically, a CICSplex has at least one terminal-owning region (TOR), more than one application-owning region (AOR), and may have one or more regions that own the resources accessed by the AORs
CICS Data Management Facility	The new facility to which all statistics and monitoring data is written, generally referred to as "DMF"
CICS/VSE	The CICS product running under the VSE/ESA operating system, frequently referred to as simply "CICS"
COBOL	The COBOL programming language
DB2 for VSE/ESA	Database 2 for VSE/ESA which was previously known as "SQL/DS".

Table 1 (Page 2 of 2). Commonly used words and abbreviations

Term	Definition (and abbreviation if appropriate)
ESM	ESM is used to indicate a RACROUTE-compliant external security manager that supports some or all of the following functions: • Signon security • Transaction attach security • Resource security • Command security • Non-terminal security • Surrogate user security • MRO/ISC security (MRO, LU6.1 or LU6.2) • FEPI security.
FOR (file-owning region)—also known as a DOR (data-owning region)	A CICS region whose primary purpose is to manage VSAM and DAM files, and VSAM data tables, through function provided by the CICS file control program.
IBM C for VSE/ESA	The Language Environment-conforming version of the C programming language compiler. Generally referred to as “C/VSE”.
IBM COBOL for VSE/ESA	The Language Environment-conforming version of the COBOL programming language compiler. Generally referred to as “COBOL/VSE”.
IBM PL/I for VSE/ESA	The Language Environment-conforming version of the PL/I programming language compiler. Generally referred to as “PL/I VSE”.
IBM Language Environment for VSE/ESA	The common runtime interface for all LE-conforming languages. Generally referred to as “LE/VSE”.
PL/I	The PL/I programming language
VSE/POWER	Priority Output Writers Execution processors and input Readers. The VSE/ESA spooling subsystem which is exploited by the report controller.
VSE/ESA System Authorization Facility	The new VSE facility which enables the new security mechanisms in CICS, generally referred to as “SAF”
VSE/ESA Central Functions component	The new name for the VSE Advanced Function (AF) component
VSE/VTAM	“VTAM”

Determining if a publication is current

IBM regularly updates its publications with new and changed information. When first published, both the printed hardcopy and the BookManager softcopy versions of a publication are in step, but subsequent updates are normally made available in softcopy before they appear in hardcopy.

For CICS Transaction Server for VSE/ESA Release 1 books, softcopy updates appear regularly on the *Transaction Processing and Data Collection Kit* CD-ROM, SK2T-0730-xx and on the *VSE/ESA Collection Kit* CD-ROM, SK2T-0060-xx. Each reissue of the collection kit is indicated by an updated order number suffix (the -xx part). For example, collection kit SK2T-0730-20 is more up-to-date than SK2T-0730-19. The collection kit is also clearly dated on the front cover.

For individual books, the suffix number is incremented each time it is updated, so a publication with order number SC33-0667-02 is more recent than one with order number SC33-0667-01. Updates in the softcopy are clearly marked by revision codes (usually a "#" character) to the left of the changes.

Note that book suffix numbers are updated as a product moves from release to release, as well as for updates within a given release. Also, the date in the edition notice is not changed until the hardcopy is reissued.

Road map

Table 2. Getting started road map

If you want to...	Refer to...

Chapter 1. An overview of XRF

CICS® offers an Extended recovery facility, **XRF** environment that runs with any IBM® processor operating in Virtual Storage Extended (VSE) mode to give your CICS systems improved recovery performance and 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 Transaction Server for VSE/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 Transaction Server for VSE/ESA Release 1 active system with a CICS/VSE® 2.3 alternate. Also, if the active and alternate CICS systems are running on separate VSE operating systems, it is advisable to use the same level of VSE for both.

XRF environments

An XRF complex is made up of:

- The active and alternate CICS systems.
- The associated software, including the operating system (each copy of which may be called a VSE image) with POWER and VTAM.®
- One or more IBM 3745/3725/3720 communication controllers or terminal switching units.
- 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, or a physical partition of that machine, the term "CPC" is used. CPC is short for "central processing 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 processing system into a number of logical partitions.

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 image
- Improved availability when 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 image and VTAM, resulting in a multi-VSE environment.
  - 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).

    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 CPCs as shown in Figure 1 on page 3.

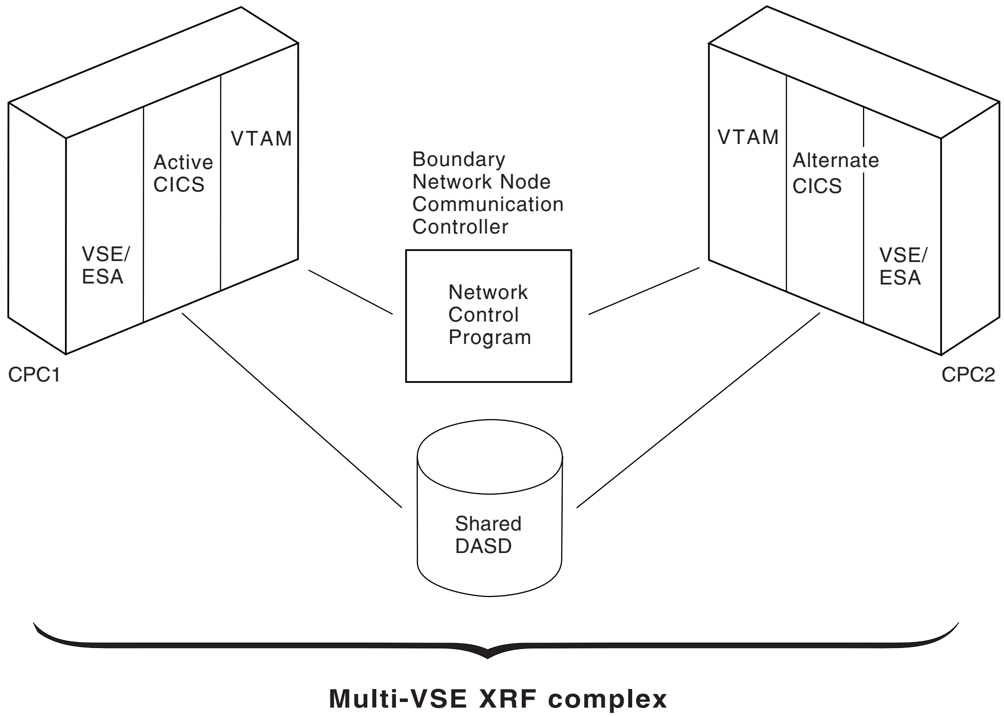

Figure 1. An XRF complex

Figure 1 illustrates the relationship between the various components of an XRF system - VSE, CICS and VTAM, each in a separate CPC, with shared DASD and the NCP connecting them.

A brief description of XRF

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

CICS 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.

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.
- 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 VSE images, the second is largely available for other work.

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 CICS system.

When VTAM terminals log on, the alternate **tracks** them, and after a takeover it tries to reestablish their sessions.

In a multi-VSE 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, end users 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 Chapter 5, "The terminal network" on page 37.

The takeover

A takeover might occur because of:

- CPC outage
- VSE outage
- VTAM outage
- CICS outage.

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

In a system running VSE/ESA™ under VM, a VM outage may be regarded as a CPC or VSE outage. VM outages are not discussed separately in this book.

In either case, XRF offers end users increased system availability. There is more information about the causes of a takeover in Chapter 2, "Types of outage handled by CICS with XRF" on page 7.

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.

This book describes 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.

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 POWER job queue.
- 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.

Chapter 2. Types of outage handled by CICS with XRF

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

The figures in this chapter show a multi-VSE environment. This environment could be provided by a single CEC or by two separate CPCs. The single CPC may be partitioned, logically (using the PR/SM feature) or physically, into a multi-VSE environment. This environment provides cover against VSE, VTAM, and CICS failures, as described in "XRF environments" on page 2. To guard against a CPC failure, you require two separate CPCs. XRF running in one VSE 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 Chapter 3, "How XRF works" on page 11.

CICS outage

Figure 2 illustrates an XRF system in which the active CICS fails, resulting in the loss of terminal sessions and the breakdown of information sent to the alternate via shared DASD.

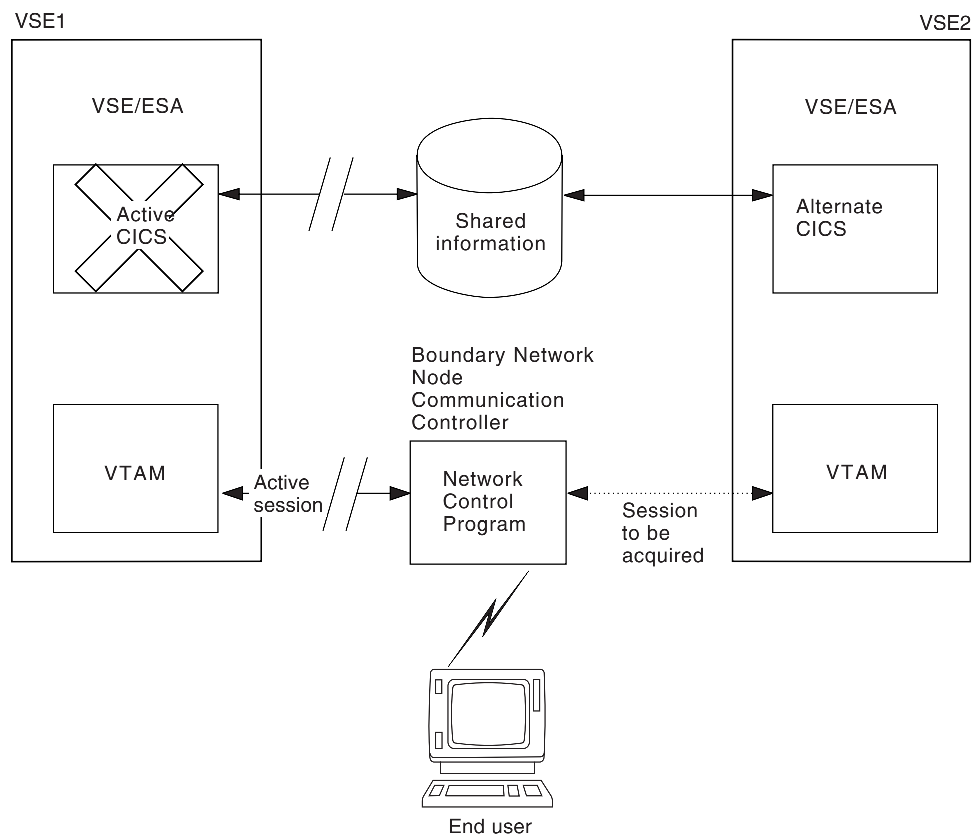

Figure 2. CICS outage

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

You do not need two VSE images to handle CICS outages. You can run XRF on a single VSE 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-VSE image, single-region XRF configuration" on page 32.

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.

VTAM outage

Figure 3 illustrates an XRF system in which the VTAM serving the active system fails, resulting in the loss of terminal sessions and the breakdown of information sent to the alternate via shared DASD.

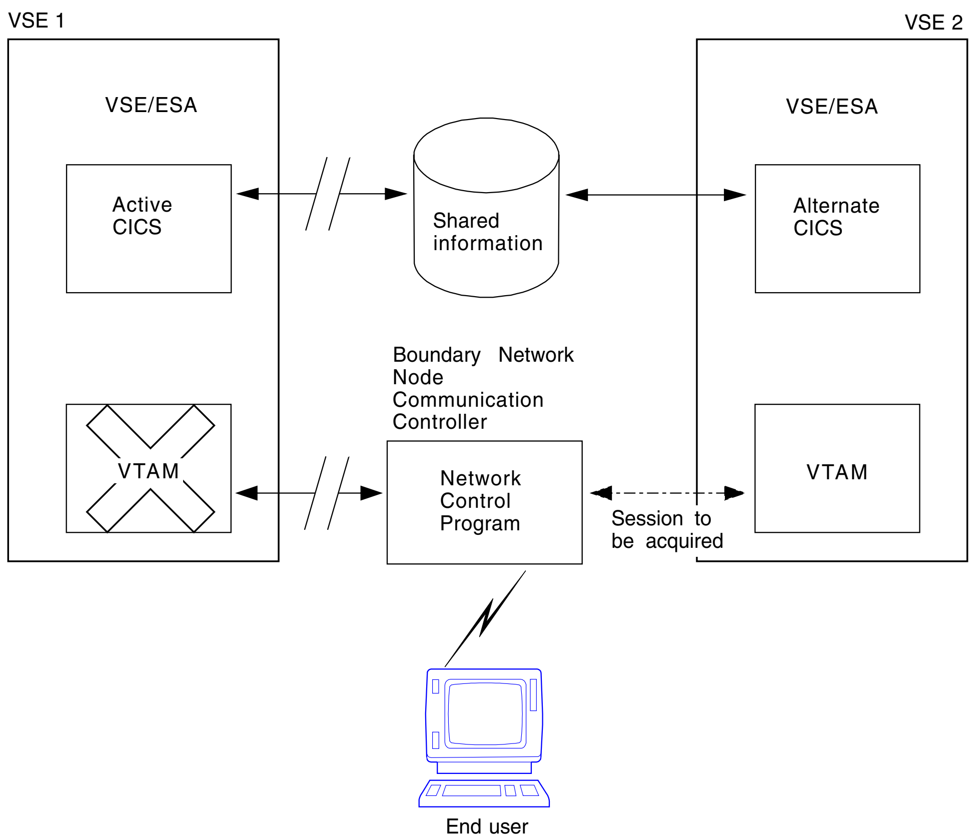

Figure 3. VTAM outage

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 a global user exit program for the XXRSTAT exit, or adding code to the overseer program to cause the takeover or other action. For more information about XXRSTAT and other global user exits, see the *CICS Customization Guide*.

If a takeover is not selected (the CICS default action), the active continues, in degraded mode.

In a multi-VSE environment, if the VTAM supporting the alternate fails, the active continues normally. Here, the alternate terminates, a new alternate can be started when VTAM has been restarted.

See “Multi-VSE, single-region XRF configuration” on page 25 and “User exit for VTAM failure” on page 62 for more information.

VSE outage

Figure 4 illustrates an XRF system in which the VSE serving the active system fails, resulting in the loss of terminal sessions and the breakdown of information sent to the alternate via shared DASD.

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

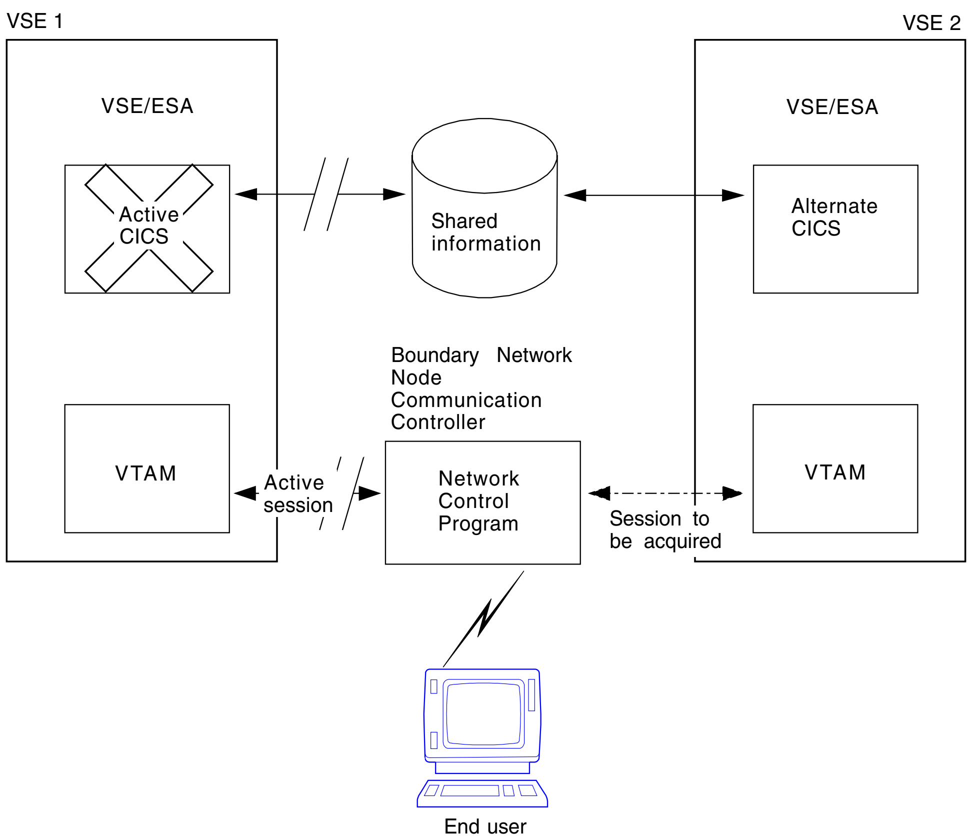

Figure 4. VSE outage

If you have two VSE images, you can run the active CICS on one VSE, and have the alternate CICS partially initialized on the other VSE. 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, a VSE (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 a VSE or hardware failure, the alternate is notified, through the CAVM, of the active’s failure and initiates a takeover.

For a VSE failure, the alternate cannot always determine the state of its active counterpart. In this case the operator confirms to the alternate that the active has failed due to VSE failure, and that a takeover can proceed. For more information, see “Checking for termination of the active” on page 19.

CPC outage

To cope with the failure of a CPC, and the other failures detailed previously, the alternate CICS has to run in a separate CPC. The second CPC 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.

For a CPC failure, like a VSE failure, 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 CPC failure and that a takeover can go ahead. For more information, see “Checking for termination of the active” on page 19.

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-VSE or two-CPC environment, XRF also helps you to take care of the maintenance of the CPCs 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 VSE image is also more restrictive, because some changes cannot be made without an IPL of VSE. This applies, for example, to maintenance of any CICS software that must reside in the SVA (shared virtual area).

For more information about the use of XRF takeovers as a maintenance aid, see Chapter 7, “XRF and other products” on page 67.

XRF gives you the flexibility, through a planned takeover, to choose when you carry out maintenance. You probably would not want to perform 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” on page 63.

Chapter 3. How XRF works

Before CICS/VSE 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.

An XRF sequence

Figure 5 on page 12 shows a possible XRF sequence. The stages in the sequence are described in the following five sections:

.
"1. Initialization" on page 13.
"2. Synchronization" on page 15.
"3. Surveillance and tracking" on page 15.
"4. Takeover" on page 16.
"5. After takeover" on page 21.

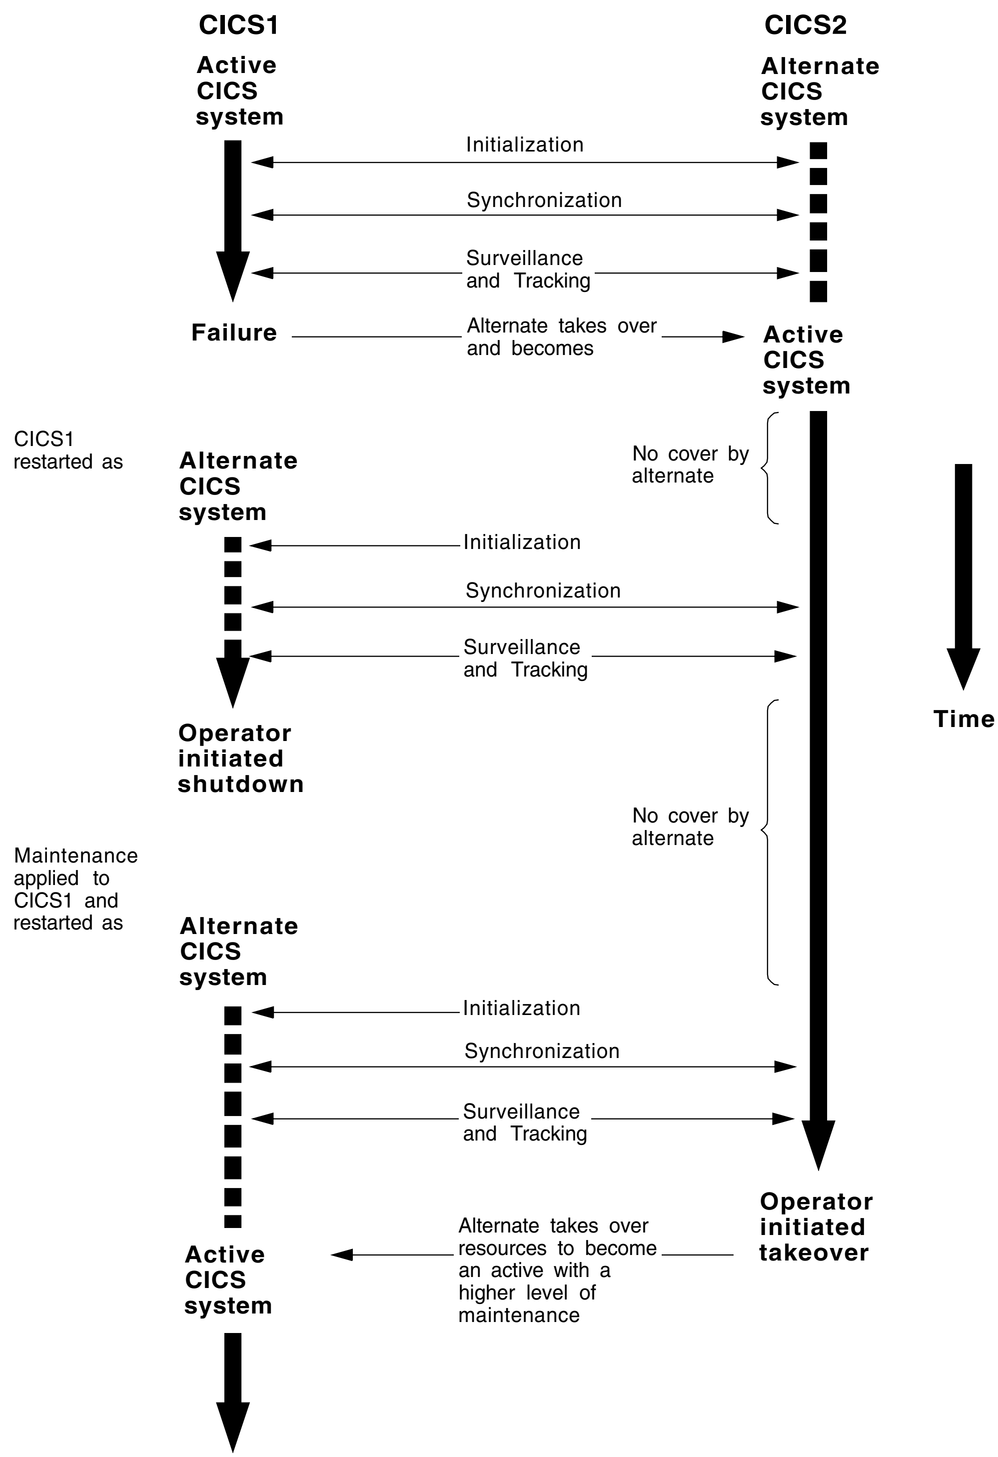

Figure 5. An XRF sequence

1. Initialization

Figure 6 illustrates the activities of the active and alternate CICS systems.

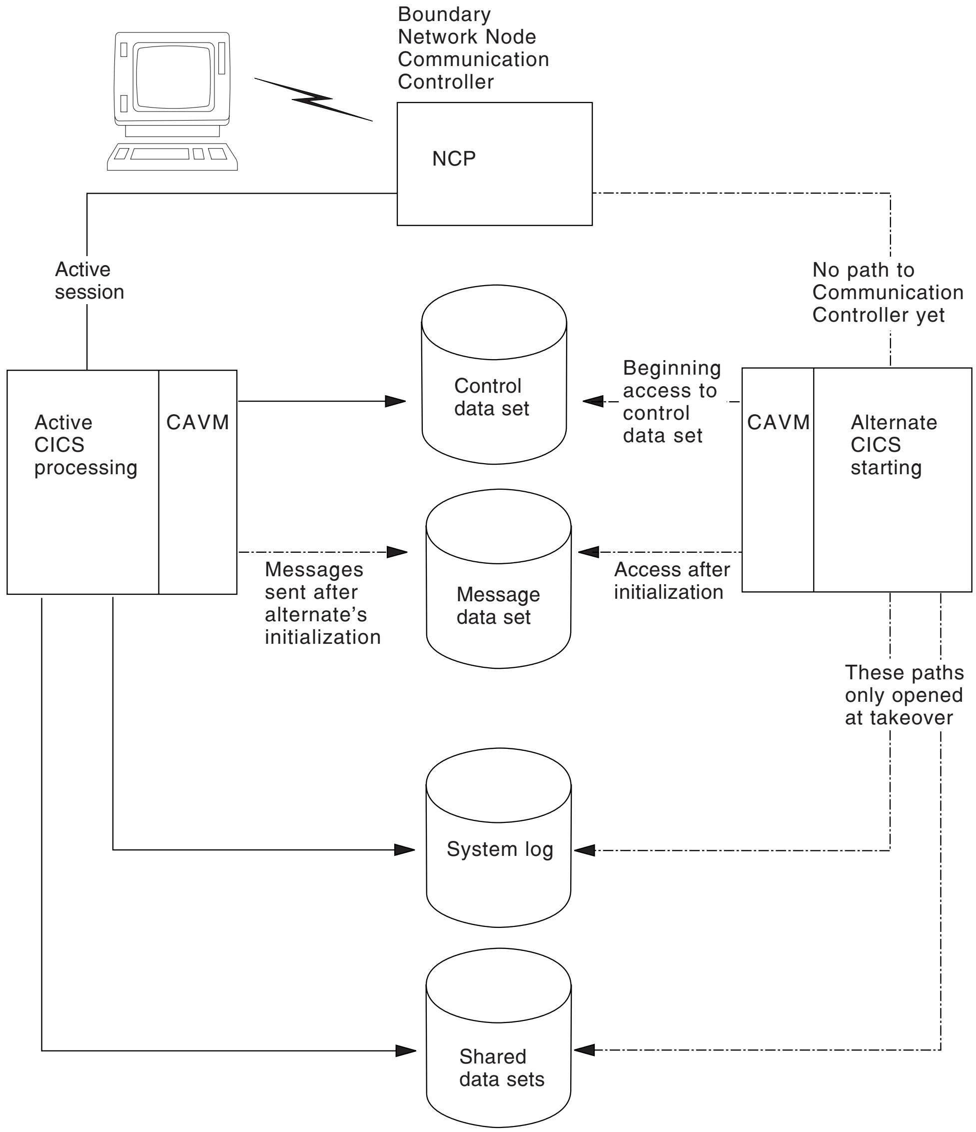

Figure 6. Initialization of alternate after active has started processing

Figure 6 shows that you need a pair of CICS systems to use XRF, the active and the alternate running in a shared POWER environment. 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.

The active and alternate systems 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 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 partition), 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 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, for example, users cannot log on to it, and it takes up very little resource. The only means of external communication with the alternate is through the VSE console communication interface or the overseer. The VSE console communication interface command is limited to a small set of CEBT commands, described in "Supplied transactions for controlling the alternate" on page 63. The overseer is described in "The overseer" on page 62. 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.

Running the active by itself

The active can run by itself without a matching alternate. This is shown in Figure 5 on page 12. 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.

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, and can track any 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. Surveillance and tracking

Most of the time, CICS with XRF is in the third stage: surveillance and tracking, as shown in Figure 7 on page 16.

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 CPC, VSE, 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.

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.

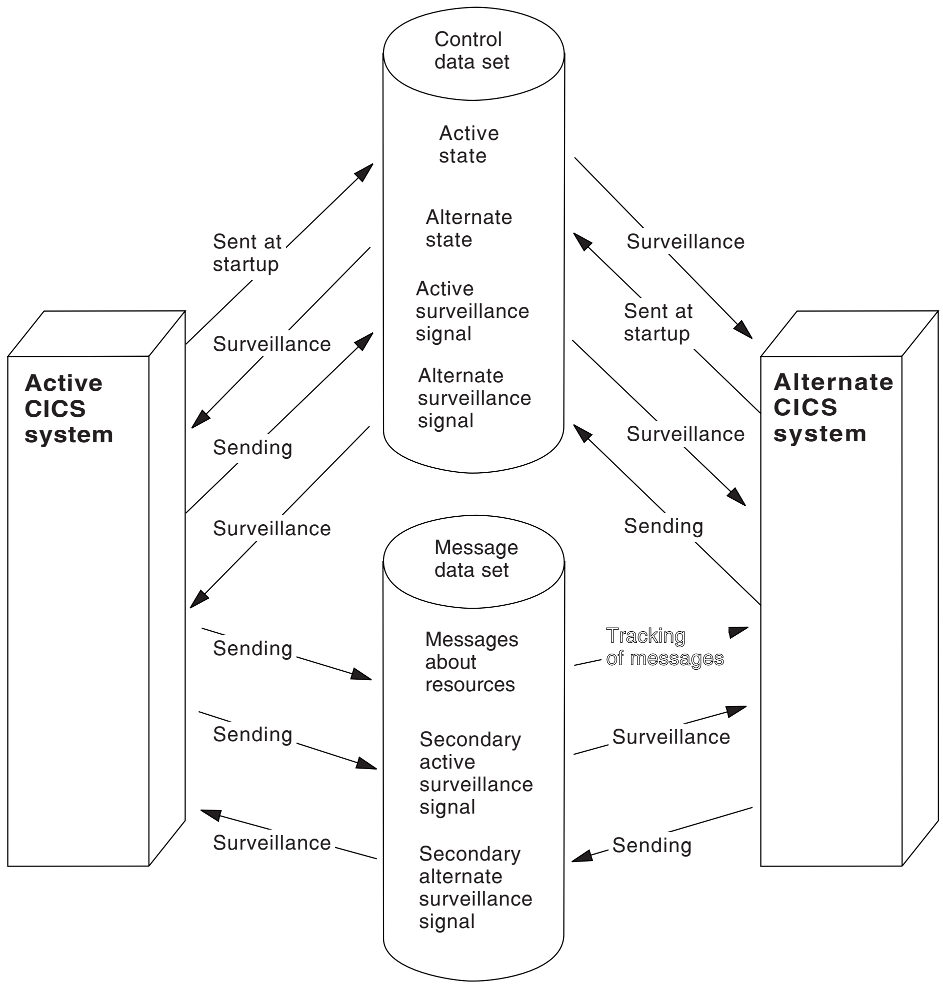

Figure 7. Use of the CAVM data sets for surveillance and tracking

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 partition 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” on page 51.

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: This command may be issued to the alternate by the operator, by another alternate taking over in a multi-VSE 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): 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.

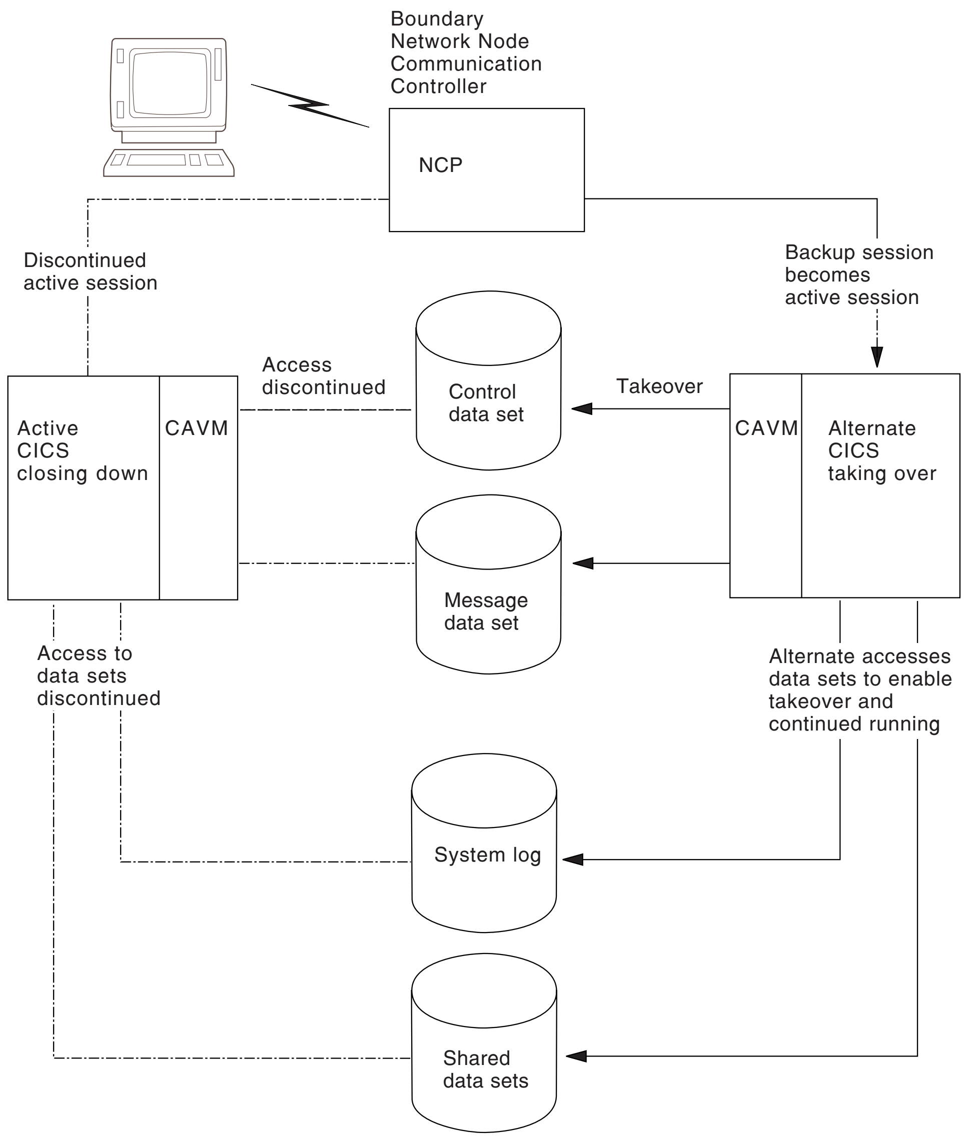

Figure 8. Takeover

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 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 VSE 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 VSE 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:

```
F NET,USERVAR,ID=generic-applid,VALUE=specific-applid
```

to redefine the CICS application name.

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 an abend (Abend Code 0206) 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 POWER routing command in a multi-VSE 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, VSE system commands, and messages to the operator. For more information, see "Command list table (CLT)" on page 54.

Checking for termination of the active

The alternate asks POWER 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 POWER replies that the job has terminated, the next phase, "Completing the takeover" on page 20, can start immediately.

If POWER replies that the job is still executing, the alternate continues to check the status until the interval defined by the XRFTODI system initialization parameter expires. After that interval, the alternate prompts the operator (with message DFHXA6561 or DFHXA6562) to investigate why the job has not stopped. There might be a POWER problem, or an authorization problem in the CLT. The alternate also offers this prompt if POWER is not running, or does not respond.

When active and alternate are running in different VSE images, POWER might continue to tell the alternate that the active job is still running even though the active's VSE or CPC 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 PCANCEL command has occurred.

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

Here, an internal record is kept that the VSE image, identified by its POWER SYSID and time and date of IPL, 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 POWER that the active job has terminated
- A previous reply to another alternate's message has already confirmed CPC or VSE failure.

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

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-VSE environment, if the time-of-day clock of the new active's VSE 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, becomes the active, and reestablishes sessions for VTAM terminals.

If the clock on the new active is fast compared to that of the new active, takeover resumes without waiting.

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 journal archiving, specified by the JOUROPT=AUTOARCH operand of the DFHJCT macro. For further guidance about automatic journal archiving, see the *CICS Operations and Utilities Guide*.

If you submit the archiving job for execution on the active's VSE, and that VSE 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 VSE.

Failure analysis

Diagnostic information about the failure of the active is provided by the termination VSE SDUMPS. 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.

CICS provides an offline dump analyzer, DFHPD410, to interpret and format the VSE SDUMP, and thereby simplify the task of problem determination. You are recommended to specify (via the JCL OPTION statement) SYSDUMP as the termination dump, both to provide adequate diagnostics, and to ensure that the active closes down as quickly as possible. For more information about how quickly the active closes down, see the ADI operand in Chapter 6, "Defining CICS for XRF" on page 49.

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

5. After takeover

In a multi-VSE environment, after the takeover, the operator manually switches any devices that need to be physically connected to the new active: perhaps local VTAM terminals, or other software outside the control of CICS.

Depending on the options you set, end users of VTAM 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.

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:

```
F NET,USERVAR,ID=generic-applid,VALUE=specific-applid
```

to change the entry of the local USERVAR. USERVAR values in remote VTAMs communicating with the local VTAM are changed by VTAM. See Chapter 5, "The terminal network" on page 37 for more information about USERVARs and applids.

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 value for the START system initialization 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.

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 page 63. The overseer, described on page 31, 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.

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” on page 47.

In a multi-VSE 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.

Performance

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

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 CPCs
- Model and characteristics of the CPCs
- Use of logical or physical partitioning
- Number of related partitions 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-VSE environment

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-VSE, large network environment.

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.

Workload on a second VSE image

This section is concerned with multi-VSE configurations. The second VSE (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 VSE" 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-CPC 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 VSE image may not be of sufficient size.) Restart in place of failing CICS partitions should be performed using the (TAKEOVR=COMMAND) system initialization parameter, but this can be automated using the overseer. See "The overseer" on page 62.

After a takeover, the new active provides the same service as the old. In a two-CPC environment, if the new active is in a CPC 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 CPC is smaller than the active's CPC. You might, for example, have to suspend some batch jobs temporarily.

If there are other subsystems running in the alternate's CPC, such as SQL/DS™, and they continue to run after takeover, performance will be degraded because the new active takes up more of the CPC's resources. A lot depends on how the VSE tuning parameters have been set. Refer to the *VSE/ESA Operation* manual, and the *VSE/POWER Administration and Operation* manual.

Chapter 4. XRF configurations

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

With each configuration diagram, there is a short explanation of the availability enhancements that each configuration provides. This chapter 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 Chapter 6, “Defining CICS for XRF” on page 49.

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

A single-VSE configuration provides protection against outages of the CICS partition. 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-CPC configuration if you want to provide protection against outages of the CPC, VSE, VTAM, and CICS.

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

Multi-VSE, single-region XRF configuration

The multi-VSE, single-region XRF configuration, shown in Figure 9 on page 26, offers increased protection against outages of VTAM, and CICS, and, in a two-CPC configuration, of the CPC. The active and alternate must be in the same equipment complex, so that they can share DASD, and they must be coupled by POWER. If a CPC 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.

VTAM is a special case. When VTAM fails, you can initiate a takeover, but 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 global 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. For more information about the XXRSTAT global user exit, see the *CICS Customization Guide*.

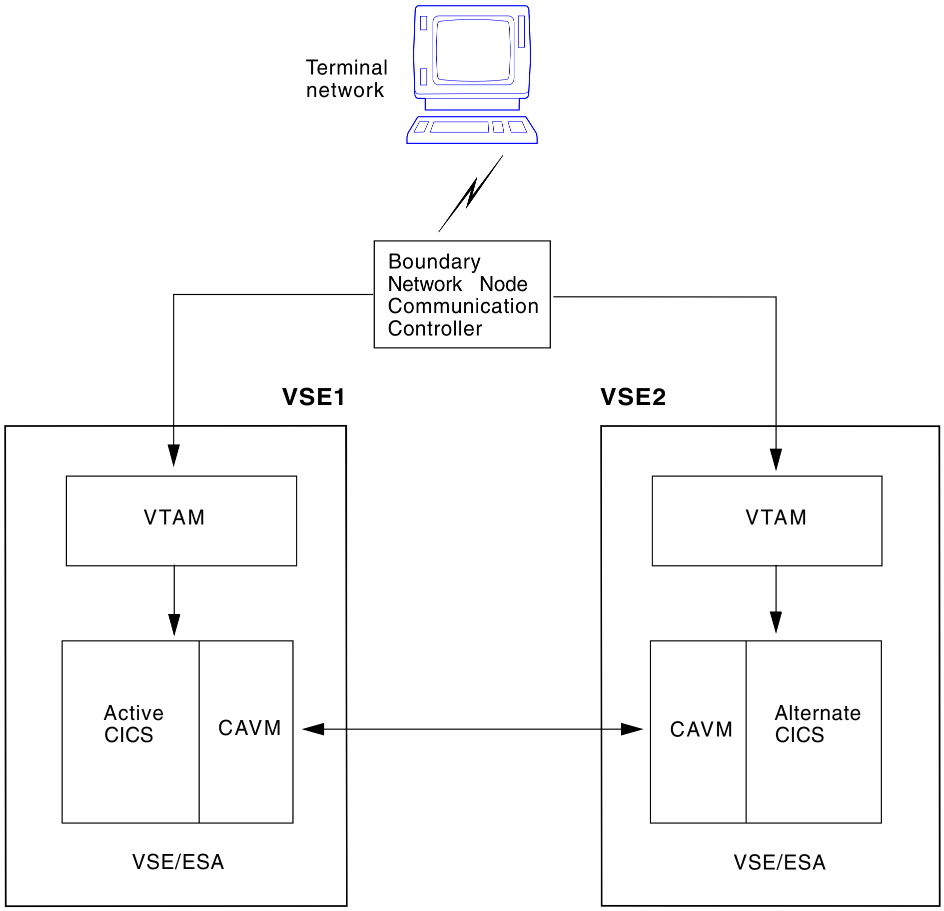

Figure 9. Multi-VSE, single-region XRF configuration

2. The overseer program, introduced more fully on page 31, 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 Chapter 6, “Defining CICS for XRF” on page 49. 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-CPC configuration) the alternate CPC 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, on page 30.

The end users of most VTAM 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.

Other terminals, such as local VTAM terminals, or remote non-SNA VTAM terminals—could also have faster recovery because of the quicker restart that XRF provides.

Multi-VSE, MRO XRF configuration

The multi-VSE MRO configuration offers increased availability against outages of VSE, VTAM, and CICS, and (in a two-CPC environment) of the CPC. The CICS system is divided into several MRO-connected active regions, each with its own alternate. Figure 10 on page 28 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 VSE images. If one alternate region takes over, all the related alternate regions must take over, because interregion communication does not operate across VSE images. A CPC or VSE failure clearly should result in a takeover of all the regions.

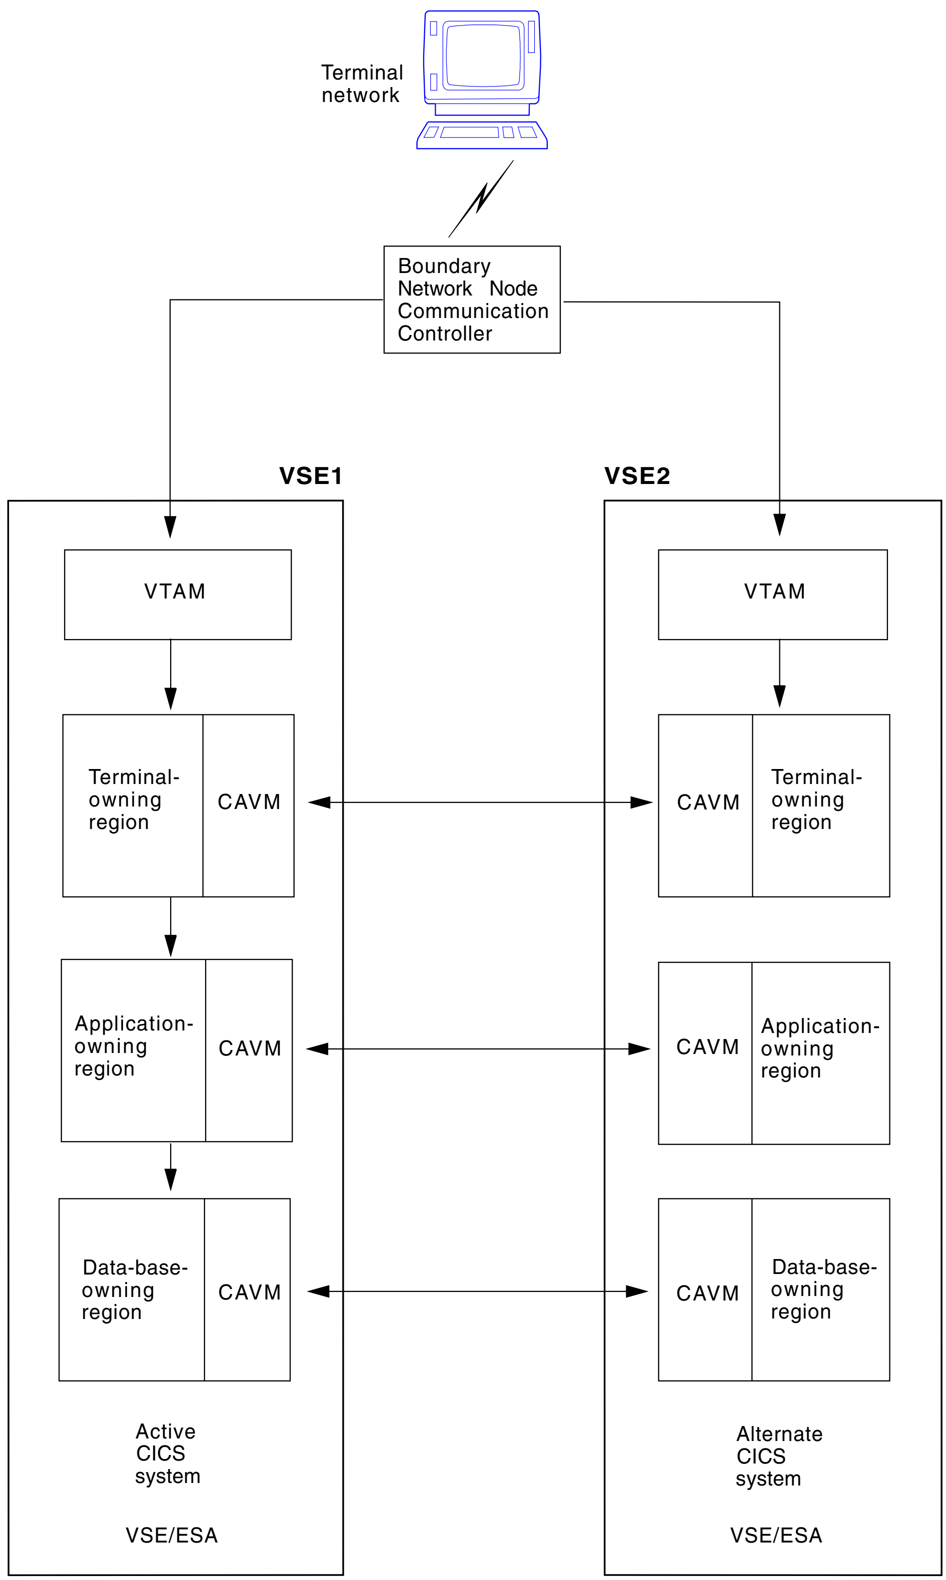

Figure 10. Multi-VSE, MRO XRF configuration

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

If a terminal-owning region experiences 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 terminal-owning regions, 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.

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 11.

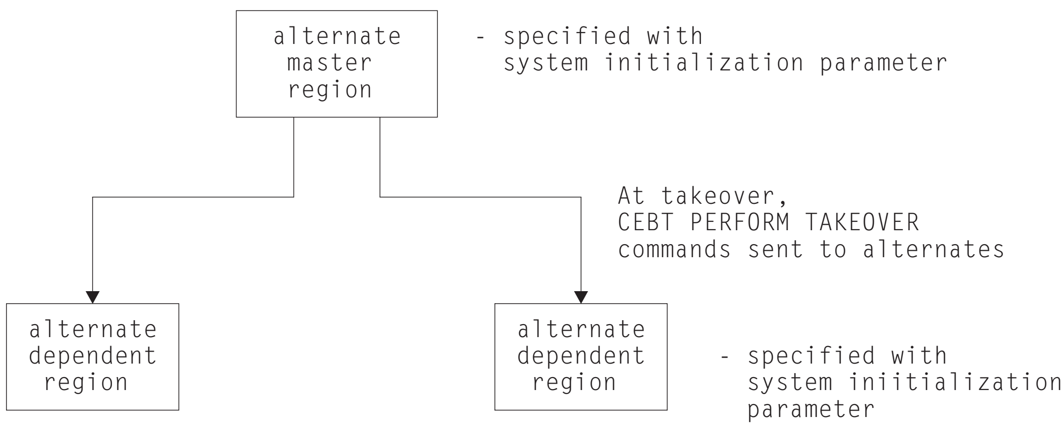

Figure 11. Hierarchy of one master and two dependent regions

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 system initialization parameter 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.

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 command list table (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 Chapter 6, "Defining CICS for XRF" on page 49.

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 VSE 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.

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

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 on page 62.

Single-VSE image, single-region XRF configuration

Figure 12 shows that, for CICS outages only, you can increase availability by using XRF in a single-VSE environment. Even if you have more than one VSE image available, you might choose a single-VSE configuration. This might be because of terminal-switching considerations, lack of capacity on the second VSE, or shared DASD limitations.

If you usually run XRF on two VSE images, but one is temporarily unavailable because of maintenance or because it has other work to do, you might choose a single-VSE 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 CPC, VSE, 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.

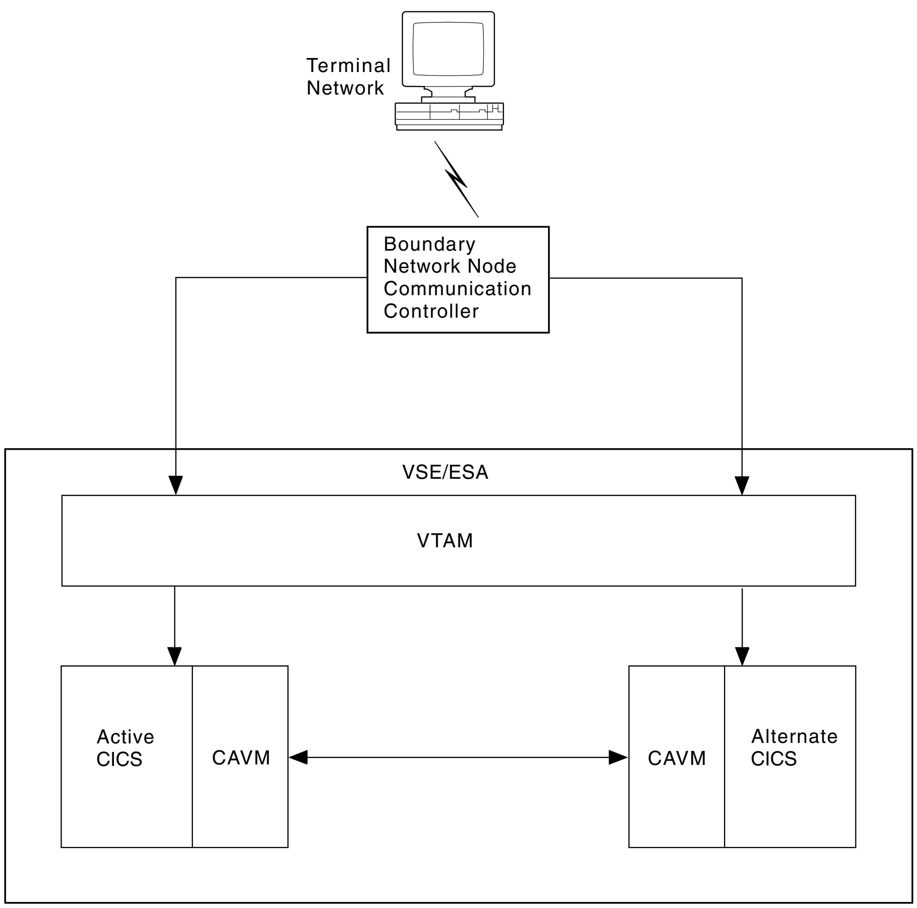

Figure 12. Single-VSE image, single-region XRF configuration

Single-VSE image, MRO XRF configuration

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

For each active region shown in Figure 13 on page 34—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-VSE, MRO XRF configuration" on page 27 (and more fully in Chapter 5, "The terminal network" on page 37). 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-VSE configurations.

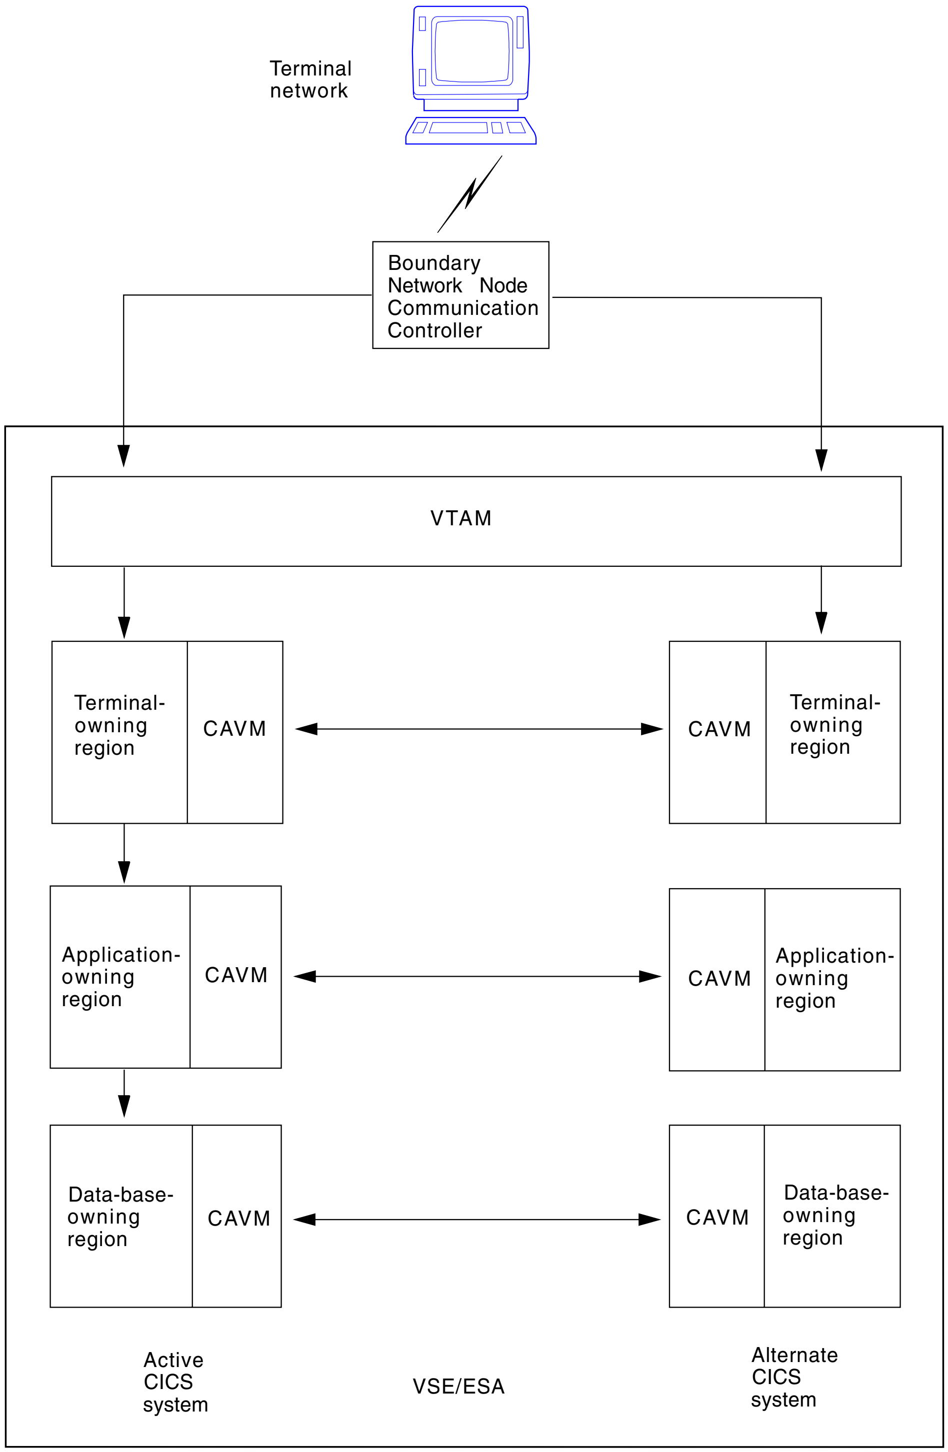

Figure 13. Single-VSE image, multiregion operation XRF configuration

Further configurations

This chapter 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 on page 31, 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-VSE configurations, but even if you are able to run XRF on two VSE images, there might be some systems that you would prefer to run with the active and alternate in the same VSE.

If you have three VSE images available, you could use the third for a new alternate CICS, if the failure of the first 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 VSE image.

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

If you have applications or databases that are rarely used, or applications that rarely fail, they could be placed in non-XRF regions. This non-XRF region could be a CICS Transaction Server for VSE/ESA system defined with XRF=NO as a system initialization parameter. 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-VSE 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.

Chapter 5. 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. XRF offers benefits to all terminals, because they may 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.

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)).

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.

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 VSE image
- How the terminal is defined by the systems programmer.

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.

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 Network Implementation Guide* and the *VTAM Operation* manual. This is particularly important if you are not accustomed to multi-VSE 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 interregion communication (IRC).

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)
CICS2 APPL AUTH=(ACQ)
```

The first part of the APPL statement defines to VTAM the specific applids (known to VTAM as the application ids).

The generic and specific applids have to be defined to CICS using the APPLID system initialization parameter. See page 49 for more information.

Controlling the use of the applids by USERVAR

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 F NET,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 14 on page 40 shows a set of definitions, with CICS1 as the active system and VTAM1 as the network owner. At startup, the active uses the:

```
F NET,USERVAR,ID=generic-applid,VALUE=specific-applid
```

command to set its specific applid (CICS1 in the Figure 14 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 "CICS" 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 command:

```
F NET,USERVAR,ID=generic-applid,VALUE=specific-applid
```

when it issues the:

```
SET LOGON START
```

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

USERVAR propagation to remote VTAMs

The USERVAR modified by a VTAM F NET 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.

Transferring a terminal session to the active

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.

Ownership of the network

In an XRF environment, terminals may be owned by a VTAM in a different VSE 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 14 on page 40)
- CICS1 may acquire terminals
- After takeover, CICS2 may acquire terminals
- New terminals may log on to CICS2.

In the example in Figure 15 on page 41, there are the following considerations:

- The ownership of the network by the VTAM in VSE1
- The cross-domain definitions of the network to VSE2
- The local definition of application CICS1 in VSE1
- The cross-domain definition of application CICS1 in VSE2
- The local definition of application CICS2 in VSE2
- The cross-domain definition of application CICS2 in VSE1.

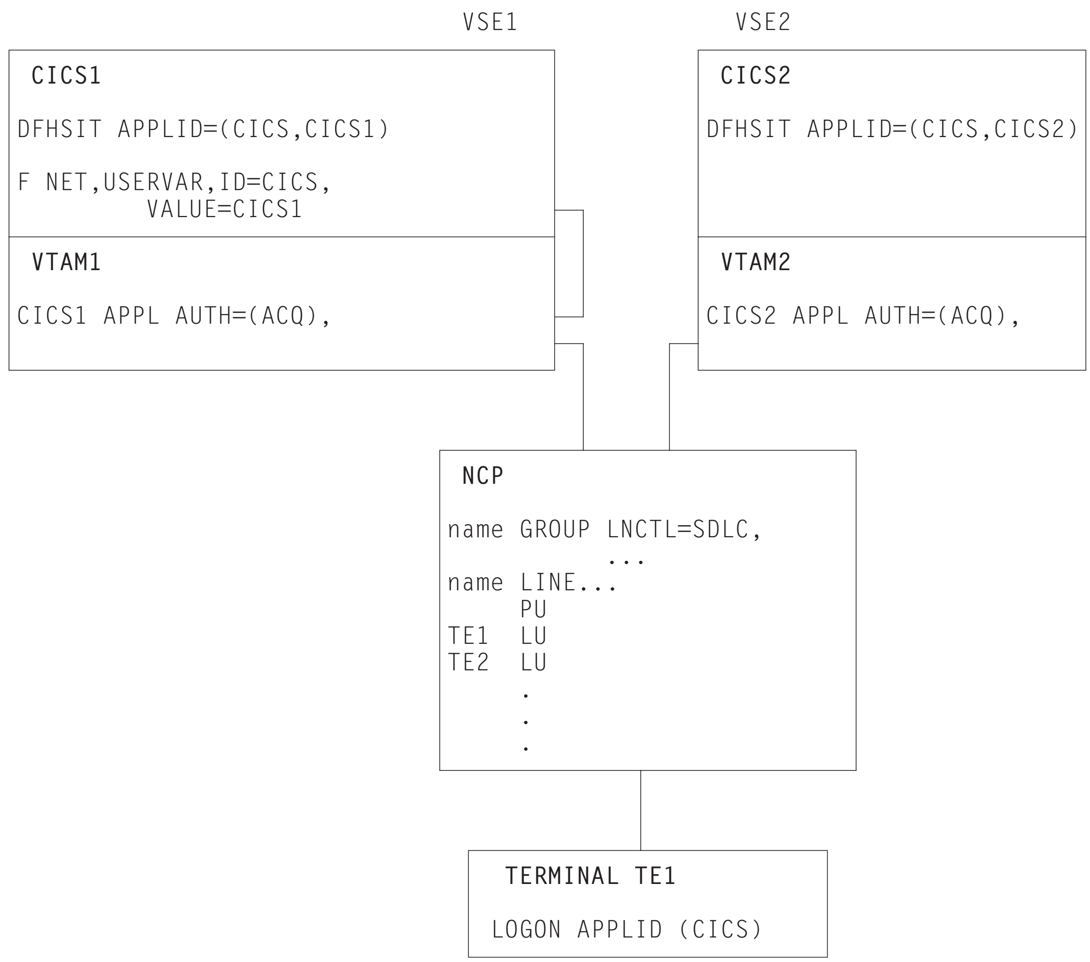

Figure 14. Logging on to the active

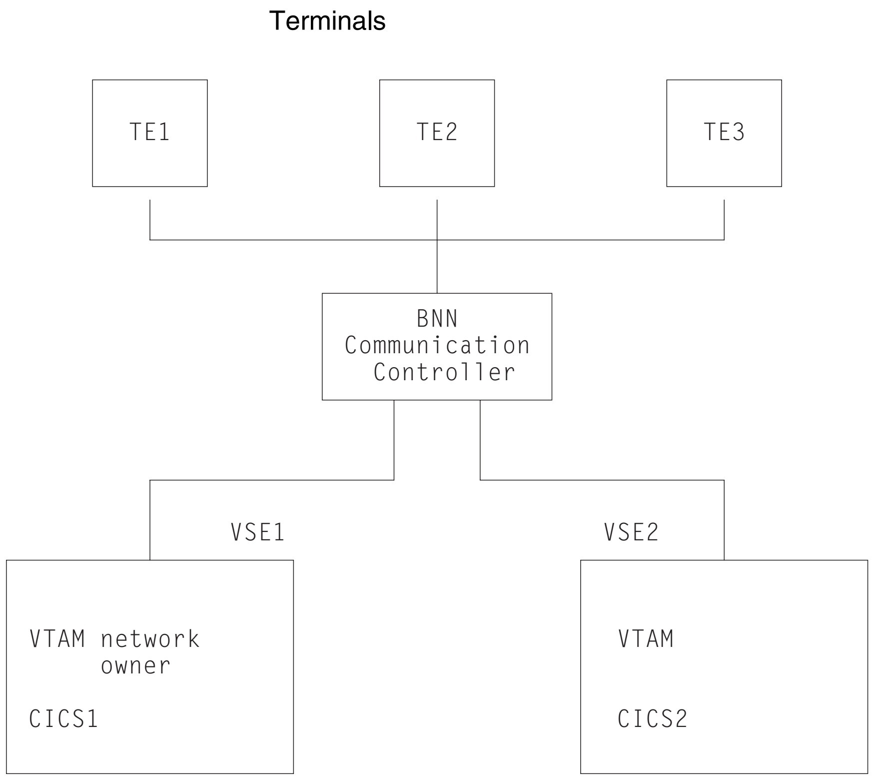

Figure 15. VTAM network ownership

The following partial NCP definition defines VSE1 as the network owner, and the terminals in that network:

```
            BUILD......,BACKUP=350
            GROUP....,LNCTL=SDLC,....,OWNER=VSE1
            LINE...
            PU...
TE1         LU...
TE2         LU...
TE3         LU...
```

The following partial definition defines CICS1 on VSE1, with a cross-domain definition for CICS2:

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

CICS2       CDRSC CDRM=VSE2
```

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

Here is a cross-domain definition in VSE2 for the terminals:

```
            VBUILD TYPE=CDRSC
TE1         CDRSC CDRM=VSE1
TE2         CDRSC CDRM=VSE1
TE3         CDRSC CDRM=VSE1
```

The following partial definition defines CICS2 to run on VSE2, with a cross-domain definition for CICS1:

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

CICS1       CDRSC CDRM=VSE1
```

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" on page 39.

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 Release Guide*.

Table 3 describes the two classes of terminals in an XRF environment, how XRF supports them, and what the user can expect at a takeover.

Table 3. Terminal support

Terminal class	**How XRF supports terminals at logon**	**How XRF supports terminals at takeover**	**How takeover affects terminal user**
Tracked (class 2)	No change to normal CICS support.	Alternate tries to reestablish session.	Brief delay in service while alternate acquires session.
Untracked (class 3)	No change to normal CICS support.	No change to normal CICS emergency restart procedures.	User loses service. Operator must reestablish session.

Note: VTAM under VSE/ESA does not support Class 1 terminals.

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 terminal definition keyword 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" on page 44.

Tracked (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-VSE 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.

Class 2 terminals 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 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 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. 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” on page 45.

You must consider how your operations staff will transfer class 2 terminals from one VSE to another in a multi-VSE environment. In a single-VSE 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 page 64.

Untracked (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.

In a multi-VSE 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.

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

Defining the recovery process

You can use RDO to define the recovery process for each terminal by the RECOVOPTION keyword on the RDO TYPETERM resource definition. For reference information about this keyword, see the *CICS Resource Definition Guide* 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" on page 45.

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.

Signon after takeover

Users of tracked 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 specified in an RDO TYPETERM resource definition, or in the external security manager (ESM) CICS segment for the signed on user.

If you specify NOFORCE, a specification of FORCE on an RDO TYPETERM definition or in the ESM CICS segment can be used to make smaller groups of terminals sign on again. The system initialization parameters are described further in "System initialization parameters" on page 49.

RDO 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 ESM 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 ESM entry to force a terminal or group of terminals to be signed off.

External security manager CICS segment, XRFSOFF=FORCE|NOFORCE
The lowest level at which you can force a terminal to be signed off is in the associated users ESM CICS segment. One ESM CICS segment could apply to a number of terminals. For more information, see the *CICS 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 16.

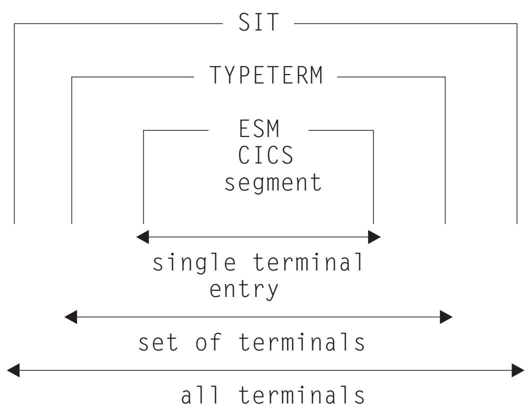

Figure 16. Signoff levels

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 ESM CICS segment 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.

You must consider the effect of the system initialization AUTCONN parameter. AUTCONN delays the reconnection of terminals (see "Starting the alternate" on page 51), 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 ESM resource profiles have been rebuilt. ESM 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.

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.

LUTYPE6 ISC application-to-application sessions

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.

CICS-to-CICS communication

An active can communicate, using ISC, with:

- A CICS/VSE Version 2 system
- A CICS/ESA Version 3 system
- A CICS/ESA Version 4 system
- A CICS Transaction Server for OS/390 Version 1 system
- A CICS OS/2™ Version 2 system
- A CICS for OS/2 Version 2.0.1 system
- A CICS/VM™ system
- A CICS 400® system
- A CICS/6000® Version 1.0 system
- A CICS for Windows NT system
- CICS on Open Systems:
  - CICS/6000 Version 1.2
  - CICS for DEC OSF/1AXP
  - CICS for HP 9000.

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.

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.

XRF SNA flows

Figure 17 shows a representative sequence of SNA flows for:

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

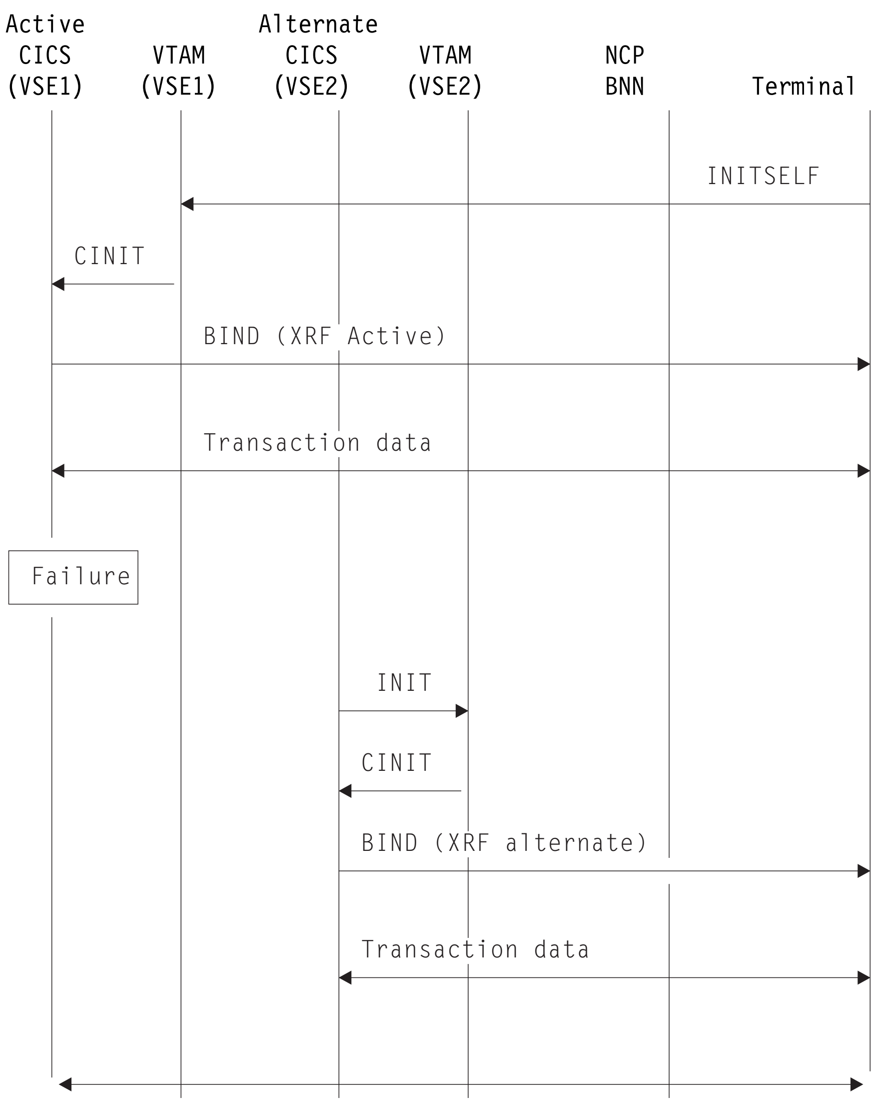

Figure 17. Abbreviated XRF SNA flows

Chapter 6. Defining CICS for XRF

This chapter 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"
- "Command list table (CLT)" on page 54
- "User exit for VTAM failure" on page 62
- "The overseer" on page 62
- "Supplied transactions for controlling the alternate" on page 63
- "Sharing data sets" on page 65
- "Storage protection considerations" on page 65.

For reference information for tables, see the *CICS Resource Definition Guide* manual. For system initialization, see the *CICS System Definition Guide*. Two specific sample implementations are given in Appendix B, "Sample XRF implementations" on page 75.

Advice about terminal operands that can influence the takeover characteristics for individual terminals is given in Chapter 5, "The terminal network" on page 37.

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 system initialization 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.

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
XSWITCH=(0-254,progname,{A|B})
```

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" on page 37 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 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 RDO TYPETERM definition or the ESM CICS segment.

For more information, see "Signon after takeover" on page 45.

XSWITCH=(0-254,progname,{A|B})

XSWITCH defines a programmable terminal switching unit, that can be used with midrange 2-CPC XRF systems, instead of using a communication controller. The program specified on this parameter instructs the unit to switch terminal lines to the active's CPC at startup and to the alternate's CPC during takeover.

The number in the range 0-254 specifies the logical unit to which the switch is assigned.

progname identifies the user-written program that will issue commands to the switching unit.

A/B identifies the CPC to which the terminal lines are to be directed.

For more information about switching units, contact your IBM technical support representative.

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
XRFTODI=30|decimal-value
AUTCONN=0|hhmmss
XRFSTME=nn|5
XSWITCH=(0-254,progname,{A|B})
```

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)” on page 54.

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 “Supplied transactions for controlling the alternate” on page 63. 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” on page 63 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-VSE environment, it could also be useful if activity on the active CICS VSE (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 4 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.

Table 4. 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	Unconditional takeover	Conditional takeover	No takeover
Operator issues a CEMT transaction	Unconditional takeover	Unconditional takeover	No takeover

Note: If the active CICS VSE 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 on page 64, can prevent the alternate from reacting to the disappearance of the active's surveillance signal while, for example, the VSE image of the active CICS is stopped.

Unpredictable, temporary stoppages of the active CICS can occur (for example, when an unrelated address space in its VSE image issues an SDUMP). You should take this into account when choosing your *ADI* value.

You should also consider how to avoid some of the causes of unnecessary takeovers.

AUTCONN=0|hhmmss

Delays the reconnection, after a takeover, of tracked terminals in session at the time of failure. The default is zero.

You might set a delay to allow the operator to do some manual switching of lines.

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.

XRFSTME=nn|5

This operand has already been described on page 46. 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.

XSWITCH=(0-254,progname,{A|B})

This option, described more fully in “Starting the active” on page 50, defines a programmable terminal switching unit. The unit may be operated, using the program defined in this option, to switch terminal lines to the alternate's CPC during takeover.

XRFTODI=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 POWER is unable to inform the alternate that the active has stopped. The *XRFTODI* value might have to be a compromise, as follows:

- A low *XRFTODI* 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.

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

A CICS failure, on the other hand, can usually be handled automatically if the POWER systems can access the shared spool. A low *XRFTODI* value would result in requests for operator action even though VSE 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 (with message DFHXA6561 or DFHXA6562) the alternate continues to ask POWER for the status of the active job. If it discovers from POWER 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 VSE or CPC has failed. If the operator replies “CPC” to the first alternate system that takes over, any other alternates taking over from actives that have failed on that VSE image do not have to ask for operator intervention, and their takeovers proceed without interruption.

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.

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 VSE CANCEL command to remove the active CICS job. If the active and alternate are running in different VSE images, the CAVM uses POWER facilities to send the CANCEL command to the destination VSE.

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

The CLT—background information

The CLT applies only to XRF. It is used only by the alternate; every alternate must have a CLT. The authenticity of the information in the CLT must be guaranteed because the integrity and security of the entire VSE system might be compromised if an alternate could be made to use data supplied by an unauthorized person.

This information is therefore placed in the CLT. 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, CICS loads the CLT, 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 cancelled.

- Routing information needed to send CANCEL commands to the appropriate target VSE system (in a multi-VSE environment). You do not need this information in a single-VSE environment.
- VSE 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 for VSE/ESA. 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” on page 75. For reference information about the CLT, see the *CICS Resource Definition Guide*.

The CLT in a single CICS configuration

Figure 18 on page 56 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.

```
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',
               MF=L

         DFHCLT TYPE=LISTEND

         DFHCLT TYPE=FINAL

         END
```

Putting together the macros described in the *CICS Resource Definition Guide* manual, the sample CLT following the figure defines the CICS2 system illustrated in Figure 18 on page 56.

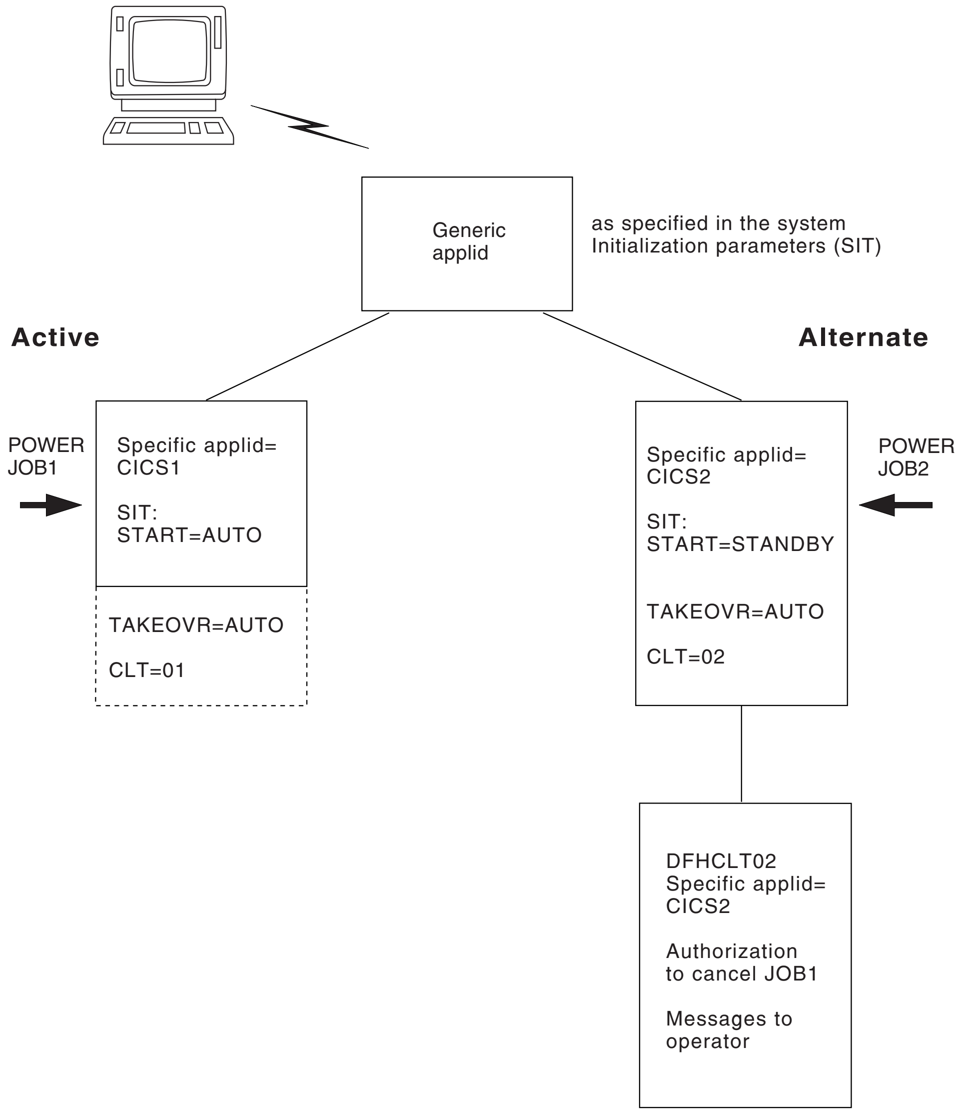

Figure 18. System initialization parameters and CLT working together

The CLT in a multi-VSE, 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-VSE, MRO configuration, when there is a takeover of one region to the second VSE, all the alternates must take over from their active counterparts to retain communication between the regions. This is because MRO does not operate across VSE 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 system initialization parameter 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 takeover.
- The takeover of a failing master region forces the takeover of all communicating regions to the alternates in the second VSE 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-VSE 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 19 on page 58, there are two active regions, connected by MRO, in a multi-VSE 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 cancellation of the active master job, and the alternate dependent region's CLT authorizes the cancellation of the active dependent job.

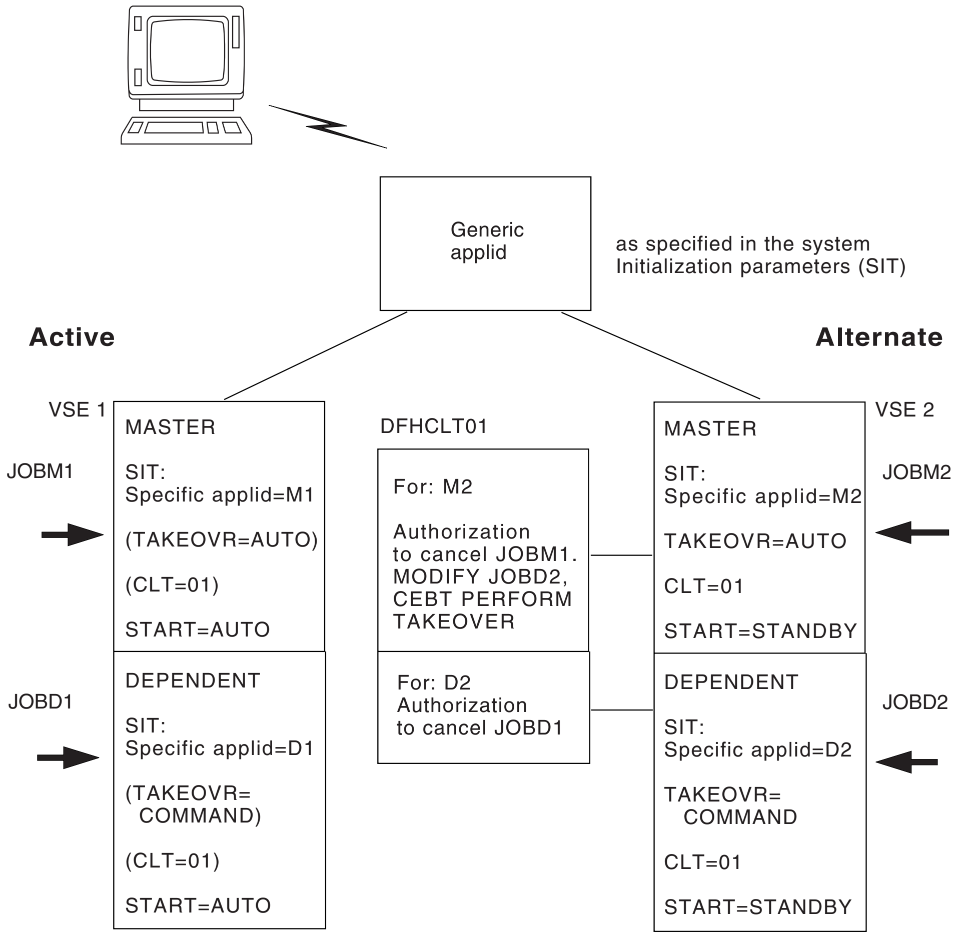

Figure 19. System initialization parameters and CLT in an MRO configuration

Figure 19 illustrates the relationship between the relevant system initialization parameters and the CLT.

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” on page 63.

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 VSE1. This would avoid the disruption in the service provided by the master region that would occur on a takeover to VSE2. 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” on page 30.

This is relevant to individual CICS failures. If the CEC or VSE failed, all regions would have to be taken over to the other VSE. 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.

```
DFHCLT01 DFHCLT01 DFHCLT TYPE=INITIAL, SUFFIX=01

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',
               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 VSE 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.

Use of the coordinator

In a large multi-VSE, MRO configuration, you might have more than one master region and any number of dependent regions. Figure 20 on page 61 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).

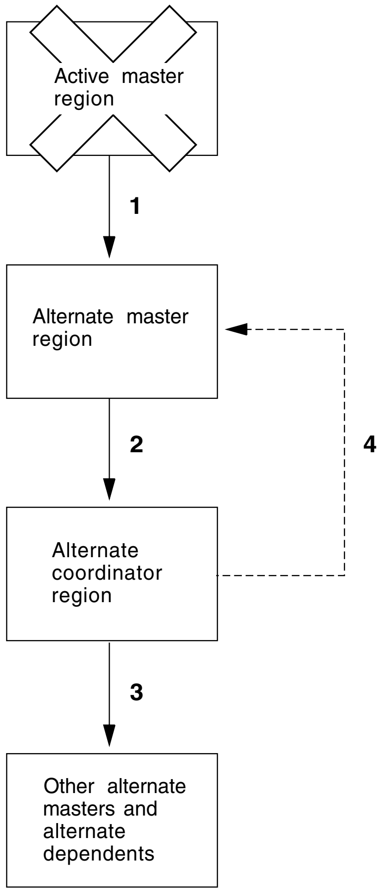

Figure 20. Flow of control and the coordinator region

See the following notes that apply to Figure 20.

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" on page 75.

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 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 terminates by itself if its VTAM fails. In a multi-VSE 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.

The overseer

The overseer was introduced on page 31. 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 source is named DFH$AXRO, and is supplied in the VSE/ESA sublibrary PRD1.BASE. There is also a pregenerated version ready to use. See the *CICS 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.

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 2.3 systems running in XRF mode in the same VSE images as a CICS Transaction Server for VSE/ESA 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.

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-Supplied Transactions* manual.

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.

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 CICS 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 VSE 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 VSE image, when it has been recovered after a failure. If you want to move a set of MRO regions from one VSE 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 the TAKEOVR=AUTO system initialization parameter 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" on page 49.) 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 CICS.
- 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 VSE image, and not cause a takeover. When the VSE 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.

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, all sessions are terminated immediately.

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

Data sets that are shared, passively or actively, such as user VSAM data or DL/I data sets, must be placed on shared volumes or VSAM spaces. For more information about data sets, see the *CICS System Definition Guide*.

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 VSE, 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 values for DSALIM and EDSALIM.

Chapter 7. XRF and other products

This chapter describes briefly some of the other products that work with CICS in an XRF environment.

- “DB2 for VSE/ESA”
- “DL/I VSE”
- “NetView”
- “VM” on page 69

DB2 for VSE/ESA

CICS with XRF supports the use of DB2® for VSE/ESA databases.

This support is not described in this book. For guidance information about DB2 for VSE/ESA, see the DB2 for VSE/ESA library.

Note that after a takeover you can automatically initiate the CIRB transaction (required for the DB2 for VSE/ESA online resource manager), by using CICS sequential device support. Sequential device support is described in the *CICS Resource Definition Guide*, and the *CICS Application Programming Guide*.

DL/I VSE

CICS with XRF supports DL/I VSE

This support is not described in this book. For guidance information about the DL/I DOS/VS interface, see the *DL/I VSE Release Guide*.

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” on page 37

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 21 on page 68 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 21 illustrates the sequence of events from the failure of the NCP, through VTAM, NCCF, the message table and a CLIST, to the sending of a message to the operator.

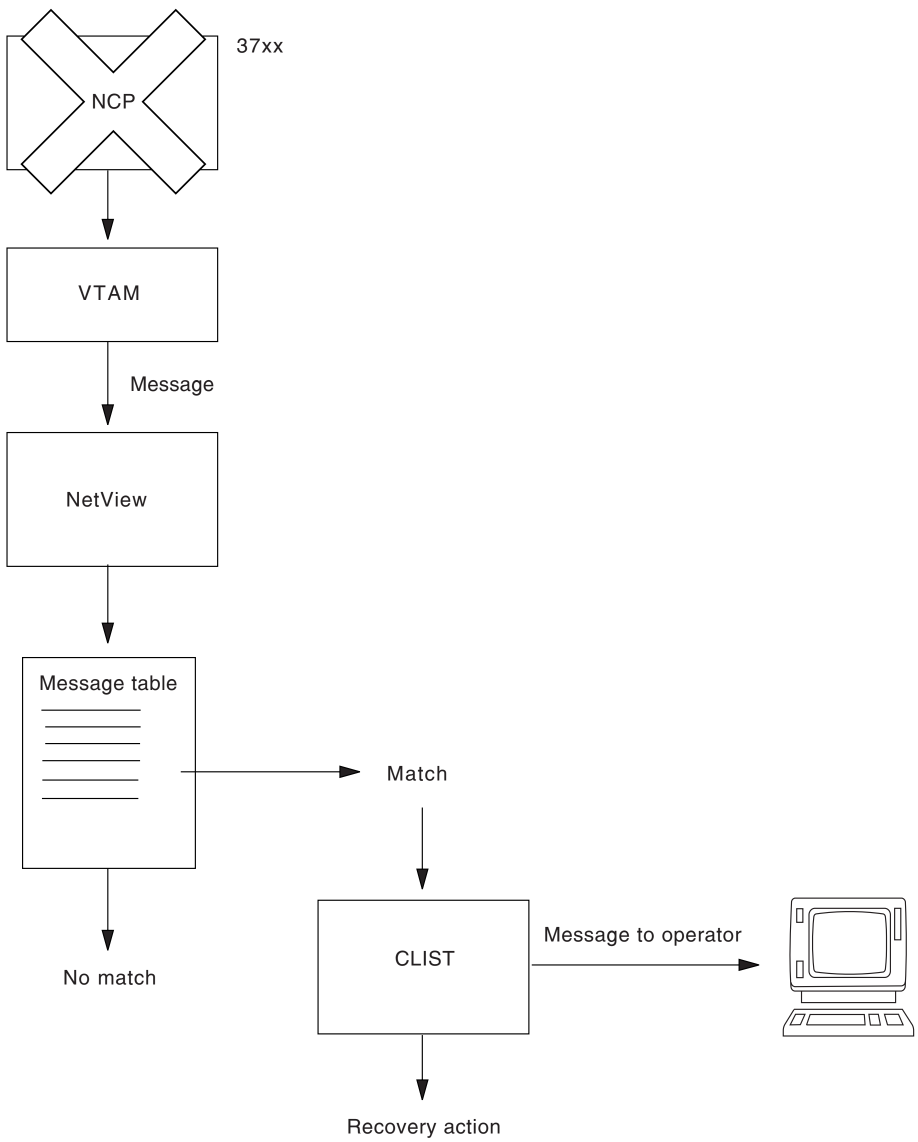

Figure 21. Automating 37xx recovery with NetView

VM

CICS with XRF will work under 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.

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/VSE Version 2. Release 3 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.

CPC-dead-data anchor block
: The module DFHCDDAN must be loaded into the SVA in a dual-CPC environment. For more information about loading modules into the SVA, see the *CICS System Definition Guide*.

DL/I VSE
: Ensure that table definitions, shared DASD, and system logging are suitable for DL/I VSE databases. For more information, see *The DL/I VSE Release Guide*.

Dump
: Determine if you need a dump of a failing active.

  Make sure that you initialize CICS, with an appropriate system initialization ADI value to avoid unnecessary takeovers. See page 53.

NCP
: Define NCP for XRF.

Node error program
: The *CICS Customization Guide* contains 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 Operations and Utilities Guide*.

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

POWER jobnames
: The POWER jobnames must be unique when running XRF.

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 Customization Guide*.

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

Shared DASD
Many data sets for XRF must be on shared DASD, in particular the CAVM data sets. The *CICS System Definition 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 “Signon after takeover” on page 45 for more information.

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 Operations and Utilities 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 system initialization overrides
- CLT
- RDO TYPETERM options.

There is some guidance about definitions in this book. For more details, see the *CICS Resource Definition Guide*

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

Time-of-day clock
The setting of the clocks in a multi-VSE 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.

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-VSE operations, you need cross domain definitions for CICS systems and logical units. These enable LUs owned by either VSE 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" on page 37.

Workload on second VSE image

Consider the effects of the workload on the second VSE after a takeover. For more information, see "Workload on a second VSE image" on page 23.

XXRSTAT exit

Create a user exit program for the XXRSTAT exit, if required. For more information, see "User exit for VTAM failure" on page 62. For the definitive product-sensitive programming interface information about global user exits, see *CICS Customization Guide*

Appendix B. Sample XRF implementations

In this appendix there are two sample implementations:

1. A single CICS region with an alternate in a second VSE image
2. An MRO configuration, with a dependent region, a master region, and a coordinator region, with actives and alternates in separate VSE images.

This appendix gives an overview of the SIT and SIT system initialization overrides, and CLT definitions. If you need more information about the SIT and CLT, see Chapter 6, "Defining CICS for XRF" on page 49. The *CICS System Definition Guide* contains a sample startup job stream.

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

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 system initialization table (SIT).

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

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)
    ,XRFTODI=35               /* (Alternate only)
    ,AUTCONN=0
    ,AIRDELAY=700             /* (Active 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 VSE1 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 SYSIPT data set, because START=STANDBY has been coded in the SIT table AA.

```
*  $$ JOB  JNM=JOB1,CLASS=2,DISP=L
 ...
// EXEC DFHSIP,SIZE=DFHSIP,PARM=' ....,SI',OS390
              ....
             ,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.

```
*  $$ JOB JNM=JOB2,CLASS=2,DISP=L
 ...
// EXEC DFHSIP,SIZE=DFHSIP,PARM=' ....,SI',OS390
               ....
              ,SIT=AA
              ,APPLID=(CICS,CICS2)
              ,.....
```

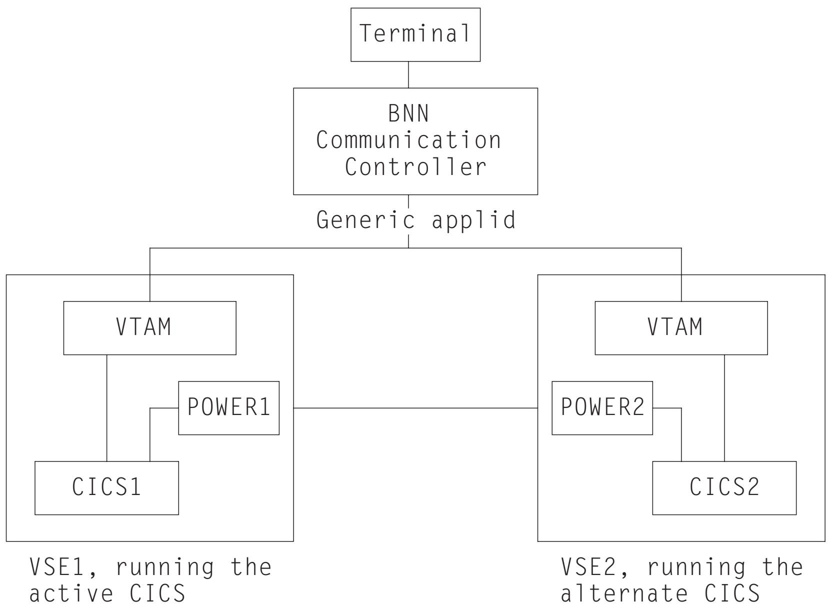

Figure 22. Sample single CICS implementation

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,
               SUFFIX=01               CLT suffix
*
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',
               MF=L

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

MRO CICS implementation

In this example, shown in Figure 23 on page 80, 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.

Note: POWER1 and POWER2 share the spool and DASD.

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, system initialization overrides are used to tailor the SIT.

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
    ,XRFTODI=25
    ,AUTCONN=0
    ,AIRDELAY=700
    ,XRFSOFF=NOFORCE
    ,XRFSTME=5
    ,.....
```

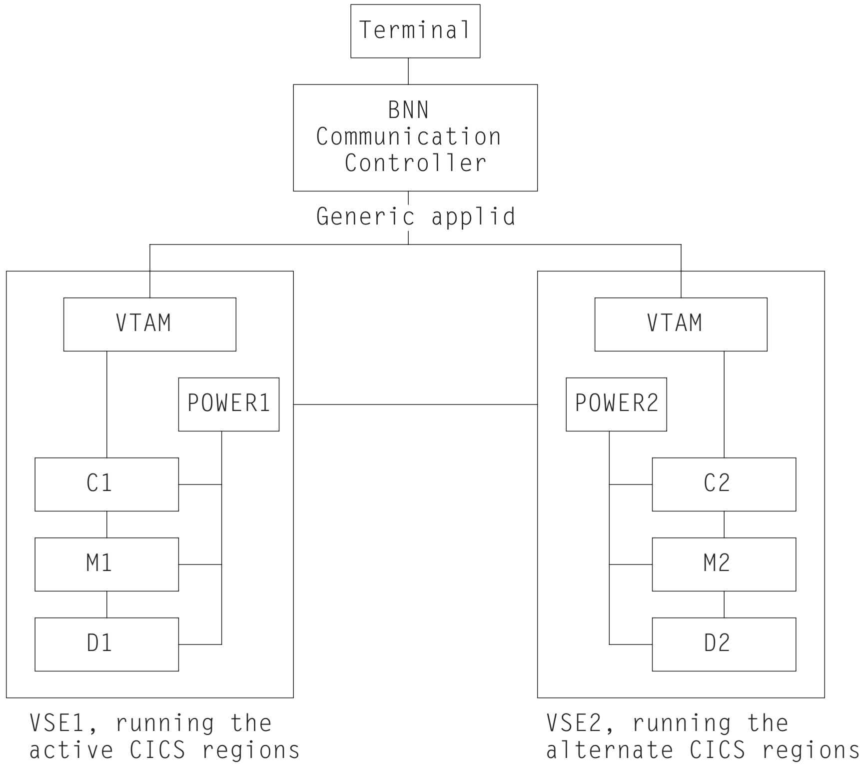

Figure 23. Sample MRO CICS implementation

CICS job JOBC1: The following SIT overrides are required to initialize the active coordinator region on VSE1.

```
*  $$ JOB  JNM=JOBC1,CLASS=2,DISP=L

// EXEC DFHSIP,SIZE=DFHSIP,PARM=' ....,SI',OS390
        ...
              ....
             ,SIT=CO
             ,START=AUTO
             ,APPLID=(C,C1)
             ,.....
```

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 because START=STANDBY is coded in each SIT.

CICS job JOBC2

```
*  $$ JOB  JNM=JOBC2,CLASS=2,DISP=L
 ...
// EXEC DFHSIP,SIZE=DFHSIP,PARM=' ....,SI',OS390
              ....
             ,SIT=CO
             ,APPLID=(C,C2)
             ,.....
```

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
    ,XRFTODI=25
    ,AUTCONN=0
    ,AIRDELAY=700
    ,XRFSOFF=NOFORCE
    ,XRFSTME=5
    ,.....
```

CICS job JOBM1

```
*  $$ JOB  JNM=JOBM1,CLASS=2,DISP=L
 ...
// EXEC DFHSIP,SIZE=DFHSIP,PARM=' ....,SI',OS390
              ....
             ,SIT=MA
             ,START=AUTO
             ,APPLID=(M,M1)
             ,.....
```

CICS job JOBM2

```
*  $$ JOB  JNM=JOBM2,CLASS=2,DISP=L
 ...
// EXEC DFHSIP,SIZE=DFHSIP,PARM=' ....,SI',OS390
               ....
              ,SIT=MA
              ,APPLID=(M,M2)
              ,.....
```

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
    ,XRFTODI=25
    ,AUTCONN=0
    ,AIRDELAY=700
    ,XRFSOFF=NOFORCE
    ,XRFSTME=5
    ,.....
```

CICS job JOBD1

```
*  $$ JOB  JNM=JOBD1,CLASS=2,DISP=L
 ...
// EXEC DFHSIP,SIZE=DFHSIP,PARM=' ....,SI',OS390
               ....
              ,SIT=DE
              ,START=AUTO
              ,APPLID=(D,D1)
              ,.....
```

CICS job JOBD2

```
*  $$ JOB  JNM=JOBD2,CLASS=2,DISP=L
 ...
// EXEC DFHSIP,SIZE=DFHSIP,PARM=' ....,SI',OS390
               ....
              ,SIT=DE
              ,APPLID=(D,D2)
              ,.....
```

CLT for MRO-connected regions

This sample CLT, shown in Figure 24 on page 84, 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,                                    *
            SUFFIX=02            CLT suffix (CLT02 for both VSEs
*
*-----------------------------------------------------------------
*
*The following CLT entries govern a takeover of the MRO group
*from C1, M1, D1 running on one VSE to C2, M2, D2 running on the
*other VSE
*
*-----------------------------------------------------------------
*
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 VSE
               COMMAND='MODIFY USERJOB,USER COMMAND'
*
         DFHCLT TYPE=WTO,              Put out a console message  *
               WTOL=MSG1
MSG1     WTO 'NOTE TAKEOVER TO NUMBER 2 REGIONS',                 *
               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
```

Figure 24 (Part 1 of 2). Sample CLT

```
*-----------------------------------------------------------
*
*The following CLT entries govern a takeover of the MRO group
*from C2, M2, D2 running on one VSE to C1, M1, D1 running on the
*other VSE
*
*------------------------------------------------------------
*
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 VSE
               COMMAND='MODIFY USERJOB,USER COMMAND'
*
         DFHCLT TYPE=WTO,            Put out a console message
               WTOL=MSG2

MSG2     WTO 'NOTE TAKEOVER TO NUMBER 1 REGIONS',
               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
```

Figure 24 (Part 2 of 2). Sample CLT

Bibliography

CICS Transaction Server for VSE/ESA Release 1 library

Evaluation and planning	
Release Guide	GC33-1645
Migration Guide	GC33-1646
Report Controller Planning Guide	SC33-1941
General	
Master Index	SC33-1648
Trace Entries	SX33-6108
User's Handbook	SX33-6101
Glossary (softcopy only)	GC33-1649
Administration	
System Definition Guide	SC33-1651
Customization Guide	SC33-1652
Resource Definition Guide	SC33-1653
Operations and Utilities Guide	SC33-1654
CICS-Supplied Transactions	SC33-1655
Programming	
Application Programming Guide	SC33-1657
Application Programming Reference	SC33-1658
Sample Applications Guide	SC33-1713
Application Migration Aid Guide	SC33-1943
System Programming Reference	SC33-1659
Distributed Transaction Programming Guide	SC33-1661
Front End Programming Interface User's Guide	SC33-1662
Diagnosis	
Problem Determination Guide	GC33-1663
Messages and Codes Vol 3 (softcopy only)	SC33-6799
Diagnosis Reference	LY33-6085
Data Areas	LY33-6086
Supplementary Data Areas	LY33-6087
Communication	
Intercommunication Guide	SC33-1665
CICS Family: Interproduct Communication	SC33-0824
CICS Family: Communicating from CICS on System/390	SC33-1697
Special topics	
Recovery and Restart Guide	SC33-1666
Performance Guide	SC33-1667
Shared Data Tables Guide	SC33-1668
Security Guide	SC33-1942
External Interfaces Guide	SC33-1669
XRF Guide	SC33-1671
Report Controller User's Guide	SC34-5688
CICS Clients	
CICS Clients: Administration	SC33-1792
CICS Universal Clients Version 3 for OS/2: Administration	SC34-5450
CICS Universal Clients Version 3 for Windows: Administration	SC34-5449
CICS Universal Clients Version 3 for AIX: Administration	SC34-5348
CICS Universal Clients Version 3 for Solaris: Administration	SC34-5451
CICS Family: OO programming in C++ for CICS Clients	SC33-1923
CICS Family: OO programming in BASIC for CICS Clients	SC33-1671
CICS Family: Client/Server Programming	SC33-1435
CICS Transaction Gateway Version 3: Administration	SC34-5448

Books from VSE/ESA 2.5 base program libraries

VSE/ESA Version 2 Release 5

Book title	Order number
Administration	SC33-6705
Diagnosis Tools	SC33-6614
Extended Addressability	SC33-6621
Guide for Solving Problems	SC33-6710
Guide to System Functions	SC33-6711
Installation	SC33-6704
Licensed Program Specification	GC33-6700
Messages and Codes Volume 1	SC33-6796
Messages and Codes Volume 2	SC33-6798
Messages and Codes Volume 3	SC33-6799
Networking Support	SC33-6708
Operation	SC33-6706
Planning	SC33-6703
Programming and Workstation Guide	SC33-6709
System Control Statements	SC33-6713
System Macro Reference	SC33-6716
System Macro User’s Guide	SC33-6715
System Upgrade and Service	SC33-6702
System Utilities	SC33-6717
TCP/IP User's Guide	SC33-6601
Turbo Dispatcher Guide and Reference	SC33-6797
Unattended Node Support	SC33-6712

High-Level Assembler Language (HLASM)

Book title	Order number
General Information	GC26-8261
Installation and Customization Guide	SC26-8263
Language Reference	SC26-8265
Programmer’s Guide	SC26-8264

Language Environment for VSE/ESA (LE/VSE)

Book title	Order number
C Run-Time Library Reference	SC33-6689
C Run-Time Programming Guide	SC33-6688
Concepts Guide	GC33-6680
Debug Tool for VSE/ESA Fact Sheet	GC26-8925
Debug Tool for VSE/ESA Installation and Customization Guide	SC26-8798
Debug Tool for VSE/ESA User's Guide and Reference	SC26-8797
Debugging Guide and Run-Time Messages	SC33-6681
Diagnosis Guide	SC26-8060
Fact Sheet	GC33-6679
Installation and Customization Guide	SC33-6682
LE/VSE Enhancements	SC33-6778
Licensed Program Specification	GC33-6683
Programming Guide	SC33-6684
Programming Reference	SC33-6685
Run-Time Migration Guide	SC33-6687
Writing Interlanguage Communication Applications	SC33-6686

VSE/ICCF

Book title	Order number
Adminstration and Operations	SC33-6738
User's Guide	SC33-6739

VSE/POWER

Book title	Order number
Administration and Operation	SC33-6733
Application Programming	SC33-6736
Networking Guide	SC33-6735
Remote Job Entry User's Guide	SC33-6734

VSE/VSAM

Book title	Order number
Commands	SC33-6731
User's Guide and Application Programming	SC33-6732

VTAM for VSE/ESA

Book title	Order number
Customization	LY43-0063
Diagnosis	LY43-0065
Data Areas	LY43-0104
Messages and Codes	SC31-6493
Migration Guide	GC31-8072
Network Implementation Guide	SC31-6494
Operation	SC31-6495
Overview	GC31-8114
Programming	SC31-6496
Programming for LU6.2	SC31-6497
Release Guide	GC31-8090
Resource Definition Reference	SC31-6498

Books from VSE/ESA 2.5 optional program libraries

C for VSE/ESA (C/VSE)

Book title	Order number
C Run-Time Library Reference	SC33-6689
C Run-Time Programming Guide	SC33-6688
Diagnosis Guide	GC09-2426
Installation and Customization Guide	GC09-2422
Language Reference	SC09-2425
Licensed Program Specification	GC09-2421
Migration Guide	SC09-2423
User's Guide	SC09-2424

COBOL for VSE/ESA (COBOL/VSE)

Book title	Order number
Debug Tool for VSE/ESA Fact Sheet	GC26-8925
Debug Tool for VSE/ESA Installation and Customization Guide	SC26-8798
Debug Tool for VSE/ESA User's Guide and Reference	SC26-8797
Diagnosis Guide	SC26-8528
General Information	GC26-8068
Installation and Customization Guide	SC26-8071
Language Reference	SC26-8073
Licensed Program Specifications	GC26-8069
Migration Guide	GC26-8070
Migrating VSE Applications To Advanced COBOL	GC26-8349
Programming Guide	SC26-8072

DB2 Server for VSE

Book title	Order number
Application Programming	SC09-2393
Database Administration	GC09-2389
Installation	GC09-2391
Interactive SQL Guide and Reference	SC09-2410
Operation	SC09-2401
Overview	GC08-2386
System Administration	GC09-2406

DL/I VSE

Book title	Order number
Application and Database Design	SH24-5022
Application Programming: CALL and RQDLI Interface	SH12-5411
Application Programming: High-Level Programming Interface	SH24-5009
Database Administration	SH24-5011
Diagnostic Guide	SH24-5002
General Information	GH20-1246
Guide for New Users	SH24-5001
Interactive Resource Definition and Utilities	SH24-5029
Library Guide and Master Index	GH24-5008
Licensed Program Specifications	GH24-5031
Low-level Code and Continuity Check Feature	SH20-9046
Library Guide and Master Index	GH24-5008
Messages and Codes	SH12-5414
Recovery and Restart Guide	SH24-5030
Reference Summary: CALL Program Interface	SX24-5103
Reference Summary: System Programming	SX24-5104
Reference Summary: HLPI Interface	SX24-5120
Release Guide	SC33-6211

PL/I for VSE/ESA (PL/I VSE)

Book title	Order number
Compile Time Messages and Codes	SC26-8059
Debug Tool For VSE/ESA User's Guide and Reference	SC26-8797
Diagnosis Guide	SC26-8058
Installation and Customization Guide	SC26-8057
Language Reference	SC26-8054
Licensed Program Specifications	GC26-8055
Migration Guide	SC26-8056
Programming Guide	SC26-8053
Reference Summary	SX26-3836

Screen Definition Facility II (SDF II)

Book title	Order number
VSE Administrator's Guide	SH12-6311
VSE General Introduction	SH12-6315
VSE Primer for CICS/BMS Programs	SH12-6313
VSE Run-Time Services	SH12-6312

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

The following paragraph does not apply in the United Kingdom or any other country where such provisions are inconsistent with local law:
INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore this statement may not apply to you.

This publication could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact IBM United Kingdom Laboratories, MP151, Hursley Park, Winchester, Hampshire, England, SO21 2JN. Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Programming License Agreement, or any equivalent agreement between us.

Trademarks and service marks

The following terms, used in this publication, are trademarks or service marks of IBM Corporation in the United States or other countries:

CICS,	IBM,
CICS/ESA,	NetView,
CICS/MVS,	Processor Resource/Systems Manager,
CICS/VSE,	VSE/ESA,
DB2 for VSE/ESA,	VTAM,
DL/I VSE,	3090

Other company, product and service names may be the trademarks or service marks of others.

Index

Numerics

A

B

C

D

E

F

G

H

I

L

M

N

O

P

Q

R

S

W

X

Sending your comments to IBM

CICS® Transaction Server for VSE/ESA™

XRF Guide

SC33-1671-01

If you want to send to IBM any comments you have about this book, please use one of the methods listed below. Feel free to comment on anything you regard as a specific error or omission in the subject matter, and on the clarity, organization or completeness of the book itself.

To request additional publications, or to ask questions or make comments about the functions of IBM products or systems, you should talk to your IBM representative or to your IBM authorized remarketer.

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

You can send your comments to IBM in any of the following ways:

- By mail:

  IBM UK Laboratories
  Information Development
  Mail Point 095
  Hursley Park
  Winchester, SO21 2JN
  England

- By fax:
  - From outside the U.K., after your international access code use 44 1962 870229
  - From within the U.K., use 01962 870229
- Electronically, use the appropriate network ID:
  - IBM Mail Exchange: GBIBM2Q9 at IBMMAIL
  - IBMLink: HURSLEY(IDRCF)
  - Email: idrcf@hursley.ibm.com

Whichever method you use, ensure that you include:

- The publication number and title
- The page number or topic to which your comment applies
- Your name and address/telephone number/fax number/network ID.

Program Number: 5648-054

Printed in the United States of America
on recycled paper containing 10%
recovered post-consumer fiber.

SC33-1671-01

Spine information:

IBM CICS TS for VSE/ESA XRF Guide *Release 1*