

Server Time Protocol Planning Guide

Concepts and terminology

Configuration planning

Operations and recovery
planning



Octavian Lascu
Hans-Peter Eckam
George Kozakos
Paulo Vitor Pereira

Redbooks



International Technical Support Organization

Server Time Protocol Planning Guide

June 2013

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

Fourth Edition (June 2013)

This edition applies to Server Time Protocol Facility on IBM System z10 Enterprise Class (z10 EC), System z10 Business Class (z10 BC), System z9 Enterprise Class (z9 EC), System z9 Business Class (z9 BC), and zSeries z990 and z890 servers.

© Copyright International Business Machines Corporation 2006, 2013. 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	vii
Trademarks	viii
Preface	ix
Authors	ix
Now you can become a published author, too!	x
Comments welcome	xi
Stay connected to IBM Redbooks	xi
Summary of changes	xiii
June 2013, Fourth Edition	xiii
Chapter 1. Introduction to Server Time Protocol	1
1.1 Introduction to time synchronization	2
1.1.1 Time-of-Day clock	3
1.1.2 Time synchronization in a Parallel Sysplex	3
1.1.3 Overview of External Time Reference	4
1.2 Overview of Server Time Protocol	6
1.3 STP concepts and terminology	8
1.3.1 STP facility	8
1.3.2 TOD clock synchronization	9
1.3.3 STP servers	10
1.3.4 Coordinated Timing Network	11
1.3.5 STP stratum	12
1.3.6 Server roles in an STP-only CTN	14
1.4 Timing network types	14
1.4.1 Mixed CTN	14
1.4.2 STP-only CTN	18
1.5 Coupling links	22
1.5.1 Link redundancy	23
1.5.2 Coupling link requirements in a non-sysplex configuration	23
1.6 External time source	25
1.7 Summary of STP terminology	27
Chapter 2. Planning hardware and software	31
2.1 Planning hardware	32
2.1.1 Servers	32
2.1.2 Hardware Management Console	33
2.1.3 Support Elements	36
2.1.4 EC levels and MCLs	37
2.1.5 Installing the STP feature	37
2.1.6 z/BX considerations	38
2.2 Planning connectivity	38
2.2.1 Planning for coupling links for STP	38
2.2.2 Planning for timing-only links	44
2.2.3 Considerations for multi-site sysplex	46
2.3 External time source	46
2.3.1 Dial-out on the Hardware Management Console	47
2.3.2 NTP server	48

2.3.3	NTP server with pulse per second	54
2.4	Internal Battery Feature (IBF)	57
2.5	Planning z/OS software	57
2.5.1	z/OS requirements for STP	57
2.5.2	CLOCKxx	59
2.6	Planning for z/VM	67
2.6.1	STP state flow in z/VM	68
2.6.2	Updated CP commands	69
2.6.3	z/VM IOCP	69
2.6.4	VM guests	69
2.6.5	z/TPF	70
2.7	Special considerations for a single server STP-only CTN	70
2.7.1	Software planning: z/OS currently running in a sysplex	72
2.7.2	Software planning: z/OS systems currently not in a sysplex	73
2.7.3	Software planning: z/OS systems currently in a sysplex using ETR	74
2.8	Planning for zBX	74
2.9	Operating systems with no STP support	75
2.9.1	Linux for System z	75
2.9.2	z/VSE	75
Chapter 3.	Operations	77
3.1	Displaying and monitoring the environment	78
3.1.1	CPC Detail panel	78
3.1.2	System (Sysplex) Time panel	79
3.1.3	z/OS display commands	86
3.1.4	z/OS messages	90
3.2	Configuring the Coordinated Timing Network	92
3.2.1	Configuring a Mixed CTN	92
3.2.2	Configuring an STP-only CTN	94
3.3	Managing the time	106
3.3.1	STP time adjustment	106
3.3.2	STP offset adjustments	108
3.3.3	Changes in local time	115
3.3.4	Logical partition time offset	118
3.3.5	Parallel Sysplex and multiple time zones	120
3.4	Managing the CTN configuration	125
3.4.1	Changing the CTN ID	125
3.4.2	Single-server or dual-server auto resume after power-on reset	125
3.4.3	Changing the Current Time Server	127
3.4.4	Changing the server roles	129
3.4.5	Deconfiguring a CTN	133
3.5	Operational considerations	134
3.5.1	Migrating timing-only links to coupling links	134
3.5.2	Coupling link path validation	135
3.5.3	Last timing link validation and operational safeguards	136
3.5.4	Last timing link condition examples	137
3.5.5	Disruptive actions on the Current Time Server	143
3.5.6	Disruptive actions on the BTS or Arbiter	145
3.5.7	Restarting a CTN power outage or power-on reset (unplanned)	145
3.5.8	Planned single-site maintenance	148
3.6	MES upgrade considerations	148
3.6.1	Mixed CTN	149
3.6.2	STP-only CTN	149

Chapter 4. Recovery concepts	153
4.1 Terminology overview	154
4.2 Freewheel interval	154
4.3 Server offline signal	155
4.4 Going Away Signal	155
4.5 Arbiter-assisted recovery	156
4.6 Console-assisted recovery	158
4.6.1 Console-assisted recovery in a CTN with BTS	158
4.6.2 Console-assisted recovery in a CTN with BTS and Arbiter	158
4.7 Island condition	159
4.8 Switch to local timing mode	159
4.9 External Time Source (ETS)	160
Appendix A. Migration planning	163
Preparation and installation	164
Hardware installation	164
z/OS preparation and installation	166
Server roles in an STP-only CTN	166
Plan for high availability	168
Migration paths	169
Migration between an ETR network and a Mixed CTN	172
Migration from an ETR network to a Mixed CTN	172
Migration from a Mixed CTN to an ETR network	176
Migration between a Mixed CTN and an STP-only CTN	177
Migration from a Mixed CTN to an STP-only CTN	178
Migration from an STP-only CTN to a Mixed CTN	182
Migration from an ETR network to an STP-only CTN	184
Installation of a new STP-only CTN	189
Preparation	190
Activation	190
Installation of a single-server STP-only CTN	191
Preparation	192
Activation	193
Enhancing a single-site ETR network to a Mixed CTN across two sites	194
Appendix B. Server Time Protocol (STP) messages in z/OS	199
Supervisor messages	200
XCF messages	201
XES messages	201
Automating-related messages	201
Related publications	203
IBM Redbooks publications	203
Other publications	203
Online resources	204
How to get Redbooks publications	205
Help from IBM	205
Index	207

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 illustrate 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.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

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

BladeCenter®	POWER7®	System x®
CICSplex®	PowerVM®	System z10®
CICS®	PR/SM™	System z9®
DataPower®	Processor Resource/Systems Manager™	System z®
ESCON®	Redbooks®	z/Architecture®
FICON®	Redpapers™	z/OS®
GDPS®	Redbooks (logo)  ®	z/VM®
Geographically Dispersed Parallel Sysplex™	Resource Link®	z/VSE®
Global Technology Services®	RMF™	z10™
IBM®	S/390®	z9®
MVS™	Sysplex Timer®	zEnterprise®
Parallel Sysplex®	System Storage®	zSeries®

The following terms are trademarks of other companies:

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

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

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

Preface

Server Time Protocol (STP) is a server-wide facility that is implemented in the Licensed Internal Code (LIC) of the IBM® zEnterprise® EC12 (zEC12), IBM zEnterprise 196 (z196), IBM zEnterprise 114 (z114), IBM System z10® Enterprise Class (z10 EC), System z10 Business Class (z10 BC), IBM System z9® Enterprise Class (z9 EC), and System z9 Business Class (z9 BC). It provides improved time synchronization in a sysplex or non-sysplex configuration.

This IBM Redbooks® publication is intended for infrastructure architects and system programmers who need to understand the IBM STP functions. Readers are expected to be generally familiar with System z® technology and terminology.

This book provides planning information for Server Time Protocol functions and associated software support. For more detailed installation, operation, and recovery information, refer to the companion books *Server Time Protocol Implementation Guide*, SG24-7281, and *Server Time Protocol Recovery Guide*, SG24-7380.

Authors

This book was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

Octavian Lascu is a Project Leader at the International Technical Support Organization, Poughkeepsie Center. He writes extensively and teaches IBM classes worldwide on areas such as cross-platform clustering and High Performance Computing. Before joining the ITSO 10 years ago, Octavian worked in Global Technology Services®, Romania, as Senior IT Consultant.

Hans-Peter Eckam is a Support Center Specialist in the IBM Hardware Support Center in Mainz, Germany. Hans-Peter joined IBM in 1986 and has 22 years of experience in supporting System z customers. His areas of expertise include System z, Sysplex Timer®, Server Time Protocol, and connectivity such as ESCON®, FICON®, and InfiniBand.

George Kozakos is a Senior IT Specialist with IBM Australia. He has more than 25 years of experience in MVS™ system programming. George's areas of expertise include Server Time Protocol and GDPS®. He holds degrees in Computing Science and Pure Mathematics.

Paulo Vitor Pereira is an IT Specialist in STG Lab Services in São Paulo, Brazil. He has joined IBM in 2009 and has 6 years of experience in System z. His areas of expertise include, z/OS®, WLM, z/BX and System z Hardware.

Thanks to the following people for their contributions to this project:

William G White
International Technical Support Organization, Poughkeepsie Center

Donald J Smith, Richard M Hamilton, Steve Cherniak, Glen Poulsen
IBM Poughkeepsie

Gregory K. Hutchison
IBM Gaithersburg

Thanks to the authors of the previous editions of this book:

- ▶ Authors of third edition of *Server Time Protocol Planning Guide*, published in September 2010 were:
Noshir Dhondy, Hans-Peter Eckam, Andy Kilhoffer, Jeremy Koch, Hao Shen, Martin Söllig, Ray Takabe, Octavian Lascu.
- ▶ Authors of the second edition, *Server Time Protocol Planning Guide*, published in June 2009, were:
Noshir Dhondy, Hans-Peter Eckam, Marcel Jaeckel, Jeff McDonough, George Ng, Frank Packheiser, and Bernard Tanguy
- ▶ Authors of the first edition of *Server Time Protocol Planning Guide*, published in December 2006 and *Server Time Protocol NTP Client Support*, published in September 2007, were:
Noshir Dhondy, Greg Hutchison, Paul Hagen, Dick Jorna, George Kozakos, and Iain Neville.

Thanks to the ITSO support team:
Robert Haimowiz, Donald R Brennan and Roy P Costa
International Technical Support Organization, Poughkeepsie Center

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author - all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

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 books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

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

ibm.com/redbooks

- ▶ Send your comments in an e-mail to:

redbooks@us.ibm.com

- ▶ Mail your comments to:

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

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:

<http://www.facebook.com/pages/IBM-Redbooks/178023492563?ref=ts>

- ▶ Follow us on twitter:

<http://twitter.com/ibmredbooks>

- ▶ Look for us on LinkedIn:

<http://www.linkedin.com/groups?home=&gid=2130806>

- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>

- ▶ Stay current on recent Redbooks publications with RSS Feeds:

<http://www.redbooks.ibm.com/rss.html>

Summary of changes

This section describes the technical changes made in this edition of the book. This edition may also include minor corrections and editorial changes that are not identified.

Summary of Changes
for SG24-7280-03
for Server Time Protocol Planning Guide
as created or updated on April 11, 2016.

June 2013, Fourth Edition

This revision reflects the addition, deletion, or modification of new and changed information described below.

New information

- ▶ STP Version 4 support updates
- ▶ Information update includes IBM zEnterprise Servers (zEC12, z196, and z114).
- ▶ Dual-server auto resume after power-on reset.
- ▶ External Time Source: Hardware Management Console acting as NTP server.
- ▶ Removal of HMC dial out support



Introduction to Server Time Protocol

This chapter discusses time synchronization and introduces the Server Time Protocol (STP) facility. The following topics are covered:

- ▶ Introduction to time synchronization
- ▶ Overview of Server Time Protocol
- ▶ STP concepts and terminology
- ▶ Timing network types
- ▶ Coupling links
- ▶ External time source
- ▶ Summary of STP terminology

1.1 Introduction to time synchronization

Historically, the most important requirement for highly accurate time is for navigational purposes. For applications such as very precise navigation and satellite tracking, which must be referenced to the earth's rotation, a time scale that is consistent with the earth's rotation must be used.

In the information technology world, time synchronization has become a critical component for managing the correct order of the events in distributed applications (transaction processing, message logging), especially for audit and legal purposes.

Today, this time scale is known as *Universal Time 1* (UT1). UT1 is computed using astronomical data from observatories around the world. It does not advance at a fixed rate, but speeds up and slows down with the earth's rotation rate. While UT1 is actually measured relative to the rotation of the earth with respect to distant stars, it is defined in terms of the length of the mean solar day. This makes it more consistent with civil, or solar, time.

Until 1967, the *second* was defined on the basis of UT1. Since 1967, the internationally accepted definition of the second has been "9 192 631 770 periods of the radiation corresponding to the transition between two hyperfine levels of the ground state of the cesium-133 atom." In 1967, this definition was already 1000 times more accurate than what could be achieved by astronomical methods, and today it is even more accurate. The atomic definition of the second is primarily aimed at providing an accurate measure of time intervals. At the same time, the need for an accurate time-of-day measure was recognized, leading to the adoption of the following two basic scales of time:

- ▶ International Atomic Time (TAI¹)

TAI is based solely on an atomic reference. It provides an accurate time base that increases at a constant rate with no discontinuities.

- ▶ Coordinated Universal Time (UTC² - Universal Time, Coordinated)

UTC is derived from TAI, and is adjusted to keep reasonably close to UT1. UTC is the official replacement for (and generally equivalent to) the better known Greenwich Mean Time (GMT).

Since January 1, 1972, an occasional correction of exactly one second, called a *leap second*, has been inserted into the UTC time scale. It kept UTC time within ± 0.9 seconds of UT1 at all times. These leap seconds have always been positive (in theory they can be positive or negative). Correction is coordinated under international agreement by the Bureau International des Poids et Mesures (BIPM) in Paris. This adjustment occurs at the end of a UTC month, which is normally on June 30 or December 31. The last minute of a corrected month can, therefore, have either a positive adjustment to 61 seconds or a reduction to 59 seconds. Twenty-five positive leap seconds have been introduced into the UTC time scale as of July 1, 2012.

Note: The effect of a leap second is the introduction of an irregularity into the UTC time scale, so exact interval measurements are not possible using UTC, unless the leap seconds are included in the calculations. After every positive leap second, the difference between TAI and UTC increases by one second.

¹ From the French name Temps Atomique International

² Compromise between CUT for *coordinated universal time* and TUC for *temps universel coordonné* (French)

1.1.1 Time-of-Day clock

The Time-of-Day (TOD) clock was introduced as part of the System/370 architecture to provide a high-resolution measure of real time, suitable for the indication of date and time of day. The cycle of the clock is approximately 143 years and wraps on September 18, 2042.

In July 1999, the extended TOD clock facility was announced. The TOD clock was extended by 40 bits on the right. This 104-bit value, along with 8 zero bits on the left and a 16-bit programmable field on the right, can be stored by the problem program instruction STORE CLOCK EXTENDED (STCKE). With proper support from the operating system, this provides for unique timestamps across the sysplex. The value of the TOD clock is directly available to application programs by using the instructions STORE CLOCK (STCK), STORE CLOCK FAST (STCKF), and STORE CLOCK EXTENDED (STCKE). These instructions store the value of the clock in a storage location specified by the instruction. Figure 1-1 shows the format of the TOD clock.




Figure 1-1 TOD format

Conceptually, the TOD clock is incremented so that a 1 (one) is added into bit position 51 every microsecond. Actual TOD-clock implementations might not provide a full 104-bit counter, but maintain an equivalent stepping rate by incrementing a different bit at such a frequency that the rate of advancing the clock is the same as though a 1 were added in bit position 51 every microsecond.

Figure 1-1 shows the stepping rate (rate at which the bit positions change) for selected TOD clock bit positions. A carry-out of bit 32 of the TOD clock occurs every 2^{20} microseconds (1.048576 seconds). This interval is sometimes called a *mega-microsecond*.

The use of a binary counter for the time of day, such as the TOD clock, requires the specification of a time origin, or epoch. *Epoch* is the time at which the TOD clock value would have been all zeros. The z/Architecture®, ESA/390, and System/370 architectures established the epoch for the TOD clock as January 1, 1900, 0 a.m. GMT.

1.1.2 Time synchronization in a Parallel Sysplex

In z/Architecture, the STCK and STCKE instructions provide a means by which programs can both establish time-of-day and unambiguously determine the ordering of serialized events, such as updates to a database, a log file, or another data structure. The architecture requires that the TOD clock resolution be sufficient to ensure that every value stored by a STCK or STCKE instruction is unique. Consecutive STCK or STCKE instructions executed, possibly on different CPUs in the same server, must always produce increasing values. Thus, the time stamps can be used to reconstruct, recover, or in many different ways ensure the ordering of these serialized updates to shared data.

Parallel Sysplex® extended this requirement to the scope of an entire sysplex. Specifically, in a Parallel Sysplex, the processes are multisystem processes executing on different servers in the same sysplex. Therefore, time consistency must be maintained across servers. To accommodate this requirement for a Parallel Sysplex, the TOD clock architecture was extended by introducing two new major components:

- ▶ External Time References (ETR) architecture

The Sysplex Timer provides an external master clock that can serve as the primary time reference, with a link connecting each server in the Parallel Sysplex to the Sysplex Timer. The External Time References architecture link is used for both setting the local TOD clocks and maintaining time synchronization between the clocks.

- ▶ TOD-Clock Synchronization Facility

The TOD-Clock Synchronization Facility provides an interface between the operating system (OS) and the Sysplex Timer in order to allow the OS to:

- Coordinate setting the local clocks with the Sysplex Timer.
- React to losses in synchronization through an external interruption so that data integrity is maintained.

The process to achieve time consistency is:

1. Server A executes a STCK instruction (time stamp x), which places the clock contents in storage.
2. Server A signals server B.
3. Server B, on receipt of the signal, immediately executes STCK (time stamp y).

For time stamp x and time stamp y to reflect the fact that y is later than x, the two TOD clocks must agree within the time required to send the signal. The consistency required is limited by the time required for signalling between the coupled servers and the time required by the STCK instruction itself.

In practical terms, this means that the Sysplex Timer is required to keep the TOD clocks of all participating systems synchronized with each other to within a small number of microseconds, dictated by the fastest possible passing of data from one system to another through a coupling facility (CF) structure.

1.1.3 Overview of External Time Reference

Note: For more details on Sysplex Timer refer to *S/390® Time Management and IBM 9037 Sysplex Timer*, SG24-2070.

The IBM External Time Reference architecture provides a means of synchronizing TOD clocks in different servers with a centralized time reference, which in turn can be set accurately based on an international time standard. The architecture defines a time-signal protocol and a distribution network, called the ETR network, which permits accurate setting and maintenance of consistency of TOD clocks.

An ETR network consists of the following three elements, which are configured in a network:

- ▶ Sysplex Timer sending unit
- ▶ Sysplex Timer link
- ▶ Sysplex Timer receiving unit

The *sending unit* is the centralized, external time reference that transmits ETR signals over dedicated Sysplex Timer links. It also provides a means by which Sysplex Timer time can be accurately maintained with respect to external standard time services.

The *receiving unit* in each server receives the Sysplex Timer signals and includes the means by which the TOD clocks are set and maintained consistent with Sysplex Timer time. Figure 1-2 shows a typical ETR network, which connects the sending unit to all servers in an installation. The ETR network might comprise one or more sysplexes and servers not belonging to a sysplex.

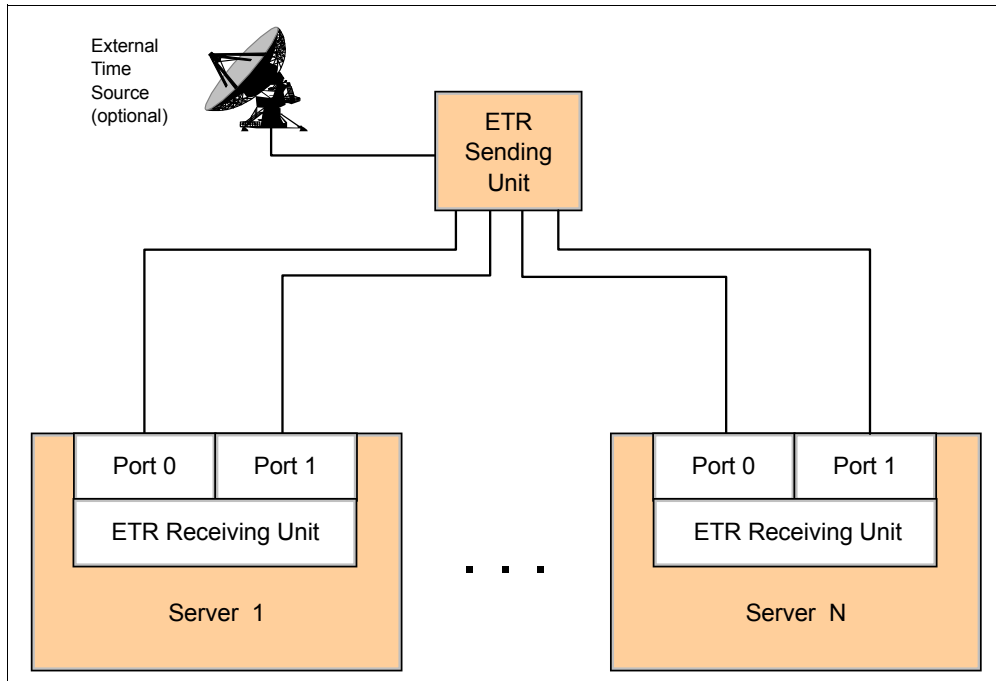


Figure 1-2 A Typical ETR network

A fault-tolerant configuration can be provided by coupling and synchronizing two sending units with each other so that each unit transmits consistent Sysplex Timer timing information. Figure 1-3 shows a typical fault-tolerant ETR network.

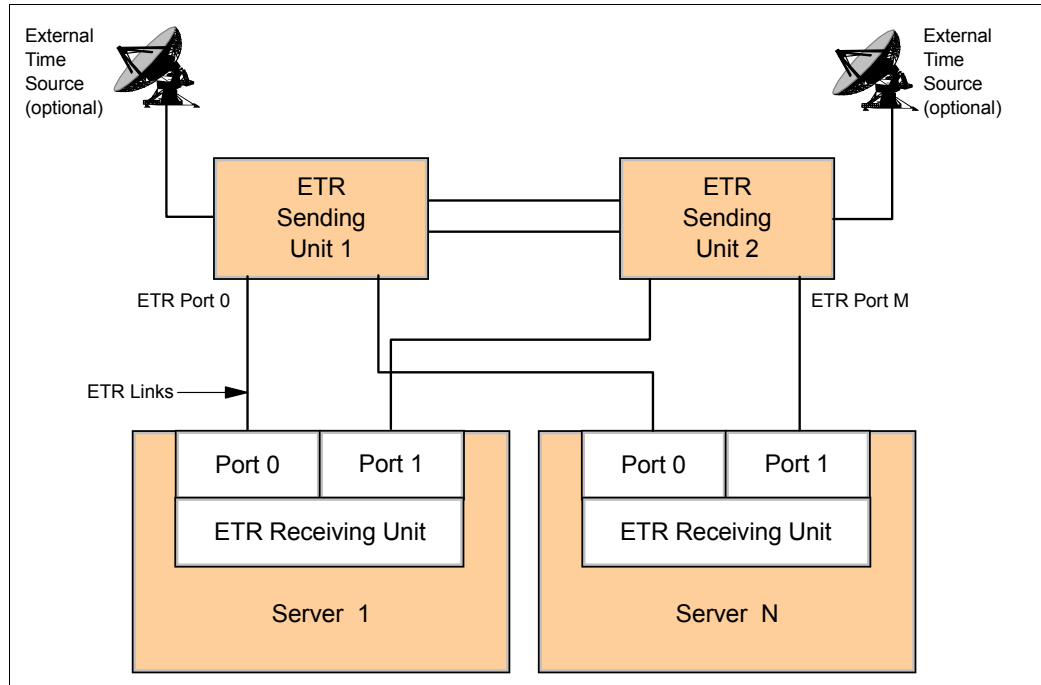


Figure 1-3 Fault-tolerant ETR network

The receiving unit on each server has two ports. Each port is normally connected to a different Sysplex Timer sending unit of a coupled pair in the same network. This fully duplicated structure minimizes the potential that a single failure can adversely impact the ETR network.

1.2 Overview of Server Time Protocol

The IBM 9037 Sysplex Timer is the ETR sending unit described in 1.1.3, “Overview of External Time Reference” on page 4. It provides the synchronization for the TOD clocks of multiple servers, and thereby allows events occurring on different servers to be properly sequenced in time. When multiple servers update the same database and database reconstruction is necessary, all updates are required to be time-stamped in the proper sequence.

Some of the key limitations of the Sysplex Timer are:

- ▶ Distance limitations of the Sysplex Timer constrain the distance separation of a multi-site Parallel Sysplex.

In information technology, two important objectives for survival are:

- Systems designed to provide continuous availability
- Near transparent disaster recovery (DR)

Systems designed to provide continuous availability combine the characteristics of high availability and continuous operations to deliver high levels of service around the clock.

To attain these objectives, solutions such as GDPS³ are based on geographical clusters (such as Parallel Sysplex) and remote data mirroring across two or more data centers. An increasing number of enterprises require that the geographical cluster or Parallel Sysplex be dispersed over distances of 100 km or more to mitigate the risk that a single disaster could impact multiple data centers.

Today, in order to disperse a Parallel Sysplex over a fiber distance of 100 km, it is required that the second Sysplex Timer in a Sysplex Timer Expanded Availability Configuration be located in an intermediate site between the data centers. It should not be more than 40 km (fiber distance) away from one of the primary data centers. This is because the two timers in this configuration cannot maintain synchronization when the distance between these timers exceeds 40 km. In many cases, finding an intermediate site is not a practical solution. This limits the effective dispersion of GDPS sites to only 40 km, so that one timer can be placed in each of the two sites, with no intermediate site.

- ▶ The Sysplex Timer is required to keep the TOD clocks of all participating servers synchronized with each other to within a small number of microseconds. This value is dictated by the fastest possible passing of data from one server to another through the coupling facility structure. With the processor and coupling link technology having improved over the years, this constraint has become steadily more rigorous, dropping from a few tens of microseconds to (now) on the order of 8–10 microseconds. So far the synchronization accuracy of the Sysplex Timer has been able to stay ahead of the requirements, and thus preserve the validity of the architecture and programming model, both within a single footprint and across servers in a Parallel Sysplex. Unfortunately, this situation is now nearing an end as server and coupling link technologies keep improving.

Notes: The Sysplex Timer Model 2 was withdrawn from marketing effective December 31, 2006. Enterprises that do not have a Sysplex Timer need a new methodology for time synchronization across multiple servers.

The service for the Sysplex Timer Model 2 will continue beyond the withdrawal from marketing.

An up-to-date time synchronization must provide the following functionality:

- ▶ Improved time synchronization (compared to Sysplex Timer) for System z10 and z9 servers in a sysplex or non-sysplex configuration.
- ▶ Scale with distance. Servers exchanging messages over fast, short links require more stringent synchronization than servers exchanging messages over long distances.
- ▶ Scale with server and coupling link technologies. For example, the solution is re-used with appropriate changes when coupling link technologies change in the future.
- ▶ Support a multi-site sysplex of at least 100 km without requiring an intermediate site, and should not preclude going to longer distances in the future.
- ▶ Provide at least all the functions provided by Sysplex Timer console functions.
- ▶ Allow concurrent migration from the ETR network.
- ▶ Allow coexistence with the ETR network.
- ▶ Allow implementation into an existing Network Time Protocol (NTP) network, which may or may not utilize a pulse per second (PPS) signal.

The new time synchronization infrastructure developed by IBM for the System z environment is called *Server Time Protocol*.

³ Geographically Dispersed Parallel Sysplex™

Server Time Protocol

Server Time Protocol (STP) is a server-wide facility that is implemented in the Licensed Internal Code (LIC) of the zEC12, z196, z114, z10 EC, z10 BC, z9 EC, z9 BC, z990, and z890 servers. STP presents a single view of time to Processor Resource/Systems Manager™ (PR/SM™) and is designed to provide the capability for multiple STP configured servers to maintain time synchronization with each other. It is the follow-on to the Sysplex Timer. STP is designed to allow events occurring in different servers to be properly sequenced in time.

STP is designed for servers that have been configured in a Parallel Sysplex or a basic sysplex (without a coupling facility), as well as servers that are not in a sysplex but that must be time synchronized.

STP is a message-based protocol in which timekeeping information is passed over data links between servers. The timekeeping information is transmitted over externally defined coupling links. Coupling links that can be used to transport STP messages are the Inter System Channel-3 (ISC-3) links configured in peer mode, Integrated Cluster Bus-3 (ICB-3) links, Integrated Cluster Bus-4 (ICB-4) links, and InfiniBand (IFB) links.

STP provides the following functionality:

- ▶ Allow clock synchronization for the zEC12, z196, z114, z10 EC, z10 BC, z9 EC, z9 BC, z990, and z890 servers and CFs, without requiring the Sysplex Timer.
- ▶ Support a multi-site timing network of up to 100 km (62 miles) over fiber optic cabling, allowing a Parallel Sysplex to span these distances with no intermediate site requirement.
- ▶ Potentially reduce the cross-site connectivity required for a multi-site Parallel Sysplex.
- ▶ Coexist with an ETR network.
- ▶ Allow use of an external time source to set the time to an international time standard, such as Coordinated Universal Time, as well as adjust to the time standard on a periodic basis.
- ▶ Allow setting of local time parameters, such as time zone and daylight saving time (DST).
- ▶ Allow automatic updates of daylight saving time.

Note: The z12EC is the last server to support connections to a Mixed CTN. However, it cannot be configured in the same CTN with any of the following servers: z9 EC, z9 BC, z990 or z890.

Software support for STP is available in all supported z/OS releases with applicable PTFs.

1.3 STP concepts and terminology

This section provides an overview of some of the new concepts and terminology related to STP that did not exist with the ETR network. Understanding this new vocabulary will assist you in successfully planning a timing network based on STP.

1.3.1 STP facility

STP provides the means by which the TOD clocks in various systems can be synchronized using messages transported over links. STP operates in conjunction with the TOD-clock steering facility, providing a new timing mode, timing states, external interrupts, and machine check conditions.

TOD-clock steering facility

The TOD-clock steering facility provides a means to change the apparent stepping rate of the TOD clock without changing the physical hardware oscillator that steps the physical clock. This is accomplished by means of a *TOD-offset register* that is added to the physical clock to produce a logical-TOD-clock value.

TOD-clock steering permits the timing facility control program to adjust the apparent stepping rate of the TOD clock. In normal operation, the update is performed frequently so that the effect, as observed by the program, is indistinguishable from a uniform stepping rate.

The total steering rate is made up of two components:

- ▶ Fine-steering rate

The fine-steering rate is used to correct the inaccuracy in the local oscillator, which is stable over a relatively long period of time. The value normally is less than the specified tolerance of the local oscillator. The change occurs infrequently (on the order of once a day to once a week) and is small.

- ▶ Gross-steering rate

The gross-steering rate is used as a dynamic correction for all other effects, the most predominant being to synchronize time with an external time source or with other clocks in the timing network. The value normally changes frequently (on the order of once per second to once per minute).

1.3.2 TOD clock synchronization

This section provides definitions of the timing mode, timing state, and STP clock source state when the STP feature is installed.

Timing mode

The timing mode specifies the method by which the TOD clock is maintained for purposes of synchronization within a timing network. A TOD clock operates in one of the following timing modes:

- ▶ Uninitialized

When the server is in uninitialized timing mode, the TOD clock has not been initialized. The server is not part of a synchronized timing network.

- ▶ Local

When the server is in local timing mode, the TOD clock is initialized to a local time and is stepped at the rate of the local hardware oscillator. The server is not part of a synchronized timing network.

- ▶ ETR

When the server is in ETR timing mode, the TOD clock is initialized to the Sysplex Timer and is being stepped by stepping signals from the Sysplex Timer. To be in ETR timing mode, the server must be part of an ETR network.

- ▶ STP

When the server is in STP timing mode, the TOD clock is initialized to Coordinated Server Time (CST) and is stepped at the rate of the local hardware oscillator. In STP timing mode, the TOD clock is steered so as to maintain, or attain, synchronization with CST. To be in STP-timing mode, the server must be part of an STP network. Refer to 1.3.4, “Coordinated Timing Network” on page 11, for an explanation of CST.

Timing states

The *timing state* indicates the synchronization state of the TOD clock with respect to the timing network reference time. The timing states are:

► Synchronized

When a server is in the synchronized timing state, the TOD clock is in synchronization with the timing-network reference time as defined here:

- If the server is in ETR timing mode, the server is synchronized with the Sysplex Timer.
- If the server is in STP timing mode, the server is synchronized with CST.
- A server that is in the local timing or uninitialized timing mode is never in the synchronized state.

► Unsynchronized

When a server is in the unsynchronized timing state, the TOD clock is not in synchronization with the timing network reference time as defined here:

- If the server is in ETR timing mode, the server has lost synchronization with the Sysplex Timer.
- If the server is in STP timing mode, the server has lost or has not been able to attain synchronization with Coordinated Server Time. The server is out of synchronization with CST when the TOD clock differs from CST by an amount that exceeds a model-dependent synchronization threshold value.

► Stopped

When a server is in the stopped state, the TOD clock is either in the stopped state or TOD-clock recovery is in progress. When the TOD-clock recovery is completed, the TOD clock enters either the synchronized or unsynchronized state.

STP clock source state

The STP clock source state indicates whether a usable STP clock source is available. The STP clock source is used to determine the Coordinated Server Time required to be able to synchronize the TOD clock.

► Not usable

The *not usable* STP clock source state indicates that a usable STP clock source is not available to STP. When a usable STP clock source is not available, CST cannot be determined.

► Usable

The *usable* STP clock source state indicates that a usable STP-clock source is available to STP. When a usable STP-clock source is available, CST has been determined and can be used to synchronize the TOD clock to the STP network.

1.3.3 STP servers

In this section we discuss STP servers.

STP capable

An STP-capable server is a server that has all required LIC to support STP installed. IBM System zEC12, z196, z114, z10 EC, z10 BC, z9 EC, z9 BC, and zSeries® 990 and 890 with the required LIC installed are STP capable.

STP enabled

An STP-enabled server is an STP-capable server that has the STP function enabled. Even after LIC to support STP is installed on a server, the STP function cannot be used until it is enabled.

STP configured

An STP-configured server is a server that has been configured to be part of a CTN by assigning it a CTN ID. When the STP network ID field is not specified, the server is not configured to be part of a CTN, and is therefore not an STP-configured server. CTN and CTN ID are discussed in the next section.

1.3.4 Coordinated Timing Network

A Coordinated Timing Network (CTN) contains a collection of servers that are time synchronized. They are time synchronized to a time value called Coordinated Server Time (CST). The CST represents the time for the entire network of servers.

The servers that make up a CTN are all configured with a common identifier, referred to as a CTN ID. Only servers with the same CTN ID are allowed to become members of the CTN. Servers with different CTN IDs can be members of different Coordinated Timing Networks. All servers in a CTN maintain an identical set of time control parameters that are used to coordinate the TOD clocks.

A CTN can be configured as either:

- ▶ STP-only CTN

An STP-only CTN (described later in 1.4.2, “STP-only CTN” on page 18) is a timing network in which all servers are configured to be in STP timing mode. It can only be configured with STP-capable servers, and none of the servers can be in ETR timing mode.

- ▶ Mixed CTN

A Mixed CTN (described later in 1.4.1, “Mixed CTN” on page 14) allows the coexistence of servers and CFs synchronized in an ETR network, with servers and CFs that are synchronized with CST. The Sysplex Timer provides the timekeeping information in a Mixed CTN.

CTN ID

The CTN ID is an identifier that is used to indicate whether the server has been configured to be part of a CTN and, if so configured, identifies the CTN. The CTN ID is made up of two fields:

- ▶ A field that defines the STP network ID
- ▶ A field that defines the ETR network ID

These fields are of the format STP network ID and ETR network ID, as illustrated on Figure 1-4 (the hyphen is used to identify each field. it is not part of the CTN ID).




Figure 1-4 Coordinated Timing Network ID (CTN ID)

The Hardware Management Console (HMC) is used to enter the CTN ID. See Figure 1-5.

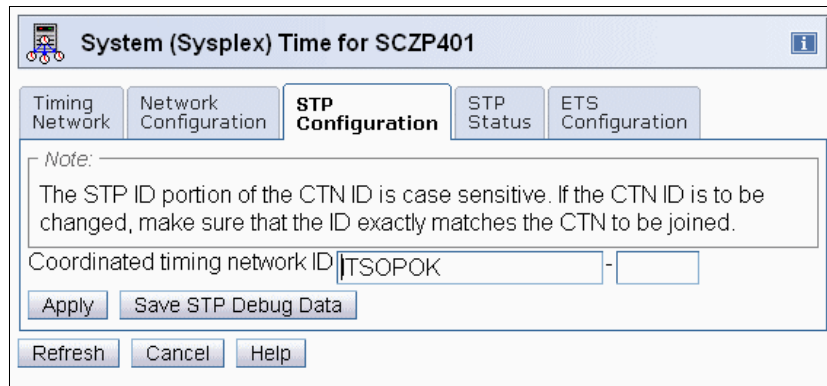


Figure 1-5 CTN ID

When both fields are specified, the CTN is referred to as a Mixed CTN. When the STP network ID field is specified but the ETR network ID field is not specified (set to null or blank), the CTN is referred to as an STP-only CTN.

For servers that do have the External Time Reference feature for Sysplex Timer attachment installed (z10 EC, z10 BC, z9 EC, z9 BC, z990 and z890), the ETR network ID is configured in the ETR Configuration tab. The ETR Configuration tab and the ETR Status tab are not available for servers zEC12, z196 and z114 (as seen in Figure 1-5) since the ETR feature is not installed, nor available in these servers.

1.3.5 STP stratum

The term *stratum* is found throughout this document and within the Hardware Management Console panels.

The Sysplex Timer distributes time to multiple servers in a star pattern. That is, the Sysplex Timer is the star, and its time signals emanate out from it to all attached servers. The signals from the Sysplex Timer are used to increment or step the TOD clocks in the attached server.

Unlike the Sysplex Timer, STP distributes time messages in layers, or stratum. The top level (stratum 1) distributes time messages to the layer immediately below it (stratum 2). Stratum 2 in turn distributes time messages to stratum 3. More layers are conceivably possible, but the current STP implementation is limited to three layers. Figure 1-6 illustrates the stratum concept.

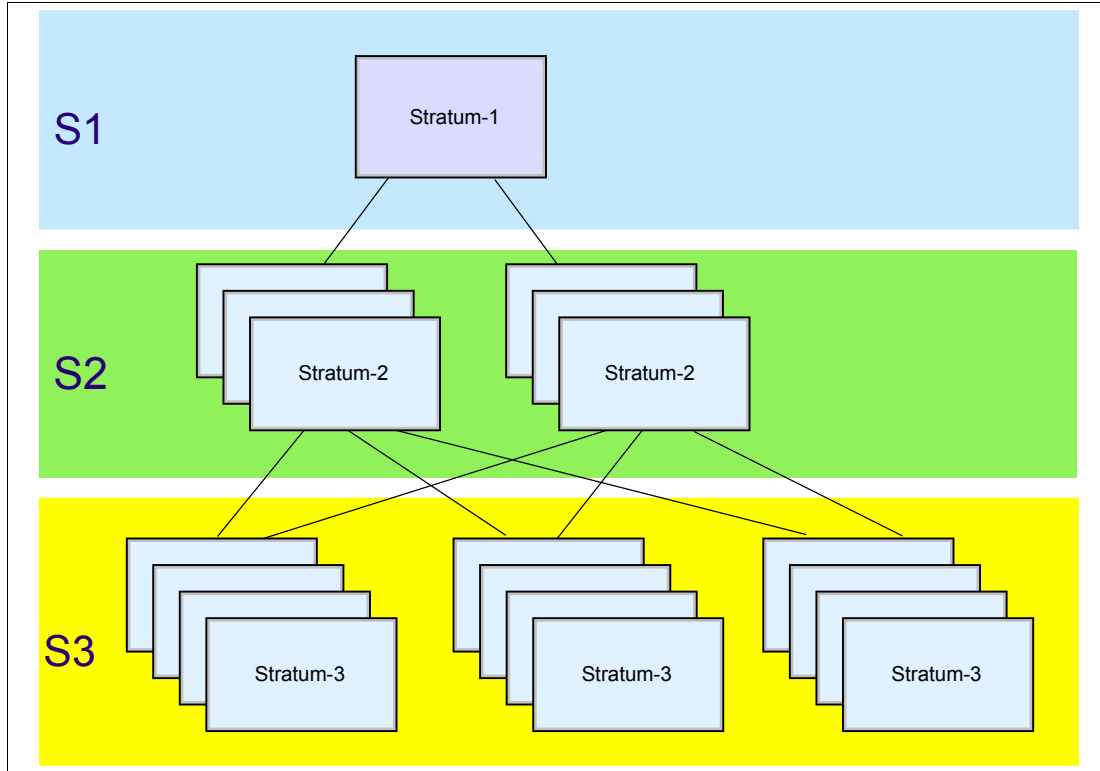


Figure 1-6 STP stratum levels

STP Stratum levels

In a timing network based on STP, stratum are used as a means to define the hierarchy of a server in the timing network. A stratum 1 server is the highest level in the hierarchy in the STP network.

The stratum 1 level is determined indirectly in one of the following ways:

- ▶ In a Mixed CTN, any STP-configured server synchronized to the Sysplex Timer is a Stratum 1 server. A Mixed CTN is allowed to have multiple Stratum 1 servers. Refer to 1.4, “Timing network types” on page 14 for more details.
- ▶ In an STP-only CTN, the roles of Preferred and Current Time Servers are assigned, and the Current Time Server becomes the active stratum 1 server. Details regarding Preferred and Current Time Servers are in 1.3.6, “Server roles in an STP-only CTN” on page 14.

Stratum 2 and stratum 3 levels are determined by how many stratum levels they are away from stratum 1. A server that uses STP messages to synchronize to a stratum 1 server is referred to as a stratum 2 server. Similarly, a server that uses STP messages to synchronize to a stratum 2 server is referred to as a stratum 3 server.

Note: Throughout the book you will also find STP stratum 0 (S0). This refers to the unsynchronized state of a server (CEC).

1.3.6 Server roles in an STP-only CTN

Unlike a Mixed CTN, which can have multiple stratum 1 servers, an STP-only CTN can have only one active stratum 1 server. The following definitions explain the roles that must be assigned (through the HMC interface) for certain servers in an STP-only CTN.

Preferred Time Server (PTS)

Using the STP panels provided at the HMC, a server must be assigned that has preference to be the stratum 1 server of an STP-only CTN. This is the Preferred Time Server. This server should have connectivity to the Backup Time Server and the Arbiter, as well as to all servers that are planned to be stratum 2 servers. The connectivity can be either ISC-3 links in peer mode, ICB-3 links, ICB-4 links, or IFB links.

Backup Time Server (BTS)

Although optional, we strongly recommend also assigning a Backup Time Server whose role it is to take over as the stratum 1 server should the PTS fail. The Backup Time Server is a stratum 2 server that should have connectivity to the Preferred Time Server and the Arbiter, as well as to all other stratum 2 servers that are connected to the Preferred Time Server.

Current Time Server (CTS)

The Current Time Server is the active stratum 1 server in an STP-only CTN.

Note: There can only be one active stratum 1 server in an STP-only CTN, and only the Preferred Time Server or the Backup Time Server can be assigned to be the active Stratum 1 server.

The Current Time Server can be assigned (only) to either the Preferred Time Server or the Backup Time Server. We recommend that the Current Time Server be assigned to the Preferred Time Server when the configuration is being initialized. Subsequently, if there is a need to reassign the roles, the Current Time Server can be concurrently assigned to the Backup Time Server. This action can be part of a planned reconfiguration of the Preferred Time Server as long as the planned action is not *disruptive*.

Arbiter

Optionally, a server might be assigned the role of *Arbiter*. The Arbiter provides additional means for the Backup Time Server to determine whether it should take over as the Current Time Server in the case of unplanned events that affect the CTN. The Arbiter is a stratum 2 server that must have connectivity to both the PTS and the BTS.

1.4 Timing network types

- ▶ STP-only CTN, which does not require a Sysplex Timer

Refer to 1.1.3, “Overview of External Time Reference” on page 4, for a brief description of an ETR network.

1.4.1 Mixed CTN

As mentioned earlier, it is essential that STP supports a concurrent migration from an existing ETR network to an STP-based one. STP is implemented in the LIC of the zEC12, z196, z114, z10 EC, z10 BC, z9 EC, z9 BC, z990, and z890 servers.

It is also crucial that STP allow servers and coupling facilities that are synchronized with the Sysplex Timer (ETR network) and servers and CFs that are synchronized with Coordinated Server Time to be synchronized with each other in the same network.

In order to meet these two important requirements, the STP design allows the formation of a Mixed Coordinated Timing Network (Mixed CTN). The Sysplex Timer provides timekeeping information in a Mixed CTN.

The minimum requirements to configure a Mixed CTN are:

- ▶ There must be at least one STP-enabled server. To reiterate, an STP-capable server that has the STP function enabled is referred to as an STP-enabled server.

Important: In order to avoid a single point of failure in a Mixed CTN, the recommendation is to have a minimum of two STP-configured servers stepping to timing signals received from the Sysplex Timers before configuring any server using STP messages for synchronization.

- ▶ The STP-enabled servers must be attached to a Sysplex Timer or a pair of Sysplex Timers configured in an Expanded Availability configuration. Since the Sysplex Timer provides the timekeeping information in a Mixed CTN, at least one STP-enabled server in the Mixed CTN must be stepping to timing signals provided by the Sysplex Timer before additional servers or CFs can be STP configured.
- ▶ The STP-enabled servers must be configured with the same CTN ID, with the format [STP network ID] - [ETR network ID]. The STP network ID can be from 1 to 8 characters long. Valid characters are A–Z, a–z, 0–9, '-', and '_'. The ETR network ID is the same as the ETR network ID of the Sysplex Timers to which the servers are attached. After the CTN ID is assigned, the servers are now STP-configured servers.

Note: A CTN is configured only after the servers become STP-configured.

Figure 1-7 shows an example of a Mixed CTN consisting of two STP-configured servers, a z10 EC, and a z9 BC. Both servers are synchronized by one of the attached Sysplex Timers, whose ETR network ID is 31.

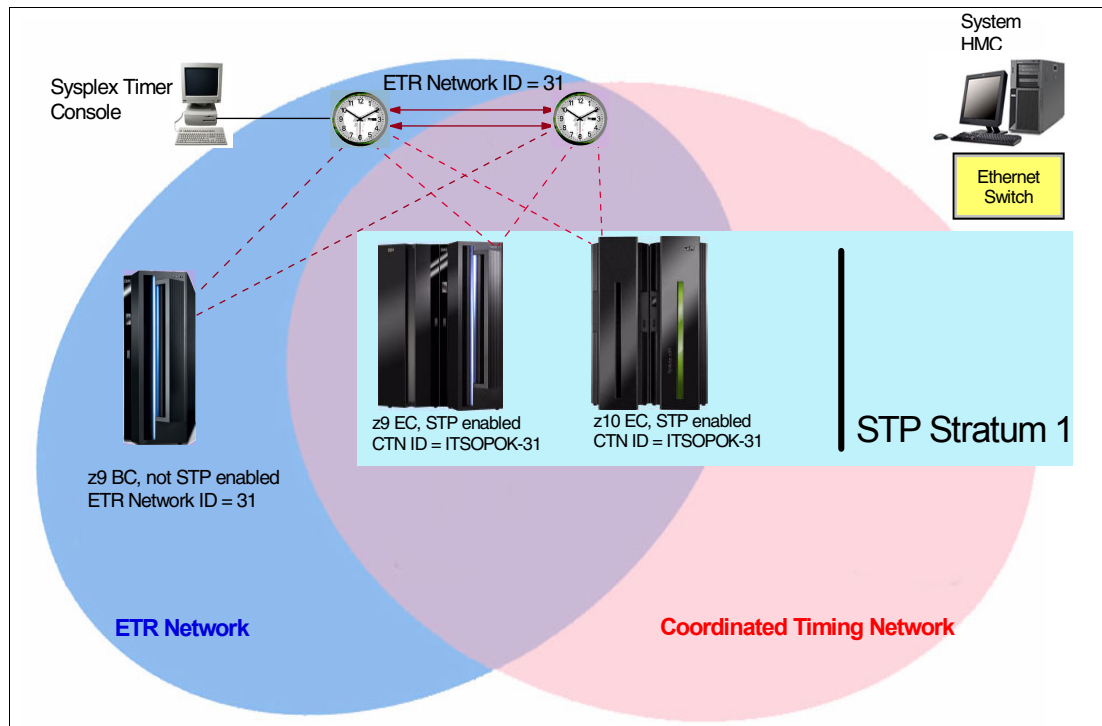


Figure 1-7 Mixed CTN, stratum 1 servers only

In this example, the z10 EC and the z9 EC have been assigned the same CTN ID = [ITSOPK] - [31] and are therefore in a Mixed CTN, with the Sysplex Timer providing the timekeeping information. The servers are in ETR timing mode. If coupling links are configured between the z10 EC and the z9 EC, then STP messages are exchanged by the two servers. As described in “STP Stratum levels” on page 13, the z10 EC and the z9 EC become stratum 1 servers in the Mixed CTN. The z9 BC has not been STP-enabled, however, it is attached to the same pair of Sysplex Timers (ETR Network ID =31). The z9 BC is therefore in an ETR network and can coexist with the servers that are in the Mixed CTN.

Each Stratum 1 server receives time-control parameters from the attached Sysplex Timer (or timers) and propagates the information to stratum 2 servers, which in turn propagate the information to stratum 3 servers. Time-control parameters for a configuration include:

- ▶ Total-time offset (the combination of time-zone offset and DST offset)
- ▶ Leap seconds offset

In a Mixed CTN, the Sysplex Timer console provides the user interface for all time-related functions, such as time initialization, time adjustment, and so on. This is not different from the operations in an ETR-only network. However, the HMC must be used to initialize or modify the CTN ID. In a Mixed CTN if there is a requirement for time accuracy, the external time source (ETS) function is provided by attaching the ETS device to the Sysplex Timer.

Servers configured in STP timing mode

Once the minimum requirements described in the previous section have been met, additional servers can be added to the Mixed CTN without these servers requiring connectivity to the Sysplex Timers. Servers in the CTN that are synchronized using STP messages, and not stepping to the Sysplex Timer, are stratum 2 or stratum 3 servers.

If you have more than two stratum 1 servers in a Mixed CTN configuration (all stratum 1 are directly connected to the Sysplex timers), to migrate these servers to STP timing mode, disable their ETR ports. These servers maintain time synchronization using STP protocol via coupling link connectivity to the remaining stratum 1 servers.

The servers that are planned to be configured in STP timing mode must meet the following requirements:

- ▶ They must be configured with the same CTN ID as the servers stepping to the Sysplex Timer. These are the stratum 1 servers.
- ▶ They must have coupling link connectivity (ISC-3 Peer, ICB-3, ICB-4, or IFB) to the servers to which they are planned to be synchronized.

In Figure 1-8, an additional server is configured with the same CTN ID as the Mixed CTN shown in Figure 1-7 on page 16.

The STP-enabled z10 BC is assigned a CTN ID = [ITSOPOK] - [31]. Also, ISC-3 Peer links are configured between the z10 BC and the z10 EC as well as between the z10 BC and the z9 EC. STP algorithms determine whether the z10 BC will synchronize to the z10 EC or the z9 EC. Redundancy of stratum 1 servers is provided to allow the z10 BC server to stay within the Mixed CTN in the event that either the z10 EC or z9 EC has a planned or unplanned outage.

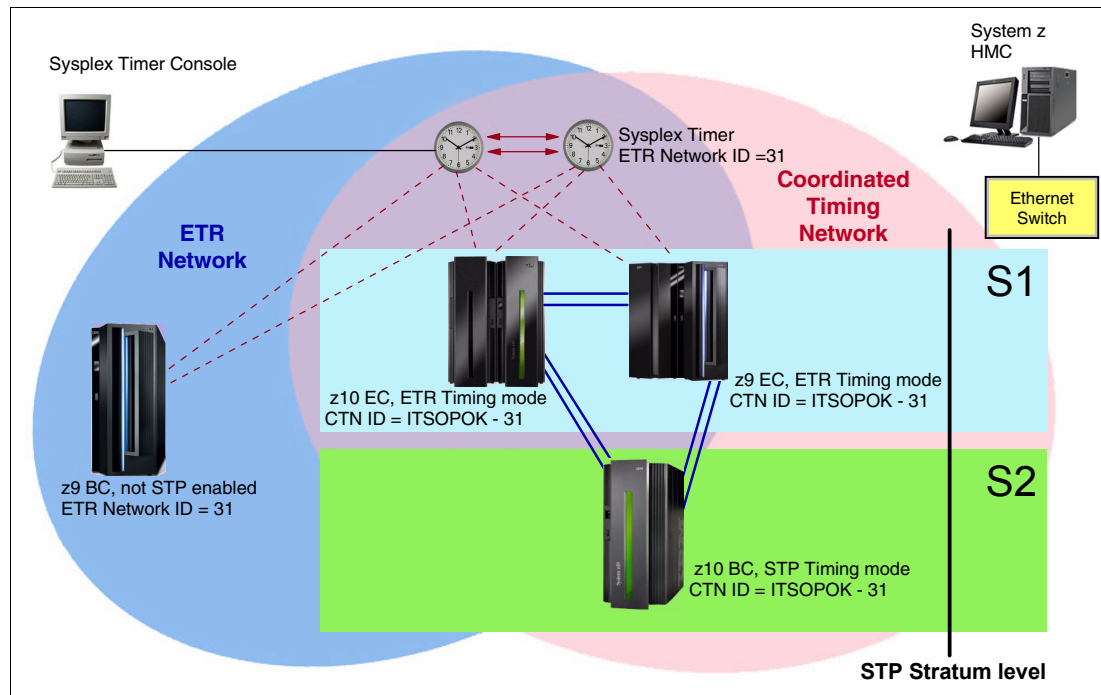


Figure 1-8 Mixed CTN, stratum 1 and stratum 2 servers

When the Mixed CTN configuration is active, the Sysplex Timer provides the timekeeping information. The z10 EC and z9 EC are synchronized to the Sysplex Timers via the ETR links and are in ETR timing mode. The z10 BC is synchronized to either the z10 EC or the z9 EC by exchanging STP messages over the coupling links, and is in STP timing mode.

Note: zEC12 servers are not supported in the same CTN with z9 or earlier servers.

Parallel Sysplex configuration in a Mixed CTN

A Parallel Sysplex configuration requires that all servers in the sysplex are time coordinated. Figure 1-9 shows a Parallel Sysplex comprising z/OS images P1, P2, P3, P4, and P5, along with two Internal coupling facilities (ICFs).

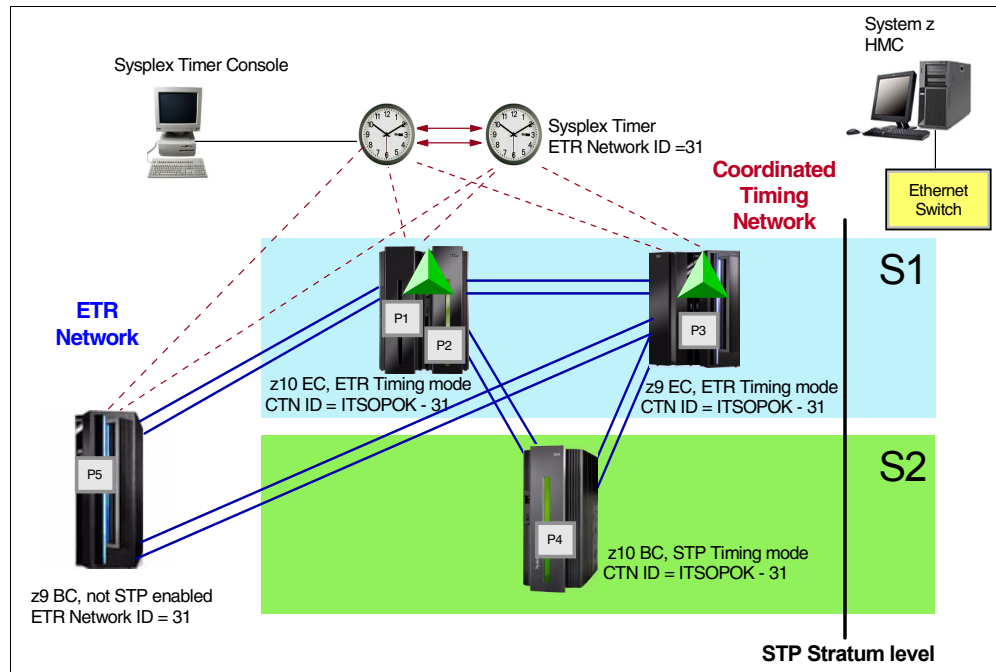


Figure 1-9 Mixed CTN Parallel Sysplex example

Note: zEC12 servers are not supported in the same CTN with z9 or earlier servers.

The coupling links used for CF messaging between servers, as shown in Figure 1-9, can also be used for STP message exchanges. It is possible to configure the CTN so that additional coupling links are not required to specifically transport STP messages. The same coupling links are also used for STP message exchanges. The z/OS image P5, located in the z9 BC server that is not STP-enabled, maintains time synchronization from the Sysplex Timer.

Note: Since STP configuration allows stratum 3 servers in a CTN (even though Figure 1-8 and Figure 1-9 show only two stratum levels), the CTN can be further extended by synchronizing a stratum 3 to the stratum 2 server.

1.4.2 STP-only CTN

An STP-only CTN is a timing network in which all servers are configured to be in STP timing mode. None of the servers in an STP-only CTN can be in ETR timing mode. Therefore, an STP-only CTN can only be configured with STP-enabled servers, and the Sysplex Timer does not provide stepping signals to any servers in this network. As a minimum, before an STP-only CTN configuration can become active and servers can exchange timekeeping messages, the following conditions must exist:

- There must be at least one STP-enabled server in order to configure an STP-only CTN. To reiterate, an STP-capable server that has STP function enabled is referred to as an STP-enabled server.

- ▶ The STP-enabled servers must be configured with the same CTN ID, with the format [STP network ID] - [ETR network ID]. The STP network ID can be from 1 to 8 characters long. The valid characters are A–Z, a–z, 0–9, '-', and '_'. In an STP-only CTN, the ETR network ID must *not* be specified. For example, a valid CTN ID could be [STPCONFG] - []. After the CTN ID is assigned, the servers are now considered to be STP configured.

TIP: The STP network ID is case-sensitive.

- ▶ Assign the Preferred Time Server and the Current Time Server. The server that is assigned the role of the Current Time Server becomes the active stratum 1 when an STP-only CTN configuration is completed. In most cases, the Preferred Time Server is assigned to also be the Current Time Server.
- ▶ Optionally, a Backup Time Server can be assigned. If a Backup Time Server is assigned, an Arbiter can be assigned as well if there are three or more servers in the STP CTN.

The servers that are planned to be synchronized to either the Preferred Time Servers or the Backup Time Servers need coupling link connectivity (ISC-3 Peer, ICB-3, ICB-4, or IFB). All servers in an STP-only CTN are in STP timing mode.

Note: ISC links are not supported on zEC12 servers.

In an STP-only CTN, the HMC provides the user interface for all time-related functions, such as time initialization, time adjustment, offset adjustment, and so on. The HMC must also be used to initialize or modify the CTN ID. For time accuracy, the external time source (ETS) function is also provided. In certain cases, the HMC is also used to provide Console Assisted Recovery (CAR).

Preferred Time Server only

Figure 1-10 shows an STP-only CTN where the zEC12 is assigned as the Preferred Time Server. A Backup Time Server is not assigned, and therefore an Arbiter cannot be assigned either. Since only the Preferred Time Server or the Backup Time Server can be the Current Time Server, in this example the Preferred Time Server must be the Current Time Server. Redundant coupling link connectivity is provided as shown from each of the other servers (z196 and z114) to the zEC12.

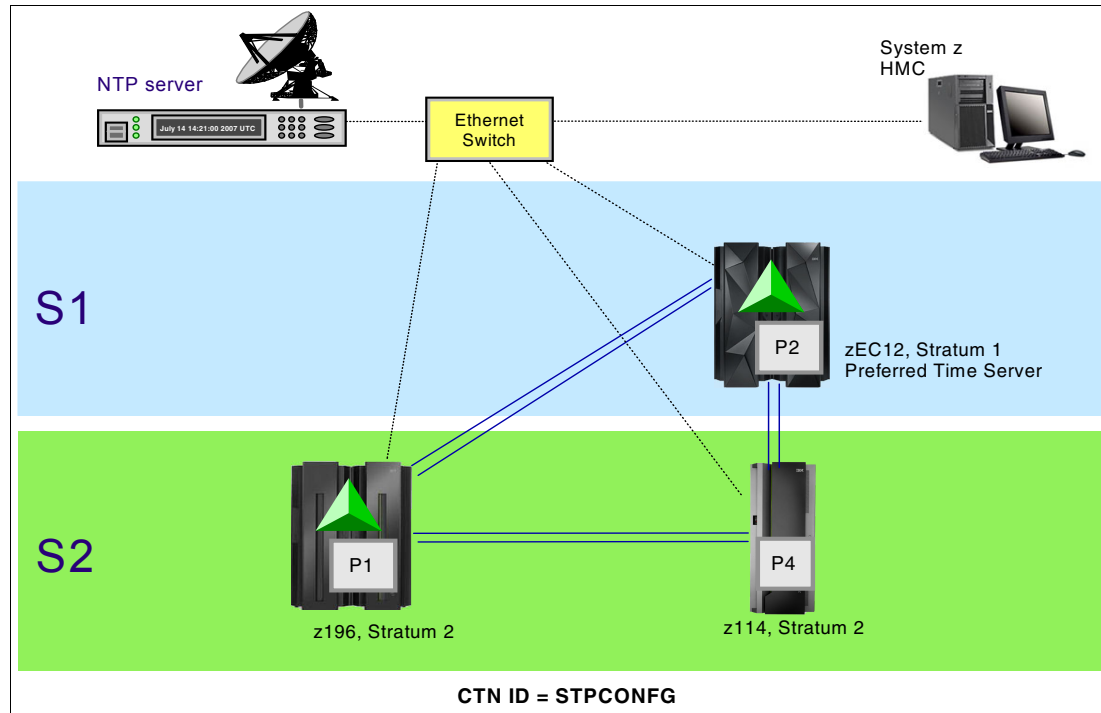


Figure 1-10 STP-only CTN (Preferred Time Server only)

When an STP-only CTN configuration becomes active:

- ▶ The zEC12 becomes the stratum 1 server.
- ▶ The z196 and z114 become stratum 2 servers.

Note: In this example, without a Backup Time Server assigned, the loss of the Preferred Time Server results in the loss of synchronized time in the CTN until the Preferred Time Server is restored or a new configuration is defined.

In an STP-only CTN, time-control parameters are provided via the Hardware Management Console or Support Element panels. The time-control parameters must be sent to the Current Time Server in the CTN. The Current Time Server receiving the parameters propagates the parameters to each server in the CTN. Time-control parameters for a configuration include:

- ▶ Leap seconds offset
- ▶ Time zone offset
- ▶ Daylight saving time offset
- ▶ Scheduled changes to any of these offsets

Preferred and Backup Time Servers

Figure 1-11 shows an example of an STP-only CTN where the zEC12 is assigned as the Preferred Time Server, the z114 is assigned as the Backup Time Server, and an Arbiter is not assigned. Either the Preferred Time Server or the Backup Time Server is assigned as the Current Time Server. In this example, the Preferred Time Server is assigned as the Current Time Server (CTS).

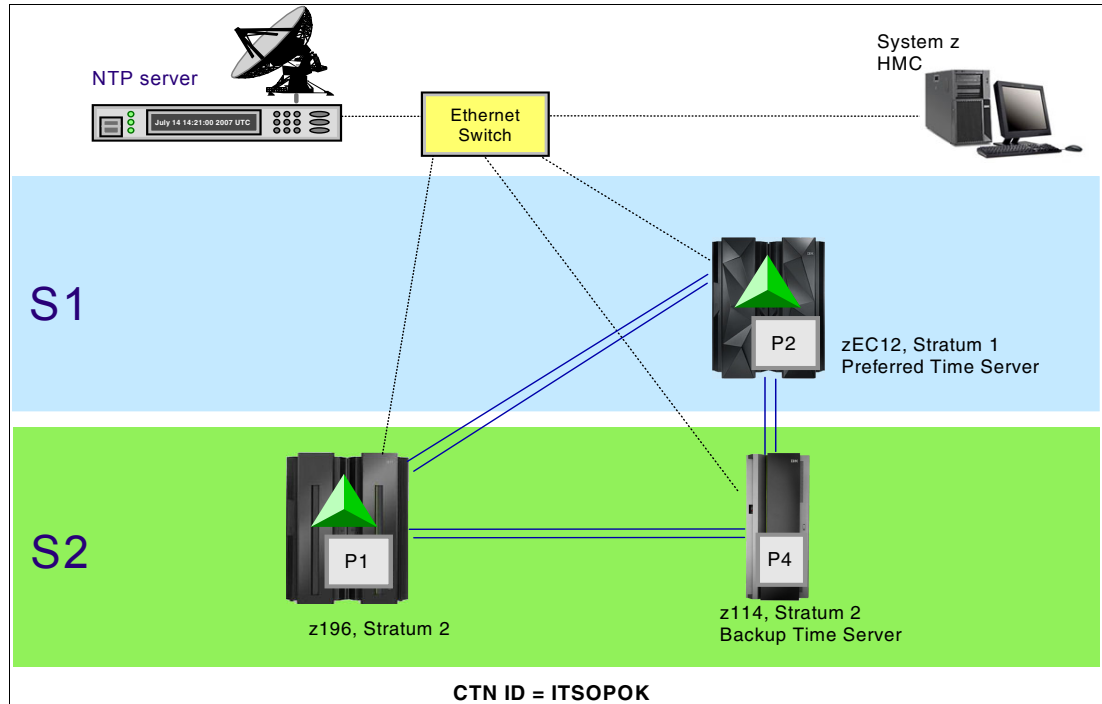


Figure 1-11 STP-only CTN (Preferred Time Server and Backup Time Servers)

In this example, redundant coupling links should be provided from the z196 to not only the zEC12, but also to the z114, which has been assigned as the Backup Time Server. By providing this connectivity, the z196 can maintain synchronization when the PTS (zEC12) fails and the BTS (z114) takes over as the CTS.

Preferred, Backup, and Arbiter servers

Figure 1-12 shows an example of an STP-only CTN where the zEC12 is assigned as the Preferred Time Server, the z114 is assigned as the Backup Time Server, and the z196 is assigned as an Arbiter. Either the Preferred Time Server or the Backup Time Server can be assigned as the Current Time Server. In this example, the Preferred Time Server is assigned as the Current Time Server. The Arbiter provides additional means for the Backup Time Server to determine whether it should take over as the Current Time Server in the event of unplanned exception conditions.

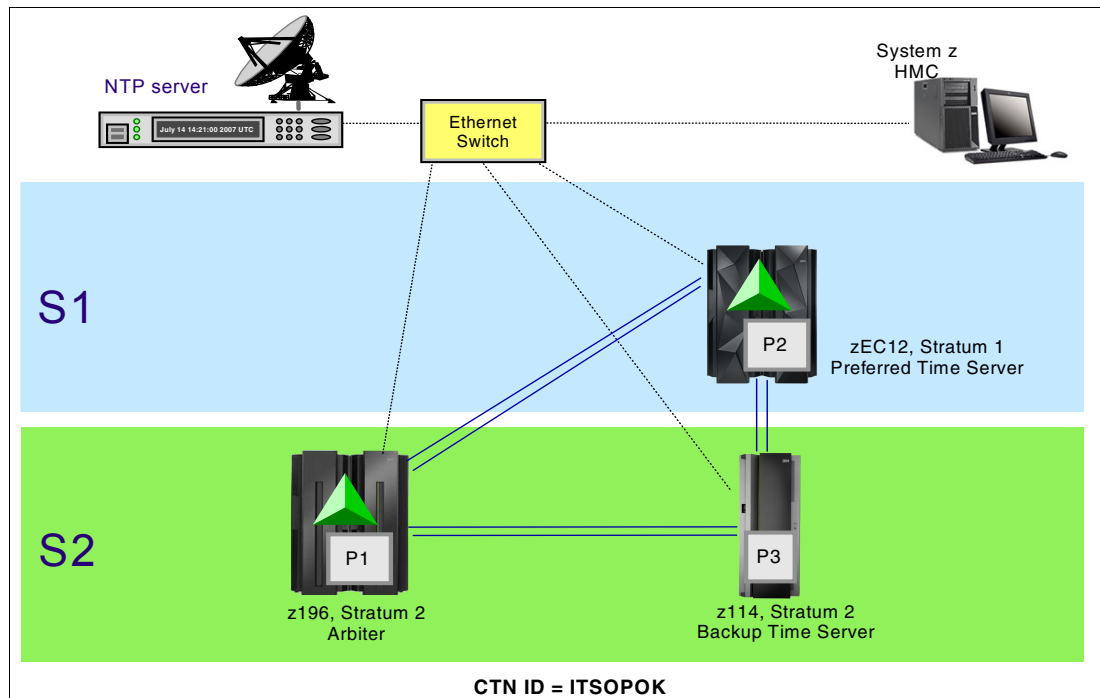


Figure 1-12 STP-only CTN (Preferred, Backup, Arbiter)

Figure 1-12 also shows three z/OS systems configured in a Parallel Sysplex: P1, P2, and P3.

Since the two CFs reside on the zEC12 and z196, the coupling links provide connectivity for both the CF and STP message exchanges.

1.5 Coupling links

STP is a message-based protocol in which time-keeping information is passed over data (coupling) links between servers. The time-keeping information provides the means for STP to determine an estimate of the time-of-day clock for the CTN, referred to as CST. The time-keeping information is transmitted over externally defined coupling links. Coupling links provide the transmission characteristics required for STP synchronization requirements. Coupling links that can be used to transport STP messages are the InterSystem Channel-3 (ISC-3) links in Peer mode, the Integrated Cluster Bus (ICB-4) or (ICB-3) links, or InfiniBand (IFB) links.

The advantages of using coupling links to exchange STP messages are:

- By using the same links to exchange timekeeping information and CF messages in a Parallel Sysplex, STP can meet the requirement stated in 1.2, “Overview of Server Time Protocol” on page 6, that STP must scale with distance. Servers exchanging messages

over ICB-3, ICB-4, and IFB links can now meet more stringent synchronization requirements than servers exchanging messages over long ISC-3 links (distances up to 100 km), where the synchronization requirements are less stringent.

- ▶ Coupling links also provide the connectivity needed in a Parallel Sysplex. Therefore, there is a potential benefit of minimizing the number of cross-site links required in a multi-site Parallel Sysplex.

The following types of STP messages are carried over the coupling links:

- ▶ Time parameter messages:
Time parameter messages are used to exchange the timekeeping information needed to determine Coordinated Server Time at each server. The time parameter messages are used to exchange timestamps, time-keeping information, and CTN-parameter information between two directly attached servers. The information in the message response is used by the message originator to calculate the round-trip delay, offset, and dispersion values that are used by STP clock filtering and selection algorithms to select a clock source.
- ▶ Control messages:
Control messages are used to set and modify various Coordinated Timing Network parameters required by servers in the CTN. They are used to request CTN parameter updates, to establish and remove STP paths, and to read configuration information from attached servers.

1.5.1 Link redundancy

We recommend that at least two coupling links are configured for communication between any two servers that are intended to exchange STP messages. This prevents the loss of one link causing the loss of STP communication between the servers. There is no architectural limit to the maximum number of links that can be defined. The limit is based on what is supported by each server in the configuration.

The maximum number of attached servers supported by any STP-configured server in a CTN is equal to the maximum number of coupling links supported by the servers in the configuration. Not considering redundancy recommendations, this value is equal to the maximum number of combined ISC-3, ICB-3, ICB-4, and IFB links. Up to 128 combined coupling links are supported on the zEC12, z196, and z114, while up to 64 combined coupling links are supported on the z10 EC, z10 BC, z9 EC, z9 BC, z990, and z890.

1.5.2 Coupling link requirements in a non-sysplex configuration

Time coordination is required by other applications besides the sysplex and Parallel Sysplex configurations. One example of time coordination requirements when a Parallel Sysplex configuration does not necessarily exist is the remote copy technology referred to as *z/OS Global Mirror* (previously called Extended Remote Copy (XRC)). *z/OS Global Mirror* is a combined hardware and software asynchronous remote copy solution. The application I/O is signaled complete when the data update to the primary storage is completed.

Subsequently, a DFSMSdfp component called *System Data Mover* (SDM) asynchronously offloads data from the primary storage subsystem's cache and updates the secondary disk volumes. Typically, the production systems are in site 1 and the data is being mirrored to a Disaster Recovery site (site 2). The production systems located in site 1 can be a single system, multiple systems sharing disk, or a base or Parallel Sysplex cluster. For more information about IBM DFSMSdfp, see:

<http://www-03.ibm.com/systems/storage/software/sms/dfp/index.html>

Data consistency across all primary and secondary volumes spread across any number of storage subsystems is essential for providing data integrity and the ability to do a normal database restart in the event of a disaster. Data consistency in a XRC environment is provided by the Consistency Group (CG) processing performed by the SDM. The CG contains records that have their order of update preserved across multiple Logical Control Units within a storage subsystem and across multiple storage subsystems. CG processing is possible only because each update on the primary disk subsystem has been time-stamped. If multiple systems on different servers are updating the data, time coordination is required across the different servers in each site. This requirement can be met either by the Sysplex Timer or STP. Time coordination is not required across the two sites that have the primary and secondary disk subsystems.

Figure 1-13 shows a GDPS/XRC configuration. More information about XRC and GDPS/XRC, a solution that automates the recovery of production systems when a disaster takes place, can be found at:

<http://www-03.ibm.com/systems/z/advantages/gdps/index.html>

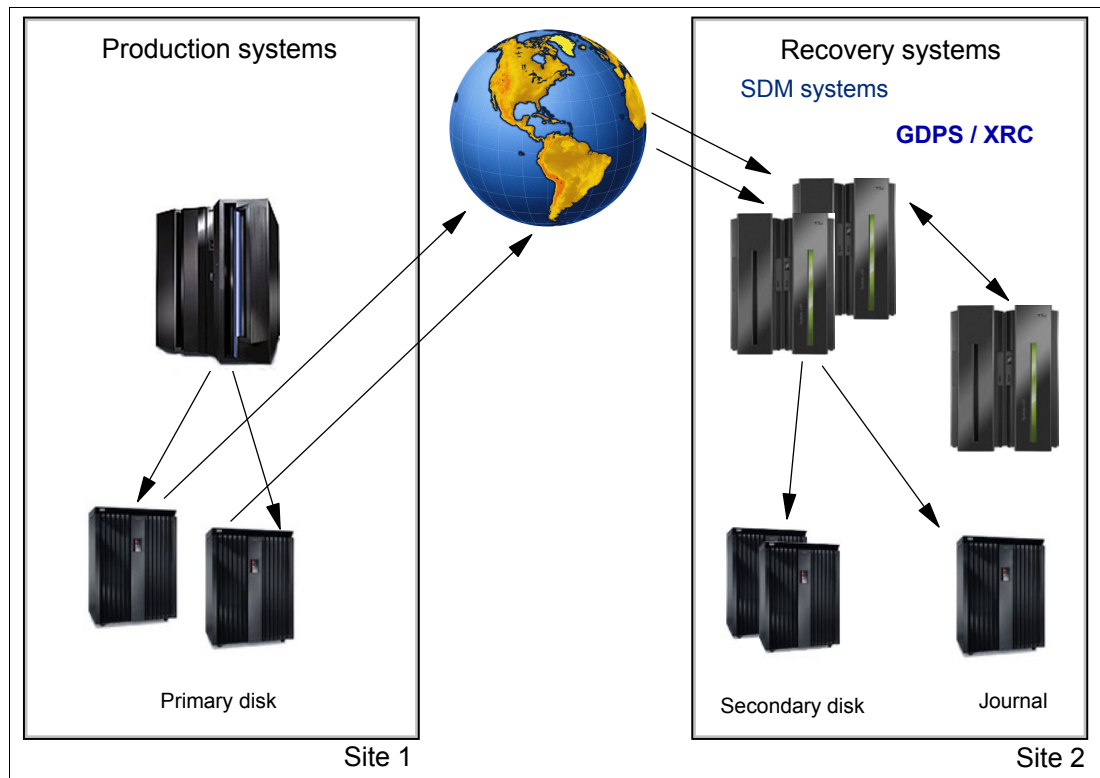


Figure 1-13 GDPS/XRC configuration

For a server that is not part of a Parallel Sysplex but that must be in the same CTN, additional coupling links must be configured in order for the server to be configured in the CTN. Hardware configuration data (HCD) has been enhanced to permit special *timing-only* links and control units to be defined. The timing-only links can be of type CFP for ISC-3 in peer mode, or CBP for ICB-3/ICB-4, or CIB for IFB (Example 1-1). The control unit type is STP and it has no devices defined to it. If a coupling facility link already exists between two servers, the STP control unit type cannot be defined. Likewise, if the STP control unit type is defined, a new CF Link cannot be defined unless the STP control unit is removed. For more details refer to 2.2.2, "Planning for timing-only links" on page 44.

Example 1-1 Coupling links

-----Source-----						-----Destination-----				-CU-	-#-	
/	CHPID	CF	Type	Mode	Occ	Proc.CSSID	CHPID	CF	Type	Mode	Type	Dev
_	83	Y	CIB	SHR	N	SCZP401.0	B2	Y	CIB	SHR	CFP	7
_	88	Y	CIB	SHR	N	SCZP401.0	B6	Y	CIB	SHR	CFP	7
_	A2	N	CFP	SHR	N	SCZP201.2	A7	N	CFP	SHR	STP	
_	A3	N	CFP	SHR	N	SCZP201.2	AC	N	CFP	SHR	STP	

1.6 External time source

If there are specific requirements to provide accurate time relative to an external time standard for data-processing applications, then consider using the external time source (ETS) function of STP. For servers in an ETR network or a Mixed CTN, the ETS function is provided by the Sysplex Timer. Refer to *Planning for the 9037 Model 2, SA22-7233*, for details. For servers in an STP-only CTN the external time source can be configured in one of the following ways:

- ▶ Dial out from the HMC to an available time service such as Automated Computer Time Service (ACTS). Accuracy of 100 milliseconds to the time provided by the time service can be obtained. The dial out option is available up to HMC Version 2.11.1 - starting with HMC Version 2.12.0 the dial out option is no more available since this HMC does not support a modem connection.
- ▶ Interface with an NTP server to maintain Coordinated Server Time accuracy of 100 milliseconds to the time provided by the NTP server. Simple Network Time Protocol (SNTP) client support is available on the IBM System zEC12, z196, z114, z10, z9 EC and z9 BC Support Elements to interface with NTP servers. It is also possible to configure the HMC as an NTP server. HMC Authentication Support is available starting with HMC Version 2.12.0 and provides NTP security either per Symmetric key (NTP Version 3 and 4) or Autokey (NTP Version 4).
- ▶ Interface with an NTP server with a pulse per second (PPS) output option. STP has been designed to track to the highly stable, accurate PPS signal from NTP servers that provide this output and maintain accuracy of 10 microseconds as measured at the PPS input of the System z server.

NTP client support addresses the requirements of customers who need to provide the same time across heterogeneous platforms in an enterprise. Dialing out provides time accuracy for the System z platform only, whereas attaching to an NTP server allows time accuracy as well as the same reference time source across heterogeneous platforms. In order to provide customers with the capability to enhance the STP time accuracy to an external time source, an option is available to configure an NTP server with a pulse per second output at the ETS. This type of ETS device is available worldwide from several vendors that provide network timing solutions. NTP client support and support of PPS are available on the zEC12, z196,

z114, z10 EC, z10 BC, z9 EC, and z9 BC configured in an STP-only Coordinated Timing Network.

The z990 and z890 servers do not support NTP (with or without the PPS option) as an external time source, however, they can use dial-out from the HMC. They can also participate in an STP-only CTN that has a IBM System z10 EC, Z10 BC or z9 configured to use NTP as an external time source.

Restriction: When ordering a new server (at the time of this publication zEC12, z196, or z114 can be ordered), the default HMC shipped with the order is 2.12.0 or above. If you still plan to use dial-out to a time service, you need to use an HMC that has not been upgraded to this level. In this case, you need to plan your CTN accordingly.

The Network Time Protocol was introduced in the early 1980s as part of the early development of Internet technology. NTP is used to synchronize the time of a computer to a time server or another time source reference. Developed over the last three decades, it provides accuracies typically within a millisecond on Local Area Networks (LANs) and up to a few tens of milliseconds on Wide Area Networks (WANs) relative to UTC. A timing signal can be obtained from a Global Positioning Service (GPS) receiver, a radio service, or another highly accurate time source⁴. A complete description of NTP can be found at the Web site:

<http://www.ntp.org>

NTP stratum

In an NTP-based network the servers are defined by stratum and are hierarchically arranged. In a CTN only three levels of stratum are supported or available, while the servers in an NTP V4 network can have an upper limit of 15.

A reference clock (for example, cesium clock, Global Positioning System, terrestrial broadcasts) itself is defined as stratum 0. A server directly attached to a reference clock is a stratum 1.

Attention: The *NTP stratum* level is completely independent of the *STP stratum* level. Although both STP and NTP use the term stratum, there is no direct relationship between the two.

For clarity, subsequently in this book we refer to *STP stratum* for server (CEC) stratum level in a CTN, and *NTP stratum* for referring to the NTP servers used as External Time Source.

⁴ From Web site <http://toi.iriti.cnr.it/>

Another requirement is high availability. NTP stratum 2 servers receive their time from one or more stratum 1 servers and are, according to the NTP Public Service Project, not significantly less accurate than a stratum 1 server. See Figure 1-14.

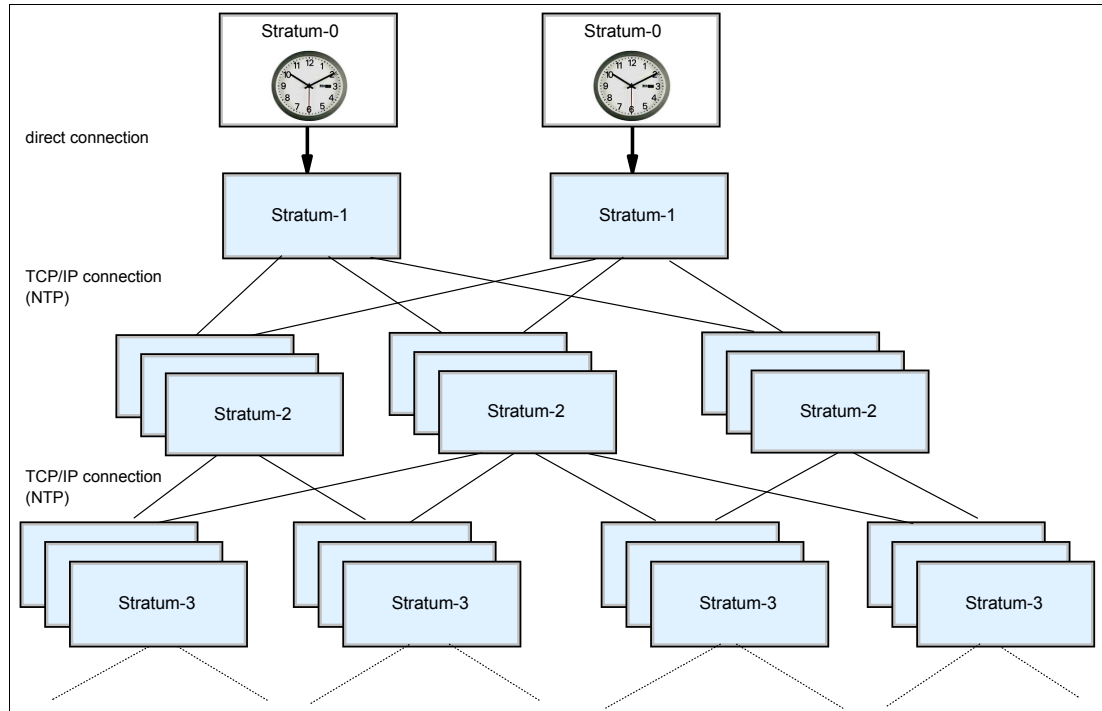


Figure 1-14 Stratum levels in an NTP network

For information about NTP servers see the NTP Public Services Project Web page:

<http://support.ntp.org>

1.7 Summary of STP terminology

Table 1-1 provides a summary of the STP terminology, acronyms that are used in this document, and some usage notes that facilitate understanding of the terminology.

Table 1-1 Summary of STP terminology

Name	Acronym	Description	Usage notes
Arbiter		Server assigned by the customer to provide additional means for the Backup Time Server to determine whether it should take over as the Current Time Server.	A stratum 2 server that must be attached to both the Preferred Time server and the Backup Time Servers.
Backup Time Server	BTS	Server assigned by the customer to take over as the Current Time Server (stratum 1 server) because of a planned or unplanned reconfiguration.	A stratum 2 server that must be attached to the Preferred Time Server and the Arbiter, as well as all the servers that are attached to the Preferred Time Server.
Coordinated Server Time	CST	The Coordinated Server Time in a CTN represents the time for the CTN. CST is determined at each server in the CTN.	

Name	Acronym	Description	Usage notes
Coordinated Timing Network	CTN	A network that contains a collection of servers that are time synchronized to CST.	<ul style="list-style-type: none"> ▶ All STP-configured servers in a CTN must have the same identifier, referred to as a CTN ID. ▶ Use CTN when there is no need to distinguish between a Mixed CTN and an STP-only CTN.
Coordinated Timing Network Identifier	CTN ID	Identifier that is used to indicate whether the server has been configured to be part of a CTN and, if so, identifies that CTN.	<p>Made up of two fields:</p> <ul style="list-style-type: none"> ▶ Field that defines STP network ID ▶ Field that defines ETR network ID
Current Time Server	CTS	A server that is currently the stratum 1 for an STP-only CTN.	Only the Preferred Time Server or the Backup Time Server can be a Current Time Server.
ETR network		A network that contains a collection of servers that are time synchronized to the Sysplex Timer.	All servers in the ETR network must have the same ETR network ID.
ETR network ID		An identifier assigned to a pair of synchronized Sysplex Timers. It is also the second field of a CTN ID whose low order 5 bits define the ETR network for which the server is configured.	ETR network ID 00 to 31, if the server is part of the ETR network, and not specified when STP-only CTN is configured.
ETR timing mode		The TOD clock has been initialized to the Sysplex Timer and is being stepped by stepping signals from the Sysplex Timer.	The server can be part of an ETR network or a stratum 1 server in a Mixed CTN.
External time source	ETS	The time source used to set and maintain the Coordinated Server Time close to an international time standard such as UTC.	The ETS can be a dial-out time service, an NTP server, or an NTP server with pulse per second output.
Local timing mode		The TOD clock has been initialized to a local time and is being stepped at the rate of the local hardware oscillator.	The server is <i>not</i> part of a synchronized timing network.
Mixed Coordinated Timing Network	Mixed CTN	The timing network that contains a collection of servers and has at least one STP-configured server stepping to timing signals provided by the Sysplex Timer. STP-configured servers in the Mixed CTN not stepping to the Sysplex Timer achieve synchronization by exchanging STP messages.	The CTN ID must be a valid STP-network ID and the ETR Net ID must be in the range 00–31.
Network Time Protocol	NTP	A protocol used to synchronize timekeeping among a set of time servers and clients.	See RFC-1305.
Preferred Time Server	PTS	The server assigned by the customer to be the preferred stratum 1 server in an STP-only CTN.	Usually the Current Time Server in an STP-only CTN. Planned or unplanned reconfiguration can result in the Backup Time Server being the Current Time Server.
Pulse per second	PPS	A highly stable and accurate output of some NTP servers to provide enhanced time accuracy for the CST.	The PPS input port is located on the ETR card of the System z10 and z9 servers.
Simple Network Time Protocol	SNTP	An adaptation of the Network Time Protocol.	See RFC-4330.

Name	Acronym	Description	Usage notes
STP network ID		The first field of the CTN ID that defines the STP network that the server is configured for.	When the STP-network ID is not specified, the server is not configured to be part of a CTN.
STP timing mode		The TOD clock has been initialized to CST and is being stepped at the rate of the local hardware oscillator, and is steered so as to maintain, or attain, synchronization with CST.	The server must be part of an STP CTN. In a Mixed CTN, only stratum 2 or 3 servers can be in STP timing mode.
STP-capable server		The server that has Licensed Internal Code (LIC) to support STP protocol.	The zEC12, z196, z114, z10 EC, z10 BC, z9 EC, z9 BC, z990, and z890 servers can be STP capable.
STP-enabled server		The STP-capable server that has STP function enabled.	
STP-configured server		The STP-enabled server that has a CTN ID defined.	
STP-only Coordinated Timing Network	STP-only CTN	The timing network that contains a collection of servers configured to be in STP timing mode. None of the servers in an STP-only CTN can be in ETR timing mode.	The CTN ID must be a valid STP-network ID and the ETR Net ID must not be specified. None of the servers are configured to be part of an ETR network.
Stratum		A means to identify the hierarchy of the server within the STP CTN or NTP server timing network.	The customer does not assign the stratum directly on a panel. Refer to <i>STP stratum</i> for server (CEC) stratum level in a CTN, and <i>NTP stratum</i> for the NTP servers used as External Time Source.
Stratum 0 server		An STP stratum level of zero is used to indicate that the stratum level is undefined in an STP-only CTN. This convention differs from an NTP stratum 0 server, which identifies a physical primary time source that is assumed to be accurate and has little or no delay associated with it.	Generally, to indicate that the CTN ID of a server is still undefined. If the CTN ID is defined, it means that the server does not have a usable clock source.

Name	Acronym	Description	Usage notes
Stratum 1 server		STP stratum 1 is the highest level in the hierarchy of a timing network that uses STP messages for synchronization. This convention differs from an NTP Stratum 1 server, which identifies a server that is attached to an NTP Stratum 0 server and typically acts as a server for timing requests from NTP stratum 2 servers via NTP protocol.	<p>In a Mixed CTN:</p> <ul style="list-style-type: none"> ▶ Any STP-configured server synchronized to the Sysplex Timer is a Stratum 1 server. ▶ Can have multiple Stratum 1 servers. <p>In an STP-only CTN, the Current Time Server is the Stratum 1 server.</p>
Stratum 2 server		System z server that uses STP messages to synchronize to an STP Stratum 1 server. This convention differs from an NTP Stratum 2 server, which identifies a server that sends requests to an NTP Stratum 1 server and typically acts as a server for timing requests from NTP Stratum 3 servers via NTP.	
Stratum 3 server		System z server that uses STP messages to synchronize to an STP Stratum 2 server. This convention differs from an NTP Stratum 3 server, which identifies a server that sends requests to an NTP Stratum 2 server and can act as a server for lower strata, potentially up to 16 levels, as per NTP protocol specification.	



Planning hardware and software

This chapter discusses the following planning concepts and considerations:

- ▶ Planning hardware
- ▶ Planning connectivity
- ▶ External time source
- ▶ Internal Battery Feature (IBF)
- ▶ Planning z/OS software: CLOCKxx
- ▶ Planning for z/VM®
- ▶ Special considerations for a single server STP-only CTN
- ▶ Planning for zBX
- ▶ Operating systems with no STP support

2.1 Planning hardware

STP is a server-wide facility that is implemented in the LIC of the zEC12, z196, z114, z10 EC, z10 BC, z9 EC, and z9 BC servers. This section discusses the following planning topics:

- ▶ Installing the Hardware Management Console (HMC) required to support STP. Support Element (SE) planning considerations are also covered.
- ▶ Installing the Licensed Internal Code (LIC) on the appropriate servers in the configuration to render them STP capable.
- ▶ Installing Feature Code 1021 on STP-capable servers and CFs to render them STP enabled.

2.1.1 Servers

An STP-capable server is a server that has the correct STP feature code installed. Only STP-capable servers can be STP enabled. LIC requirements to make servers STP capable and STP enabled are described later in this chapter. The following servers can be made STP capable:

- ▶ IBM zEnterprise EC12
- ▶ IBM zEnterprise z196
- ▶ IBM zEnterprise z114
- ▶ IBM System z10 Enterprise Class (z10 EC)
- ▶ IBM System z10 Business Class (z10 BC)
- ▶ IBM System z9 Enterprise Class (z9 EC)
- ▶ IBM System z9 Business Class (z9 BC)

Important: A zEC12 cannot be in the same CTN as z9 or previous servers.

A Mixed CTN allows servers that do not support STP to be synchronized in the same network with servers that support STP. However, the servers that do not have the STP feature installed do not support connection to Stratum 2 servers. The Sysplex Timer provides the timekeeping information in a Mixed CTN.

Note: zEC12 will be the last high-end server to support connections to a Mixed CTN.

Servers used as stand-alone coupling facilities can also be made STP capable. The servers that are used as stand-alone coupling facilities that can be made STP capable are the zEC12, z196, z114, z10 EC, z10 BC, z9 EC, and z9 BC. Table 2-1 provides a summary of the servers and the minimum Coupling Facility Control Code (CFCC) level required. Depending on their server, coupling facilities can be configured in an ETR network, a Mixed CTN, or an STP-only CTN.

Table 2-1 Coupling facility requirements

Server	CFCC level required on the coupling facility ^a
z9 EC and z9 BC	CFLEVEL 14 and higher
z10 EC and z10 BC	CFLEVEL 15 and higher
z196 and z144	CFLEVEL 17 and higher
zEC12	CFLEVEL 18 and higher

a. zEC12 cannot be in the same CTN as z9 or earlier servers.

2.1.2 Hardware Management Console

The HMC provides the user interface to manage either a Mixed CTN or an STP-only CTN.

HMC functionality in a Mixed CTN

In a Mixed CTN, the HMC is used to:

- ▶ Initialize or modify the CTN ID and ETR port states.
- ▶ Monitor the status of the CTN.
- ▶ Monitor the status of the coupling links initialized for STP message exchanges.

HMC functionality in an STP-only CTN

In an STP-only CTN, the HMC is used to:

- ▶ Initialize or modify the CTN ID.
- ▶ Initialize the time manually or by using an external time source to keep the Coordinated Server Time (CST) synchronized to the time source provided by the external time source (ETS).
- ▶ Initialize the time zone offset, daylight saving time offset, and leap second offset.
- ▶ Configure the HMC as an NTP server.
- ▶ Configure access to an ETS so that the CST can be steered to an external time source:
 - To schedule periodic access to a dial-out external time source ¹
 - To configure the connection from the NTP client on the SE to NTP servers
 - To configure the connection from the NTP client on the SE to NTP servers with pulse per second (PPS)
- ▶ Assign the roles of Preferred, Backup, and Current Time Servers, as well as Arbiter.
- ▶ Adjust time by up to +/- 60 seconds.
- ▶ Schedule changes to offsets listed previously. STP can automatically schedule daylight saving time based on the selected time zone.
- ▶ Monitor the status of the CTN.
- ▶ Monitor the status of the coupling links initialized for STP message exchanges.

Eligible Hardware Management Consoles

At a minimum, management of STP-enabled servers requires HMC application level V2.9.2 with the latest MCL. Some functions are added or removed in newer levels. For example, NTP server function is only available in HMC application level V2.10.1 and later, and HMC application level V2.12.0 and later no longer support modem (no dial-out capability).

The HMC application level must be equal to or higher than the level of Support Element consoles that it manages. For example, if a zEC12 is part of the CTN, HMC application V2.12.0 or later is required. Table 2-2 shows the major STP-related changes and the HMC Application levels that introduce them.

Table 2-2 HMC STP-related changes

Change	HMC application level	Supported servers
STP support	V2.9.2	z9 EC, z9BC and previous
NTP server in the HMC	V2.10.1 and V2.10.2	z10 EC, z10 BC and previous

¹ Dial-out not supported for HMC 2.12.0 or later.

Change	HMC application level	Supported servers
Removed ETR support	V2.11.0 and V2.11.1	z196, z114 and previous
HMC NTP authentication Removed dial-out support	V2.12.0	zEC12 and previous

Previously installed HMCs can be upgraded by requesting your IBM Service Representative to order the appropriate ECA for each server that has one or more HMC features that must be upgraded. zEC12 requires feature code 0091 and HMC Application level V2.12.0.

Setting up the HMC for local and remote operation

Because of the functionality provided by the HMC, it is required that all CPCs that need role assignment in the STP-only CTN be defined objects to any HMC used to manage the CTN. Furthermore, we strongly recommend that all CPCs in the STP-only CTN or Mixed CTN be defined objects to any HMC in order to have the capability for CTN reconfigurations (for instance, within a recovery scenario). Even remote HMCs should define all CPCs in the CTN. See Figure 2-1.

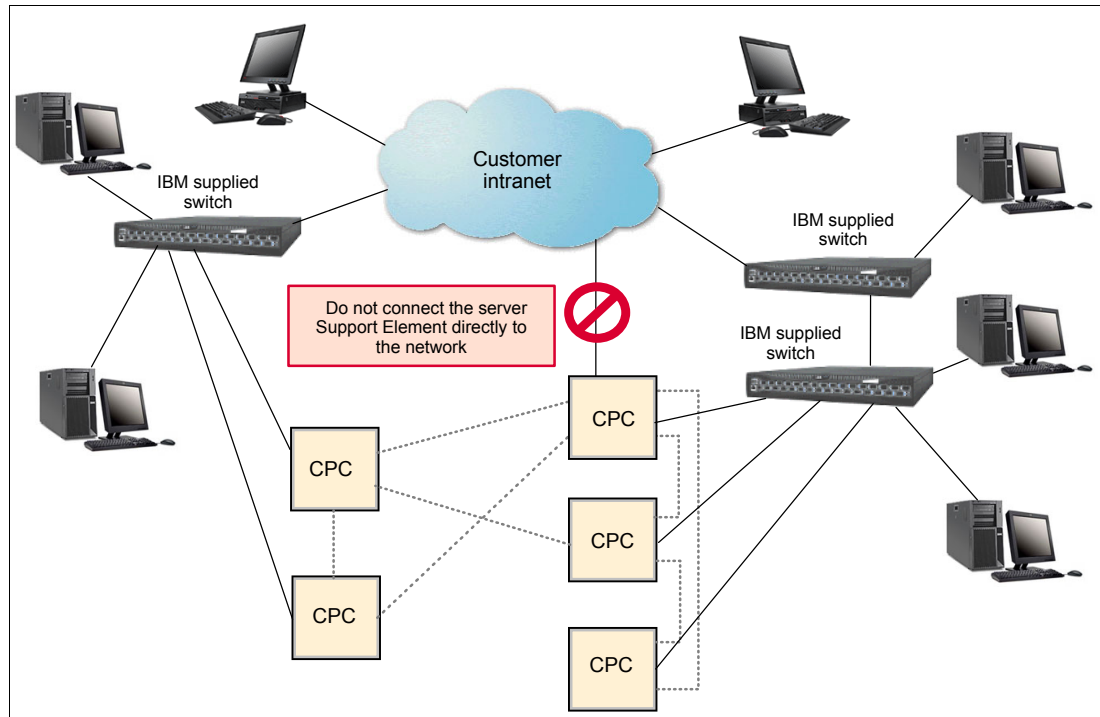


Figure 2-1 Hardware Management Console: local and remote operations

Important: CPCs that are not defined objects to the HMC being used to manage the CTN will not appear in the server pull-down list of the network configuration tab, except when they already have a role assigned (PTS, BTS, or Arbiter). If a CTN reconfiguration is required, it will not be possible to assign a role to the undefined servers. For a new role assignment, one HMC needs access to both the CTS and the target server.

There are two methods of defining server (CPC) objects to HMCs:

- ▶ When the HMC is on the same local LAN as the servers, local HMCs automatically detect the presence of Support Elements and automatically set up all the necessary internal configuration information for communication without additional information from the users.
- ▶ When the HMC is on a remote LAN, defining the servers requires manual action.

For details and guidance consult the manual *Hardware Management Console Operations Guide*, SC28-6919. The guide has an appendix that explains the process of defining servers to remote HMCs. This publication can be found on the IBM ResourceLink website at:

<http://www.ibm.com/servers/resourceLink/>

Note: In order to define the CPCs to the HMC, the HMC and the SEs must be part of the same HMC domain security.

Always use the switches supplied with the IBM servers to connect HMCs to SEs within a site and between sites. Remote HMCs should be connected via the same switches, and then connected from the IBM switch to the network devices. Figure 2-1 depicts an example of a network with multiple HMCs configured for both local and remote operations.

HMC application programming interface (API)

Specific STP-related attributes have been added to the Defined CPC objects starting with HMC Version 2.10.1. The following STP commands were added in order to manage the CTN or to be able to perform automated recovery actions such as reassigning the PTS, BTS, or Arbiter role in case of a server or site failure:

- ▶ **Swap Current Time Server**
- ▶ **Set STP Configuration**
- ▶ **Change STP-only CTN**
- ▶ **Join STP-only CTN**
- ▶ **Leave STP-only CTN**

For object command details refer to *Application Programming Interfaces*, SB10-7030.

HMC security

Security for remote HMC connectivity is provided by the HMC user logon procedures in the same way that it is for a local HMC. As with a local HMC, all communication between a remote HMC and each SE is encrypted. Certificates for secure communications are provided and can be changed by the user if desired. TCP/IP access to the remote HMC is controlled through its internally managed firewall and is limited to HMC-related functions.

Note: HMC Domain Security can be used to isolate servers on a common LAN or to provide additional security. Individual remote users can be configured to have restricted access in the same way that they can be configured on a local HMC.

For a complete description of HMC and Support Element connectivity options refer to the server installation manuals:

- ▶ *zEnterprise EC12 Installation Manual for Physical Planning*, GC28-6914
- ▶ *zEnterprise 196 Installation Manual for Physical Planning*, GC28-6897
- ▶ *zEnterprise 114 Installation Manual for Physical Planning*, GC28-6907
- ▶ *System z10 EC Installation Manual for Physical Planning*, GC28-6864
- ▶ *System z10 BC Installation Manual for Physical Planning*, GC28-6875
- ▶ *System z9 EC Installation Manual for Physical Planning*, GC28-6844
- ▶ *System z9 BC Installation Manual for Physical Planning*, GC28-6855

2.1.3 Support Elements

The Support Elements are important components of STP. The SE's disk storage provides the means for STP to initialize the configuration and timing parameters. STP configuration information is mirrored to the alternate SE in case the primary SE encounters a problem.

The STP ID, ETR ID, and ETR Port States are stored on the SE and preserved through a power-on reset (POR) and power off/on sequence. Server roles configuration (PTS, BTS, Arbiter) does not remain valid across a power-on reset or power off/on sequence of the CTS unless the CTN is configured with the **Only allow the server(s) specified above to be in the CTN** option selected.

Note: Configuration of a CTN that has *Only allows the servers specified above to be in the CTN* selected is described in detail in 3.4.2, "Single-server or dual-server auto resume after power-on reset" on page 125. This is also known as *saving CTN configuration across power-on reset operations*.

Disruptive actions, such as power-on reset, are only allowed for servers without roles, or for a server in a single-server CTN.

The server roles must be reassigned after:

- ▶ The CTS in a single server CTN has been power-on reset or has gone through a power off/on sequence and has not been specified to be the only member of the CTN. Refer to 3.4.2, "Single-server or dual-server auto resume after power-on reset" on page 125.
- ▶ The PTS and BTS, if assigned, in a two-server CTN experience a power outage at the same time (for example, a data center power outage) and they have not been specified to be the only members of the CTN. See 3.4.2, "Single-server or dual-server auto resume after power-on reset" on page 125.
- ▶ The PTS and BTS, if assigned, in a CTN with more than two servers experience a power outage at the same time (for example, a data center power outage).
- ▶ An STP-only CTN is deconfigured intentionally by the user.

In an STP-only CTN dial-out is possible from a HMC with HMC Application level V2.11.1 or previous, to adjust Coordinated Server Time (CST) to a time standard such as UTC. The SE supports automatic retrieval of time from a time service according to a user defined schedule. The dial-out schedule must be set up using the SE's scheduled operations task. We recommend that manual actions are used to mirror the dial out schedule immediately to the alternate SE. See the appropriate manual for instructions:

- ▶ *zEnterprise Support Element Operations Guide V2.11.1, SC28-6906*
- ▶ *System z10 Support Element Operations Guide V2.10.2, SC28-6882*
- ▶ *System z9 Support Element Operations Guide V2.9.2, SC28-6860*

Dial-out support removal: As with HMC version 2.12.0 dial-out support has been removed. This also impacts SE version 2.12.0. See:

- ▶ *zEnterprise Support Element Operations Guide V2.12.0, SC28-6920*

By default, Support Element data is mirrored automatically each day or when internal code changes are installed through single-step internal code changes.

When the external time source is configured to use an NTP server, with or without pulse per second, the NTP client code is on the Support Element. Access to the NTP server is automatic and does not have to be scheduled, as described in 2.3, "External time source" on page 46.

2.1.4 EC levels and MCLs

The STP Licensed Internal Code is already installed in zEnterprise servers, while System z10 and System z9 servers can be made STP capable by installing the STP LIC. The LIC can be installed concurrently. The STP LIC requirements are as follows:

- ▶ System z9
 - The System z9 must be at EC Driver level 67L or higher.
 - Prior to installing the STP feature, the Service Support Representative will download and install the latest MCLs available for the driver. Several STP prerequisite MCLs will be installed as a result.
 - After the MCLs are installed, the z9 EC or z9 BC become STP-capable servers. STP functions cannot be used until the STP feature is installed.
- ▶ System z10
 - The System z10 must be at EC Driver level 79D or higher.
 - Prior to installing the STP feature, the Service Support Representative will download and install the latest MCLs available for the driver. STP functions cannot be used until the STP feature is installed.

Attention: When a CEC is replaced via an MES, the new CEC is at the latest level. In an STP-only CTN the remaining CECs might be back level (MCLs) due to long intervals between update sessions. This can result in an instability in the CTN due to the fact that certain fixes or enhancements could be set only in the new box.

Also, when adding a new CEC to an existing CTN:

- ▶ Ensure that the STP feature is installed.
- ▶ Install the latest MCLs on the new and the existing CECs.

2.1.5 Installing the STP feature

The STP feature (Feature Code 1021) must be ordered for each server on which STP is to be enabled. This is a chargeable feature for each server. Feature code 1021 provides:

- ▶ Installation instructions
 - The IBM Service Support Representative (SSR) performs this part of the installation.
- ▶ A required Sysplex Timer EC LIC disk
 - The Sysplex Timer's Licensed Internal Code (LIC) has been upgraded to support using STP in a Mixed Coordinated Timing Network (Mixed CTN). The required Sysplex Timer LIC must be installed by the IBM SSR prior to a concurrent migration from a Sysplex Timer-based ETR network to a Mixed CTN.
- ▶ A required server diskette or CD
 - This feature should be installed on any server that will be configured in a Mixed CTN or an STP-only CTN.

STP is enabled without disruption to current operations. This is done through the server SE.

2.1.6 z/BX considerations

The zEnterprise servers' SE incorporate an NTP server for the use of NTP clients running on zBX firmware components. The goal for zEnterprise is to keep the time stamps in logs on the zBX components in reasonable time coordination with SE logs.

NTP time synchronization is provided also to the hypervisors running on the zBX Power, System x®, and DataPower® blades. For more information, see 2.8, "Planning for zBX" on page 74.

2.2 Planning connectivity

This section focuses on hardware connectivity requirements, offering planning information and recommendations.

2.2.1 Planning for coupling links for STP

STP identifies and maintains a list of all data links on the server that are capable of STP message communication. Coupling-peer data links are the only data links considered to be STP-capable data links.

Three types of external coupling links are available to exchange STP messages to keep servers and CFs synchronized:

- ▶ InfiniBand (IFB) links
- ▶ Integrated Cluster Bus (ICB)
- ▶ Inter System Channel (ISC) links

There are significant differences between the three with respect to the speed of the link and the maximum supported distance. From an availability point of view, all types should be considered to be equivalent, with the same considerations and recommendations. ICB links are used when the servers to be connected are within 7 meters (10 meters cable length). IFB links can be used for distances up to 150 meters, whereas IFB LR (long reach) links can be used for unrepeated distances up to 10 km, or 100 km with repeaters. ISC links can be used for unrepeated distances up to 10 km, or 20 km when running at 1 Gbps.

IFB links

IFB links are high-speed external coupling links available on zEC12, z196, z114, z10 and z9 servers. IFB links consist of a Host Channel Adapter (HCA) coupled directly to the processor I/O interface.

The IBM System z10 and zEnterprise servers support a 12x (12 lanes of fiber in each direction) InfiniBand-Double Data Rate (IB-DDR) coupling link designed to support a total link data rate of 6 Gbps. Another option is the IFB LR link, which supports a 1x (one lane of fiber in each direction) InfiniBand-Double Data Rate (IB-DDR) or InfiniBand-Single Data Rate coupling link for a distance of up to 10 km and is capable of a data rate of 5 Gbps. More information about planning for configuration of coupling links is available in *IBM System z Connectivity Handbook*, SG24-5444, and *Implementing and Managing InfiniBand Coupling Links on System z*, SG24-7539.

A 12x InfiniBand-Single Data Rate (IB-SDR) coupling link is available between System z10 and System z9 servers. This coupling link is designed to support a total link data rate of 3 Gbps. Any other IFB link cannot be used with the System z9 servers.

By feature code, InfiniBand coupling links are:

- ▶ FC 0167 (HCA1-O) is for the 12x IB-SDR links between z9 and z10 servers only, on z9.
- ▶ FC 0163 (HCA2-O) is for the 12x IB-DDR/SDR links on z10 and zEnterprise servers.
- ▶ FC 0168 (HCA2-O LR) is for the 1x IB-DDR/SDR links on z10 and zEnterprise servers.
- ▶ FC 0170 (HCA3-O LR) is for the 1x IB-DDR/SDR links on zEnterprise servers.
- ▶ FC 0171 (HCA3-O) is for the 12x IB-DDR/SDR links on zEnterprise servers.

Adapter ID assignment and VCHIDs

An adapter ID (AID) is assigned to every IFB link fanout at installation time. It is unique for the CPC. There is only one AID per fanout, so all ports on the fanout share the same AID. The adapter ID is:

- ▶ A number between 00 and 1F on z196, z10 EC, and z9 EC
- ▶ A number between 00 and 0B on z114
- ▶ A number between 00 and 05 on z10 BC
- ▶ A number between 08 and 0F on a z9 BC

In the I/O configuration program (IOCP or HCD), the AID and port number are used to connect the assigned CHPID to the physical location of the fanout.

A physical channel identifier (PCHID) normally has a one-to-one relationship between the identifier and a physical location in the machine, however a PCHID in the range from 0700 to 07FF lacks the one-to-one relationship between the identifier and the physical location, either because they do not have a physical card (like Internal Coupling connections (ICPs)), or because they are administered through different identifiers (as for IFB links, with the AIDs). No one-to-one relationship is possible due to the capability to define more than one CHPID for a physical location. Therefore, these are sometimes referred to as Virtual Channel Path Identifiers (VCHIDs).

VCHIDs are assigned automatically by the system and are not defined by you in the IOCDs. A VCHID is also not permanently tied to an AID. Therefore the VCHID assignment can change after a Power-On Reset (POR) if the hardware configuration changed (if a HCA was added or removed, for example). Due to the automatic assignment of the VCHID at every POR, the client or SSR needs to make sure that the correlation for the channel that they intend to manipulate has not changed. The VCHID that is currently associated with a coupling CHPID can be found by issuing an `MVS D CF,CFNM=xxxx` command for the associated CF.

You can also find the association between a CHPID and a VCHID, which is still referred to as a PCHID by the HMC and the SE, by displaying the VCHID details using the SE. In the HMC you can display the VCHIDs connected between the servers of your CTN, using the STP Status tab, in the System (Sysplex) Timer task. In this tab the VCHIDs will be organized between brackets and parenthesis. The ones that are between parenthesis are assigned to the same physical port in the adapter, and the ones between brackets are assigned to the same physical adapter. For example, in the Figure 2-2 the VCHIDs 704 and 725 are assigned to the same port in the same adapter, while 737 and 738 are assigned to different ports but are in the same physical adapter.

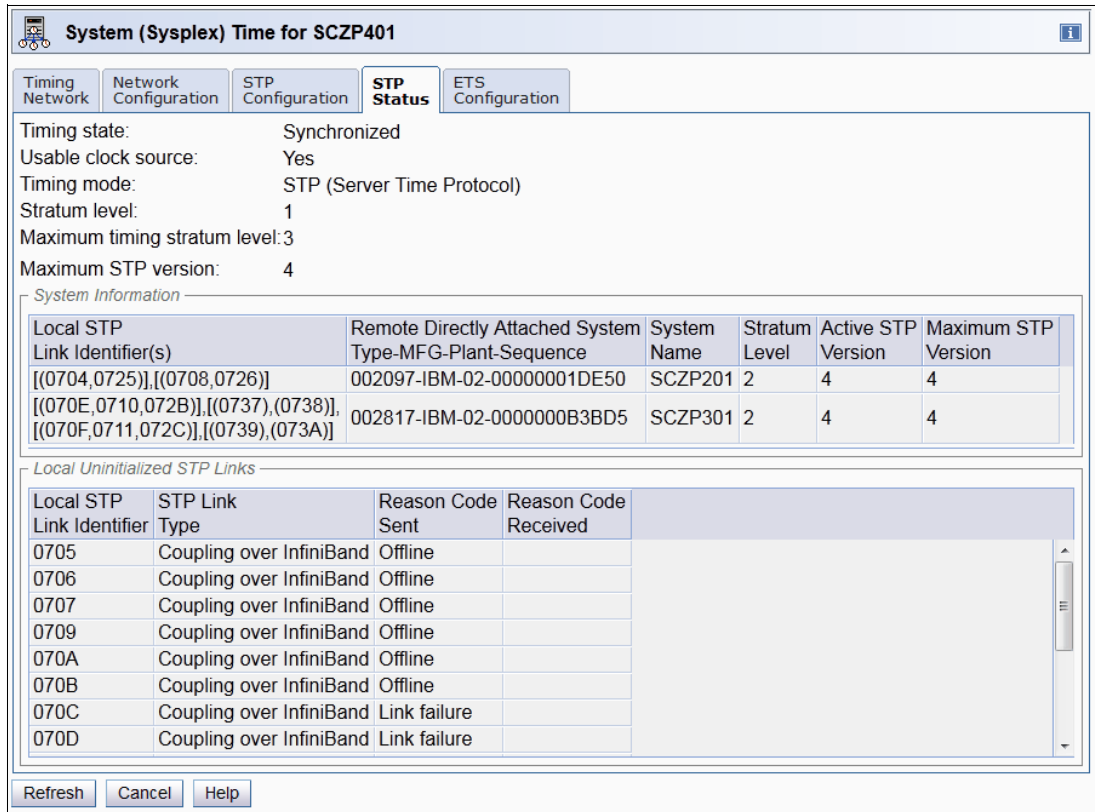


Figure 2-2 Displaying VCHID information (Local STP Link Identifier(s))

IFB cables

IFB links use a multi-pair optical fiber cable. This cable contains 12 pairs of multimode optical fibers, terminated at each end with a Multiple Fiber Push On (MPO) connector. This cable supports the InfiniBand specifications for a *12x* or *twelve lane* connection. The maximum cable length is 150 meters. The cable is not supplied with the IFB feature.

IFB LR links use a one-pair optical fiber cable. This cable contains one pair of single mode optical fibers, terminated at each end with a LC Duplex connector. This cable supports the InfiniBand specifications for a *1x* or *one lane* connection. The maximum cable length is 10 kilometers. The cable is not supplied with the IFB LR feature.

IFB3 Protocol

With the introduction of the HCA3-O fanout, it is possible to utilize an improved IFB protocol which is called IFB3. This new protocol provides improved service times for the IFB 12X link. However, certain conditions must be met to utilize it:

- ▶ The IFB3 protocol will only be used when both ends of the link are connected to an HCA3-O fanout.
- ▶ The IFB3 protocol will only be used if a maximum of four CHPIDs are *defined* per HCA3-O fanout port for all logical partitions (LPAR) combined.

The IFB link will automatically detect if the given requirements are met and will auto-negotiate the use of the IFB3 protocol. The two ports of an HCA3-O fanout are able to work in different protocol modes. For more information about InfiniBand links refer to *IBM System z Connectivity Handbook*, SG24-5444, and *Implementing and Managing InfiniBand Coupling Links on System z*, SG24-7539

ICB links

ICB links are high-speed external coupling links for short-distance connections and are not available for zEnterprise servers. ICB links consist of a 10-meter copper cable (up to seven meters between servers) that connects directly to the Self-Timed Interconnect (STI) I/O interface.

There are two different types of ICBs (ICB-3 and ICB-4). Which one is most appropriate for a given situation depends on the types of servers being connected:

- ▶ ICB-3 (FC 0993) supports a link data rate of 1 Gbps and is used to connect z9 EC, z9 BC, at a distance of 7 meters.
- ▶ ICB-4 (FC 3393) supports a link data rate of 2 Gbps and is used to connect z10 EC, z10 BC, z9 EC, and z9 BC servers at a distance of 7 meters.

Notes:

- ▶ ICB-3 links are not supported on System z10, and will be supported on System z9 through the lifecycle of System z9.
- ▶ ICB-4 links are not supported on zEnterprise servers and will be supported on System z10 through the lifecycle of System z10.

ICB cables

The total number of servers and ICB features determine the number of ICB cables to be ordered. For example, if two servers with four ICB features are being ordered and enabled, the total number of cables required is two. Proper planning ensures ordering the correct number of cables.

ICB cables are typically delivered with the server. Cables for ICB-3 and ICB-4 are not interchangeable. An ICB-3 cannot use an ICB-4 cable and vice versa. ICB-4 cables from System z9 are not interchangeable with cables for System z10. Installing ICB-4 links with System z10 requires new cables orderable with the System z10.

ICB-3 and ICB-4 cables are orderable as follows:

- ▶ ICB-3 Cable (FC 0227)
- ▶ ICB-3 Cable for Non Raised Floor (FC 0231)
- ▶ ICB-4 Cable (FC 0228)
- ▶ ICB-4 Cable for Non Raised Floor (FC 0232)
- ▶ ICB-4 Cable for System z10 to System z10 (FC 0230)
- ▶ ICB-4 Cable for System z10 to System z9 (FC 0229)

ISC-3 links

ISC-3 links in peer mode support a link data rate of 2 Gbps and are used to connect zEC12, z196, z114, z10, z9, z990, and z890 servers for distances up to 10 km (6.2 miles) without repeaters (WDMs are considered a repeater) using 9-micron single mode fiber optic cables. ISC-3 links are supported exclusively in peer mode, CHPID type CFP.

- ▶ FC 0217 - ISC-M (ISC Mother Card)
- ▶ FC 0218 - ISC-D (ISC Daughter Cards)
- ▶ FC 0219 - ISC-3 Port (2 ports per Daughter card, LIC enabled)

ISC-3 links running at 1 Gbps in Peer Mode (RPQ 8P2197) are supported by STP. This RPQ allows ISC-3 links to be extended to a maximum distance of 20 km without repeaters.

Important: zEC12 is the last server to support ISC-3 features, and can only be configured using carry-forward. It also requires an I/O Drawer or I/O Cage, and RPQ 8P2602 requires a 3 or 4 books configuration. ISC-3 links in compatibility mode, and older links prior to ISC-3 cannot be used to exchange STP messages.

ISC-3 cables

The ISC-3 uses a 9-micron single mode (SM) fiber and an LC Duplex connector, shown in Figure 2-3.




Figure 2-3 Single Mode Fiber: LC Duplex

Enhanced Book Availability (EBA)

The zEnterprise and System z servers are designed to allow a single book, in a multi-book server, to be concurrently removed from the server and reinstalled during an upgrade or repair action. During the process, it continues to provide connectivity to the server I/O resources using a second path from a different book. With enhanced book availability, and with proper planning to ensure that all resources are still available to run the critical applications in a (N-1) book configuration, planned outages can be avoided.

With Server Time Protocol, it is necessary to ensure that coupling links are not impacted during the removal of a book. This is especially critical where ICB and IFB links are concerned. ICB and IFB links are connected directly to books.

Based on the design of redundant I/O interconnect in I/O cages and drawers that allow for connection through multiple paths, ISC-3 connectivity will not be interrupted. Ensure that all the ICB and IFB links have redundant paths from different books. STI Rebalance feature (FC 2400) might need to be ordered for a System z10 or System z9 after a concurrent book add to provide even distribution of the connections across books.

This planning should be included as part of the initial installation and any follow-on upgrade that modifies the operating environment. The e-config report can be used to determine the number of books, active PUs, memory configuration, and the channel layout.

For more information about Enhanced Book Availability refer to:

- ▶ *IBM zEnterprise EC12 Technical Guide, SG24-8049*
- ▶ *IBM zEnterprise 196 Technical Guide, SG24-7833*
- ▶ *IBM zEnterprise z114 Technical Guide, SG24-7954*
- ▶ *IBM System z10 Enterprise Class Technical Guide, SG24-7516.*
- ▶ *IBM System z9 Enterprise Class Technical Guide, SG24-7124*

Maximum number of links

The maximum number of attached servers supported by any STP server in a CTN is equal to the maximum number of coupling links supported. The number of links that can be installed varies by server type.

Important: You must provide at least two coupling links between any two servers that are intended to exchange STP messages for synchronization. This prevents the loss of one link causing the loss of STP communication between the servers.

The maximum number of servers that can be attached to a Sysplex Timer is 24. In a high-availability configuration, the Sysplex Timer Expanded Availability (EA) configuration is installed, whereby each server attaches to the two timers in an Expanded Availability configuration.

Table 2-3 indicates the maximum number of coupling links each server type supports.

Table 2-3 Maximum number of coupling links

Link type	Maximum number of supported links ^a						
	zEC12	z196	z114	z10 EC	z10 BC	z9 EC	z9 BC
ISC-3	48 ^b	48	48	48	48	48	48
ICB-3	n/a	n/a	n/a	n/a	n/a	16	16
ICB-4	n/a	n/a	n/a	16	12	16	16 ^c
HCA1-O (IFB)	n/a	n/a	n/a	n/a	n/a	16 ^d	12
HCA2-O (IFB)	32 ^{be}	32 ^f	16 ^g	32 ^h	12	n/a	n/a
HCA2-O LR (IFB LR)	32 ^{be}	32 ^f	12 ^g	32 ^h	12	n/a	n/a
HCA3-O (IFB)	32 ^e	32 ^f	16 ^g	n/a	n/a	n/a	n/a
HCA3-O LR (IFB LR)	64 ⁱ	48 ⁱ	32 ^j	n/a	n/a	n/a	n/a

- a. The maximum number of coupling links combined (ICB-3, ICB-4, IFB, active ISC-3 links, and Internal Coupling) for System z10 and System z9 servers is 64, and for zEnterprise servers is 128.
- b. zEC12 supports ISC-3, HCA2-O and HCA2-O LR for carry forward only.
- c. A z9BC capacity setting A01 supports a maximum of eight ICB-4s.
- d. A z9 EC S08 supports a maximum of 12 IFB links.
- e. A zEC12 H20 supports a maximum of 16 HCA2-O, HCA2-O LR or HCA3-O IFB links.
- f. A z196 M15 supports a maximum of 16 HCA2-O, HCA2-O LR or HCA3-O IFB links.
- g. A z114 M05 supports a maximum of 8 HCA2-O, HCA2-O LR or HCA3-O IFB links.
- h. A z10 EC E12 supports a maximum of 16 IFB links.
- i. A zEC12 H20 and a z196 M15 support a maximum of 32 HCA3-O LR IFB links each.
- j. A z114 M05 supports a maximum of 16 HCA3-O LR IFB links

For System z10 and z9 servers, there is a maximum of 64 coupling link CHPIDs per server (ICs, ICB-4s, active ISC-3 links, and IFBs). For zEnterprise servers this maximum is 128 coupling link CHPIDs per server.

The maximum number of external coupling links combined cannot exceed the values shown in Table 2-4. Internal coupling channels (ICs) are not used by STP.

Table 2-4 Maximum number of external coupling links

Machine/Model	H66, H89 & HA1	H43	H20
zEC12	104	100	72
Machine/Model	M49, M66 & M80	M32	M15
z196	104	100	72
Machine/Model	M10	M05	---
z114	72	56	---
Machine/Model	EC	BC	---
z10	64	56	---

Coupling link CHPIDs must be defined in peer mode. They cannot be used for STP when defined in compatibility mode.

2.2.2 Planning for timing-only links

Timing-only links are coupling links that allow two servers to be synchronized using STP messages when a CF does not exist at either end of the coupling link. Refer to 1.5.2, “Coupling link requirements in a non-sysplex configuration” on page 23, for an example of time coordination requirements in a non-sysplex configuration.

HCD has been enhanced to define timing-only links and a new STP control unit. Installation of the HCD/IOCP/HCM support is required to obtain the enhancement.

Hardware Configuration Dialog: timing-only links

Hardware Configuration Dialog (HCD) and Hardware Configuration Manager (HCM) support the definition of timing-only links. The timing-only links bring STP capability into a basic-mode sysplex, or a non-sysplex set of systems that do not currently have coupling links in use the way that a Parallel Sysplex, by definition, would.

A timing-only link is established through the CF Channel Path Connectivity List panel in HCD. This is shown in Figure 2-4. Defining timing-only links is very much like defining other coupling links. IFB links are defined as CIB, ISC-3 links are defined as CFP, and ICB-3 or ICB-4 links are defined as CBP using the normal CF connection dialogs. As with coupling links, timing-only links can be defined as dedicated, reconfigurable, shared, or spanned. Multiple Image Facility (MIF) can be used to share the timing-only link between z/OS images on a server, but not between CF images on a server.

```

Connect to CF Channel Path

Specify the following values.

Source processor ID . . . . . : SCZP301
Source channel subsystem ID . . : 3
Source channel path ID . . . . . : A0
Source channel path type . . . . : CFP

Destination processor ID . . . . . SCZP201  +
Destination channel subsystem ID . . 0  +
Destination channel path ID . . . . A3  +

Timing-only link . . . . . Yes

```

Figure 2-4 HCD: timing-only link definition

The HCD dialog contains an option where the user can specify whether the CF connection is for a timing-only link. The default is no. If the user selects a timing-only link, HCD generates an STP control unit on each side of the connection, and no devices. The last line (see Figure 2-4) to define a timing-only link is added to the dialog window.

Note: The addition of timing-only links is fully supported by the Dynamic I/O Reconfiguration process so that they can be added non-disruptively.

For a timing-only link, HCD performs the following checks:

- ▶ The CHPIDs of a timing-only link must belong to different servers. A timing-only link cannot be looped back into the same server.
- ▶ If a CF connection already exists between two servers, HCD does not allow defining a timing-only connection between these servers. Also, if a timing-only link already exists between two servers, HCD will not allow defining a CF connection between these servers. The existing connection type must be broken before the other connection type can be established.

Note: CF messages cannot be transferred over timing-only links.

Timing-only links are supported between all System z and zEnterprise servers, *except* between a zEC12 and a z9². The status of these links can be displayed from the STP Status panel within the System (Sysplex) Time task in the HMC workplace. It reflects which links are available and initialized.

² System z9 and earlier systems are not supported in the same sysplex or CTN with zEnterprise EC12.

Standalone coupling facilities and timing-only links

Dynamic I/O reconfiguration supports the transition from timing-only to standard coupling links. However, a stand-alone coupling facility has no z/OS image to facilitate the initiation of a Dynamic I/O reconfiguration. A power-on reset is necessary to make the change at the coupling facility end of the link.

Important: We recommend that a single IODF be used for all servers in the installation.

2.2.3 Considerations for multi-site sysplex

In current business and IT environments, two important objectives for survival are systems designed to provide continuous availability and near transparent DR. Systems that are designed to deliver continuous availability combine the characteristics of high availability and near continuous operations to deliver high levels of service—targeted at 24x7. To attain high levels of continuous availability and near-transparent DR, the solution should be based on geographical clusters such as a multi-site sysplex and data mirroring. Geographically Dispersed Parallel Sysplex (GDPS) is an industry-leading solution that provides this capability. For more information about GDPS see:

<http://www-03.ibm.com/systems/z/advantages/gdps/index.html>

If a sysplex across two or more sites is configured, it is necessary to synchronize servers in multiple sites. Consequently, you might plan to extend the ISC-3 or IFB links beyond the distance supported without repeaters. IBM supports Wavelength Division Multiplexing (WDM) products qualified by IBM for use in multi-site sysplex solutions such as GDPS.

Attention: To transmit timing information, STP can choose any defined coupling link that is online between two zEnterprise or System z servers. If coupling links are configured over WDM equipment, *all* coupling links must use specific WDM hardware (optical modules^a, transponders^b, TDM modules^c) with interface cards qualified by IBM for STP.

If both coupling facility data and STP timing information must be transmitted between two servers, you cannot select a subset of coupling links to be used only for STP timing information.

- a. Including, but not limited to optical multiplexers/demultiplexers, optical add-drop modules (OADM), small form-factor pluggable (SFP) transceivers, dispersion compensation modules, optical amplifiers, and so forth
- b. Wavelength-converting transponders: converts client layer signal to WDM wavelength
- c. Time-Division Multiplexing modules: used to transmit multiple data links over the same WDM channel

The latest list of qualified WDM vendor products and links to corresponding Redpapers™ for each product can be found on the Resource Link® website at:

<https://www-1.ibm.com/servers/resourceLink/>

They are listed under Hardware products for servers on the Library page.

2.3 External time source

If there are specific requirements to provide accurate time relative to an external time standard for data processing applications, then consider using the ETS function.

In a Mixed CTN, the ETS function is provided by the Sysplex Timer. The planning considerations are the same as those for an ETR network.

In an STP-only CTN, the ETS function is provided by the following options:

- ▶ The dial-out service on the HMC
- ▶ An NTP server
- ▶ An NTP server with a pulse per second output option

The Current Time Server (CTS) is the only server that adjusts the Coordinated Server Time (CST) by steering it to the time obtained from an external time source.

Note: The HMC application level 2.12.0, required for zEC12, does *not* support dial-out.

2.3.1 Dial-out on the Hardware Management Console

When dial-out³ is configured from the HMC to an available time service, the following three protocols are supported:

- ▶ NIST⁴ Automated Computer Time Service (ACTS)
- ▶ IEN⁵ Telephone Date Code (CTD)
- ▶ NRC⁶ Canadian Time Service

A modem attached to the HMC is required to perform this function. The Coordinated Server Time can be initialized for the STP-only CTN to within +/- 100 ms of the time provided by the configured dial-out time service. Once the Coordinated Server Time has been initialized, a periodic dial-out to the time service, either manually or automatically, is needed in order to maintain the accuracy of the CST.

Note: The actual time accuracy, relative to UTC, depends on the accuracy provided by the dial-out time service.

Using dial-out to telephone time service consists of two steps:

1. Configure an HMC to dial out to telephone time services. At the External Time Source tab of the Customize Outbound Connectivity task on the HMC, check both the **Enable local system as a call-home server** and the **Allow external time source dialing using the local modem** check boxes to enable the configuration of the external time source. Configuration of the ETS should be done on multiple HMCs in order to have an alternate HMC for ETS queries.
2. Add scheduled operations to access the ETS from the Preferred Time Server (PTS) and the Backup Time Server (BTS), if applicable. Scheduled operations are included on the BTS for situations when the BTS becomes the CTS.

The HMC and SE support automatic retrieval of the time from a time service and automatic update of CST on a scheduled basis.

Setting up a schedule to dial out to the time service automatically can be done from the HMC. Select a server as the target object, then select **Operational Customization** \emptyset **Customize Scheduled Operations** \emptyset **Options** \emptyset **New** \emptyset **Access external time source**. One of the following can be requested:

- A single scheduled operation at a specified date and time

³ Dial-out function is not supported on HMC Application level 2.12.0 and later.

⁴ National Institute of Standards and Technology (USA)

⁵ Istituto Elettrotecnico Nazionale (Italy)

⁶ National Research Council (Canada)

- A recurring scheduled operation that occurs at a specified frequency

Scheduled operations that are to be executed by a server that is not the Current Time Server are ignored. At the scheduled time the SE requests the HMC to dial out to the time service. The HMC sends the difference with the time obtained from the service to the STP facility, which will make gradual adjustments by steering the CST to the time obtained from the external time source.

Note: The HMC Application level 2.12.0, required for zEC12, does *not* support dial-out.

2.3.2 NTP server

The use of a Network Time Protocol server as an external time source addresses the requirement of customers who need an accurate time source for the zEnterprise or System z server, or those who need to use an accurate common time reference across heterogeneous platforms. These customers, in most cases, purchase an NTP server that will obtain the time using high-precision time technology, such as GPS signals. Actual time accuracy, relative to UTC, depends on how accurate the NTP server time is with respect to UTC. Proper NTP server selection is more of a business decision than a technical requirement.

Note: Like STP, NTP is a hierarchical protocol with servers organized in stratum levels. However, there are significant differences between NTP and STP stratum, so these terms cannot be intermixed.

For further information about NTP refer to the RFCs at the following web pages:

- ▶ NTPv3: <http://www.faqs.org/rfcs/rfc1305.html>
- ▶ NTPv4: <http://www.faqs.org/rfcs/rfc5905.html>

For further information about NTP and the NTP Public services project refer to the following websites:

<http://www.ntp.org>
<http://support.ntp.org>

NTP client support on the Support Element

The NTP client on System z and zEnterprise servers can provide accuracy for the Coordinated Server Time by accessing an NTP server. An accuracy of 100 milliseconds to the time provided by the NTP server is possible. Simple Network Time Protocol (SNTP) client support is added to the STP code on the SE to interface with those NTP servers. Access to the NTP server from the CTS is initiated and controlled by the Support Element.

Planning for NTP client and LAN connectivity

The implementation of the SNTP client on the Support Element employs:

- ▶ NTP V3 (RFC-1305), NTP V4 (RFC-5905), or SNTP V4 (RFC-4330) on zEC12, z196, z114, System z10, and System z9 at Driver 67L
- ▶ IPv4 or IPv6 on zEC12, z196, z114, and System z10
- ▶ IPv4 on System z9 at Driver 67L

The NTP server can be either an external time source device available from several timekeeping device manufacturers, a local NTP server, or an NTP server configured on the HMC. The NTP traffic between the SNTP client (running on the Support Element) and the NTP server is not encrypted.

Several manufacturers market NTP servers that contain a 9037 Sysplex Timer connection in addition to NTP outputs. This option allows customers to migrate from an existing ETS connection on the 9037 to an STP-only CTN with NTP, without having to invest in a new NTP server. The solution provides time continuity (to the external time source provided by the unit) throughout the entire STP migration path. This might have value for customers with time accuracy requirements when migrating from an ETR configuration with an external time source to a Mixed CTN, and then to an STP-only CTN with NTP.

NTP server on the HMC

The NTP server configured as the ETS must be attached directly to the SE LAN. The SE LAN is considered in many configurations to be a private (dedicated) LAN and should be kept as isolated as possible. Allowing the NTP client running on the SE to directly communicate to an NTP server that is on the corporate or the external network (Internet) introduces potential security issues (even if the SE LAN is behind a firewall).

Note: For security reasons, we do not recommend configuring the SNTP client on the Support Element to directly access NTP servers outside its local area network.

The ability to define an NTP server on the HMC addresses potential security concerns because the HMC is normally attached directly to the SE LAN. The NTP server on the HMC can access another NTP server via a separate LAN connection to obtain its time reference (Figure 2-5).

Note: The NTP server function on the HMC does not provide a pulse per second output.

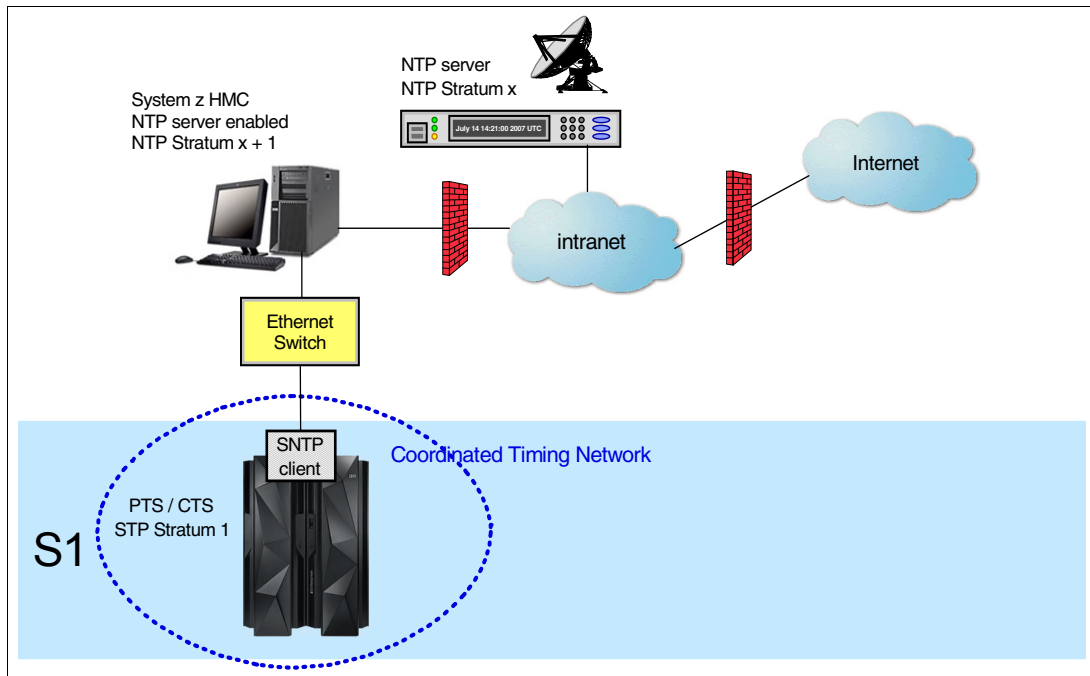


Figure 2-5 HMC acting as NTP server

HMC NTP broadband authentication support

The possibility to use HMC NTP authentication is added with HMC level 2.12.0. To exploit this option on the SE, configure the HMC with NTP authentication as the NTP server for the SE. The following scenarios describe the situations in which you could exploit this option.

► **Authentication support with a proxy**

Some configurations use a proxy to access outside the corporate data center. NTP requests are UDP socket packets and cannot pass through the proxy. The proxy must be configured as an NTP server to get to target servers on the web. Authentication can be set up on the proxy to communicate to the target time sources.

► **Authentication support with a firewall**

Some configurations use a firewall. HMC NTP requests can pass through the firewall. When using this configuration, you should use the HMC authentication to ensure un-tampered with time stamps.

► **Symmetric key and autokey authentication**

With the symmetric key and autokey authentication, the highest level of NTP security is available. HMC application level 2.12.0 provides panels that accept and generate key information to be configured into the HMC NTP configuration, and offer the possibility to issue NTP commands.

- Symmetric key (NTPv3-NTPv4) authentication
Symmetric key authentication, as described in RFC-1305 (made available in NTPv3), uses the same key for both encryption and decryption. Users exchanging data keep this key to themselves. Messages encrypted with a secret key can only be decrypted with the same secret key. Symmetric does support network address translation (NAT).
- Autokey (NTPv4) authentication
Autokey authentication, as described in RFC-5906 (made available in NTPv4), uses public key cryptography. The key generation for the HMC NTP is performed by clicking the “Generate Local Host Key” button on the Autokey Configuration panel. Pressing this button will issue the “ntp-keygen” command to generate the specific key and certificate for this system. Autokey authentication is *not* available with NAT firewall.
- NTP commands
NTP command support is also added to display the status of remote NTP servers and the current NTP server (HMC).

Planning for NTP server configuration

There are many possible configuration variations, depending on the type of NTP server equipment used, redundancy and security requirements, and existing LAN topology.

In the examples presented in this document, firewalls are indicated to draw the reader’s attention to the potential network security concerns.

Note: Configuration of network components, such as routers or firewall rules, goes beyond the scope of this document. Refer to *Hardware Management Console Operations Guide*, SC28-6919, and to the *Installation Manual for Physical Planning* (select the manual specific to your server) for security guidelines and recommendations for the HMC/SE network. For IBM zEnterprise and System z hardware documentation, see:

<https://www-01.ibm.com/servers/resourceLink/hom03010.nsf/pages/library?OpenDocument&login>

Example: Single NTP server (typically from a timekeeping device manufacturer)

Figure 2-6 on page 51 shows an example that has an NTP server attached directly to the SE LAN. Note that in this case, the NTP server is a single point of failure.

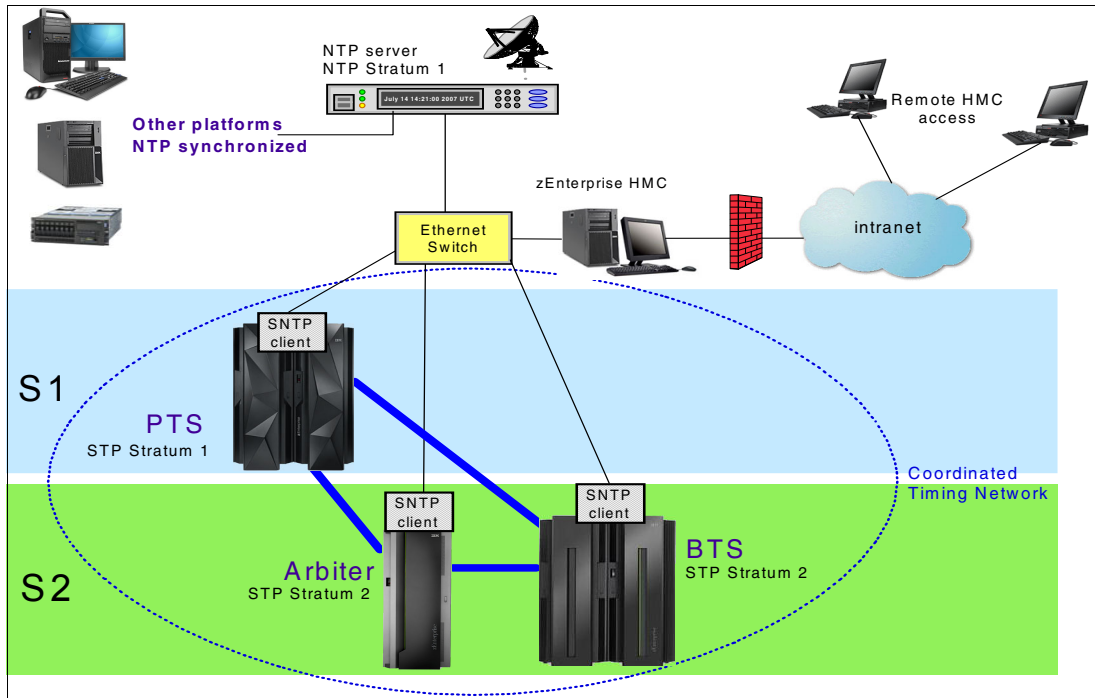


Figure 2-6 zEnterprise and other platforms configured with NTP server

One of the NTP output ports of the NTP server is connected directly to the HMC/SE local area network, allowing the SNTP client on the Support Element to obtain the external time from the NTP server. Other NTP output ports of this NTP server can be connected to other platforms that need the same accurate time reference.

In this example, there is no ETS (NTP server, in this case) redundancy. To provide redundancy, additional NTP servers can be installed and configured, as shown in the following section.

NTP server redundancy

The CTS is the only server that adjusts the CST by steering it to the time obtained from an external time source. To provide ETS redundancy, the user should consider configuring two or more NTP servers.⁷

Up to two NTP servers can be configured on each server in the STP-only CTN. When two NTP servers are configured, the user is responsible for selecting the preferred NTP server. This NTP server is called the selected NTP server. The other is called the non-selected NTP server. Normally, the SNTP client uses the time information from the selected NTP server to perform the time adjustment. The SNTP client also compares the quality of both NTP servers and informs the user in case either of the following conditions is detected:

- ▶ The selected NTP server has a stratum level that is lower in the hierarchy than the non-selected NTP server. (A stratum 1 server is a better choice than a stratum 2.)
- ▶ The time obtained from the selected NTP server has less accuracy than the non-selected NTP server.

⁷ The SE can handle up to two NTP servers. However, the HMC can configure 16 NTP servers. Three or more NTP servers is recommended for ETS redundancy for the CTN.

The following planning considerations should be made to provide NTP server redundancy:

- ▶ When two NTP servers are configured on the server that has the PTS/CTS role, STP will automatically access the second NTP server configured on the PTS/CTS if the selected NTP server fails.
- ▶ When NTP servers are configured on the server with the BTS role, the following is provided:
 - Access to an NTP server when the BTS becomes the CTS as the result of planned or unplanned recovery
 - Time adjustments to the STP-only CTN when the PTS/CTS cannot access any of its NTP servers

Recommendation: Configure at least one unique NTP server on both the PTS and the BTS. For redundancy, up to two separate NTP servers can be configured for both the Primary Time Server (PTS) and the Backup Time Server (BTS).

Example: Using the HMC to provide NTP server redundancy

A second example is illustrated in Figure 2-7. One of the NTP ports of NTP server 1 is connected directly to the HMC/SE local area network, as in the previous example. However, this example also introduces using the HMC as an NTP server to provide NTP server redundancy.

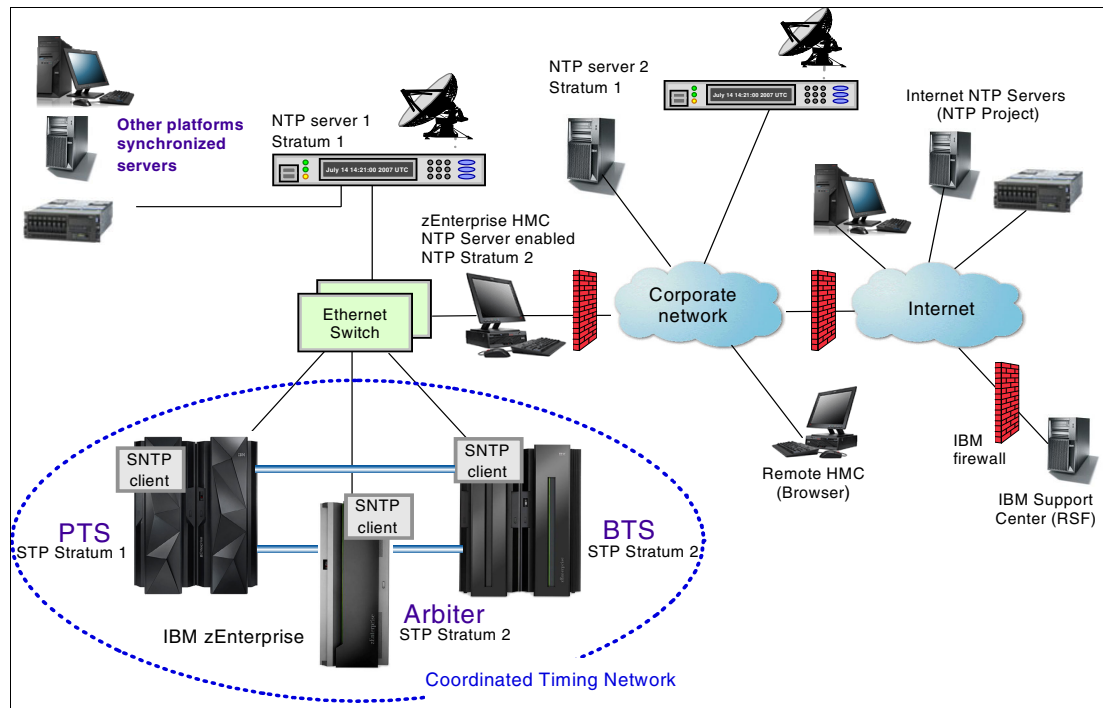


Figure 2-7 zEnterprise and other platforms configured with NTP servers: Example 2

Since the HMC is normally connected to the SE LAN, the capability to configure an NTP server on the HMC should be considered. The HMC can connect using a different network adapter to a remote NTP server located on the corporate intranet, or even on the Internet to access its time reference (NTP server 2 in Figure 2-7). Such a configuration requires UDP port 123 being enabled on the customer firewall to allow NTP traffic between the local NTP server (running on HMC) and remote NTP servers.

Alternatively, a separate NTP server (for example, running on UNIX or Linux) can be configured similar to the HMC (with dual Ethernet adapters) and connected directly to the SE LAN. This provides separation between the HMC/SE LAN and the external networks.

This configuration offers redundancy because the SNTP client on the SE allows configuration of two NTP servers. If the preferred NTP server (NTP Server 1 in this example) cannot be accessed, time adjustments can still be made by accessing the NTP server on the HMC. Refer to *Server Time Protocol Recovery Guide*, SG24-7380, for a more detailed discussion.

As in the first example, browser access to the HMC might be available using the corporate network through the customer firewall.

Example: NTP server redundancy on PTS and BTS

Figure 2-8 shows a third example. It involves multiple sites with NTP servers accessing remote NTP servers on the corporate network or the Internet. This example is different from the previous ones because the corporate network is always used as an access path to the time source. In addition, even though the example uses multiple sites, this configuration can be used in a single site as well.

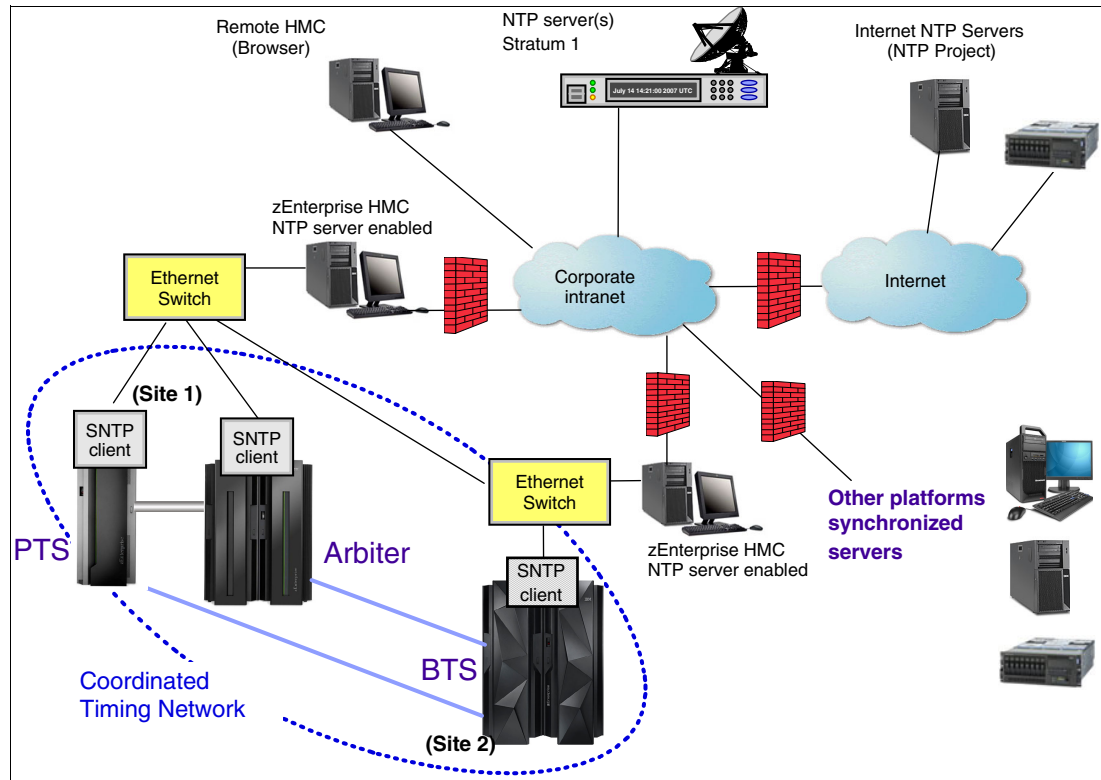


Figure 2-8 zEnterprise and other platforms configured with NTP servers: Example 3

This configuration shows that continuous availability of NTP servers can be obtained by configuring a different NTP server in site 1 (on the PTS/CTS) and site 2 (on the BTS). In each site a different HMC is providing the NTP server function. The HMCs in turn can get their time reference from the corporate NTP server 1. The BTS in site 2 accesses the NTP server in site 2 on a periodic basis, calculates time adjustments as required, and propagates them to the PTS in site 1. If the PTS/CTS cannot access the NTP server in site 1, it is able to use the adjustments calculated by the BTS to perform the necessary time adjustment steering for the entire CTN.

This example also requires UDP port 123 to be enabled in the customer firewall for access to NTP servers on the corporate network or on the Internet.

The examples presented illustrate that many variations are possible. The choice of one configuration versus another depends on the security and networking requirements of the enterprise.

For more information about NTP server recovery refer to *Server Time Protocol Recovery Guide*, SG24-7380.

2.3.3 NTP server with pulse per second

The time accuracy of an STP-only CTN has been improved by adding the capability to configure an NTP server that has a PPS output signal as the ETS device. This type of ETS device is available worldwide from several vendors that provide network timing solutions.

Pulse per second is an electrical signal that very precisely indicates the start of a second. STP has been designed to track to the highly stable, accurate PPS signal from the NTP server and maintain an accuracy of 10 microseconds as measured at the PPS input of the zEnterprise or System z server. A number of variables such as accuracy of the NTP server to its time source (GPS radio signals, for example) and the cable used to connect the PPS signal will determine the ultimate accuracy of STP relative to Coordinated Universal Time (UTC).

In comparison, if STP is configured to use a dial-out time service or an NTP server without pulse per second, it provides only a time accuracy of 100 milliseconds to the ETS device.

Planning for NTP server with PPS

To configure an NTP server with PPS, the NTP output of the NTP server must be connected to the Support Element LAN, and the PPS output of the same NTP server must be connected to the PPS input provided on the External Clock Facility (ECF) card of the zEnterprise or System z server.

The NTP information is provided via the HMC/SE local area network. However, the NTP server propagates its pulse per second signal via a coaxial cable. The signals are carried over a copper cable, which is limited in length. The cable length limitation depends on vendor specifications, and certain vendors might offer a fiber optic cable connection as an option. The cable is connected to a coaxial connector, which resides on the ECF card.

As a minimum, one NTP server with a pulse per second output should be configured at the server that has the Current Time Server role.

NTP server with PPS redundancy

Each zEnterprise or System z server has been equipped with two ECF cards, providing the capability of attaching one or two NTP servers with pulse per second. Attaching a second NTP server with pulse per second gives the redundancy that might be required in certain configurations.

When two NTP servers with PPS are configured, the user is responsible for selecting the preferred NTP server. This NTP server is called the selected NTP server with PPS. The other is called the non-selected NTP server with PPS. Normally, STP uses the time information from the selected NTP server to perform the time adjustment. The SNTP client also compares the quality of the NTP data received from both NTP servers and informs the user if either of the following conditions is detected:

- ▶ The selected NTP server with PPS has a stratum level that is lower in the hierarchy than the non-selected NTP server. (A stratum 1 server is a better choice than a stratum 2.)
- ▶ The time obtained from the selected NTP server with PPS has less accuracy than the non-selected NTP server.

The following planning considerations should be made to provide NTP servers with PPS redundancy:

- ▶ When two NTP servers with PPS are configured on the server that has the PTS/CTS role, STP will automatically access the PPS signal from the non-selected NTP server configured on the PTS/CTS if the PPS signal from the selected NTP server fails. Figure 2-9 is an example of this redundancy.
- ▶ When NTP servers with PPS are configured on the server with the BTS role, the following is provided:
 - Access to an NTP server with PPS when the BTS becomes the CTS as the result of planned or unplanned recovery
 - Time adjustments to the STP-only CTN when the PTS/CTS does not receive PPS signals from all of its configured NTP servers with PPS

Figure 2-10 on page 56 illustrates this type of redundancy.

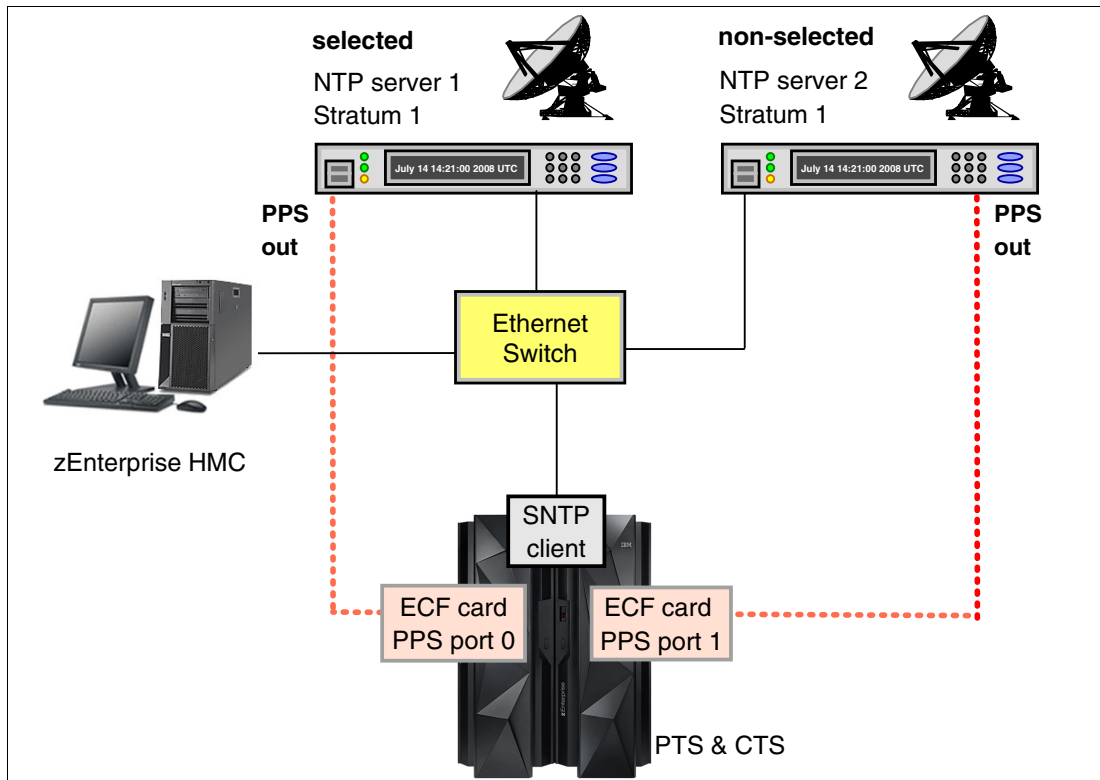


Figure 2-9 zEnterprise with NTP pulse per second

Recommendation: Configure at least one unique NTP server on both the PTS and the BTS. For redundancy, up to two separate NTP servers can be configured for both the PTS and the BTS.

Example: Two NTP servers with PPS configured on PTS/CTS

In this example, two NTP servers with PPS are configured on the PTS/CTS (Figure 2-9). The NTP outputs of both NTP servers are connected to the SE LAN via the Ethernet switch. The PPS output of NTP server 1 is connected to PPS port 0, and the PPS output of NTP server 2 is connected to PPS port 1.

Assume that the PTS/CTS is tracking to the PPS signal received on PPS port 0. If there is a failure associated with this PPS signal, STP will switch to the PPS signal received on PPS port 1 and can continue tracking to an ETS, thus maintaining time accuracy.

Refer to *Server Time Protocol Recovery Guide*, SG24-7380, for a more detailed recovery discussion.

Example: Different NTP servers with PPS on PTS and BTS

The configuration shown in Figure 2-10 allows continuous availability of ETS by configuring a different NTP server in site 1 (NTP server A on the PTS/CTS) and site 2 (NTP server B on the BTS). The PPS output of each NTP server is connected to PPS port 0 on the respective ECF card.

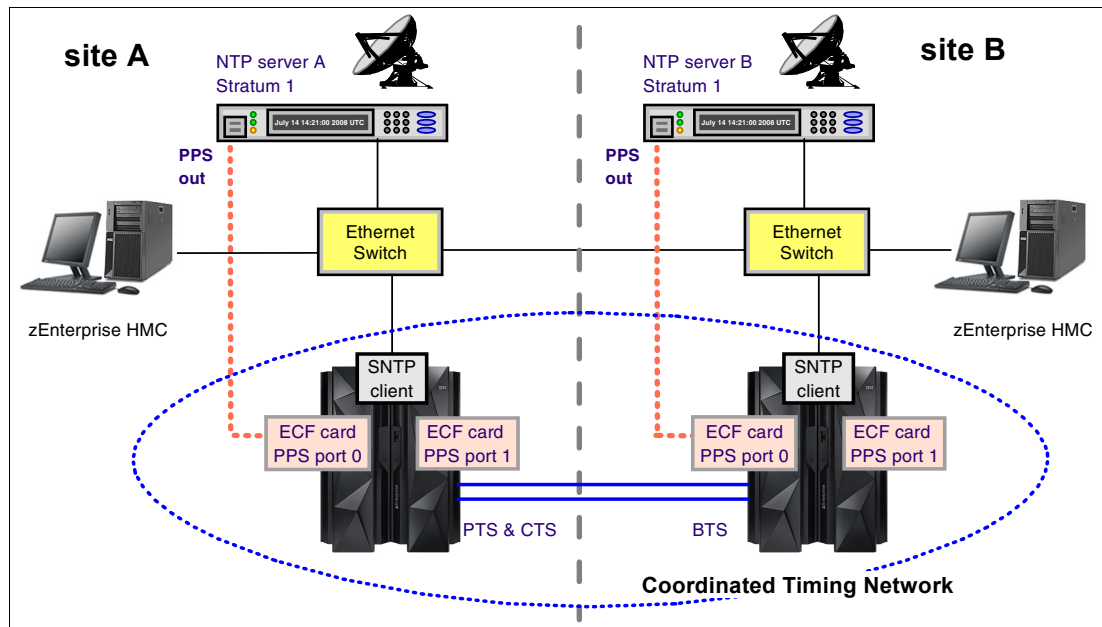


Figure 2-10 NTP servers with pulse per second on PTS and BTS

During normal operation, the STP-only CTN will track to the PPS signals received on PPS port 0 of the PTS/CTS. Also, the BTS in site 2 calculates time adjustments based on NTP data and PPS signals that it receives from NTP server B, and propagates them to the PTS/CTS in site 1. The data from the BTS is not used for steering until there is a planned or unplanned recovery action as follows:

- If the BTS takes over the role of the CTS due to a recovery action, then the STP-only CTN can still continue to track to PPS signals and NTP data received from NTP server B.

- ▶ If there is a failure that results in PPS signals from NTP server A not being received by the PTS/CTS, it is able to use the adjustments calculated by the BTS to perform the necessary time adjustment steering for the entire CTN.

For more information about NTP server with PPS recovery, refer to *Server Time Protocol Recovery Guide*, SG24-7380.

2.4 Internal Battery Feature (IBF)

In order to plan for improved STP recovery when power has failed for a single server (PTS/CTS) or when there is a site power outage in a multi-site configuration, we recommend that you install the Internal Battery Feature on one or more servers in the CTN. If an Internal Battery Feature is installed on your zEnterprise or System z server, STP now has the capability of receiving notification that customer power has failed and that the IBF is engaged. When STP receives this notification from a server that has the role of the PTS/CTS, STP can automatically reassign the role of the CTS to the BTS, thus automating the recovery action and improving availability.

If the entire CTN is located in a single data center and only has two servers (PTS and BTS), then we recommend installing the IBF on both the PTS and the BTS. This should provide recovery capabilities when the server that is the CTS experiences a power failure. If the CTN in a single data center has three or more servers, we recommend assigning the Arbiter, in which case the IBF does not provide any additional benefit for server power outages.

If the CTN spans two data centers, we recommend installing the IBF on the servers that will be assigned the roles of PTS, BTS, and Arbiter. This should provide recovery capabilities when the site where the CTS is located experiences a power failure.

For a more detailed recovery discussion when the IBF is installed, refer to *Server Time Protocol Recovery Guide*, SG24-7380.

2.5 Planning z/OS software

When migrating from an ETR network to a Mixed CTN or an STP-only CTN, or when a new STP-only CTN is being planned, there are z/OS considerations that must be planned for. Some key planning considerations are:

- ▶ Apply STP-related maintenance to all z/OS systems that need multisystem time synchronization.
- ▶ Update SYS1.PARMLIB(CLOCKxx) (optional).

2.5.1 z/OS requirements for STP

z/OS releases V1.7 and later support STP.

All z/OS images, whether participating in an STP-only CTN, a Mixed CTN, an ETR timing network, or in an unsynchronized mode, should be migrated to a supported release. This is particularly important when making a change to the timing network.

z/OS maintenance

Even though all supported releases of z/OS have the function to support STP, additional maintenance might be required. Maintenance being installed on the z/OS system image can

be done using the *rolling IPL* process, which avoids a sysplex-wide outage. See also “Requirement to re-IPL z/OS” on page 59.

In a sysplex configuration, if there is any coupling facility or z/OS image running on an STP-enabled server, STP maintenance is required on all z/OS images in the sysplex before the server is STP-configured.

The STP-related maintenance can be found via the appropriate FIXCAT for your server, for example `IBM.Device.Server.zEC12-2827.ServerTimeProtocol` contains STP-related maintenance for zEC12. To ensure you have the latest maintenance applied, download the current HOLDDATA from:

<http://service.software.ibm.com/holdata/390holdata.html#download>

After downloading, use the SMP/E RECEIVE HOLDDATA command to receive it to your GLOBAL zone, and then generate a report of missing fixes using the SMP/E commands shown in Example 2-1, replacing *targetzone* with your z/OS target zone.

Example 2-1 Generating fix report

```
SET BOUNDARY(GLOBAL).
REPORT MISSINGFIX FIXCAT(IBM.DEVICE.SERVER.*.SERVERTIMEPROTOCOL)
ZONES(targetzone).
```

The generated report will contain all the fixes that belong to this FIXCAT that have not been applied, and will indicate whether they have been received, as shown in Example 2-2.

Example 2-2 Identifying fixes

MISSING FIXCAT SYSMOD REPORT FOR ZONE targetzone

FIX CATEGORY	FMID	HOLD CLASS	MISSING APAR	HELD SYSMOD	RESOLVING SYSMOD		
					NAME	STATUS	RECEIVED
IBM.Device.Server.zEC12-2827.ServerTimeProtocol							
	HBB7780		AA39334	HBB7780	UA66257	GOOD	NO
			AA39522	HBB7780	UA66476	GOOD	NO
			AA39956	HBB7780	UA66079	GOOD	YES

Identify the SYSMODs that have not been received, download and receive them, and then apply them using the SMP/E APPLY command by specifying their names, or the FIXCAT as their SOURCEID as shown in Example 2-3, replacing the *targetzone* with your z/OS target zone and removing the CHECK keyword after verifying the HOLD information.

Example 2-3 Applying fixes

```
SET BDY(targetzone).
APPLY SOURCEID(IBM.DEVICE.SERVER.*.SERVERTIMEPROTOCOL)
BYPASS(HOLDERERROR) GROUPEXTEND CHECK.
```

Note: The addition of FIXCAT (fix category) information to Enhanced HOLDDATA and the REPORT MISSINGFIX function introduced in z/OS V1.10 replace the Enhanced PSP Tool function.

For information about SMP/E commands syntax or on receiving and applying SYSMODs, refer to *SMP/E for z/OS User's Guide, SA22-7773* or *SMP/E for z/OS Reference, SA22-7772*.

Requirement to re-IPL z/OS

In any of the following three situations you must *re-IPL* your z/OS images:

- ▶ After installing FC 1021 (STP code)
- ▶ After applying z/OS-related maintenance (as required)
- ▶ When changing CLOCKxx parameters

2.5.2 CLOCKxx

The z/OS system data set SYS1.PARMLIB contains many members that are used during IPL to determine how the z/OS system should be configured. One member is CLOCKxx, where xx is a 2-character suffix. The CLOCKxx member performs the following functions:

- ▶ Allows you to specify the difference between the local time and UTC if you do not want to use the time zone used by ETR or STP
- ▶ Provides the means of specifying that z/OS should make use of the Sysplex Timer and its time zone parameters
- ▶ Provides the means to specify that z/OS should make use of STP and its time zone parameters

In a multi-system environment, the objective is to share one SYS1.PARMLIB concatenation with as many z/OS images as possible. This simplifies system management by helping to present a single system image to system programmers so that changes must be made in one place rather than in every system image in the sysplex.

For more information about CLOCKxx and the Sysplex Timer see *z/OS MVS Setting Up a Sysplex*, SA22-7625.

The following CLOCKxx statements have been added in z/OS in support of STP:

- ▶ STPMODE
- ▶ STPZONE
- ▶ TIMEDELTA

Syntax format of CLOCKxx

Example 2-4 shows the CLOCKxx parameters format.

Example 2-4 Syntax format of CLOCKxx parameters and values

```
OPERATOR {PROMPT } | {NOPROMPT}
TIMEZONE d.hh.mm.ss
ETRMODE {YES} | {NO }
ETRDELTA nn
ETRZONE {YES} | {NO }
SIMETRID nn
STPMODE {YES} | {NO }
TIMEDELTA nn
STPZONE {YES} | {NO }
```

IBM supplied default for CLOCKxx

The IBM supplied default PARMLIB member of SYS1.PARMLIB is CLOCK00. If the STP-related maintenance is installed, CLOCK00 contains the parameters shown in Example 2-5.

Example 2-5 Default parameters and values for CLOCKxx

```
OPERATOR NOPROMPT
TIMEZONE W.00.00.00
ETRMODE YES
ETRZONE YES
STPMODE YES
STPZONE YES
TIMEDELTA 10
```

Statements and parameters for CLOCKxx

This section identifies the CLOCKxx statements and provides a short description of each statement.

Note: When planning to concurrently migrate from an ETR network to a mixed CTN or from a mixed CTN to an STP-only CTN, leave the ETR parameters (ETRMODE, ETRZONE) unchanged and add the STP parameters (STPMODE, STPZONE^a), after the ETR parameters. Special attention is required for ETRDELTA and TIMEDELTA. This approach facilitates both forward and reverse migrations. The ETR parameters can be removed after the migration to STP has been completed and the Sysplex Timers have been removed.

- a. CLOCKxx parameter ETRDELTA *should be replaced with* TIMEDELTA. ETRDELTA and TIMEDELTA are aliases for the same CLOCKxx parameter, and only one should be specified. If both are specified, the first one is ignored, and the following message will be generated:

```
IEA385I  CLOCKxx ETRDELTA & TIMEDELTA BOTH SPECIFIED.  yyyyyyy IGNORED
```

OPERATOR {PROMPT | NOPROMPT}

This specifies whether the operator is to be prompted to set the TOD clock during system image initialization.

- ▶ PROMPT
Specifies that the system image is to prompt the operator during TOD initialization.
- ▶ NOPROMPT
Specifies that the system image is not to prompt the operator during TOD initialization unless the TOD clock is not set. This is the default.

Notes

- ▶ If ETRMODE YES or STPMODE YES is specified, the system image ignores the OPERATOR parameter.
- ▶ OPERATOR PROMPT and SIMETRID are mutually exclusive keywords.
Specify *either* OPERATOR PROMPT or SIMETRID, but not both. If both are specified, the system image rejects the CLOCKxx member during system image initialization and issues a message to prompt the operator for one of the following:
 - A valid CLOCKxx member
 - End of Block (EOB) (by pressing the Enter key on the console) to use the defaultsOtherwise, the operator must re-IPL the system image.

ETRMODE {YES | NO}

This specifies whether z/OS should use ETR timing mode. ETRMODE YES must be coded if this system image is a member of a multisystem sysplex in an ETR network or a mixed CTN.

If all z/OS images in the sysplex run in logical partitions or under z/VM on a single physical processor, SIMETRID can be specified instead of ETRMODE YES.

▶ YES (default)

This specifies that z/OS is to use ETR timing mode, if available. If ETRMODE YES is specified and an operational Sysplex Timer is not available, the operator will be prompted to set the TOD clock during system image initialization.

▶ NO

This specifies that z/OS is not to use ETR timing mode.

Notes

- ▶ If the z/OS system image is running on a server that is in ETR timing mode, and both STPMODE and ETRMODE have been specified as YES, z/OS uses ETRMODE YES.
- ▶ If the z/OS system image is running on a server that is in STP timing mode, and both STPMODE and ETRMODE have been specified as YES, z/OS uses STPMODE YES.
- ▶ When planning to concurrently migrate from an ETR network to a Mixed CTN, or from a Mixed CTN to an STP-only CTN, it is important that both ETRMODE YES and STPMODE YES are specified prior to the migration for systems that are running on servers that are in ETR timing mode.

ETRDELTA nn

This indicates the greatest difference (in seconds), after IPL, between the system image's TOD and the Sysplex Timer TOD by which the system image adjusts its TOD, when necessary, to match the Sysplex Timer TOD.

If the difference between the system image's (CEC) TOD and the Sysplex Timer's TOD exceeds the ETRDELTA, the result is:

- ▶ If the system image is in ETR synchronization mode and is part of a multisystem sysplex, the WTOR message IEA015A is issued.
- ▶ If the system image is not part of a multisystem sysplex, processing continues, but the Sysplex Timer is not used for the remainder of that IPL.

If a value of 0 seconds is selected for ETRDELTA, no deviation between the processor TOD clock and the Sysplex Timer clock can be corrected by z/OS, so no TOD adjustment is possible. When a synchronization check is recognized in the ETRDELTA 0 case, one of the two previously described actions will result.

- ▶ Value range: 0–99 seconds
- ▶ Default: 10 seconds

Note: The default of 10 seconds is strongly recommended if you are planning a migration from an ETR network to an STP-only CTN.

Selecting an ETRDELTA value of 10 seconds does not mean that the CEC TOD and the Sysplex Timer TOD can be out of synchronization by as much as 10 seconds. The two TODs are synchronized to a tolerance of a few microseconds and the value of ETRDELTA has no effect on the synchronization tolerance between the processor TOD and the Sysplex Timer TOD.

To understand the significance of ETRDELTA and whether you should consider changing the default value of 10 seconds, it is important to first understand what an ETR sync check is. The Sysplex Timer transmits an ETR on-time signal, referred to as the OTE (on time event). It occurs every megamicrosecond - Mµs (1.048576 seconds), and corresponds to the carry-out

of bit 32 of the Sysplex Timer TOD. It is initially used as the starting signal for CPC TOD clocks in the clock-setting process, and later to monitor continued synchronization.

The CPC hardware performs this monitoring by comparing the time difference between the received OTE from the Sysplex Timer and the carry signal (bit 32 to bit 31) from the CPC TOD. An ETR sync check condition is generated when this difference is greater than the ETR sync check tolerance.

This tolerance ensures the correct sequence of execution of STCK instructions by different CPCs in the sysplex. The tolerance is CPC-model-dependent, with a typical value of 1 microsecond.

Some of the possible events that could result in a sync check are:

- ▶ z/OS is IPLed without the Sysplex Timer, but with ETRMODE=YES. In this case, z/OS uses the current value in the processor TOD clock during IPL. Later, when the Sysplex Timer is attached or powered on, it starts sending data to the attached CPCs. There is a very high probability that a sync check will be detected.
- ▶ For certain CPC models, a sync check condition might be detected when a port switch from the active to the alternate port occurs.
- ▶ There is a hardware malfunction in either the CPC or Sysplex Timer hardware.

When an ETR sync check condition occurs, the normal recovery action performed by z/OS is to resynchronize the TOD clock to the Sysplex Timer TOD. This is when the value assigned to ETRDELTA has significance. A sync check results in a machine check interrupt, at which time z/OS takes control of all the processor units in the CPC in order to attempt a resynchronization. z/OS compares the difference between the Sysplex Timer time and the TOD clock time, and if the time difference is within the NN value in ETRDELTA, it allows resynchronization and sets the TOD clock with the time value from the Sysplex Timer.

If the resynchronization requires setting the TODs *backward* in time, z/OS keeps all the processors in the CPC in a disabled spin loop until the new TOD value catches up in time. It is important that no application observes repeated time stamps. In essence, the ETRDELTA value is the maximum time that a spin loop might be observed when z/OS is attempting this resynchronization.

If the resynchronization requires setting the TODs *forward* in time, z/OS does not have to keep the processors in the CPC in a disabled spin loop and can set the new TOD value immediately.

Should the ETRDELTA default value be changed

A default value of 10 seconds was selected for ETRDELTA because it was felt to be a reasonable amount of time that would not cause system operators or users undue alarm if the system locked up for this duration because it was in a spin loop.

If you select a larger value for ETRDELTA (99 seconds, for example), operators and users are probably going to notice when the CPC is in a spin loop, and unnecessary recovery actions might be taken.

At the other extreme, if you select the smallest value for ETRDELTA, 0 seconds, no deviation between the processor TOD clock and the Sysplex Timer can be corrected by z/OS, so this is not a good idea. There can be small deviations between the two clocks (microseconds) in the normal course of operation and these deviations can be corrected by the software. There is no sense in putting processors in the sysplex in a wait state for deviations that can be corrected and are not significant to the operation of the sysplex. Any value less than 10

seconds will also limit the capability of attaching a Sysplex Timer after z/OS has been IPLed with ETRMODE=YES.

In summary, 10 seconds is a good working value for ETRDELTA and should be used unless there is a specific requirement to make the number smaller or larger. The point to keep in mind is that there is no reason for the processor TOD clock and the Sysplex Timer to be out of synchronization by more than a few microseconds.

See also “TIMEDELTA nn” on page 64. TIMEDELTA and ETRDELTA are basically aliases, and are not dependent on whether the server is in ETR or STP timing mode. If both are specified, z/OS will use the second one—whichever one that is—and reject the first one.

Example 2-6 shows a case when TIMEDELTA is specified after ETRDELTA.

Example 2-6 TIMEDELTA and ETRDELTA as aliases

```
TIMEZONE E.11.00.00
ETRMODE YES
ETRZONE YES
ETRDELTA 10
STPMODE YES
STPZONE YES
TIMEDELTA 10
```

TIMEDELTA will be used and the following message will be issued:

```
IEA385I CLOCKxx ETRDELTA & TIMEDELTA BOTH SPECIFIED, ETRDELTA IGNORED
```

Recommendation: We strongly recommend that the value applied to ETRDELTA and TIMEDELTA be left at the default of 10 (seconds). This recommendation is even more important for customers planning a reverse migration. A reverse migration refers to when an STP-only CTN is concurrently migrated to a Mixed-CTN.

ETRZONE {YES | NO}

This parameter specifies whether the system image is to get the time zone constant from the Sysplex Timer. The time zone constant specifies the difference between the local time and the UTC.

► **YES**

This specifies that the system image is to use the Sysplex Timer to set the time zone constant. If ETRZONE YES is specified and an operational Sysplex Timer is not available, the system image uses the time zone constant specified on the TIMEZONE parameter.

► **NO**

This specifies that the system image is to use the time zone constant specified on the TIMEZONE parameter. Also, it is possible to use the SET CLOCK command to set the local time and date. Specifying NO makes it more difficult to handle Daylight Saving Time.

► **Default**

- YES when ETRMODE is set to YES
- NO when ETRMODE is set to NO

STPMODE {YES | NO}

This specifies whether z/OS should use STP timing mode.

► **YES (default)**

This specifies that z/OS is to use STP timing mode if STP is configured.

► NO

This specifies that z/OS is not to use STP timing mode. Either ETR or local z/OS timing mode will be used depending on the value of the ETRMODE keyword and whether the ETR is usable.

Notes: If the z/OS system image is running on a server that is in ETR timing mode, and both STPMODE and ETRMODE have been specified as YES, z/OS uses ETRMODE YES.

If the z/OS system image is running on a server that is in STP timing mode, and both STPMODE and ETRMODE have been specified as YES, z/OS uses STPMODE YES.

When planning to concurrently migrate from an ETR network to a Mixed CTN, or from a Mixed CTN to an STP-only CTN, it is important that both ETRMODE YES and STPMODE YES are specified prior to the migration for systems that are running on servers that are in ETR timing mode.

TIMEDELTA nn

Important: TIMEDELTA and ETRDELTA are basically aliases, and are not dependent on whether the server is in ETR or STP timing mode. If both are specified, z/OS will use the second one—whichever one that is—and reject the first one.

This indicates the greatest difference, after IPL, between the system's TOD and the Coordinated Server Time (the STP stratum 1 server's time value), by which the system will adjust its TOD, when necessary, to match the stratum 1 server's TOD.

Note: Selecting a TIMEDELTA value of 10 seconds does not mean that the processor TOD and STP CST can be out of synchronization by as much as 10 seconds. The two are synchronized to a tolerance of a few microseconds and the value of TIMEDELTA has no effect on the synchronization tolerance between the processor TOD and CST.

The discussion on "ETRDELTA nn" on page 61 also applies to TIMEDELTA. This will help you understand the significance of TIMEDELTA and whether it should be changed from its default value.

If the difference between the system's TOD and the Coordinated Server Time (CST) exceeds the TIMEDELTA, the result is:

- If the system is in STP synchronization mode and is part of a multisystem sysplex, the WTOR message IEA394A is issued.
- If the system is not part of a multisystem sysplex, processing continues, but STP is not used for the remainder of that IPL.

If a value of 0 seconds is selected for TIMEDELTA, no deviation between the processor TOD clock and the Stratum 1 Server's TOD can be corrected by z/OS, so no TOD adjustment is possible. When a synch check is recognized in the TIMEDELTA 0 case, one of the two previously described actions will result.

- Value range: 0 to 99 seconds
- Default: 10 seconds

Note: The default of 10 seconds is strongly recommended if you are migrating from an ETR network to an STP-only CTN.

STPZONE {YES | NO}

This specifies whether the system image is to get the time zone constant from the STP. The time zone constant specifies the difference between the local time and the UTC.

- ▶ YES

This specifies that the system image is to use STP to set the time zone constant. If STPZONE YES is specified and STP is not enabled, the system image uses the time zone constant specified on the TIMEZONE parameter.

- ▶ NO

This specifies that the system image is to use the time zone constant specified on the TIMEZONE parameter.

- ▶ Default

- YES when STPMODE is set to YES
- NO when STPMODE is set to NO

SIMETRID nn

This specifies the simulated Sysplex Timer identifier. SIMETRID allows z/OS images running on the same server, in native mode in logical partitions, or as VM guests, to participate in a multisystem sysplex when no real Sysplex Timer or STP is available. In these environments, the z/OS TOD clocks are synchronized by PR/SM or the z/VM host. If a Sysplex Timer or STP is available when z/OS is not being run as a VM guest, IBM suggests using appropriate statements described previously instead of SIMETRID.

Do not use SIMETRID to coordinate z/OS images running on different servers. Instead, use a real Sysplex Timer together with ETRMODE YES, or run in STP timing mode with STPMODE YES.

SIMETRID is only available when z/OS runs in a logical partition, or under VM.

Each image participating in the same sysplex and using a simulated Sysplex Timer or a simulated STP for time synchronization must specify the same simulated Sysplex Timer identifier (nn).

When the logical TOD clock has been set in a logical partition by an IPL, the logical partition must be reactivated to reset the time again.

Note: OPERATOR PROMPT and SIMETRID are mutually exclusive keywords. Specify either OPERATOR PROMPT or SIMETRID, but not both. If both are specified, the system image rejects the CLOCKxx member during initialization and issues a message to prompt the operator for one of the following:

- ▶ A valid CLOCKxx member
- ▶ EOB (by pressing the Enter key on the console) to use the defaults

Otherwise, the operator must re-IPL the system image.

- ▶ Value range: a 2-digit hexadecimal number (X'00-1F')
- ▶ Default: None

This value is optional. If specified, ETRZONE and ETRMODE are ignored and the time zone constant specified by TIMEZONE is used.

TIMEZONE d.hh.mm.ss

This specifies the difference between the local time and the UTC. If ETRMODE YES and ETRZONE YES are specified and an operational Sysplex Timer is available, the system image ignores the TIMEZONE parameter. Similarly, if STPMODE YES and STPZONE YES are specified and the server is being synchronized using STP messages, the system image ignores the TIMEZONE parameter.

► *d*

This specifies the direction from UTC.

- Value range: E for East of UTC or W for West of UTC
- Default: W

► *hh.mm.ss*

This specifies the number of hours (*hh*) minutes (*mm*) and seconds (*ss*) that the local time differs from UTC.

- Value range: The value for *hh* must be between 00 and 15. The value for *mm* and *ss* must be between 00 and 59 (*mm.ss* values are optional). In addition, the combined *hh.mm.ss* value must be within the range (00:00:00–15:00:00). This means that a value like 15.59.59 is not valid because it is outside the range, even though the *hh* portion is between 00 and 15 and the *mm* and *ss* portions are between 00 and 59. If the *mm* portion or *ss* portion or both are omitted, a default value of 00 is applied to the omitted portion and appears in message IEA598I at IPL time. For example, if CLOCKxx contains TIMEZONE W.15, then at IPL time, message IEA598I indicates:

```
IEA598I TIME ZONE = W.15.00.00
```

- Default: 00.00.00

General rules and recommendations

Note: The synchronization mode for the server and the synchronization mode for a z/OS image running on the server are usually the same but can be different. If the server is in ETR timing mode, a z/OS image could be running in local, ETR, or SIMETR mode. If the server is in STP timing mode, a z/OS image could be running in local, STP, or SIMETR mode.

The following general rules apply to the creation of CLOCKxx:

- The following combinations of parameters are not valid and will be rejected:
 - ETRMODE=NO and ETRZONE=YES
 - STPMODE=NO and STPZONE=YES

Recommendation: STPMODE YES and STPZONE YES are the defaults with z/OS. Letting the STP settings default gives the most configuration flexibility. When both ETRMODE YES and STPMODE YES are in effect, the z/OS synchronization mode is determined by the timing mode of the server.

- If both ETRDELTA and TIMEDELTA are specified in the same CLOCKxx member, z/OS will use the second entry, whichever one that is, and flag the first entry as an error. Thus, ETRDELTA *should be removed and replaced with* TIMEDELTA.
- The CLOCKxx member for a system image that is a member of a multisystem sysplex must contain a specification of ETRMODE YES or STPMODE YES. The system image then uses the Sysplex Timer or STP timing signals to synchronize itself with the other members of the sysplex. The system image uses a synchronized time stamp to provide appropriate sequencing and serialization of events within the sysplex.

2.6 Planning for z/VM

z/VM now exploits the Server Time Protocol (STP) facility to generate time stamps for guest and system DASD write I/O operations, allowing these I/O operations to be synchronized with those of other systems. This support allows data used by z/VM and its guests to be replicated asynchronously over long distances by IBM System Storage® z/OS Global Mirror (formerly known as Extended Remote Copy, or XRC). For example, this allows z/VM to participate in a Geographically Dispersed Parallel Sysplex Metro/z/OS Global Mirror (GDPS/MzGM) environment. This satisfies the statement of direction in the XRC Services announcement dated March 23, 2010.

This baseline STP support is available in z/VM versions 6.2, 6.1 and 5.4, with the last two requiring the PTFs for APARs VM64814 and VM64816, and provides the following functions:

- ▶ Synchronizes z/VM TOD clock with STP server at IPL. Before the introduction of baseline STP support it was up to z/VM, under operator control, to change the TOD clock during IPL by the operator prompt for date and time, entering YES to the Change TOD clock {YES/NO} prompt during IPL. With STP support enabled in z/VM, you will not receive any prompts to change the TOD clock; instead, STP is automatically used to initialize the TOD clock with the CST and set the same time zone used by the CTN, displaying this message:

```
HCP986I "TOD Clock Synchronized via STP."
```

- ▶ Maintains a delta value of TOD changes over the lifetime of the z/VM IPL.
- ▶ Supports STP time zone management.
- ▶ CPC must be either a member of an STP-only CTN, or a stratum 2 or higher member of a mixed-CTN.

To enable this support, new FEATURES statements for SYSTEM CONFIG were added:

- ▶ **STP_Timestamping:** Timestamps will be added to write channel programs issued to all DASD devices that have the XRC LIC installed.
- ▶ **STP_TIMEZone/STP_TZ:** System time zone will be derived from the STP server.
- ▶ **XRC_OPTIONal:** System behaves differently when STP is suspended. Instead of deferring all I/O that is to be timestamped until STP sync is restored, it stops timestamping but continues issuing I/O.
- ▶ **XRC_TEST:** Only allowed on 2nd-level z/VM. This will enable STP_Timestamping without STP availability. Manually-specified TOD value is used for timestamping, intended for vendor test support.

For instance, Example 2-7 show the FEATURES section of a SYSTEM CONFIG that has STP_TZ enabled only.

Example 2-7 z/VM SYSTEM CONFIG with STP enabled

```
/******  
/*  
/* Features Statement */  
/******  
Features ,  
  Disable , /* Disable the following features */  
    Set_Privclass , /* Disallow SET PRIVCLASS command */  
    Auto_Warm_IPL , /* Prompt at IPL always */  
    Clear_TDisk , /* Don't clear TDisks at IPL time */  
  Enable,  
    STP_TZ , /* timezone from STP on CPC */
```

```

Retrieve ,                /* Retrieve options          */
  Default 20 ,           /* Default.... default is 20 */
  Maximum 255 ,         /* Maximum.... default is 255 */
MaxUsers noLimit ,      /* No limit on number of users */
Passwords_on_Cmnds ,    /* What commands allow passwords? */
  Autolog yes ,         /* ... AUTOLOG does          */
  Link yes ,            /* ... LINK does              */
  Logon yes ,           /* ... and LOGON does, too    */
Vdisk Userlim 144000 blocks /* Maximum vdisk allowed per user */

```

For more details about the FEATURES statements, refer to *CP Planning and Administration*, SC24-6178.

2.6.1 STP state flow in z/VM

Depending on what FEATURES statements have been used and on the STP status, you might have different results:

- ▶ During IPL:
 - If STP_Timestamping, STP_TZ, or both are specified: Perform activation by querying STP and setting the TOD clock to match the STP TOD value.
 - If STP_TZ is specified: Set the system timezone to match the STP timezone.
 - If either of these fails, STP will enter SUSPENDED state. Otherwise, STP activation completes successfully and STP is considered ACTIVE.
- ▶ When STP is ACTIVE:
 - If STP_Timestamping is specified: I/O to XRC-capable DASD will be timestamped (when required) and I/O to non-XRC-capable DASD will be unchanged.
- ▶ When STP is SUSPENDED:
 - If STP_Timestamping is specified and XRC_OPT is not specified: I/O to XRC-capable DASD that must be timestamped will be deferred until STP becomes ACTIVE again, while I/O that does not need to be timestamped and I/O to non-XRC-capable will continue to be issued.
 - If STP_Timestamping is specified and XRC_OPT is also specified: I/O to all DASD will be issued without a timestamp until STP becomes ACTIVE again.

Some events in the system might cause STP state changes:

- ▶ From ACTIVE to SUSPENDED: Occurs when an STP machine check is received that informs CP that the TOD value must be synchronized.
- ▶ From SUSPENDED to ACTIVE: In response to machine check/external interrupts received, CP will attempt to re-synchronize with the STP server. A successful resynch will *not* change the system TOD value, but will update the delta between the system time and the STP TOD value and will update the system time zone (if STP_TZ is enabled).
- ▶ From ACTIVE to ACTIVE: External interrupts might be received that require CP to query STP time zone information (for example, the time zone was changed via the HMC). This does not cause STP to go SUSPENDED, but will cause the system time zone to change.

2.6.2 Updated CP commands

With the baseline STP support to provide accurate XRC timestamping associated with VM guests and host I/O operations, some CP commands were added or updated:

- ▶ **QUERY STP:** This command has been added, and is used to display the Server Time Protocol status. Example 2-8 shows a possible response.

Example 2-8 Query STP information in z/VM

```
Q STP
Server Time Protocol synchronization activated.
Ready; T=0.01/0.01 17:30:45
```

- ▶ **SET TIMEZONE:** This command is used to change the system's time zone ID and time zone offset. It has been updated to display the response shown in Example 2-9 if z/VM is configured to use STP time zone.

Example 2-9 Setting the time zone

```
SET TIMEZONE EST
HCPTZN987E SET TIMEZONE not valid - STP timezone in use
Ready(00987); T=0.01/0.01 17:28:55
```

- ▶ **QUERY TIMEZONES:** This command is used to display the list of active and inactive time zone definitions on your system. It has been updated to show the STP time zones when z/VM is configured to use it.

For a full description of possible responses, refer to *CP Commands and Utilities Reference*, SC24-6175.

2.6.3 z/VM IOCP

If z/VM is used to create IOCP decks for the z/OS logical partitions (not z/OS Guest machines), z/VM can be used to create timing-only links. To define timing-only links, simply define the peer channel (CFP or CBP) with a control unit of UNIT=STP, and no devices. See 2.2.2, "Planning for timing-only links" on page 44.

2.6.4 VM guests

A z/OS sysplex can be established among z/OS images running as guests under z/VM, but the guests still cannot take advantage of STP or the Sysplex Timer. Only their write I/O to an XRC-enabled storage will have the STP timestamps if STP_Timestamping is specified. z/OS sysplex guests use z/VM timer services and must specify SIMETRID in the parmlib CLOCKxx member to synchronize their virtual TOD clocks.

To be able to change the date and time of a guest system running under z/VM, an OPTION TODENABLE statement must be added in the directory entry for the guest virtual machine. Then change the system date and time for the guest system using the procedures of the guest system. TODENABLE must not be specified for guest machines that will be specifying SIMETRID. This support requires that all of the TODs for the system be synchronized, and if TODENABLE is specified, the operator can subsequently change the guest TOD clock.

The three z/OS guest images in Figure 2-11 on page 70 belong to the same sysplex. The guest images must use SIMETRID in their CLOCKxx parmlib members. The TOD clock for a

VM guest must not be altered from the value set by LPAR during logical partition activation, either by prior systems IPLed in the partition or z/VM during its initialization.

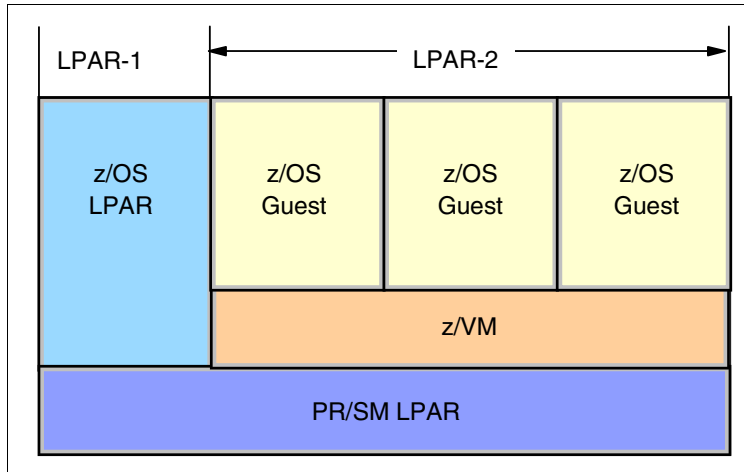


Figure 2-11 A sysplex within a logical partition and VM guests

2.6.5 z/TPF

Server Time Protocol (STP) is supported in z/TPF with APAR PJ36831. PJ36831 was released on PUT 07. Information about STP support in z/TPF can be found in the TPF Information Center:

<http://publib.boulder.ibm.com/infocenter/tpfhelp/current/index.jsp>

For more information about z/TPF, refer to the TPF web site at:

<http://www.ibm.com/tpf>

Or contact tpfqa@us.ibm.com.

2.7 Special considerations for a single server STP-only CTN

You might want to plan a single-server STP-only CTN because of one or more of the following reasons:

- ▶ You have a requirement to coordinate time accuracy across zEnterprise, System z and other platforms of your enterprise. As discussed in “NTP server” on page 48, you can use the ETS option of using an NTP server to accomplish this.
- ▶ You have a single server and want to implement a sysplex made up of multiple members. Even though this can be accomplished without using STP and coding SIMETRID nn for the CLOCKxx member of each z/OS image, there are reasons, mentioned in the following paragraphs, that you might want to implement this using STP and coding STPMODE YES in the CLOCKxx member.

You might have both z/OS systems as well as other operating systems that do not support STP running on the single server. In the following text, OS-xyz applies to any operating system that does not support STP, for example z/VSE® or Linux, as well as z/OS systems that do not have STPMODE YES coded in the CLOCKxx member. Even though STP is only supported by z/OS 1.7 and higher, all active LPARs on servers in the STP-only CTN can also

benefit from the time accuracy provided by using the ETS function. All LPARs (OS-xyz included) can still maintain time within 100 ms of the External Time Source.

There are a few instances that will cause a larger difference between the CEC time and the ETS time that OS-xyz running on a STP configured server will not be able to handle. These limitations do not apply to images that are at z/OS 1.7 and higher that have STPMODE YES coded in the CLOCKxx member. For details, refer to:

<http://www-03.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FLASH10631>

Note: The sync check limitation discussed in the referenced document does not apply to a single server CTN, since a stratum 1 server (PTS/CTS) is the source of time and cannot get sync checks.

Regardless of the reason you are planning a single server CTN, the hardware and ETS planning will be the same. The software planning in each case might be unique and will determine whether you can do a concurrent implementation or not.

Hardware and ETS planning for single server configuration

Hardware planning includes installing FC 1021 on an STP-capable server. There are no considerations for coupling links as this will be a single-server CTN (no external timing links required).

For planning your ETS configuration refer to 2.3, “External time source” on page 46.

Note: Since you are implementing a single server CTN, your server will be the Primary Time Server and Current Time Server (PTS/CTS). There is no Backup Time Server (BTS) or Arbiter in this configuration. However, all rules for providing ETS redundancy to the PTS/CTS apply here.

It is highly recommended that you save the STP configuration by selecting the **Only allow the server(s) specified above to be in the CTN** check box in the Network Configuration menu, as shown in Figure 2-12 on page 72.

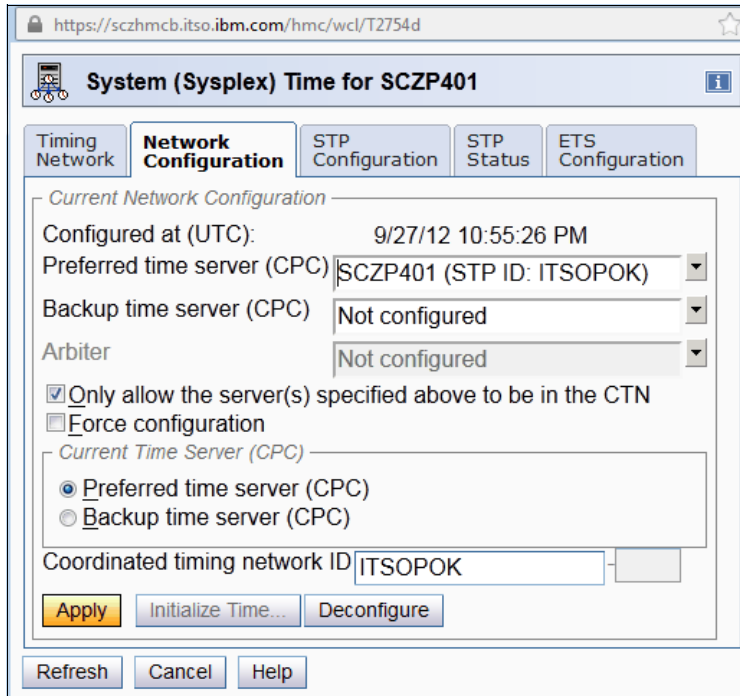


Figure 2-12 Enabling configuration save for single GEC CTN

Note: After a power off/on sequence, the TOD value in the Current Time Server is initialized from the TOD values stored in its Support Element. If you have previously configured an external time source, STP code does not perform an *initialize time* using the ETS, but automatically accesses the configured ETS and starts adjusting the time so as to maintain time accuracy. STP allows time adjustments of up to +/- 60 seconds.

2.7.1 Software planning: z/OS currently running in a sysplex

In this section, the assumption is that you have running z/OS systems in a sysplex with SIMETRID coded in the CLOCKxx member prior to configuring an STP-only CTN. You do not have a Sysplex Timer in your configuration attached to the single server.

The main software planning consideration is whether you should continue to run with SIMETRID or whether you should update the CLOCKxx member to use STPMODE YES and STPZONE YES or NO.

Continuing to run with SIMETRID will allow you to configure the CTN concurrently without stopping production. This is because in SIMETRID simulation mode, ETR and STP interrupts and machine checks are not applicable.

For configuring the CTN, refer to “Installation of a new STP-only CTN” on page 189. When the procedure gets you to the initialize time step for a new STP-only CTN, the TOD value can be changed since the z/OS systems are not running in STP or ETR synchronization mode. Appropriate adjustments to the LPAR offset of each z/OS image running with SIMETRID will be made by PR/SM code. There are no WTOR messages posted to the z/OS Console.

As mentioned in 2.5, “Planning z/OS software” on page 57, if a Sysplex Timer or STP is available when z/OS is not being run as a VM guest, IBM recommends using appropriate CLOCKxx statements instead of SIMETRID⁸ because of the following advantages:

- ▶ If you update the CLOCKxx member to use STPMODE YES and STPZONE YES, then Daylight Savings Time (DST) changes can be handled without using manual procedures you might be using today to change the time zone, such as using the SET TIMEZONE command.
- ▶ z/OS messages will be posted to the z/OS console when z/OS detects certain conditions with STP such as messages related to ETS. See Appendix B, “Server Time Protocol (STP) messages in z/OS” on page 199.
- ▶ If Leap Seconds are important to you, z/OS will recognize them when you specify Leap seconds during the Initialize Time procedure as well as later when Leap seconds are adjusted. See “Leap second offset adjustment” on page 110.

Attention: In order to change from z/OS systems running with SIMETRID to STPMODE YES, you must plan to take an outage of the z/OS systems to activate the STP mode. Once the STP-only CTN is configured, you can shut down and re-IPL the z/OS systems using the updated CLOCKxx member.

2.7.2 Software planning: z/OS systems currently not in a sysplex

In this section, the assumption is that you have running z/OS systems not in a sysplex prior to configuring an STP-only CTN. You do not have a Sysplex Timer in your configuration attached to the single server.

You might already be running with the defaults in CLOCKxx member, namely STPMODE YES and STPZONE YES. You might be disregarding the messages z/OS posts at IPL time and z/OS has been completing its IPL in local timing mode. Since the z/OS systems are currently running in local mode, you can configure an STP-only CTN without planning to take an outage and shutting down the z/OS systems.

After the STP-only CTN is configured (see “Installation of a new STP-only CTN” on page 189), plan a *rolling re-IPL* of your z/OS systems using either the defaults in CLOCKxx or updating CLOCKxx to use STPMODE YES and STPZONE YES or NO.

As each z/OS system is re-IPLed, it will now start operating in STP synchronization mode, allowing you to get the advantages stated in the previous section.

Note: Once the z/OS systems are running in STPMODE YES, you are now enabled to follow the procedures to set up a sysplex as well.

⁸ When a synch check occurs, LPAR takes an offset equal to the synch check and generates an STP machine check. z/OS systems running in STP timing mode would process the sync check by clearing the offset. However, SIMETRID systems are not enabled for STP sync checks and will not get an STP sync check machine check. This means the LPAR will have an offset (aka an EPOCH) which will not be cleared. The logical TOD for the SIMETRID LPARs will drift away from the STP time of the CEC as each sync check occurs. With OA26800 (see <http://www-01.ibm.com/support/docview.wss?uid=isg1OA26800>), the LPAR offset is not cleared at IPL and so there is no mismatch when systems are re-IPLed due to the LPAR EPOCH. However, when an LPAR is reactivated, the LPAR EPOCH is cleared. Now when z/OS is IPLed in that reactivated LPAR it will not be able to join the sysplex due to the LPAR EPOCH mismatch as the running system has a non-zero LPAR EPOCH, while the system being IPLed has a zero EPOCH. This means if one LPAR in the sysplex needs to be reactivated, then all LPARs in the sysplex must be reactivated to ensure there is no mismatch of LPAR EPOCHs.

2.7.3 Software planning: z/OS systems currently in a sysplex using ETR

In this section it is assumed that you might have consolidated a multi-server Parallel Sysplex configuration to a single-server configuration and you still have a Sysplex Timer attached to the single server in your configuration. Therefore, we assume that you have ETRMODE YES and ETRZONE as either YES or NO.

The software planning considerations are the same as outlined in 2.5, "Planning z/OS software" on page 57.

You can perform a concurrent migration from ETR to Mixed CTN and from Mixed CTN to STP-only CTN as described in Appendix A, "Migration planning" on page 163.

2.8 Planning for zBX

NTP clients on the BladeCenter® infrastructure, and DataPower appliances, access the SE NTP server every hour to keep time coordinated. The time adjustment is steered to correct time differences.

PowerVM blades

PowerVM® hypervisor accesses the SE NTP server every hour to keep time coordinated. The time adjustment is steered to correct for time differences. During the activation of the hypervisor, a hard adjustment ("jump") is done if the time delta is greater than +-15 seconds, otherwise the time difference is normally adjusted.

Each time a zBX virtual server (VS) on POWER7® is activated, it inherits the system time of its hypervisor (because we delete and re-create the LPAR). From then on, the system clock in the VS runs independently from the hypervisor, and is not corrected as the system clock in the VIOS partition is. All virtual servers on the same hypervisor will have the same skew rate. Since each VS is expected to be activated at different times, each VS might have different accumulated skew values.

Important: Once the VS has been activated, if you need to maintain time coordination, you should configure the OS to synchronize its time to an available NTP server.

System x blades

The integrated x hypervisor (xHyp) accesses the SE NTP server every hour to keep time coordinated. The time adjustment is steered to correct for time differences. During the activation of the hypervisor, a hard adjustment ("jump") is done if the time delta is greater than +-15 seconds, otherwise the time is normally adjusted.

Each time a zBX virtual server (VS) on System x blade is activated, it inherits the system time of its hypervisor. From then on, the system clock in the virtual server runs independently from the hypervisor, and is not corrected as the system clock is. All virtual servers on the same hypervisor will have the same skew rate. Since each VS is expected to be activated at different times, each VS might have different accumulated skew values.

Important: Once the virtual server has been activated, if you require to maintain time coordination with other servers in your environment, you must configure the operating system running on the virtual server to synchronize its time to an available NTP server.

2.9 Operating systems with no STP support

There are considerations to be aware of when running an operating system that does not have STP support on an STP-configured server. This is the case for logical partitions running z/VSE or Linux for System z. Consider the following information:

- ▶ The logical partition TOD is set to the processor TOD clock at the time that the partition is activated. When using an external time source, either a dial-out from the HMC or an NTP server, logical partitions on servers configured in an STP-only CTN will also benefit from the time adjustment from the ETS function.

An operating system that uses the hardware clock during IPL will obtain the time set in the logical partition TOD clock.

- ▶ No STP information will display through the operating system messages. STP hardware messages will still be available on the HMC.
- ▶ An STP sync check will not be recognized if one happens and will result in the partition's logical TOD being offset from the processor TOD clock.
- ▶ To reset the logical partition TOD you must deactivate and activate the logical partition, then re-IPL the operating system.
- ▶ Leap seconds will not be recognized.

For more detailed information refer to the Techdocs library at:

<http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/FLASH10631>

2.9.1 Linux for System z

The Linux operating system does not support STP or the Sysplex Timer. However, it tolerates their existence. That is, if Linux is running on a server configured with STP with a Sysplex Timer, Linux continues to operate normally and has no awareness of STP or the Sysplex Timer.

The server's TOD is synchronized to STP or the Sysplex Timer by the LPAR LIC. The logical partition has a logical TOD offset from the processor TOD, which can be set by Linux during IPL. If a synchronization check occurs, the LPAR LIC adjusts the logical TOD offset for each logical partition. To zero out the offset, the logical partition must be deactivated and activated again.

The Linux kernel keeps track of time independently from the hardware clock. During the boot process, Linux sets its own software-managed clock to the same time as the hardware clock. Alternatively, the clock can be set through operator commands. After the boot, the Linux clock runs independently of the hardware clock.

Because Linux manages its own software clock, Linux frequently uses Network Time Protocol to keep its time synchronized to an external source. If the Linux image (logical partition or guest) has a Local Area Network (LAN) connection to an NTP server, Linux can take advantage of it.

2.9.2 z/VSE

The z/VSE operating system does not support STP or the Sysplex Timer. However, if z/VSE is running on an STP-enabled server or with a Sysplex Timer, z/VSE continues to operate normally and has no awareness of STP or the Sysplex Timer.

The server's TOD is synchronized to STP or the Sysplex Timer by the LPAR LIC. The logical partition has a logical TOD that can be set by z/VSE during IPL. If a synchronization check occurs, the LPAR LIC adjusts the logical TOD offset for each logical partition. To zero out the offset, the logical partition must be deactivated and activated again.

It is up to z/VSE (under operator control) to change the TOD clock during IPL by the operator prompt for date and time. The z/VSE operator can change the TOD clock during IPL by specifying new values via the SET DATE and SET CLOCK command. As a VM guest, z/VSE can modify only the virtual TOD clock of the virtual machine. See the operating system publications for more information.



Operations

This chapter discusses considerations when planning for the operational management of the Coordinated Timing Network (CTN) once your configuration is in a Mixed or STP-only CTN. This chapter covers the following topics:

- ▶ “Displaying and monitoring the environment” on page 78
- ▶ “Configuring the Coordinated Timing Network” on page 92
- ▶ “Managing the CTN configuration” on page 125
- ▶ “Operational considerations” on page 134
- ▶ “MES upgrade considerations” on page 148

Note: Configuring the Coordinated Timing Network is covered in the *Server Time Protocol Implementation Guide*, SG24-7281.

3.1 Displaying and monitoring the environment

This section discusses the facilities available for monitoring your Mixed and STP-only CTNs. The facilities covered are:

- ▶ CPC Details panel
- ▶ System (Sysplex) Time panel
- ▶ z/OS commands and messages
- ▶ CFCC commands and messages

3.1.1 CPC Detail panel

The CPC Details panel on the Hardware Management Console (HMC) or Support Element (SE) is the first panel to look at for timing information about a server, as shown in Figure 3-1. If your server has the STP Feature Code #1021 installed, there will be a tab entitled STP Information. This is the easiest way to check whether your servers have the STP feature code installed. The number of fields shown on the STP Information tab varies depending on the type of timing network. This is a display-only panel. It does not allow for any modification of the timing network. Any changes to the CTN require you to cancel out of the panel and return (there is no refresh button).

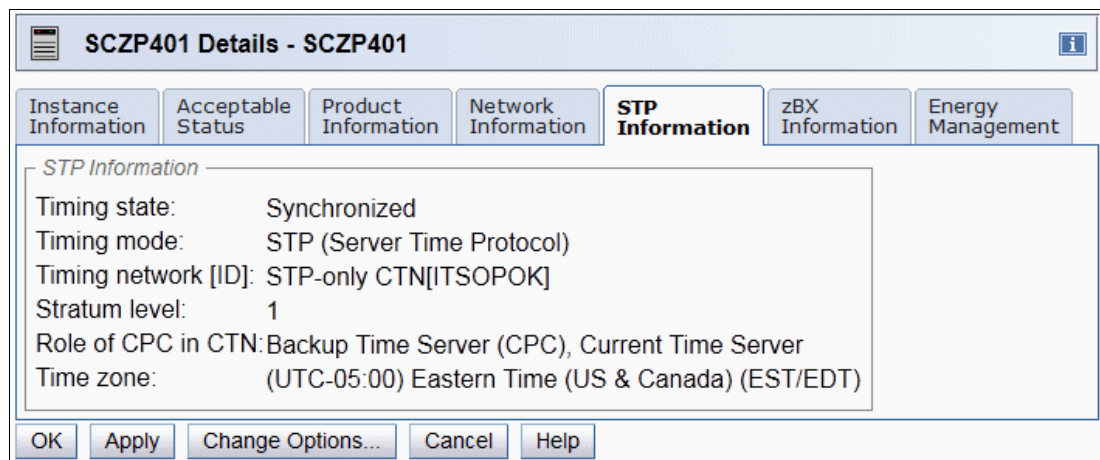


Figure 3-1 CPC Details panel from HMC Workplace

The left Context-Sensitive Help area is a newly added function for HMC and SE Version 2.10.0 that provides brief descriptions of the action buttons available on each individual panel. The buttons can be sequentially highlighted by pressing the Tab key and thus getting the information for all of them. The help function itself is activated or deactivated by clicking the **i** button in the upper right corner of the opened windows (Figure 3-1).

Note: Ensure that the HMC that you intend to use as an operational focal point for CTN management is at the latest driver level with the latest maintenance change levels (MCLs).

3.1.2 System (Sysplex) Time panel

When the System (Sysplex) Time icon is selected on the HMC, with a single CPC targeted, initially you will see the Secondary Object Notification for Disruptive Task message, as shown in Figure 3-2. This message reminds the user that time configuration commands can potentially be disruptive to operating system images on this CPC. If your server is not equipped with the ETR feature or the STP feature #1021, then this task is not available.

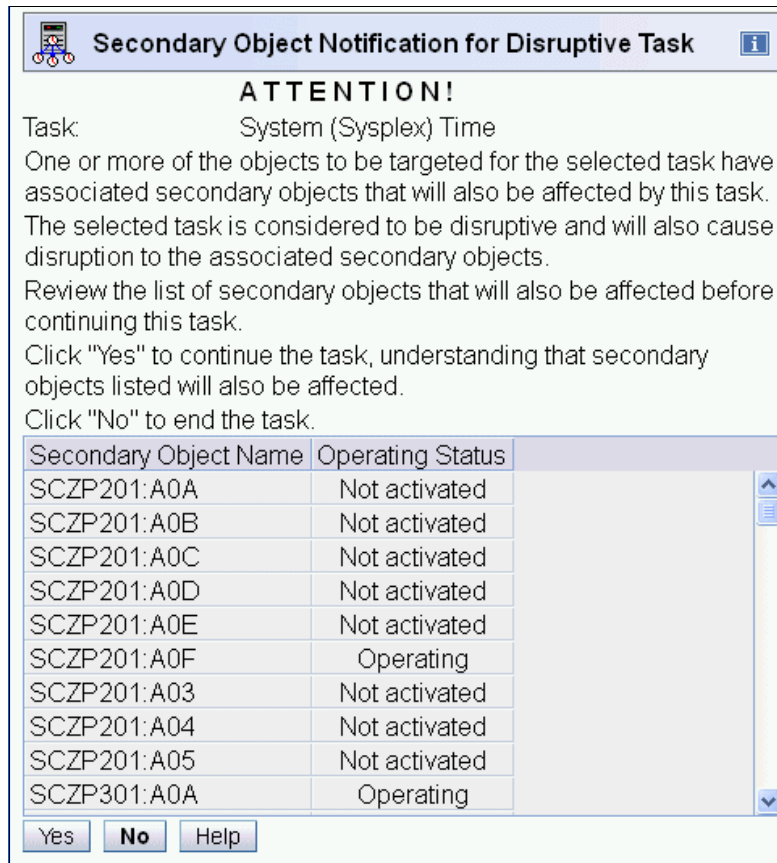


Figure 3-2 Secondary Object Notification for Disruptive Task

Note: This notification window serves as a reminder that changes made in this task can potentially be disruptive to the secondary object names listed.

However, if no changes are made in the Sysplex Timer related tasks (view only) there is no risk of disruption.

After selecting **Yes**, you will see the panel displayed in Figure 3-3.

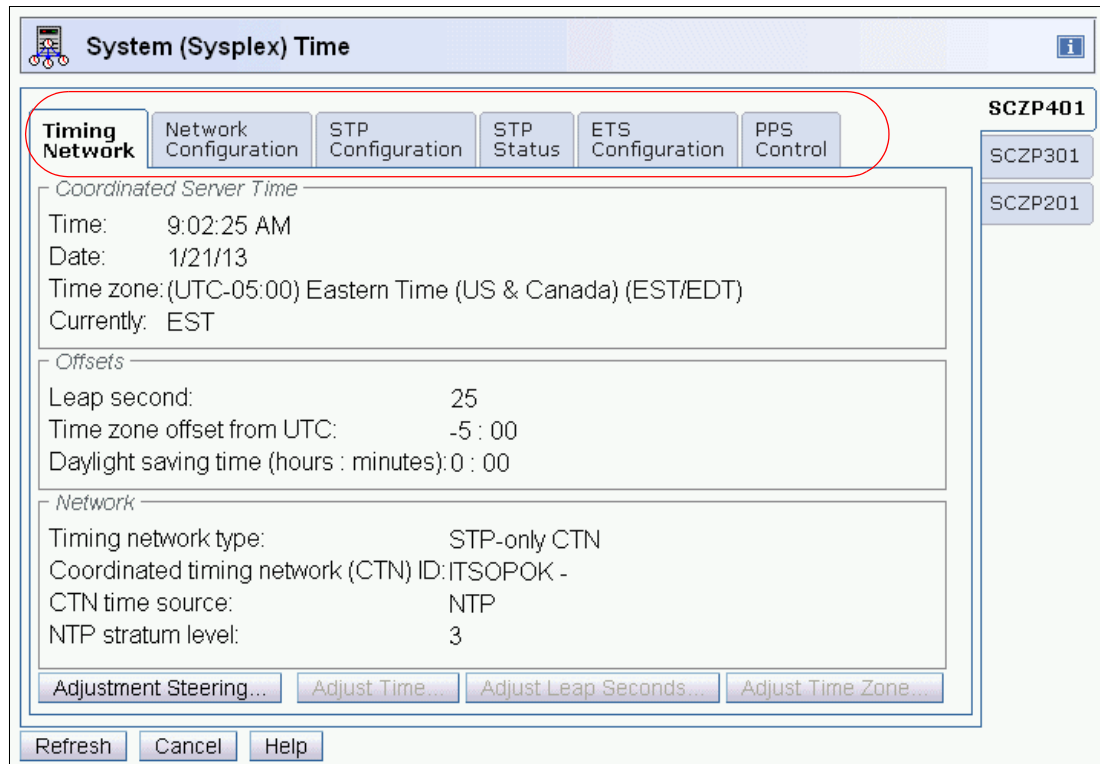


Figure 3-3 System (Sysplex) Time panel for single-selected CPC: Timing Network tab

The number of tabs on the panel (Figure 3-3) varies depending on the type of server and what features are installed on it, as shown in Table 3-1. If your environment consists of different server types, ensure that your operations staff recognizes the differences when displaying these panels.

Table 3-1 The tabs displayed on the server

Server	Timing network	Network config	ETR config	ETR status	STP config	STP status	ETS config	PPS control
zEC12 (STP only)	✓	✓			✓	✓	✓	✓
z196 and z114 (STP only)	✓	✓	✓		✓	✓	✓	✓
z9/z10 ^a with ETR and STP features	✓	✓	✓	✓	✓	✓	✓	
z990/z890 with ETR and STP features	✓	✓	✓	✓	✓	✓		
z9 with STP feature only	✓	✓			✓	✓	✓	
z990/890 with STP feature only	✓	✓			✓	✓		
All Servers with ETR feature only			✓	✓				

a. z10 includes the ETR feature by default.

Timing network type operations summary

Table 3-2 summarizes what operation is performed from each tab in the various configurations. For details about how to perform a particular operation on a tab, refer to Part 2 of the *Server Time Protocol Implementation Guide*, SG24-7281.

Table 3-2 STP tabbed display operations summary

Type	Timing network	Network configuration	ETR configuration	ETR status	STP configuration	STP status	ETS configuration
ETR	Display only	N/A	Manage ports and ETR ID	Display only	Join a Mixed CTN	Display only	N/A
Mixed CTN	Display only	N/A	Manage ports and ETR ID	Display only	Join a Mixed CTN, change to STP, or change to ETR ^a	Display only	N/A
STP-only CTN	Time/offset adjustments (enabled on CTS only) ^b	Role assignment ^c (enabled on CTS only)	N/A	Display only ^d	Join an STP CTN or change to Mixed CTN ^e	Display only	Define ETS (PTS and BTS)

- a. Changing to a different Mixed CTN is disruptive for z/OS.
- b. This includes Adjustment Steering (display only), Adjust Time and Adjust Time zone, and Leap second offsets.
- c. This is a global change to the CTN.
- d. N/A for zEnterprise Servers.
- e. Changing to a different STP-only CTN is disruptive for z/OS.

In the following sections we provide a brief description of the seven tabs. The most relevant tabs in this planning phase are the Timing Network, ETR Status, and STP Status. The seven tabs are covered in detail in the migration chapters in *Server Time Protocol Implementation Guide*, SG24-7281.

Timing Network tab

The information shown on the Timing Network tab (Figure 3-3 on page 80) displays overall timing information for the CTN, including the current date and time, local time offsets, and general timing network information. The information displayed here is identical on each server in the CTN.

The Coordinated Server Time represents the date and time for the CTN. The Offsets section displays the leap second offset and other offset information that it varies depending on whether it is an ETR network, Mixed or STP-only CTN, and whether a time zone offset has been defined.

The Network section displays general information about the CTN, including timing network type, CTN ID, and CTN time source. If the CTN time source is an NTP server, as shown in Figure 3-3 on page 80, its NTP stratum level and source ID also will be displayed.

Network Configuration tab

The information shown on the Network Configuration tab depends on the network type configured. By looking on the Network Configuration tab for an existing configuration, you can tell whether you are in an ETR network, Mixed CTN, or STP-only CTN. Figure 3-4 shows the Network Configuration tab for the SCSZP201 server, which is not part of any timing network.¹

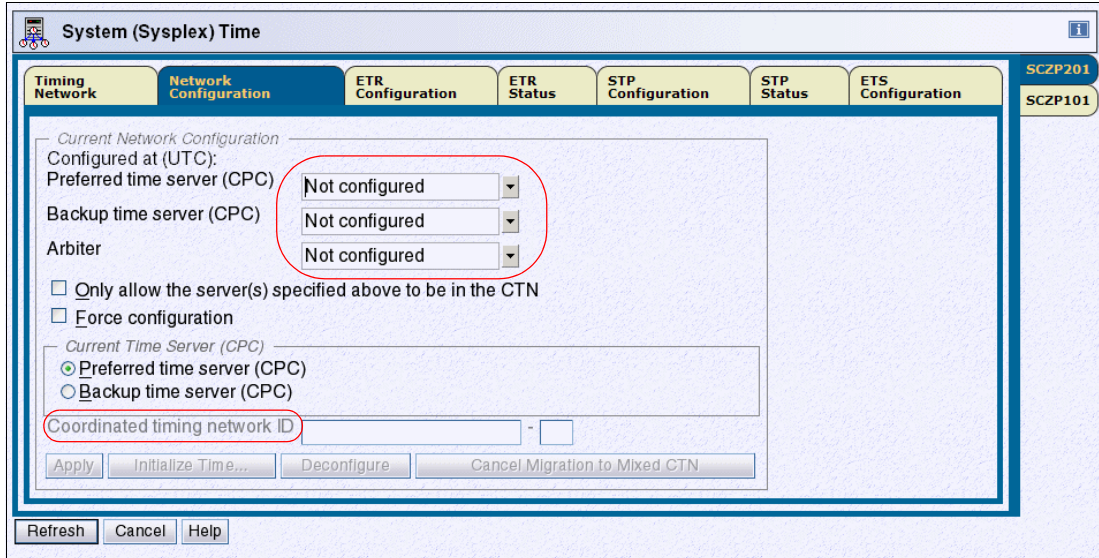


Figure 3-4 Network Configuration tab: No timing network configured

ETR Configuration tab

This tab allows users to view and modify the ETR network ID and ETR ports status (Figure 3-5). Information and data on this panel is not applicable to STP-only CTN.

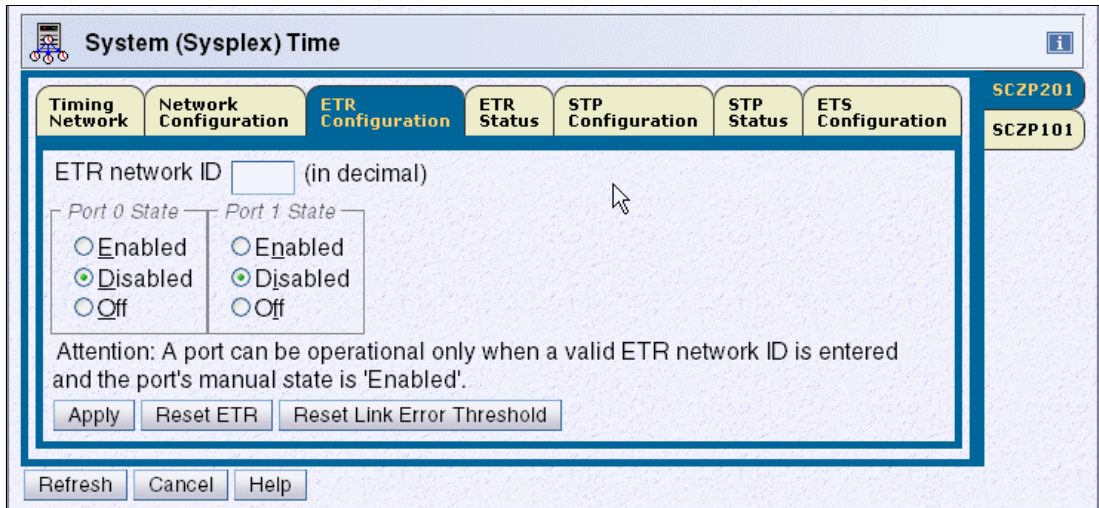


Figure 3-5 ETR Configuration tab

¹ This panel is for GLOBAL changes (affecting all servers in the CTN).

ETR Status tab

The ETR Status tab is used to display the configuration and operational state of the connections to the Sysplex Timer. The ETR Status tab displays information about the connected Sysplex Timer units and the server's TOD clock. Figure 3-6 shows an example of the ETR Status tab for a server in STP timing mode.

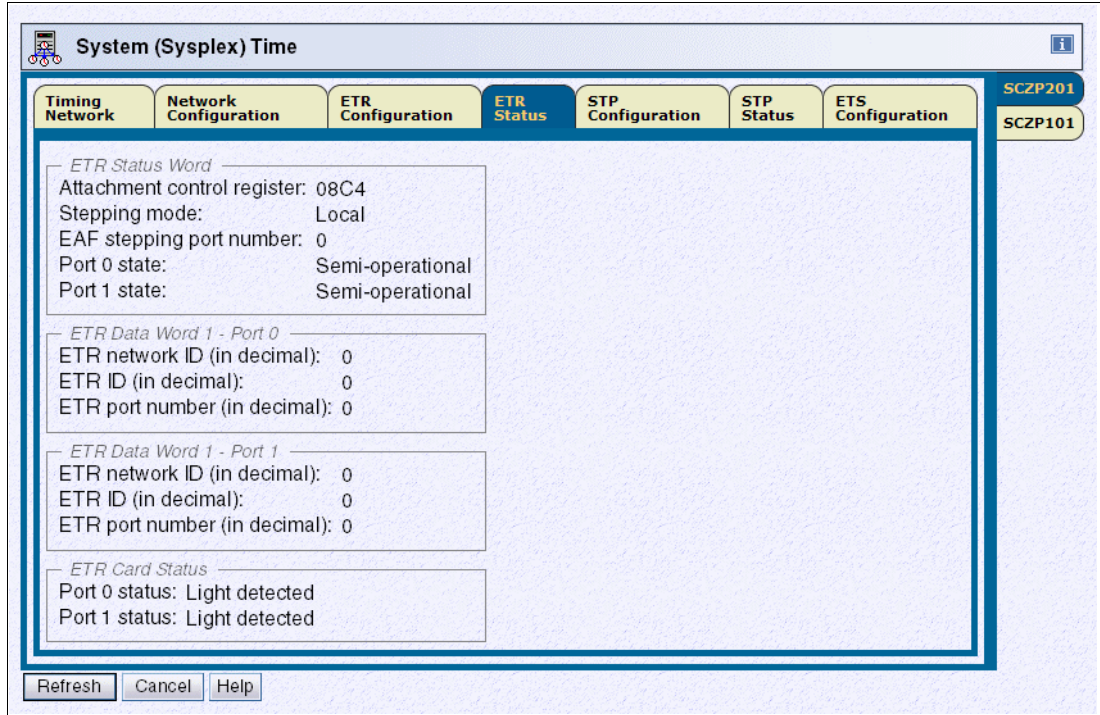


Figure 3-6 System (Sysplex) Time panel for multiple-selected CPCs: ETR Status tab

Note: If the ETR feature is actively connected to an NTP server through the PPS port, status information is not applicable.

For each ETR port connected to a Sysplex Timer, the following information received from the attached unit is displayed:

- ETR network ID** Network ID of the Sysplex Timer to which the ETR port is connected
- ETR ID** Identifies the Sysplex Timer that sends the ETR data
- ETR port number** Port number of the Sysplex Timer output port that sends the ETR data

If the configuration has been properly set, both ports display the same ETR network ID as the value entered in the ETR Configuration tab.

Important: If the values displayed are not identical, the state indicates semi-operational.

The ETR network ID, ETR ID, and ETR port number data is for information only. The values are obtained from the Sysplex Timer unit and can only be modified from the Sysplex Timer Console application.

If you are planning for a migration from ETR to Mixed-CTN or Mixed-CTN to STP-only, use this panel to map out the connectivity between your server and the Sysplex Timers.

STP Configuration tab

The STP Configuration tab allows you to enter and modify the timing network name, depending on the changes that you want to make (Figure 3-7).²

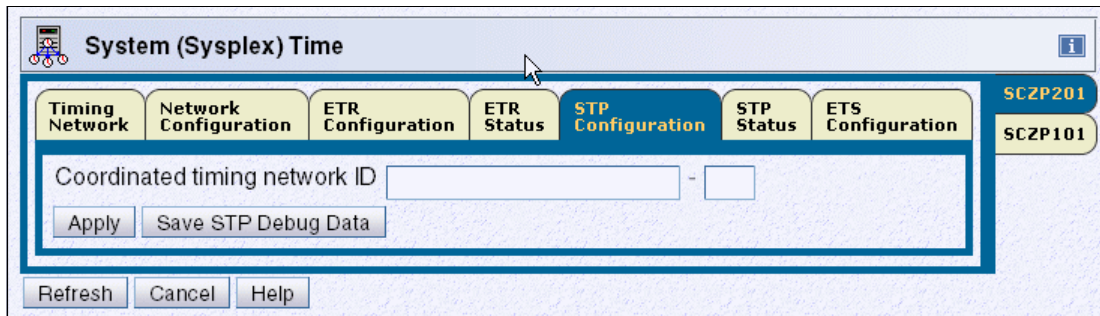


Figure 3-7 STP configuration tab

STP Status tab

The STP Status tab on the System (Sysplex) Time panel displays STP information from the point of view of the selected server. This tab can be used to identify the topology of the CTN from the perspective of the selected server. Figure 3-8 shows an example of the STP Status tab for a server in STP timing mode.

The timing state is of particular importance. It can be unsynchronized, synchronized, or stopped. If this value is anything other than synchronized then this server is not actively participating in the CTN.

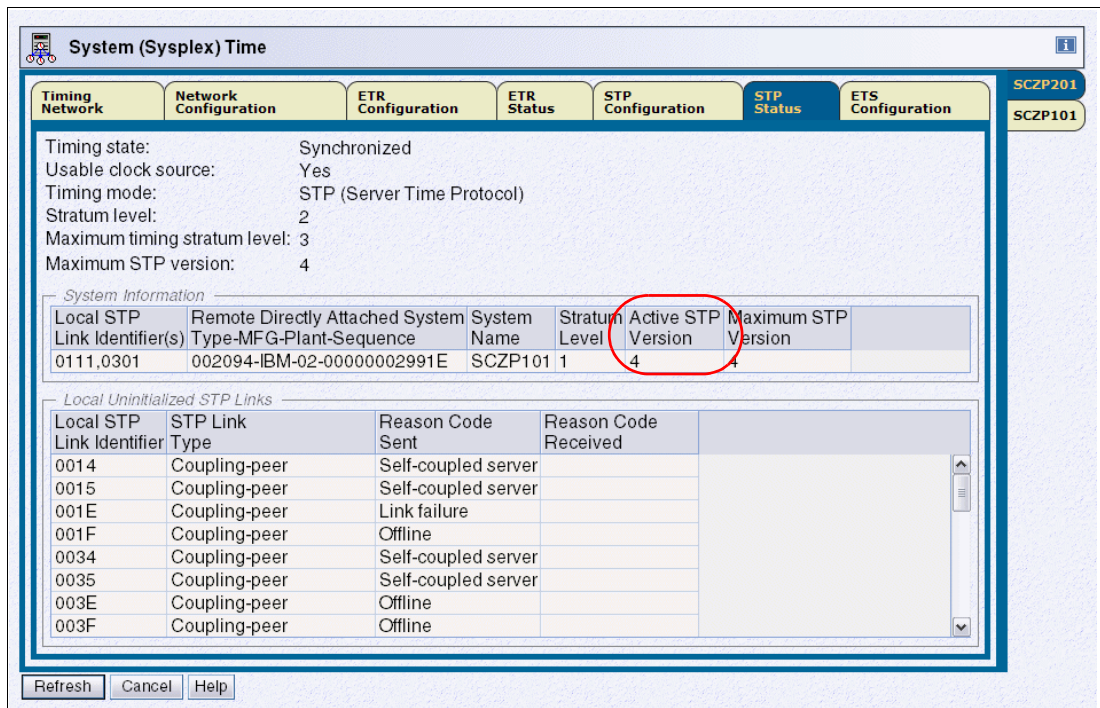


Figure 3-8 STP Status tab

The System Information section identifies the remote servers (CPCs) that are directly attached to this server for STP purposes. As soon as the STP portion of the CTN ID is applied

² This panel is for LOCAL changes (affecting only the selected server).

to this server, coupling links that are initialized to transport STP messages *for this CTN ID* are listed using the PCHID addresses and are grouped according to the system that is directly attached to the links. Additionally, the stratum level, active STP version, and maximum STP version for each directly attached system are shown. The information provided in this section can be used to build a topology diagram of the CTN.

Servers with different maximum STP versions can coexist in the same CTN. When two attached servers are operating different STP versions, communication between the two servers uses the lowest version number installed. The version number used is indicated in the active STP version column.

Tip: In conjunction with the Timing Network tab, a CTN topology diagram can be built using the information from the stratum level of the current server and the stratum level of each of the directly attached servers.

The Local Uninitialized STP Links table identifies all coupling links defined in the active IOCDs that can be used by this server. Each entry in this table represents a link that is currently in an uninitialized state. The Reason Codes sent and received provide further information to assist in determining why the link is not initialized for STP.

ETS Configuration tab

On the ETS Configuration tab you can configure your ETS for an STP-only CTN, for both PTS and BTS servers.

Although by default only CTS/PTS adjusts CST to the ETS, we recommend that the BTS also be ETS configured. If an ETS is configured, the ETS device is monitored and errors are logged. If an ETS is configured as NTP or NTP with pulse per second (PPS), it provides redundancy of the ETS for the STP-only CTN.

PPS Control tab

On zEnterprise servers, through the PPS Control tab, you can test and control the use of the PPS signal (if PPS port is connected to an NTP server with an active PPS port), as shown in Figure 3-9.

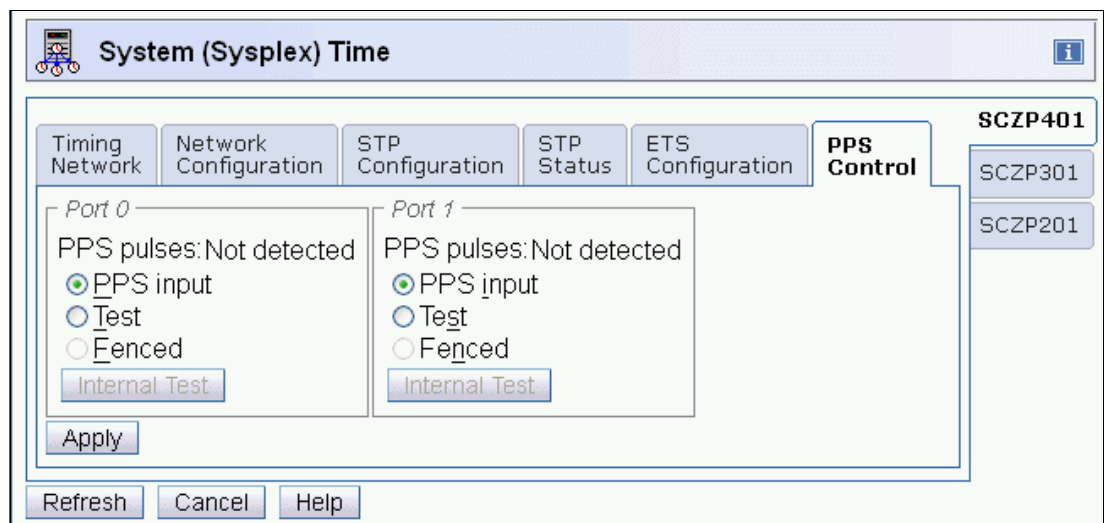


Figure 3-9 PPS Control tab (zEnterprise servers only)

3.1.3 z/OS display commands

There are a number of z/OS commands that display information about the timing mode or the status of the CTN as viewed by z/OS itself. The output from these commands can change depending on the CTN type and the timing mode that the server is currently using.

DISPLAY ETR

Prior to STP, the DISPLAY ETR command was used to display the synchronization mode and the status of the ETR ports as seen by z/OS. With STP support, the command itself has not changed. However, the output has been updated to display STP-related information where applicable.

In an ETR network

When the server is in an ETR network, it is in ETR timing mode synchronized to a Sysplex Timer. Because no CTN ID has been defined, all of the information is ETR-related. See Example 3-1.

Example 3-1 z/OS DISPLAY ETR in an ETR network

```
-D ETR
IEA282I 09.11.21 TIMING STATUS 010
SYNCHRONIZATION MODE = ETR
  CPC PORT 0 <== ACTIVE      CPC PORT 1
  OPERATIONAL                OPERATIONAL
  ENABLED                    ENABLED
  ETR NET ID=31              ETR NET ID=31
  ETR PORT=21                ETR PORT=21
  ETR ID=13                  ETR ID=12
```

In a Mixed CTN on a server in ETR timing mode

If a server in a Mixed CTN is synchronized to the Sysplex Timer, it is in ETR timing mode and is a stratum 1 in the Mixed CTN. In this case, both Sysplex Timer-related information and CTN ID information are displayed. The format is similar to that in Example 3-1. However, an extra line is added to display the CTN ID, as shown in Example 3-2.

Example 3-2 z/OS DISPLAY ETR in a Mixed CTN with ETR timing mode

```
D ETR
IEA282I 14.48.26 TIMING STATUS 951
SYNCHRONIZATION MODE = ETR
  CPC PORT 0 <== ACTIVE      CPC PORT 1
  OPERATIONAL                OPERATIONAL
  ENABLED                    ENABLED
  ETR NET ID=31              ETR NET ID=31
  ETR PORT=01                ETR PORT=01
  ETR ID=01                  ETR ID=00
THIS SERVER IS PART OF TIMING NETWORK ITSOP0K -31
```

In a Mixed CTN on a server in STP timing mode

It is also possible to have servers in a Mixed CTN that are synchronized using STP messages with no active connection to a Sysplex Timer. They are stratum 2 or stratum 3 servers in STP timing mode.

Example 3-3 shows that the server is in a Mixed CTN, because the CTN ID contains both STP ID and ETR network ID components.

Example 3-3 z/OS DISPLAY ETR in a Mixed CTN with STP timing mode

```
-D ETR
IEA386I 14.40.19 TIMING STATUS 337
SYNCHRONIZATION MODE = STP
  THIS SERVER IS A STRATUM 2
  CTN ID = ITSOP0K -31
  NUMBER OF USABLE TIMING LINKS = 4
```

STP-only CTN, all servers in STP timing mode

In an STP-only CTN the TOD clock is being steered to the time provided by the Current Time Server (CTS). No reference to a Sysplex Timer is displayed. See Example 3-4.

Example 3-4 z/OS DISPLAY ETR in an STP-only CTN

```
-D ETR
IEA386I 11.25.04 TIMING STATUS 787
SYNCHRONIZATION MODE = STP
  THIS SERVER IS A STRATUM 3
  CTN ID = ITSOP0K
  THE STRATUM 1 NODE ID = 002097.S18.IBM.02.000000099999
  NUMBER OF USABLE TIMING LINKS = 4
  THIS STP TIMING NETWORK HAS NO SERVER TO ACT AS ARBITER
```

The DISPLAY ETR output in an STP-only CTN incorporates additional information in comparison to STP mode within a Mixed CTN. Node ID information for the CTS is displayed, in addition to a number of extra optional lines that show a different output depending on the CTN topology and whether this server has been assigned a Preferred Time Server, Backup Time Server, or Arbiter Server role. More DISPLAY ETR commands are illustrated in the following section.

DISPLAY ETR examples

Determining the Current Time Server can be done by using the DISPLAY ETR command and checking which server is performing the stratum 1 role, as shown in Example 3-5.

Example 3-5 z/OS DISPLAY ETR command: identifying CTS (stratum 1)

```
D ETR
IEA386I 14.50.16 TIMING STATUS 931
SYNCHRONIZATION MODE = STP
  THIS SERVER IS A STRATUM 2
  CTN ID = ITSOP0K
  THE STRATUM 1 NODE ID = 002097.E26.IBM.02.000000088888
  THIS IS THE BACKUP TIME SERVER
```

Example 3-6 shows checking CTN roles using the z/OS DISPLAY ETR command. Observe that there is currently no Backup Time Server configured and, implicitly, there is no Arbiter.

Example 3-6 z/OS command DISPLAY ETR: no BTS configured

```
D ETR
IEA386I 15.12.27 TIMING STATUS 975
SYNCHRONIZATION MODE = STP
  THIS SERVER IS A STRATUM 1
  CTN ID = hmctest
  THE STRATUM 1 NODE ID = 002097.E26.IBM.02.00000001DE50
```

```
THIS IS THE PREFERRED TIME SERVER
THIS STP NETWORK HAS NO BACKUP TIME SERVER
THIS STP NETWORK HAS NO SERVER TO ACT AS ARBITER
```

Once a Backup Time Server has been configured, the absence of the line pertaining to the Backup Time Server indicates that the CTN now has a BTS configured. However, the CTN has no Arbiter defined, as shown in Example 3-7.

Example 3-7 z/OS command DISPLAY ETR: CTN has PTS and BTS, no Arbiter

```
D ETR
IEA386I 15.15.43 TIMING STATUS 982
SYNCHRONIZATION MODE = STP
  THIS SERVER IS A STRATUM 1
    CTN ID = hmctest
    THE STRATUM 1 NODE ID = 002097.E26.IBM.02.00000001DE50
    THIS IS THE PREFERRED TIME SERVER
    THIS STP NETWORK HAS NO SERVER TO ACT AS ARBITER
```

As with the Backup Time Server, issue the z/OS DISPLAY ETR command to determine whether the STP-only CTN has an Arbiter defined. Absence of the line pertaining to the Arbiter indicates that the Arbiter has been configured, as shown in Example 3-8.

Example 3-8 z/OS command DISPLAY ETR: both BTS and Arbiter configured

```
D ETR
IEA386I 15.18.08 TIMING STATUS 993
SYNCHRONIZATION MODE = STP
  THIS SERVER IS A STRATUM 1
    CTN ID = hmctest
    THE STRATUM 1 NODE ID = 002097.E26.IBM.02.00000001DE50
    THIS IS THE PREFERRED TIME SERVER
```

The number of usable timing links account for the number of STP-initialized links that supply *or can supply* timing information to this server. That is why this line does not appear in the D ETR output of a server in ETR timing mode or a Current Time Server because it is the source of timing information. A stratum 2 server includes links to another stratum 2 server since if it transitions to a stratum 3 server those links become sources of timing signals.

SETETR PORT=n

The SETETR command can only be used to enable Sysplex Timer ports that have been previously disabled by z/OS as a consequence of hardware error. A Sysplex Timer port disabled by a hardware problem can be enabled after the problem has been corrected.

The syntax for the SETETR command is SETETR PORT=*n*, where *n* specifies the number of the ETR port to be enabled (Example 3-9).

Example 3-9 SETETR command

```
-D ETR
IEA282I 09.34.52 TIMING STATUS 031
SYNCHRONIZATION MODE = ETR
  CPC PORT 0 <== ACTIVE      CPC PORT 1
  OPERATIONAL                OPERATIONAL
  ENABLED                    DISABLED
  ETR NET ID=31              ETR NET ID=31
  ETR PORT=21                ETR PORT=21
  ETR ID=13                  ETR ID=12
  THIS SYSTEM IS PART OF TIMING NETWORK ITSOP0K -31
```

```
-SETETR PORT=1
IEA283I ETR PORT 1 IS ENABLED
```

The SETETR command cannot be used when the server is in STP timing mode (Example 3-10).

Example 3-10 z/OS SETETR in STP timing mode

```
IEA384I SETETR COMMAND IS NOT VALID IN STP TIMING MODE
```

DISPLAY XCF,SYSPLEX,ALL

The DISPLAY XCF,SYSPLEX,ALL command displays the system status and the last recorded system status monitor time stamp for each system in the sysplex (Example 3-11).

Example 3-11 z/OS DISPLAY XCF,SYSPLEX,ALL in a Mixed CTN

```
D XCF,S,ALL
IXC335I 17.31.31 DISPLAY XCF 375
SYSTEM  TYPE SERIAL LPAR STATUS TIME      SYSTEM STATUS
SC80    2097 8888 01   11/22/2007 17:31:28 ACTIVE      TM=STP
SC74    2094 9999 01   11/22/2007 17:31:29 ACTIVE      TM=STP
SC75    2084 7777 15   11/22/2007 17:31:30 ACTIVE      TM=ETR
```

The output from this command incorporates information about the timing mode in effect.

DISPLAY CF

The DISPLAY CF command does not directly provide information regarding the CTN type or timing mode of the server. However, the output does display the CF Request Time Ordering status. If the CTN is a Mixed or STP-only CTN in a Parallel Sysplex configuration, the requirement is that all servers support CF Request Time Ordering.

Note: CF Request Time Ordering is also referred to as the Message Time Ordering Facility (MTOF).

The DISPLAY CF command can be used to verify whether CF Request Time Ordering is required and enabled, as shown in Example 3-12.

Example 3-12 z/OS DISPLAY CF

```
D CF
IXL150I 14.55.46 DISPLAY CF 975
COUPLING FACILITY 002094.IBM.02.000000088888
                    PARTITION: 2C CPCID: 00
                    CONTROL UNIT ID: FFFC
NAMED CF7B
COUPLING FACILITY SPACE UTILIZATION
. . . . .

CFCC RELEASE 15.00, SERVICE LEVEL 02.09
BUILT ON 04/30/2009 AT 14:18:00
COUPLING FACILITY HAS 1 SHARED AND 0 DEDICATED PROCESSORS
DYNAMIC CF DISPATCHING: ON
```

CF REQUEST TIME ORDERING: REQUIRED AND ENABLED

.

STP requires that each coupling facility within a Parallel Sysplex is enabled for CF Request Time Ordering before any server within the Parallel Sysplex can be defined in a Mixed or STP-only CTN.

Example 3-13 displays the other CF Request Time Ordering messages that can appear in the output of the DISPLAY CF.

Example 3-13 CF Request Time Ordering messages

```
CF REQUEST TIME ORDERING: NOT-REQUIRED AND ENABLED
CF REQUEST TIME ORDERING: NOT-REQUIRED AND NOT-ENABLED
CF REQUEST TIME ORDERING: REQUIRED AND NOT-ENABLED
REASON: FUNCTION NOT INSTALLED ON THIS SYSTEM
REASON: ETR NOT CONNECTED TO COUPLING FACILITY
REASON: REQUEST TIME ORDERING FUNCTION FAILURE
REASON: ETR NETID MISMATCH - CF ETR NETID: etr netid
REASON: CTNID MISMATCH - CF CTNID: ctnid
```

There are two time-related CFCC commands, which are described in “Coupling facility commands” on page 117.

3.1.4 z/OS messages

Ensure that your operations staff reviews all STP-related messages and plan for which ones they would like automation (or console operator staff) to take action on. Note that some messages will be issued on every member of the sysplex with STPMODE YES specified, which might cause automation to take multiple or redundant actions.

Refer to Appendix B, “Server Time Protocol (STP) messages in z/OS” on page 199, for all STP-related messages.

New z/OS messages in April 2009 STP enhancements

Prior to z/OS 1.11, messages related to the following STP events were only posted as hardware messages on the HMC:

- ▶ ETS failure or status change events
- ▶ CTS role changing from PTS to BTS or vice versa because of STP recovery

To improve the delivery of important information to the operator and to better integrate with system automation tools, z/OS 1.11 (also rolled back to 1.9 and 1.10 with PTFs) adds new STP-related messages in addition to the hardware messages posted on the HMC.

Note: These messages will be posted by every z/OS sysplex member using STPMODE YES in the CTN. The operations staff or automation tools can monitor these messages even if there are stand-alone CFs in the CTN assigned the special roles of PTS and BTS.

To address the CTS change a new z/OS message was created:

```
IEA395I THE CURRENT TIME SERVER HAS CHANGED TO THE cccccc
```

Where *ccccc* is BACKUP or PREFERRED. This informational message does not require any action, but ensures that the operational staff responsible for STP is aware of the change.

When an ETS was configured with the Sysplex Timers it provided a means to notify operations of ETS-related alerts, for example:

```
IEA272I ETR SERVICE IS REQUESTED. REASON CODE=045
```

Where 45 is The external time source is not responding”.

The following new message is created for all ETS-related alerts if an ETS is configured with STP:

```
IEA031I STP ALERT RECEIVED. STP ALERT CODE=cc
```

Refer to Appendix B, “Server Time Protocol (STP) messages in z/OS” on page 199, for a list of STP-related z/OS messages.

z/OS messages context

z/OS messages might help you identify various configuration issues and status changes for the servers in your CTN. However, consider planning your servers roles and changing your CTN configuration considering the following aspects:

- ▶ Although stand-alone CFs will typically provide best connectivity to other servers requiring time synchronization, stand-alone CFs do not produce z/OS messages for operations or interception by automation routines. For example, the following information messages at IPL or interrupt time might not be displayed:

```
IEA380I THIS SYSTEM IS NOW OPERATING IN STP TIMING MODE.
```

```
IEA381I THE STP FACILITY IS NOT USABLE. SYSTEM CONTINUES IN LOCAL TIMING MODE
```

- ▶ If the condition being raised relates to connectivity between two servers, the information might be available to a z/OS system image at the other end of the link. However, if both ends of the link are CF partitions, no warning message is available to the user.

- ▶ There are several IEAxxx and IXCxxx messages that report current and changed timing status. As an example, the following message reports the result of a successful migration from a mixed CTN to an STP-only CTN:

```
IXC438I COORDINATED TIMING INFORMATION HAS BEEN UPDATED
FOR SYSTEM: scpz101
PREVIOUS CTNID: ITSOPOK-31
CURRENT CTNID: ITSOPOK
```

- ▶ In general, there are no z/OS messages that are posted only on the PTS, BTS, or Arbiter.
 - Certain messages do not appear on the CTS since it is the time source:

```
IEA382I THIS SERVER HAS ONLY A SINGLE LINK AVAILABLE FOR TIMING PURPOSES.
IEA383I THIS SERVER RECEIVES TIMING SIGNALS FROM ONLY ONE OTHER NETWORK
NODE.
IEA281I STP SYNC CHECK THRESHOLD EXCEEDED. CPC CONTINUES IN LOCAL MODE.
IEA390I TOD CLOCKS DYNAMICALLY ADJUSTED TO MAINTAIN STP SYNCHRONISM.
```

- The following message might not appear on certain special role servers:

```
IEA388I THIS SERVER HAS NO CONNECTION TO THE nnnnnnnnnn
Where nnnnnnnnnnn = 'PREFERRED ' | 'BACKUP ' | 'ARBITER '
```

- ▶ For example, the following message never appears on a z/OS system running on the BTS:


```
IEA388I THIS SERVER HAS NO CONNECTION TO THE BACKUP
```

3.2 Configuring the Coordinated Timing Network

This section describes how to configure either a Mixed CTN or an STP-only CTN.

3.2.1 Configuring a Mixed CTN

In this section we discuss configuring a Mixed CTN.

Setting the CTN ID

Applying the STP Network ID as the first portion of the CTN ID on an STP-enabled server activates a Mixed CTN. The STP Network ID is entered on the STP Configuration tab as illustrated in Figure 3-10. The same CTN ID must be entered on every server that will participate in the Mixed CTN.

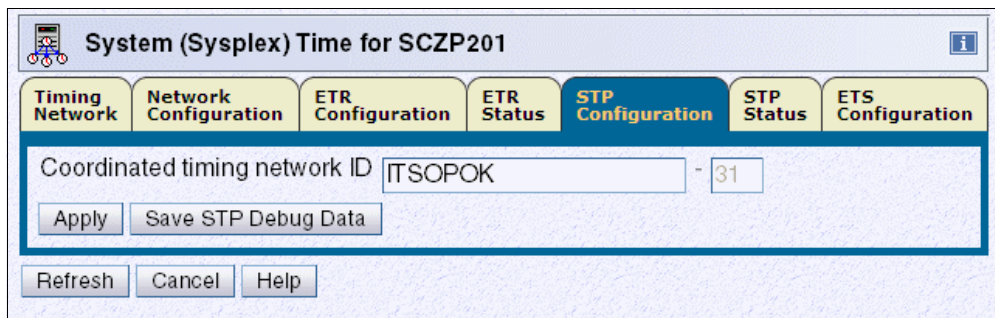


Figure 3-10 System (Sysplex) Time: STP Configuration tab

The response to the z/OS DISPLAY ETR command, message IEA282I now includes the CTN ID (Example 3-14).

Note: The menu shown in Table 3-9 on page 124 is for a server that has been previously configured in an ETR network. If you add a new server to a Mixed CTN you must enter both the CTN ID and the ETR ID.

Example 3-14 D ETR command and response

```

D ETR
IEA282I 14.39.57 TIMING STATUS 292
SYNCHRONIZATION MODE = ETR
  CPC PORT 0  <== ACTIVE      CPC PORT 1
  OPERATIONAL                                OPERATIONAL
  ENABLED                                       ENABLED
  ETR NET ID=31                               ETR NET ID=31
  ETR PORT=05                                 ETR PORT=05
  ETR ID=00                                   ETR ID=01
  THIS SERVER IS PART OF TIMING NETWORK ITSOPOK -31
  
```

Because the Sysplex Timer is the time source in a Mixed CTN, at least two servers or CFs should be in ETR Timing mode (stratum 1) in order to avoid a single point of failure.

For the same reason, if stratum 2 servers or CFs are configured in the Mixed CTN, each should be connected to at least two stratum 1 servers or CFs, with at least two coupling links to each stratum 1 server or CF.

The same recommendation applies to any planned stratum 3 servers in the Mixed CTN. Each stratum 3 server should be connected to at least two stratum 2 servers or CFs.

When more than two servers or CFs are configured in a Mixed CTN, servers beyond the two recommended stratum 1 servers can be individually configured in STP timing mode (stratum 2). Changing the timing mode for a given server or CF is done by disabling its ETR ports using the **System (Sysplex) Time** \emptyset **ETR Configuration** tab, as shown in Figure 3-11. Enabling the ETR ports switches the server or CF back to ETR Timing mode.

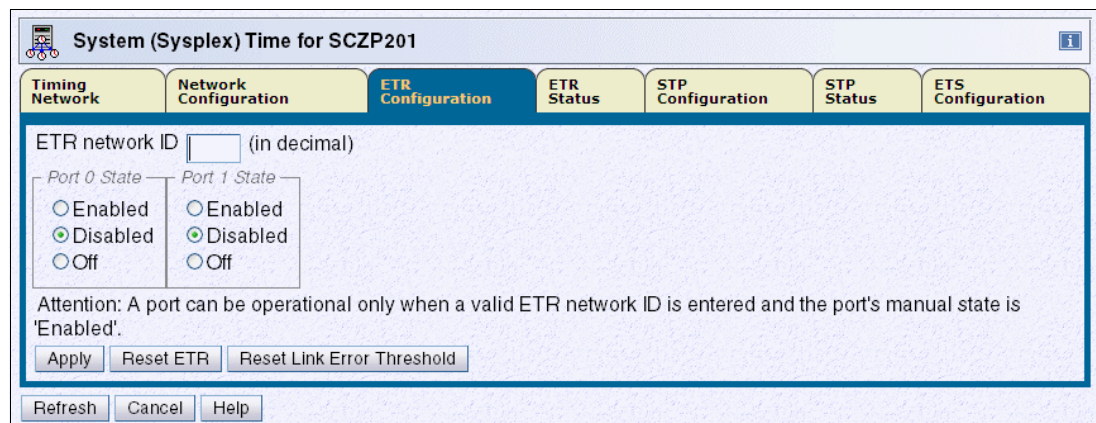


Figure 3-11 ETR Configuration tab: port enabled or disabled

3.2.2 Configuring an STP-only CTN

The sequence of steps in configuring an STP-only Coordinated Timing Network depends on whether it is a new STP-only installation or a migration from a Mixed to an STP-only CTN. The overall procedure consists of the following steps:

1. Configure an external time source (optional).

This optionally applies to both a new install and a migration from a Mixed to an STP-only CTN.

2. Set the CTN ID.

This applies only to an STP-only new installation.

3. Initialize the time.

This is only necessary for an STP-only new install. It includes:

- Set time zone.
- Set leap seconds.
- Set date and time.

4. Activate the STP-only CTN.

This is necessary for both an STP-only new install and a migration from a Mixed to an STP-only CTN.

Configuring the external time source (ETS)

To maintain time accuracy, both the Sysplex Timer and STP support connection to an external time source. Using an ETS, regular adjustments can be performed either manually or automatically.

In a Mixed CTN, if the Sysplex Timer has an external time source configured, it will be used and timing accuracy will be maintained as per the ETR network that was in place prior to the migration.

In an STP-only CTN the Sysplex Timer is no longer used. If an external time source is to be used, it must first be configured.

The System z10 and System z9 with driver 67L with the latest MCL level offer three choices for the ETS:

- ▶ A dial-out function with an HMC connected modem
- ▶ Synchronization to an NTP server connected to the Support Element LAN
- ▶ Synchronization to an NTP server with pulse per second (PPS) output connected to the SE LAN and the ETR card

The System z9 pre-driver 67L, z990, and z890 can only use the dial-out function, but can participate in an STP-only CTN where the CTS is synchronized to an NTP server.

Dial out from Hardware Management Console

Figure 3-12 shows the ETS Configuration tab from the System (Sysplex) Time task when the **Use dial out if configured on Hardware Management Console** radio button is selected.

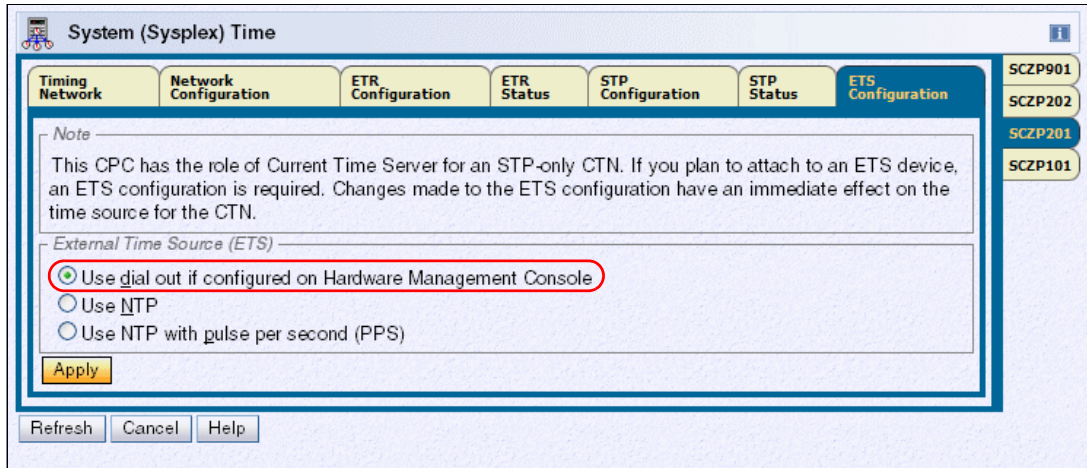


Figure 3-12 ETS configuration tab: use dial out on the HMC selected

The HMC dials out to the time service selected. To configure the dial-out telephone connection on the HMC select **Service Management group** \emptyset **Customize Outbound Connectivity**. Check the **Enable local system as call-home server** box.

On the External Time Source tab, check **Allow external time source dialing using the local modem** and configure the protocol (Figure 3-13).

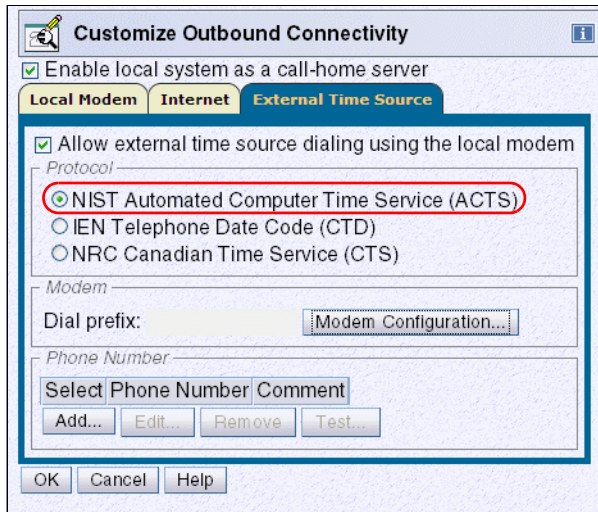


Figure 3-13 Customize Outbound Connectivity: External Time Source tab

To perform adjustments to the dial-out ETS automatically on a scheduled basis, the scheduled operations for the specific server must be customized. The function is established on the Hardware Management Console using customize scheduled operations in the Operational Customization menu. Selecting **New** on the Option tab guides you to a selection menu, shown in Figure 3-14.

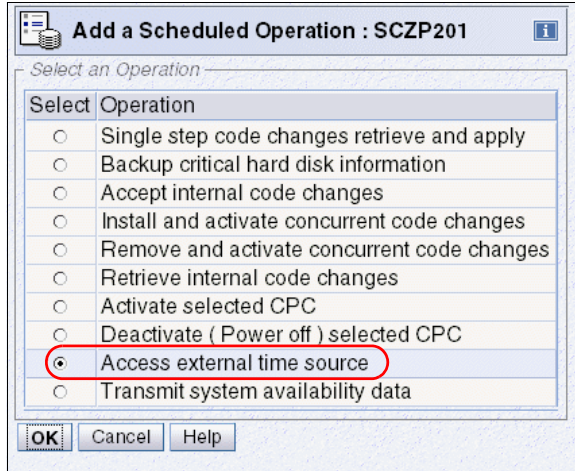


Figure 3-14 Add a Scheduled Operation for external time source access

On the Access external time source panel, specify when the operation is to take place. Set up a repeated scheduled operation. Depending on the time accuracy requirement, select one or more days that the ETS should be contacted. Use the option to repeat indefinitely to make sure that the scheduled operation never expires.

Note: In the case of an ETS dial-out failure a hardware message will be posted on the HMC.

Network Time Protocol (NTP) server

Selecting the **Use NTP** radio button displays the NTP Time Server Information table, shown in Figure 3-15. When at least one Configured check box is selected, it indicates that the IP address or web address entered should be used as the NTP server address.

System (Sysplex) Time for SCZP101

ETS Configuration

Note
This CPC does not have the role of Current Time Server for an STP-only CTN. However, it is configured to become the Current Time Server, if necessary. If you plan to attach to an ETS device, an ETS configuration is recommended. If an ETS is configured, the ETS device is monitored and errors are logged. If an ETS is configured as NTP or NTP with pulse per second (PPS), it provides redundancy of the ETS for the STP-only CTN. Changes made to the ETS configuration for this CPC may affect the time source for the CTN.

External Time Source (ETS)

Use dial out if configured on Hardware Management Console

Use NTP

Use NTP with pulse per second (PPS)

NTP Time Server Information

Select	Configured	NTP Time Server	Stratum	Source	Status
<input checked="" type="radio"/>	<input checked="" type="checkbox"/>	10.1.6.48	2	128.138.140.44	Success
<input type="radio"/>	<input checked="" type="checkbox"/>	10.1.6.46	2	128.138.140.44	Success

Apply Query

Refresh Cancel Help

Figure 3-15 ETS Configuration tab: NTP w/o PPS

The use of a web address requires that the SE be customized with Domain Name Services (DNS) enabled. Use the task **Customize Network Settings** \emptyset **Name services** in the Support Element Console Application. Refer to the manual *Support Element Operations Guide* for each server for information about how to customize network settings for the SE.

The **Query** button tests the IP connectivity and fills in the Stratum, Source, and Status table fields for the corresponding NTP server. The NTP configuration takes effect when **Apply** is clicked.

The Source field contains a description of the time source for the NTP server provided as information by the NTP server.

When two NTP servers are configured, it is the user's responsibility to choose the preferred NTP server by selecting the appropriate **Select** radio button. This NTP server is called the selected NTP server. The other NTP server is called the non-selected NTP server. The SNTP client compares the quality of both NTP servers and informs the user in case the selected NTP server has a stratum level that is lower in the hierarchy than the non-selected NTP server (a stratum 1 server is usually a better choice than a stratum 2, and so forth), or if the time obtained from the selected NTP server is less accurate than the non-selected NTP server.

Network Time Protocol (NTP) server with pulse per second (PPS)

Selecting the **Use NTP with pulse per second (PPS)** radio button also displays the NTP Time Server Information table, shown in Figure 3-16. In addition, a PPS Port column is shown, indicating the NTP server to PPS port correlation.

System (Sysplex) Time

ETS Configuration

Note
This CPC has the role of Current Time Server for an STP-only CTN. If you plan to attach to an ETS device, an ETS configuration is required. Changes made to the ETS configuration have an immediate effect on the time source for the CTN.

External Time Source (ETS)

Use dial out if configured on Hardware Management Console
 Use NTP
 Use NTP with pulse per second (PPS)

NTP Time Server Information

Select	Configured	PPS Port	NTP Time Server	Stratum	Source	Status
<input checked="" type="radio"/>	<input checked="" type="checkbox"/>	0	10.1.6.48	1	PPS	Success
<input type="radio"/>	<input checked="" type="checkbox"/>	1	10.1.6.46	1	hPPS	Success

PPS Port Status
Port 0: Tracking to PPS signal
Port 1: Capable of tracking to PPS signal

Apply Query

Refresh Cancel Help

Figure 3-16 ETS Configuration: NTP with pulse per second

PPS Port 0 corresponds to the NTP server, defined in the upper row of the NTP Time Server Information table. The PPS output of this NTP server must be connected to PPS port 0. PPS Port 1 corresponds to the NTP server, defined in the lower row of the NTP Time Server Information table. The PPS output of this NTP server must be connected to PPS port 1 on the System z server.

The physical location of ETR card 0 and ETR card 1 is documented in the *Installation Manual Physical Planning* documentation for each server (see “Related publications” on page 203).

The user is responsible for entering the IP address matching the PPS output port on the NTP server unit.

The NTP with pulse per second configuration takes effect when **Apply** is clicked. STP creates a time adjustment to steer the Current Server Time to the time provided by the NTP server and the pulses received by the corresponding PPS port.

On the ETS configuration tab, the PPS Port Status section is displayed below the NTP Time Server Information. Information in this section is only displayed for servers configured with the PTS or BTS role. When all four PPS ports are configured, the status of all four reveals which port is tracking to the PPS signal, which ports are capable of tracking to the PPS signal, and which ports are in a state that do not allow use of PPS. The information is valid at the time that the status is displayed and must be refreshed if conditions change.

Setting the CTN ID

When an STP-only CTN is implemented, entering a CTN ID is required to define the CTN. The STP Network ID portion of the CTN ID is entered in the STP Configuration tab, as shown

in Figure 3-17. The same CTN ID must be entered on every server that will participate in the STP-only CTN.

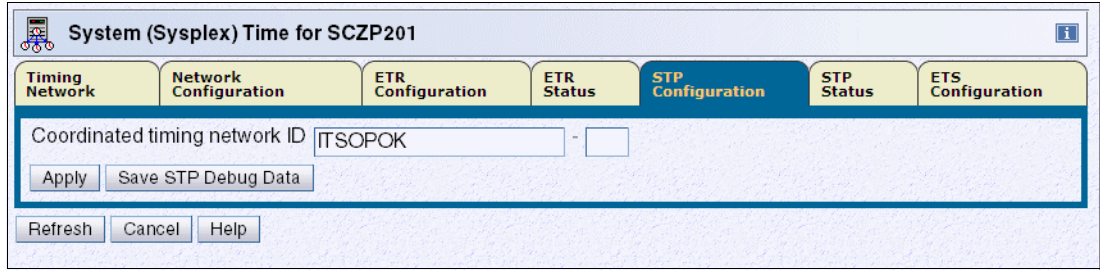


Figure 3-17 System (Sysplex) Time: STP Configuration tab

Initializing the time

When migrating from a Mixed CTN to an STP-only CTN, there is no need to initialize the time because the CTN inherits the following timing information from the Sysplex Timer:

- ▶ Date and time
- ▶ Leap second offset
- ▶ Time zone offset, including Daylight Saving adjustment

However, the time zone algorithm must still be defined before a daylight saving time (DST) schedule can be selected (see “Time zone offset adjustment” on page 112).

When an STP-only CTN is implemented without migrating from a Mixed CTN, time initialization is required to complete the CTN configuration process.

Important: Initializing the time must be done on the server that will become the Current Time Server for the STP-only CTN.

The Initialize Time button is accessed from the Network Configuration tab shown in Figure 3-42 on page 146. The Initialize Time button is only enabled prior to configuring the STP-only CTN. This is the case when an STP-only CTN is set up for the first time or if an STP-only CTN was previously deconfigured.

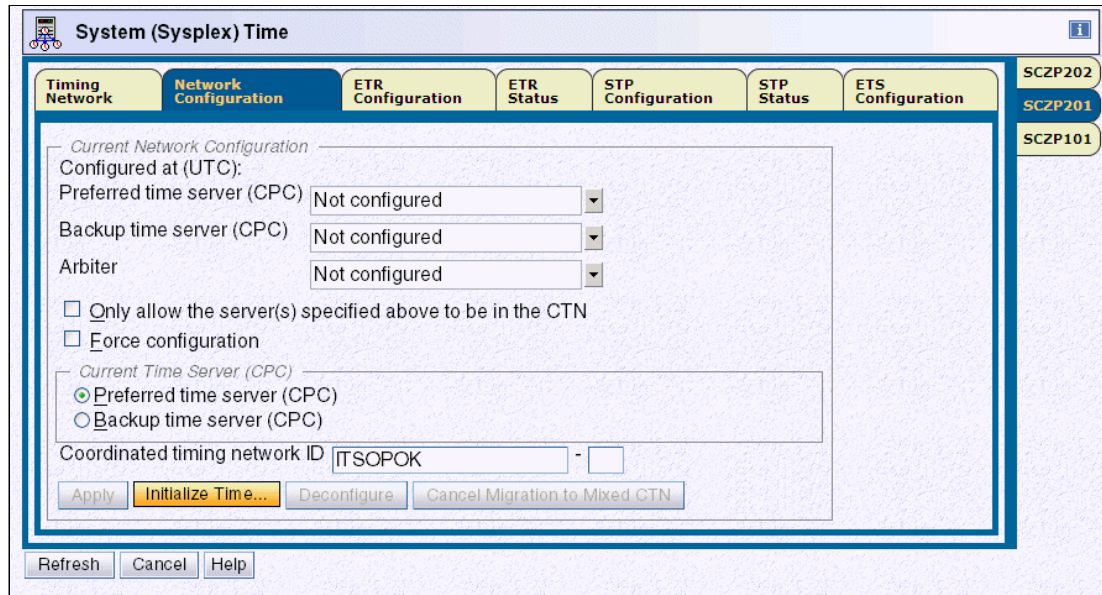


Figure 3-18 HMC workplace: System (Sysplex) Time, Network Configuration

Clicking **Initialize Time** displays the Initialize Time panel displayed in Figure 3-19. There are three radio buttons on the panel, each representing a task that must be completed before a Network Configuration can be applied for an STP-only CTN. The three tasks related to initializing the time are:

- ▶ Set leap seconds.
- ▶ Set time zone.
- ▶ Set date and time.

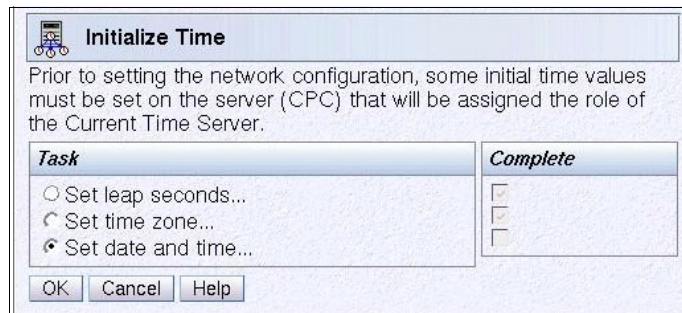


Figure 3-19 HMC workplace: Initialize Time

Set leap seconds

Introduction of the ETR architecture provided a means whereby TOD clocks could be set and stepped very accurately on the basis of an external UTC time source. As time becomes increasingly more accurate, consideration of the usage of leap seconds should be more carefully evaluated.

Since January 1, 1972, occasional corrections of exactly one second, called a leap second, have been inserted into the UTC time scale to keep UTC time within 0.9 second of UT1 at all times.

Depending on the applications and business requirements, leap seconds should be evaluated given the following circumstances:

- ▶ If there are specific accuracy requirements to provide UTC or GMT to the very second, at any instant, then leap seconds should be considered.

Some examples of such specific requirements might be legal or contractual requirements for time stamps to be within some tolerance of UTC Time, or if time stamps are used for time-dependent banking, scientific, or navigational purposes.

To account for leap second corrections, the total accumulated number of leap seconds since January 1972 *must* be entered when setting the time.

- ▶ Most sites have little awareness of leap seconds and ongoing leap second adjustments, and therefore this setting can probably be ignored. If there are no specific requirements for leap seconds then specify a leap second value of zero.

When a leap second is announced, it is inserted as the 61st second (or subtracted as the 60th second) of the last minute of the month of June or December. When a leap second is added, the sequence of time stamps looks like the one illustrated in Example 3-15.

Example 3-15 Leap second sequence: changing from 23 to 24

STCK (STP time)	UTC time (= STCK - leap seconds)
00:00:20	23:59:57
00:00:21	23:59:58
00:00:22	23:59:59 -> Dec 31st
00:00:23	23:59:60 -> Dec 31st leap second
>> leap seconds changed from 23 to 24 <<	
00:00:24	00:00:00 -> Jan 1st
00:00:25	00:00:01
00:00:26	00:00:02

When the leap second offset occurs, z/OS is interrupted just as it is for Daylight Saving Time offset changes if STPMODE YES was specified in the CLOCKxx member of PARMLIB. The offset value is updated if STPZONE YES was also specified in the CLOCKxx member.

Applications do not see a time stamp of 23:59:60, but the extra second is added to the day, thus keeping the z/OS image aligned instantly with the world time standard, UTC.

Note: Making a positive leap second change causes z/OS to spin for the amount of the leap second change. This is to avoid duplicate UTC time stamps. However, a large positive change in leap seconds can cause a system outage, as z/OS must spin for that amount of time to avoid duplicate UTC time stamps.

Most sites have little awareness of leap seconds and ongoing leap second adjustments. Therefore, this setting can probably be ignored. If there are no specific requirements for leap seconds, then you should specify a leap second value of zero.

An external time source (ETS) supplies UTC time. When STP synchronizes with the ETS it steers STP time such that UTC time is synchronized with the ETS. If leap seconds are not used or a leap second adjustment is not scheduled when a leap second is inserted into UTC then the UTC time of the CTN will be wrong by one second. STP will steer to the correct UTC time but steering takes seven hours to complete. Sites that have a requirement for correct UTC time must use leap seconds to avoid this seven hour period while UTC time gets steered to the correct UTC time when a leap second is inserted.

STP accepts local time provided by the user, the UTC time from an external time source (ETS), or a delta value, and calculates the value of the Coordinated Server Time using the leap seconds and time zone offset values specified. The leap seconds value is sent to STP when **OK** is clicked. Table 3-3 lists the current leap second offset.

Table 3-3 Current leap second offset

Leap second offset	Date applied
+23 seconds	December 31, 2005 (scheduled +1 increment)
+24 seconds	December 31, 2008 (scheduled +1 increment)
+25 seconds	June 30, 2012 (scheduled +1 increment)

If a leap seconds value is specified during the *initialize time* process, the appropriate international time bureau should be monitored to determine when future leap seconds might occur. As they occur, the accumulated leap seconds value should be changed on the HMC to maintain accurate UTC time without having to steer for a new leap second over a seven hour period.

Important: Table 3-3 contains leap second offset values in place at the time of writing of this book. (The most recent leap second change was June 30st, 2012, as indicated in Table 3-3.) Check with the official agencies regarding the current leap second offset and future leap second scheduling before making any changes to the system.

The agency that is responsible for making leap seconds announcements is International Earth Rotation and Reference Systems (IERS) ServiceOne of the official agencies is the United States Naval Observatory. See:

<http://www.iers.org>

You can also subscribe for leap seconds announcements at:

http://hpiers.obspm.fr/eop-pc/products/bulletins/bulletin_registration.html

Set time zone

The panel shown in Figure 3-20 is used to set initial time zone parameters for the CTN. The current time zone is set by selecting an entry from the Time zone drop-down list. The values entered are sent to STP when **OK** is clicked.

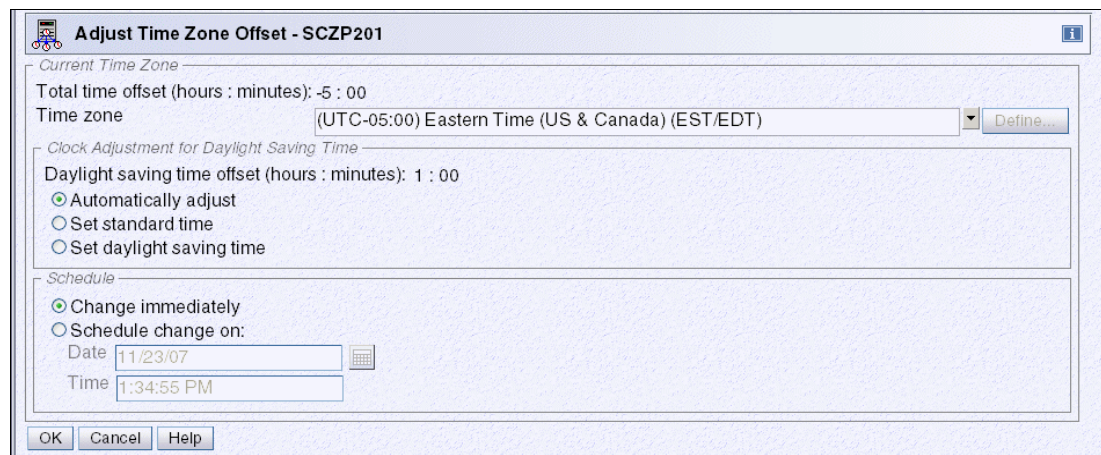


Figure 3-20 HMC workplace: Adjust Time Zone Offset

Clicking the Time zone drop-down arrow displays a list of all the supported time zones. Each of the supported time zone entries has a defined offset from UTC and the Daylight Savings Time offset for that entry if applicable.

The following Daylight Savings Time information is obtained by STP when you select a time zone entry:

- ▶ Daylight Saving time offset
- ▶ Daylight Saving automatic adjustment information (optional)
 - Daylight Saving date and time start algorithm
 - Daylight Saving date and time end algorithm

Automatically adjust is selected by default when the time zone selected supports automatic adjustment of daylight saving time. Otherwise, this button is disabled. Even if automatic adjustment is supported, the user still has the option of selecting Set standard time or Set daylight saving time.

If automatic adjustment for daylight saving time is not supported by the selected time zone, you must decide whether the current date and time are in a daylight saving period and select the Set standard time or Set daylight saving time radio buttons accordingly.

User-defined time zone

If a supported time zone entry that meets the requirements cannot be found, then one of the five user-defined time zones (that is, UD1 to UD5) can be used to define the desired time zone. If a user-defined time zone entry is selected, the Define button is enabled. It is used to display the Define Time Zone Algorithm panel shown in Figure 3-21.

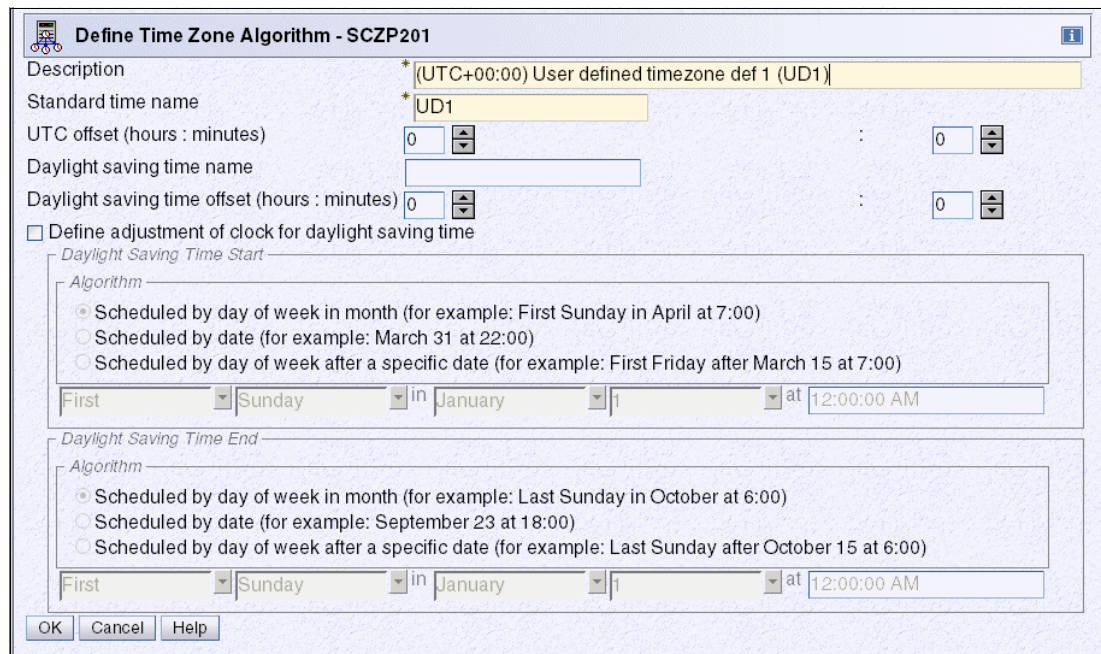


Figure 3-21 HMC workplace: Define Time Zone Algorithm

The Description (maximum of 80 characters) and Standard time name fields (maximum of 4 characters) must be entered, otherwise an error message appears when **OK** is clicked. The standard time name is an abbreviation displayed on various panels to differentiate standard time from daylight saving time.

The UTC offset must be entered in +/- hours and minutes and ranges from -14 to +14 hours.

Also, if the time zone is subject to daylight saving adjustments, then the daylight saving time name and daylight saving offset must be specified. Optionally, algorithms for daylight saving time start and daylight saving time end can be defined to support automatic clock adjustment by selecting the **Define adjustment of clock for daylight saving time** option. The algorithm is saved when **OK** is clicked, but it is not sent to STP until **OK** is clicked on the Adjust Time Zone Offset panel.

Set date and time

The final task in the sequence is to initialize the date and time. Three different methods are provided. The most appropriate method for the site should be chosen.

- ▶ If an ETS is not planned for the CTN, **Set date and time** is clicked, allowing you to specify the local date and time.
- ▶ An external time source can be used to set the time to one of the ETS options configured for STP. Once configured, the ETS is used to calculate the difference between the server's time and UTC. The resulting delta value is not displayed to the user, but instead passed directly to STP when **OK** is clicked.
- ▶ A delta value can be specified through selection of **Modify time by delta to set date and time**. The value specified is either positive (default) or negative and is entered in the *+/-hh:mm:ss.mmm* format.

Irrespective of the method chosen, STP uses the information to calculate the Coordinated Server Time and set the server TOD clock when **OK** is clicked.

At this point all the tasks have a check mark in the Complete column. On the Network Configuration tab, the **Apply** button is enabled, allowing the network configuration to be defined.

The Coordinated Server Time will be passed to other participating servers in the CTN when the server roles and the Current Time Server are configured and activated using the Network Configuration tab. Once the STP-only CTN is configured, the Initialize Time button is disabled.

Activating the STP-only CTN

This is necessary for both an STP-only new install and a migration from a Mixed to an STP-only CTN.

Activation of an STP-only CTN is done by assigning a Current Time Server.

The configuration change is done from the Network Configuration tab, shown in Figure 3-22. The change must be initiated from the server that will become the CTS. Although not mandatory, we recommend that the Backup Time Server and Arbiter also be assigned at this stage. If the STP network consists of fewer than three servers, the Arbiter remains unassigned.

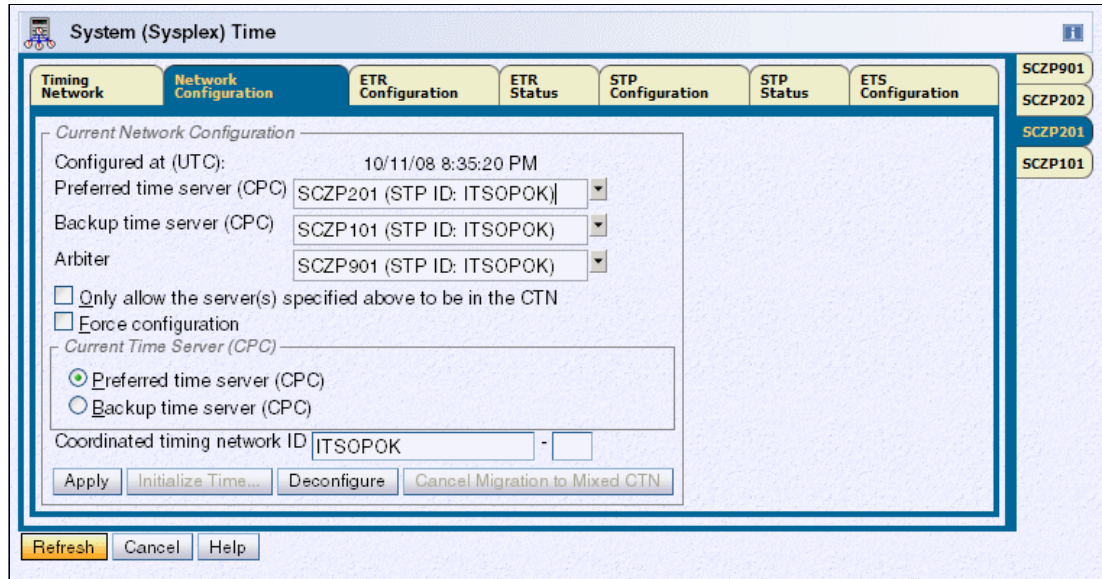


Figure 3-22 System (Sysplex) Time: Network Configuration tab

When the configuration is applied, the assignment of the Current Time Server (CTS) globally transitions all servers with the same CTN ID to STP timing mode. The Current Time Server specified must be either the Preferred Time Server or the Backup Time Server. Either role is able to provide backup for the CTS. Furthermore, the PTS has the capability to automatically re-takeover the CTS role within a recovery scenario once the PTS is back in operational state. For this reason it is a good idea to assign the Current Time Server to the Preferred Time Server because the Backup Time Server cannot do an automatic re-takeover after a recovery action.

If the new STP-only CTN consists of only one server, it is a good idea to select the **Only allow the server(s) specified above to be in the CTN** option. This prevents the server from getting deconfigured when doing a power-on reset or power off/on. However, this option also limits the CTN to one server, unless the option is deselected, which can be done concurrently at any time.

Note: The Only allow option also applies to a two-CEC CTN. It prevents servers from being deconfigured when a site power off affects both the PTS and the BTS.

If the request is for a new STP-only installation, the **Force configuration** check box must be selected (see “Force configuration” on page 130).

If the request includes a migration from a Mixed CTN, the following changes are also triggered when the CTS assignment is applied:

- ▶ The ETR Network ID is removed from the CTN ID for all servers in this CTN. This operation migrates the Mixed CTN to an STP-only CTN with immediate effect.
- ▶ All ETR ports are disabled for servers that were still in ETR timing mode. The operation transitions all servers in this CTN into STP timing mode.

When the activation is complete, the z/OS DISPLAY ETR command now identifies the timing network as an STP-only CTN, as shown in Example 3-16.

Example 3-16 DISPLAY ETR command

```
D ETR
IEA386I 18.06.54 TIMING STATUS 216
SYNCHRONIZATION MODE = STP
  THIS SERVER IS A STRATUM 2
  CTN ID = ITSOP0K
  THE STRATUM 1 NODE ID = 002097.E26.IBM.02.00000001DE50
  THIS IS THE BACKUP TIME SERVER
  NUMBER OF USABLE TIMING LINKS = 7
```

3.3 Managing the time

This section discusses the following time modifications within a CTN and the operational requirements:

- ▶ Time adjustment in an STP-only CTN
- ▶ Offset adjustments in an STP-only CTN
- ▶ Changes in local time
- ▶ Sysplex and multiple time zones

3.3.1 STP time adjustment

Without regular adjustment, the time on the Current Time Server drifts from its initial setting because of the frequency drift of the oscillator used to step the TOD clock. This might not meet time accuracy requirements. Therefore, adjustments might need to be made on a regular basis.

Adjustments of time can be performed at the Current Time Server through use of the Adjust Time button on the Timing Network tab. This can be done manually or by referencing an external time source.

In a Mixed CTN, time adjustment is the same as in an ETR network. The Sysplex Timer adjusts the time of the stratum 1 servers in the Mixed CTN. Since stratum 2 servers in the Mixed CTN are synchronized to the stratum 1 servers, and stratum 3 servers are synchronized to the stratum 2 servers, time adjustments are gradually propagated through the entire Mixed CTN.

In an STP-only CTN, time adjustments are only permitted from the Current Time Server, which propagates the adjustments through the STP-only CTN.

The following adjustments are possible:

► Adjustment steering

STP supports adjustment steering, which allows the time at the Current Time Server to be changed by up to +/- 60 seconds. This is an improvement over the Sysplex Timer that supported a maximum adjustment of +/- 4.999 seconds. Adjustments greater than 60 seconds can be implemented in multiple increments of +/- 60 seconds. This can take considerable elapsed time to achieve.

The offset specified is gradually incorporated into the standard timing messages in small enough increments or decrements that the operating systems, subsystems, and applications are unaware that time is speeding up or slowing down.

The input of the offset to be *steered out* is done either manually or through the ETS.

Note: In an STP-only CTN, the adjustment steering rate is approximately one second every 7 hours. This compares to an ETR network or Mixed CTN where the Sysplex Timer provides a steering rate of approximately one second every 11 hours.

► Base steering

There is another time-steering method built into the STP facility that works in a similar fashion to the adjustment steering. This is an automatic function requiring no user control. This is known as base steering, which is performed at the Current Time Server and requires an ETS.

By comparing the time obtained from subsequent accesses of a configured ETS with the corresponding Coordinated Time Server values, STP can compute the amount of drift that has occurred between the events. This represents the inherent inaccuracy of the Current Time Server oscillator over time. With this information, STP can automatically introduce a compensation offset into the STP timing messages by additional steering to counter the drift. As a result, the Current Time Server self-corrects over time so that the offset returned from future ETS accesses approaches zero as greater accuracy is achieved.

Base steering is only performed if the NTP server is a stratum 1. The stratum level of the NTP server can be seen in the Network section of the Timing Network panel as shown in Figure 3-23 on page 108.

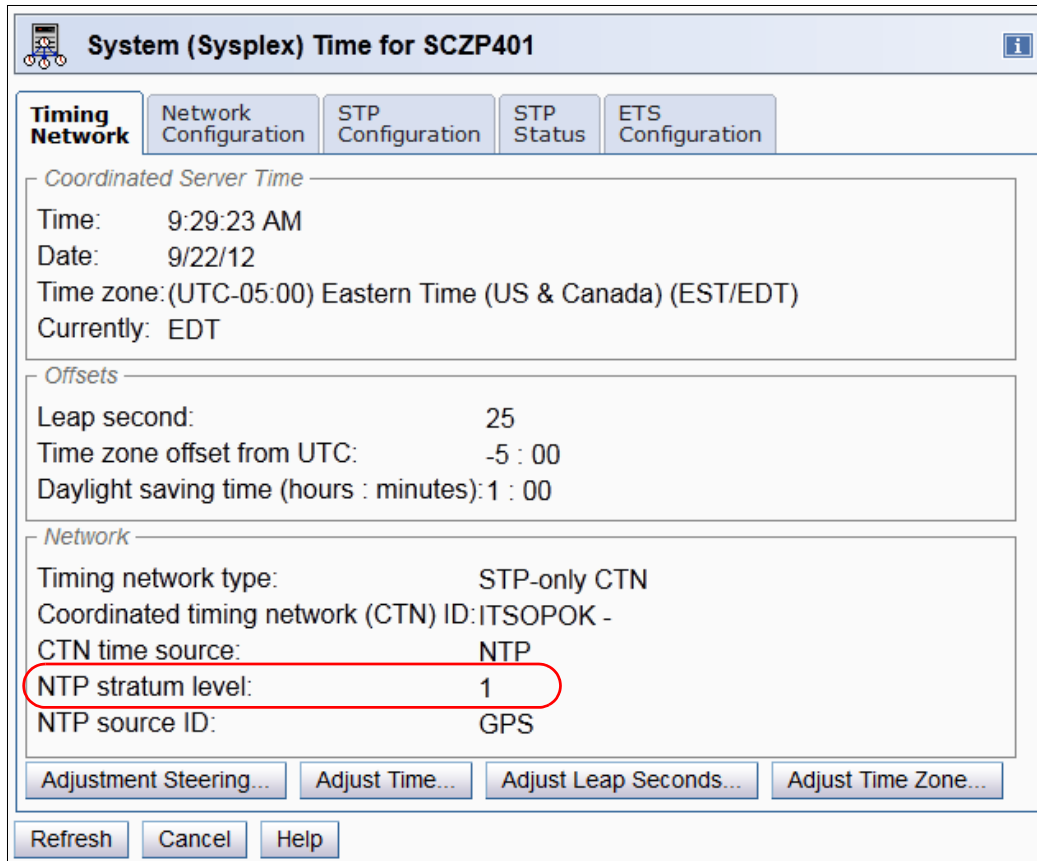


Figure 3-23 NTP server stratum level

Table 3-4 provides a summary of STP clock adjustment.

Table 3-4 STP clock adjustment summary

	Dial-out or NTP server access adjustment	Manual adjustment	Manual set
Adjustment steering	Yes	Yes	No
Base steering	Yes, after multiple ETS accesses	No	No

Note: Once a CTN is configured, the only way to set the date and time is to deconfigure the CTN and then go through the process to initialize the time again. Refer to the *Server Time Protocol Implementation Guide*, SG24-7281, for details.

The Adjust Time button is not available when the CTN time source is NTP with pulse per second. Since STP utilizes PPS input every second, manual adjustment is disabled.

3.3.2 STP offset adjustments

The STP timing message includes the following:

- ▶ Coordinated Server Time
- ▶ Leap second offset

- ▶ Time zone offset
- ▶ Daylight saving time offset

These values are transmitted from the Current Time Server to all the servers in the CTN. How the z/OS system image uses these values is dependent on options specified in the TIME macro in combination with options specified in CLOCKxx at IPL.

Different time results can be received depending on the options specified in the TIME macro, as shown in Table 3-5.

Table 3-5 TIME macro options

	TIME macro with ZONE=LT	TIME macro with ZONE=UTC	TIME macro with STCK
Include TOD in result	Yes	Yes	Yes
Include leap second offset	Yes	Yes	No
Include time zone offset	Yes	No	No

In addition, the parameters specified in the CLOCKxx member at IPL determine where these values are obtained. Table 3-5 indicates the behavior of a z/OS system on a server in ETR timing mode based on the CLOCKxx settings.

Table 3-6 z/OS system on a server in ETR timing mode

	ETRMODE=NO ETRZONE=NO	ETRMODE=YES ETRZONE=NO	ETRMODE=YES ETRZONE=YES
Step TOD to Sysplex Timer.	No	Yes	Yes
Include time zone offset from ETR.	No	No	Yes
Include time zone offset from CLOCKxx.	Yes	Yes	No
Allow local time adjustment via z/OS SET commands.	Yes	Yes	No

Table 3-7 shows comparable details for z/OS systems on a server in STP timing mode.

Table 3-7 z/OS system on a server in STP timing mode

	STPMODE=NO STPZONE=NO	STPMODE=YES STPZONE=NO	STPMODE=YES STPZONE=YES
Step TOD to Current Time Server.	No ^a	Yes	Yes
Include time zone offset from Current Time Server.	No	No	Yes
Include time zone offset from CLOCKxx.	Yes	Yes	No
Allow local time adjustment via z/OS SET commands.	Yes	Yes	No

a. The base TOD steering affects all logical partitions regardless whether they operate in STP mode or not because it changes the H/W TOD gradually without notifying it to operating systems. The STP mode makes the difference when the local H/W TOD drifts and differs from Stratum 1 (CTS/CST) resulting in an STP synch check. The STP synch check is to be presented only to logical partitions running in STP mode. Several other STP machine checks and external interruptions are also to be presented only to those logical partitions running in STP mode.

Leap second offset adjustment

Adjustments are only necessary (if leap seconds are used) if the applications require precise synchronization accuracy to UTC. Examples of such specific requirements might be legal or contractual requirements for time stamps to be within tolerance of UTC Time, or if time stamps are used for time-dependent banking, scientific, or navigational purposes.

Periodically, the International Earth Rotation and Reference Systems Service in Paris advises that a leap second adjustment, which might be either positive or negative, is required to be introduced into civil time. At the time of this writing, there have been 24 leap seconds introduced into civil time since 1972, the last being December 31, 2008. Each of the leap seconds introduced so far has been positive. However, it is theoretically possible to have a negative leap second adjustment.

Note: Leap seconds are automatically built into UTC time obtained from an external time source. Any leap second offset that is defined is taken into account when calculating the delta between the Current Time Server and the time received from the ETS. This is required to prevent double accounting.

If the ETS is used to incorporate an additional leap second into the Current Time Server TOD, then be aware that the requirement that leap second adjustments should occur at the same time worldwide has been breached. This is because TOD adjustments via the ETS are implemented over a period of time via adjustment steering rather than immediately.

If leap seconds are used, the *adjust leap seconds* facility should be used to ensure that the new offset is applied as a single adjustment at the correct time.

Typically, adjustments to the leap second offset are scheduled to be applied on either June 30th or December 31st. However, the offset adjustment might be scheduled for a particular date depending upon the requirements.

A positive leap second is inserted after UTC 23:59:59, as UTC 23:59:60, then UTC 00:00:00 would occur as normal. Note that 23:59:60 will never be made visible to users of the TOD

clock. A negative leap second is potentially possible (however, it has not occurred to date) and is implemented by excluding UTC 23:59:59.

When the leap second offset occurs, z/OS is interrupted just as it is for DST offset changes if STPMODE YES was specified in the CLOCKxx member of SYS1.PARMLIB. The offset value is updated if STPZONE YES was also specified.

The Adjust Leap Second panel is used to apply the leap second offsets when they become available. Clicking **Adjust Leap Seconds** from the Timing Network tab on the Current Time Server displays the panel shown in Figure 3-24.



Figure 3-24 System (Sysplex) Time, Timing Network, Adjust Leap Second Offset

Important: Never make leap seconds adjustments with more than one second at a time. A large positive change in leap seconds can cause a system outage, as z/OS must spin for that amount of time to avoid duplicate UTC time stamps.

When the Adjust Leap Second Offset panel is displayed, the current leap second offset in effect for the STP-only CTN value is displayed in the Offset box. This value is either inherited from the Sysplex Timer during the migration process from an ETR network or set during the Initialize Time function.

Important: Leap second adjustments occur at the same time worldwide, so it must be understood when UTC 23:59:59 will occur in the particular time zone, because this is when a scheduled leap second offset adjustment will be applied.

Time zone offset adjustment

The Offsets section of the Timing Network tab displays the current time zone offset information for an STP-only CTN.

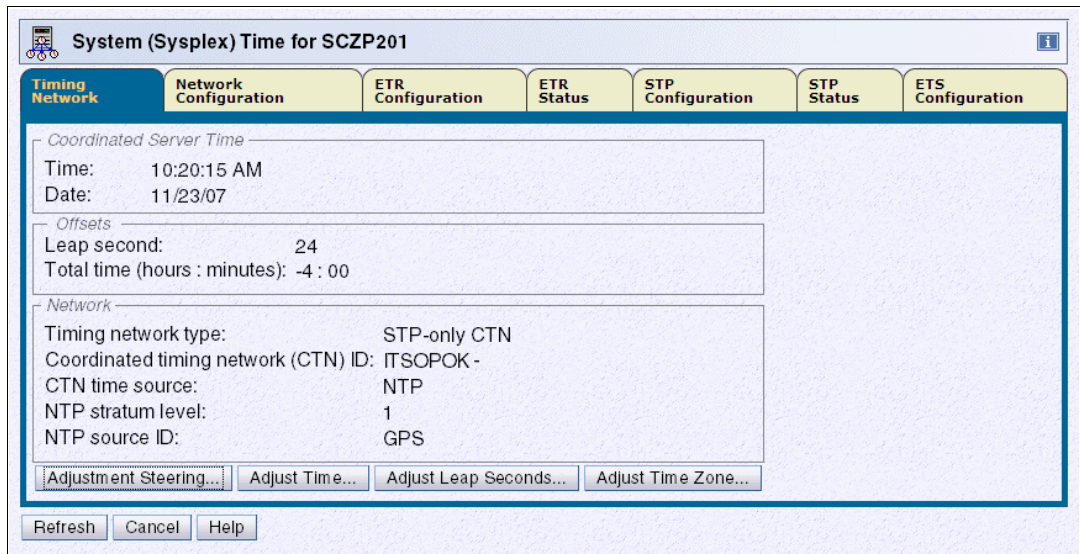


Figure 3-25 Timing Network tab following a migration from Mixed CTN to STP-only

If the Offsets section displays a Total time (hours : minutes) field as shown in Figure 3-25, a specific time zone is not set. This only occurs when the time zone information (incorporating a Daylight Saving offset, if any) has been inherited from a Sysplex Timer during a migration process. The value displayed in the Total time (hours : minutes) field is retained until a time zone is selected.

The Total time offset field is no longer displayed once a time zone selection has been applied (Figure 3-26).

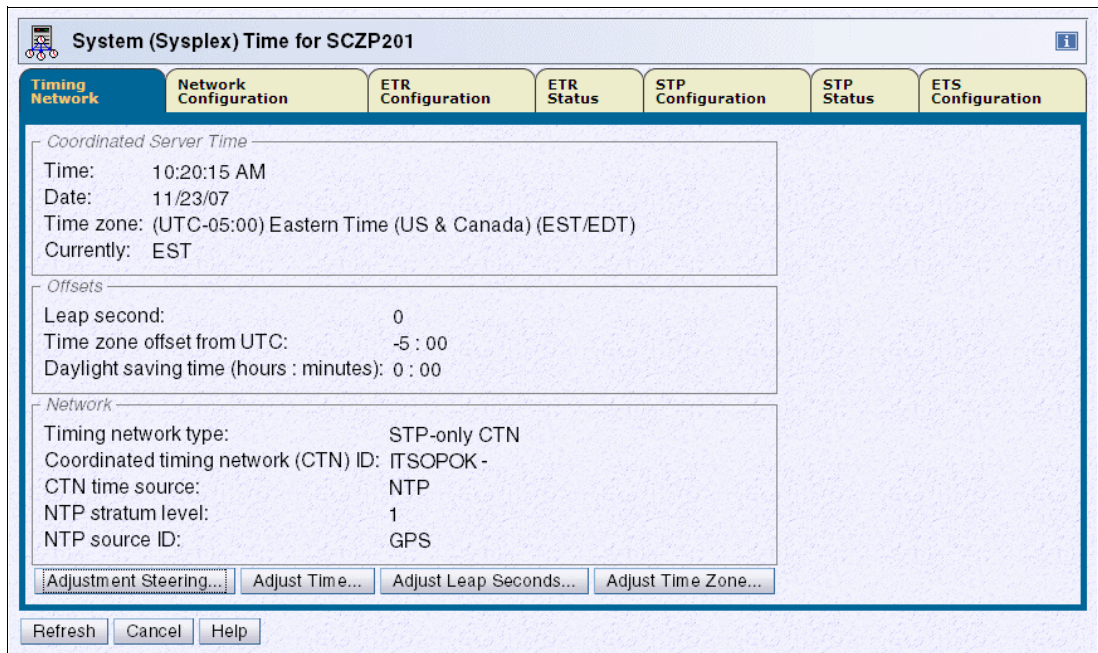


Figure 3-26 Timing Network tab; after adjusting the time zone

Following a migration to STP-only CTN, a Current Time Zone must be defined before the automatic clock adjustment for daylight saving time can be enabled.

Note: During a migration process from Mixed to STP-only CTN, only the time zone offset (including daylight saving offset adjustment) is inherited from a Sysplex Timer, but no time zone definition is made.

The Adjust Time Zone button on the Timing Network tab is used to define or modify a time zone setting applicable for the environment. This displays the Adjust Time Zone Offset panel shown in Figure 3-27.

Adjust Time Zone Offset - SCZP901

Current Time Zone
Time zone (UTC-05:00) Eastern Time (US & Canada) (EST/EDT) Define...

Clock Adjustment for Daylight Saving Time
Daylight saving time offset (hours : minutes): 1 : 00
 Automatically adjust
 Set standard time
 Set daylight saving time

Schedule
 Change immediately
 Schedule change on:
Date 10/4/09
Time 2:47:39 PM

Scheduled Clock Adjustment for Daylight Saving Time
Local time name: EST
Offset (hours : minutes): 0 : 00
Scheduled time (UTC): 11/1/09 6:00:00 AM

OK Cancel Help

Figure 3-27 HMC workplace: System (Sysplex) Time, Timing Network, Adjust Time Zone Offset

Select one of the supported time zones that are provided by default or use one of the five user-defined time zones to customize an entry to specifically meet the requirements.

Supported time zone offsets

A number of supported time zone entries are provided by default. Each of these entries has a defined offset from UTC.

Note: The time zone on the System z10 or System z9 Support Element is independent of the time zone set for the CTN. The same is also true for the HMC. It is possible to have a different time zone at each Support Element, the HMC, and for the CTN.

In addition, each entry can optionally have a time zone algorithm defined that contains the following daylight saving information:

- ▶ Daylight saving offset
- ▶ Daylight saving automatic adjustment information (optional):
 - Daylight saving date and time start algorithm
 - Daylight saving date and time end algorithm

If the selected time zone supports automatic adjustment by providing a time zone algorithm with the necessary start and end information, then the Automatically adjust radio button can be used to activate automatic adjustment for the site.

Alternatively, the time zone might not support automatic adjustment, or is handled manually. In this case select either **Set standard time** or **Set daylight saving time** to indicate whether a daylight saving period is active when the selected time zone is activated.

The Scheduled Clock Adjustment for Daylight Saving Time section is displayed if an adjustment has been scheduled either automatically as the result of an Adjust Time Zone offset algorithm or manually through the *schedule change on* facility.

Tip: There is no direct way to review the time zone algorithm that might optionally be part of a supported time zone entry. Therefore, there is no direct way to verify that the daylight saving time offset entry for the time zone will meet the requirements. However, this can be done indirectly by applying the time zone offset entry for the time zone with *automatically adjust* and then reviewing fields in the Schedule Clock Adjustment for Daylight Saving Time section for applicability correctness in relation to the next time zone adjustment.

DST changes implemented by a government authority

An example of this is the USA Energy Policy Act of 2005 that changed daylight saving time in the USA in 2007.

For an ETR network or a Mixed CTN, a DST offset must be scheduled at the appropriate time at the Sysplex Timer console. For an STP-only CTN, which has the option of automatically scheduling DST, support for the new daylight saving time start and end dates are available provided that the server has the correct level of microcode.

Future DST changes implemented by any governmental authority might require MCL updates. Make sure to keep the server MCLs as up to date as possible.

User-defined time zone offsets

If a supported time zone entry that meets the requirements cannot be found, then one of the five user-defined time zones (that is, UD1 to UD5) can be used to define the desired time zone. If a user-defined time zone entry is selected, the Define button is enabled. It is used to display the Define Time Zone Algorithm panel shown in Figure 3-28.

Define Time Zone Algorithm - SCZP901

Description * (UTC+00:00) User defined timezone def 1 (UD1)

Standard time name * UD1

UTC offset (hours : minutes) -5 : -30

Daylight saving time name

Daylight saving time offset (hours : minutes) 0 : 0

Define adjustment of clock for daylight saving time

Daylight Saving Time Start

Algorithm

Scheduled by day of week in month (for example: First Sunday in April at 7:00)

Scheduled by date (for example: March 31 at 22:00)

Scheduled by day of week after a specific date (for example: First Friday after March 15 at 7:00)

Third Sunday in March at 12:00:00 AM

Daylight Saving Time End

Algorithm

Scheduled by day of week in month (for example: Last Sunday in October at 6:00)

Scheduled by date (for example: September 23 at 18:00)

Scheduled by day of week after a specific date (for example: Last Sunday after October 15 at 6:00)

Fourth Sunday in November at 12:00:00 AM

OK Cancel Help

Figure 3-28 HMC workplace: Define Time Zone Algorithm

The Description (maximum of 80 characters) and Standard time name fields (maximum of 4 characters) must be entered, otherwise an error message appears when **OK** is clicked. The standard time name is an abbreviation displayed on various panels to differentiate standard time from daylight saving time.

The UTC offset must be entered in +/- hours and minutes and ranges from -14 to +14 hours.

Also, if the time zone is subject to daylight saving adjustments, then the daylight saving time name and daylight saving offset must be specified. Optionally, algorithms for daylight saving time start and daylight saving time end can be defined to support automatic clock adjustment by selecting the **Define adjustment of clock for daylight saving time** option. The algorithm is saved when **OK** is clicked, but it is not sent to STP until **OK** is clicked on the Adjust Time Zone Offset panel.

3.3.3 Changes in local time

z/OS allows the user to obtain either STCK time, UTC time, or local time depending on their requirements. The difference between UTC time and local time is usually the time zone offset under normal circumstances. The time zone offset can be managed at the z/OS level by specifying the ETRZONE=NO and STPZONE=NO options in the CLOCKxx member. The relevant option applies depending on whether the server is in ETR timing mode or STP timing mode.

When this is done, the TIMEZONE parameter in the CLOCKxx member is used to set the time zone offset at IPL, and a number of z/OS SET commands can be used to dynamically

adjust the offset when required. Similarly, the coupling facility supports the concept of time zone offset and allows dynamic modification of the time zone offset via a command.

z/OS commands

On a z/OS system, local date and time can be modified dynamically. The ability to do this is dependent on what options have been specified in the CLOCKxx member at IPL, as shown in Table 3-8.

Table 3-8 z/OS time adjustment through command cross-reference

	Adjust time via z/OS command		
	Local time ZONE=LT	UTC time ZONE=UTC	STCK time STCK
ETRMODE=NO, ETRZONE=NO	Yes	Yes ¹	Yes ¹
ETRMODE=YES, ETRZONE=NO	Yes	No	No
ETRMODE=YES, ETRZONE=YES	No	No	No
STPMODE=NO, STPZONE=NO	Yes	Yes ¹	Yes ¹
STPMODE=YES, STPZONE=NO	Yes	No	No
STPMODE=YES, STPZONE=YES	No ²	No	No

1. UTC time cannot be changed by a z/OS command. It can be modified during the IPL process in response to message
 IEA888A [UTC DATE=yyyy.ddd,CLOCK=hh.mm.ss] LOCAL
 DATE=yyyy.ddd,CLOCK=hh.mm.ss

REPLY U, OR UTC/LOCAL TIME

which is issued at IPL when OPERATOR PROMPT is specified in CLOCKxx with ETRMODE NO, STPMODE NO and SIMETRID not specified.

2. Any attempt to change the local time or date when the server is operating in STP mode generates message IEA279I: IEA279I ALL CLOCK RELATED SET COMMANDS ARE IGNORED WHEN IN STP MODE.

DISPLAY TIME

This command can be used to display the local time and date, and the UTC time and date, as seen in Example 3-17.

Example 3-17 z/OS DISPLAY TIME command

```
D T
IEE136I LOCAL: TIME=hh.mm.ss DATE=yyyy.ddd UTC: TIME=hh.mm.ss DATE=yyyy.ddd
```

Under normal circumstances the difference between local time and UTC time is the time zone offset (incorporating daylight saving time offset, if any) applicable to the time zone.

SET DATE

This command is used to change the local date. The syntax is:

```
SET DATE=yyyy.ddd
```

It has the following restrictions:

- ▶ yyyy is the year. It must be in the range 1900–2042. The value specified must consist of four digits and must be within 70 years of the UTC date, otherwise the SET command is ignored.

- ▶ *ddd* is the day. It must be in the range 001–366 and meet leap year restrictions.
- ▶ The maximum date that can be specified is 2042.260.

SET CLOCK

This command is used to change the local time. The syntax is:

```
SET CLOCK=hh.mm.ss
```

This command is used in conjunction with the SET DATE command to set a maximum value of 23.53.47 on 2042.260. The server's TOD clock is not updated by this command, nor is the logical TOD of the logical partition this z/OS image is operating on. The change made by this command is effective for the duration of this IPL only.

Also, z/OS does not change the date when the new time implies a change of date, so either use the DATE parameter or wait for time to pass midnight if the new time is for tomorrow.

SET RESET

This command causes the time zone offset to be reset to the value that was read in from the CLOCKxx member during IPL, causing the local date and time to be changed accordingly. The syntax is:

```
SET RESET
```

This annuls all previous SET DATE, SET CLOCK, and SET TIMEZONE commands and reestablishes the relationship local date and time = UTC date and time + time zone offset.

SET TIMEZONE

This command can be used to change the time zone offset to a different value from that specified at IPL via the TIMEZONE parameter in CLOCKxx. This automatically adjusts local date and time accordingly. The syntax is:

```
SET TIMEZONE={W|E}.hh[.mm]
```

The time zone offset direction is West (W) or East (E). West is defaulted if not specified. The value for *hh* must be between 00 and 15 and the value for *mm* must be between 00 and 59.

The daylight saving time changes can be handled manually using the SET CLOCK command rather than having it done automatically via Server Time Protocol (STP) or the Sysplex Timer. Using this method there is always a degree of error because the difference between the local time and UTC time will not exactly match the time zone offset that would have been achieved by updating the TIMEZONE statement in CLOCKxx and IPLing.

The new z/OS SET TIMEZONE command overcomes this problem by applying the correct offset value in the CVTTZ field, causing an exact time zone offset to be applied.

Tip: If the SET CLOCK command is used to change local time for daylight saving offset purposes, then use the SET TIMEZONE command instead for better accuracy.

Remember that if the time zone offset for daylight saving purposes is changed dynamically using either the z/OS SET CLOCK or the SET TIMEZONE commands, the TIMEZONE statement in CLOCKxx must be updated so that the new offset is not regressed at the next IPL.

Coupling facility commands

In a Parallel Sysplex environment, coupling facilities require time awareness to support CF request time ordering. The server TOD is used for this purpose.

Coupling facilities also support the concept of time zone offset, which is used only for the purpose of timestamping messages that are displayed on the console.

There is no CFCC command available to display time. However, all messages that appear on the CF console include a time stamp in local time format, which is the server TOD with the time zone offset applied. Therefore, the current local date and time at the CF console can be indirectly determined by entering any command (valid or invalid) and reviewing the time stamp in the resulting response.

Because the CF supports a local time format that incorporates the time zone offset, it also provides methods to both display the current time zone offset setting and to change it if required.

DISPLAY TIMEZONE

Use the CFCC DISPLAY TIMEZONE command to display the current time zone offset being used by the coupling facility. The syntax is:

```
DISPLAY TIMEZone
```

This produces a single line indicating how many hours and minutes the current time zone is east or west of Greenwich Mean Time (Example 3-18).

Example 3-18 CFCC DISPLAY TIMEZONE command

```
2006272 11:06:47 => display timezone
2006272 11:06:47 CF0271I Timezone is 04:00 West of Greenwich Mean Time
```

TIMEZONE

The CFCC supports a command allowing the time zone offset to be changed, if this is a requirement. The syntax is:

```
TIMEZone {0|hh|hh:mm|:mm} {East|West}
```

Use the command in Example 3-19 to adjust the local time as displayed in messages on the coupling facility console for the onset and removal of daylight saving time.

Example 3-19 CFCC TIMEZONE command

```
2006272 11:17:31 => timezone 05:00 west
2006272 11:17:31 CF0271I Timezone is 05:00 West of Greenwich Mean Time
```

Coupling facility implications at daylight saving time changes

When a CF image partition is activated and the server is using Server Time Protocol or a Sysplex Timer source, the CFCC uses only one of the following time offset options:

- ▶ The logical partition time offset specified in the image profile
- ▶ The TIMEZ offset

The TIMEZ offset overrides the logical partition time offset. Use the TIMEZ command for DST changes as described in the document at the following website:

<http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/TD103077>

3.3.4 Logical partition time offset

IBM System z servers support a function called *logical partition time offset*. Logical partition time offset³ support provides for the optional specification of a fixed time offset (specified in days, hours, and quarter hours) for each logical partition activation profile. The offset, if

specified, is applied to the time that a logical partition receives from Server Time Protocol or a Sysplex Timer.

It is sometimes necessary to run multiple Parallel Sysplexes with different local times and run with the time set to TOD Clock LOCAL. This causes the results returned in the store clock (STCK) instruction to reflect local time. With logical partition time offset support, logical partitions on each server in a Parallel Sysplex that must do this can specify an identical time offset that shifts time in the logical partition sysplex members to the desired local time. The logical partition time offset value, if specified, is applied to the time value that a logical partition receives from the time source. The time zone offset and DST offset are independent of this parameter.

Remaining logical partitions on the servers can continue to participate in current date production Parallel Sysplexes utilizing the same STP messages or Sysplex Timers with the time provided by STP or the Sysplex Timers. This function is supported by all supported releases of z/OS.

Open the image profile of the image to be modified. From the General tab, select **Logical Partition Time Offset** from the Clock Type Assignment area (Figure 3-10).

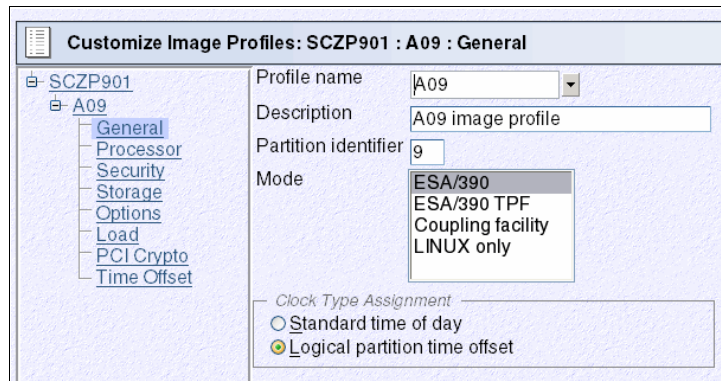


Figure 3-29 Clock Type Assignment: Logical partition time offset

A new selection becomes available on the left side of the panel, called Time Offset. Use the Time Offset window to set the offset, as seen in Figure 3-30.

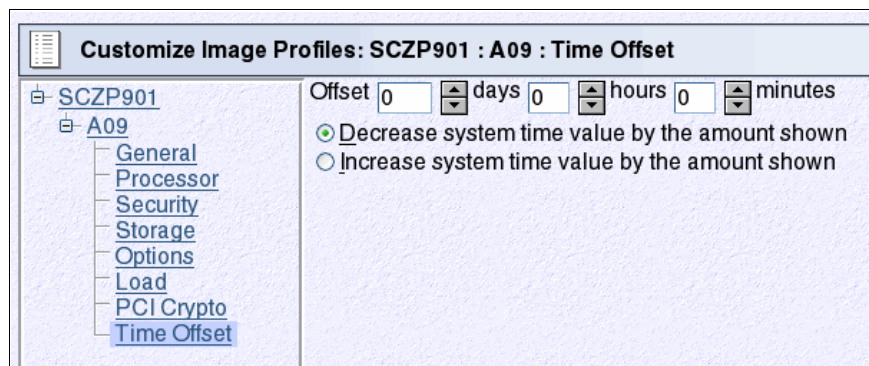


Figure 3-30 Time Offset

³ In earlier server documentation this function was called *logical partition Sysplex Timer offset*.

The settings available on the Time Offset window have the following meanings:

▶ **Offset.**

Type or spin to the number of days, hours, and minutes to be set as the offset from the time of day supplied by its time source. The offset can be in the following range:

- 0 to 999 days
- 0 to 23 hours
- 0, 15, 30, or 45 minutes

▶ **Decrease system time value by the amount shown.**

Select this choice to set the logical partition's clock back from the time of day supplied by its time source by the number of days, hours, and minutes in the offset. Use this setting to provide a local time zone *west* of UTC.

▶ **Increase system time value by the amount shown.**

Select this choice to set the logical partition's clock ahead of the time of day supplied by its time source by the number of days, hours, and minutes in the offset. Use this setting to provide a local time zone *east* of UTC or a date and time in the future.

3.3.5 Parallel Sysplex and multiple time zones

Some Parallel Sysplex installations have a requirement where different z/OS sysplex members are IPLed in different time zones. This is accomplished by:

- ▶ Having all sysplex members on a common time source with STPMODE YES or ETRMODE YES
- ▶ Not using STP or the Sysplex Timer to obtain time zone information, by specifying STPZONE NO or ETRZONE NO and TIMEZONE d.hh.mm.ss.

Alternatively, one group of sysplex members in a common time zone could use STP or the Sysplex Timer as the time zone source, using STPZONE YES or ETRZONE YES, while other sysplex members do not.

This section discusses how to create and manage the *multiple time zone* environment.

Create the multiple time zone environment

The example in this section reflects a consolidation of multiple sysplexes into a single sysplex. This in itself is not remarkable. What is unusual is the fact that the sysplex is supporting applications intended to operate in images that support different time zones.

The following issues are addressed when creating the environment:

- ▶ Software considerations
- ▶ Vendor software
- ▶ Staggered time zone setup considerations

Software considerations

Producing a sysplex with staggered time zone offsets introduces many peculiarities into the environment that not every installation or department is happy with. However, there might be business reasons why this configuration is favored, such as sysplex consolidation efforts or business consolidations. Some, but not all, of the peculiarities are:

- ▶ SMF considerations

SMF records contain local time. If SMF records from two systems are used for accounting purposes, the data is not accurate. For example, assume that a job is submitted on logical partition 1A, getting a job reader start time in one time zone, and the job executes on another image, getting job start and end times in the other time zone.

- ▶ RMF™ considerations

RMF displays and reports are suspect for the same reasons cited under SMF.

- ▶ OPERLOG considerations

Although the records are ordered by UTC (GMT) time stamps, they are displayed with local time offsets. For example, in a merged OPERLOG, messages from the two systems described appear to be out of sequence.

- ▶ JES2 time-sensitive command considerations

Some JES2 commands, such as \$TA and \$PJQ, have time-sensitive parameters that would cause actions to be taken at different times. For example, the commands would execute N hours earlier on one system than the other.

The system will function properly, but the external view will complicate the role of some people in the enterprise.

Vendor software

Since a sysplex allows work to be routed to any machine in the configuration, programs asking for local time can get a variety of answers depending on the server where the work was routed. To help in this situation there is software available that will provide an *altered* date and time to users, user applications, subsystems software, CICS® regions, and even specific transactions or terminal IDs.

One example is Application Time Facility for z/OS (formally TIC-TOC from ISOGON, Inc). To evaluate whether it is appropriate, start by consulting the following website:

<http://www.ibm.com/software/awdtools/tictoc/>

Configuration considerations

To set up the sysplex with different time zones, there are considerations for z/OS and middleware. In-depth details are beyond the scope of this book, but some examples are mentioned here:

- ▶ CICS complex (CICSplex)

Although originally a full-function CICS ran in a single address space (region) in the z/OS environment, most CICS users are more accustomed to running multiple, interconnected CICS regions. Using the CICS Multi-region Operation (MRO) Intercommunication Facility, users can combine CICS regions into a complex of subsystems.

Using multi-region operations, CICS functions can be separated into individual regions, with the different types of CICS regions being classified as *resource managers*. With the latest enhancements to MRO, these CICS resource manager regions can reside in one or more z/OS images.

► CICSplex® system manager

If there are servers in different geographical locations, are there connections between those processors, or are they managed as separate entities, each with its own workload? If these separate units exist in the enterprise, it is likely that there is a need to define multiple CICSplexes, and so manage the enterprise CICS systems as though they belonged to more than one enterprise.

Much of CICSplex SM's activity is time-dependent. For example, it can be specified that a monitor definition or an analysis definition is to be active during a particular time period. CICSplex SM does not require every instance in a single CICSplex to be running in the same time zone, and so must be able to accommodate any time zone differences between entities. Therefore:

- Whenever a time period definition (using the CICSplex SM PERIODEF view) is created, a time zone must be specified in the definition. For example, a time period definition called MORNING can be created for the hours 08:00 through 11:59 Eastern standard time.
- A time zone must be specified for each CICSplex SM address space (CMAS) in its data repository initialization job, EYU9XDUT. A permanent change to the CICSplex SM address space time zone value can be made, even while the CMAS is running, via the CICSplex SM user interface.
- A time zone must be established for each managed CICS system. When a CICS system is defined in a CICSplex SM, the time zone in which the system is running can be specified. Alternatively, if no time zone is specified in the CICS system definition, the CICS system is assumed to be running in the time zone that is the default for the CICSplex SM address space to which it connects. Allow the time zone of a managed CICS system to default to that of its CICSplex SM address space. The time zone of a managed CICS system can be altered subsequently while the CICS system is running. Any change made in this way lasts for the lifetime of the CICS system, or until it is next changed, whichever is sooner.
- A time zone must be specified for every CICSplex when it is first defined. This time zone is used by the CICSplex SM monitor function to determine the actual time at which the monitor interval for the CICSplex expires. The CICSplex time zone can be altered via the CICSplex SM user interface.

Time zones are specified using single-character codes in the range B through Z. For example, code S represents Mountain Standard Time, code T represents Central Standard Time, and code C represents Eastern Europe Time. A complete list of the codes can be found in CICSplex SM Setup and Administration - Volume 1. CICSplex SM allows offsets (known as *time zone adjustments*) in the range 0 through 59 minutes to be specified to accommodate regions that are not running in the standard time zones. Also, daylight saving time can be specified.

Because multiple CICSplex SM entities require a time zone to be specified, there is obvious potential for *conflicting* time zones to be specified. For example, it is possible that a CMAS and a CICS instance in the same CICSplex could be in different time zones. CICSplex SM always honors the time zone of the CICS-managed application system. Suppose the following conditions exist:

- The time period definition time zone is S.
- The CICSplex SM address space time zone is B.
- The CICS-managed application system time zone is C.

Time zone C is used by the CICS-managed application system, and the CMAS makes any necessary adjustments between time zones B, C, and S to ensure that the time zone is honored.

Manage the multiple time zone environment

Note: In this section, it might be preferable to use the new z/OS SET TIMEZONE command instead of the z/OS SET CLOCK command. Refer to “SET TIMEZONE” on page 117 for more information.

Since each sysplex member is in a different time zone, it is necessary to stagger Daylight Saving Time changes. Daylight saving time changes could be accomplished individually in each sysplex member using the SET CLOCK command from the z/OS console.

The cities used in the examples that follow include Sydney in the southern hemisphere (during summertime) and London, Atlanta, and Los Angeles in the northern hemisphere (using standard time).

Operational considerations

In Figure 3-31, each of the three cities accurately reflects the number of hours from the Greenwich meridian when Daylight Saving Time would be in effect in the southern hemisphere. That is, Sydney's TIMEZONE parameter is shown while Daylight Saving Time is in effect. London and Atlanta TIMEZONE parameters are shown with standard time in effect. Each hemisphere is changing the offset in different directions. When the southern hemisphere is changing from Daylight Saving Time to standard time (back), the northern hemisphere is changing from standard time to Daylight Saving Time (a forward change).

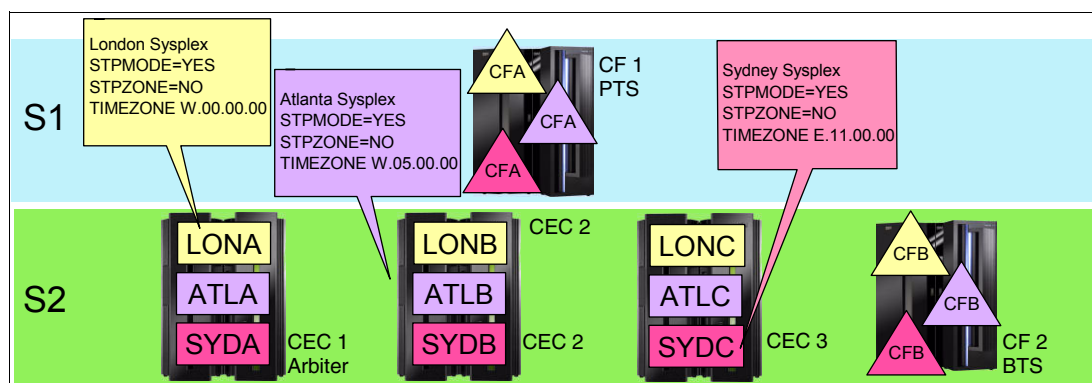


Figure 3-31 Staggered offset changes: one sysplex supporting multiple time zones

When planning for Daylight Saving Time changes in different countries, remember that not all countries agree on which dates the switch will occur. In fact, each country's government might change the Daylight Saving Time switch dates from year to year. A country that has multiple time zones might have a different change date in each zone. In our example, however, assume that all countries are going either to Daylight Saving Time (London and Atlanta) or to standard time (Sydney) on the same arbitrary day.

Just to keep things interesting, assume that this sysplex is physically located in Los Angeles, California. Any z/OS images representing Pacific Standard Time in the USA are not depicted in Figure 3-31. However, assume that those images contain STPZONE YES and that STP is being used to obtain Daylight Saving Time offsets.

Note: The world time zones shown in Table 3-9 represent only the local standard time offset from the Greenwich meridian and do not include Daylight Saving Time offsets in either hemisphere.

Table 3-9 Local standard time offsets

Sydney	London	Atlanta	Los Angeles
02:00 AM	04:00 PM	11:00 AM	08:00 AM
12:00 PM	02:00 AM	09:00 PM	06:00 PM
05:00 PM	07:00 AM	02:00 AM	11:00 PM
08:00 PM	10:00 AM	05:00 AM	02:00 AM

Which sysplex members (cities) change their offsets first? The answer is Sydney: SYDA, SYDB, and SYDC. Since a new day will dawn in Sydney prior to in London and Atlanta, Sydney changes first. The operators in Los Angeles must be at Sydney's z/OS console to enter the SET TIMEZONE command at the correct instant.

What is the local time in Los Angeles when Sydney's DST offset changes? Remember that Sydney is on DST in the southern hemisphere. Assuming that Sydney wants the offset changed at 02:00 a.m. Sydney time, the Los Angeles operator must enter the command at 08:00 a.m. After the offset change, Sydney's local time will be changed from 02:00 a.m. to 01:00 a.m. (back one hour).

Several hours later, the same thing must be done for the London sysplex members. Again, assuming that London will change the offset at 02:00 a.m. London time, the Los Angeles operator must enter the command at 06:00 p.m. The only difference is that the northern hemisphere is changing the offset in the opposite direction from Sydney. London is going from standard time to Daylight Saving Time (forward).

Sometime later, it is Atlanta's turn. Atlanta is also going from standard time to Daylight Saving Time. The Los Angeles operator must enter the SET CLOCK command on Atlanta's z/OS console at 11:00 p.m.

Finally, the offset for Los Angeles can be changed through the Server Time Protocol, Adjust time zone offset panel. These images have STPZONE YES in their CLOCKxx members.

Note: Each customer situation is different. Certain ones are much more complicated, and others less so. Planning and excellent communications with the user community are essential when Daylight Saving Time schedules is a prime concern.

For some countries, time zones are not aligned on hourly divisions. For example, the time zone might be E.07.30.00 or W.08.45.00.

After a SET CLOCK command has been entered, change the TIMEZONE parameter in the CLOCKxx member in each affected image to reflect the new DST offset. The next IPL uses the new TIMEZONE value, and z/OS recognizes the correct local time for that image.

3.4 Managing the CTN configuration

This section discusses tasks associated with operational changes to a CTN. It discusses the following topics:

- ▶ Changing the CTN ID
- ▶ Changing the Current Time Server
- ▶ Changing the server roles
- ▶ Deconfiguring a timing network

3.4.1 Changing the CTN ID

STP supports a change to the CTN ID in either a Mixed or STP-only CTN from one value to another. If, for example, the CTN ID contains a value that is no longer relevant to the company it can be changed dynamically without a CTN outage.

Considerations for a CTN ID change in a Mixed CTN

In a Mixed CTN, a CTN ID change can be made using the STP Configuration tab. However, the only field within the CTN ID that can be changed dynamically is the STP portion of the CTN ID. The second field (the ETR portion) requires an outage of all systems within the CTN to force this change. An individual change is necessary for every server in the CTN. No global change is possible.

Important: XCF tolerates a mismatch of the CTN ID between systems for a matter of seconds only. When the time limit is reached, those systems with inconsistent CTN IDs are varied out of the sysplex. Consequently, where multiple servers are involved, this is not a recommended operation because the time to individually perform the sequential actions might exceed the XCF time limit.

Considerations for a CTN ID change in an STP-only CTN

This is a global change made from the Network Configuration tab of the Current Time Server and applies to the entire CTN. The new CTN ID is propagated in a coordinated fashion to all servers within the CTN. The change must be made from the Current Time Server. Otherwise, the request is rejected.

During a CTN ID change, z/OS systems and coupling facilities running on the servers might not all recognize the CTN ID change at exactly the same time. This can cause an inconsistent timing source scenario to occur between various components in the sysplex until the new CTN ID has been fully implemented across the CTN.

This temporary condition is recognized, tolerated, and automatically resolved. During this period various timing-related error messages can be sent to the console and syslog by both XCF and XES as inconsistent CTN IDs between z/OS logical partitions and coupling facilities are detected, respectively.

3.4.2 Single-server or dual-server auto resume after power-on reset

One of the reasons to implement STP, even on a single System z server, is to provide the same time across heterogeneous platforms using NTP client support. When an STP-only CTN is formed of a single server, it is possible as part of normal operation to deactivate or power-on reset the single server that is in the CTN.

In such a configuration, the STP configuration is saved across PORs to eliminate the need to initialize the time and reassign the PTS role after each POR. This can be accomplished by selecting the **Only allow the server(s) specified above to be in the CTN** option, as shown in Figure 3-32.

Note: For brevity's sake, the capability implemented by selecting **Only allow the server(s) specified above to be in the CTN** is referred to as the *save configuration* feature and CTNs for which this capability is selected are referred to as *bounded* CTNs.

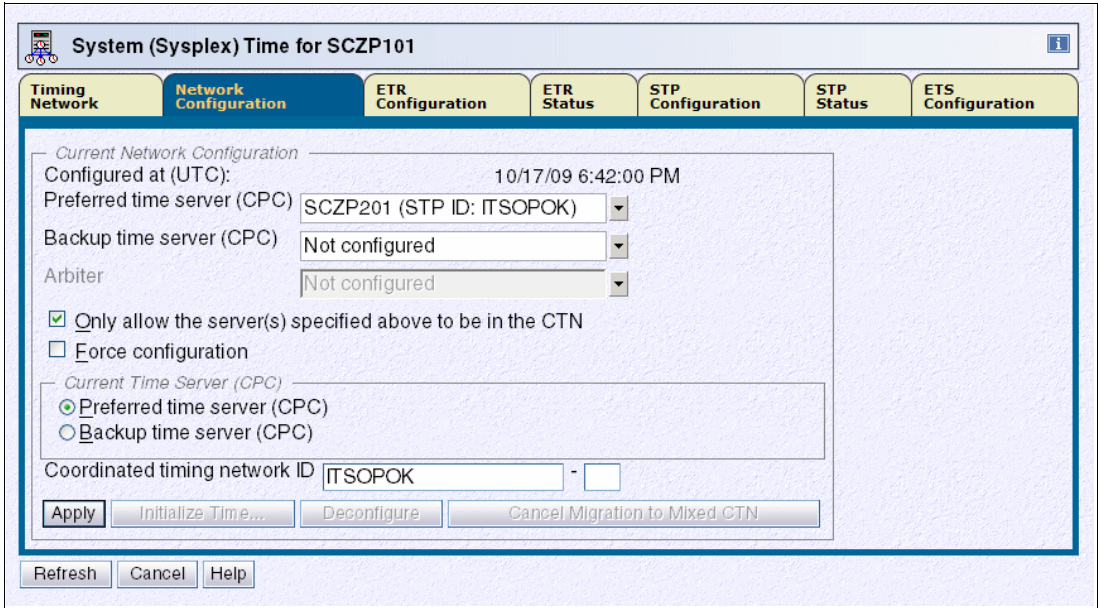


Figure 3-32 Server SCZP201 configured in a single server STP-only CTN

The option is available in a CTN consisting of either a single server or two servers. Figure 3-33 shows the Network Configuration tab for a dual-server CTN.

Note: It should be noted that after a power-off/on sequence, the TOD value in the Current Time Server is initialized from the TOD values stored in its SE. If you previously configured an ETS, STP code does not perform an initialize time using the ETS, but automatically accesses the configured ETS and starts adjusting the time so as to maintain time accuracy. STP allows time adjustments up to +/- 60 seconds.

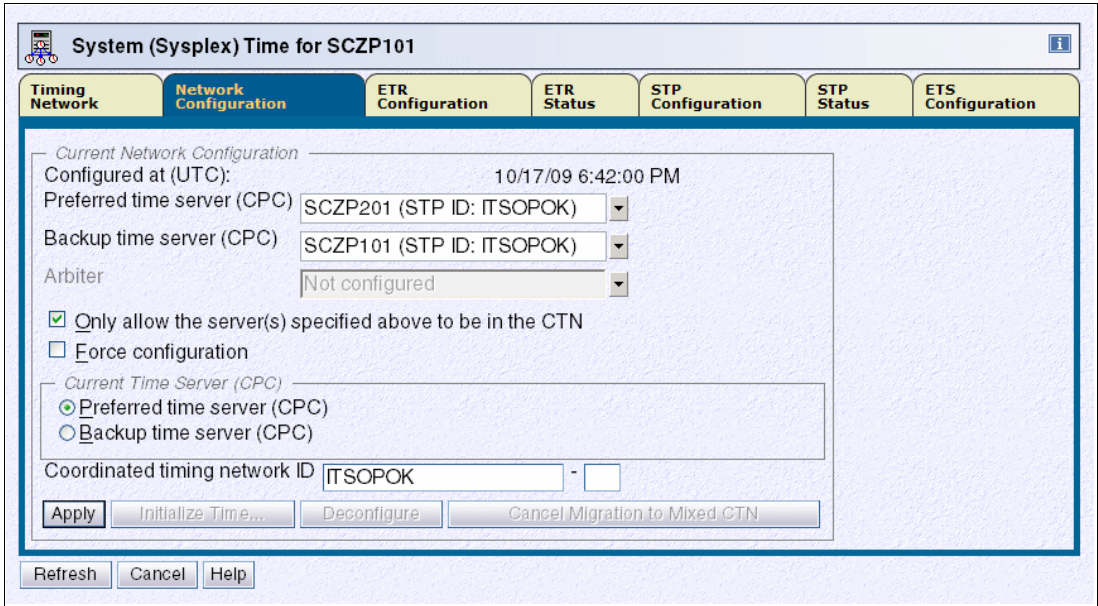


Figure 3-33 System (Sysplex) Time for SCZP101

When the option is checked, STP prevents any other server from joining the CTN. If you want to expand the single-server or dual-server CTN to three or more, you must deselect this option and click **Apply** before attempting to configure additional servers.

3.4.3 Changing the Current Time Server

STP supports dynamically changing the Current Time Server to a different server in an STP-only CTN, for example, for planned maintenance when there is a need to perform a disruptive action on the Current Time Server.

Before changing the CTS role, consider whether the Internal Battery Feature (IBF) is present. The CTS can take advantage of the IBF to provide improved recovery against a power outage (see 2.4, “Internal Battery Feature (IBF)” on page 57). If the IBF is installed on one or more servers in the CTN, reconfigure the CTS to a server that has the IBF installed, unless it interferes with other availability rules.

Changing the Current Time Server is done without disruption to the z/OS and coupling facility components within the CTN. The message IEA395I is issued (z/OS V1.11, rolled back to z/OS 1.10 and 1.9 with PTFs):

```
IEA395I THE CURRENT TIME SERVER HAS CHANGED TO THE cccccccccc
```

Where ccccccc is BACKUP or PREFERRED.

From z/OS you can determine which server is the Current Time Server by issuing a DISPLAY ETR command and checking which server is performing the stratum 1 role, as shown in Example 3-20.

Example 3-20 z/OS DISPLAY ETR command

```
D ETR
IEA386I 14.50.16 TIMING STATUS 931
SYNCHRONIZATION MODE = STP
  THIS SERVER IS A STRATUM 2
  CTN ID = ITSOP0K
  THE STRATUM 1 NODE ID = 002097.E26.IBM.02.000000088888
  THIS IS THE BACKUP TIME SERVER
```

From the HMC, the Network Configuration tab within the System (Sysplex) Time panel is used to change the CTS assignment from one server to another (Figure 3-34).

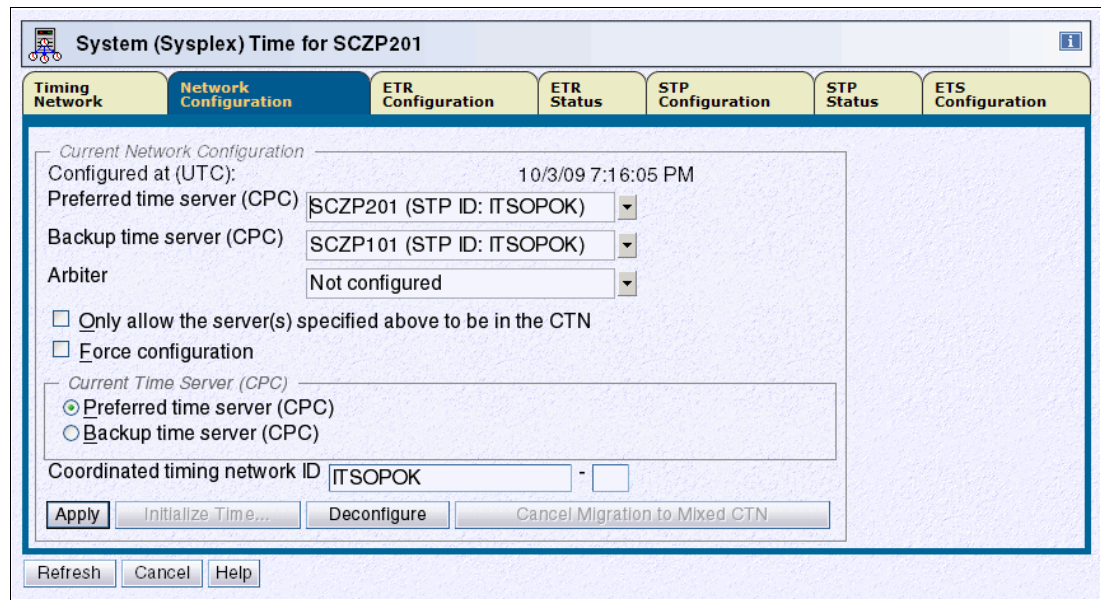


Figure 3-34 HMC workplace: System (Sysplex) Time, Network Configuration

We recommend that you have the Preferred Time Server assigned the Current Time Server role under nominal conditions because the Backup Time Server can only take over the CTS role under PTS-*failed* conditions and therefore cannot re-assume the CTS role after recovery.

However, there might be specific situations, such as maintenance, that require you to reassign the Current Time Server. You can use any of three different methods to move the Current Time Server role to another server in the CTN, depending upon the requirements:

- ▶ Switch the Current Time Server from the Preferred Time Server to the Backup Time Server.

This facility is provided specifically for a scenario in which you want to remove the Preferred Time Server for maintenance purposes and must relocate the Current Time Server function for the duration. Once the maintenance has been completed, reverse the process to restore the original CTN configuration.

- ▶ Assign the Preferred Time Server role to a new server and make it the Current Time Server.

This method utilizes a server role change (see “Reconfiguring Preferred, Backup, and Arbiter Time Servers” on page 130 for details) in order to change the Current Time Server assignment.

This can be used in a scenario where permanent changes are made to the CTN configuration. Depending on the requirements, a number of incremental changes could be made or combined into a single network configuration change.

- ▶ Switch the Current Time Server from the Preferred Time Server to the Backup Time Server and at the same time reconfigure the CTN to define a new Backup Time Server.

This is a combination of the previous two options, and can be done as a single reconfiguration change. This can be used in a situation where a permanent CTN configuration change is made but the server that will become the Preferred Time Server is not ready to assume the Current Time Server role at the time the change is made.

Regardless of which method is chosen to move the Current Time Server, there is one rule that applies in all network reconfiguration circumstances: the CTN configuration changes must be done from the server that will become the Current Time Server when the reconfiguration is complete. An attempt to make the change from any other server is rejected.

3.4.4 Changing the server roles

The STP-only CTN roles of Preferred Time Server, Backup Time Server, and Arbiter are displayed and modified using the Network Configuration tab within the System (Sysplex) Time task at the HMC.

The restrictions associated with changing the CTN roles are the same as when initially converting to an STP-only CTN:

- ▶ Only the Preferred Time Server role must be defined. This server defaults to become the Current Time Server. Running without a Backup Time Server is not advised because the Preferred Time Server becomes a single point of failure in the CTN.
- ▶ An Arbiter can only be defined if a Backup Time Server has also been defined. The roles of Backup Time Server and Arbiter can be removed by making these *not configured*.
- ▶ No server can assume multiple roles.

Before changing the CTS role, consider whether the Internal Battery Feature is present. The CTS can take advantage of the IBF to provide improved recovery against a power outage (see 2.4, “Internal Battery Feature (IBF)” on page 57). If the IBF is installed on one or more servers in the CTN, reconfigure the CTS to a server that has the IBF installed, unless it interferes with other availability rules.

Note: The same rule applies to changing the CTN server roles as to changing the Current Time Server: The CTN configuration changes must be done from the server that *will become* the Current Time Server when the reconfiguration is complete.

Reconfiguring Preferred, Backup, and Arbiter Time Servers

All the existing server roles within an STP-only CTN can be redefined in a single reconfiguration. There is no requirement to stage changes in increments, although this is possible. The reconfiguration must be done from the server that will become the Current Time Server once the reconfiguration is complete.

Force configuration

There is, however, a situation where a reconfiguration might not be allowed to proceed due to connectivity requirements not being satisfied.

STP checks that the following connectivity is in place before allowing a reconfiguration to be applied:

- ▶ Timing link connectivity between the new Preferred Time Server and the new Backup Time Server (if defined)
- ▶ Timing link connectivity between the new Preferred Time Server and the new Arbiter (if defined)
- ▶ Timing link connectivity between the new Backup Time Server (if defined) and the new Arbiter (if defined)

There might be certain situations where the connectivity requirements are not being met at the time that a reconfiguration must be performed. This could be the result of a hardware maintenance outage that has temporarily made part of the CTN unusable.

In such situations use the **Force configuration** check box to bypass this checking.

Attention: Use the Force configuration option with care because it might unintentionally implement an STP-only CTN that is not tolerant of Current Time Server failure.

The Force configuration option only bypasses the connectivity checking as outlined here. It does not allow reconfigurations that are invalid, such as a CTN with an Arbiter defined but no Backup Time Server specified.

Configuring a Backup Time Server

Definition of a Backup Time Server is optional. Defining a Backup Time Server is done from the Network Configuration tab by selecting the required server in the Backup time server (CPC) box and clicking **Apply**.

Because this is a network configuration change, it must be done from the Current Time Server.

Recommendation: Configuring a Backup Time Server is optional, but recommended, because otherwise the Preferred Time Server is a single point of failure in the CTN.

A Backup Time Server requires direct timing link connectivity to the Preferred Time Server as well as to the Arbiter, if defined. This is verified as part of the reconfiguration process unless the Force configuration option has been selected.

There are no messages sent to the z/OS console during the definition of the Backup Time Server, in contrast to removal of the Backup Time Server, which causes message IEA389I to be displayed. However, as an alternative, use the z/OS DISPLAY ETR command before the reconfiguration to verify that the CTN does not have a Backup Time Server (Example 3-21).

Example 3-21 z/OS command DISPLAY ETR

```
D ETR
IEA386I 15.12.27 TIMING STATUS 975
SYNCHRONIZATION MODE = STP
  THIS SERVER IS A STRATUM 1
  CTN ID = hmctest
  THE STRATUM 1 NODE ID = 002097.E26.IBM.02.00000001DE50
  THIS IS THE PREFERRED TIME SERVER
  THIS STP NETWORK HAS NO BACKUP TIME SERVER
  THIS STP NETWORK HAS NO SERVER TO ACT AS ARBITER
```

Once the reconfiguration has completed, the absence of the line pertaining to the Backup Time Server indicates that the CTN now has a BTS configured, as shown in Example 3-22.

Example 3-22 z/OS command DISPLAY ETR

```
D ETR
IEA386I 15.15.43 TIMING STATUS 982
SYNCHRONIZATION MODE = STP
  THIS SERVER IS A STRATUM 1
  CTN ID = hmctest
  THE STRATUM 1 NODE ID = 002097.E26.IBM.02.00000001DE50
  THIS IS THE PREFERRED TIME SERVER
  THIS STP NETWORK HAS NO SERVER TO ACT AS ARBITER
```

Configuring an Arbiter

Definition of an Arbiter is optional. Defining an Arbiter is done from the Network Configuration tab by specifying the required server in the Arbiter box and clicking **Apply**.

Because this is a CTN configuration change it must be done from the Current Time Server.

Recommendation: Configuring an Arbiter is optional, but recommended, to enhance the failure detection and recovery capabilities of an STP-only CTN.

If you intend to include an Arbiter in an STP-only CTN, ensure that a Backup Time Server is also defined. Otherwise, the reconfiguration is rejected.

The Arbiter's role is to assist in reconfiguring the Current Time Server role from the Preferred Time Server to the Backup Time Server in recovery scenarios. The Arbiter needs direct timing link connectivity to the Preferred Time Server as well as to the Backup Time Server. This is verified as part of the reconfiguration process. This connectivity check can be bypassed if required by selecting **Force configuration** on the Network Configuration tab. However, this should not be used unless absolutely necessary.

No z/OS messages are produced during definition of an Arbiter. As with the Backup Time Server, issue the z/OS DISPLAY ETR command to determine whether the STP-only CTN has an Arbiter defined (Example 3-23).

Example 3-23 z/OS command DISPLAY ETR

```
D ETR
IEA386I 15.18.08 TIMING STATUS 993
SYNCHRONIZATION MODE = STP
  THIS SERVER IS A STRATUM 1
  CTN ID = hmctest
  THE STRATUM 1 NODE ID = 002097.E26.IBM.02.00000001DE50
  THIS IS THE PREFERRED TIME SERVER
```

Again, absence of a line pertaining to the Arbiter is an indication that an Arbiter is defined.

Configuring the same server to multiple CTN server roles

Each of the server roles in the CTN must be allocated to a different server or set to a value of Not Configured, except for the Preferred Time Server, which must be defined.

If the configuration request tries to allocate the same server to multiple roles in an STP-only CTN, the request is rejected.

Removing the Preferred Time Server

All STP-only CTNs must have a Preferred Time Server defined. Any attempt to set the Preferred Time Server to Not Configured in an initialized CTN is rejected as an invalid configuration. The only time that the Preferred Time Server role can be removed is when the CTN has been de-configured.

Removing the Backup Time Server

There might be a requirement to remove the Backup Time Server role from the CTN. This can be achieved by setting the Backup time server field to Not Configured on the Network Configuration tab.

The removal of the BTS fails if:

- ▶ An Arbiter is defined, because it is an unsupported configuration to have an Arbiter without a Backup Time Server.
- ▶ The Backup Time Server is performing the Current Time Server role, because this would leave an STP-only CTN without a clock source.

In both of these cases, additional CTN reconfiguration must be done before the Backup Time Server can be removed.

If the reconfiguration is successful, message IEA389I is displayed on each z/OS console, warning that the CTN no longer has a Backup Time Server available (Example 3-24).

Example 3-24 z/OS message IEA389I

```
IEA389I THIS STP NETWORK HAS NO SERVER TO ACT AS BACKUP
```

Use the DISPLAY ETR command to determine whether the CTN currently has a Backup Time Server defined (Example 3-25).

Example 3-25 z/OS command DISPLAY ETR

```
D ETR
IEA386I 15.12.27 TIMING STATUS 975
SYNCHRONIZATION MODE = STP
  THIS SERVER IS A STRATUM 1
  CTN ID = hmctest
  THE STRATUM 1 NODE ID = 002097.E26.IBM.02.00000001DE50
  THIS IS THE PREFERRED TIME SERVER
  THIS STP NETWORK HAS NO BACKUP TIME SERVER
  THIS STP NETWORK HAS NO SERVER TO ACT AS ARBITER
```

The presence of the last two messages in Example 3-25 indicates that both Backup Time Server and Arbiter roles are not currently defined. Note that having no Backup Time Server implies that there is no Arbiter.

Removing the Arbiter

The Arbiter can be removed from an STP-only CTN. Unlike the Backup Time Server, the Arbiter can be removed without regard to the Current Time Server role.

If the reconfiguration is successful, a message is displayed on each z/OS console warning that the CTN no longer has an Arbiter available (Example 3-26).

Example 3-26 z/OS message IEA389I

```
IEA389I THIS STP NETWORK HAS NO SERVER TO ACT AS ARBITER
```

Use the DISPLAY ETR command to determine whether the CTN currently has an Arbiter defined (Example 3-27).

Example 3-27 z/OS command DISPLAY ETR

```
D ETR
IEA386I 15.23.29 TIMING STATUS 985
SYNCHRONIZATION MODE = STP
  THIS SERVER IS A STRATUM 1
  CTN ID = hmctest
  THE STRATUM 1 NODE ID = 002097.E26.IBM.02.00000001DE50
  THIS IS THE PREFERRED TIME SERVER
  THIS STP NETWORK HAS NO SERVER TO ACT AS ARBITER
```

3.4.5 Deconfiguring a CTN

The HMC provides a facility to deconfigure the STP-only CTN, which removes the roles of Preferred Time Server, Backup Time Server, and Arbiter. This facility is accessed from the Network Configuration tab and must be used on the Current Time Server. In essence, this operation makes all servers stratum 0. It results in the loss of the clock source for all servers in the STP-only CTN, which in turn causes the loss of all active sysplex images and access to coupling facilities using CF request time ordering.

3.5 Operational considerations

The loss of timing synchronization for servers in a Coordinated Timing Network might have grave consequences for system images running in a Parallel Sysplex. Successful operation of a CTN requires thoughtful consideration of several time-related decisions that include:

- ▶ Migrating timing-only links to coupling links
- ▶ Validating coupling link paths
- ▶ What to expect when attempting to disable a server's last coupling link
- ▶ The effect of performing disruptive actions on the CTS
- ▶ Guidelines to follow when restarting the CTN after a power outage

Safeguards have been added to prevent:

- ▶ An accidental deconfiguration of the last timing link between any two servers
- ▶ Execution of a disruptive task on a server that is configured as Current Time Server and has STP connectivity to stratum 2 servers
- ▶ Execution of a disruptive task on a server that is configured as the BTS or Arbiter.

With these new safeguards, the user must complete additional steps prior to executing the disruptive action. New messages draw the attention of the user when certain STP configuration conditions are not met. These conditions are described in the following sections.

3.5.1 Migrating timing-only links to coupling links

STP messages are transferred over external coupling channels as described in 2.2.1, "Planning for coupling links for STP" on page 38. These can be defined in one of two ways:

- ▶ Timing-only links: Only STP messages are transferred across the link.
- ▶ Coupling (CF) links: STP and CF messages are transferred across the link.

In the event of a requirement to concurrently migrate timing-only links to coupling links, for example, Parallel Sysplex configuration changes, there are certain guidelines that must be adhered to:

- ▶ A CF partition must be in the access list of the CHPID at one end of the link. It is not valid to define a coupling link with only z/OS logical partitions connected.
- ▶ The configuration change must be made by HCD for servers at both ends of the link. If these servers use different IODFs, both IODFs require modification.
- ▶ The IODF is changed by initially disconnecting the timing-only link via HCD, to remove the STP control unit at each end of the link.

Important: We recommend having a consolidated IODF to manage all servers. This shows the end-to-end connection for these links. When separate IODFs are used, dummy processor definitions must be used to generate the CNTLUNIT and DEVICE definitions for timing-only and coupling links.

When reconnecting the channel path, the *timing-only link* field must be left as No to define a coupling link.

- ▶ A production IODF must be built to reflect the changes for each server.

A dynamic reconfiguration change is required at each end of the link. This activation of the new IODF can be done concurrently. The channels that will be changed from timing-only to coupling links can remain online because the channel type will not change, so there is no

interruption to STP. A change at only one end of the link puts the link into a transient state that remains until both ends are of the same control unit type.

Important: A dynamic IODF activation causes the links to go offline/online. This means a dynamic activation on the PTS will cause the BTS to take the S1 role momentarily, and cause a sync check when the PTS takes back the S1 role. The links going offline can result in IEA382I, IEA383I and IEA388I messages.

It is also possible to change coupling links to timing-only links with a dynamic reconfiguration change. When activating this IODF configuration the 'FORCE'-option must be used because CF devices will be deleted.

3.5.2 Coupling link path validation

The CHPID definition of a coupling link only requires one logical partition in its access list to establish an STP communication path to another server. Although more links are required for redundancy, a configuration with only one coupling link defined provides the foundation for an entire server to send and receive STP information over that link. In this discussion, no distinction is made between coupling links that transmit both CF and STP messages or STP messages only.

STP messages flow from one server to another but are not directed to any logical partition. As a consequence, STP messaging requires the CHPIDs to be online at both ends of the link but does not require any logical partition in the CHPID access list or candidate list to be activated.

The path is available for STP messaging as long as the CHPID is online to any logical partitions in its access list, whether the partition is activated or not. The path becomes unavailable when one CHPID at the end of the link is configured offline to the last logical partition in its access list or candidate list.

Irrespective of the number of coupling links available between a pair of servers, only one is used at any time. Multiple coupling links are configured for availability. There is no performance benefit gained by adding more than one. In any case, the volume of STP messages transferred between servers is small and will not drive high channel utilization.

A coupling link can be initialized for STP timing or actively used for STP timing:

- ▶ A link is *initialized* for STP timing when it has matching CTNIDs configured on both ends of the link, but it is not actively sending or receiving STP timing messages.
- ▶ A link is *actively used* for STP timing when it has matching CTNIDs configured on both sides and STP signals are being exchanged over the link.

Note: Remember also that when a CHPID is configured OFF to an inactive logical partition, the effect is immediate. However, when it is configured ON to an inactive logical partition, the LPAR will queue the request and the CHPID will only be brought online the next time that the logical partition is activated.

The STP Status panel on the HMC displays the initialized links for each server. No indication is provided of which link is actively used. If the active link fails or is removed, STP automatically switches to another available initialized link. This is an automatic process without impact to operating systems.

3.5.3 Last timing link validation and operational safeguards

The loss of a server's last coupling link might cause the loss of timing synchronization for all system images on the server. To prevent against an accidental misconfiguration of the last usable coupling link between any two servers, additional safeguards are provided to prevent operational errors.

The last coupling link that is being used to deliver STP timing between any two servers is also referred to as the *last timing link*. A last timing link condition occurs when the following two conditions are true:

- ▶ The path is initialized for STP timing messages, whether it is defined as a coupling link or as a timing-only link. The associated PCHID is listed on the Hardware Management Console STP Status tab.

To identify a physical channel, a physical channel identifier (PCHID) is assigned to each possible location that can support a channel card or that can provide I/O connectivity (for example, ESCON or FICON, networking, and coupling interfaces) or a logical interface, such as cryptographic attachments. PCHIDs represent the physical location of the I/O ports within the server. CHPIDs are then mapped to PCHIDs.

Note: The PCHID is initialized for STP timing messages, whether it is defined to HCD as a coupling link or as a timing-only link.

- ▶ It is the last online coupling link CHPID between any two servers participating in a Mixed CTN or an STP-only CTN.
 - If the CHPID is not shared or spanned, it is online to only one logical partition.
 - If the CHPID is not shared or spanned but is reconfigurable, it can be moved across logical partitions, but is online to only one logical partition.
 - If the CHPID is shared or spanned, more than one logical partition can access it concurrently, as determined by the CHPID access and candidate lists. In this situation the last path condition is raised when the CHPID remains online to only one logical partition.

Note: A CIB^a, CBP^b, or CFP^c channel path (CHPID) can be shared or spanned but can only be online to one active coupling facility logical partition at a time. If such a channel path is targeted to a coupling facility partition but is already online to another coupling facility partition, then it is removed (deconfigured) from the activating logical partition.

- a. Coupling over InfiniBand (PSIFB)
- b. Cluster Bus Coupling (ICB-3, ICB-4)
- c. Coupling Facility Peer (ISC-3)

While XCF communication requires an active logical partition at each end of the link, either z/OS image or CFCC, STP messaging does not.

Note: The status of the logical partitions defined in the CHPID access list is not used to determine whether a last timing link condition exists. STP messages will flow on any coupling link as long as CHPIDs at both ends of the link are online, regardless of whether the logical partition that the CHPIDs are defined to is active.

Unfortunately, the CTN topology is not readily available in one single place in the system. It must be reconstructed from various HMC and z/OS sources (see Appendix B, "Server Time

Protocol (STP) messages in z/OS” on page 199, for information regarding how to draw a CTN topology).

To evaluate and plan for operations that deal with last timing link situations, a CTN configuration diagram with PCHID and CHPID information is a good starting point. It can be done manually or be derived from HCD or HCM if the installation uses a single IODF. However, given the real-life constraints of any IT installation, a diagram might not always reflect the status of the CHPID access lists at that precise moment in time.

Attention: Before attempting to override a last timing link condition, review and understand the DISPLAY ETR output from ALL z/OS system images on each server, and check the STP connectivity information using both the STP Status tab and the CHPID access list information from the Configure Channel Path On/Off panels.

3.5.4 Last timing link condition examples

Considering the Mixed CTN illustrated in Figure 3-35, the coupling link between server A and server C is a potential last path. The coupling link between server B and server C is also a potential last path. This means that servers A and B each have one potential last path. Also, whereas server C has two coupling links, each can also be considered a potential last path.

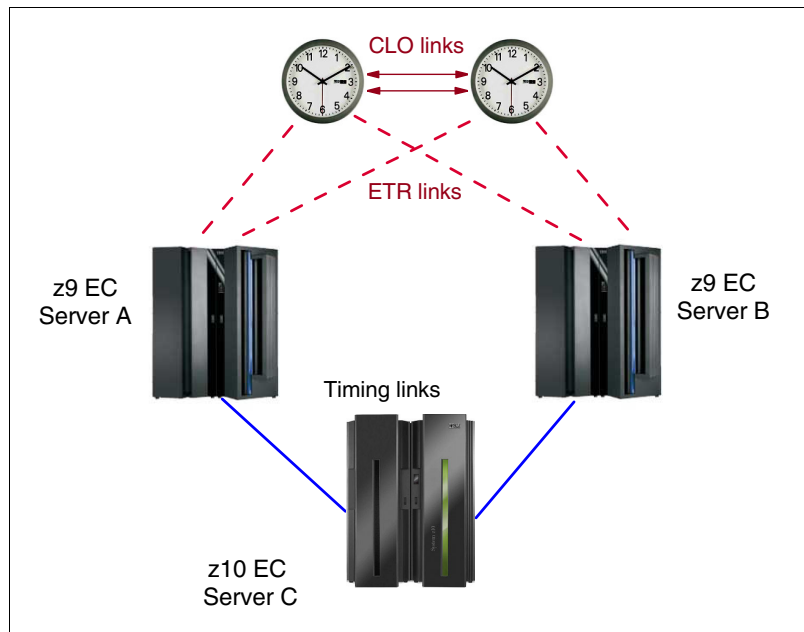


Figure 3-35 Last path validation of coupling links

When the last path condition is raised, any attempt to remove the path from a z/OS image from either end of the link using the z/OS command CONFIG CHP(xx),OFF results in the request being rejected and message IEE148I being issued:

```
IEE148I CHP(xx) NOT RECONFIGURED - WOULD REMOVE A CPC-CRITICAL STP TIMING LINK
```

Note: Remember also that when a CHPID is configured OFF to an inactive logical partition, the effect is immediate. However, when it is configured ON to an inactive logical partition, the LPAR queues the request and the CHPID is brought online the next time that the logical partition is activated.

The Config Off request is also rejected if the FORCE option is specified.

If a last path condition exists on all three servers shown in Figure 3-35, the IEE148I message is issued whether the reconfiguration request is initiated from a z/OS image on server A, on server B, or on server C. If the request is initiated from a z/OS image on server C, message IEE148I is issued for the A–C path, and also for the B–C path.

Message IEE148I is intended as a warning that the action can potentially disrupt the entire CTN and the request is rejected. However, removing the last coupling link between two servers or CFs might be a valid request, for example, for scheduled maintenance or reconfiguration of a server.

When a last path condition exists, the configuration change can only be accomplished from the Hardware Management Console or the Support Element. From the HMC, the first attempt to configure the CHPID Off is rejected with a warning message, but an override option is available. If the override option is selected, the request can be reissued and the CHPID reconfiguration proceeds. However, the decision to override should not be made lightly. The following example illustrates why:

1. Reconfiguration of coupling link B–C is requested, either from server B or server C in Figure 3-35 on page 137.

In this case, the CHPID reconfiguration can be done from the HMC without any damaging consequences to the server C timing. STP messages still flow using the A–C link. However, the resulting configuration, shown in Figure 3-36, now leaves only one coupling link available to server C.

2. If an attempt is made later to configure off the coupling link A–C, whether from server A or server C, the same warning message is issued again.

However, the circumstances are now different. In this situation, if the warning is ignored and the CHPID reconfiguration proceeds, server C loses time synchronization and switches to an unsynchronized state (STP stratum 0). All active z/OS images on server C participating in a Sysplex post WTOR IEA394A. See *STP Recovery Guide*, SG24-7380, for a description of z/OS message IEA394A. If the server has been assigned a role in the CTN, this could have consequences for the recovery of the entire CTN.

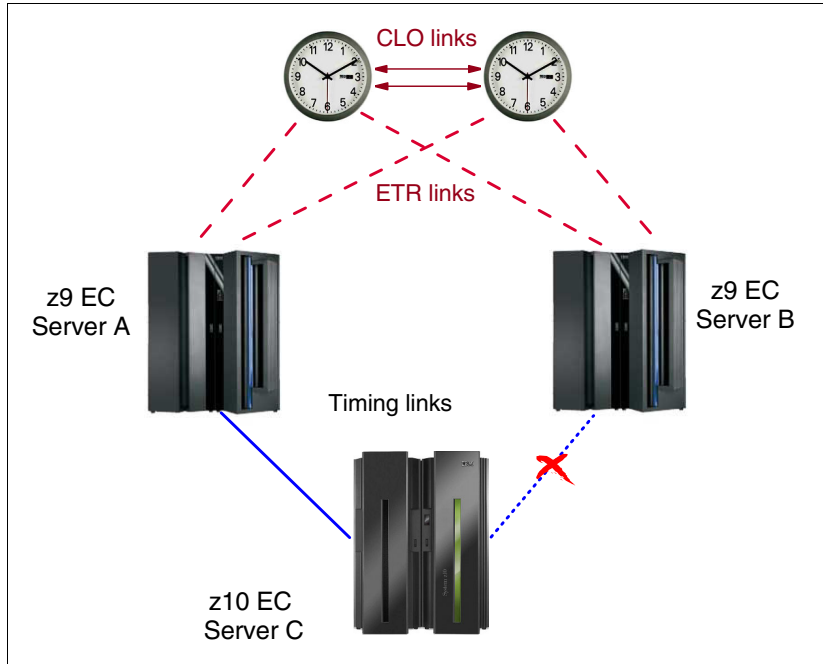


Figure 3-36 Last path validation of coupling links: one remaining link

Over time, the list of channel paths targeted to be configured online or offline to a given logical partition can be changed by system control program configuration commands, configure CHPID On and Off tasks, or dynamic I/O configuration commands. Similarly, reconfigurable unshared channel paths can be moved from one logical partition to another using the same commands, changing the owner of the unshared channel path. Configuration changes are remembered by the Support Element on an IOCDS basis.

To determine whether it is safe to continue with a last path reconfiguration request, the system programmer must assess the consequences to the configuration. The decision requires an understanding of the CTN topology and precise knowledge of what the current timing distribution is at that specific moment in time.

In the STP-only CTN configuration displayed in Figure 3-37, the PCHIDs are indicated so that the reader can relate the STP Status information to the CTN topology. In this example, the coupling link B between server CEC1 and server CEC3 is a potential last timing link. The coupling link A between server CEC3 and server CEC2 is also a potential last timing link.

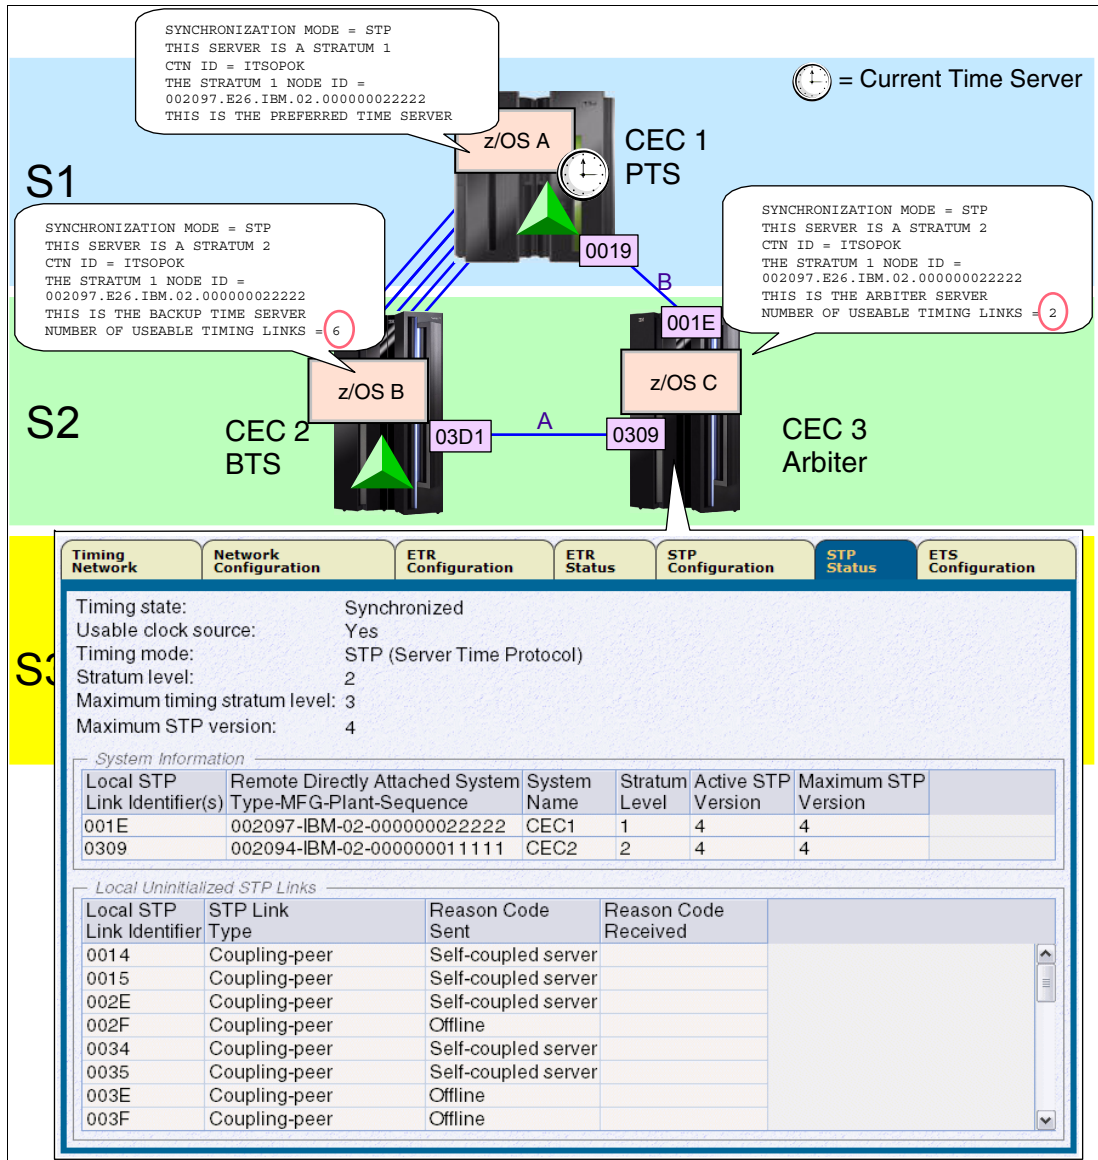


Figure 3-37 Last timing link example

This means that servers, CEC1 and CEC2, each have one potential *last timing link*. While server CEC3 has two coupling links available for STP messages, each can raise a last timing link condition because each of these links is the last link that is being used to deliver STP timing between two servers.

When a last timing link condition is detected, any attempt to remove the channel path from a z/OS image at either end of the link using the z/OS command CONFIG CHP(xx),OFF is rejected.

In the example shown in Figure 3-38, an attempt to configure a channel path off from a z/OS image on CEC3 results in message IEE148I:

IEE148I CHP(CO) NOT RECONFIGURED - WOULD REMOVE A CPC-CRITICAL STP TIMING LINK

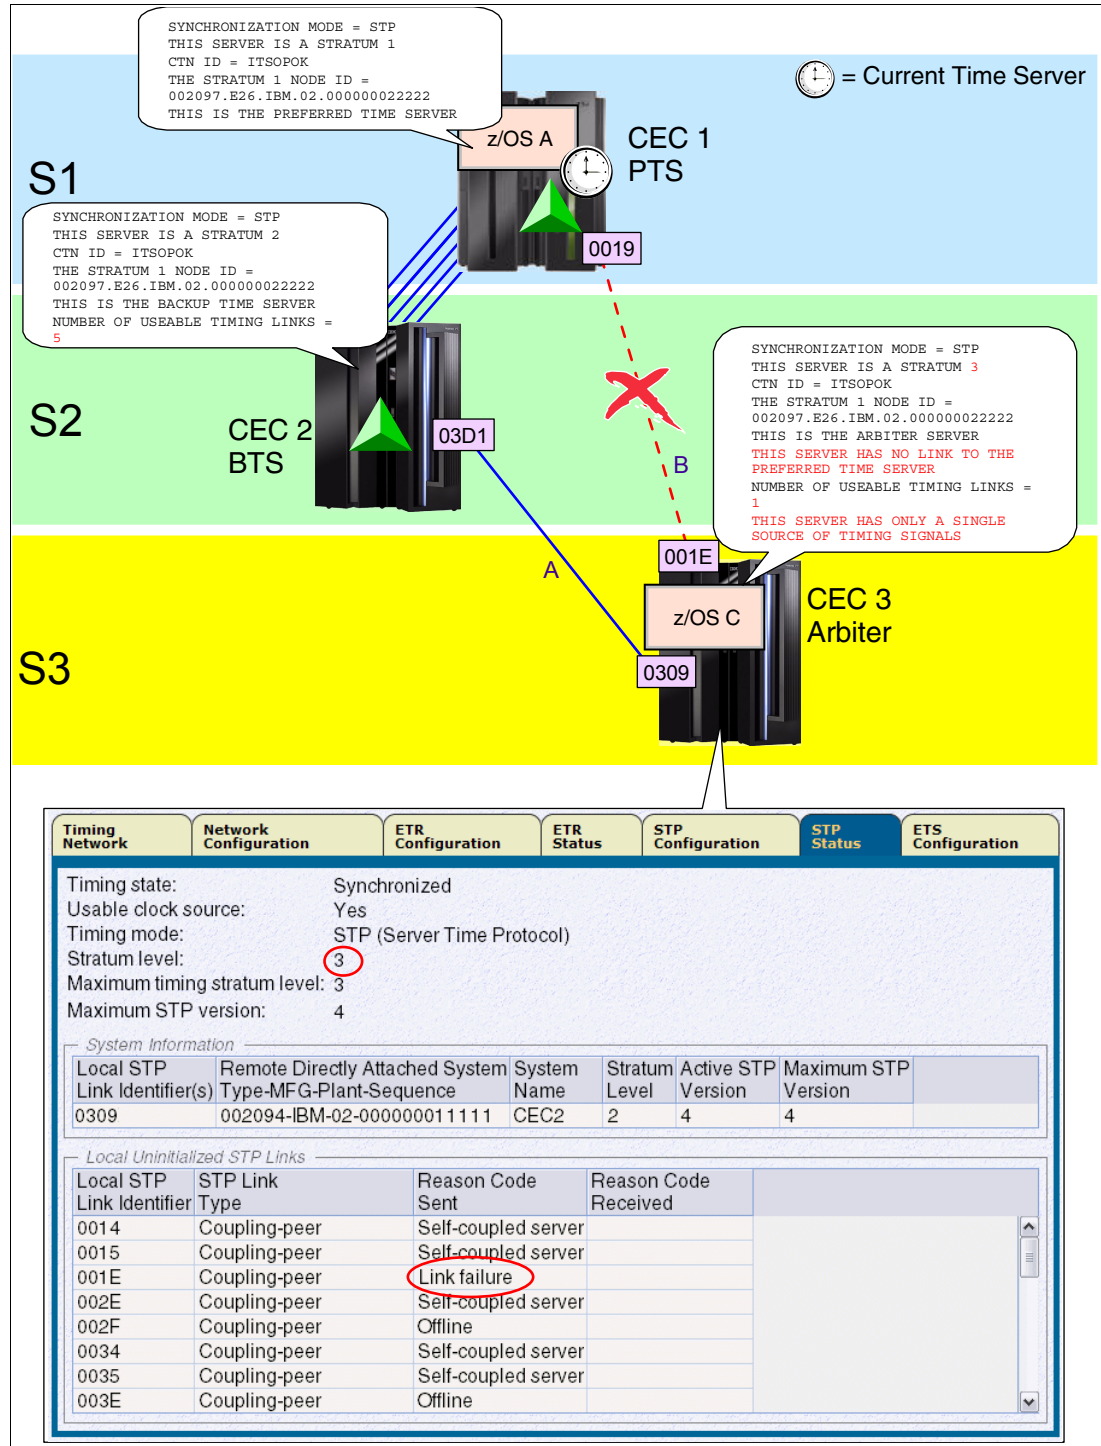


Figure 3-38 Last timing link example: link failure

The Config Off request is also rejected if the FORCE option is specified on the z/OS or CFCC command.

When a last timing link condition exists, the configuration change must be completed on the HMC or Support Element. This procedure is described in detail in the *Server Time Protocol Implementation Guide*, SG24-7281. Assessing the consequences of forcing a CHPID reconfiguration in a last timing link scenario requires the user to understand the CTN topology at that precise moment in time.

As shown in Figure 3-38 on page 141, there was no hard consequence because STP automatically adjusted to use alternate links. However, CEC3 is now a stratum 3 server and has only one single source and one single link remaining to receive STP messages. The resulting configuration now contains single points of failure.

The primary source of information for operations is z/OS messages. However, z/OS messages provide no information about the last timing link condition because a z/OS system image has no visibility to CHPIDs defined to partitions other than its own.

The output of the z/OS DISPLAY ETR command identifies coupling links (PCHIDs), and only as viewed from the z/OS system image.

- ▶ z/OS counts physical links (PCHIDs).
 - In the configuration shown in Figure 3-37 on page 140, on CEC3 the DISPLAY ETR output on image LPC identifies no single point of failure:

```
NUMBER OF USABLE TIMING LINKS = 2
```

Still, links A and B could raise a last timing link condition if only one online CHPID remains mapped to the PCHID on each server.
 - In Figure 3-38 on page 141, on CEC3, the DISPLAY ETR information about image LPC identifies a single point of failure:

```
NUMBER OF USABLE TIMING LINKS = 1  
THIS SERVER HAS ONLY A SINGLE SOURCE OF TIMING SIGNALS
```

The user must go to the STP Status tabs to map the link to PCHIDs 03D1 to 0309. Again, this does not automatically imply a last timing link condition on CEC3 (or at the other end of the link, on CEC1) because the last timing link condition depends on the contents of the CHPID access list.
- ▶ The DISPLAY ETR output provides information only for PCHIDs that are input to the server hosting the z/OS image. It does not consider coupling links to dependent systems. For example, consider the configuration shown in Figure 3-38 on page 141:
 - On CEC1, which has the Current Time Server role, the number of usable timing links is not indicated. Since the CTS owns the Current Server Time and does not rely on external links, links are not considered a single point of failure for images on the server.
 - On CEC2, the number of usable timing links at this point in time is now five, because z/OS only counts the number of PCHIDs from its host server to the Current Time Server, CEC1.

Notice the difference with the configuration in Figure 3-37 on page 140. Prior to the link failure, the z/OS image running on CEC2 recognized six usable timing links (PCHIDs) because, at the time, CEC3 had connectivity to the Current Time Server and was a possible source of STP timing messages.
- ▶ The DISPLAY ETR information is not available for stand-alone coupling facilities since the CFCC has no DISPLAY ETR equivalent.

3.5.5 Disruptive actions on the Current Time Server

A disruptive action on the CTS could jeopardize the entire CTN. Protection has been added to prevent a disruptive task from being performed on a server that:

- ▶ Is configured as the Current Time Server
- ▶ Has initialized STP links to stratum 2 servers

When these conditions are met, the server is locked for disruptive tasks such as activate, deactivate, power off, power-on reset, disruptive switch to alternate SE, and so on.

The user is also prevented from modifying the server STP ID and ETR Network ID on the STP Configuration tab, and modification of the ETR Network ID on the ETR Configuration tab.

For example, as illustrated in Figure 3-39, server SCPZ201 satisfies the two conditions. SCPZ201 is stratum 1 in the CTN and is the time source for other servers in the CTN.

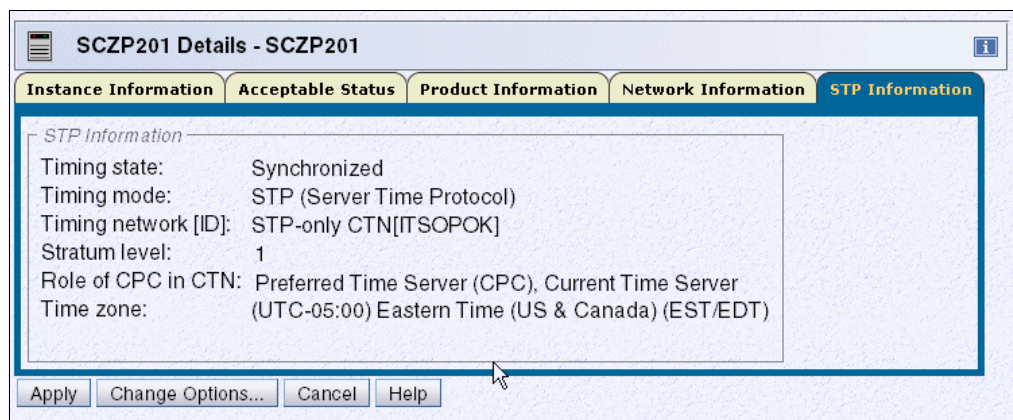


Figure 3-39 CPC Details: STP Information tab

The STP Status tab (Figure 3-40) shows that SCZP201 has initialized STP links to SCZP101, another server in the same CTN.

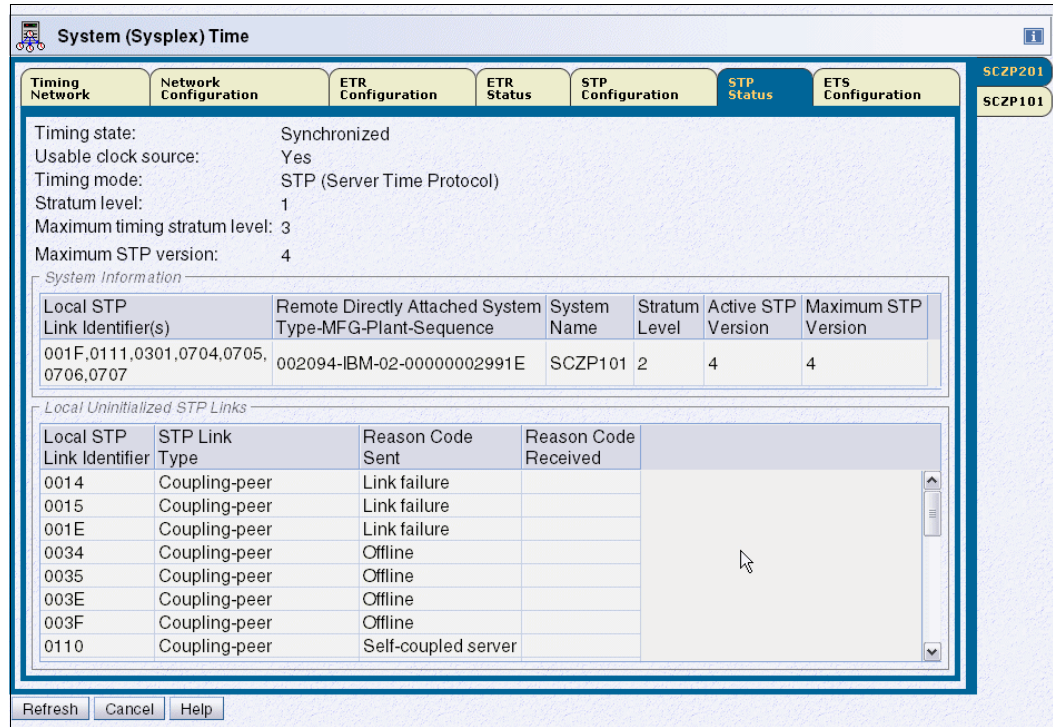


Figure 3-40 STP Status tab

Special checking is only done for the Current Time Server in an STP-only CTN. Only the combination of CTS role and active STP link conditions is checked. The disruptive-action detection does not check for other server roles in the CTN.

An attempt to deactivate the CTS returns a failed status, as shown in Figure 3-41.

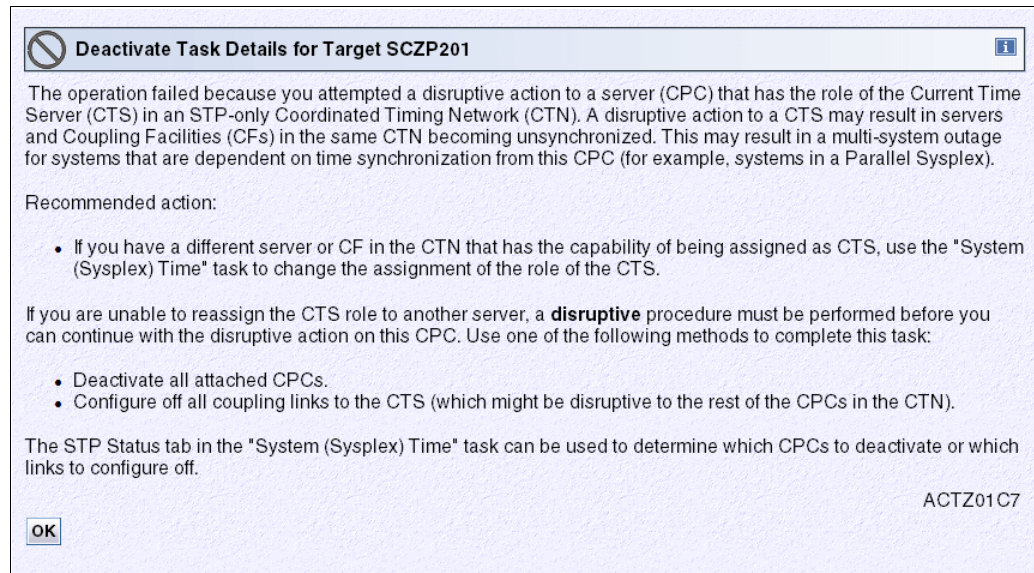


Figure 3-41 Deactivate Task details: message ACTZ01C7

No alternative action is possible from this panel. The user can only acknowledge that the request has failed and click **OK**.

If the request really must be pursued, the solution is to first change the STP configuration by reconfiguring the Current Time Server to another server from the Network Configuration page of the server that will become the new CTS. Only then can the deactivate request be accepted.

If the CTS *cannot* be re-assigned, the alternatives will be *disruptive* to maintaining synchronization for other servers in the CTN. The alternatives are to deactivate all the stratum 2 and stratum 3 servers in the CTN first. Once that is done the disruptive action on the CTS can be performed.

Configure off all the coupling link CHPIDs from the CTS to all its attached servers. The special procedure that must be used to configure off all, or more precisely the last CHPID used for STP timing is described in the following section. The disruptive action on the CTS can be performed after this action is completed.

3.5.6 Disruptive actions on the BTS or Arbiter

Disruptive actions are now blocked for the BTS and Arbiter. This new function was introduced with the MCL levels shown in Table 3-10.

Table 3-10 Blocked disruptive actions

Driver/Server	MCL	Bundle	Release Date
D86E / z196	N29809.277	45	Sept 8, 2011
D86E / z196	N29802.420	45	Sept 8, 2011
D79F / z10	N24415.078	50	Sept 28, 2011
D79F / z10	N24409.184	50	Sept 28, 2011
D93G / z114 and z196 GA2	Integrated	N/A	Sept 9, 2011

This provides the same safeguards for the BTS and Arbiter as described in 3.5.5, “Disruptive actions on the Current Time Server” on page 143. This prevents a disruptive action causing the CTS to give up the S1 role and go S0. For example, if the PTS has the CTS role and then the Arbiter (or BTS) has a planned or unplanned outage, a disruptive action on the BTS (or Arbiter) causes the CTS to give up the S1 role and go S0 as it loses communication to both the BTS and Arbiter.

Attention: This recent restriction that a disruptive action cannot be taken against the server assigned with the BTS role has changed how the “Only allow the server(s) specified above to be in the CTN” button is used because it requires that a BTS is assigned in a two server CTN. This means that the BTS server has to be removed from the CTN when the BTS role is removed to allow the “Only allow the server(s) specified above to be in the CTN” button to be selected when the PTS is PORed.

3.5.7 Restarting a CTN power outage or power-on reset (unplanned)

A fundamental aspect of an STP-only CTN is that the network design must make sure that it does *not* create an *island* condition, where two active stratum 1 servers exist within the same CTN. This is because a stratum 1 server does not synchronize to any other servers’ time in the CTN, but is the one with which all other servers are synchronized. If two stratum 1 servers

were allowed to exist in the same CTN, the software and applications running on them would consider the two stratum 1 servers to be synchronized, when in fact they most likely are not.

After a power outage is restored, when an individual server is powered up, it does not know whether the entire data center has been shut down, whether there has been a physical or logical reconfiguration during the shutdown, or whether the customer might not want to re-instate the same roles to the servers that existed prior to the power outage. Therefore, in order to make sure that an island condition is not inadvertently created after a power outage, the STP design:

- ▶ Deconfigures the CTN if the power outage has affected both the PTS and the BTS. This results in all servers in the CTN becoming unsynchronized stratum 0 servers.
- ▶ Requires an STP-configured server that is being powered-on to obtain its configuration information from the CTN that is already up and running, if one exists, rather than making an independent decision about its role in the CTN. So, if the power outage affects only one time server (PTS or BTS) and recovery actions have permitted the CTN to stay operational, the affected server can rejoin the CTN without requiring a restart of the CTN.

The server roles therefore must be reassigned after:

- ▶ The CTS in a single-server CTN has been power-on reset or has gone through a power-off/on sequence and has not been specified to be the only member of the CTN.
- ▶ The PTS (and the BTS if assigned) in a two-server CTN has been power-on reset or has gone through a power-off/on sequence and has not been specified to be the only members of the CTN.
- ▶ The PTS (and the BTS if assigned) in a three or more server CTN experience a power outage (for example, a data center power outage).
- ▶ An STP-only CTN is deconfigured intentionally by the user.

When the site power is restored or after an STP-only CTN is deconfigured, the CTN remains deconfigured, all participating servers being stratum 0.

The Network Configuration tab shows that the PTS, BTS, and Arbiter are not configured, as shown in Figure 3-42. Servers rely on explicit instructions from the console to re-establish a new time source and reconfigure the CTN.

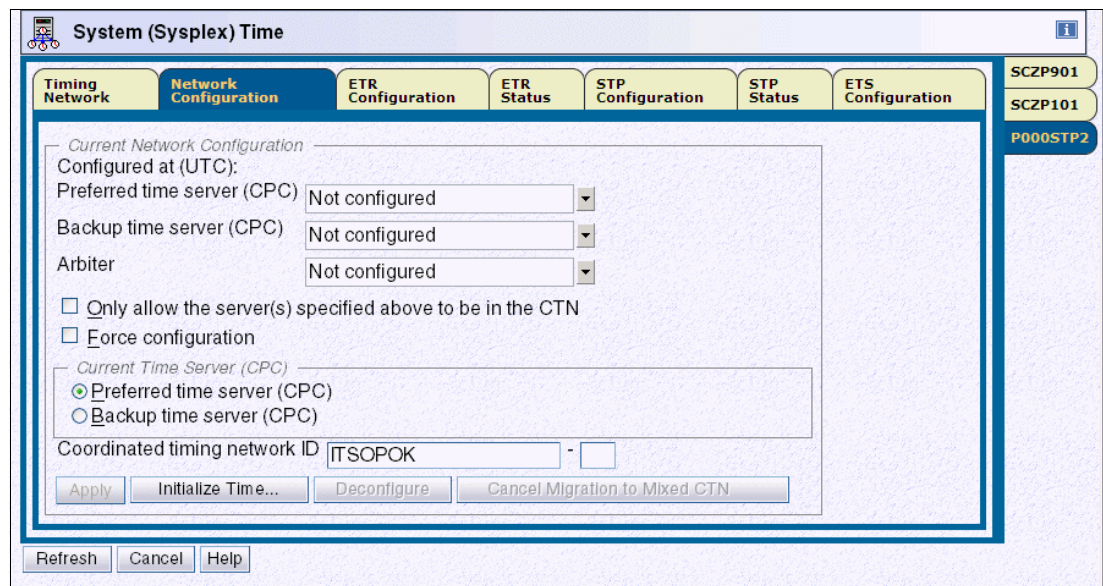


Figure 3-42 HMC workplace: System (Sysplex) Time, Network Configuration

This is a change from the restart procedures used in a similar situation when using a Sysplex Timer configuration. Operating procedures should be modified accordingly. The sequence to restart an STP-only CTN is:

1. Set the time zone offset, leap seconds, and date and time.
2. Assign the CTN roles.

Set the time zone offset, leap seconds, and date and time

The three tasks related to setting the time are:

- ▶ Set leap seconds.
- ▶ Set time zone.
- ▶ Set date and time.

This is accomplished by clicking **Initialize Time** on the Network Configuration tab, shown in Figure 3-42. The button becomes enabled when the CTN has been deconfigured as a result of a site power failure or POR of a single server CTN.

Following a power outage or POR of the CTS, the date and time can be set depending on the following:

- ▶ If the ETS is configured and is already available at the time of the recovery, the preferred method is to use the Use External Time Source option.
- ▶ For those customers who do *not* have an ETS, select **Modify time by delta to set date and time** and enter a delta value of zero.

Assign the CTN roles

Once the initialize time task has been completed, the Apply button on the Network Configuration tab becomes enabled and server roles can be assigned. Under normal circumstances, we recommend performing the configuration in two distinct steps. However, following a site power outage, it might be convenient to force the configuration in one step, because not all servers might be in POR-complete state.

The roles are assigned from the Network Configuration tab from the server that will become the Current Time Server.

1. Assign the PTS role, the BTS if there are two or more servers, and the Arbiter if the CTN has three or more participating servers.

For each of the roles (PTS, BTS, and Arbiter) there are drop-down boxes listing the servers currently available. Only STP-enabled servers with their Support Element visible to the HMC are selectable. In a normal situation, when the HMC has connectivity to the server, the STP ID information is shown if there is an initialized STP link to the target server and the server participates to the same CTN. In a recovery situation, a server's Support Element might be visible to the HMC, but the STP ID is not available. This means that either the activation of this particular server is not complete or coupling link connectivity has not yet been re-established.

2. Assign the CTS role.
3. Click the **Force configuration** check box. The option bypasses a number of validity checks on server connectivity, allowing the configuration of servers that might not be in POR-complete state or do not yet have coupling link connectivity to the selected CTS.
4. Click **Apply**. The assignment of the CTS globally configures into the CTN all servers that:
 - Are already powered-on and in POR-complete state
 - Have the same CTN ID as the CTS
 - Have coupling link connectivity to the CTS

Other servers that are not yet in POR-complete state, or that have not yet re-established connectivity to the CTS, automatically join the CTN and recognize their role assignments once their operating conditions return to normal.

After all conditions have returned to normal, we recommend that the server roles be reconfigured to the configuration that was in place before the outage.

3.5.8 Planned single-site maintenance

This section presents an alternate way to shut down/power on the CECs in a single site configuration when site power or building maintenance is scheduled.

The automatic activation of CECs is allowed using the proper sequence:

1. When initiating the maintenance:
 - a. Deactivate every CEC but the PTS/CTS.
 - b. Remove the BTS and Arbiter roles
 - c. On the PTS/CTS, remove the Arbiter role, then click the **Only allow the server(s) specified above to be in the CTN** box to keep the configuration at next power on.
 - d. Deactivate the PTS/CTS (which is now not synchronizing with any other CEC, so the deactivate is not blocked).
2. After the maintenance, take the following steps.

Note: After a power off/on sequence, the TOD value in the CTS is initialized from the TOD values stored in its SE. If you previously configured an ETS, STP does not perform an *initialize time* using the ETS, but automatically accesses the configured ETS and starts adjusting the time so as to maintain time accuracy. STP allows time adjustments of up to +/- 60 seconds.

- a. Activate the PTS/CTS first. It becomes stratum 1 (CTN configuration kept).
- b. From the PTS/CTS, deselect the **Only allow the server(s) specified above to be in the CTN** option.
- c. Activate the other CECs. They become stratum 2 or stratum 3, as appropriate, during their automatic activation.
- d. From the PTS/CTS, reassign the BTS and Arbiter roles.

This speeds up the reactivation of the CECs and simplifies the customer procedures for this scenario.

See also *Server Time Protocol Implementation Guide*, SG24-7281.

3.6 MES upgrade considerations

If you are upgrading a machine with a different footprint (often referred to as a push/pull MES) the machine type and possibly the serial number change. A server is deactivated, powered off, uncabled, and then moved aside for the new server to take its place. In other cases the new server is powered up alongside the one that it is replacing and is configured for a workload migration.

Configuring a new server might include IPLing a stand-alone image in order to load Crypto keys. As mentioned in 1.5.2, “Coupling link requirements in a non-sysplex configuration” on page 23, if this new server will be accessing any primary volumes that are part of an XRC relationship, your new server must observe the same time reference used by existing servers using those volumes. This can be accomplished by either joining the new server to the same CTN ID as the other servers or creating a new CTN but using the same ETS to initialize the time.

Depending on your current CTN configuration you might need to plan for certain actions prior to the MES. The issues discussed are categorized between:

- ▶ Mixed CTN
- ▶ STP-only CTN

3.6.1 Mixed CTN

If CEC1 is being replaced (Figure 3-43), since the recommendation is to always have two STP timing sources for any stratum 2 servers, CEC2’s ETR ports (if still cabled to enabled ports on the Sysplex Timers) should be enabled to make it a stratum 1 server. Ensure beforehand that CEC2 has initialized STP links to the other stratum 2 servers.




Figure 3-43 Replacing CEC1 in a Mixed CTN

If a new server will be temporarily joining the CTN you must plan for timing links for the new server. If either of the stratum 1 servers have unused CFP CHPIDs currently defined to and online to at least one CF LPAR, those links can be used as a timing source. If there are none you might have to dynamically define timing or timing-only links from one of the other servers.

3.6.2 STP-only CTN

If your STP-only CTN consists of one or two servers and you are only allowing these servers to be in the CTN per the option on the Network Configuration tab (see “Activating the STP-only CTN” on page 104, regarding this option), ensure that this option is *deselected* before powering off the server being replaced. Otherwise, the new server’s node ID will not be recognized and will not be allowed to join the existing CTN.

When replacing the PTS/CTS in a multi-server CTN, always move the PTS/CTS role to a server that is not being replaced. Similarly, if the BTS is being replaced, assign another server

in the CTN as a BTS so that recovery for a loss of the CTS is possible. Refer to section 3.4 of the *Server Time Protocol Implementation Guide*, SG24-7281, for more information.

Also in a multi-server CTN, using your CTN topology diagram, ensure that the server being replaced does not leave you with a last timing link condition upon deactivation. To illustrate, in Figure 3-44 we swap out CEC2 from a z9 to a z10.




Figure 3-44 Replacing CEC2 from a z9 to a z10

After completion of structure rebuilds from CFB to CFC and moving the workload off of z/OS B, there are no operational safeguards to prevent you from deactivating CEC2 if a last timing link condition exists. This would be the case if CEC3 was not part of the CTN or did not provide alternate STP timing connectivity to CEC4 (see Example 3-28).

Example 3-28 Messages revealing single points of failure for CEC4

```

CNZHF0005I One or more consoles are configured to receive messages
intended only for programmers.
15.18.09 SC81 STC03801 HZS0002E CHECK(IBMxcf,XCF_CF_PROCESSORS):
IXCH0444E Coupling facility processor configurations in use by the
local system may result in degraded response time and throughput
for coupling facility requests.
15.18.09 SC81 STC03801 HZS0001I CHECK(IBMcatalog,
CATALOG_IMBED_REPLICATE):
IGGHC104E The CATALOG_IMBED_REPLICATE check has detected one or
more catalogs defined with the IMBED and/or REPLICATE attributes.
15.18.09 SC81 STC03801 HZS0001I CHECK(IBMCSV,CSV_LNKLST_SPACE):
CSVH0980E Some LNKLST sets include data set(s) allocated with
secondary space defined.
*15.18.10 SC81 STC03801 *HZS0003E CHECK(IBMxcf,XCF_CDS_SPOF):
*IXCH0242E One or more couple data sets have a single point of failure.
15.18.22 SC81 STC03801 HZS0001I CHECK(IBMCSV,CSV_APF_EXISTS):
CSVH0957E Problem(s) were found with data sets in the APF list.

```

15.18.26 SC81 STC03801 HZS0001I CHECK(IBMUSS,USS_PARMLIB):
BPXH046E Syntax error(s) were found in the parmlib members.
*15.18.27 SC81 STC03801 *HZS0003E CHECK(IBMRA CF,RACF_SENSITIVE_RESOURCES):
*IRRH204E The RACF_SENSITIVE_RESOURCES check has found one or
*more potential errors in the security controls on this system.
15.42.59 SC80 ICH408I USER(HAIMO) GROUP(SYS1) NAME(HAIMO



Recovery concepts

Time synchronization supports two types of network configurations (Coordinated Timing Network (CTN)):

- ▶ Mixed CTN
- ▶ STP-only CTN

A Mixed CTN allows the use of both Sysplex Timer signals and STP messages. An STP-only CTN uses only STP messages to keep servers synchronized.

In this chapter we present a high-level overview of STP recovery concepts and definitions, and how they apply to each type of CTN. Later in this chapter we describe how recovery is achieved for different types of failure.

For a complete description of the concepts refer to *Server Time Protocol Recovery Guide*, SG24-7380.

4.1 Terminology overview

This section provides a review of STP terminology:

- ▶ Sysplex Timer offline sequence
 - Sysplex Timer detects a failure. The failing Sysplex Timer transmits an offline sequence symbol on the Control Link Oscillator (CLO) links to signal the other Sysplex Timer in an expanded availability configuration that it is going offline.
 - If Sysplex Timers in an expanded availability configuration lose the capability to synchronize:
 - The primary Sysplex Timer continues to transmit ETR signals whether or not an offline sequence is received.
 - The secondary Sysplex Timer becomes the primary Sysplex Timer if it receives an offline sequence.
 - The secondary Sysplex Timer discontinues transmission of ETR signals if it does not receive an offline sequence.
- ▶ Synchronization check threshold
 - The server or coupling facility (CF) is considered to be in synchronized state if its time of day (TOD) clock is within the synchronization check threshold of the Coordinated Server Time (CST).
 - The STP synchronization check threshold is 50 microseconds. If a server's TOD clock differs from the Coordinated Server Time by more than +/- 50 microseconds, the server or CF can become unsynchronized and therefore becomes a *stratum 0*.
 - A synchronization check can be recoverable when it happens and STP can re-establish synchronization.

4.2 Freewheel interval

The freewheel interval is the amount of time that a stratum 2 or stratum 3 server can remain synchronized without receiving messages from its clock source. It is approximately 1 second for a Mixed CTN or 10 seconds for an STP-only CTN.

Freewheel only occurs if a server has connectivity to at least one potential clock source, stratum 1 or 2, but valid messages are not being received from the clock source. However, if a server loses connectivity to all potential clock sources (link failures, for example), it immediately initiates recovery.

For stratum 2 and stratum 3 servers, the freewheel interval is entered when STP timing messages from the selected time source are not received. If no alternative time source can be found, the server will become unsynchronized at the end of the freewheel interval and switch to stratum 0. If an alternative time source is available, a switch to a different stratum level might be required in order to receive STP timing messages.

4.3 Server offline signal

An offline signal (OLS) is transmitted by the server to indicate that the channel is going offline. This is independent of STP. Conditions when an OLS is issued by the server include:

- ▶ Server or LPAR dump
- ▶ Server power off
- ▶ CHPID configure off
- ▶ Channel failure (channel checkstop)

An offline signal might not be transmitted for certain failures, such as:

- ▶ Channel subsystem failure
- ▶ System Assist Processor (SAP) recovery
- ▶ Site or server power outage
- ▶ Link failure

In an STP-only CTN, when a BTS is configured but an Arbiter is not, STP recovery checks whether offline signals have been transmitted on the initialized coupling links between the PTS and the BTS. Furthermore, with STP Version 2, OLSs are considered only if multiple links that went down in the last two seconds have received the offline signal indication:

- ▶ If the BTS receives OLS on multiple links, including the last link to the CTS within a two-second interval, the BTS will take over as CTS.
- ▶ If the CTS has sent OLS on multiple links, including the last link to the BTS within a two-second interval, the CTS will release its CTS role.

If, when the last link goes down, all other links went down more than two seconds ago, then the link failure is considered a single link case and OLS rules cited previously do not apply. Also, STP recovery does not use OLS in a CTN where an Arbiter is configured.

When a BTS is configured but an Arbiter is not, STP recovery uses offline signals in conjunction with console-assisted recovery to make the final determination as to whether the BTS initiates a take over as the CTS. See 4.6, “Console-assisted recovery” on page 158.

4.4 Going Away Signal

The *Going Away Signal (GAS)* is a reliable unambiguous signal to indicate that the CPC is about to enter a check stopped state. When a GAS from the CTS is received by the BTS, it safely takes over as CTS.

GAS has priority over OLS in a CTN where an Arbiter has not been assigned. The BTS can also use GAS to take over as CTS for CTNs with an Arbiter assigned without communicating with the Arbiter. This is in contrast to OLS, where OLS is ignored for CTNs with an Arbiter assigned.

GAS removes the dependency on OLS and CAR in a CTN without an Arbiter assigned and the dependency on BTS to Arbiter communication for CTNs with an Arbiter assigned.

GAS is sent on InfiniBand (IFB) links using HCA3-O to HCA3-O - 12x IFB or 12x IFB3 or HCA3-O LR to HCA3-O LR - 1x IFB for z196 GA2 and later machines.

The current recovery design is still used when GAS is not received by BTS and for other failure types.

4.5 Arbiter-assisted recovery

Arbiter-assisted recovery is applicable when both a BTS and an Arbiter are assigned. The BTS does not invoke OLS rules because the Arbiter provides additional means to determine whether the BTS can take over.

If the BTS loses communication on all of its established paths to the CTS, it attempts to determine the status of the CTS through the Arbiter:

- ▶ If both the BTS and the Arbiter cannot communicate with the CTS, then the BTS takes over as the CTS and becomes the Stratum 1.
- ▶ If the CTS is still alive after losing communication with both the Arbiter and the BTS, it will switch to Stratum 0 so that the CTN does not end up with two Stratum 1 servers.
- ▶ If the Arbiter can communicate with the CTS, then the BTS will not take over, but instead transition to Stratum 3 and get its timing signals from a Stratum 2 server, for example the Arbiter. The Arbiter takes note that the BTS no longer has connectivity to the CTS and, should it subsequently lose contact with the CTS, the Arbiter will inform the BTS accordingly, causing the BTS to proceed with taking over as the CTS.
- ▶ If the BTS is unable to communicate with the Arbiter due to connectivity failure, console-assisted recovery is invoked by the BTS as an alternate method for determining the status of the CTS.

Note: In a two-site CTN, the location of the Arbiter is critical: it will determine which site remains operational after a loss of communication between the PTS and BTS. Arbiter location considerations are discussed in Server Time Protocol Recovery Guide, SG24-7380, Section 2.4.8, “Two-site considerations”.

Blocking disruptive actions on STP role servers

Disruptive actions such as POR are blocked for the CTS. New function has now been added to block disruptive actions on any of the STP role servers, PTS, BTS, or Arbiter. This prevents a disruptive action causing the CTS to give up the S1 role and go S0. For example, if the PTS has the CTS role and then the Arbiter (or BTS) has a planned or unplanned outage, a disruptive action on the BTS (or Arbiter) causes the CTS to give up the S1 role and go S0 as it loses communication to both the BTS and Arbiter.

This new function was introduced with the MCL levels shown in Table 4-1.

Table 4-1 Blocking disruptive actions

Driver/Server	MCL	Bundle	Release Date
D86E / z196	N29809.277	45	Sept 8, 2011
D86E / z196	N29802.420	45	Sept 8, 2011
D79F / z10	N24415.078	50	Sept 28, 2011
D79F / z10	N24409.184	50	Sept 28, 2011
D93G / z114 and z196 GA2	Integrated	N/A	Sept 9, 2011

Arbiter-assisted recovery enhancements

The current STP Arbiter-assisted recovery design handles recovery of single failures in an STP-only CTN. Enhancements have now been made to handle planned and unplanned actions that could affect two of the three STP role servers.

This includes safeguarding against the following potential hazards, which could result in a CTN-wide failure:

- ▶ Planned disruptive actions on the BTS and Arbiter in parallel as part of the same task. Note that disruptive actions on any of the STP role servers will be blocked via the enhancement described previously.
- ▶ Unplanned failure of a second STP role server when the STP role for the first unplanned failure is not reassigned or removed.
- ▶ Failure to remove the STP role of a server being upgraded to a new machine type. If the STP role is not removed prior to the upgrade, the new node descriptor prevents the server from reassuming the STP role. This puts the CTN in a state equivalent to that when the server with that STP role has an unplanned outage.

Arbiter-assisted recovery is enhanced so that a degraded state is entered when any two of the three STP role servers (PTS, BTS or Arbiter) agree that they cannot communicate with the third STP role server. A degraded state is entered when:

- ▶ The BTS and Arbiter can communicate but cannot communicate with the PTS/CTS. The BTS will take over as the CTS and then Arbiter-assisted recovery is disabled.
- ▶ The PTS and BTS can communicate but cannot communicate with the Arbiter.
- ▶ The PTS and Arbiter can communicate but cannot communicate with the BTS.

This new function was introduced with the MCL levels shown in Table 4-2.

Table 4-2 Arbiter-assisted recovery enhancements

Driver/Server	MCL	Bundle	Release date
D79F / z10	N24406.094	50	Sept 28, 2011
D86E / z196	N29799.110	44	Aug 24, 2011
D93G / z114 and z196 GA2	Integrated	N/A	Sept 9, 2011
Note: The PTS, BTS, and Arbiter servers must all be at the required MCL level for Arbiter-assisted recovery to be disabled.			

Disabling Arbiter-assisted recovery provides safeguards against:

- ▶ Planned disruptive actions done sequentially on the BTS or PTS (whichever is not the CTS) and the Arbiter
- ▶ Double failures, unplanned or a combination of planned and unplanned, of the BTS or PTS (whichever is not the CTS) and the Arbiter

It does not provide safeguards against planned disruptive actions initiated as part of the same HMC task on both the PTS or BTS (whichever is not the CTS) and the Arbiter.

While Arbiter-assisted recovery is disabled:

- ▶ The BTS cannot take over as CTS using Arbiter-assisted recovery.
- ▶ The CTS will not surrender its role when it loses attachment to the remaining special role server.
- ▶ The BTS can still take over as CTS using either Console-assisted recovery (CAR) or the STP Going Away Signal (GAS) transmitted from the CTS.

Note: After Arbiter-assisted recovery has been disabled it will not be reenabled until there is full connectivity between the PTS, BTS, and Arbiter.

4.6 Console-assisted recovery

Console-assisted recovery uses the HMC in an attempt to determine the status of the PTS (when initiated by the BTS) or the status of the BTS (when initiated by the PTS). Console-assisted recovery helps to determine whether the BTS can take over as CTS, or the PTS can take back its role as a CTS.

4.6.1 Console-assisted recovery in a CTN with BTS

In a CTN that does not have an Arbiter configured, console-assisted recovery is used to determine whether the BTS should take over the CTS role. The BTS initiates console-assisted recovery when the BTS has lost communication with the CTS.

If the failure has already been detected through OLS, the BTS has taken over the CTS role, and console-assisted recovery is used to confirm that the PTS has failed.

If the CTS failure has not yet been recognized through OLS, for example, because the failure involved a single link, the BTS takes over if console-assisted recovery confirms that the CTS has failed.

When it initiates console-assisted recovery, the BTS sends a command to its Support Element (SE) to determine the state of the CTS by communicating through the HMC.

- ▶ If the response indicates that the CTS has failed, the BTS can take over as the new CTS.
- ▶ If the response indicates that the status of the CTS is either good or indeterminate, the BTS cannot take over as the new CTS and becomes stratum 0.

Analysis of the error by offline signals or console-assisted recovery is made at the time that each process is invoked. In most cases, OLS and console-assisted recovery are processed almost simultaneously and only one recovery situation is visible to the user. However, if the error conditions change between the time that the OLS check is made and console-assisted recovery is run, the final STP recovery decision is based on the analysis of conditions at the time that console-assisted recovery is run.

4.6.2 Console-assisted recovery in a CTN with BTS and Arbiter

In a CTN that has an Arbiter configured, offline signal indications are not used. Console-assisted recovery can be initiated by the BTS or the PTS.

- ▶ Console-assisted recovery is initiated by the BTS when it has lost communication with *both* the CTS and the Arbiter. In this case, because the Arbiter is configured, though not available, the offline signals are ignored. Only console-assisted recovery is used.

The BTS sends a command to its SE to determine the state of the CTS by communicating with the CTS through the HMC.

- If the response indicates that the CTS has failed, the BTS can take over the CTS role.
- If the response indicates that the status of the CTS is either good or indeterminate, the BTS cannot take over as the new CTS.
- ▶ Console-assisted recovery is initiated by the PTS when it has lost communication with *both* the BTS and Arbiter.

The STP-only CTN must only have one CTS, so the status of the BTS must be determined.

The PTS first surrenders its role of CTS as soon as connectivity to the BTS and Arbiter is lost. Then it uses console-assisted recovery in an attempt to determine whether it should take back the CTS role. The PTS sends a command to its SE to determine the state of the BTS by communicating with the BTS through the HMC.

- If the response indicates that the BTS has failed, the PTS re-takes over its CTS role.
- If the response indicates that the BTS has taken over the CTS role, or is inconclusive, the PTS does not take back the CTS role.

4.7 Island condition

An island condition occurs when a server detects that one or more servers might be operating as a separate timing network with the same CTN ID, but have a different definition of which servers are performing the PTS, CTS, and Arbiter roles.

4.8 Switch to local timing mode

When a server becomes unsynchronized and transitions to stratum 0, the resident z/OS system images running in STP timing mode switch to local timing mode.

For z/OS systems running on the server, the impact of switching to local timing mode depends on the PLEXCFG parameter in IEASYSxx, and the ETRMODE or STPMODE specified in CLOCKxx.

- ▶ If a system running in either ETR or STP timing mode loses its time source, then:
 - If `plexcfg=xcflocal` or `monoplex`, then the system will continue in LOCAL timing mode. In ETR mode, message IEA261I NO ETR PORTS ARE USABLE. CPC CONTINUES TO RUN IN LOCAL MODE. is issued while in STP mode, message IEA381I THE STP FACILITY IS NOT USABLE. SYSTEM CONTINUES IN LOCAL MODE. is issued.
 - If `plexcfg=multisystem` or `any`, the IEA015A is issued in ETR timing mode or IEA394A is issued in STP timing mode. In a Mixed CTN, the S1 z/OS servers issue IEA015A while the S2 z/OS servers issue IEA394A.
- ▶ z/OS systems that specify `PLEXCFG=MULTISYSTEM` or `PLEXCFG=ANY` in IEASYSxx, and `ETRMODE YES` or `STPMODE YES` in CLOCKxx, issue WTOR message IEA015A or IEA394A to allow operator intervention to resolve the problem before a wait state is loaded.
- ▶ z/OS system images that are using a Sysplex timer as their timing source issue WTOR message IEA015A (Example 4-1). The Sysplex timer connectivity needs to be reestablished before a reply with RETRY will be accepted.

Example 4-1 WTOR message IEA015A

IEA015A THIS SYSTEM HAS LOST ALL CONNECTION TO THE SYSPLEX TIMER.

IF THIS EVENT OCCURRED ON SOME, BUT NOT ALL SYSPLEX MEMBERS THE LIKELY CAUSE IS A LINK FAILURE. TO FIX, ENSURE THAT EACH AFFECTED SYSTEM HAS AT LEAST ONE CORRECTLY CONNECTED AND FUNCTIONAL LINK.

IF THIS EVENT OCCURRED ON ALL SYSPLEX MEMBERS, THEN THE LIKELY CAUSE IS A SYSPLEX TIMER FAILURE. TO FIX, REFER TO THE MESSAGE IEA015A DESCRIPTION IN MVS SYSTEM MESSAGES.

AFTER FIXING THE PROBLEM, REPLY "RETRY" FROM THE SERVICE CONSOLE

(HMC). IF THE PROBLEM WAS NOT CORRECTED, THIS MESSAGE WILL BE REISSUED AND YOU MAY TRY AGAIN. REPLY "ABORT" TO EXIT MESSAGE LOOP. PROBABLE RESULT: 0A2-114 WAITSTATE.

- ▶ z/OS system images that are using a Stratum 1 or Stratum 2 server as timing source issue WTOR message IEA394A (Example 4-2). Once the CEC is Stratum 1, Stratum 2 or Stratum 3 again, a reply of RETRY will be accepted.

Example 4-2 WTOR message IEA394A

IEA394A THIS SERVER HAS LOST CONNECTION TO ITS SOURCE OF TIME.

IF THIS EVENT OCCURRED ON SOME, BUT NOT ALL NETWORK SERVERS THE LIKELY CAUSE IS A LINK FAILURE. TO FIX, ENSURE THAT EACH AFFECTED SERVER HAS AT LEAST ONE CORRECTLY CONNECTED AND FUNCTIONAL LINK.

IF THIS EVENT OCCURRED ON ALL NETWORK SERVERS, THEN THE LIKELY CAUSE IS A TIMING NETWORK FAILURE. TO FIX, REFER TO THE MESSAGE IEA394A DESCRIPTION IN MVS SYSTEM MESSAGES.

AFTER FIXING THE PROBLEM, REPLY "RETRY" FROM THE SERVICE CONSOLE (HMC). IF THE PROBLEM IS NOT CORRECTED, THIS MESSAGE WILL BE REISSUED AND YOU MAY TRY AGAIN. REPLY "ABORT" TO EXIT MESSAGE LOOP. PROBABLE RESULT: 0A2-158 WAITSTATE.

4.9 External Time Source (ETS)

In an STP-only CTN, the ETS function is available using three different options:

- ▶ Using dial-out on the hardware management console (HMC)
- ▶ Using Network Time Protocol (NTP) client support on the Support Element
- ▶ Using NTP client support on the Support Element along with a pulse per second input on the ETR cards

There are no specific recovery actions when the ETS is configured to use a dial-out service. Depending of the ETS configuration, there are two recovery concepts available when NTP with or without PPS is being used:

- ▶ NTP server availability: Two NTP servers configured for one System z server
- ▶ Continuous NTP server availability: NTP server configured for both the PTS and the BTS

Note: If STP loses connectivity to all its NTP servers (with or without PPS), all servers in the CTN remain time synchronized. The CST might drift away from the NTP time source until NTP server communication is re-established.

The following list explains terms that are used within the NTP recovery sections:

- ▶ Network Time Protocol server
This provides the capability to keep the CST synchronized to the ETS to within 100 milliseconds.
- ▶ NTP server with pulse per second (PPS)

This provides the capability to keep the CST synchronized to the ETS to within 10 microseconds. Therefore, a highly stable and accurate PPS output, provided by an NTP server, is utilized.

- ▶ Selected and non-selected NTP server

The user is responsible for selecting the preferred NTP server. The preferred NTP server is called the *selected* NTP server. If two NTP servers have been configured, the second one is called the *non-selected* NTP server.

For detailed planning information about how CTN reacts to ETS changes, see *Server Time Protocol Recovery Guide*, SG24-7380.



Migration planning

This chapter provides guidance for configuration planning. The topics covered are:

- ▶ “Preparation and installation” on page 164
- ▶ “Migration between an ETR network and a Mixed CTN” on page 172
 - “Migration from an ETR network to a Mixed CTN” on page 172
 - “Migration from a Mixed CTN to an ETR network” on page 176
- ▶ “Migration between a Mixed CTN and an STP-only CTN” on page 177
 - “Migration from a Mixed CTN to an STP-only CTN” on page 178
 - “Migration from an STP-only CTN to a Mixed CTN” on page 182
- ▶ “Migration from an ETR network to an STP-only CTN” on page 184
- ▶ “Installation of a new STP-only CTN” on page 189
- ▶ “Installation of a single-server STP-only CTN” on page 191
- ▶ “Enhancing a single-site ETR network to a Mixed CTN across two sites” on page 194

Preparation and installation

In most cases, Server Time Protocol (STP) configuration planning involves the migration from an External Time References (ETR) network to an STP-only Coordinated Timing Network (CTN). Although it is recognized that there can be an intermediate phase during which servers in both ETR timing mode and STP timing mode will coexist in a Mixed CTN, the final stage is expected to be an STP-only CTN. A new configuration that is not migrating from an ETR network can be directly configured as an STP-only CTN.

Before starting a migration, you must complete the installation tasks that are associated with hardware and software components. It is important to complete all hardware and software tasks for the intended Mixed CTN prior to enabling STP.

Hardware installation

To make the servers STP-capable, hardware maintenance change levels (MCLs) must be installed on the servers that will be configured in a Mixed or STP-only CTN. In essence, this means that you install all MCLs that are available for all STP-capable servers and coupling facilities (CFs):

- ▶ The z990 and z890 must be at EC Driver level 55K or later, in addition to the latest MCLs that are available for the driver. STP on z990 or z890 does not support using a Network Time Protocol (NTP) server as an External Time Source (ETS).
- ▶ The z9 EC and z9 BC must be at EC Driver level 67L or later, in addition to the latest MCLs. The NTP server is supported as an ETS.
- ▶ The z10 EC and z10 BC must be at Driver level 76D or later, in addition to the latest MCLs. The NTP server is supported as an ETS.

Table A-1 shows the correspondence between the driver level and the Support Element (SE) console application level for servers that can install the STP feature.

Table A-1 EC Driver and SE console application levels

Server	EC driver level	Support Element console application	includes NTP client support	HMC Application
z10 EC, z10 BC	79F	2.10.2	Yes	2.10.2
z10 EC, z10 BC	76D	V2.10.1	Yes	V2.10.1
z9 EC, z9 BC	67L	V2.9.2	Yes	V2.9.2

The Hardware Management Console (HMC) Application level must be equal to or higher than the level of the Support Element console that it manages.

STP functions cannot be used until the STP feature (Feature Code 1021) is installed. The feature should be installed on any server or coupling facility that is to be used in a Mixed or STP-only CTN. The STP feature is installed via the server Support Element without disruption to current operations. After the MCLs are installed, all eligible servers are STP-capable servers.

After installation of the STP feature, the server can be configured for STP. Note the following:

- ▶ The System (Sysplex) Time task is available on the SE and HMC to configure a server to use STP.
- ▶ Coupling facility partitions automatically recognize the new functions. However, z/OS system images must be IPLed.

Sysplex Timer considerations

The Sysplex Timer's Licensed Internal Code (LIC) has been upgraded to support using STP in a Mixed Coordinated Timing Network (Mixed CTN). The required Sysplex Timer LIC is shipped along with the STP feature, and must be installed by the IBM Service Support Representative prior to migrating from a Sysplex Timer-based ETR network to any Coordinated Timing Network.

The Sysplex Timer provides the timekeeping information in a Mixed CTN. The Sysplex Timer consoles continue to be used for all timing-related functions, such as initializing time, scheduling daylight saving time, adjusting time, and so forth. Plan to keep access to the Sysplex Timer consoles for as long as the Sysplex Timer remains operational in the timing configuration.

Continuous availability of timers

The two timers should be physically separated. The better the separation, the lower the probability that an environmental failure will affect both timers. Note the following considerations:

- ▶ If an external time source is used, an ETS should be attached to each timer for continuous ETS availability.
- ▶ An active and a stand-by console should be configured.
- ▶ Each timer unit should be connected to a separate AC power source.
- ▶ When extenders are used, connect at least one pair of the extenders used for the CLO links to a power source that is separate from the power source used by the primary Sysplex Timer. Doing so greatly improves the probability that the OLS from the primary Sysplex Timer will be received by the secondary Sysplex Timer.
- ▶ If available, protect the power connections with an uninterruptible power supply (UPS) for each Sysplex Timer and extender.
- ▶ Route each CLO link between timers and each fiber link between timer and CPCs across separate paths.

Availability of an STP-capable HMC

Either upgrade an existing HMC or install a new HMC capable of running the HMC application V2.9.2 or later. The HMC Application level must be equal to or later than the level of that Support Element console that it manages (Table A-1 on page 164).

All servers in the CTN must be defined to the HMC. We strongly recommend this even if the servers are in a multi-site configuration. If at least one HMC does not have connectivity to all servers, it is not possible to freely configure all roles in the CTN.

The HMC workstation is a closed environment dedicated to the HMC Application. No other application can be installed.

z/OS preparation and installation

Even though z/OS supports STP, additional maintenance may be required. STP-related maintenance may be required for all z/OS systems that require time synchronization with a z/OS image or CF on a server configured in a Mixed or STP-only CTN. In a sysplex configuration, if any coupling facility or z/OS image is running on an STP-configured server, STP maintenance may be required on *all* z/OS images in the sysplex.

A re-IPL is required for the z/OS maintenance to take effect. Before performing the IPL, the CLOCKxx member in the PARMLIB might need to be edited to contain the required STP statements (see 2.5.2, “CLOCKxx” on page 59).

Important: If the z/OS maintenance is applied before installing the STP feature, z/OS must be re-IPLed after the STP feature is installed.

z/OS images that do not require multisystem time synchronization can be at any supported release regardless of the timing mode of the server.

PSP buckets

Consult the PSP Buckets for STP-related maintenance. Look in the associated hardware Preventive Service Planning subset. The upgrade is xxxxDevice and the subset is xxxx/ZOS where xxxx is one of the machine types that supports STP (that is, 2094, 2096, 2084, or 2086).

To simplify the identification of PTFs needed to support Server Timer Protocol, a functional Preventive Service Planning (PSP) bucket has been created. Use the Technical Help Database for Mainframe Preventive Service Planning Buckets and the Enhanced PSP Tool to reconcile the recommended service for STP and the hardware device against your systems' target zones. The database is available at the following Web site:

<http://www14.software.ibm.com/webapp/set2/psp/srchBroker>

This Web site provides a search capability, as well as pull-down lists for types of PSP buckets. Scroll to the “Find the bucket by Type, Category, and Release” heading, then select **Function** for the Type field and **STP** for the Category field.

Server roles in an STP-only CTN

It is important to decide what role each server in the CTN will play, as discussed here:

- ▶ One server will be the Preferred Time Server *and* should be assigned as the Current Time Server (CTS, stratum 1).

Note: Only the PTS can automatically re-take over the CTS role after a recovery action when BTS has taken over as the CTS.

- ▶ Assign a second server, if available, the role of the Backup Time Server (BTS).
- ▶ Assign a third server, if available, the role of Arbiter. The Arbiter provides additional means for the BTS to determine whether it should take over as the Current Time Server in the event of planned or unplanned outages.

Do not make the choice of Preferred and Current Time Server lightly. This is the preferred server for good reason. Keep in mind the following considerations when assigning this role:

► The location of the server

In a multi-site configuration, is the Preferred Time Server located in the most important site (if relevant)? Is it a server running critical production? Is the server part of a disaster recovery setup?

Note: If disks are being mirrored, locate the PTS in the same site as the primary disk subsystems.

► The technology level of the server

A new STP function will be introduced to the latest technology, not necessarily to older server generations. For example, if NTP client support is planned, a z990 or z890 cannot be chosen as PTS or BTS.

► Maintenance considerations

A server that requires frequent scheduled maintenance is not the best choice. Disruptive actions are not permitted on a server that is assigned the role of Current Time Server. The CTS role must be moved in advance.

► Connectivity considerations

- With three or more servers, the PTS, BTS, and Arbiter must connect to each other with coupling links forming a triangle.
- The BTS connectivity to other servers should be such that all servers remain synchronized when the PTS fails and the BTS takes over the CTS role.
- All stratum 3 servers should connect to at least two stratum 2 servers. This allows the stratum 3 server to continue in the synchronized state in the event that one of the stratum 2 servers fails or is unavailable due to maintenance. Since STP does not support four stratum levels, a stratum 4 server would be placed in the unsynchronized timing state (stratum 0).
- Also pay attention not to exceed the maximum number of CF links that can be connected to a single CEC when planning for additional connectivity.
- A standalone coupling facility typically has coupling link connectivity to other servers requiring time coordination. This existing connectivity can be exploited for STP message exchanges without having to deploy additional coupling links. However, you must also pay attention to the connectivity rules listed above and the best practices recommendations for the reassignment of servers listed in , “Plan for high availability” on page 168.
- A standalone coupling facility does not produce z/OS messages for operations or interception by automation routines. Although CTN information is available from z/OS images on other servers, not all information specific to the server role assigned to the coupling facility is available. This includes:
 - Information messages at IPL or interrupt time, or Display ETR output that specifically identifies a server role. However, CTN information is available from other z/OS images in the CTN.

IEA380I THIS SYSTEM IS NOW OPERATING IN STP TIMING MODE.

IEA381I THE STP FACILITY IS NOT USABLE. SYSTEM CONTINUES IN LOCAL TIMING MODE

IEA386I TIMING STATUS SYNCHRONIZATION MODE

NUMBER OF TIMING LINKS=nn

THIS IS THE PREFERRED / BACKUP / ARBITER SERVER

THIS SERVER HAS NO LINK TO THE PREFERRED TIME SERVER

THIS SERVER HAS NO LINK TO THE ARBITER SERVER
THIS SERVER HAS ONLY A SINGLE SOURCE OF TIMING SIGNALS
IEA390I TOD CLOCKS DYNAMICALLY ADJUSTED TO MAINTAIN STP SYNCHRONISM.

- Warning messages identifying a condition being associated to a given server. If the condition being raised relates to connectivity between two servers, the information might be available to a z/OS system image at the other end of the link. However, if both ends of the link are CF partitions, no warning message is available to the user.

IEA382I THIS SERVER HAS ONLY A SINGLE LINK AVAILABLE FOR TIMING PURPOSES

IEA383I THIS SERVER RECEIVES TIMING SIGNALS FROM ONLY ONE OTHER NETWORK NODE

IEA388I THIS SERVER HAS NO CONNECTION TO THE BACKUP / ARBITER

IEA394A THIS SERVER HAS LOST CONNECTION TO ITS SOURCE OF TIME. IF THIS EVENT OCCURRED ON SOME, BUT NOT ALL NETWORK SERVERS, THE LIKELY CAUSE IS A LINK FAILURE. TO FIX, ENSURE THAT EACH AFFECTED SERVER HAS AT LEAST ONE CORRECTLY CONNECTED AND FUNCTIONAL LINK. IF THIS EVENT OCCURRED ON ALL NETWORK SERVERS, THEN THE LIKELY CAUSE IS A TIMING NETWORK FAILURE. TO FIX, REFER TO THE MESSAGE IEA394A DESCRIPTION IN MVS SYSTEM MESSAGES. AFTER FIXING THE PROBLEM, REPLY "RETRY" FROM THE SERVICE CONSOLE (HMC). IF THE PROBLEM WAS NOT CORRECTED, THIS MESSAGE WILL BE REISSUED AND YOU MAY TRY AGAIN. REPLY "ABORT" TO EXIT THE MESSAGE LOOP. PROBABLE RESULT: 0A2-158 WAITSTATE.

► Recovery considerations

Keep in mind that STP recovery and sysplex recovery are two different things. When the CTS fails, not only STP must recover, but also z/OS-images or coupling facilities, or even both, are unavailable, if they were located on the failing CEC. Therefore, even if all remaining servers stay in synch due to STP recovery, recovery actions for the sysplex might still be required.

Once these points have been considered for the PTS/CTS, a similar review should be made for the Backup Time Server. In a two-site configuration the PTS and BTS should be in different sites.

The Arbiter must be assigned to a server that has coupling connectivity to both the PTS and the BTS. Also, for a two-site-configuration, the Arbiter must be in the same site as the PTS, unless specific reasons suggest otherwise (for example, GDPS with Freeze policy).

Note: No single recommendation for server role is appropriate for all possible situations. Each customer must evaluate which factors are important in his operational environment.

Plan for high availability

Regarding high availability:

- In a Mixed CTN there must be at least one STP-configured server connected to a pair of Sysplex Timers. The server steps to signals received from one of the Sysplex Timers.
- In order to avoid single points of failure, there should be at least two stratum 1 servers in a Mixed CTN (stepping to timing signals received from the Sysplex Timers) before configuring any additional servers that use STP messages for synchronization.
- A Mixed or STP-only CTN should be configured so that every stratum 3 server is attached to at least two stratum 2 servers. This allows the stratum 3 server to continue in the synchronized state in the event that one of the stratum 2 servers fails or is unavailable due to maintenance. Since STP does not support four stratum levels, a stratum 4 server would be placed in the unsynchronized timing state.
- Redundant coupling links should be configured between any two servers in the CTN. Whenever possible, CHPIDs should be defined in the access list for at least two partitions.

Also consider the possibility for timing-only links to be reassigned as coupling links, or vice versa. Timing-only links must be configured when a CF does not exist at either end of the coupling link. Re-assigning timing-only links as coupling links and vice versa can be done without disruption, as described in 3.5.1, “Migrating timing-only links to coupling links” on page 134.

- ▶ For installations with more than three servers, plan for which servers CFs will be reassigned the STP roles in case one of the servers will be unavailable because of an extended planned or unplanned outage. With the reassignment the servers should be in a fault-tolerant configuration again. This is especially critical for the PTS and BTS-roles. If no reassignment is done, a subsequent failure of the CTS would lead to a CTN-wide sysplex outage, because there is no other server to take over the CTS role. Thus, no STP-recovery can be done.
- ▶ If the Arbiter is unavailable because of an outage, the rules for Arbiter-assisted recovery still apply and the OLS is ignored. In this case it might be reasonable to either move the Arbiter role to a different server or (if that is not possible) to temporarily remove the Arbiter role from the CTN. For a two-site-configuration, also plan ahead which servers to reassign the roles in case a site failure occurs.

Migration paths

After the preparation and installation steps are completed, the migration to a Coordinated Timing Network (CTN) can begin.

Definition and configuration actions are performed from the Hardware Management Console using the System (Sysplex) Time task. When properly planned and executed, all operations are concurrent to server operations and nondisruptive to the z/OS images. Certain operations affect an individual server. Other operations affect all servers defined to the CTN. Figure A-1 summarizes the possible paths.

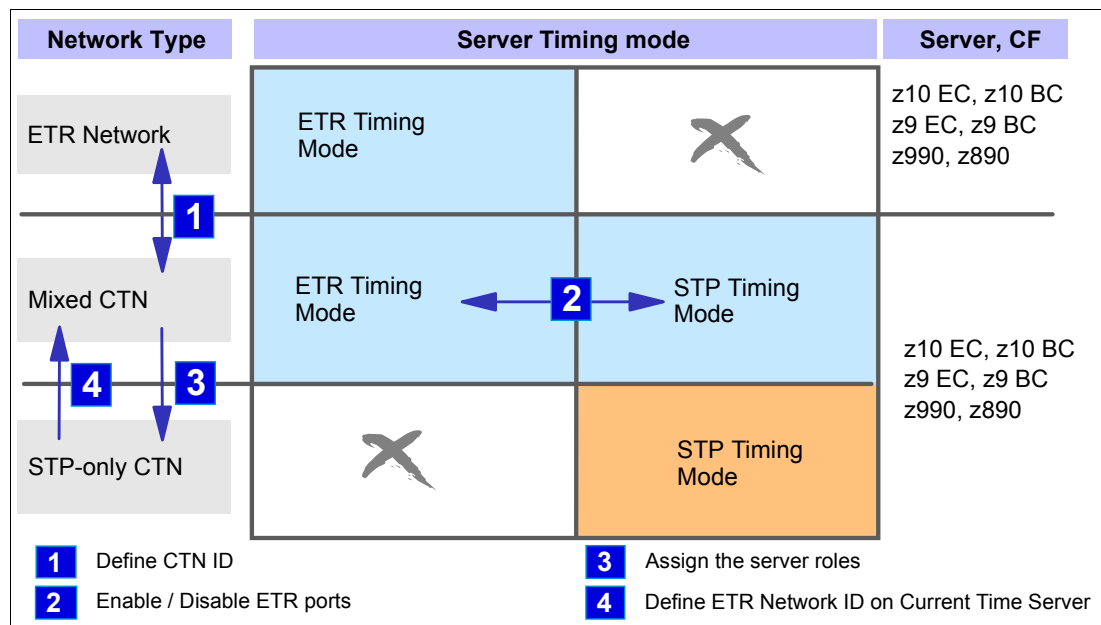


Figure A-1 Overview of possible migration paths

The installation of, or migration to, a Coordinated Timing Network is a combination of one or more of the four steps illustrated, as described here:

- 1.** Migrate from an ETR network to a Mixed CTN, or from a Mixed CTN to an ETR network.

This is a local action, one server or CF at a time. Definition of the CTN ID in the first server *creates* the Coordinated Timing Network, and removal of the last server *deletes* the Coordinated Timing Network.

- 2.** Migrate a specific server or CF from ETR to STP timing mode in a Mixed CTN, or vice versa, migrate from STP timing mode to ETR timing mode.

This is a local action, one server at a time.

- 3.** Migrate a Mixed CTN to an STP-only.

The scope is global, and the action affects all servers and CFs defined in the Mixed CTN.

- 4.** Migrate an STP-only to a Mixed CTN.

The scope of the migration is global for all servers and CFs in the STP-only CTN. The Preferred and Backup Time Servers both migrate from STP timing mode to ETR timing mode. For other servers in the CTN, timing mode changes can be requested individually.

The following sections present a selected set of examples to illustrate possible migration steps. Example A-1 lists the scenarios documented. Each scenario focuses on one part of the migration. Review all configuration rules and recommendations documented in earlier chapters.

Example: A-1 Migration examples

Section	Description	Page
, "Migration between an ETR network and a Mixed CTN" on page 172	, "Migration between an ETR network and a Mixed CTN" on page 172	, "Migration between an ETR network and a Mixed CTN" on page 172
, "Migration from an ETR network to a Mixed CTN" on page 172	, "Migration from an ETR network to a Mixed CTN" on page 172	, "Migration from an ETR network to a Mixed CTN" on page 172
, "Migration from a Mixed CTN to an ETR network" on page 176	, "Migration from a Mixed CTN to an ETR network" on page 176	, "Migration from a Mixed CTN to an ETR network" on page 176
, "Migration between a Mixed CTN and an STP-only CTN" on page 177	, "Migration between a Mixed CTN and an STP-only CTN" on page 177	, "Migration between a Mixed CTN and an STP-only CTN" on page 177
, "Migration from a Mixed CTN to an STP-only CTN" on page 178	, "Migration from a Mixed CTN to an STP-only CTN" on page 178	, "Migration from a Mixed CTN to an STP-only CTN" on page 178
, "Migration from an STP-only CTN to a Mixed CTN" on page 182	, "Migration from an STP-only CTN to a Mixed CTN" on page 182	, "Migration from an STP-only CTN to a Mixed CTN" on page 182
, "Migration from an ETR network to an STP-only CTN" on page 184	, "Migration from an ETR network to an STP-only CTN" on page 184	, "Migration from an ETR network to an STP-only CTN" on page 184
, "Installation of a new STP-only CTN" on page 189	, "Installation of a new STP-only CTN" on page 189	, "Installation of a new STP-only CTN" on page 189
, "Installation of a single-server STP-only CTN" on page 191	, "Installation of a single-server STP-only CTN" on page 191	, "Installation of a single-server STP-only CTN" on page 191
, "Enhancing a single-site ETR network to a Mixed CTN across two sites" on page 194	, "Enhancing a single-site ETR network to a Mixed CTN across two sites" on page 194	, "Enhancing a single-site ETR network to a Mixed CTN across two sites" on page 194

Migration from an ETR network to a Mixed Coordinated Timing Network involves tasks that fall into the following broad categories:

- ▶ Preparation: tasks planned and implemented prior to enabling STP
- ▶ Installation: tasks that must be completed prior to configuring a Mixed CTN
- ▶ Activation: steps that will migrate the ETR network to a Mixed CTN

Considerations before activating an STP-only CTN

Before an STP-only CTN configuration can become active and servers can exchange timekeeping messages, consider the following:

- ▶ Assign the CTN ID for each server that must be in the CTN.
- ▶ Assign the Preferred Time Server and the Current Time Server.

The server that is assigned the role of the Current Time Server becomes the active stratum 1 when an STP-only CTN configuration becomes active. In most cases, the Preferred Time Server is assigned also to be the Current Time Server.

- ▶ Optionally, a Backup Time Server can be assigned. This avoids a single point of failure.

The role of the Backup Time Server is to take over the role of the stratum 1 server in the event of planned or unplanned outages affecting the Preferred Time Server. Only the Preferred Time Server or the Backup Time Server can be the Current Time Server (that is, the active stratum 1 server) at any instant of time.

- ▶ If the CTN configuration comprises three or more servers, assign an Arbiter.

Migration between an ETR network and a Mixed CTN

This section describes a migration from an ETR network to a Mixed CTN, and then going back from a Mixed CTN configuration to an ETR network.

Migration from an ETR network to a Mixed CTN

Example scenarios are introduced in order of increasing difficulty, starting with the simplest. This scenario assumes the following:

- ▶ All servers and CFs are STP capable and are already synchronized in an ETR network.
- ▶ Coupling links connectivity is already in place for Parallel Sysplex, and the existing links can be used for STP messaging.
- ▶ The HMC on the operations LAN has access to all servers (not shown in Figure A-2 on page 173). All servers in this example are located within a single data center.

This starting configuration is illustrated in Figure A-2 on page 173. The STP-capable servers CEC1, CEC2, and CEC3 are configured in an ETR network, synchronized by the Sysplex Timers. Logical partitions P1, P2, and P3 are z/OS images in a Parallel Sysplex, using internal coupling facilities on servers CEC1 and CEC2.

The objective is to migrate to a Mixed CTN with the following configuration:

- ▶ CEC1, CEC2, and CEC3 participating in the Mixed CTN.
- ▶ CEC1 and CEC2 still in ETR timing mode, connected to the Sysplex Timers. Having two servers connected allows for high availability.
- ▶ CEC3 switched into STP timing mode, synchronized to CEC1 or CEC2.

Figure A-4 on page 175 illustrates the target configuration. Note that the flexibility of the Mixed CTN allows for multiple variations, and this example only illustrates one possible example of what can be achieved. Other variations are possible.

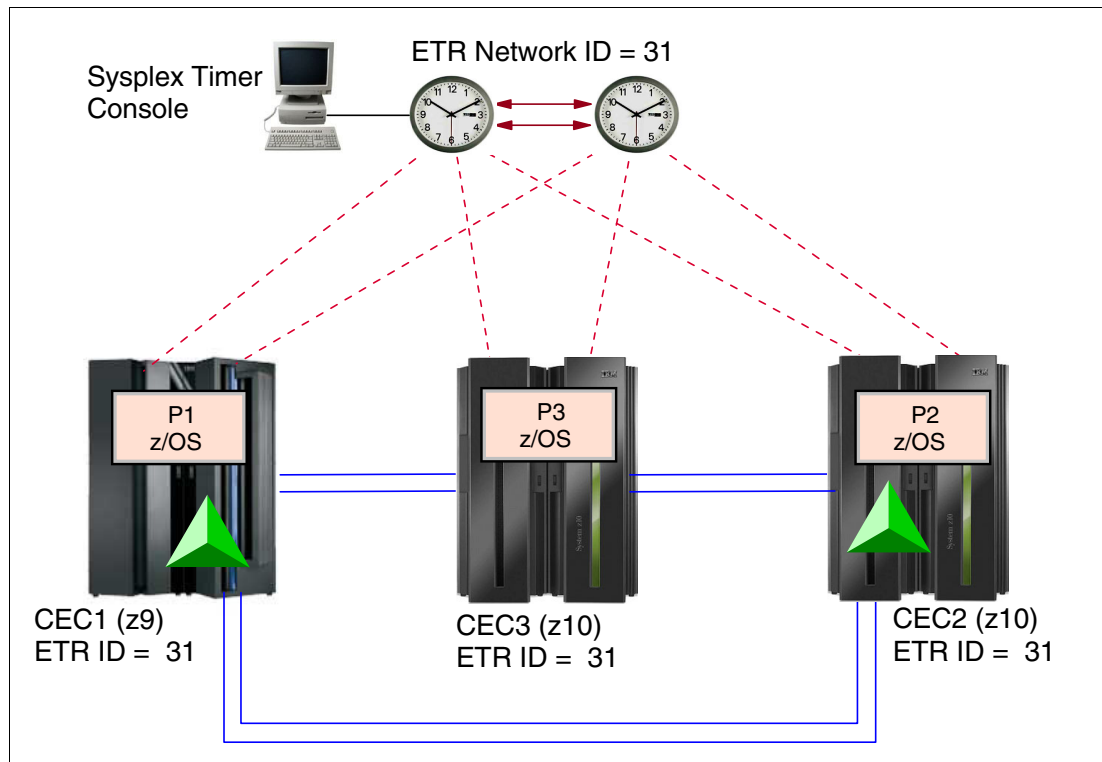


Figure A-2 Migration from an ETR network to a Mixed CTN: start scenario

Preparation

Complete the preparation tasks for the intended Mixed CTN prior to performing either the hardware or the z/OS installation tasks.

In this example there are no z/OS preparation tasks. However, you must install hardware MCLs on CEC1, CEC2, and CEC3. After the MCLs are installed, CEC1, CEC2, and CEC3 are STP-capable, but at this point STP is *not* enabled.

Installation

For installation:

- ▶ The STP feature must be installed on each server (CEC1, CEC2, and CEC3). Installation of the STP feature is concurrent.

After installation, the System (Sysplex) Time task is available on the SE and HMC to configure a server to use STP. The servers are now STP-enabled.

- ▶ Install the STP-related z/OS maintenance on all logical partitions (P1, P2, and P3), and update CLOCKxx, if needed.

A re-IPL of the z/OS images is required for the maintenance to take effect and for z/OS to recognize the newly installed STP feature.

Activation of the Mixed CTN

After the preparation and installation steps are completed, the CTN configuration and activation steps can be initiated. Perform all actions from the Hardware Management Console, using the System (Sysplex) Time task.

Note: When properly planned and executed, all operations from this point onward are concurrent to server operations and nondisruptive to the z/OS images.

Assign a CTN ID

The CTN ID is of the format [STP network ID] - [ETR network ID]. The ETR network ID already assigned to the ETR network should not be changed because it reflects the existing Sysplex Timer configuration. Only the STP network ID portion of the CTN ID should be changed.

- ▶ In this example, the STP ID selected is ITSOPK. The value is entered on each server so that the same CTN ID is assigned to CEC1, CEC2, and CEC3.
- ▶ When completed, all servers are defined to a Mixed CTN with CTN ID = [ITSOPK] - [31]. All servers are STP-configured. They are still connected to the Sysplex Timer and running in ETR timing mode (Figure A-3).

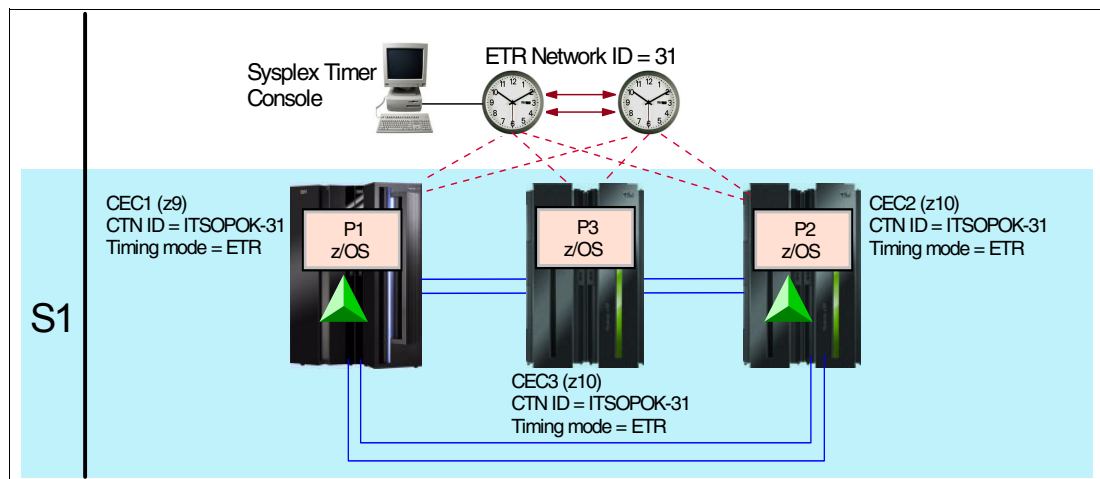


Figure A-3 Migration from an ETR network to a Mixed CTN: all servers in ETR timing mode

Migrate a server to STP timing mode

This step can be considered as optional and should not be used for the last (or preferably the last two) servers in the Mixed CTN to avoid losing the clock source for z/OS images with STPMODE YES or ETRMODE YES. If the Mixed CTN consists of one or two servers only, the next step might be to continue the migration to an STP-only CTN, which is described in , “Migration between a Mixed CTN and an STP-only CTN” on page 177.

Since the Mixed CTN has been created the user may want to switch an individual server to STP timing mode. However, this step does not migrate into an STP-only CTN. It simply switches a server into STP timing mode, forcing this individual server to use STP messaging through existing coupling links.

Important: The user should *not* disable the ETR ports on the last stratum 1 server because this removes the time source for the entire Mixed CTN.

In our example, the next step is to change the timing mode on CEC3 from ETR timing mode to STP timing mode. We consider this step because there are still two other servers (CEC1 and CEC2) connected to the Sysplex Timer and CEC3 does have coupling links to these servers. If the server being transitioned to STP timing mode does not have coupling links to at least one existing stratum 1 server the operation will not be successful> it will transition to local mode.

Follow these steps:

1. Verify that the STP connectivity is established between CEC3 and the two other servers. From the STP Status tab, the System Information section must show one line for each attached server.
2. From the ETR configuration panel of CEC3, disable its ETR ports.
3. CEC3 recognizes the loss of the Sysplex Timer signals and automatically switches to STP timing mode. CEC3 now synchronizes to either CEC1 or CEC2 using STP messages.

Note: This technique should *not* be used for the last two servers stepping to the ETR signals because the Sysplex Timer is the time source for the CTN. In a Mixed CTN, two stratum 1 servers should be configured in order to avoid a single point of failure.

The final configuration is shown in Figure A-4.

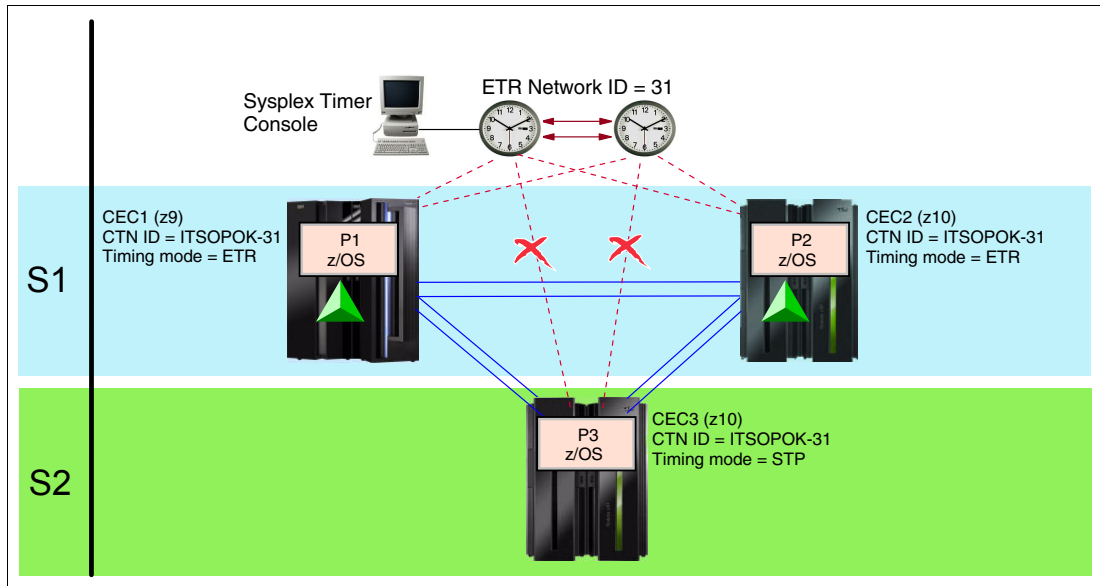


Figure A-4 Migration from an ETR network to a Mixed CTN: end scenario

Migration from a Mixed CTN to an ETR network

This scenario illustrates a migration going from a Mixed CTN back to an ETR network. This capability could be useful if, for example, you want to use a change window to do STP testing and then go back to a previous configuration concurrently. Figure A-5 provides an example. It assumes that all servers are already time synchronized using a Mixed CTN configuration. Logical partitions P1, P2, and P3 are three z/OS images participating in a Parallel Sysplex.

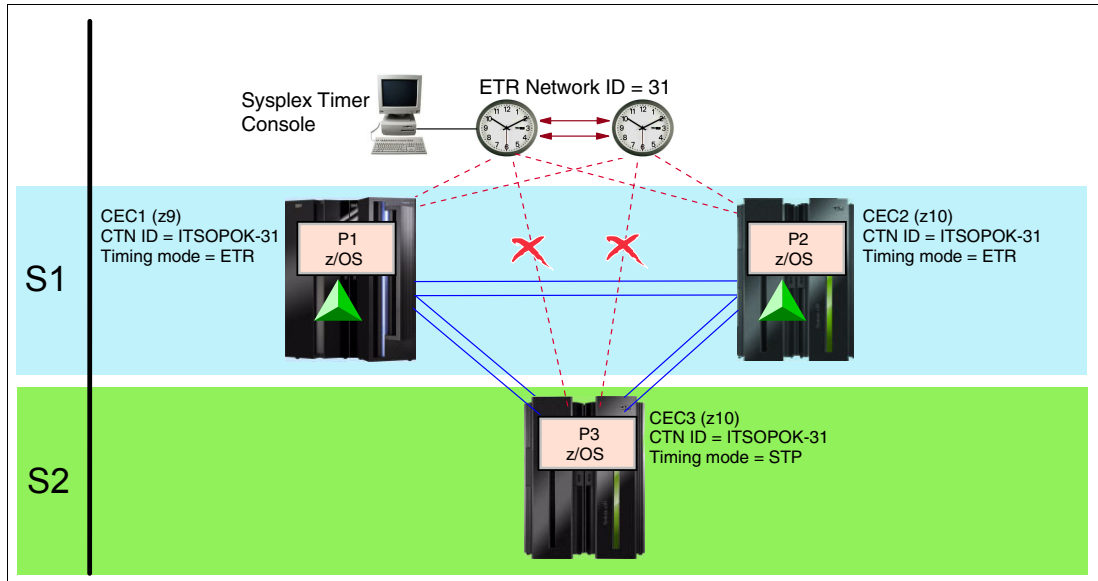


Figure A-5 Migration from a Mixed CTN to an ETR network: start scenario

The configuration planning considerations are:

- ▶ Since the servers are already in a Mixed CTN, there is no need to be concerned about enabling the hardware or software.
- ▶ When the migration to the ETR network is complete, the objective is to have all servers back into an ETR Network.

The CLOCKxx statements, such as ETRMODE and STPMODE, are available. Both are already specified as YES by default in all logical partitions participating in this sysplex example.

The resulting configuration is illustrated in Figure A-6 on page 177.

Preparation and installation

To prepare the hardware for migration, the CEC3 connections to the Sysplex Timer must still be in place. Ensure that the links are still available and that the Sysplex Timer ports are configured online on the Sysplex Timer console.

Activation

After the preparation and installation step is completed, the CTN configuration and activation steps can be initiated. All actions are now performed from the Hardware Management Console using the System (Sysplex) Time task.

When properly planned and executed, all operations from here on are concurrent to server operations and nondisruptive to the z/OS images.

Migrate a server to ETR timing mode

Enable the ETR ports. In this example, the CEC3 server ETR ports are enabled. This changes CEC3 from a stratum 2 to a stratum 1. It returns the server to ETR timing mode.

At this point, the CEC1, CEC2, and CEC3 are now in ETR timing mode. They still participate in a Mixed CTN with CTN ID = [ITSOPOK] - [31].

Modify the CTN ID

The next step is to modify the CTN ID. When both the STP network ID and ETR network ID are specified, the resulting CTN is referred to as a Mixed CTN:

- ▶ Modify the CTN ID on each STP-configured server. From the STP Configuration panel on the HMC, delete the STP ID portion of the CTN ID.
 - Do *not* change the ETR network ID, which identifies the ETR network.
- ▶ All servers in the configuration (CEC1, CEC2, and CEC3) are synchronized to the Sysplex Timer. They are part of the ETR network.

The migration from a Mixed CTN to an ETR network is complete. Figure A-6 shows the resultant ETR network.

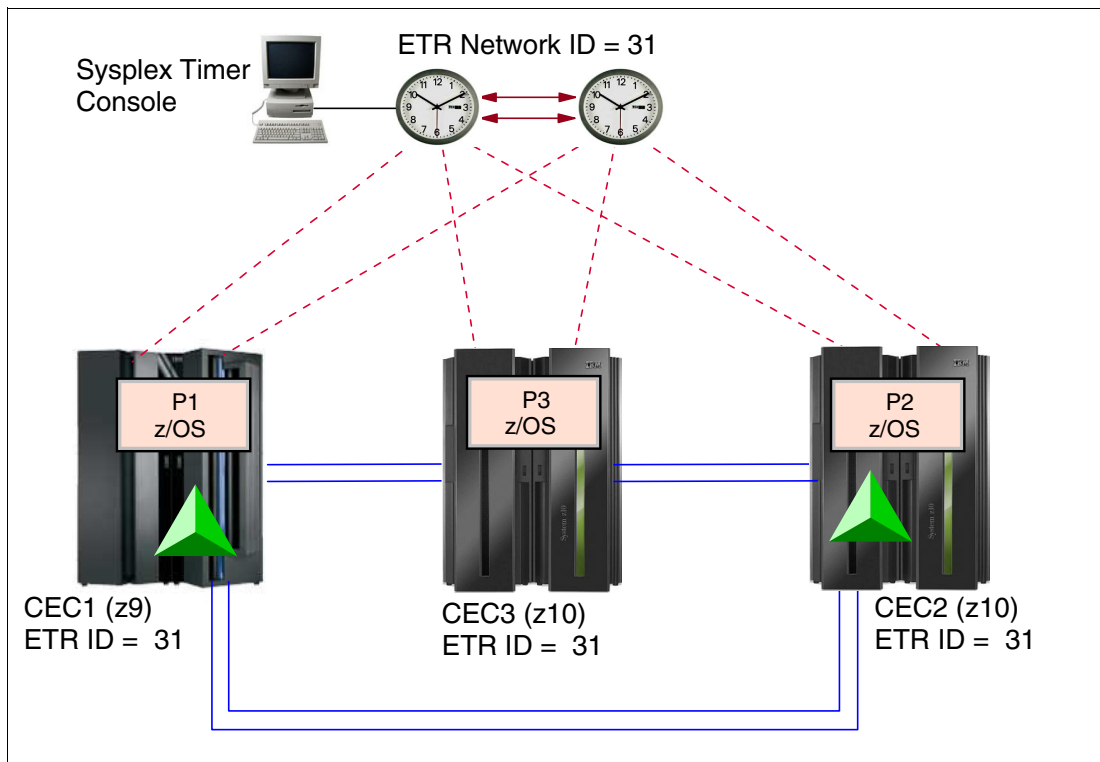


Figure A-6 Migration from a Mixed CTN to an ETR network: end scenario

Migration between a Mixed CTN and an STP-only CTN

This section describes a migration from a Mixed CTN to an STP-only CTN, and then going back from an STP-only configuration to a Mixed CTN.

Note: Keep the sysplex timer and the STP-only CTN in sync until you are sure that no migration back from the STP-only CTN to a mixed CTN must occur. This can be accomplished by synchronizing both the ETR and the CTN to the same external time source. Before the reverse migration can successfully complete, the ETR and CTN must have the same time. Otherwise, adjustment steering will occur, which may take several days to complete.

Migration from a Mixed CTN to an STP-only CTN

This scenario is an example of a migration from a Mixed CTN to an STP-only CTN. This scenario assumes the following:

- ▶ All servers and CFs that are to participate in the CTN are STP-configured and already synchronized in a Mixed CTN.

There may be other non-STP-capable servers present, but they cannot be considered for the migration to an STP-only CTN. They remain either synchronized to a different ETR network or unsynchronized.

- ▶ All z/OS images that require time synchronization, such as those in a Parallel Sysplex, already have STP-related maintenance applied.
- ▶ Coupling link connectivity is already in place for Parallel Sysplex. The existing links are available for STP messaging and satisfy redundancy requirements. The STP messaging capability and redundancy was verified from the STP STATUS panel from each CEC, System Information section.
- ▶ The HMC on the operations LAN has access to all servers.

During a migration from a Mixed to STP-only CTN, the assigned Current Time Server inherits the time definitions from the ETR. No time initialization is required.

All the servers in this example are located within a single data center. Figure A-7 illustrates an example of such a configuration. All servers are synchronized in a Mixed CTN, and the CTN ID is [ITSOPOK] - [31].

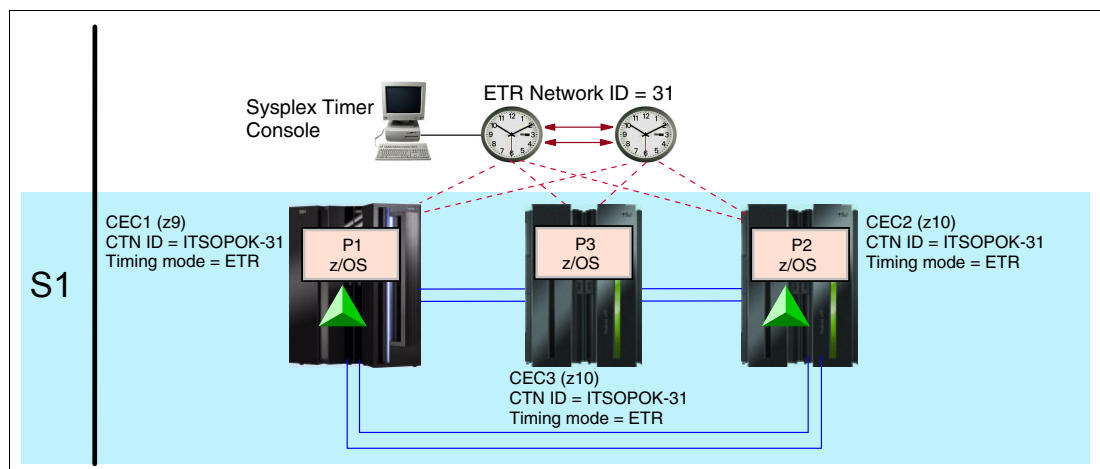


Figure A-7 Migration from a Mixed CTN to an STP-only CTN: start scenario

When the migration to an STP-only CTN is completed, the objectives are to have:

- ▶ CEC1, CEC2, and CEC3 configured in an STP-only CTN.
 - The CTN ID contains an STP network ID part but no ETR network ID. A Preferred Time Server, a Backup Time Server, and an Arbiter are defined.
- ▶ CEC1, CEC2, and CEC3 are not synchronized to the Sysplex Timer.
 - All ETR ports on the servers are disabled. However, leave the ETR links in place for a certain period of time to allow for a backout plan. Depending on individual change management constraints, the ETR links can be disconnected once a migration back to a Mixed CTN is not being considered anymore.

Figure A-9 on page 181 illustrates the target configuration.

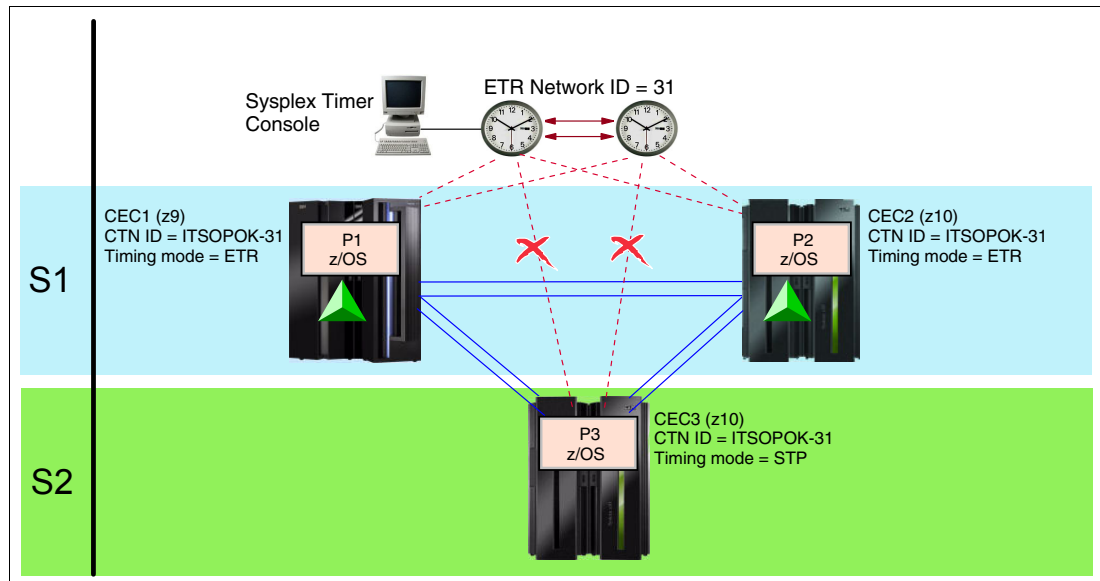


Figure A-8 Migration from a Mixed CTN to an STP-only CTN: phased approach

Preparation

To prepare for the migration, the following is required:

- ▶ At least one HMC must be able to access each server participating in an STP-only CTN. In an STP-only CTN, the HMC is used to modify the time zone, daylight saving time, and possibly leap second offsets. If using dial-out ETS, at least one HMC must be chosen that has the ability to dial out via an analog phone line.
- ▶ If using an ETS, dial-out or NTP, access to the external time source must be configured and tested. Refer to “External time source” on page 46.
- ▶ The target CTN in this example has three servers, so plan to assign roles. A Preferred Time Server, Backup Time Server, and an Arbiter are selected. Although optional, the Arbiter provides additional means for the Backup Time Server to determine whether it should take over in the event of unplanned outages.

Activation

When properly planned and executed, all operations are concurrent to server operations and nondisruptive to the z/OS images. There are two approaches possible:

- ▶ A phased approach, *one* server at a time. This approach entails:
 - Migration of each individual server to STP timing mode, except the last two.
 - Definition of the server roles and assignment of the Current Time Server.
- ▶ A global approach.

This is a one-step operation that transitions *all* STP-configured servers in the Mixed CTN into STP timing mode. This is done by defining the server roles and assigning the Current Time Server.

Next we explain each approach in more detail.

Phased approach

A phased approach migrates individual servers to STP timing mode.

1. Change a server to STP timing mode.

Each server but the last two is individually changed to STP timing mode. This is done by disabling the ETR ports from the ETR Configuration panel.

Tip: Do *not* disable ETR ports via the Sysplex Timer console. The server ETR ports are disabled via the System (Sysplex) Time task available on the SE and HMC.

This approach preserves synchronization to the ETR until the last server is switched. Each server individually switched to STP timing mode becomes a stratum 2 and continues to use a stratum 1 as its time source.

This approach cannot be used against the last server with Sysplex Timer connectivity (or rather, with the last two servers, because in a Mixed CTN it is good practice to preserve two stratum 1 servers until the final step of the migration). When the ETR ports are disabled on the *last* server connected to the Sysplex Timer, *all* servers in the CTN lose their timing source and become unsynchronized. Unsynchronized z/OS images issue WTOR message IEA394A.

In this example, only CEC3 is switched into STP timing mode. This is achieved by going to the ETR Configuration panel for CEC3 and disabling the ETR ports. CEC3 is synchronized to CEC1 or CEC2 using STP messages.

The next step in this migration path is to define the server roles and assign the Current Time Server.

2. Assign server roles.

This step completes the transition of all servers to STP timing mode. It migrates the network to an STP-only CTN. This is done from the Network Configuration tab. It is the last step in a phased approach.

Important: No manual action is needed to update the CTN ID. The change is made by STP when the Current Time Server is established. The operation must be initiated from the server that will become the Current Time Server as a result of the operation.

In this example, CEC2 is assigned the role of Preferred Time Server (PTS). It is also the Current Time Server (CTS). The Backup Time Server (BTS) and Arbiter roles are defined at this point. In this example CEC1 is assigned as the BTS and CEC3 as the Arbiter. The

change is initiated from the Network Configuration tab on CEC2 that is to become the Current Time Server.

When the role assignment is done and the **Apply** button is clicked, the following actions are taken:

- The CEC2 becomes the Current Time Server, stratum 1, and broadcasts its role to the other servers in the network.
- The ETR ports on *all* the servers in the Mixed CTN are disabled.
- The ETR network ID is removed from the CTN ID for all servers in the CTN. The resulting CTN ID in this example is [ITSOPOK] - []. The new CTN ID is broadcast to all other STP-configured servers in the CTN.

The operation migrates the Mixed CTN to an STP-only CTN with immediate effect.

Figure A-9 shows the resulting configuration.

Global approach

This approach migrates the network to an STP-only CTN in only one step. This is done from the Network Configuration tab.

In this example, the CEC2 is assigned the role of Preferred Time Server (PTS). It is also the Current Time Server (CTS). The Backup Time Server (BTS) and Arbiter roles are defined at this point. CEC1 is assigned as the BTS and CEC3 as the Arbiter. The change is initiated from the Network Configuration tab on CEC2 that is to become the Current Time Server.

When the role assignment is done and **Apply** is clicked, the following actions are taken:

- ▶ CEC2 becomes the Current Time Server, stratum 1, and broadcasts its role to the other servers in the network.
- ▶ The ETR ports on *all* the servers in the Mixed CTN are disabled.
- ▶ The ETR network ID is removed from the CTN ID for all servers in the CTN. The resulting CTN ID in this example is [ITSOPOK] - []. The new CTN ID is broadcast to all other STP-configured servers in the CTN.

The operation migrates the Mixed CTN to an STP-only CTN with immediate effect. Figure A-9 shows the resulting configuration.

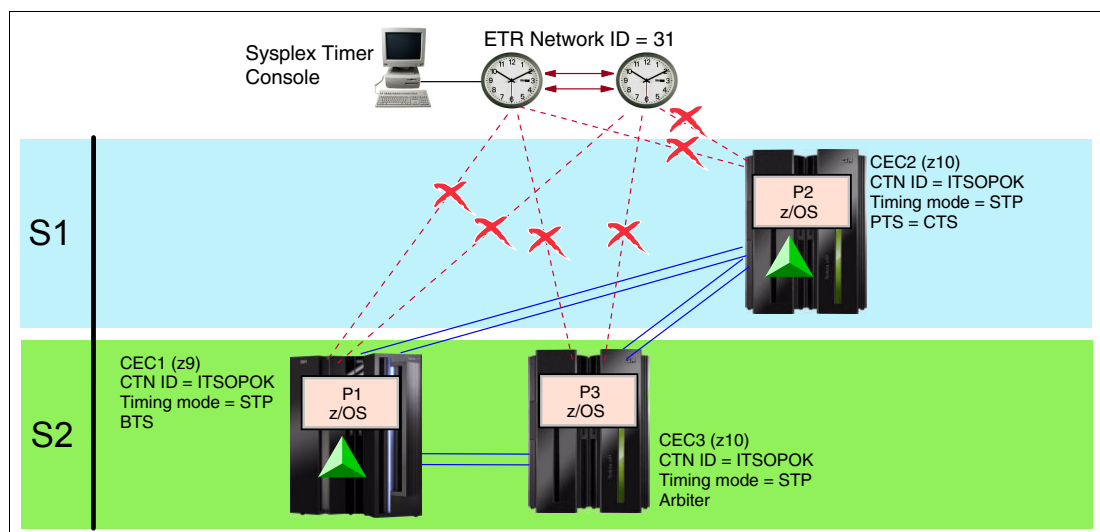


Figure A-9 Migration from a Mixed CTN to an STP-only CTN: end scenario

Note: Disable the ETR ports to the CECs in STP Timing Mode on the sysplex timer after a successful migration. Otherwise, during a POR the CEC will initially adjust its internal clock using the timing signals from the sysplex timer and will later adjust it to the CTN time. Depending on the discrepancy this adjustment might take some time.

Remember to enable the ports on the sysplex timer again when you plan a reverse migration.

Migration from an STP-only CTN to a Mixed CTN

This is a reverse scenario from the one described in the previous section. This capability could be useful if, for example, you want to use a change window to do STP testing and then go back to a previous configuration concurrently. Figure A-10 shows an example for migration back to a Mixed CTN.

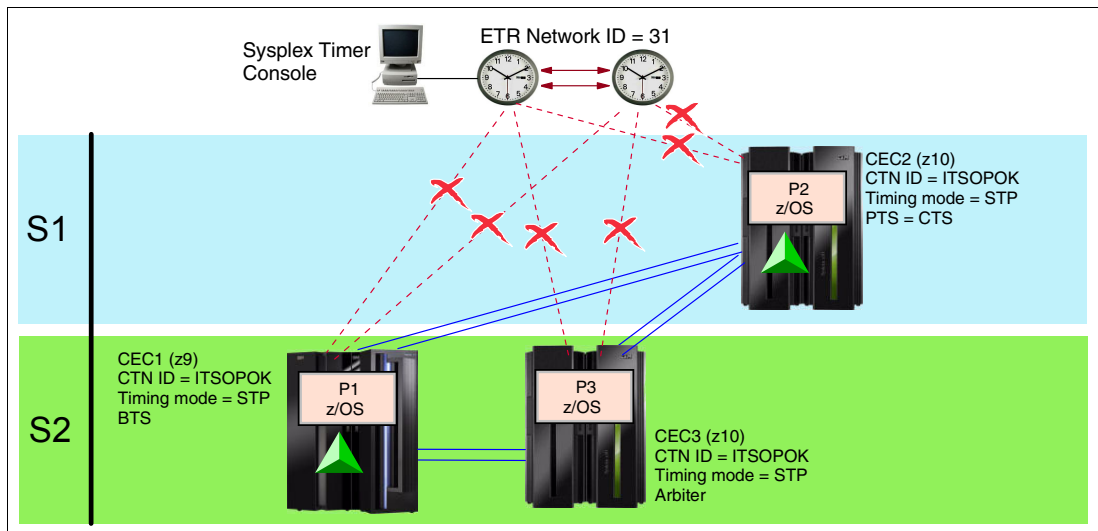


Figure A-10 Migration from an STP-only CTN to a Mixed CTN: start scenario

We assume that all operating z/OS images that require multisystem synchronization already have the STP-related maintenance applied.

The objectives for the migration back to a Mixed CTN are:

- ▶ To have all servers in the timing network end up with a CTN ID = [ITSOPK] - [31].
- ▶ Both the Preferred and the Backup Time Server (CEC2 and CEC1, in this example) are established as stratum 1, stepping to the Sysplex Timers.
- ▶ CEC3 as stratum 2 is in STP timing mode, synchronized to a stratum 1 server.

Figure A-11 illustrates the target configuration.

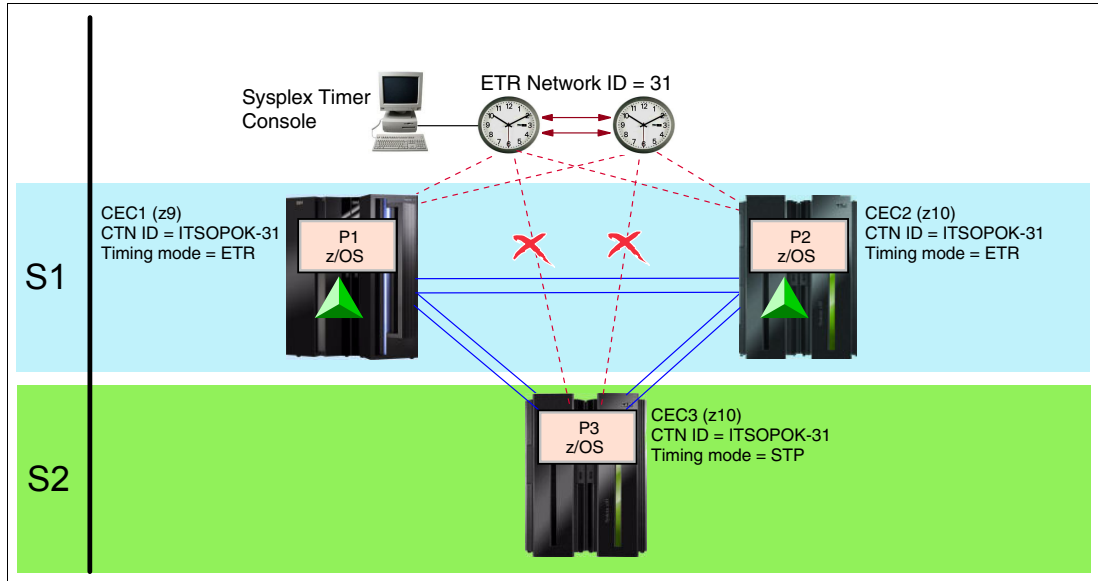


Figure A-11 Migration from an STP-only CTN to a Mixed CTN: end scenario

Preparation

If the Sysplex Timers shown in Figure A-10 on page 182 have not been changed since the migration to STP-only CTN and all ETR cabling is still intact, there is no special preparation.

If the Sysplex Timer configuration has been changed, review the following preparation steps and perform those that are applicable:

1. Verify that an operational Sysplex Timer console is attached to the Sysplex Timers. A Sysplex Timer console is required for all timing-related functions in a Mixed CTN. This console is used for all timing-related functions, such as initializing time, scheduling daylight saving time, adjusting time, and so on.
2. Ensure that all requisite Licensed Internal Code (LIC) has been applied to the Sysplex Timers to support a Mixed CTN configuration.
3. Ensure that the Sysplex Timer ports attached to the Preferred and Backup Time Servers are online and transmitting ETR data.

For high availability, we recommend a minimum of two stratum 1 servers stepping to the Sysplex Timers in an expanded availability configuration.

All z/OS images running in the STP-only CTN have the STP-related maintenance already applied. The Mixed CTN has no additional requirements.

Installation

If the Sysplex Timer configuration has not been changed, there are no special installation tasks. If the Sysplex Timer is a new configuration, refer to the following publications, supplied with the Sysplex Timer, for detailed instructions:

- ▶ *Planning for the 9037 Model 2 Sysplex Timer, SA22-7233*
- ▶ *Using the 9037 Model 2 Sysplex Timer, SA22-7230*

Activation

After the preparation and installation steps are completed, the activation steps can be initiated. All actions are now performed from the Hardware Management Console using the System (Sysplex) Time task.

When properly planned and executed, all operations are concurrent to server operations and nondisruptive to the z/OS images.

Change the CTN ID

A change to the CTN ID can be either GLOBAL or LOCAL. The migration from STP-only to Mixed CTN is initiated via a GLOBAL change to the CTN ID. A GLOBAL change applies to all servers in the CTN. It is initiated from the Network Configuration panel. This can only be done from the Current Time Server or the request is rejected.

A valid ETR network ID (31, in this case) is entered on the Network Configuration panel on the Current Time Server. CEC2 is the Current Time Server in this example.

The action initiates the migration back to a Mixed CTN. As part of the migration, the CTS first steers the CST until it is synchronized with the Sysplex Timer time.

It is most likely that the Sysplex Timer time is out of synch with Coordinated Server Time. An estimated completion time message is displayed, indicating how long it will take for the migration to complete. The amount of time depends on how close the Sysplex Timer time is to the CST. If the time estimated to complete the migration is acceptable, then respond to the message accordingly. Otherwise, indicate that the migration should not proceed. If accepted, the Current Time Server enables the ETR ports for the Preferred and Backup Time Servers if previously assigned (CEC2 and CEC1, in this example). The Coordinated Server Time will start to steer back to the time at the Sysplex Timer. During the migration, any task that involves changing the configuration is not allowed. Migration is only completed when the steering is complete.

Note: To minimize the migration duration, a good practice is to synchronize the Sysplex Timers to the same external time source used by the CTN prior to starting the migration.

All other servers leave the ETR ports disabled. Other servers in the timing network remain in STP timing mode, receiving their timing signals from the stratum 1 servers.

Figure A-11 on page 183 shows the final configuration.

Enable ETR ports for additional servers

Both the Preferred and Backup Time Server automatically have their ETR ports enabled, and switch from STP to ETR timing mode. Other servers in the CTN remain in STP timing mode.

In our example, CEC3 can have its ETR ports enabled manually to transition to ETR timing mode. The change is made from the ETR Configuration panel and is local to the server.

Migration from an ETR network to an STP-only CTN

A migration from an ETR Timing Network to an STP-only CTN is not possible in one step. The existing ETR network must first be migrated to a Mixed CTN, and then the Mixed CTN must be migrated to an STP-only CTN. However, the two steps can be combined such that they appear as a single migration.

Note: Keep the sysplex timer and the STP-only CTN in sync until you are sure that no migration back from the STP-only CTN to a mixed CTN or ETR network must occur. This can be accomplished by synchronizing both the ETR and the CTN to the same external time source. Before the reverse migration can successfully complete the ETR and CTN must have the same time. Otherwise, adjustment steering occurs, which may take several days to complete.

Because the target configuration is an STP-only CTN, your planning should include all of the requirements for an STP-only CTN, as described here. All servers in the final configuration must be STP-capable. All z/OS images that require synchronization must have the STP maintenance applied.

This scenario assumes that:

- ▶ All servers and CFs that are to participate in the CTN are STP-capable and are already synchronized in an ETR network.
There may be other non-STP-capable servers present but they cannot be considered for the migration to an STP-only CTN. They remain either synchronized to a different ETR network or unsynchronized.
- ▶ Coupling link connectivity is already in place for STP messaging and satisfies redundancy requirements for STP.
- ▶ The HMC on the operations LAN has access to all servers in the intended STP-only CTN. All the servers in this example are located within a single data center.

During a migration from a Mixed to STP-only CTN, the assigned Current Time Server inherits the time definitions from the ETR. No time initialization is required. However, the time zone offset must be configured to complete the migration.

Figure A-12 shows an example. STP-capable servers CEC1, CEC2, and CEC3 are configured in an ETR network, synchronized by the Sysplex Timers. Logical partitions P1, P2, and P3 are z/OS system images in a Parallel Sysplex, using internal coupling facilities (ICFs) on servers CEC1 and CEC2. All the servers in this example are located within a single data center.

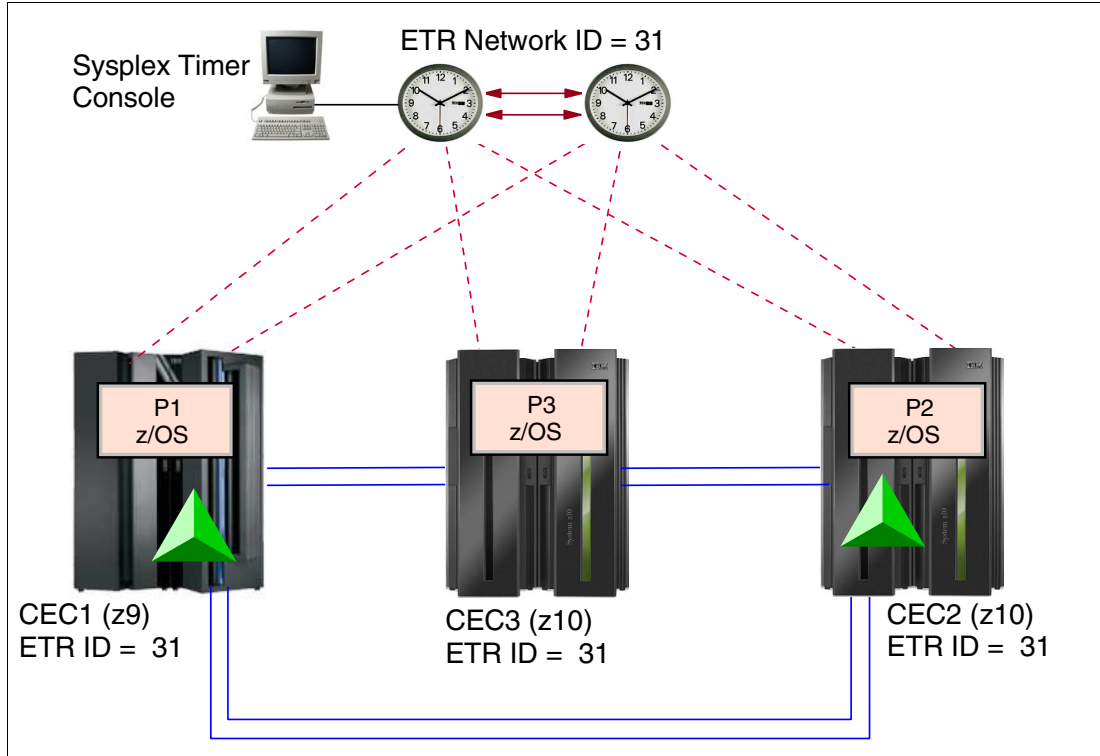


Figure A-12 Migration from an ETR network to an STP-only CTN: start scenario

When the migration to STP-only CTN is complete:

- ▶ CEC1, CEC2, and CEC3 are configured in an STP-only CTN. The CTN ID contains only an STP network ID and no ETR network ID. A Preferred Time Server, a Backup Time Server, and an Arbiter are defined.
- ▶ CEC1, CEC2, and CEC3 has no active connection to the Sysplex Timer. ETR ports on all servers are disabled. However, the ETR links should remain in place for a certain period of time to allow for a backout plan. Depending on individual change management constraints, the ETR links can be disconnected once a migration back to a Mixed CTN is no longer being considered.

Note: Disable the ETR ports to the CECs in STP timing mode on the Sysplex Timer after a successful migration. Otherwise, during a power-on reset (POR) the CEC initially adjusts its internal clock using the timing signals from the sysplex timer and will later adjust it to the CTN time. Depending on the discrepancy between the two clock sources, this adjustment might take some time.

Remember to enable the ports on the sysplex timer again when you plan a reverse migration.

Preparation

There are preparation tasks associated with hardware and software components. It is important to complete all hardware and z/OS preparation tasks for the intended STP-only CTN prior to performing either the hardware or z/OS installation tasks.

Prepare the hardware for migration by configuring your system as follows:

- ▶ At least one HMC in an STP-only CTN must be configured with connectivity to each server in the CTN. The HMCs must be running the HMC Application V2.9.2 or later.

In an STP-only CTN, the HMC is used to modify the time zone, daylight saving time, and possibly leap second offsets. If using an external time source, an HMC must be chosen that has the ability to dial-out via an analog phone line.

- ▶ Hardware MCLs must be installed on CEC1, CEC2, CEC3, and the sysplex timer.
- ▶ If using an external time source, dial-out or NTP, access to the ETS must be configured and tested. Refer to “External time source” on page 46.
- ▶ The target CTN in this example has three servers, so plan to assign roles. A Preferred Time Server, Backup Time Server, and Arbiter are selected. Although optional, the Arbiter provides additional means for the Backup Time Server to determine whether it should take over in the event of unplanned outages.

Installation

For installation:

- ▶ Install the STP feature on CEC1, CEC2, and CEC3. Installation of the STP feature is concurrent. After installation of the STP feature, the server can be configured for STP.
 - The System (Sysplex) Time task is available on the SE and HMC to configure a server to use STP. Various functions, such as initialization of the CTN ID, can now be performed.
 - Coupling facility partitions automatically recognize the new functions.
- ▶ Install the STP-related maintenance on all z/OS images that require time synchronization. In the example, the STP-related maintenance must be installed on all images (P1, P2, and P3).

A re-IPL of the z/OS images is required to recognize the newly installed STP feature and for the maintenance to take effect.

Activation

After you complete the preparation and installation tasks, you can initiate the CTN configuration and activation steps.

When properly planned and executed, all operations from here on are concurrent to server operations and nondisruptive to the z/OS images. All actions are initiated from the HMC using the System (Sysplex) Timer task.

Assign a CTN ID

Within the Network Configuration panel, do *not* change the ETR network ID already assigned to the ETR network. Only the STP network ID portion of the CTN ID should be specified.

In this example, the value is entered for each server so that the same CTN ID is assigned to CEC1, CEC2, and CEC3. When done, CEC1, CEC2, and CEC3 are all running in ETR timing mode and are defined in a common Mixed CTN with CTN ID = [ITSOPOK] - [31].

Assign the server roles

This step transitions all servers to STP timing mode. It migrates the network to an STP-only CTN. This is done from the Network Configuration tab. However, start by verifying the STP messaging capability and Inks redundancy from the STP STATUS tab.

In this example, CEC2 is assigned the role of Preferred Time Server (PTS). It is also the CTS. The Backup Time Server and Arbiter roles are also defined at this point. In this example CEC1 is assigned as the BTS and CEC3 as the Arbiter. The change is initiated from the Network Configuration tab on CEC1 that is to become the Current Time Server.

Important: No manual action is needed to update the CTN ID. The change is made by STP when the Current Time Server is established. The operation must be initiated from the server that will become the Current Time Server as a result of the operation.

When the role assignment is done and **Apply** is clicked, the following actions are taken:

- ▶ CEC2 becomes the Current Time Server, stratum 1, and broadcasts its role to the other servers in the network.
- ▶ The ETR ports on all the servers in the Mixed CTN are disabled.
- ▶ The ETR network ID is removed from the CTN ID for all servers in the CTN. The resulting CTN ID in this example is [ITSOPOK] - [].

The new CTN ID is broadcast to all other STP-configured servers in the CTN.

Figure A-13 shows the resulting configuration.

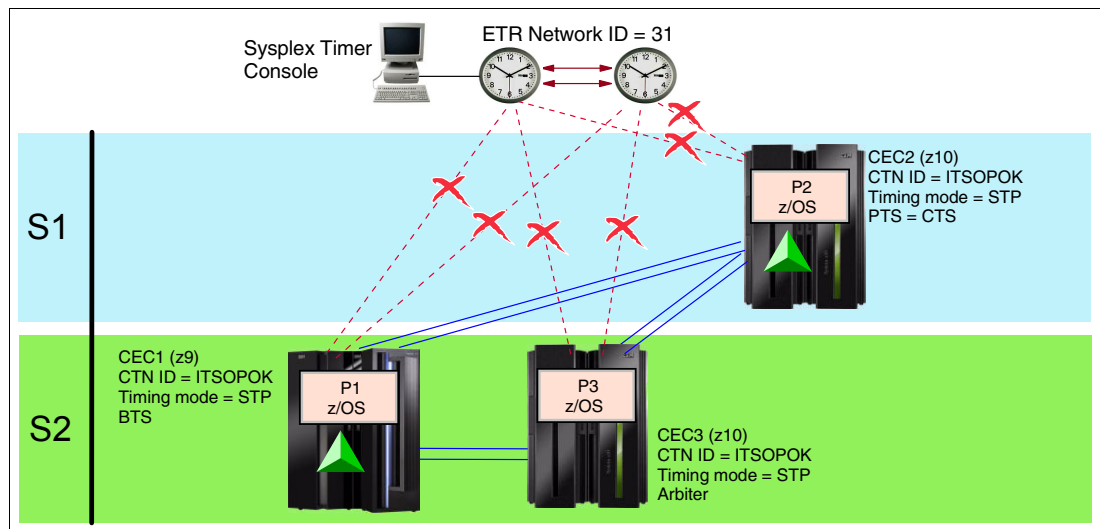


Figure A-13 Migration from an ETR network to an STP-only CTN: end scenario

Note: Disable the ETR ports to the CECs in STP timing mode on the Sysplex Timer after a successful migration. Otherwise, during a POR the CEC initially adjusts its internal clock using the timing signals from the sysplex timer and later adjusts it to the CTN time. Depending on the discrepancy, this adjustment might take some time.

Remember to enable the ports on the sysplex timer again when you plan a reverse migration.

Installation of a new STP-only CTN

Because the target configuration is an STP-only CTN, your planning should include all the requirements for an STP-only CTN. This scenario assumes that:

- ▶ All servers and CFs that are to participate in the CTN are STP capable.
 - There may be non-STP-capable servers present, but they are not considered for the installation of an STP-only CTN. They can only remain unsynchronized.
- ▶ There are no coupling facilities and no coupling facility links.
- ▶ The HMC on the operations LAN has access to all servers. All the servers in this example are located within a single data center.
- ▶ There are no Sysplex Timers in this configuration. Therefore, there is no need for migration.

Figure A-14 shows an example of the configuration. The example assumes that there is no timing network installed and that the servers are not yet time synchronized. All the servers in this example are located within a single data center.




Figure A-14 Installation of a new STP-only CTN: start scenario

When the initialization of the STP-only CTN is complete, the objective is to have CEC1, CEC2, and CEC3 configured as members of an STP-only CTN. Each server is assigned a specific role in the CTN (Preferred Time Server, Backup Time Server, or Arbiter).

Figure A-15 illustrates the target configuration.

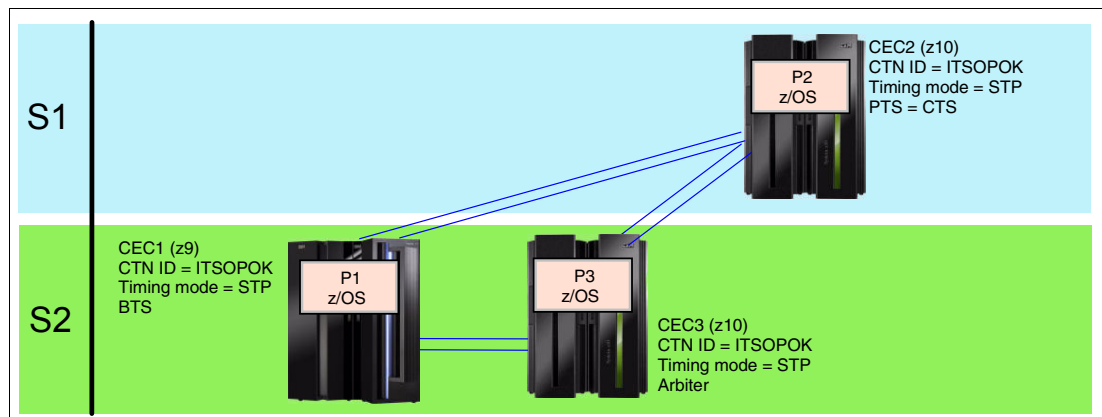


Figure A-15 Installation of a new STP-only CTN: end scenario

Preparation

Prepare the hardware for migration by configuring your environment as follows:

- ▶ Ensure availability of an STP-capable HMC. The HMCs must be running the HMC Application V2.9.2 or later. The HMC is used to establish the time zone, daylight saving time, and possibly leap second offsets. If using dial-out ETS, at least one HMC must be chosen that has the ability to dial out via an analog phone line.
- ▶ Install hardware MCLs to support STP on the STP-capable servers. In this example, this would be CEC1, CEC2, and CEC3.
- ▶ If using an external time source, dial-out or NTP, access to the ETS must be configured and tested. Refer to 2.3, “External time source” on page 46.
- ▶ Assign roles

CEC2 in this example is selected as the Preferred Time Server (stratum 1) and is assigned as the Current Time Server.

CEC1 is synchronized to CEC2. Thus, it is a Stratum 2. It is assigned the role of Backup Time Server.

CEC3 is also a stratum 2 synchronized to CEC2. The CEC3 is the Arbiter.
- ▶ Order coupling link features on each STP capable server. In this example, we want to add redundant coupling links running between all STP-capable servers (CEC1, CEC2, and CEC3). Because there are no coupling facilities in this example, redundant timing-only links must be defined.

z/OS preparation

To perform z/OS preparation apply STP-related maintenance. This is necessary because servers will be in STP timing mode and time synchronization will be used. A re-IPL of the images is required after installation of the maintenance.

STP feature installation

This task requires the installation of the STP feature on the STP-capable servers. In the example, these are CEC1, CEC2, and CEC3. Three STP features are required.

Installation of the MCL is concurrent. After installation of the STP feature, the servers can be configured for STP.

- ▶ STP panels are available on the SE and HMC. Various functions such as initialization of the CTN ID can now be performed.
- ▶ z/OS system images must be IPLed.

Activation

After the preparation and installation tasks are completed, the CTN can be defined and activated. When properly planned and executed, all operations are concurrent to server operations and nondisruptive to the z/OS images. The activation steps are:

1. Assign a CTN ID to each server.
2. Initialize the time, including offsets.
3. Assign the Current Time Server.

Assign a CTN ID

Each server must be assigned the same CTN ID. Because this is an STP-only CTN and there is no Sysplex Timer synchronization, there is no ETR network ID present. Only the STP network ID portion of the CTN ID is defined.

The task must be performed for each server in the network from the STP Configuration tab, one server at a time. After all of the servers intended to be configured in the STP-only CTN are assigned the same CTN ID, the System (Sysplex) Time task is available on the SE and HMC.

Initialize the time

Because there is no pre-existing ETR network, there is no inheritance of existing time parameters such as time, time zone offsets, or leap seconds. The information must be defined from the server that will be assigned the Current Time Server, CEC2 in this example.

Specific server in an STP-only CTN

If the STP-only CTN to be activated is a single-server CTN the consideration should be taken to also limit the CTN to a one-server CTN. This is to avoid an automatic CTN deconfiguration during power-on reset or power off/on. This can be achieved by selecting the **Only allow the server(s) specified above to be in the CTN** option. See 3.4.2, “Single-server or dual-server auto resume after power-on reset” on page 125, for details. The same applies for a two-server CTN.

Assign the Current Time Server

Up to this point, the servers are still unsynchronized (stratum 0) because no time source has been assigned yet.

In this example, CEC2 is to be the stratum 1. CEC2 is assigned the roles of the Preferred Time Server and the Current Time Server. Thus, CEC2 must be the server that assigns the roles. CEC2 can also be used to assign the roles of the Backup Time Server and the Arbiter. These roles can be applied one at a time if desired, or globally. If the Current Time Server role is to change later, the change must be driven by the server that will become the Current Time Server.

In this example, CEC1 is the Backup Timer Server and CEC3 is the Arbiter. The roles are assigned via the System (Sysplex) Time panels.

The operation initializes an STP-only CTN. The change is initiated from CEC2, which becomes the Current Time Server and broadcasts its role to the other servers in the CTN.

Figure A-15 on page 189 shows the resulting STP-only CTN.

Installation of a single-server STP-only CTN

This section describes the migration process for a single-server environment with no existing time synchronization network.

The migration process is concurrent to server operations but disruptive to the z/OS images.

This scenario assumes that:

- ▶ The server is STP-capable.
- ▶ The HMC on the operations LAN has access to the server.
- ▶ The server is not yet time synchronized.

Figure A-14 on page 189 illustrates an example of the configuration.




Figure A-16 Installation of a new STP-only CTN: start scenario

When the initialization of the STP-only CTN is complete, the objective is to have CEC1 configured as an STP-only CTN. Since it is the only server in the CTN, CEC1 is assigned the roles of Preferred Time Server and Current Time Server.

Preparation

Prepare the hardware for migration by configuring your environment as follows:

- ▶ Ensure availability of an STP-capable HMC. The HMCs must be running HMC Application V2.9.2 or later. The HMC is used to establish the time zone, daylight saving time, and possibly leap second offsets. If using dial-out ETS, at least one HMC must be chosen that has the ability to dial out via an analog phone line.
- ▶ Install hardware MCLs to support STP on the STP-capable server, CEC1.
- ▶ If using an external time source, dial-out or NTP, access to the ETS must be configured and tested. Refer to 2.3, “External time source” on page 46.
- ▶ Assign timing roles. In this example, CEC1 is selected as the Preferred Time Server (stratum 1) and is assigned as the Current Time Server.

z/OS preparation

To perform z/OS preparation apply STP-related maintenance. This is necessary because the server is in STP timing mode and time synchronization is used. A re-IPL of the images is required after installation of the maintenance.

STP feature installation

This task requires the installation of the STP feature on the STP-capable server, CEC1.

Installation of the MCL is concurrent. After installation of the STP feature, the server can be configured for STP. STP panels are available on the SE and HMC. Various functions such as initialization of the CTN ID can now be performed. z/OS system images must be IPLed.

Activation

After the preparation and installation tasks are completed, the CTN can be defined and activated. When properly planned and executed, all operations are concurrent to server operations and nondisruptive to the z/OS images. The activation steps are:

1. Assign a CTN ID to each server.
2. Initialize the time, including offsets.
3. Assign the Current Time Server.

Assign a CTN ID

Because this is an STP-only CTN and there is no Sysplex Timer synchronization, there is no ETR network ID present. Only the STP network ID portion of the CTN ID is defined.

After the server has the CTN ID assigned, the System (Sysplex) Time task is available on the SE and HMC.

Initialize the time

Because there is no pre-existing ETR network, there is no inheritance of existing time parameters such as time, time zone offsets, or leap seconds. The information must be defined from the server that will be assigned the Current Time Server (CEC1 in this example). See the “Configuring an STP-only CTN” section in the *Server Time Protocol Implementation Guide*, SG24-7281, for details.

Specific server in an STP-only CTN

Since the STP-only CTN to be activated is a single-server CTN, the option should be exercised to limit the CTN to a one-server CTN. This is to avoid an automatic CTN deconfiguration during power-on reset or power off/on. This can be achieved by selecting the **Only allow the server(s) specified above to be in the CTN** option. See 3.4.2, “Single-server or dual-server auto resume after power-on reset” on page 125, for details.

Assign the Current Time Server

Up to this point, the server is still unsynchronized (stratum 0) because no time source has been assigned yet.

In this example, CEC1 is to be the stratum 1 and is assigned the roles of the Preferred Time Server and the Current Time Server. The roles are assigned via the System (Sysplex) Time panels. The operation initializes an STP-only CTN.

Figure A-15 on page 189 shows the resulting STP-only CTN.




Figure A-17 Installation of a new STP-only CTN: end scenario

Enhancing a single-site ETR network to a Mixed CTN across two sites

This scenario provides an example of a migration from an ETR network to a Mixed CTN within two physical data centers. This scenario has the following assumptions:

- ▶ Two sites:
 - Site 1 looks similar to the example in , “Migration from an ETR network to a Mixed CTN” on page 172.
 - Site 2 has two servers, with no timing network and z/OS images at z/OS V1.11 levels.
- ▶ At least two of the servers in the ETR network that are being considered for the migration must be STP-capable.
- ▶ z/OS images that require time synchronization, such as those in a Parallel Sysplex, are at z/OS V1.11 levels.

The planning considerations associated with the migration are a combination of the considerations previously discussed in this chapter, with the added complexity of a two-site situation.

Figure A-18 shows an example of such a configuration. This example assumes that all servers at site 1 are already time synchronized using an Expanded Availability Sysplex Timer configuration, with ETR network ID=31. Logical partitions P1, P2, and P3 are three z/OS images participating in a Parallel Sysplex.

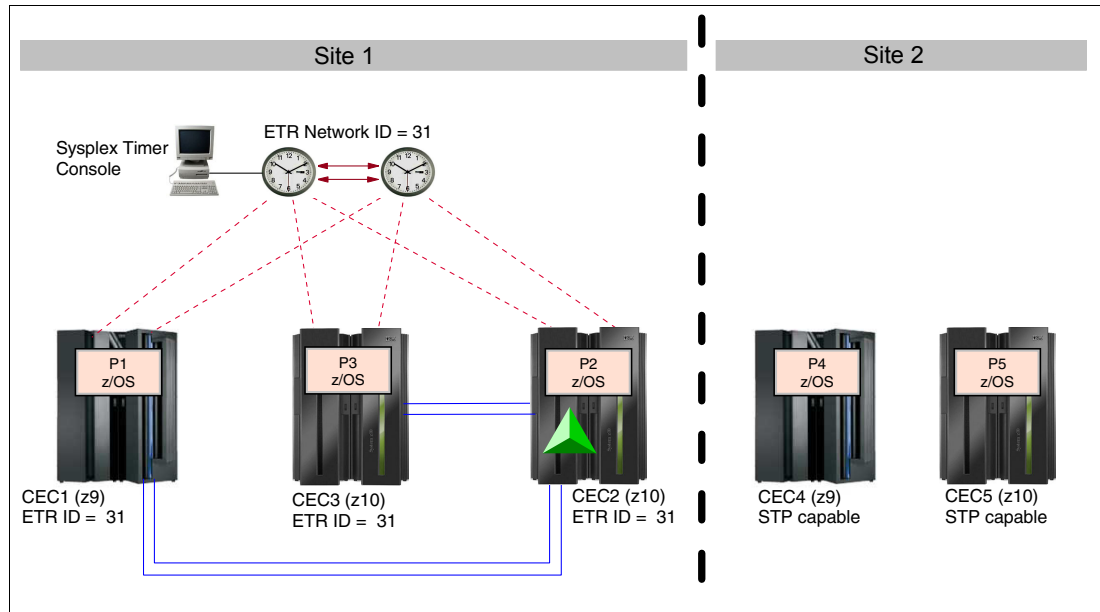


Figure A-18 Enhancing a single-site ETR network to a Mixed CTN across two sites: start scenario

The images are connected to the two CFs with sufficient links to avoid a single point of failure.

Site 2 servers are not attached to the Sysplex Timers in site 1. They have coupling link connections to two stratum 1 servers at site 1 to become part of the CTN.

The target configuration is a Mixed CTN. This also implies that the Sysplex Timer must remain as a part of the timing configuration.

When the migration to the Mixed CTN is complete, the objective is to have:

- ▶ At site 1, CEC1, CEC2, and CEC3 configured in a Mixed CTN. CEC3 is synchronized to the CEC1 or CEC2 using STP and is no longer directly synchronized to the Sysplex Timer.
- ▶ At site 2, CEC4 and CEC5 are configured for STP as members of the same Mixed CTN and are connected to site 1 using new Dense Wave Division Multiplexors (WDM) for coupling link and timing-only link connectivity.
- ▶ For redundancy, at least two STP-configured servers remain connected to the Sysplex Timer. They are CEC1 and CEC2 at site 1.
- ▶ z/OS images at site 2 are not part of the sysplex configured at site 1, although this would be possible.

Figure A-20 on page 198 illustrates the target configuration.

Preparing to migrate from an ETR network to a Mixed CTN

In this section we discuss preparing to migrate from an ETR network to a Mixed CTN.

Hardware preparation

Prepare your environment by doing the following:

- ▶ Ensure availability of an STP-capable HMC. You must plan to either upgrade the existing HMC or purchase a new HMC capable of running the HMC Application V2.10.0 to have visibility to *all* servers in the CTN. The Support Elements participating in the CTN configuration are CEC4 and CEC5 in site 2, running at console application V2.10.0.

Servers and coupling facilities at both sites should be able to be controlled by a Hardware Management Console. If there are HMCs at both sites, they must be able to access all servers in the Mixed CTN.

In a Mixed CTN, the Sysplex Timer consoles continues to be used for all timing-related functions, such as initializing time, scheduling daylight saving time, adjusting time, and so on. You must keep access to the Sysplex Timer consoles for as long as the Sysplex Timer remains operational in the timing configuration.

- ▶ Install hardware MCLs on the STP-capable servers. In the example, this would be CEC1, CEC2, and CEC3 at site 1, and CEC4 and CEC5 at site 2.
- ▶ Install LIC 3.0 on the Sysplex Timers.
- ▶ Ensure that your DWDM is certified for STP. See:
<http://www.ibm.com/systems/z/advantages/gdps/qualification.html>
- ▶ High-availability considerations: For high availability, the Mixed CTN should be configured so that every stratum 2 server is attached to at least two stratum 1 servers. You must decide which two servers at site 1 will continue to be synchronized to timing signals provided by the Sysplex Timers in the Mixed CTN.

At site 1, CEC1 and CEC2 continue to step to the Sysplex Timer. CEC3 and both servers at site 2 will be synchronized using STP messages.

We assume that the Mixed CTN configuration is a temporary step towards an STP-only CTN. If the plan is to stay in a Mixed CTN for a long time, it would be advisable to place one Sysplex Timer at site 1 and another at site 2 if fiber distance between sites is less than 40 km. This provides more resiliency and redundancy to the overall configuration.

- ▶ Have a plan for timing-only links that might be required. Timing-only links must be configured when a CF does not exist at either end of the coupling link. The CEC3 server at site 1 is the server that will synchronize to either CEC1 or CEC2. Coupling links already exist between CEC1 and CEC2 and can be used to synchronize CEC1 to CEC2 in a future configuration change.

However, there are no coupling links between the CEC1 and CEC3 because only CEC2 has a CF configured. Therefore, it is necessary to order and install coupling links that will be configured as timing-only links between CEC1 and CEC3. The timing-only links can be either the ICB-3, ICB-4, or ISC-3 (peer) links. Figure A-19 on page 197 illustrates the addition of these timing-only links for STP messaging.

CEC4 and CEC5 at site 2 would synchronize to either CEC1 or CEC2 at site 1 using STP timing messages. Coupling links to CEC1 at site 1 must be defined as timing-only links. Coupling links to the CEC2 in site 1 (with a CF) are defined as coupling links.

Figure A-19 illustrates the addition of the coupling links and timing-only links for STP messaging via the DWDMs.

z/OS preparation

Install z/OS maintenance related to STP. In the example discussed here, z/OS images P1, P2, P3, P4, and P5 require the software maintenance. Figure A-19 illustrates the resulting configuration.

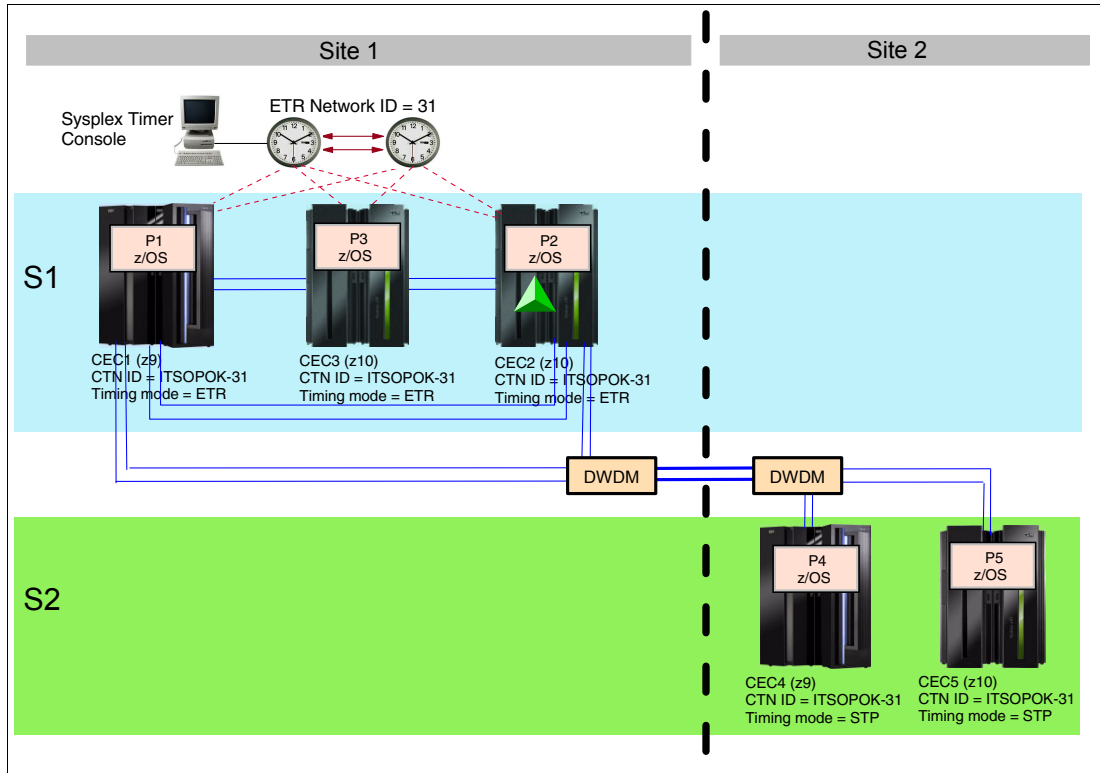


Figure A-19 Enhancing single-site ETR network to Mixed CTN across two sites: synchronized servers

Installation

In this section we discuss installation.

STP feature installation

This task requires the installation of the STP feature on the STP-capable servers. In this example, these are CEC1, CEC2, and CEC3 at site 1, and both servers at site 2.

After installation these servers can now be configured for STP.

- ▶ STP panels are available on the Support Element and HMC, and various functions, such as initialization of the CTN ID, can now be performed.
- ▶ Coupling facility partitions automatically recognize the new functions.

z/OS images with STP-related maintenance applied, however, require an IPL. In this example, logical partitions P1, P2, P3, P4, and P5 have been upgraded with STP maintenance. An IPL of the images is required to recognize the STP feature.

Activating the configuration for migration

After the preparation and installation tasks are completed, the CTN configuration and activation steps can now be initiated. When properly planned and executed, all operations are concurrent to server operations and nondisruptive to the z/OS images.

Assigning the CTN ID

The Hardware Management Console (HMC) is used to enter the CTN ID. For servers at site 1, the ETR network ID already assigned to the ETR network must *not* be changed. A new value must be entered for the STP network ID portion of the CTN ID. In this example, a CTN ID = [ITSOPOK] - [31] is assigned to CEC1, CEC2, and CEC3 at site 1.

The same CTN ID must be assigned to the servers at site 2. In this case, both the STP network ID and the ETR network ID must be specified to exactly match the CTN ID at site 1. Assuming that the coupling links are online, CEC4 and CEC5 at site 2 automatically become stratum 2 servers after the CTN ID [ITSOPOK] - [31] is specified. See Figure A-20.

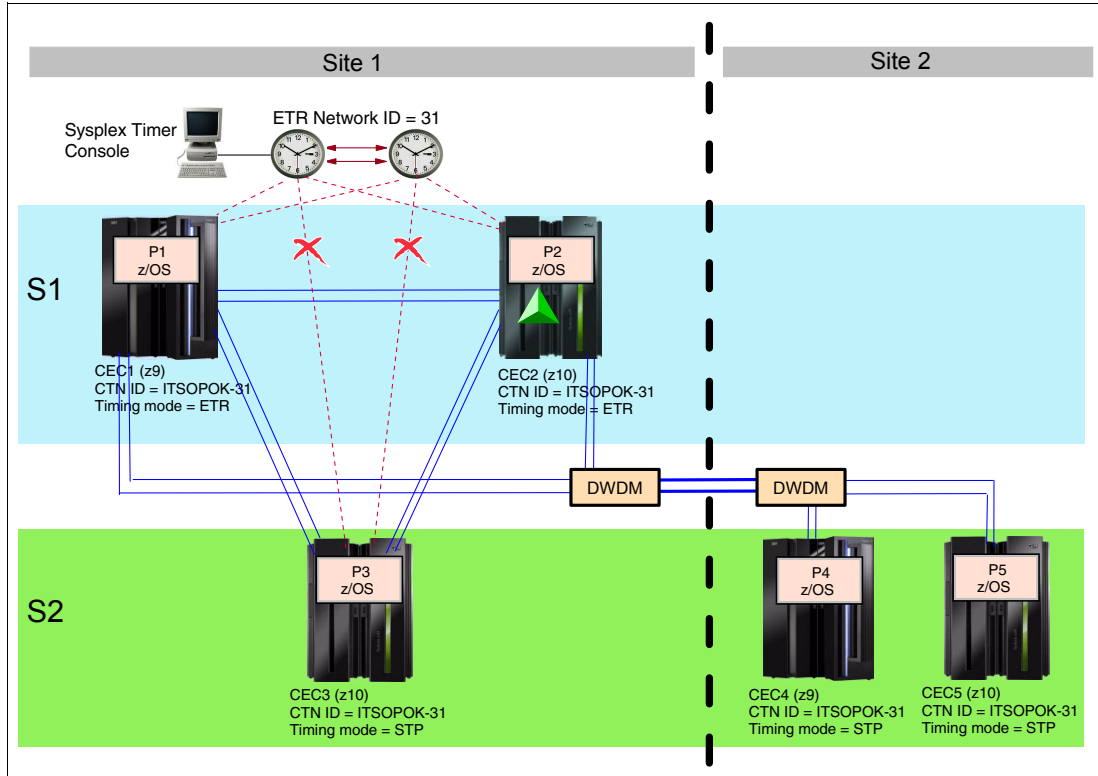


Figure A-20 Enhancing a single-site ETR network to a Mixed CTN across two sites: end scenario

Migrate a server to STP timing mode

The objective is to have CEC3 at site 1 and the two servers at site 2 synchronized to either the CEC1 or CEC2 using STP messages.

You must change the timing mode of CEC3 from ETR timing mode to STP timing mode. The change is made from the ETR Configuration panel of the CEC3 by disabling its ETR ports.

Note: The STP messaging capability and links redundancy must first be verified from STP STATUS or CEC3 will go to stratum 0.

In this two-site scenario, one Sysplex Timer configuration has been considered. If each site had its own Sysplex Timer configuration, it would have made a nondisruptive migration more challenging, if not impossible. A method has been developed to merge two Sysplex Timer networks, but this procedure usually requires a sysplex outage. Therefore, the scenario that is provided seems to be a more likely approach.



Server Time Protocol (STP) messages in z/OS

This appendix provides a summary of the z/OS STP-related messages. For detailed message information, refer to the following manuals:

- ▶ *z/OS - MVS System Messages Volume 10 (IXC - IZP)*, SA22-7640-24
- ▶ *z/OS - MVS System Messages Volume 6 (GOS - IEA)*, SA22-7636-23

Supervisor messages

IEA031I STP ALERT RECEIVED. STP ALERT CODE = nn

IEA265I UNABLE TO SYNC TOD CLOCKS TO timesource. CPC CONTINUES IN LOCAL MODE.

IEA279I ALL CLOCK RELATED SET COMMANDS ARE IGNORED WHEN IN xxx MODE.

IEA380I THIS SYSTEM IS NOW OPERATING IN STP TIMING MODE.

IEA381I THE STP FACILITY IS NOT USABLE. SYSTEM CONTINUES IN LOCAL TIMING MODE

IEA382I THIS SERVER HAS ONLY A SINGLE LINK AVAILABLE FOR TIMING PURPOSES.

IEA383I THIS SERVER RECEIVES TIMING SIGNALS FROM ONLY ONE OTHER NETWORK NODE.

IEA384I SETETR COMMAND IS NOT VALID IN STP TIMING MODE.

IEA385I CLOCKxx ETRDELTA & TIMEDELTA BOTH SPECIFIED. yyyyyyy IGNORED.

IEA386I hh.mm.ss TIMING STATUS
 SYNCHRONIZATION MODE=mode
 THIS SERVER IS A STRATUM n
 CTNID=ctnid
 THE STRATUM 1 NODE ID=cpcid
 NUMBER OF TIMING LINKS=ll
 optional text

IEA387I STP DATA CANNOT BE ACCESSED. SYSTEM CONTINUES IN yyyyyy TIMING MODE.

IEA388I THIS SERVER HAS NO CONNECTION TO THE nnnnnnnnnn

IEA389I THIS STP NETWORK HAS NO SERVER TO ACT AS nnnnnnnnnn

IEA390I TOD CLOCKS DYNAMICALLY ADJUSTED TO MAINTAIN STP SYNCHRONISM.

IEA391W TWO STP TIMING NETWORKS WITH THE SAME CTN ID HAVE BEEN DETECTED

IEA392I STP TIME OFFSET CHANGES HAVE OCCURRED.

IEA393I ETR PORT n IS NOT OPERATIONAL. THIS MAY BE A CTN CONFIGURATION CHANGE.

IEA394A THIS SERVER HAS LOST CONNECTION TO ITS SOURCE OF TIME. IF THIS EVENT OCCURRED ON SOME, BUT NOT ALL NETWORK SERVERS, THE LIKELY CAUSE IS A LINK FAILURE. TO FIX, ENSURE THAT EACH AFFECTED SERVER HAS AT LEAST ONE CORRECTLY CONNECTED AND FUNCTIONAL LINK. IF THIS EVENT OCCURRED ON ALL NETWORK SERVERS, THEN THE LIKELY CAUSE IA A TIMING NETWORK FAILURE. TO FIX, REFER TO THE MESSAGE IEA394A DESCRIPTION IN MVS SYSTEM MESSAGES. AFTER FIXING THE PROBLEM, REPLY "RETRY" FROM THE SERVICE CONSOLE (HMC). IF THE PROBLEM WAS NOT CORRECTED, THIS MESSAGE is REISSUED AND YOU MAY TRY AGAIN. REPLY "ABORT" TO EXIT THE MESSAGE LOOP. PROABLE RESULT: 0A2-158 WAITSTATE.

IEA395I THE CURRENT TIME SERVER HAS CHANGED TO THE cccccccccc

IEA503W MULTIPLE CPUS FAILED DURING xxx SYNC-CHECK PROCESSING.

XCF messages

IXC410E SYSTEM sysname LOST SYNCHRONIZATION WITH THE TIMING NETWORK. LOCAL TIMING MODE WILL BE ALLOWED FOR 80 MINUTES.

IXC434I sysname HAS TIMING DEFINITIONS THAT ARE NOT CONSISTENT WITH THE OTHER ACTIVE SYSTEMS IN SYSPLEX sysplex - EFFECTIVE CLOCK VALUES ARE NOT CONSISTENT
text

IXC435I ALL SYSTEMS IN SYSPLEX sysplex ARE NOW SYNCHRONIZED TO THE SAME TIME REFERENCE.
text

IXC436W THIS SYSTEM HAS LOST TIME SYNCHRONIZATION WITH THE OTHER SYSTEMS IN THE SYSPLEX AND HAS BEEN PLACED INTO A NON-RESTARTABLE WAIT STATE CODE: X'0A2'
REASON CODE: X'15A'

IXC437I SYSTEMS CAN NOW ENTER THE SYSPLEX USING COORDINATED SERVER TIME
CTNID=stpid[-etrid]

IXC438I COORDINATED TIMING INFORMATION HAS BEEN UPDATED FOR SYSTEM sysname
PREVIOUS timing
CURRENT timing

IXC439E ALL SYSTEMS IN THE SYSPLEX sysplex ARE NOT SYNCHRONIZED TO THE SAME TIME REFERENCE. THE FOLLOWING SYSTEM IS NOT PART OF THE TIMING NETWORK AND WILL BE REMOVED.
text

IXC468W XCF IS UNABLE TO ACCESS THE CTN AND HAS PLACED THIS SYSTEM INTO A NON-RESTARTABLE WAIT STATE CODE: X'0A2' REASON CODE: X'158'

XES messages

IXL160E CF REQUEST TIME ORDERING: REQUIRED AND NOT-ENABLED text reason

IXL161I CF REQUEST TIME ORDERING: text1 text2

IXL162E CF REQUEST TIME ORDERING: REQUIRED AND WILL NOT BE ENABLED text reason

Automating-related messages

IEA031I STP ALERT RECEIVED. STP ALERT CODE = nn

IEA265I UNABLE TO SYNC TOD CLOCKS TO timesource. CPC CONTINUES IN LOCAL MODE.

IEA381I THE STP FACILITY IS NOT USABLE. SYSTEM CONTINUES IN LOCAL TIMING MODE

IEA382I THIS SERVER HAS ONLY A SINGLE LINK AVAILABLE FOR TIMING PURPOSES.

IEA383I THIS SERVER RECEIVES TIMING SIGNALS FROM ONLY ONE OTHER NETWORK NODE.

IEA387I STP DATA CANNOT BE ACCESSED. SYSTEM CONTINUES IN yyyy TIMING MODE.

IEA388I THIS SERVER HAS NO CONNECTION TO THE nnnnnnnnn

IEA389I THIS STP NETWORK HAS NO SERVER TO ACT AS nnnnnnnnnn

IEA394A THIS SERVER HAS LOST CONNECTION TO ITS SOURCE OF TIME. IF THIS EVENT OCCURRED ON SOME, BUT NOT ALL NETWORK SERVERS, THE LIKELY CAUSE IS A LINK FAILURE. TO FIX, ENSURE THAT EACH AFFECTED SERVER HAS AT LEAST ONE CORRECTLY CONNECTED AND FUNCTIONAL LINK. IF THIS EVENT OCCURRED ON ALL NETWORK SERVERS, THEN THE LIKELY CAUSE IS A TIMING NETWORK FAILURE. TO FIX, REFER TO THE MESSAGE IEA394A DESCRIPTION IN MVS SYSTEM MESSAGES. AFTER FIXING THE PROBLEM, REPLY "RETRY" FROM THE SERVICE CONSOLE (HMC). IF THE PROBLEM WAS NOT CORRECTED, THIS MESSAGE IS REISSUED AND YOU MAY TRY AGAIN. REPLY "ABORT" TO EXIT THE MESSAGE LOOP. PROBABLE RESULT: 0A2-158 WAITSTATE.

IEA503W MULTIPLE CPUS FAILED DURING xxx SYNC-CHECK PROCESSING.

IXC410E SYSTEM sysname LOST SYNCHRONIZATION WITH THE TIMING NETWORK. LOCAL TIMING MODE WILL BE ALLOWED FOR 80 MINUTES.

IXC434I sysname HAS TIMING DEFINITIONS THAT ARE NOT CONSISTENT WITH THE OTHER ACTIVE SYSTEMS IN SYSPLEX sysplex - EFFECTIVE CLOCK VALUES ARE NOT CONSISTENT
text

IXC436W THIS SYSTEM HAS LOST TIME SYNCHRONIZATION WITH THE OTHER SYSTEMS IN THE SYSPLEX AND HAS BEEN PLACED INTO A NON-RESTARTABLE WAIT STATE CODE: X'0A2'
REASON CODE: X'15A'

IXC439E ALL SYSTEMS IN THE SYSPLEX sysplex ARE NOT SYNCHRONIZED TO THE SAME TIME REFERENCE. THE FOLLOWING SYSTEM IS NOT PART OF THE TIMING NETWORK AND WILL BE REMOVED.
text

IXC468W XCF IS UNABLE TO ACCESS THE CTN AND HAS PLACED THIS SYSTEM INTO A NON-RESTARTABLE WAIT STATE CODE: X'0A2' REASON CODE: X'158'

IXL160E CF REQUEST TIME ORDERING: REQUIRED AND NOT-ENABLED text reason

IXL162E CF REQUEST TIME ORDERING: REQUIRED AND WILL NOT BE ENABLED text reason

Related publications

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

IBM Redbooks publications

For information about ordering these publications, see “How to get Redbooks publications” on page 205. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *IBM zEnterprise EC12 Technical Introduction*, SG24-8050
- ▶ *IBM zEnterprise EC12 Technical Guide*, SG24-8049
- ▶ *IBM zEnterprise EC12 Configuration Setup*, SG24-8034
- ▶ *IBM zEnterprise 196 Technical Guide*, SG24-7833
- ▶ *IBM zEnterprise System Technical Introduction*, SG24-7832
- ▶ *IBM zEnterprise 196 Configuration Setup*, SG24-7834
- ▶ *IBM zEnterprise 114 Technical Guide*, SG24-7954
- ▶ *IBM System z10 Enterprise Class Technical Introduction*, SG24-7515
- ▶ *IBM System z10 Enterprise Class Technical Guide*, SG24-7516
- ▶ *IBM System z10 Business Class Technical Overview*, SG24-7632
- ▶ *IBM System z9 Enterprise Class Technical Guide*, SG24-7124
- ▶ *IBM System z9 Business Class Technical Introduction*, SG24-7241
- ▶ *Getting Started with InfiniBand on System z10 and System z9*, SG24-7539
- ▶ *IBM System z Connectivity Handbook*, SG24-5444
- ▶ *Server Time Protocol Implementation Guide*, SG24-7281
- ▶ *Server Time Protocol Recovery Guide*, SG25-7380
- ▶ *IBM System z10 Enterprise Class Configuration Setup*, SG24-7571
- ▶ *S/390 Time Management and IBM 9037 Sysplex Timer*, SG24-2070

Other publications

These publications are also relevant as further information sources:

- ▶ *z/Architecture Principles of Operation*, SA22-7832
- ▶ *zEnterprise EC12 Installation Manual for Physical Planning*, GC28-6914
- ▶ *zEnterprise 196 Installation Manual for Physical Planning*, GC28-6897
- ▶ *zEnterprise 114 Installation Manual for Physical Planning*, GC28-6907
- ▶ *IBM System z10 Business Class Installation Manual for Physical Planning*, GC28-6875
- ▶ *IBM System z10 Enterprise Class Installation Manual for Physical Planning*, GC28-6864

- ▶ *IBM System z10 Enterprise Class Processor Resource/Systems Manager Planning Guide*, SB10-7153
- ▶ *IBM System z10 Enterprise Class System Overview*, SA22-1084
- ▶ *IBM System z10 Business Class System Overview*, SA22-1085
- ▶ *IBM System z10 Enterprise Class Service Guide*, GC28-6866
- ▶ *System z9 Enterprise Class System Overview*, SA22-6833
- ▶ *System z9 Business Class System Overview*, SA22-1083
- ▶ *System z9 EC Installation Manual for Physical Planning*, GC28-6844
- ▶ *System z9 BC Installation Manual for Physical Planning*, GC28-6855
- ▶ *System z9 Processor Resource/Systems Manager Planning Guide*, SB10-7041
- ▶ *System z9 Support Element Operations Guide Version 2.9.1*, SC28-6858
- ▶ *zSeries 990 Installation Manual for Physical Planning*, GC28-6824
- ▶ *zSeries 890 Installation Manual for Physical Planning*, GC28-6828
- ▶ *zSeries 890 and 990 Processor Resource/Systems Manager Planning Guide*, SB10-7036
- ▶ *z890 and z990 Support Element Operations Guide Version 1.8.2*, SC28-6831
- ▶ *Hardware Management Console Operations Guide Version 2.12.0*, SC28-6919
- ▶ *zEnterprise System Support Element Operations Guide Version 2.12.0*, SC28-6920
- ▶ *System z CHPID Mapping Tool User's Guide*, GC28-6900
- ▶ *Hardware Management Console Operations Guide Version V2.10.1*, SC28-6873
- ▶ *Hardware Management Console Operations Guide Version 2.10.0*, SC28-6867
- ▶ *Support Element Operations Guide V2.10.1*, SC28-6879
- ▶ *Support Element Operations Guide V2.10.0*, SC28-6868
- ▶ *Planning for the 9037 Model 2*, SA22-7233
- ▶ *Model 5900 and 4900 ESCON Server Installation and User's Guide*, GA22-1082
- ▶ *z/OS MVS Setting Up a Sysplex*, SA22-7625
- ▶ *IOCP User's Guide*, SB10-7037
- ▶ *Stand-Alone Input/Output Configuration Program User's Guide*, SB10-7152
- ▶ *Planning for Fiber Optic Links*, GA23-0367
- ▶ *Common Information Model (CIM) Management Interfaces*, SB10-7154

Online resources

These Web sites are also relevant as further information sources:

- ▶ IBM Resource Link
<http://www.ibm.com/servers/resourceLink/>
- ▶ Server Time protocol Web page
<http://www-03.ibm.com/servers/eserver/zseries/pso/stp.html>
- ▶ GDPS
<http://www.ibm.com/systems/z/gdps/>

How to get Redbooks publications

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

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Index

A

accuracy 48
active Stratum 1 13–14, 171
ACTS 25, 47
adjust
 Daylight Saving offset 103
 Leap Seconds 111
 Time 36
 time 33, 61, 106
 Time Zone 113, 124
Adjustment steering 108
Arbiter 14, 19, 22, 27, 87, 129, 131, 133, 166
Arbiter connectivity 131
Arbiter-assisted recovery 156
Automated Computer Time Service 25
Automatically adjust, radio button 114

B

Backup Time Server 14, 19, 27, 87, 130, 132, 166
Base steering 107–108
BIPM 2
bit position 51 3
BTS 14
Bureau International des Poids et Mesures 2

C

care 100
CF 8
CF Channel Path Connectivity List 45
CF message 22, 45, 134
CF Request Time Ordering 89, 133
CFCC
 DISPLAY TIMEZONE 118
 TIMEZONE 118
CFCC level 32
CFLEVEL 32
CHPID 135–136, 168
 access list 136
CICS 122
CICS Complex (CICSplex) 121
Clock Time Assignment area 119
CLOCKxx 59, 109
 ETRDELTA 61
 ETRMODE 60, 116, 159
 ETRZONE 63
 OPERATOR 60
 SIMETRID 65
 STPMODE 63, 110, 116, 120, 159
 STPZONE 65, 101, 116
 TIMEDELTA 59, 64
 TIMEZONE 66
command
 SET TIMEZONE 73

concurrent migration 14
Configuring
 a Mixed CTN 92
 an STP-only CTN 94
Consistency Group 24
Console assisted recovery 158
Control messages 23
Coordinated Server Time 9, 27, 108, 184
Coordinated Timing Network 11, 28, 81
Coordinated Universal Time 2
Coupling Facility 8, 116
Coupling Facility Control Code 78
coupling link 8, 134
coupling link connectivity 17
CPC 78
CST 47
CTD 47
CTN 11, 164
CTN configuration 125
CTN ID 11–12, 16, 81, 86, 92, 125
CTN power outage 145
CTN topology 85
CTS 14
Current Time Server 14, 20, 28, 87, 109, 127, 143, 171
Current Time Zone 102
Customize Network Settings 97
CVTTZ 117

D

Daylight Saving Time 8, 16, 33, 114, 123
Daylight Saving Time offset 109
Define Time Zone Algorithm 103, 115
Dial out 25, 47
dial-out schedule 36
Directly Attached Systems 84
Display ETR
 ETR Network 86
 Mixed CTN 87
 STP-only CTN 87–88, 128, 132
Domain Name Services 97
Driver level 37
DST 114
DST offset 16

E

EC Driver level 164
encryption 48
Enhanced Book Availability 42
Enhanced PSP Tool 166
Epoch 3
ETR
 Configuration tab 175, 180, 184
 Network ID 12, 86, 89
 port 86, 105

- Status tab 83
- timing mode 9, 86, 109
- ETR network 4, 6–7, 184
- ETR Network ID 105
- ETR network ID 11
- ETR port 93
- ETR timing mode 17
- ETRDELTA 63
- ETRMODE 59, 109
- ETRZONE 59, 109
- ETS 16, 46, 94, 102, 107, 165
- ETS planning 71
- Expanded Availability 7
- extended TOD clock 3
- External Time Reference 4, 81
- External Time Source 102, 106, 160, 187
- External Time Source (ETS) 25

F

- FC 0217 41
- FC 0218 41
- FC 0219 41
- FC 0227 41
- FC 0228 41
- FC 0229 41
- FC 0230 41
- FC 0231 41
- FC 0232 41
- FC 0993 41
- FC 1021 32, 164
- FC 2400 42
- FC 3393 41
- figetnet 5
- fine-steering rate 9
- firewall 52, 54
- Force configuration checkbox 105, 130

G

- GDPS 7
- GDPS/XRC 24
- Global Positioning Service (GPS) 26
- Global Positioning System 26
- Greenwich Mean Time 2
- Gross-steering rate 9

H

- Hardware Management Console 33, 47, 78
- HCD 44
 - STP control unit 25
 - Timing-only link 44
- HCM 44
- high availability 168
- HMC 165
 - Dial out 95
 - Domain Security 35
 - ETR Configuration tab 82
 - ETR Status tab 83
 - ETS Configuration tab 85

- local operation 34
- modem 47
- Network Configuration 100, 104, 125, 132
- Network Configuration tab 82
- remote 34
- remote operation 34
- security 35
- STP Configuration tab 84
- STP Status tab 84
- Timing network tab 81
- HMC message
 - ACTZ01C7 144

I

- IBF 57
- IBM 9037 4
- ICB
 - cable 41
- ICB links 41
- ICB-3 8, 17
- ICB-4 8, 17
- IEASYSxx
 - PLEXCFG 159
- IEN Telephone Date Code 47
- InfiniBand 38
- Initialize Time button 100
- Initialize Time panel 100
- Integrated Cluster Bus 22, 38
- Inter System Channel 38
- Internal Battery Feature 57
- International Atomic Time 2
- InterSystem Channel-3 22
- IOCP 44
- IODF 134
- IPL 58–59, 109, 165
- ISC Daughter Card 41
- ISC Mother Card 41
- ISC-3 8
 - cable 42
 - feature 41
- ISC-3 Peer 17

L

- last path 136
- last path validation 139
- Last timing link 136
- LC Duplex connector 42
- Leap second 2, 99–100, 110
 - offset 16, 110
- leap second 101, 110
- Leap second offset 99, 108
- LIC 10
- Licensed Internal Code (LIC) 8
- Linux for System z9 75
- Linux on System z 75
- local time 81, 115
- Local timing mode 9, 28, 159
- Logical partition time offset 118

M

- mega-microsecond 3
- message
 - IEA015A 159
 - IEA279I 116
 - IEA283I 89
 - IEA384I 89
 - IEA389I 132–133
 - IEA391W 200
 - IEA598I 66
 - IEA886A 116
 - IEE136I 116
 - IEE148I 137–138, 141
- Message Time Ordering Facility 89
- Mixed CTN 11, 16, 33, 86, 92, 149, 178
- MTOF 89
- multiple time zone 120
- multiple time zones 106
- Multi-site Sysplex 46

N

- Network Configuration tab 125, 184, 187
- Network Time Protocol 7, 164
- NIST Automated Computer Time Service 47
- NRC Canadian Time Service 47
- NTP 7, 164
 - NTP clien 48
 - NTP client support 25
 - NTP Public Service Project 27
 - NTP server 25, 50, 52, 97
 - continuous availability 53
 - NTP server redundancy 51
 - NTP server selection 48
 - NTP stratum 26

O

- offset 179
- offset adjustment 106

P

- Parallel Sysplex 3–4, 120
- PCHID 136
- Peer mode 14, 25
- periodic dial-out 33, 47
- PLEXCFG 159
- plexcfg 159
 - any 159
 - monoplex 159
 - multisystem 159
 - xcflcal 159
- POR 126
- Power on reset 125
- PPS 55
- PPS port 98
- Preferred Time Server 14, 20, 87, 129–130, 133, 166
- Processor Resource/Systems Manager 8
- PSIFB 17, 25
- PSIFB cables 40

- PSP bucket 166
- PTS 14
- pulse per second 49, 54, 98
- pulse per second (PPS) 7

R

- radio button, Automatically adjust 114
- Redbooks Web site 205
 - Contact us xi
- RFC-1305 48
- RFC-4330 48
- RPQ 8P2197 41

S

- Scheduled Clock Adjustment 114
- scheduled operation 47
- Server Offline Signal 155
- server role 14, 129, 132
- Server Time Protocol 6
- Set
 - Daylight saving time 103
 - standard time 103
- Set date and time 104
- SIMETRID 61, 65, 69
- Simple Network Time Protocol 25
- Single Mode 42
- single point of failure 15, 129–130, 171
- single server CTN 70
- Source field 97
- standalone Coupling Facility 46
- star pattern 12
- STCK 3
- STCKE 3
- STCKF 3
- steering 107, 184
- STP 8, 13
 - installation 173, 187, 197
 - Network ID 87
 - Timing mode 9, 16
- STP clock source date
 - Not usable 10
 - Usable 10
- STP Configuration tab 125, 177
- STP control unit 44–45
- STP feature 37
- STP message 8, 13, 16, 44, 136, 168
- STP network ID 11–12, 29
- STP offset adjustments 108
- STP Status tab 84, 135–136
- STP stratum 12–13
- STP timing 174
- STP timing mode 17, 159
- STP-capable 10, 32, 164
- STP-capable server 29
- STP-capable servers 14
- STP-clock-source state 10
- STP-configured 11, 13
- STP-configured server 15, 29
- STP-enabled 11, 15, 92

- STP-enabled server 15, 18, 29, 58
- STPMODE 59
- STP-only CTN 11, 87, 104, 111, 149, 178
- STP-timing mode 9
- STPZONE 59
- stratum 12, 26
- stratum 0 26
- Stratum level 85
- Stratum levels 13
- Support Element 25, 35–36, 48, 138, 158
- Sysplex Timer 4, 13, 32, 154, 165
 - receiving unit 4
 - sending unit 4
- Sysplex Timer connection 49
- System (Sysplex) Timer task
 - Network Configuration 147
- System Assist Processor (SAP) 155
- System Data Mover 23
- System z10 7
- System z9 7
- System/370 architecture 3

T

- Time button 106
- TIME macro 109
- Time panel 84
- Time parameter messages 23
- time scale 2
- Time Zone 103, 113, 116, 118
 - Algorithm 103, 115
 - Offset 33, 112, 114
 - Offset adjustment 112
- Time Zone offset 113
- Time zone offset 99, 109
- TIMDELTA 59, 64
- Time-of-Day 3
- Time-of-Day clock 3
- time-zone offset 16
- Timing mode 9
- timing mode 86
- Timing Network tab 112
- Timing state 10
 - Stopped 10
 - synchronized 10
 - unsynchronized 10
- timing-only link 25, 45, 69, 134, 169, 196
- timing-only links 44
- TOD 3, 7, 12
- TOD clock 4, 69, 87, 106, 154
 - steering 9
- TOD-Clock Synchronization Facility 4
- TOD-clock-steering 8
- TOD-offset register 9

U

- UDP port 123 52, 54
- uninitialized STP Links 85
- Uninitialized timing mode 9
- Universal Time 1 2

- unsynchronized 10, 178, 189
- USA Energy Policy Act 114
- user-defined Time Zone offset 115
- UT1 2, 100
- UTC 2, 36, 110
- UTC month 2

V

- VM guest 69

W

- WDM 41, 46, 195, 197

Z

- z/OS 57, 78
 - commands 86, 116
 - DISPLAY CF 89
 - DISPLAY ETR 86–87
 - DISPLAY XCF 89
 - maintenance 57
 - SET CLOCK 117
 - SET DATE 116
 - SET RESET 117
 - SET TIMEZONE 117
- z/OS command
 - DISPLAY ETR 133
- z/OS Console 72
- z/OS Global Mirror 23
- z/OS Messages 91
- z/VSE 75
- z10 BC 8
- z10 EC 8
- z890 8, 26
- z9 BC 8
- z9 EC 8
- z990 8, 26



Server Time Protocol Planning Guide



Redbooks®

Concepts and terminology

Configuration planning

Operations and recovery

Server Time Protocol (STP) is a server-wide facility that is implemented in the Licensed Internal Code (LIC) of the System zEnterprise (zEC12, z196, z114) Servers, System z10 Enterprise Class (z10 EC), System z10 Business Class (z10 BC), IBM System z9 Enterprise Class (z9 EC), System z9 Business Class (z9 BC), and zSeries z990 and z890 servers. It provides improved time synchronization in a sysplex or non-sysplex configuration.

This IBM Redbooks publication is intended for infrastructure architects and system programmers who must understand IBM STP functions. Readers are expected to be generally familiar with System z technology and terminology.

This book provides planning information for Server Time Protocol functions and associated software support. For more detailed installation, operation, and recovery information, refer to the companion books *Server Time Protocol Implementation Guide*, SG24-7281, and *Server Time Protocol Recovery Guide*, SG24-7380.

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

SG24-7280-03

ISBN 0738438103