

Migration Planning for the 9037 Model 2 Sysplex Timer

Planning

Sysplex considerations

External Time Source
considerations



Noshir Dhondy
Frank Kyne



International Technical Support Organization

Migration Planning for the 9037 Model 2 Sysplex Timer

April 2003

Note: Before using this information and the product it supports, read the information in “Notices” on page v.

First Edition (April 2003)

This edition applies to 9037-001 and 9037-002 Sysplex Timers working with OS/390 Version 1 Release 3 and later.

© Copyright International Business Machines Corporation 2003. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Notices	v
Trademarks	vi
Preface	vii
The team that wrote this Redpaper	vii
Become a published author	vii
Comments welcome	viii
Chapter 1. Introduction	1
1.1 About this Redpaper	2
1.2 Terminology used	2
1.3 Which systems are potentially affected	2
1.4 Who should use this Redpaper	3
Chapter 2. Migration considerations	5
2.1 Overview	6
2.2 9037 Sysplex Timer configurations	6
2.3 Customer requirements during migration	9
2.4 Migration limitations	9
2.4.1 Upgrade of 9037-001 to 9037-002 not supported	9
2.4.2 Data integrity considerations	9
2.5 Recommended migration procedures	11
Chapter 3. Migration with limited outages	13
3.1 Overview	14
3.2 Service recommendations prior to performing migration	14
3.3 Planning tasks prior to migration	15
3.4 Risks during migration	17
3.5 Estimated time of limited operations	18
Chapter 4. Migration with planned outages	19
4.1 Overview	20
4.2 Service recommendations prior to performing migration	20
4.3 Planning tasks prior to migration	20
4.4 Risks during migration	20
4.5 Estimated time of outage	21
Appendix A. Connecting 9037-002s at extended distances	23
Expanded Availability configuration options	24
Related publications	27
IBM Redbooks	27
Other publications	27
How to get IBM Redbooks	27
Index	29

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.

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

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

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

@server™
@server™
Redbooks (logo) ™
ibm.com®
z/OS™
zSeries®

ESCON®
GDPS®
IBM®
MVS™
MVS/ESA™
OS/390®

Parallel Sysplex®
Redbooks™
S/390®
Sysplex Timer®

The following terms are trademarks of other companies:

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

C-bus is a trademark of Corollary, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

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

Preface

The 9037 Model 1 (9037-001) Sysplex Timer® has an announced worldwide end-of-service (EOS) date of December 31, 2003. This critical component of any multi-CPC sysplex will no longer be supported by IBM® after December 31, 2003.

IBM recommends replacing any 9037-001 models with Sysplex Timer Model 2s (9037-002). The 9037-002 is a full replacement for the 9037-001. There is no upgrade path from the 9037-001 to the 9037-002.

This Redpaper provides the information required for customers to successfully plan for the migration of their 9037-001 Sysplex Timers to 9037-002 Sysplex Timers.

The team that wrote this Redpaper

This Redpaper was produced by the following specialists working at the International Technical Support Organization, Poughkeepsie Center.

Noshir Dhondy is a Senior Engineer in the zSeries® Hardware Products organization in Poughkeepsie, New York. He is a member of the Parallel Sysplex® Product development team and the GDPS® technical marketing support team. He was the lead designer of the 9037 Sysplex Timer and continues to work on External Time Reference-related development for the zSeries Parallel Sysplex. He was one of the authors of the IBM Redbook *S/390® Time Management and IBM 9037 Sysplex Timer*. He joined IBM as an Associate Engineer in 1968 and has been associated with Parallel Sysplex since 1995. He holds degrees in Electrical Engineering from the Indian Institute of Technology, Bombay and University of Pittsburgh, USA.

Frank Kyne is a Senior IT Specialist at the International Technical Support Organization, Poughkeepsie Center. He writes extensively and teaches IBM classes worldwide on all areas of Parallel Sysplex. Before joining the ITSO five years ago, Frank worked in IBM Global Services in Ireland as an MVS™ Systems Programmer.

Thanks also to **Bob Haimowitz**, International Technical Support Organization, Poughkeepsie Center, for his contributions to this project.

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our papers to be as helpful as possible. Send us your comments about this Redpaper or other Redbooks™ in one of the following ways:

- ▶ Use the online **Contact us** review redbook form found at:

ibm.com/redbooks

- ▶ Send your comments in an Internet note to:

redbook@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYJ Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400



Introduction

This chapter describes the objective and target audience of this book. It also describes the types of systems that are likely to be impacted by the migration from 9037-001s to 9037-002s.

1.1 About this Redpaper

This Redpaper provides planning information for S/390 and zSeries sysplex customers that want to minimize the application outage that will result when their existing 9037-001 units are replaced by 9037-002 units.

This paper does not provide the detailed instructions required to perform the migration from a 9037-001 to a 9037-002. The actual migration should not be attempted without reference to the detailed instructions and a full understanding of the Sysplex Timer environment. The planning for the migration will involve your hardware engineer, systems programmers, and operations personnel. The actual migration tasks would normally be carried out by your service representative.

Your IBM representative can obtain the detailed migration procedure from the following Web site:

http://nascpok.pok.ibm.com/9037/mes/9037MIGR_Procedure.PDF

Important: If migration procedures other than those described in this document are used, there is the potential for a data integrity exposure.

Also, if the detailed instructions provided by the migration procedure are not followed, there is the potential for a total outage of all the systems in the sysplex.

1.2 Terminology used

In this paper we use the term *Sysplex Timer* to refer to either a 9037-001 or a 9037-002. If we are referring to a specific model, we will use the full model name.

Another term we use is *sysplex*. A *base sysplex* is a multi-system configuration consisting of a Sysplex Timer, CTCs for XCF communication, and a set of XCF Couple Data Sets. A *Parallel Sysplex* is similar to a base sysplex, but with the addition of a Coupling Facility.

However, from the perspective of the Sysplex Timer migration, it doesn't make any difference whether you have a base sysplex or a Parallel Sysplex. Therefore, we simply use the term sysplex throughout this book to refer to either configuration.

1.3 Which systems are potentially affected

The supported migration procedures assure that systems that are allowed to execute during the migration do not have a potential data integrity exposure and that they will be continuously available.

The following are systems that *could* potentially have a data integrity exposure if migration procedures are not adhered to:

- ▶ Systems that are members of a multi-CPC sysplex.

2.4.2, "Data integrity considerations" on page 9 describes the data integrity exposure if all the members of a multi-CPC sysplex are not using synchronized Sysplex Timers.

Therefore, at least some of the members of a multi-CPC sysplex will need to be quiesced for a time during the migration.

- ▶ Any system that requires time coordination across multiple CPCs.

An example of this is where multiple systems on different CPCs are writing to DASD that are being remote copied using XRC. XRC is used here as one example of systems that need time coordination. You need to carefully examine your environment to determine if other examples exist.

When using XRC, all updates to XRCed volumes are time stamped by the system issuing the write. The timestamps are subsequently used by the System Data Mover system to apply the changes, in time order, to the remote DASD. If the timestamps are not consistent across all the systems writing to the DASD, updates could be applied to the remote DASD in the wrong sequence.

It is important to remember that Linux on S/390 or zSeries also timestamps data updates so that Linux data can be mirrored using XRC, so this consideration does not only apply to z/OS™ or OS/390®-based systems.

In the case of XRC, the only way to guarantee the integrity of the remote XRC volumes during the migration is to quiesce the systems that write to the XRCed volumes on all CPCs except for one CPC. As you will see in Chapter 3, “Migration with limited outages” on page 13, one of the migration options allows all the systems on one CPC to continue operations—it is acceptable to leave any systems that use XRC up and running on that CPC.

There is a potential risk that all the timer signals to a CPC could be lost during the migration, if the detailed instructions in the migration procedures are not followed. The fact that the timer-related hardware is being worked on makes this a possibility. If all timer signals are lost, any system that has specified ETRMODE=YES in the CLOCKxx member of SYS1.PARMLIB will either go into a non-restartable wait state, or halt processing and issue a WTOR (depending on whether APAR OW44231 is applied).

The following are systems that have ETRMODE=YES specified and can therefore potentially experience an outage, if the provided migration procedures are not adhered to:

- ▶ Systems that are members of a multi-CPC sysplex and systems that require time coordination across multiple CPCs.
- ▶ Systems in a sysplex that are completely self-contained on one CPC (Monoplex).
- ▶ Systems that are not in a sysplex but require the same time.

1.4 Who should use this Redpaper

This Redpaper is for system planners, systems programmers, operations personnel, and engineers considering the replacement of 9037-001 units with 9037-002 units.



Migration considerations

This chapter provides a brief background of the 9037-001 and 9037-002 Sysplex Timers and discusses why it may be advantageous to upgrade to a 9037-002. It provides a short overview of the supported 9037-002 configurations, and goes on to discuss why migrating to 9037-002 timers may not be as simple as you might expect. It ends by summarizing the two options available to you.

2.1 Overview

The IBM 9037-001 Sysplex Timer was announced in 1990, to provide a common time reference across all the members of a sysplex. The IBM 9037-002 Sysplex Timer was announced and made available in 1996 as an enhancement to, and functional replacement for, the 9037-001. The 9037-001 was withdrawn from marketing in 1997, and support for service will be withdrawn on December 31, 2003.

The 9037-001 will have its 13th birthday in September 2003. It is reaching the end of its “burn-in hours” life. The components are beginning to age and will reach their maximum power-on hours capabilities and begin to fail. The components have gone end-of-life and are no longer available. The component stockpiles have been depleted and no new spare parts are being built. As a result, the stock of maintenance parts is being depleted. Based on these facts, the end-of-service date of December 31, 2003 was selected.

The Sysplex Timer is a key component when systems on multiple CPCs share access to the same data. If the Sysplex Timer fails, the sysplex will fail. For this reason, IBM is recommending that all sysplex customers currently using 9037-001s should migrate to 9037-002s before the end of 2003.

In addition to being supported, the 9037-002 delivers a number of enhancements over the 9037-001, including:

1. Increased separation between timers, from the three meters supported on the model 1, up to:
 - 3.0 Km without repeaters
 - 26.0 Km with repeaters
 - 40 Km with Dense Wavelength Division Multiplexers (DWDMs)
2. Increased number of ports from 16 up to 24
 - Port cards are hot-pluggable
3. Improved console support
 - Attachment via a Token Ring
 - Active/Backup console capability
 - Sysplex Timer console application can run on its own PC, on a Hardware Management Console (HMC), or on an ESCON® Director Console
4. Improved External Time Source (ETS) support
 - Improved tracking to ETS
 - Ability to add a backup ETS
 - Available in a configuration with only one 9037-002 (Expanded Basic configuration)

Refer to *Planning for the 9037 Model 2 Sysplex Timer*, SA22-7233 for details.

2.2 9037 Sysplex Timer configurations

The 9037-001 supports two configurations:

- ▶ The Basic configuration, consisting of a single 9037-001.
- ▶ The Expanded Availability configuration, consisting of two interconnected 9037-001s.

The Expanded Availability configuration provides improved availability via its ability to continue sysplex operation even if one of the two 9037s failed. However, the fact that the

two 9037s had to be within 3 meters of each other still effectively represented a single point of failure.

There are three possible 9037-002 configurations that will be referred to in this document. To facilitate understanding of these possibilities, we have included diagrams showing each.

Figure 2-1 shows the recommended configuration, known as the Expanded Availability configuration. This configuration consists of two 9037-002 timers, interconnected to each other. If either timer fails, or a link to any CPC fails, the CPC will automatically switch to the other timer and continue uninterrupted.

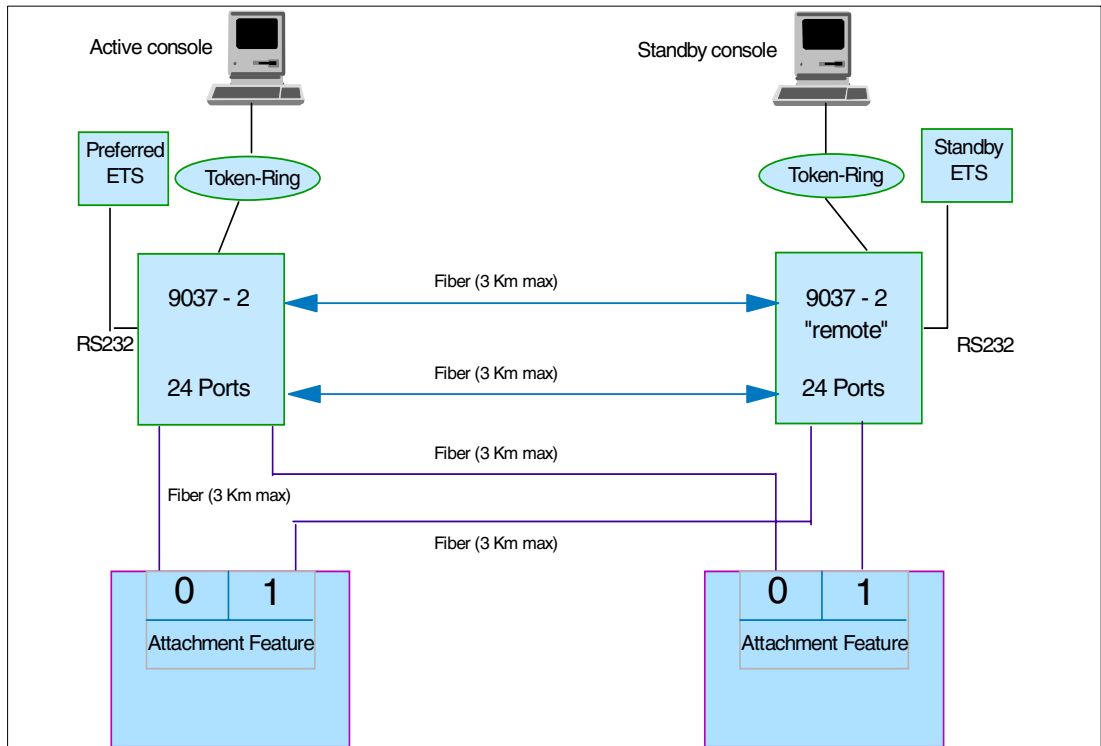


Figure 2-1 Expanded Availability configuration

Figure 2-2 on page 8 contains a diagram showing an Expanded Basic configuration. This configuration only has a single Sysplex Timer and therefore would not be a high availability configuration. The Expanded Basic configuration provides the ability to track to an External Time Source, and also includes an Expanded Availability feature, meaning that a second 9037-002 can subsequently be added nondisruptively.

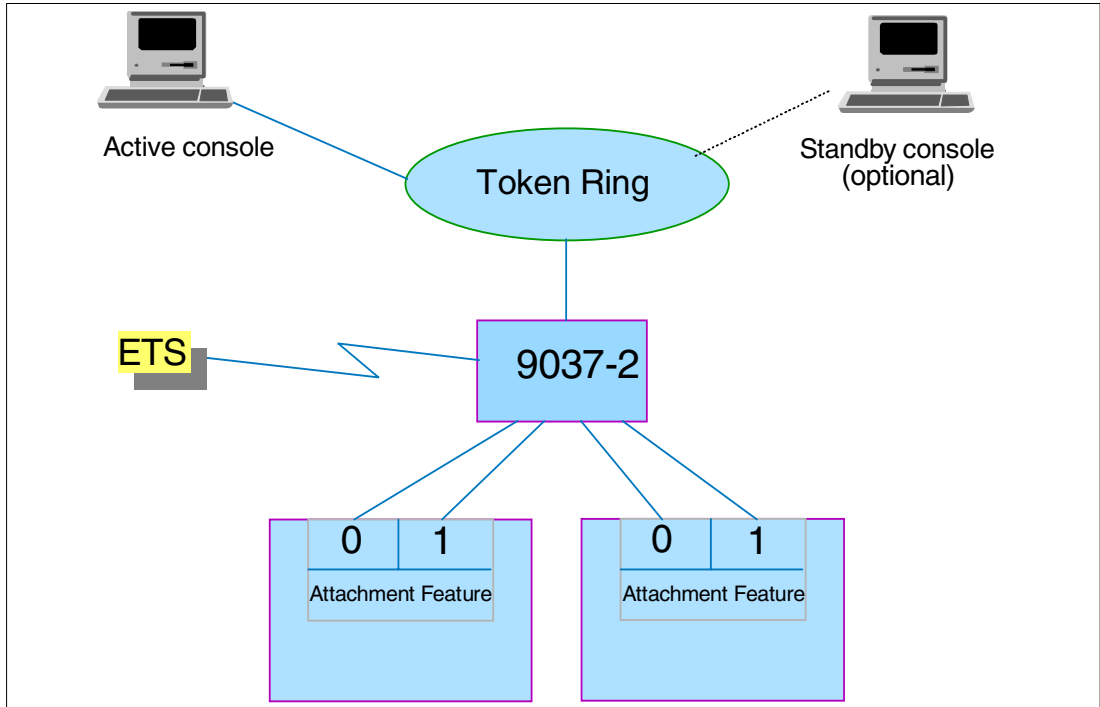


Figure 2-2 Expanded Basic configuration

The final, and least desirable option is the Basic configuration, as shown in Figure 2-3. Once again, this consists of a single Sysplex Timer and therefore would not be suitable for a high availability environment.

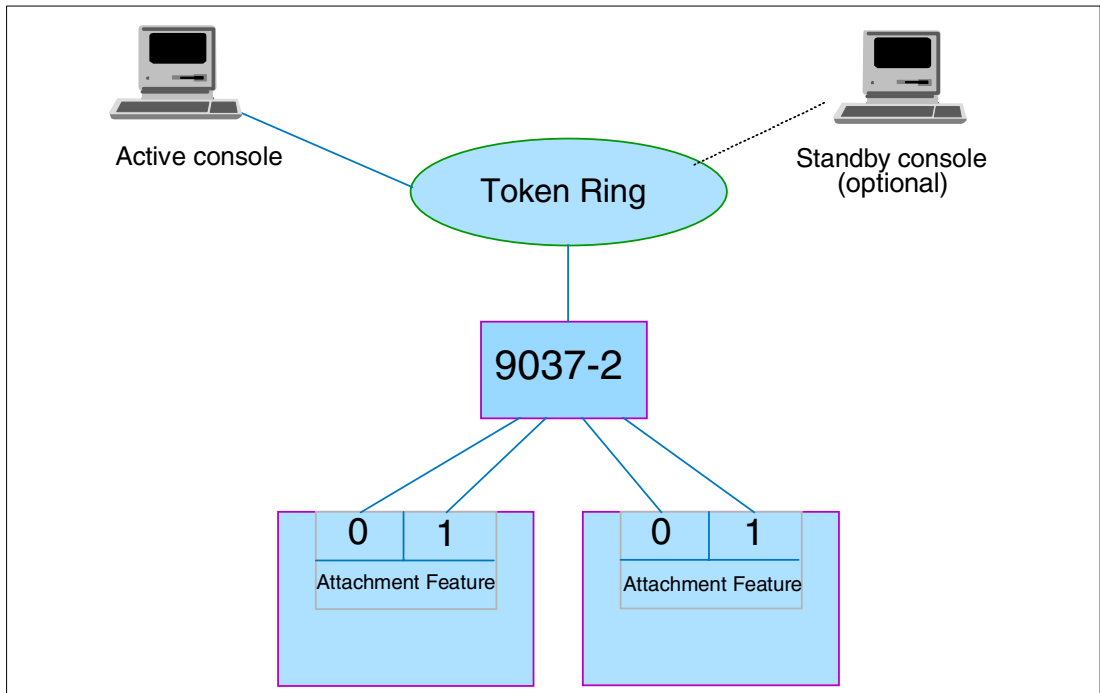


Figure 2-3 Basic configuration

2.3 Customer requirements during migration

The fundamental customer requirement is to maintain data integrity when migrating from a 9037-001 to a 9037-002. Also, since one of the key attributes of the sysplex is continuous availability, you will wish to limit the impact of the migration on your critical applications.

2.4 Migration limitations

Before proceeding, it is important to understand the limitations related to replacing an existing 9037-001 or pair of 9037-001s with a 9037-002 (or pair of 9037-002s).

2.4.1 Upgrade of 9037-001 to 9037-002 not supported

The 9037-002 is a completely redesigned unit, with a different chassis, power supply, card form factors, interfaces to the console, and so on. Therefore, you *cannot* upgrade from a 9037-001 to a 9037-002 by replacing cards or upgrading code. The migration consists of a complete unit replacement.

2.4.2 Data integrity considerations

As shown in Figure 2-4, a 9037-001 and a 9037-002 cannot be connected to each other. Therefore, there is no procedure that will safely allow all the systems that require time coordination across multiple CPCs to continue operations during the migration window. A potential data integrity risk exists if such a migration is attempted. However, a migration with a limited outage *is* possible as outlined in Chapter 3, “Migration with limited outages” on page 13.

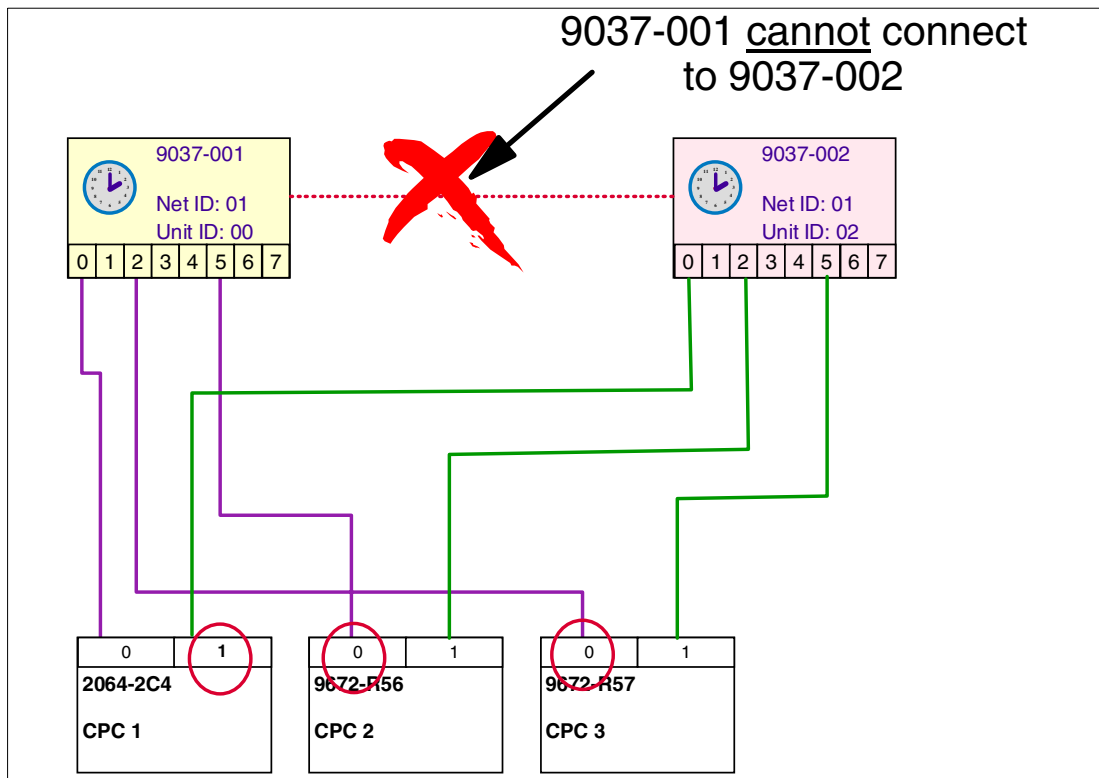


Figure 2-4 Non-configurable Expanded Availability configuration

To illustrate the potential data integrity exposure, consider the following configuration with three CPCs, each of which contains at least one system that requires time coordination. Figure 2-5 illustrates an intermediary stage during the migration procedure. At this stage of the migration, the connections between the CPCs and one of the 9037-001 units has been replaced with connections to one of the new 9037-002 units. The Time-of-Day clocks of all the CPCs will be stepping to signals received from the 9037-001 (Port 0 of each CPC), and the alternate ports of all CPCs will be connected to enabled ports from the 9037-002.

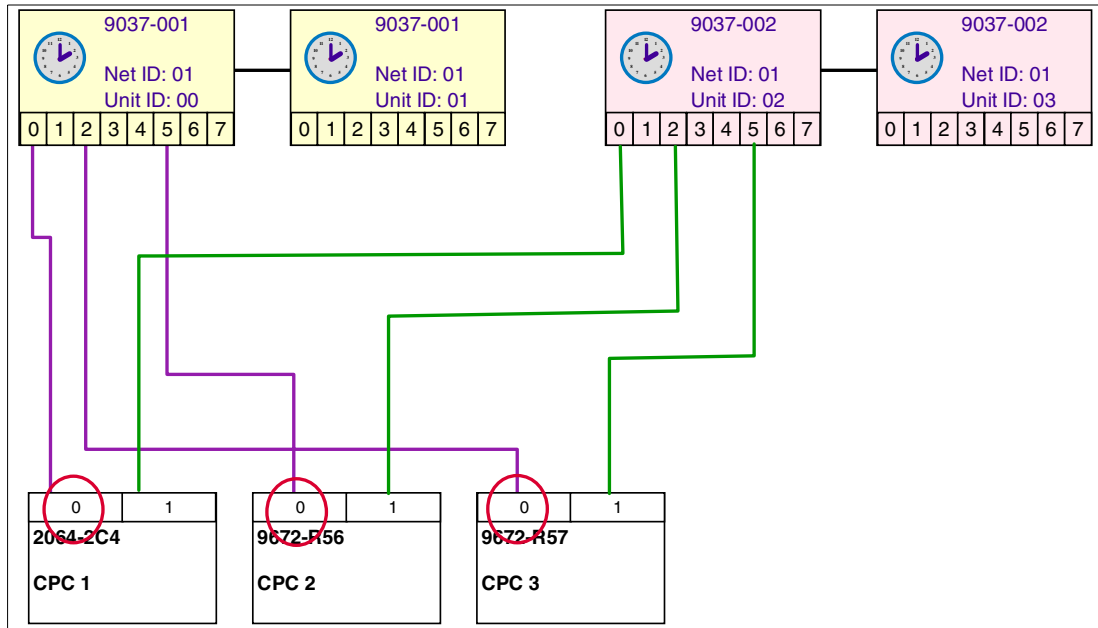


Figure 2-5 Use of non-synchronized Sysplex Timers

If at this point of the migration a failure was to occur in the connection between the 9037-001 and CPC 1 (as shown in Figure 2-6 on page 11) such that Port 0 of CPC 1 becomes non-operational, the hardware in CPC 1 will automatically switch ports and start stepping to signals received on Port 1 from the 9037-002.

At this point, we have CPC 1 receiving timer signals from the 9037-002, and CPCs 2 and 3 receiving timer signals from the unsynchronized 9037-001. Note that each CPC decides independently which timer to step to - just because CPC 1 has switched to Port 1 does not mean that the other CPCs will also switch ports.

Since there are systems on CPCs 1, 2, and 3 that require time coordination, a data integrity exposure now exists, since time stamps for systems in CPC 1 are obtained from the 9037-002, and time stamps for systems in CPCs 2 and 3 are obtained from the 9037-001, which is not synchronized with the 9037-002.

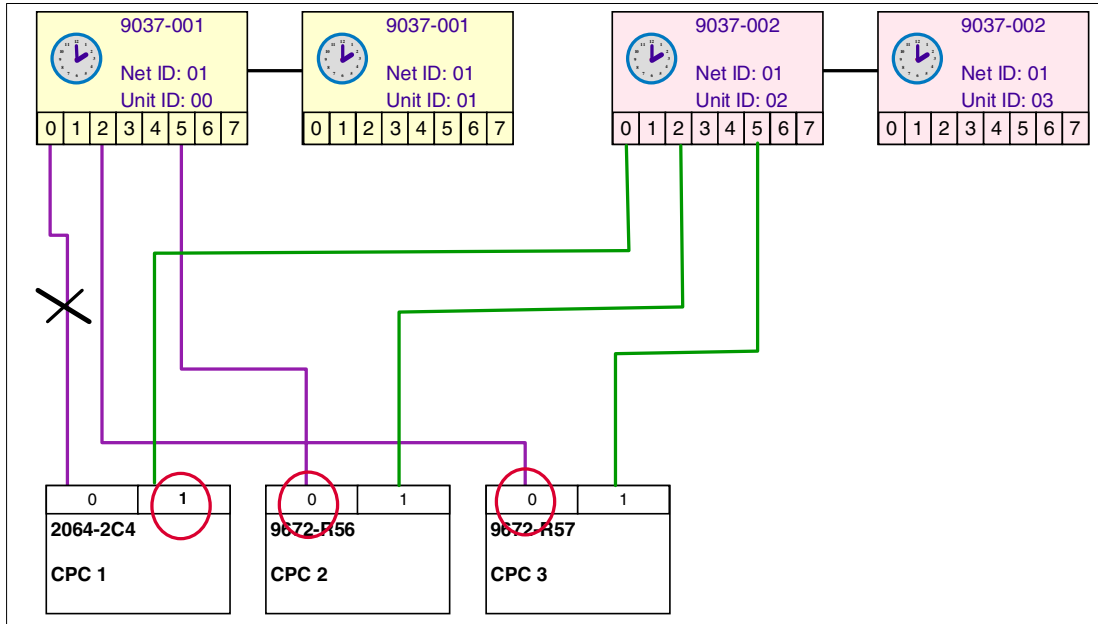


Figure 2-6 Use of non-synchronized Sysplex Timers - data integrity exposure

Therefore, to ensure data integrity, the migration procedure cannot allow systems that require time coordination across multiple CPCs to continue running on more than one CPC.

2.5 Recommended migration procedures

Two procedures have been developed that meet the fundamental requirement of maintaining data integrity during the migration. One procedure results in a limited outage—that is, a subset of the systems will need to be stopped. The other procedure involves a complete planned outage of all systems that require multi-CPC time coordination.

These procedures can be used for either a 9037-001 Basic configuration or Expanded Availability configuration.

The next two chapters describe the planning required for these two procedures.



Migration with limited outages

This chapter provides information about items that must be addressed before starting the migration. It then discusses the planning tasks that must be carried out and ends with a short discussion of how to calculate the impact to your application's availability.

3.1 Overview

This section describes the tasks necessary to plan for a 9037 migration with a limited outage. Before making a final decision on which approach to use, you should review the tasks described in this chapter to determine if this is the most appropriate one for your environment. The risks that are involved in the migration, and the estimated time for limited operations, are also discussed here.

The migration procedure allows all the systems on one of the CPCs to continue normal operations during the migration. On each of the other CPCs, the following systems must be quiesced:

- ▶ Any OS/390 or z/OS systems in multi-CPC sysplexes
- ▶ Any system requiring time coordination across multiple CPCs - XRC is an example of an application that requires this coordination.

By quiescing these systems, there is no data integrity exposure during the migration window. Any systems that do not fall into one of these categories can continue operation.

Remember, however, that if all timer signals are lost, any system that has specified ETRMODE=YES in the CLOCKxx member of SYS1.PARMLIB will either go into a non-restartable wait state, or halt processing and issue a WTOR (depending on whether APAR OW44231 is applied). However, the risk is only to the availability of those systems; there is no data integrity exposure.

Note: The details of the migration procedure are not included in this Redpaper or in *IBM 9037 Model 2 Sysplex Timer Maintenance Information, SY27-2641*, which is shipped with the 9037-002. The actual migration should not be attempted without reference to the detailed instructions and a full understanding of the Sysplex Timer environment.

The migration procedure is available to your IBM representative from the following URL:

http://nascpok.pok.ibm.com/9037/mes/9037MIGR_Procedure.PDF

If the detailed instructions provided by the migration procedure are not followed, the potential result could be a total sysplex outage.

3.2 Service recommendations prior to performing migration

If you use System Automation for OS/390, we recommend that the PTF for APAR OW21776 is installed on all OS/390 systems. This APAR (which is integrated into OS/390 1.3 and later) categorizes new 9037-002 error codes into informational and problem categories, and provides support for System Automation for OS/390.

In addition, we strongly recommend applying the PTF for APAR OW44231. The effect of this APAR is that if a system that has ETRMODE=YES specified loses all Sysplex Timer signals, the system will suspend operation and issue a WTOR. When connectivity to a Sysplex Timer is restored, the WTOR can be replied to, and operation can continue. Without the APAR, the system will enter a non-restartable wait state immediately when the loss of the last Sysplex Timer signal occurs, and an IPL is required to restart operations.

3.3 Planning tasks prior to migration

By reviewing the following tasks, you will be able to determine whether it is possible to migrate to the 9037-002 timers with a limited outage, or if you will require a complete outage of all potentially-affected systems.

1. All the systems must be at MVS/ESA™ 5.1 or higher.
2. Identify the systems that are potentially affected. See 1.3, “Which systems are potentially affected” on page 2 for more information.

Select the CPC that will continue normal processing during the migration, and identify the systems on each of the remaining CPCs that need to be quiesced during the migration. This decision could be made based on a number of factors such as where critical applications are being run, or which CPC has the most capacity. Consideration should also be given to whether subsystems that need to be quiesced can be moved to an OS/390 or z/OS image running in a logical partition on the CPC that will continue normal processing during the migration.

For the configuration example shown in Figure 3-1, CPC 1 is an IBM 2064, and CPCs 2 and 3 are S/390 9672s. A customer could select CPC 1 to continue processing and quiesce affected systems on CPCs 2 and 3 during the migration window. This may allow you to keep a significant portion of the operations available during the migration.

If your configuration or your operations do not allow you to make such a selection, your only option is to follow the procedure documented in Chapter 4, “Migration with planned outages” on page 19.

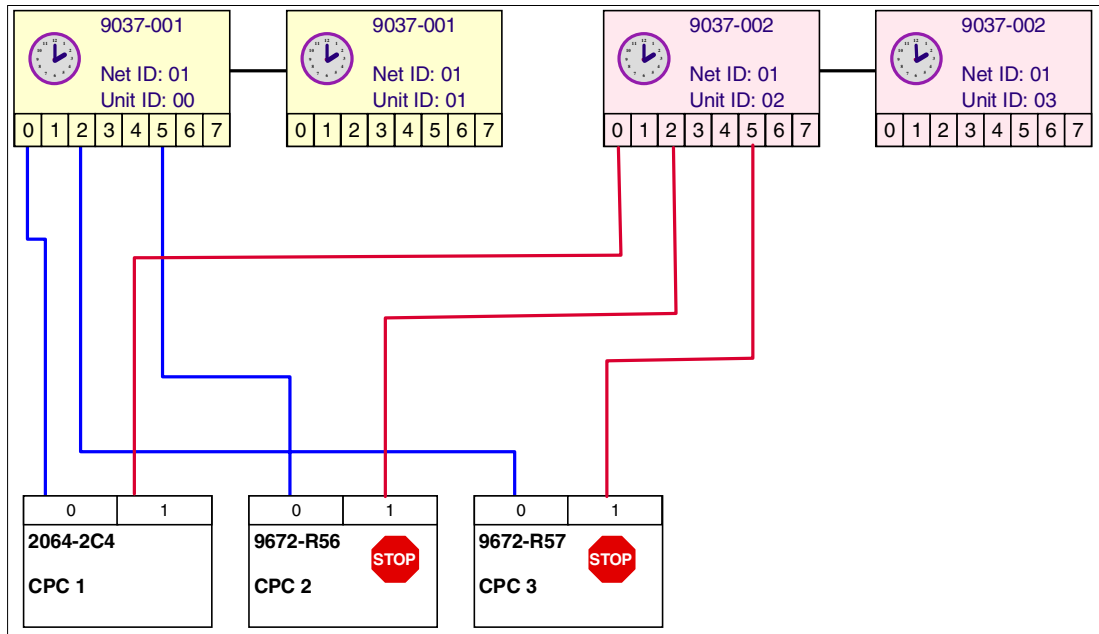


Figure 3-1 Example of Multi-CPC sysplex with three CPCs

3. If you are migrating from a 9037-001 Expanded Availability configuration to a 9037-002 Expanded Availability configuration, the following verification should be performed:
 - Determine which CPCs are attached to the 9037-001 ports being migrated.

- At the OS/390 or z/OS console, use the command **D ETR,DATA** for each system running on these CPCs.

If all the systems are in the same sysplex, you can issue this command on all systems with the Route command:

```
RO *ALL,D ETR,DATA
```

Following is an example of this display:

```
IEA282I hh.mm.ss ETR STATUS
Synchronization Mode = ETR                      CPC Side = 0
CPC Port 0 <=== ACTIVE                          CPC Port 1
Operational                                       Operational
Enabled                                           Enabled
ETR NET ID = 15                                  ETR NET ID = 15
ETR PORT = 01                                    ETR PORT = 01
ETR ID = 03                                       ETR ID = 04
```

Note: In this example, CPC Port 0 is the active port and CPC Port 1 is the alternate port.

- For each attached CPC, you need to verify that each system that has specified ETRMODE=YES (Synchronization Mode = ETR in example shown above) has an operational and enabled alternate port, with an ETR ID that is *different* than the ETR ID displayed for the active port.

Note: If there are multiple LPARs on a CPC, it is important that this verification be made for *every* system. It is possible that just one of the systems on a CPC has disabled an ETR port.

- Any systems that do not meet this criteria will either need to have the problem corrected, or they will need to be shut down. Otherwise, they will either enter a non-restartable wait state or halt processing during the migration.

4. If you are migrating from a 9037-001 Basic configuration to a 9037-002 Basic configuration or Expanded Availability configuration, the following verification should be performed:

- Determine which CPCs are attached to the 9037-001 ports being migrated.
- At the OS/390 or z/OS console, use the command, **D ETR,DATA** for each system running on these CPCs.

If all the systems are in the same sysplex, you can issue this command on all systems with the Route command:

```
RO *ALL,D ETR,DATA
```

An example of this display is shown above.

Note: In this example, CPC Port 0 is the active port and CPC Port 1 is the alternate port.

- For each attached CPC, you need to verify that each system that has specified ETRMODE=YES (Synchronization Mode = ETR in example shown above) has an operational and enabled alternate port.

Note: If there are multiple LPARs on a CPC, it is important that this verification be made for every system. It is possible that one of the systems on a CPC has disabled an ETR port.

- Systems that do not meet this criteria will either need to have the problem corrected, or they will need to be shut down. Otherwise, they will either enter a non-restartable wait state or halt processing during the migration.

5. If your 9037-001 configuration is using an External Time Source (ETS), and you are planning to use ETS for the 9037-002, you can skip task 6 on page 17 and go directly to task 7 on page 17.
6. If you will be setting the time at the 9037-002 manually, then for each image that has ETRMODE=YES, determine the setting of ETRDELTA (defined in the CLOCKxx member of SYS1.PARMLIB). Note the lowest value of ETRDELTA in all the systems in the sysplex. The migration may be difficult to perform if this value is less than 5 seconds. The default value of ETRDELTA is 10 seconds, which is the value recommended by IBM.

Note that this value cannot be changed without an IPL. If you find some systems where the value is less than 5, you should plan on increasing that value as part of any planned IPLs before the migration starts.

7. If your 9037-001 configuration is using an External Time Source (ETS), and you are planning to use ETS for the 9037-002, verify that the 9037-001 that will remain connected to the CPCs (that is, the one with Unit ID=00 in Figure 3-1 on page 15) is tracking the ETS device. To check this:
 - Select **Time** on the Main menu of the 9037-001 console
 - Select the **Adjust the Time** option.
 - Note:** Do *not* select the Set the Time option.
 - Select **F6=Calculate Adjustment**
 - If the time difference is less than 4.999 seconds, the 9037-001 is tracking the ETS.
 - Apply the adjustment if it is greater than zero but less than 4.999 seconds.
8. Install either one 9037-002 unit (if you are planning a Basic configuration or Expanded Basic configuration) or both 9037-002 units (if you are planning an Expanded Availability configuration) using the procedures provided in *IBM 9037 Model 2 Sysplex Timer Maintenance Information, SY27-2641*. This should include laying the timer-to-CPC and timer-to-timer cables.

The migration will be greatly facilitated if the 9037-002 active console is physically adjacent to the 9037-001 installation.

For a 9037-002 Expanded Availability configuration, it is recommended that the second 9037-002 unit be physically separated from the first unit in order to minimize any environmental single points of failure. If you need to relocate 9037-002 units after the migration, you will be able to do so concurrently, without any disruption to your operation, by using procedures described in *IBM 9037 Model 2 Sysplex Timer Maintenance Information, SY27-2641*.

9. It is not a requirement to keep port assignments of 9037-001 and 9037-002 ports the same. However, it is desirable not to change any patch panel routing during the migration in order to minimize the risk of mis-configuration.

If you have successfully completed these steps, you can schedule time with your service representative to perform the migration with a limited sysplex outage.

3.4 Risks during migration

There are no data integrity exposures by following this migration procedure, as long as all the potentially affected systems have been quiesced on all CPCs, except the one selected in task 2 on page 15.

Note: As stated before, if the migration procedure is not properly followed, it may result in an unscheduled outage of one or more OS/390 or z/OS images running with ETRMODE=YES. If

an OS/390 or z/OS image running with ETRMODE=YES does not receive timing signals from the attached timers, it will enter a non-restartable wait state or halt processing.

3.5 Estimated time of limited operations

If all the planning tasks described previously are performed prior to the migration, the time for limited operations will include:

- ▶ The time to shut down any systems that must be quiesced.
- ▶ Approximately one hour to perform the 9037 migration.
- ▶ The time to restart the quiesced systems.

While the time to quiesce and restart your systems will vary from installation to installation, you should allow 1 to 2 hours to perform the 9037 migration.



Migration with planned outages

This chapter discusses the considerations and planning tasks required if you decide to take the option of shutting down all affected systems during the migration window.

4.1 Overview

This section describes the tasks that can be performed prior to migrating with a planned complete outage, so that the duration of the outage is minimized.

To perform the migration using this option, you must plan an outage of all systems in CPCs connected to the affected timers that are running with ETRMODE=YES, plus any other systems that require time coordination across multiple CPCs (Linux systems using XRC, for example).

Note: The details of the migration procedure are not included in this Redpaper or in *IBM 9037 Model 2 Sysplex Timer Maintenance Information, SY27-2641*, which is shipped with the 9037-002. The actual migration should not be attempted without reference to the detailed instructions and a full understanding of the Sysplex Timer environment.

The migration procedure is available to your IBM representative from the following URL:

http://nascpok.pok.ibm.com/9037/mes/9037MIGR_Procedure.PDF

4.2 Service recommendations prior to performing migration

If you use System Automation for OS/390, we recommend that the PTF for APAR OW21776 is installed on all OS/390 systems. This APAR (which is integrated into OS/390 1.3 and later) categorizes new 9037-002 error codes into informational and problem categories, and provides support for System Automation for OS/390.

4.3 Planning tasks prior to migration

The duration of the planned outage can be minimized by installing either one 9037-002 unit (if you are planning a Basic configuration or Expanded Basic configuration) or both 9037-002 units (if you are planning an Expanded Availability configuration) using the procedures contained in *IBM 9037 Model 2 Sysplex Timer Maintenance Information, SY27-2641*.

The migration will be greatly facilitated if the 9037-002 active console and the 9037-002 unit attached to it are physically adjacent to the 9037-001 installation.

For a 9037-002 Expanded Availability configuration, it is recommended that the second 9037-002 unit be physically separated from the first unit in order to minimize any environmental single points of failure. If you need to relocate 9037-002 units after the migration, you will be able to do so concurrently, without any disruption to your operation using procedures outlined in *IBM 9037 Model 2 Sysplex Timer Maintenance Information, SY27-2641*.

When you have completed the above steps, you can schedule the planned outage to quiesce each affected system and connect the ports of the 9037-002 units to the appropriate CPCs.

4.4 Risks during migration

There are no risks to data integrity or unplanned system outages if you use this migration procedure.

4.5 Estimated time of outage

If all the planning tasks are performed prior to the migration, the actual outage could be restricted to the time required to:

- ▶ Quiesce each affected system.
- ▶ Unplug fiber optic cables from the 9037-001 units and plug them into the corresponding 9037-002 units.
- ▶ Re-IPL each image.



A

Connecting 9037-002s at extended distances

The 9037-002 Sysplex Timer has more options for extended distance operation than the 9037-001, and some customers may wish to avail of this greater flexibility as part of the migration from the 9037-001. This appendix briefly describes the options available for connecting 9037-002 timers at extended distances from each other and from the CPCs they are attached to.

Expanded Availability configuration options

When looking at the distances supported with 9037-002s, there are two things that must be considered:

- ▶ The distance from one 9037-002 to the other 9037-002.
- ▶ The distance from a 9037-002 to each of the CPCs it is attached to. Note that while the distance between the 9037s is changed compared to 9037-001s when using 9037-002s, the supported distance between the timer and each attached CPC is the same for 9037-001 as for a 9037-002.

Note: All distances referred to in this section for 9037-002 are distances of fiber links.

There are three options for configuring 9037-002s attached to each other over distances greater than the 3 meters that is supported with 9037-001s:

- ▶ For distances up to 3 Km, the 9037-002s can be connected to each other directly using either 62.5/125-micrometer or 50/125-micrometer multimode fiber, and without requiring any additional hardware or RPQs.
- ▶ For distances greater than 3 Km and less than 26 Km, you have two choices:
 - You can use 9036-003 repeaters, orderable via RPQ 8K1919. Two RPQ 8K1919s are required for *each* link (remember that there are two links connecting the 9037-002s, so you will need four RPQs to connect the two 9037-002s). This RPQ provides one 9036-003 repeater. Note that the length of each multimode cable cannot exceed 3 Km, and the distance between the repeaters cannot exceed 20 Km.
 - You can use an IBM-qualified Dense Wavelength Division Multiplexer (DWDM). Note that the length of each multimode cable cannot exceed 3 Km and the total fibre distance between the 9037-002s cannot exceed 26 Km.
- ▶ For distances greater than 26 Km and less than 40 Km, you must use an IBM-qualified DWDM. You must order RPQ 8P1955. Note that the length of each multimode cable cannot exceed 3 Km, and the total fibre distance between the 9037-002s cannot exceed 40 Km.

Note: The greatest distance supported between 9037-002s is 40 fibre Km. Extending the link distance beyond 40 Km could result in the two timers being out of synchronization with each other, thereby resulting in a *potential data integrity exposure*. The out-of-synchronization condition is not detectable during normal system operation.

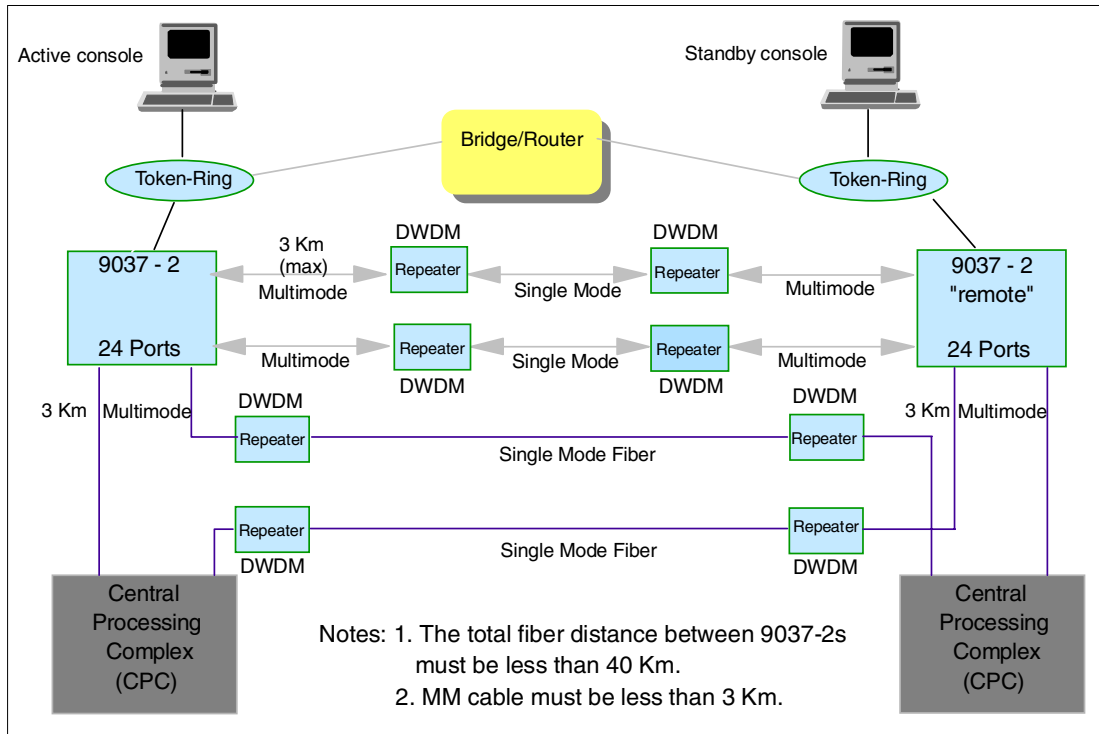


Figure 4-1 Configuring 9037-002s at extended distances

The other distance that must be considered is the distance between each 9037-002 and the CPC it is connected to. The options in this case are similar but not identical:

- ▶ For distances up to 3 Km, the Sysplex Timers can be connected to the target CPC directly using either 62.5/125-micrometer or 50/125-micrometer multimode fiber, and without requiring any additional hardware or RPQs.
- ▶ For distances greater than 3 Km and less than 26 Km, you have two choices:
 - You can use 9036 repeaters, orderable via RPQ 8K1919. Two RPQ 8K1919s are required for *each* link. This RPQ provides one 9036-003 repeater.
 - You can use an IBM-qualified Dense Wavelength Division Multiplexer (DWDM). Note that the length of each multimode cable cannot exceed 3 Km and the total fibre distance between the 9037-002s cannot exceed 26 Km.
- ▶ For distances greater than 26 Km and less than 50 Km (note that this is longer than the supported Sysplex Timer-to-Sysplex Timer distance of 40 Km), you must use an IBM-qualified DWDM. You must order RPQ 8P1955. Note that the length of each multimode cable cannot exceed 3 Km and the total fibre distance between the 9037-002 and the connected CPC cannot exceed 50 Km.

Note that the greatest distance currently supported between a 9037-002 and a CPC that it is attached to is 50 fibre Km.

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this Redpaper.

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 27. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *S/390 Time Management and IBM 9037 Sysplex Timer*, SG24-2070

Other publications

These publications are also relevant as further information sources:

- ▶ *IBM 9037 Sysplex Timer and S/390 Time Management*, GG66-3264
- ▶ *Planning for the 9037 Model 2 Sysplex Timer*, SA22-7233.
- ▶ *IBM 9037 Model 2 Sysplex Timer Maintenance Information*, SY27-2641
- ▶ *Using the 9037 Model 2 Sysplex Timer*, SA22-7230

How to get IBM Redbooks

You can search for, view, or download Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Index

Numerics

- 9036 repeaters, using to extend Sysplex Timer distances 24
- 9037-001
 - end-of-service announcement 6
 - supported configurations 6
 - upgrade support 9
- 9037-002
 - enhancements compared to 9037-001 6
 - supported configurations 7
 - supported distances between timers 6, 24

A

- affected systems 2
- APAR OW21776 14, 20
- APAR OW44231 3, 14, 18

B

- Basic configuration 8

C

- considerations for using non-synchronized Sysplex Timers 9

D

- data integrity considerations 2, 9–10, 17
- detailed migration procedure 2
 - obtaining 2
- displaying ETR port status 16
- DWDM
 - use to extend Sysplex Timer distances 25

E

- ETRDELTA recommended values 17
- ETRMODE=YES 14, 17, 20
- Expanded Availability configuration 7
- Expanded Basic configuration 7
- extended distance support 23
 - 9037-to-9037 distances 24
 - 9037-to-CPC distances 25
 - required hardware 24
- External Time Source considerations for migration 17

I

- identifying affected systems 2
- identifying the CPC that will be unaffected 15

L

- limited outage migration
 - estimating impact duration 18

- identifying which CPC to retain 15
- Linux systems considerations 3

P

- planned outage migration 20
- planning tasks 15
- port assignments 17
- prerequisite APARs 14
- prerequisites for a limited outage migration 15

R

- Redbooks Web site 27
 - Contact us viii
- RPQ 8K1919 24–25
- RPQ 8P1955 25

S

- sysplex considerations 2
- System Automation for OS/390 support for Sysplex Timer 14

X

- XRC
 - considerations for timer migration 3



Migration Planning for the 9037 Model 2 Sysplex Timer



Planning

The 9037-001 Sysplex Timer has an announced worldwide end-of-service (EOS) date of December 31, 2003. This critical component of any multi-CPC sysplex will no longer be supported by IBM after December 31, 2003.

Sysplex considerations

External Time Source considerations

IBM recommends replacing any 9037-001 models with Sysplex Timer Model 2s (9037-002). The 9037-002 is a full replacement for the 9037-001. There is no upgrade path from the 9037-001 to the 9037-002.

This Redpaper provides the information required for customers to successfully plan for the migration of their 9037-001 Sysplex Timers to 9037-002 Sysplex Timers.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks