Redbooks
ibm.com/redbooks

IBM Z Time Synchronization Implementation Guide

Bill White

Donald Crabtree

Gatto Gobehi

Steve Guendert

Clay Kaiser

Octavian Lascu

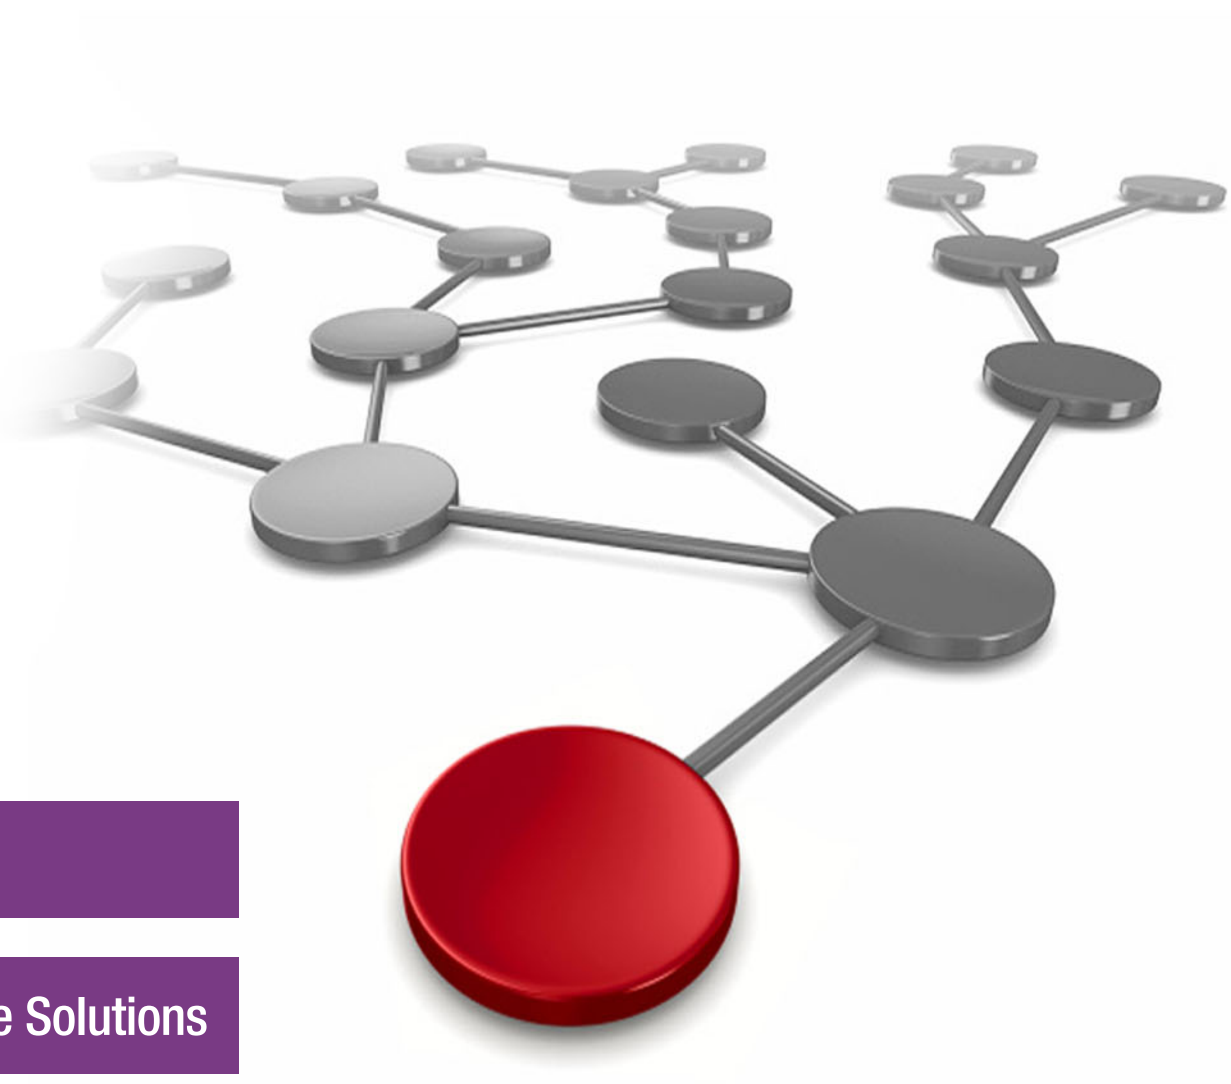

IBM Z

Infrastructure Solutions

IBM

Redbooks

IBM Redbooks

IBM Z Time Synchronization Implementation Guide

September 2025

SG24-8480-02

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

Third Edition (September 2025)

This edition applies to IBM Server Time Protocol (STP) for IBM Z platform and covers IBM z17, IBM z16, and IBM z15 server generations.

Contents

Notices v
Trademarks vi

Preface vii
Authors vii
Now you can become a published author, too! viii
Comments welcome viii
Stay connected to IBM Redbooks ix

Chapter 1. The critical role of time synchronization in IT systems 1
1.1 Understanding time synchronization and timestamps 2
1.1.1 Why accurate time matters: Compliance, security, and beyond 2
1.1.2 Clockscales, timescales, and synchronization methods 3
1.1.3 Reliable time sources and managing leap seconds 4
1.1.4 Protocols and practices for achieving common time across networks 5
1.2 Ensuring precise time synchronization on IBM Z 6
1.2.1 Using external time sources with NTP and PTP on IBM Z 6
1.2.2 IBM Z Server Time Protocol 6
1.2.3 Time synchronization: Operating system support on IBM Z 8

Chapter 2. Accurate timekeeping with IBM Z 9
2.1 IBM Z time synchronization 10
2.2 STP planning considerations 11
2.2.1 Coupling link considerations 11
2.2.2 CTN roles 12
2.3 ETS planning considerations 18
2.3.1 Network Time Protocol 18
2.3.2 Precision Time Protocol 23
2.3.3 Pulse Per Second 25
2.3.4 Network Time Security 26
2.3.5 Comparison of external time sources 26
2.3.6 ETS Network physical planning 29
2.3.7 ETS steering 31
2.4 Operating system considerations 32
2.4.1 z/OS 33
2.4.2 z/VM 35
2.5 Management and operations 38

Chapter 3. CTN configuration and operations 41
3.1 Setup new CTN 42
3.1.1 Identify active STP links 50
3.2 Adding systems to CTN 51
3.3 Modify assigned server roles 54
3.4 Remove systems from CTN 57
3.5 CTN membership restriction 61
3.6 Export CTN configuration data to file 64
3.7 Deconfiguring the CTN 64

Chapter 4. External Time Source Configuration and Management 67

4.1 Configure ETS connectivity . . . 68
4.2 ETS Network Diagnostics . . . 71
4.3 Configure ETS . . . 74
4.3.1 ETS for IBM z17 . . . 75
4.3.2 ETS for IBM z16 . . . 83
4.3.3 ETS for IBM z15 . . . 91
4.4 ETS status . . . 92
4.4.1 ETS monitoring commands . . . 92
4.4.2 Audit log . . . 99
4.5 Control Pulse Per Second signal . . . 100
4.5.1 Best practices for implementing PPS connectivity . . . 100
4.5.2 Testing and verifying PPS connectivity . . . 101
4.6 Manage CTN Certificates . . . 105

Chapter 5. Advanced CTN configurations and time management operations . . . 111
5.1 Advanced CTN Actions . . . 112
5.1.1 Join existing CTN . . . 112
5.1.2 Split to new CTN . . . 115
5.1.3 Save STP debug data . . . 119
5.1.4 Set time server power failover . . . 120
5.2 Managing the Coordinated Server Time (CST) . . . 123
5.2.1 STP time adjustment . . . 123
5.2.2 STP offset adjustments . . . 125
5.2.3 Changes in local time . . . 130
5.2.4 Logical partition time offset . . . 133
5.2.5 Parallel Sysplex and multiple time zones . . . 135

Abbreviations and acronyms . . . 139

Related publications . . . 141
IBM Redbooks . . . 141
Other publications . . . 141
Online resources . . . 141
Help from IBM . . . 142

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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 grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

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 jurisdictions 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

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

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 actual people or business enterprises 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

CICS®
CICSPlex®
FICON®
GDPS®
IBM®
IBM Z®
IBM z16®
Parallel Sysplex®
Redbooks®
Redbooks (logo) ®
Resource Link®
z/OS®
z/VM®
z15®
z16®
z17™

The following terms are trademarks of other companies:

The registered trademark Linux® is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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

Preface

The IBM® Z® time synchronization solution uses two technologies to ensure precise, consistent, and resilient timekeeping across IBM Z systems: server time protocol (STP) and external time sources (ETS). STP is a server-wide facility that is implemented in the licensed internal code (LIC) of the IBM Z platform. It provides improved time synchronization in a sysplex or non-sysplex configuration. ETS on IBM Z supports standard protocols like Network Time Protocol (NTP), Precision Time Protocol (PTP), and Pulse Per Second (PPS) to keep accurate universal time.

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

This book provides planning and implementation information for STP and ETS functions and associated software support for the IBM z17™, IBM z16®, and IBM z15® platforms.

Authors

This book was produced by working at IBM Redbooks, Poughkeepsie Center.

Bill White is an IBM Redbooks Project Leader and Senior IT Infrastructure Specialist at IBM Poughkeepsie, New York.

Donald Crabtree is a Programmer Extraordinaire for IBM Z Time Synchronization in Poughkeepsie NY. He holds a masters degree in Computer Science from the State University of NY at Albany. He has over 35 years of experience with IBM including over 20 years of working with STP.

Gatto Gobehi is an IBM Z Product Engineer working at IBM Poughkeepsie. He has over 20 years of mainframe engineering and testing experience. He specializes in IBM Z hardware (all areas) and microcode, and supports IBM Service teams and customers. He is a customer advocate for several high-profile IBM Z clients.

Steve Guendert is the chief architect for IBM Z and LinuxONE time synchronization technology and leads a team of developers focused on time synchronization and IBM STP. He represents IBM in several time synchronization standards development organizations and projects including the ITU, IEEE, and IETF. He has 48 patents in time synchronization and IBM Z I/O technology.

Clay Kaiser is a Firmware Developer for IBM Z Time Synchronization in Poughkeepsie, NY, US. He holds a degree in Electrical Engineering from the University of Illinois Urbana-Champaign. His area of focus includes security and resiliency of time synchronization.

Octavian Lascu is a Senior IT Infrastructure Architect with Inter Computer, Romania. He has over 35 years of IT experience with IBM Z, IBM Power, and IBM Storage. Octavian was an IBM Redbooks Project Leader for over 20 years, where he has co-authored many IBM Redbooks publications covering IBM Z hardware, IBM Power, and various IBM solutions..

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

Bryce Goldman, Nikolas Myer
IBM

Ewerson Palacio, Bob Haimowitz
IBM Redbooks

Thanks to the authors of the previous editions of this book:
Hans-Peter Eckam, Gatto Gobehi, Steve Guendert, Clay Kaiser, Walter Keller, Jeremy Koch, Octavian Lascu, Franco Pinto, Martin Söllig, Sebastian Zimmermann

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 IBM Redbooks 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 email to:

  redbooks@us.ibm.com

- Mail your comments to:

  IBM Corporation, IBM Redbooks
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on LinkedIn:

  https://www.linkedin.com/groups/2130806

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

  https://www.redbooks.ibm.com/subscribe

- Stay current on recent Redbooks publications with RSS Feeds:

  https://www.redbooks.ibm.com/rss.html

1

The critical role of time synchronization in IT systems

Time synchronization is critical in IT for managing the correct order of events in applications such as transaction processing, message logging, and serialization. Accurate time synchronization ensures data integrity, supports regulatory compliance, and enhances the performance of systems that require precise timing.

Without time synchronization, many systems might experience failures, and time-sensitive applications might not function properly, resulting in inconsistent and unreliable data.

When data is shared across multiple devices or systems, you must synchronize each device's clock to ensure all timestamps are aligned. Data records might drift, making it difficult to determine the correct order of events. Unsynchronized clocks also make it challenging to analyze logs, identify security incidents, and trace back the cause of problems when timestamps are inconsistent. With some systems, even slight variations in time can cause significant problems. Transactions must be recorded with precise timestamps to ensure accurate accounting and prevent discrepancies.

In summary, time synchronization is vital for maintaining security, ensuring compliance, optimizing system performance, facilitating network operations, preserving data integrity, and enabling global coordination.

This chapter includes the following topics:

- 1.1, "Understanding time synchronization and timestamps" on page 2
- 1.2, "Ensuring precise time synchronization on IBM Z" on page 6

1.1 Understanding time synchronization and timestamps

Time synchronization is the process of ensuring a device's or system's clock is aligned, consistent, and accurate. The following steps describe the typical process:

1. A device or system needing time synchronization sends a request to a reference clock for the current time.
2. The reference clock responds with a timestamp, which is used to calculate the difference between the device's or system's clock facility and the reference time.
3. When the device or system receives the response, its clock facility is adjusted to align with the reference time and includes any network latency or clock drift when performing the adjustment.

You use a timestamp to track the sequence in which applications process the data is being processed by applications. Examples include when data is created, modified, accessed, and deleted. When applications process data, they use timestamps provided by the device or system clock facility. The clock facility requires a time source when the device or system is initialized for processing data.

Also, when data is exchanged between different devices or systems, timestamps must be synchronized between them. To maintain accurate time for data processing across devices or systems, use a universal time reference source.

1.1.1 Why accurate time matters: Compliance, security, and beyond

Timestamps and regulatory compliance are intertwined. Timestamps provide a way to confirm the precise time and sequence of events, ensuring the integrity of digital records and preventing disputes. The integrity of digital records is also important in industries that are subject to regulatory oversight, where accurate record-keeping is required.

In the financial sector, for example, synchronized timestamps are essential for audit trails and legal purposes, ensuring that transactions are accurately recorded and traceable. Since the inception of the direct electronic trading of financial instruments, the speed of financial market transactions has increased at an exponential rate as CPU processing and network performance have advanced. Computer systems automatically make trades, both buying from and selling to each other based on software algorithms. The automated trades have increased in the number of exchanges competing against each other, and the amount of time required to run transactions and the *spread size* have decreased. The rapid expansion of computer-based trading with its highly sophisticated algorithms has increased the need for synchronization of trading systems and traceability to a common reference timescale to help prevent irregularities and aid forensic analysis.

Institutions that interact with financial exchanges require a high accuracy time for several reasons, which include the following examples:

- Accurate time is necessary for institutions to control their own trading traffic.
- Accurate time is also required to settle disagreements and prevent fraud.

The timestamps that are collected by each institution must be examined when there are disagreements or errors in processing trades.

Where transactions occur across multiple entities and are carried out over networks, time synchronization accuracy for business clocks must be in compliance with regulatory mandates that are established by authorized organizations. Some examples include FINRA Regulatory Notice 16-23 and Official Journal of the European Union, L 173, 12 June 2014.

The following list outlines additional reasons for synchronizing time across devices and systems:

- Security systems

  Accurate time synchronization is essential for security systems such as intrusion detection and prevention. It helps you identify and correlate security events across different systems.

- Certificate validation

  Time synchronization is tightly linked to certificate-based security protocols. Certificates used for identification and authentication rely on accurate time to determine their validity and revocation status. If a clock is significantly off, your system might misinterpret a valid certificate as expired or revoked, which can lead to security breaches.

- System performance and stability

  Time synchronization helps you maintain the performance and stability of systems and applications. It prevents time-related errors, data inconsistencies, and synchronization issues between systems.

- Network operations

  Accurate time synchronization is crucial for network operations such as routing protocols, load balancing, and monitoring tools. It ensures that network devices communicate effectively and efficiently.

- Data integrity

  Time synchronization helps you maintain data integrity by ensuring that timestamps in databases and other storage systems are consistent and accurate. This is essential for data analysis, reporting, and forensic investigations.

- Global coordination

  In a global environment, you can use time synchronization to coordinate activities and events across different time zones. It ensures that all systems and devices align with a common reference time, supporting seamless communication and collaboration.

- Emerging technologies

  Emerging technologies such as autonomous vehicles and quantum computing also require highly accurate and precise time synchronization.

1.1.2 Clockscales, timescales, and synchronization methods

Clockscales, timescales, and time synchronization help standardize time to provide consistent and accurate timekeeping.

Clockscales refer to the different time units (seconds, minutes, hours) and their relationships. A clockscale is the actual measurement of time by a clock. A timescale is a defined system for representing time, often based on the readings of multiple clocks. Essentially, a clockscale provides the raw measurements of time. A timescale offers a standardized and often more precise representation of that time.

Time synchronization is the process of ensuring multiple device or system clocks agree on a coordinated common time.

Different systems with distinct purposes are available for timekeeping, including:

- Time of Day (TOD): A clockscale, displaying time within a 24-hour period, is often used in digital clocks, timers, and other time-related applications. If a TOD system clock is not synchronized with an external time source, drift can happen, gradually deviating from the actual time and can cause issues with synchronization between different systems or components. Also, log timestamps are inaccurate, making it difficult to determine when events occurred.

- Temps Atomique International (TAI): An internationally coordinated atomic timescale based on the weighted average of many atomic clocks (over 400). TAI is a continuous, uninterrupted timescale with no leap seconds.
- Coordinated Universal Time (UTC): The primary time standard used globally as an external time source. UTC is used to calculate time zones and other time references. It is based on TAI with leap seconds incorporated (added or skipped as needed) to account for variations in the Earth's rotation. UTC also follows Universal Time 1 (UT1). UT1 is computed by 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. Although UT1 is 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, which makes it more consistent with civil, or solar, time.

To implement time synchronization, you need a reliable time source and a mechanism to distribute that time in the network to establish common time with participating systems.

1.1.3 Reliable time sources and managing leap seconds

When atomic time was adopted, some communities of users, especially of celestial navigation, requested that atomic time be synchronized with the rotation of the Earth.

To compensate for the Earth's irregular velocity of rotation, the International Telecommunication Union (ITU) defined a procedure for adding (or suppressing) a second as necessary to ensure that the difference between the international time reference and rotational time is minimal.

Leap seconds are used to keep the difference between UT1 and UTC to within ±0.9 s. The resulting timescale is UTC, the atomic timescale that is maintained at the Bureau International des Poids et Mesures (BIPM) with the contribution of 69 national institutes that operate about 400 atomic clocks.

For more information about leap seconds, see National Institute of Standards and Technology (NIST).

The International Earth Rotation and Reference Systems Service (IERS) is responsible for monitoring the Earth's rotation. It announces the dates of application of any leap seconds that are required, which usually are timed for the end of 30 June or 31 December. For more information, see IERS Bulletin C.

The leap second adjustment process is shown in Figure 1-1 on page 5.

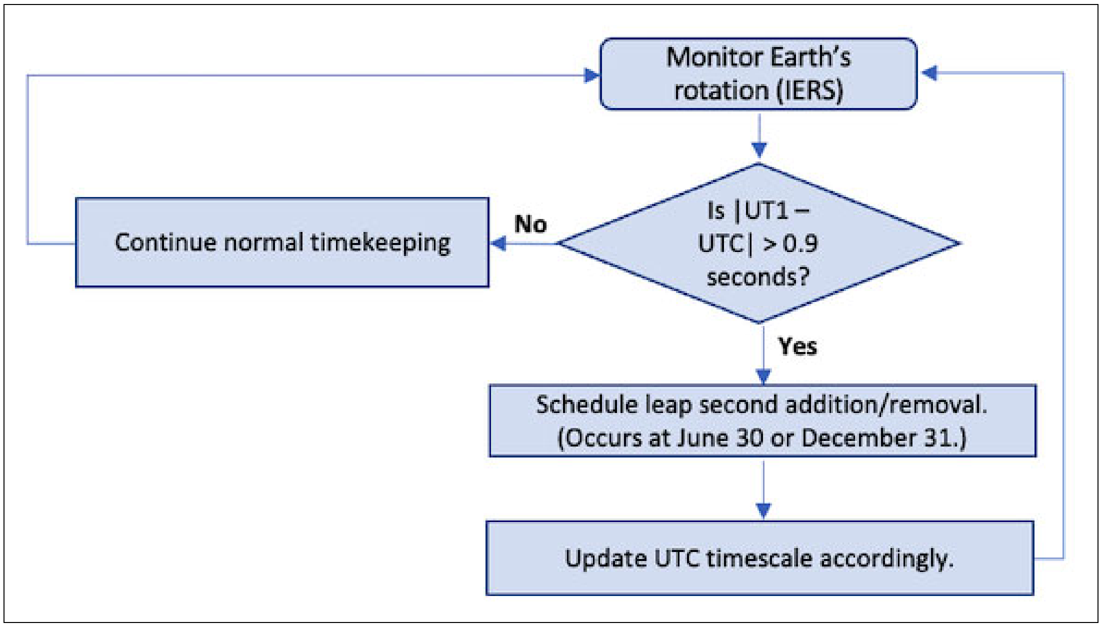

Figure 1-1 Process for leap second adjustments

Note: A leap second introduces an irregularity into the UTC timescale, so exact interval measurements are not possible by 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.

For timestamps and clock synchronization, the global reference is UTC. For more information, see BIPM technical services: Time Metrology.

1.1.4 Protocols and practices for achieving common time across networks

Various protocols and methods are used to synchronize clocks:

- Network Time Protocol (NTP), which is a widely used protocol that synchronizes clocks over networks, often by exchanging timestamps between clients and servers.
- Precision Time Protocol (PTP), which is a more precise protocol designed for high-accuracy synchronization in networks, often used in industrial settings.

Demonstrate time traceability to a common time reference by verifying the following requirements:

- Calibrate the timing equipment must be so that its unknown internal delays do not bias its time output.
- Continuously monitor the equipment so that any fault or anomaly can be detected, and the time output is not used until the equipment is working correctly again. The calibration evidence and monitoring results should be archived so that the status of the timing equipment at any point in time can be verified later.

Author Lee Segall once said, *"It is possible to own too much. A man with one watch knows what time it is, a man with two watches is never quite sure"*. Timestamps that are created by different time sources can be meaningfully compared only if they are based on the same time reference.

1.2 Ensuring precise time synchronization on IBM Z

To ensure the performance and stability of systems and applications across IBM Z platforms, time synchronization plays a crucial role. It helps prevent time-related errors, data inconsistencies, and synchronization problems between systems.

By maintaining accurate timekeeping, seamless operations and data integrity across the IBM Z infrastructure is ensured. This is particularly important for applications that rely on precise time tracking, such as transactions, logging, and monitoring tools.

Implementing robust time synchronization practices with IBM Z Server Time Protocol (STP) and an external time source can significantly enhance the overall performance and reliability of the IBM Z platform. STP synchronizes with the global timescale UTC through the use of NTP and TAI through the use of PTP.

1.2.1 Using external time sources with NTP and PTP on IBM Z

For enhanced time accuracy compared to alternative methods, you can use either NTP or PTP to communicate with an External Time Source (ETS) to meet the need for a time reference across IBM Z platforms.

NTP client support is available in the firmware partition that is running on IBM Z. The code interfaces with the NTP servers so that an NTP server can be the single time source for IBM Z and for other servers that have NTP clients.

With PTP, IBM Z acts as a TimeReceiver, getting its time from a TimeTransmitter (the source of accurate time) through messages. IBM Z firmware then uses these messages and other internal computations to adjust its own clock to achieve synchronization with the TimeTransmitter. This process enables the distribution of precise time across multiple IBM Z platforms and IBM Parallel Sysplex® environments, by using STP.

Although STP, NTP, and PTP are used for time synchronization in networked systems, they serve different purposes and operate at different levels.

NTP is a client-server-based protocol that is used to synchronize the computer clocks over a network. It is designed to synchronize time across a network to a precision within a few milliseconds of UTC over the internet. NTP is typically used for time synchronization in IP networks, including local area networks (LANs), wide area networks (WANs), and the Internet.

PTP is a network-based protocol that synchronizes computer clocks in networks to a high degree of accuracy, but unlike NTP it is in the nanosecond or even picosecond range. PTP is based on TAI, which is not subject to leap seconds, as UTC is. PTP communicates the offset between TAI and UTC so that UTC can be calculated from the PTP time.

PTP relies on hardware timestamping, which provides a more precise and accurate measurement of time compared to software-based timestamping used in NTP.

However, STP is a protocol that is used for time synchronization within a single system or for a small group of systems within a data center or a cluster of servers.

1.2.2 IBM Z Server Time Protocol

Although STP is an IBM Z protocol and technology, it is based on standards. STP is a message-based time synchronization protocol similar to NTP and PTP. It is designed to maintain time synchronization for multiple IBM Z systems to form a Coordinated Timing Network (CTN). A CTN is a collection of systems that are time synchronized to a value called Coordinated Server Time (CST), creating a single view of time with an external time reference.

Timekeeping information is transmitted over specialized physical connections called coupling links. STP maintains all of the systems in the CTN so that their timestamps stay within less than 10 microseconds of each other.

STP is implemented in firmware as a system-wide facility of IBM Z and is enabled by installing the STP feature. Extra configuration is required for an IBM Z to become a member of a CTN.

Figure 1-2 depicts a high-level overview of the IBM Z time synchronization infrastructure using STP with NTP or PPS. For greater time accuracy, the external time reference (ETS) has a direct network connection to the IBM z17 or IBM z16 CPC (Central Processing Complex) drawer. The CPC is an IBM Z system with CPU and I/O. With the IBM z15, the connection to the ETS is established through the IBM Z System Element (SE).

Configuration and management of the time synchronization infrastructure is through the IBM Z Hardware Management Console (HMC).

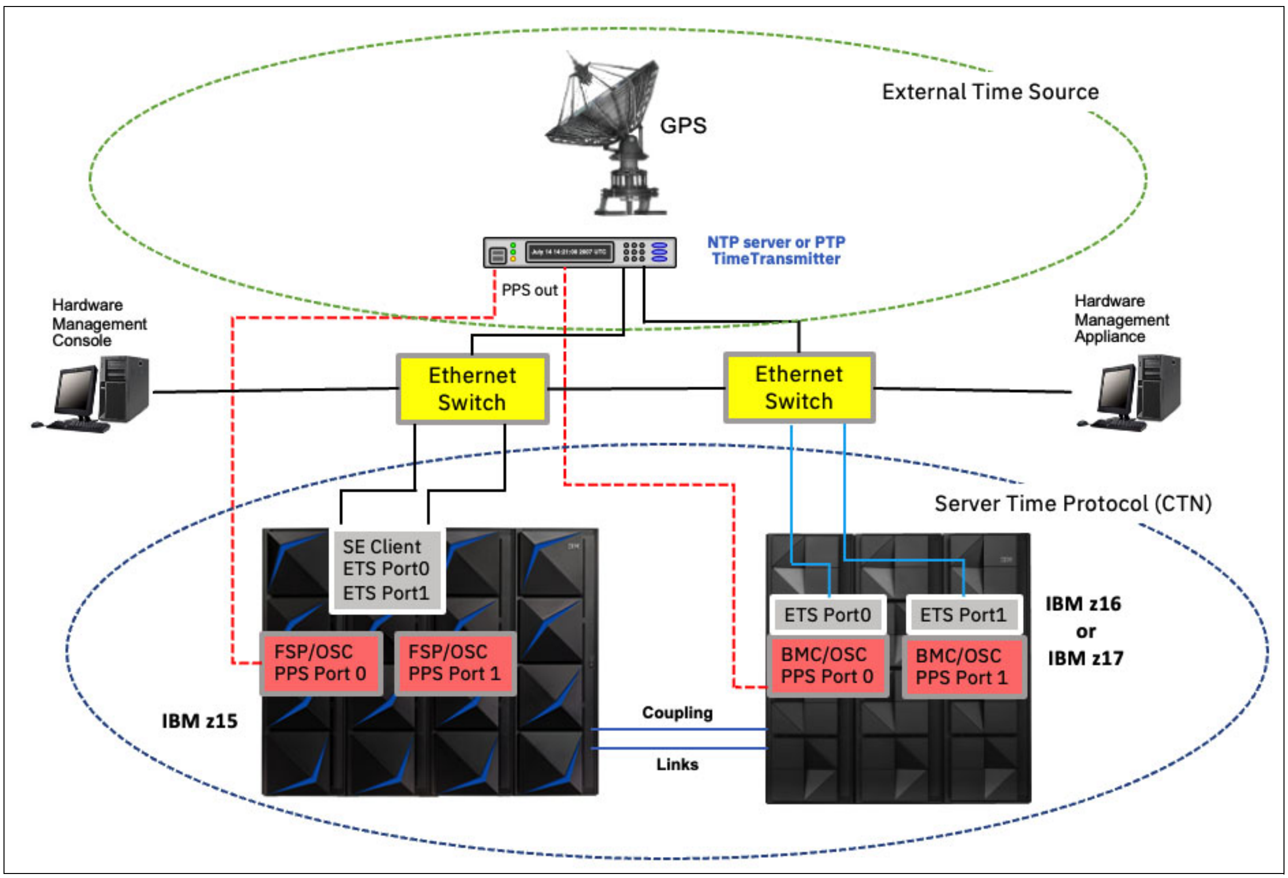

Figure 1-2 IBM Z time synchronization infrastructure

Two Oscillator Card (OSC) ports in each IBM Z provide a dual-path interface for the Pulse per Second (PPS) signal. STP tracks the highly stable and accurate PPS signal from the external time source. PPS maintains an accuracy of 10 µs as measured at the PPS input of the IBM Z. The PPS port must be connected by cable to the PPS output of an NTP server when IBM Z is configured for NTP with PPS as ETS for time synchronization.

PPS is optional for PTP but might be required for NTP to meet financial regulations regarding the accuracy of time synchronization between servers.

In addition, STP uses a hierarchical structure with Stratum 1 servers being the most accurate (directly connected to an external time source). Lower stratum servers synchronize to higher stratum servers, creating a network of time sources. STP coordinated server time ensures all systems within its network are synchronized to the same time.

STP provides the following capabilities:

- Coordination of IBM Z timestamps from external time source (or reference) in a logical construction called Coordinated Timing Network (CTN)
- Time synchronization with other IBM Z systems by using coupling links
- Time synchronization handling of events in a CTN during external time source availability issues
- Reconfiguration of the CTN in case of system issues like IBM Z power loss and various failures such as disaster recovery (DR) situations

1.2.3 Time synchronization: Operating system support on IBM Z

The IBM Z supported operating systems (IBM z/OS®, IBM z/VM®, IBM z/TPF, and Linux on IBM Z) use the time of day (TOD) clock to synchronize their time. The TOD clock is automatically set during IBM Z system initialization and synchronized with an ETS or through STP.

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

Changes to the time zone, handling of Daylight Saving Time, and time adjustments can also be performed by using operating system commands.

For more information about the components that make up the IBM Z time synchronization solution and implementation guidance, see Chapter 2, “Accurate timekeeping with IBM Z” on page 9.

2

Accurate timekeeping with IBM Z

Server Time Protocol (STP) and an External Time Source (ETS) work together to establish and maintain accurate time synchronization across IBM Z systems.

This chapter describes maintaining accurate time on the IBM Z platform. keeps precise time and provides methods to adjust or steer that time. It also covers important points about using STP and how to plan for a Coordinated Timing Network (CTN).

The chapter also describes External Time Source (ETS) options and addresses operating system (z/OS, z/VM, z/TPF, and Linux on IBM Z) integration with STP. It outlines best practices for building accurate and resilient environments with guidance on maintaining synchronized time across complex IBM Z infrastructures.

The chapter includes the following topics:

- 2.1, “IBM Z time synchronization” on page 10
- 2.2, “STP planning considerations” on page 11
- 2.3, “ETS planning considerations” on page 18
- 2.4, “Operating system considerations” on page 32
- 2.5, “Management and operations” on page 38

2.1 IBM Z time synchronization

You can use STP and ETS within a Coordinated Timing Network (CTN) to achieve precise time synchronization across IBM Z systems. These features are Licensed Internal Code (LIC) included with IBM Z systems under feature code 1021.

A CTN is a collection of connected IBM Z systems that share a common time base. STP synchronizes time within the CTN, and ETS synchronizes the CTN to an external reference.

The primary goal of STP is to maintain time coherency across IBM Z systems. By using highly efficient coupling links to accurately distribute time across the CTN, each system in the CTN reports the same clock value. STP is especially important when multiple IBM Z systems work with shared data.

ETS ensures that the clock value on any CTN member matches a global time standard, such as Coordinated Universal Time (UTC), by using standards-based synchronization protocols, including Network Time Protocol (NTP), Precision Time Protocol (PTP), and Pulse Per Second (PPS). Use ETS to match time reporting with a universal standard or with systems outside the CTN.

Figure 2-1 is an illustration of the components that make up the IBM Z time synchronization solution and the redundant physical connectivity to ensure resiliency.

.

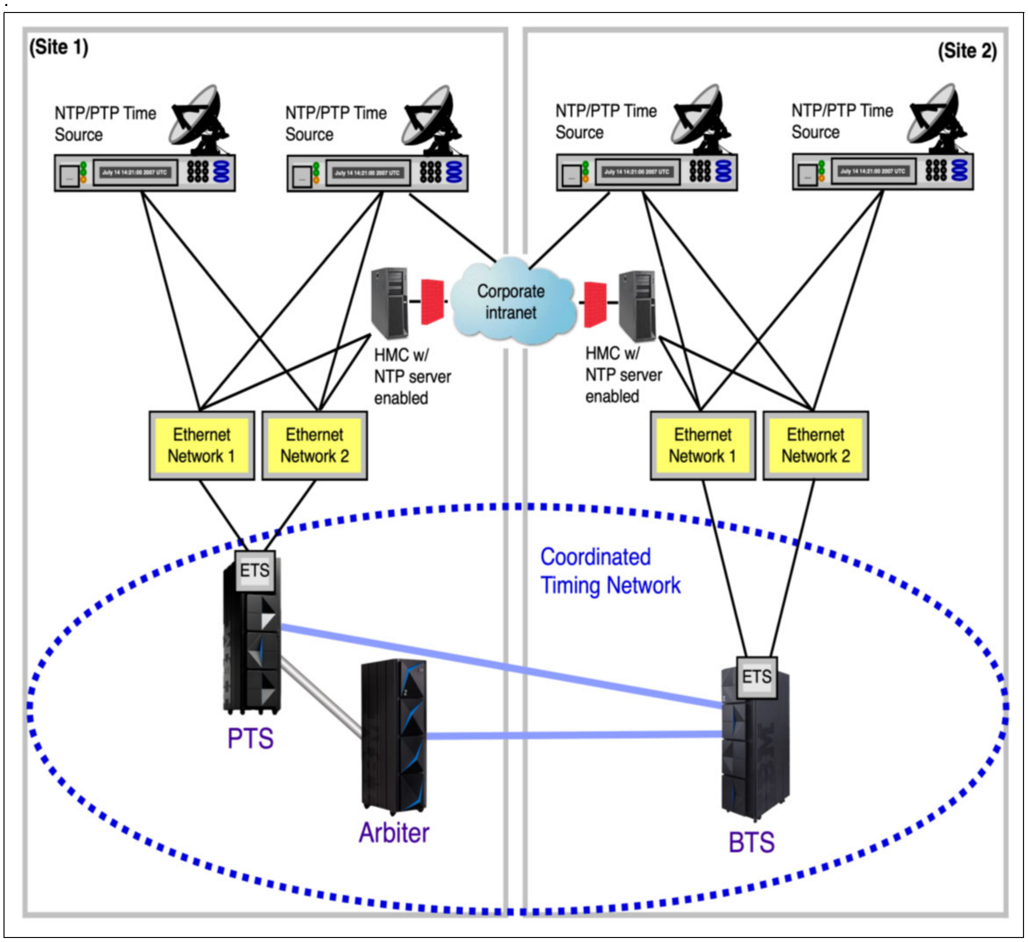

Figure 2-1 IBM Z time synchronization: Components and connectivity

At a logical level, the CTN includes the roles Current Time Server (CTS), Preferred Time Server (PTS), Backup Time Server (BTS), and Arbiter for keeping time and supporting the recovery process after an outage of one of the systems in the CTN. Also, see 2.2.2, "CTN roles" on page 12.

The IBM Z Hardware Management Console (HMC) plays a critical role with STP CTNs and an ETS. You can use the HMC to manually set the date and time, or you can initialize CST to an External Time Source (ETS). The STP CTN uses the settings from the HMC for the time zone, Daylight Saving Time (DST), and leap second offset. It also adjusts the time when needed. For more details, see 2.5, "Management and operations" on page 38.

2.2 STP planning considerations

STP is a message-based protocol in which timekeeping information is passed between IBM Z system members of a CTN through externally defined coupling links. A CTN consists of a hierarchical order of individual IBM Z systems that all share a common time. The most important aspect of STP over other timing protocols like NTP and PTP is the assurance of time consistency and continuity across the entire CTN. Therefore, if STP is synchronized, then distributed applications such as IBM Parallel Sysplex[1] always sees time move forward, regardless of which physical system is used. When an ETS is defined to the CTN, it is also possible that the entire CTN stays synchronized with the external reference, which is typically UTC through NTP or PTP.

2.2.1 Coupling link considerations

The foundation of STP is to distribute clock information via the highly specialized coupling links that is used to connect multiple IBM Z systems. Coupling links used by STP fall into two categories:

1. Long-range coupling links are used to connect to systems that are located a maximum of 10 km apart or up to 100 km using DWDM[2] technology.
2. Short range coupling links are used to connect to systems a maximum of 150 m apart (usually within the same data center). These links provide the best synchronization, so use them whenever possible.

STP communicates with other systems in the CTN via coupling links to keep the time of each system in sync with the rest of the CTN. STP does not use internal coupling links. Similar to coexistence rules for Parallel Sysplex, the N-2 generation coexistence rule applies to any CTN: Only servers that are not earlier than two generations compared to the newest servers in the CTN are supported in the same CTN.

A special type of coupling link that is used solely for STP traffic is a *timing-only link*. The intent of the offering of a timing-only link is to allow a customer to create a CTN where Parallel Sysplex and non-Parallel Sysplex systems can share the same time source. The CHPIDs of a timing-only link must belong to different IBM Z systems.

Since the initial implementation of STP, there is an increase of available external coupling links to exchange STP messages to keep servers and CFs synchronized. The N-2 coexistence rule is applicable to coupling link features and its firmware requirement.

[1] A cluster of IBM Z platforms acting together as a single system image with z/OS.
[2] DWDM - Dense Wavelength Division Multiplexing

STP identifies and maintains a list of all data links on the server that are capable of STP message communication during Power on Reset (POR). After you create the list of coupling channels that can transport STP messages, a quality of path algorithm is used to choose the STP active path to each externally connected system. This algorithm relies on the fastest link. This typically can be impacted by factors such as physical distance, channel throughput, latency through the adapter and supporting link/path logistic (point to point, or DWDM). When planning for coupling link connectivity, it is recommended to limit cable length as much as is feasible as excessive cable lengths can result in reduced accuracy. At the top of the hierarchal order of STP roles is the Primary Time Server (PTS) and Backup Time Server (BTS). Together with the Arbiter, these systems work together to resiliently control the time of the CTN. For more information, see 2.2.2, “CTN roles” on page 12.

When planning for a data center, it is recommended that all systems in a CTN have coupling link connectivity to both the PTS and BTS. It is recommended that all systems also have connectivity to the Arbiter when three or more IBM Z systems are present in the CTN in case the Arbiter must be reassigned as the PTS or BTS.

To maintain CTN connectivity and remove single points of failure, implement at least two separate coupling link connections between systems. When possible, implement the connections in different CPC drawers for Integrated Coupling Adapter Short Reach (ICA SR) or in different I/O Domains of the I/O drawer for Coupling Express LR (CE LR) features to maintain CTN connectivity if one of the coupling links fails.

When using DWDM to transport CE LR connections, it is recommended to use only IBM Z Parallel Sysplex qualified DWDM devices. For any DWDM device that is used for STP message transport, do not use transparent path changes because it can impact STP link calculations when different path lengths are used. For more information, see IBM zSystems Qualified Wavelength Disvision Multiplexer (WDM) products for Multi-site Sysplex and GDPS solutions.

For more information about coupling connectivity, see *IBM Z Connectivity Handbook*, SG24-5444.

2.2.2 CTN roles

Each CTN has a hierarchical order of IBM Z systems that adhere to a Current Time Server (CTS). The CTS is the ultimate source of time for the entire CTN, and is responsible for steering its clock, as well as the clocks of all member systems, to an ETS. To provide CTN resiliency a PTS and BTS are provided, each capable of operating as the CTS. Also, In addition, Depending on the number of systems in the CTN, an Arbiter is also available to help the PTS and BTS determine which source to use if the PTS and BTS cannot connect.

Current Time Server

The Current Time server (CTS) is mostly a flag that indicates whether the PTS or BTS is serving as the main time server for the CTN. The CTS is the system to which most commands are directed, including commands to steer out the offset to an external time source (ETS).

The term *non-CTS* is associated with PTS or BTS when it is not playing the role of CTS. Mostly, it is just a member of the CTN but remains ready to take the role of CTS in case of failure. The non-CTS forwards the steering commands to the CTS, which performs the actual steering required before distributing the time in the CTN.

Preferred Time Server (PTS)

During normal operations of the CTN, define the Preferred Time Server (PTS) server as the CTS. When it operates as the CTS, it is the main source of time for the entire CTN. Ensure that this server can connect to all other servers in the CTN and to all other role systems, such as BTS and Arbiter. In addition, on z16 and later systems, this server can connect to ETS when external synchronization is required.

Backup Time Server

When not operating as the CTS, the Backup Time Server (BTS) behaves the same as any other member of the CTN. However, when the PTS fails, the BTS is designed to take over as the CTS and then becomes the primary source of time for the CTN. For these reasons, the same recommendations provided for the PTS are relevant. Particularly, it is recommended that this server has connectivity to all other servers in the CTN and required that it has connectivity to all other role systems, PTS, and Arbiter. Also, on IBM z16 and later systems, this server must have ETS connectivity when external synchronization is required.

With the IBM z17, IBM z16 at D51 bundle s32, and IBM z15 at D41 bundle s88, steering to the external time source is received and processed from both the PTS and BTS. Before the release of those packages, steering commands were performed only when they were received by the CTS. However, if no commands are sent to the CTS within 2 hours, then it starts following commands sent to the non-CTS.

Because the CTS performs all steering to the external time source, the non-CTS forwards ETS steering commands that it receives to the CTS. When an ETS steering command is received by the CTS, it compares the dispersion of the new command to the sum of dispersion of the old plus the additional dispersion accumulated by the system since the old command was received. If the new command has a dispersion that is lower, then the command is performed, and the indicated offset is steered out. Otherwise, the command is ignored. This allows the IBM Z to ride over glitches that are occasionally seen in the external timing network.

Arbiter

The Arbiter (ARB) is another member of the CTN. In certain recovery situations, it helps determine which system is the CTS. Therfore, it is required that this server connects to both the PTS and BTS. Also, ensure that the ARB follows all coupling and ETS connectivity recommendations for the PTS and BTS so that the source can be redefined as either the PTS or BTS in the event of a planned maintenance or service window.

STP single server CTN

An STP single server CTN has only the PTS defined. If the PTS fails, there is no recovery and the entire CTN fails. The system status changes to stratum-0 (clock Uninitialized state). Therefore, use a single-system CTN only when the CTN contains only one IBM Z system.

The only reason to have more than one IBM Z system in an STP single system configuration is if the other system is frequently removed from the CTN, for example, to be rebooted or to test certain scenarios.

> **Note:** It is recommended to restrict STP single system CTNs to preserve the CTN network after Power on Reset. When a single system CTN is configured as restricted, no changes can be made to the membership of the CTN.

STP dual system CTN

An STP dual system CTN is a CTN that has both the PTS and BTS defined but no Arbiter. Most of the time, the PTS plays the role of CTS. If the PTS fails, the BTS takes the role of CTS and keeps the CTN running. Use a dual-system CTN when only two IBM Z systems are in the CTN. If there are more than two systems in the CTN, an STP Triple System configuration is recommended.

During an unplanned outage, the going away signal, n-mode power, or Console-Assisted Recovery (CAR) is used to determine that the BTS should assume the role of CTS.

Ensure that the PTS and BTS have redundant power sources so that a power failure for one does not impact the other.

When configuring a CTN in two data center sites, set up the PTS in one site and the BTS in the other so that a CTS remains if one site is lost. In this case, each site must have a managing HMC defined.

A server holding a role in a CTN cannot accidentally be removed from the CTN in a planned disruptive action unless its role has been moved to another server. If a planned disruptive action is attempted on the PTS or BTS, STP blocks that action until the role is reassigned to another server in CTN or the role has been removed from the CTN definition.

Note: It is recommended to restrict STP dual-system CTNs to preserve the CTN network after Power on Reset. When a Dual System CTN is configured as restricted, no changes can be made to the membership of the CTN

STP triple system CTN

An STP triple system CTN is a CTN that has the PTS, BTS, and ARB defined. This configuration provides the highest degree of resiliency after a failure.

When the connectivity fails between the PTS and the BTS, a majority vote with the Arbiter is performed to determine whether the PTS or BTS becomes the CTS. This recovery is called Arbiter Assisted Recovery (AAR).

A server holding a role in a CTN cannot accidentally be removed from the CTN in a planned disruptive action. If a planned disruptive action is attempted on the PTS, BTS, or Arbiter, STP blocks that action until the role is reassigned to another server in CTN or the role has been removed from the CTN definition. This block prevents you from accidentally causing a Degraded Triad State, which means that either PTS, BTS, or Arbiter have no coupling connectivity to the other two servers with roles in the CTN.

In a Degraded Triad State, any two of the special role servers (PTS, BTS, or ARB) agree that they cannot communicate with the third. In a Degraded Triad State, normal voting is disabled. Therefore, the recovery mechanisms that are used in a two-server CTN can be used in this situation and also when the coupling connection between the two surviving servers with roles breaks.

If a two-site CTN is created, configure the Arbiter on at the same site as the PTS. When communication between the two sites fails, the STP recovery begins, and the site with the PTS and ARB continues to run. The site with the BTS changes to standby (drop to stratum-0; unsynchronized state) until communication is re-established.

CTN recovery

STP is designed to sustain various types of failures. When properly configured, it can automatically recover the CTN by shifting the CTS definition between the PTS and BTS. The system is unaffected by the recovery operation.

There are multiple robust recovery paths that can be used to recover a CTN after a failure. Move the PTS, BTS, or Arbiter roles to another system before you perform any disruptive action on a role system, which includes the following examples:

- Disruptive maintenance, like CPC Drawer Services
- Replacement or repurposing of the system
- Power-on Reset or removal of links

Power failure handling (n-mode power)

A new automated failover option is available for IBM z16 and later systems. Because most data centers are equipped with an uninterruptible power supply (UPS) or dual utility line sources, complete power loss is unlikely to occur. IBM z16 and later can take automated recovery actions after N-mode is detected.

> **Important:** N-mode power provides automated CTN recovery if both PTS and BTS are IBM z16 and later and if N-mode power is enabled on both IBM Z systems.

This function is a one-time setup, but it can be changed at any time through the HMC (see 5.1.4, “Set time server power failover” on page 120). After N-mode power detection is enabled, STP can act if any power source detects a power cord or power side failure on the CTS. If N-mode power detection is enabled, any power cord or power side failure on the PTS or CTS sends a signal to the BTS to indicate that an N-mode power event happened on the CTS. If the BTS does not receive a message back within 30 seconds that indicates that the PTS is again fully redundant, the BTS takes over the CTS role. The PTS is still a member of the CTN and remains as the PTS, but it moves into STP stratum level 2. After power is restored and the PTS is fully redundant, the CTS role automatically returns to the PTS. The timing for the switch back of the CTS role is configured in the GUI in the tab Set time server power switch over. An example power connectivity configuration for IBM z17 is shown in Figure 2-2 on page 16.

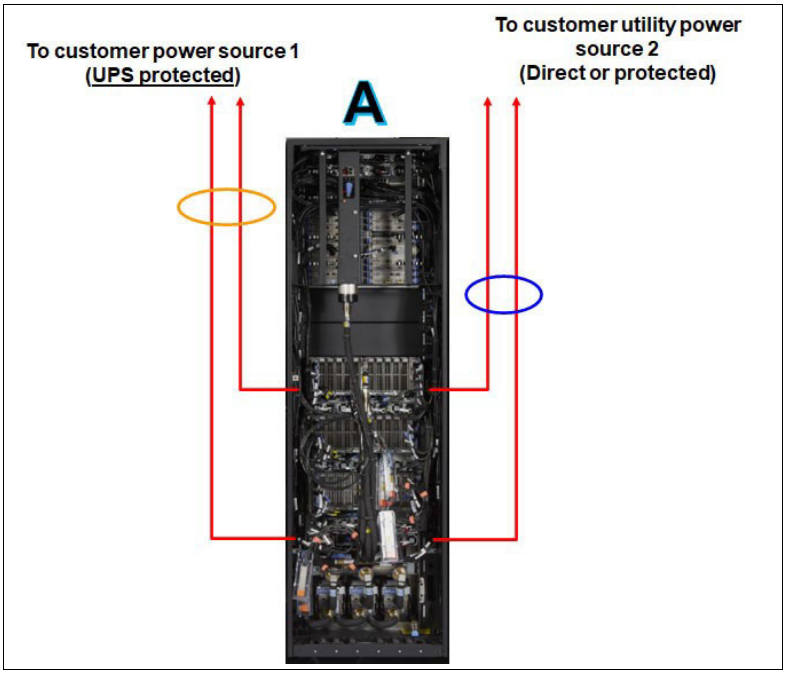

Figure 2-2 IBM z17 power connectivity for enabling N-mode power protection (PDU system}

The configuration of the Set time server power failover contains the following parameters (see also 5.1.4, “Set time server power failover” on page 120):

- A checkmark for an automated switch of the CTS to the BTS if there is a power fault detection at the PTS.
- The amount of delay for the failback of the CTS to the PTS if it is fully redundant.
- Optionally, if an automated failover occurred, the user can switch the CTS to the PTS immediately. This function is available only if there was a prior automated switch of the CTS to the BTS.

Time server power failover is available on only IBM z16 and later. In a configuration that includes anIBM z17 environment and earlier generations, consider the following aspects when you assign roles:

- Assign the PTS and BTS to an IBM z16 system or later to take advantage of the time server power failover.
- A similar function called Internal Battery Feature (IBF) is available on IBM z15 and earlier generations. This function works by assigning the PTS role to an IBM z15 that supports IBF.

Console Assisted Recovery

In a dual or triple system configuration, if the coupling connectivity between PTS and BTS is lost, STP uses Console Assisted Recovery to determine whether one of the systems is no longer operating.

> **Important:** This recovery path requires that the Service Elements (SEs) of each system are connected to the same HMC/HMA network.

The system sends a message through the SE/HMC network in both directions, from PTS to BTS and from BTS to PTS. This message indicates whether the system continues to operate or has gone offline.

If either system receives a message through the console path that a system is offline, you can transfer the CTS role to BTS or keep it with PTS in the CTN network. The system announces through coupling messages to the rest of the CTN that it has become the new CTS.

For Console Assisted Recovery to work, ensure that all SEs and the HMC are on the same network and can communicate with each other. If a power failure causes any of these systems to go offline, or if the network link between sites is lost, console recovery fails.

Triad voting

Triad voting is the most reliable recovery method, which is why you should use the STP Triple System configuration if the systems are available.

If all coupling links between PTS and BTS are lost, the systems that remain operational send a voting message to the ARB to ask whether the other system is still communicating with it. Two out of three votes determine the outcome.

If BTS loses all coupling links to PTS, it sends a message to the ARB to ask whether it is still connected to PTS. If the ARB has lost all connections to PTS, it informs BTS, and BTS assumes the CTS role for the CTN.

If the ARB remains connected to PTS, BTS does not take over. Instead, it becomes a stratum-3 system and receives time through the ARB. If the connection between PTS and ARB is lost later, the ARB notifies BTS, and BTS assumes the CTS role.

In some cases, PTS might still be operating but has lost connections to both BTS and ARB. To prevent two systems from providing different times to the CTN, PTS votes itself out and relinquishes the CTS role. It then attempts to obtain time from another system in the CTN. If no other systems are available or PTS cannot obtain time from them, it enters standby mode (drops to stratum-0; unsynchronized state) until coupling connectivity is restored.

BTS uses the same voting technique if it loses connectivity between PTS and ARB.

When a system is voted out of the CTS role, the CTN enters a degraded state in which voting is disabled. In degraded mode, the CTS role does not change due to voting until coupling connectivity is restored and the CTN exits degraded mode. This approach ensures that if the surviving system later loses connectivity to the ARB, it does not relinquish the CTS role, which would leave the CTN without a time source.

Other recovery methods, such as Going-Away Signal or Console Assisted Recovery, can still trigger a CTS role change if necessary.

When coupling connectivity is restored, the CTN exits degraded mode. If the CTS role changed during the degraded state, STP attempts to restore the CTS role to PTS. To do so, ensure that PTS has direct connections to all systems that BTS is connected to. If not, the CTS role remains with BTS. To move the CTS role back to PTS, you must either reestablish the required coupling connections or manually transfer the role through the console interface.

Going Away Signal

The Going Away Signal (GAS) is a reliable, unambiguous signal that is sent by an ICA SR coupling link to indicate that the server on which the feature is running is entering a failure mode (check stopped) state. When the GAS is sent by the CTS and received by the BTS, the BTS can safely take over as the CTS without relying on the previously mentioned recovery methods of an offline signal in a two-server CTN, or the Arbiter voting in a CTN with three or more servers.

The GAS can be sent only over coupling links that are installed directly on the CPC drawer fanouts (ICA SR). Therefore, the GAS is not supported with the CE LR coupling links in the I/O drawer.

2.3 ETS planning considerations

Use ETS to keep CTN time consistent with an external reference, such as Coordinated Universal Time (UTC), by using Network Time Protocol (NTP) or Precision Time Protocol (PTP). Pulse per Second (PPS) over a coaxial connection is also available for specific use cases.

This section compares these three technologies and explains the resiliency, accuracy, and security implications of each. It also explains physical planning and requirements, and describes how IBM Z handles corrections from ETS.

2.3.1 Network Time Protocol

The Network Time Protocol (NTP) is one of the oldest and most commonly used protocols available from the Internet. You can use NTP can as an external time source for IBM Z environments. As a best practice, it is recommended to have dedicated and redundant NTP time server appliances installed in each data center. However, depending on the redundancy, accuracy and security requirements, you can use a common enterprise reference or subscribe to NTP services from a provider over the internet.

IBM z17 and later support NTP version 4. For more information about the NTPv4 protocol and algorithms, see jRFC-9505 and NTP Project.

NTP resiliency

NTP resiliency refers to the network's ability to maintain accurate time synchronization despite network issues or failures by using multiple NTP servers, redundant pathways, and robust monitoring.

For IBM z17 and later generations, you can configure a maximum of three NTP servers on each IBM Z system in an STP-only CTN. IBM Z uses NTPv4 Peer and Poll processes and uses Clock Filter, Select, Cluster, and Combine algorithms, to determine the best time source.

Simply stated, a truechimer is a clock that maintains timekeeping accuracy according to a previously published and trusted standard, and a falseticker is a clock that does not.

For each NTP server, two processes operate in parallel: a peer process that receives and processes each packet, and a poll process that sends packets to the server at programmed or programmable intervals. The system maintains state variables and data measurements separately for each pair of these processes in a memory block containing the *peer state variables*. The peer and poll processes, and their variables, collectively form an association.

As each packet arrives, the system compares the server time to the system clock and determines an offset specific to that server. The system then modifies these offsets using the Select, Cluster, and Combine algorithms to deliver an overall correction to the NTP discipline process. The processes and algorithms perform the following actions:

- The Peer Process performs a series of error checks on received packets to validate them for further handling by the Clock Filter algorithm. An association that has valid peer variables is then called a candidate.
- The Poll Process generates output packets at regular, programmable poll intervals for the Peer Process.
- The Clock Filter Algorithm is designed to select the best sample data from the previously mentioned processes and to reject noise spikes caused by packet collisions and network congestion. It incorporates a suppressor to detect and remove transient outliers (popcorn spikes).
- The Select Algorithm determines which time sources are correct (truechimers) and which are not (falsetickers). It performs several checks to sift through the candidate sources. These checks include verifying the stratum level and root distance, and ensuring the source is reachable. The algorithm then uses correctness intervals to identify the truechimers by finding the intersection interval that contains points from the largest number of correctness intervals. A selectable is a truechimer if it is among the products of the select algorithm.
- The Cluster Algorithm sifts through the truechimers from the select algorithm to identify the sources that provide the best accuracy. It evaluates the peer jitter and select jitter from all of the truechimers, and uses a sort strategy to produce a rank order of survivors from most favored to least favored.
- The Combine Algorithm produces a weighted average of both offset and jitter from a list of surviving time sources. Each source is assigned a weight based on the reciprocal of its root distance, which is then normalized so that the sum of the weights equals unity. This combined offset is used to synchronize the system clock.

For IBM z16 and earlier generations, a maximum of two NTP servers can be configured on each IBM Z system in the STP-only CTN. When two NTP servers are configured, you must select the preferred NTP server. This NTP server is called the preferred NTP server. The other is called the non-preferred NTP server. A Simple Network Time Protocol (SNTP)[3] client (in IBM z16 and earlier generations' firmware) uses the time information from the preferred NTP server to perform the time adjustment. However, STP automatically accesses the second NTP server if the preferred NTP server becomes unavailable.

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 preferred NTP server has a stratum level that is lower in the hierarchy than the nonpreferred NTP server (NTP stratum 1 server is a better choice than NTP stratum 2).
- The time that is obtained from the preferred NTP server has less accuracy than the nonpreferred NTP server (or if a discrepancy between the time information that is provided by the two NTP servers exists).

An HMC can be used to provide added NTP resiliency. This is particularly useful when only one dedicated NTP timeserver appliance is available. An example of this method is depicted in Figure 2-3 on page 20.

[3] SNTP a simple version of NTP used to synchronize the time of computers on a network.

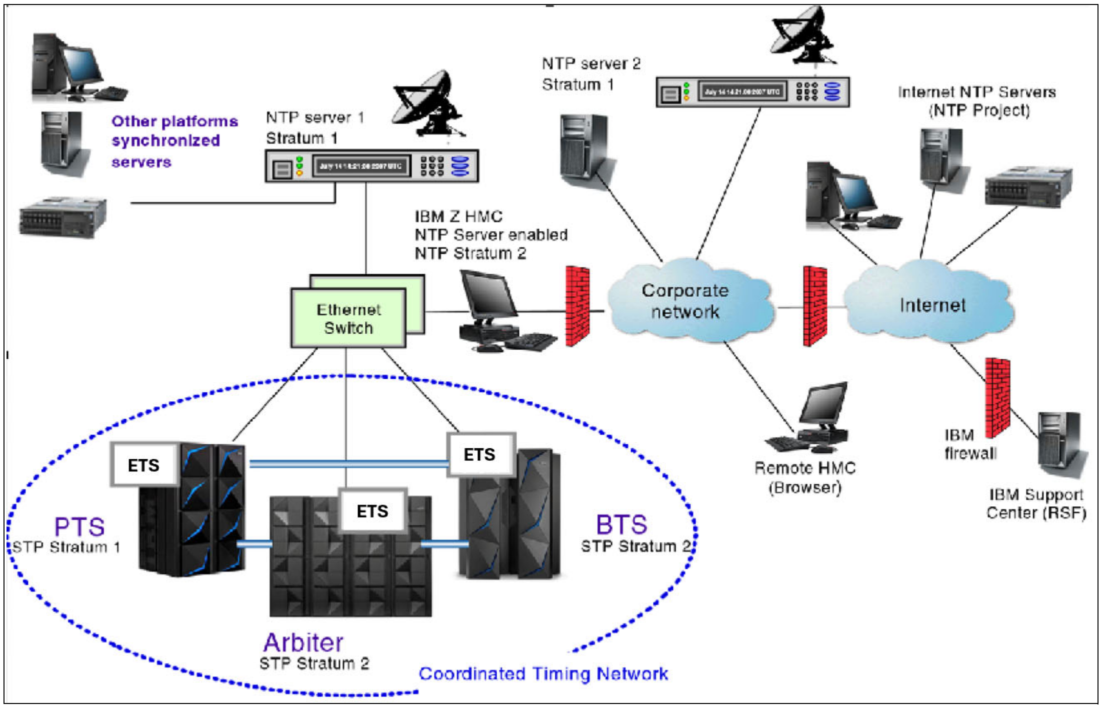

Figure 2-3 Configuration for NTP resiliency

Because the HMC is connected to the SE LAN, you can configure an NTP server on the HMC. The HMC can connect by using a different network adapter to a remote NTP server that is on the corporate intranet or on the internet to access its time reference (NTP server 2 in Figure 2-3). The use of the HMC can be especially useful on IBM z16 and earlier generations where NTP is not supported by the system NTP client and connection to only two NTP sources is possible. Although this approach enhances the resiliency and security of the ETS connectivity, it can have an accuracy penalty. Consider the tradeoffs, particularly if you must comply with strict accuracy regulatory requirements.

Alternatively, similar to the HMC, you can configure a separate NTP server (for example, running on Linux with dual Ethernet adapters) and connected directly to the SE or CPC drawer ETS LAN. This configuration provides separation between the HMC, SE, or CPC drawer ETS LAN and the external networks.

NTP accuracy

Time accuracy, relative to Coordinated Universal Time (UTC), depends on how closely the NTP server time aligns with UTC. When you require high accuracy, ensure that any timeserver appliance used as an external time source for IBM Z obtains its time directly from a primary reference clock, such as GNSS/GPS, a local atomic reference, or a subscription to a laboratory that participates in UTC time dissemination. This configuration is known as a stratum 1 NTP source.

Note: Like STP, NTP is a hierarchical protocol with servers organized into stratum levels. However, the NTP stratum and STP stratum differ significantly, so you must not interchange these terms.

The NTP protocol can suffer from link asymmetry, which occurs when the time to send an NTP packet differs from the time to receive a response. Link asymmetry typically results from network congestion and queuing delays in both the NTP client/server and the network equipment along the route. Because NTPv4 does not account for link asymmetry, IBM Z cannot remove its effects from time adjustments.

Dispersion is the defined as the uncertainty of an adjustment. Dispersion is calculated by using the total round-trip time of an NTP exchange. The longer it takes to receive a response to an NTP request, the higher the dispersion. NTP assumes zero link asymmetry when calculating the offset, so any actual asymmetry introduces error into the adjustment. IBM Z externalizes this error to you through the dispersion value.

When you use NTP on a local network with IBM z16 or later, expect NTP dispersion to be no smaller than a few hundred microseconds. If you introduce more complex NTP configurations, such as using an enterprise or public internet NTP source, you might see significantly higher dispersion. This approach should not be considered acceptable on its own when you must comply with strict regulatory requirements for time accuracy to UTC.

NTP security

As a best practice, configure NTP sources so that they connect directly to the CPC drawer (for IBM z16 and later generations) or to the SE LAN (for IBM z15 and earlier generations). Treat each of these connections as a private (dedicated) LAN in most configurations and keep them as isolated as possible.

If you connect the NTP client, which runs as an internal firmware service, directly to an NTP server on the corporate network or the internet, you introduce security risks, even if the local network is protected by a firewall.

The ability to define an NTP server on the HMC addresses potential security concerns because the HMC is attached to both the SE network and intranet (see Figure 2-4).

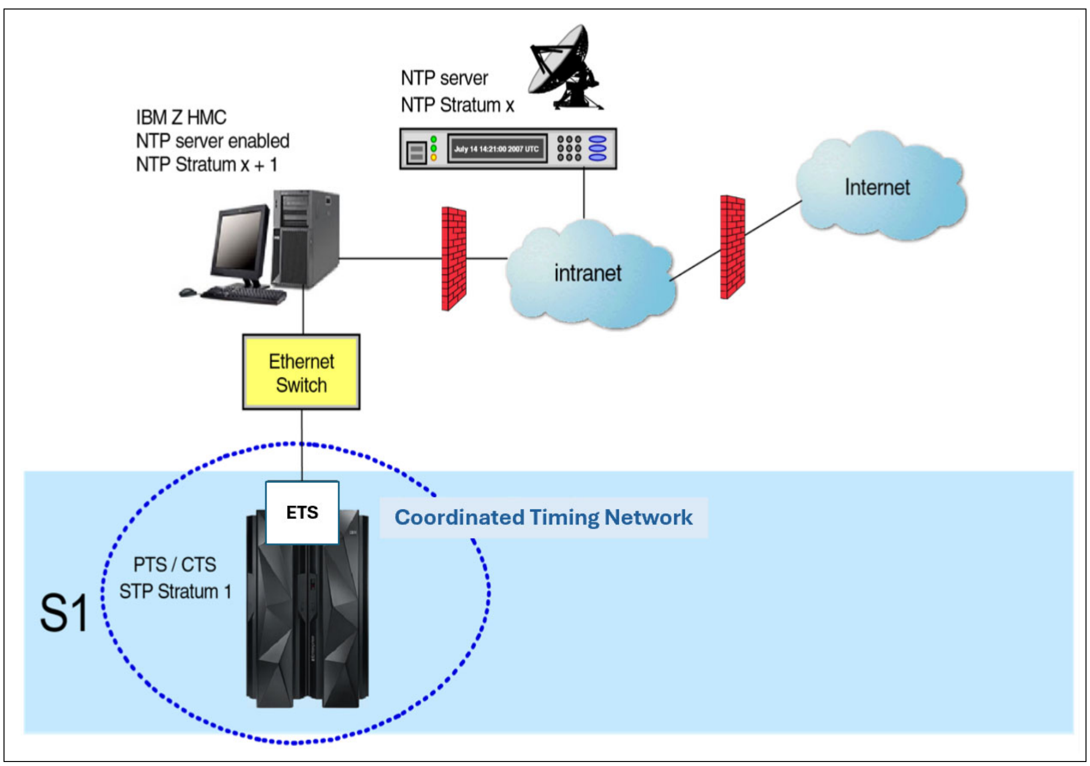

Figure 2-4 HMC acting as an NTP server

The HMC has two separate LAN ports that are physically isolated from each other:

- First port: Connects to the HMC, SE, or CPC drawer ETS LAN. This network can provide an NTP server for IBM Z systems.
- Second port: Used by the NTP client on the HMC to get time from another external NTP source.

With this configuration, the NTP server on the HMC can access a different NTP server through its own dedicated LAN connection to keep accurate time.

Note: The NTP server function on the HMC does not provide a PPS output.

HMC NTP Authentication support

You can use HMC NTP authentication. To use this option, configure the HMC as the NTP server for CTN. The following scenarios describe the situations in which you might use this option:

- NTP authentication support with a proxy

  Some configurations use a proxy to gain 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 the target servers on the web. Authentication can be configured on the proxy to communicate to the target time sources.

- NTP authentication support with a firewall

  Some configurations use a firewall. HMC NTP requests can pass through the firewall. When you use this configuration, use HMC authentication to ensure that the timestamps are not tampered with.

- NTP-integrated security

  With the symmetric key and NTS authentication, the highest level of NTP security is available. HMC 2.17.0 and later provides windows that accept and generate key information to be configured in the HMC NTP configuration:

  - Symmetric key 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 that are encrypted with a secret key can be decrypted only with the same secret key.

  - Network Time Security (NTS) authentication

    Network Time Security for NTP, as described in RFC-8915, specifies a mechanism for using Transport Layer Security (TLS) and Authenticated Encryption with Associated Data (AEAD) to provide cryptographic security for the client/server mode of the Network Time Protocol (NTP). This is the ideal solution for securing NTP traffic because it is designed specifically for the needs of time protocol traffic and uses standard building blocks such as TLS.

Note: The configuration of network components, such as routers or firewall rules, is beyond the scope of this document. For more information about security guidelines and recommendations for the HMC or SE network, see Hardware Manage Console Security, Hardware Management Console Operations Guide and Installation Manual for Physical Planning at IBM Z documentation and select the manual specific to your IBM Z model.

2.3.2 Precision Time Protocol

Precision Time Protocol (PTP) is a newer standard (IEEE-1588). PTP offers significantly higher accuracy (nanoseconds or microseconds) compared to NTP (milliseconds) by using hardware timestamping, whereas NTP relies on software processing.

The usage of PTP as an ETS for IBM Z addresses the requirement of a highly accurate time source. PTP requires the supporting infrastructure that can transmit PTP signaling to the IBM Z system, including a PTP timeserver appliance and network infrastructure with integrated PTP support.

For more information on PTP, see IEEE 1588 Working Group.

PTP resiliency

Redundant PTP sources are recommended for ETS resiliency. Ideally, allocate two dedicated PTP compatible LANs. Connect these LANs to the ETH3 and ETH4 interfaces on an IBM z15 SE or to the ETS1 and ETS2 interfaces on IBM z16 or later CPC drawer.

PTP relies on the Best Time Transmitter Clock Algorithm (BTCA) to provide resiliency within a PTP domain. This works by using a voting algorithm to determine which PTP source to follow out of all PTP Grandmasters in a particular PTP domain. This voting is done in real-time, so when the best source goes offline, another source can take its place automatically.

With the simplest PTP network in multicast mode, a domain is contained to a specific LAN. When using this simple configuration, it is recommended that multiple PTP Grandmasters be added to each domain in each LAN to eliminate any single point of failure. When only two PTP sources are available, ensure that both sources are connected to both ETS LANs to take full advantage of the BTCA in each LAN.

You can design more sophisticated PTP domains that span multiple LANs by routing multicast PTP traffic across multiple LANs, by using PTP boundary clock support in LAN gateways, or by using PTP Unicast mode (in IBM z17 and later generations only). When you use one of these special configurations, be sure to consider the implications on resiliency. For more information, contact your IBM Support Representative.

In IBM z17 and later generations, PTP domains are run through the same select and combine algorithms that are used for NTP sources. For more information on these algorithms, see 2.3.1, “Network Time Protocol” on page 18. When designing and configuring highly resilient ETS connectivity, it is recommended to configure not only multiple PTP sources but also multiple NTP sources. When more sources are used, it is easier to detect issues and determine consensus on the correct time. When you design an ETS network that contains only two timeserver appliances, it is still useful to connect to each source using both NTP and PTP. The connection provides protocol redundancy but also allowing for combination of integrated security support from NTP and enhanced accuracy from PTP.

An example of PTP connectivity (as ETS) is shown in Figure 2-5 on page 24.

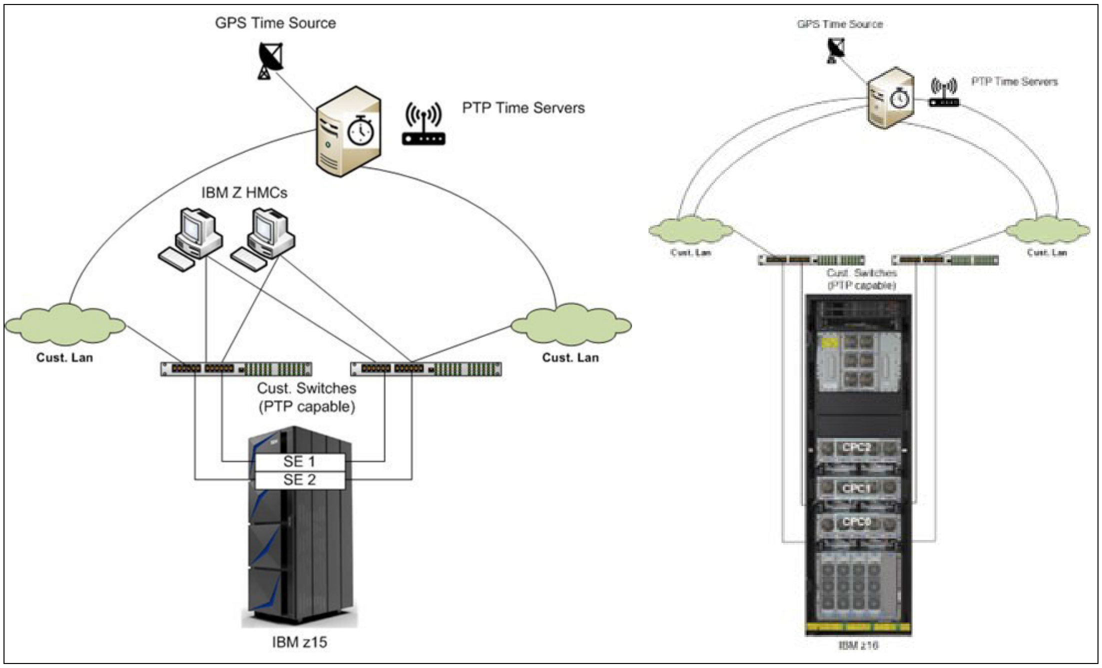

Figure 2-5 PTP connectivity diagram: IBM z15 (left) versus IBM z16 (right)

PTP accuracy

Similar to NTP, ensure that any PTP timeserver appliance that is used as an external time source by IBM Z obtains its time directly from a primary reference clock, such as GNSS/GPS, a local atomic reference, or from a subscription to a laboratory that participates in UTC time dissemination.

Unlike NTP, PTP accounts for queuing delays in network infrastructure by using integrated hardware support. This capability is commonly known as hardware timestamping. The PTP timeserver appliance, IBM Z ETS interfaces, and all connecting network infrastructure take timestamps when PTP messages enter and exit each device. The residence time, which is the time that the message resides on the device, is calculated and added to the correction field. This process removes all queuing delays from the time delta calculation, which greatly improves accuracy and can eliminate link asymmetry. As a result, PTP can provide synchronization in the nanosecond range without requiring special equipment.

Typically, you can upgrade existing timeserver appliances that serve NTP to support PTP. Depending on the vendor, this upgrade might require purchasing a new license or installing a dedicated PTP interface card. On-path devices such as routers and switches do not require PTP support, but enabling it is highly recommended to achieve the accuracy benefits of PTP. There are two types of PTP support: PTP transparent clock and PTP boundary clock.

PTP transparent clocks update the correction field as described earlier. This update is does not affect accuracy. PTP boundary clocks work similarly to NTP servers. Each boundary clock on the path increases the step count and might introduce some error, depending on the implementation. Read your equipment documentation to determine whether this feature is available and how to enable it.

PTP security

As a best practice, attach PTP sources directly to ETS LANs that connect to the CPC drawer (IBM z16 and later generations) or the SE interfaces (IBM z15 and earlier generations). Many configurations consider these LAN connections to be private (dedicated) LANs. Keep them as isolated as possible. Integrated PTP security is experimental. Allowing the PTP client that runs as an internal firmware service to communicate with a PTP source on the corporate or external network (including the internet) introduces potential security risks, even if the LAN is behind a firewall.

For IBM z17 and later generations, configure External Time Source connectivity with multiple PTP and secured NTP sources. The system runs the select and combine algorithms against all External Time Sources. This combined configuration provides the added accuracy of PTP along with the added security of NTP when you use Network Time Security.

2.3.3 Pulse Per Second

Pulse Per Second (PPS) uses an electrical signal that consists of a square wave with a one-second interval. The signal travels over a copper coaxial cable that connects to a BNC connector on the IBM Z CPC drawer. PPS signals rise at the exact beginning of each second as determined by the time source. This rise allows PPS clients such as IBM Z to calculate their offset by measuring the difference from the beginning of the last second each time the signal goes high. PPS does not provide any other data. This means that PPS alone cannot provide traceability to UTC. Although you know when the last second occurred, you do not know the UTC timestamp for that second. For these reasons, you must use NTP or PTP from the same time source in addition to PPS to provide UTC timestamps and traceability to the UTC reference clock. Correlate a PPS source with an NTP or PTP source when you create an ETS configuration. You cannot confirm this information unless the PPS signal does not generally agree with the NTP or PTP source.

The coaxial cables that transport PPS are limited in length because of voltage loss over the copper medium. You must calculate this length carefully when you consider PPS. Because of the complexity of designing, procuring, installing, and maintaining a network of coaxial cables, use PPS only when high accuracy is required and PTP is not feasible.

When you consider PPS connectivity, understand that timeserver appliances typically have multiple BNC connectors that support various input and output signal types. These ports look the same except for labeling, either physically or through a user interface. It is extremely important to inspect these ports carefully and confirm connectivity to the 1 PPS (1 Hz) output port before you connect the coaxial cable to IBM Z.

Figure 2-6 on page 26 is an example of this redundancy for PPS.

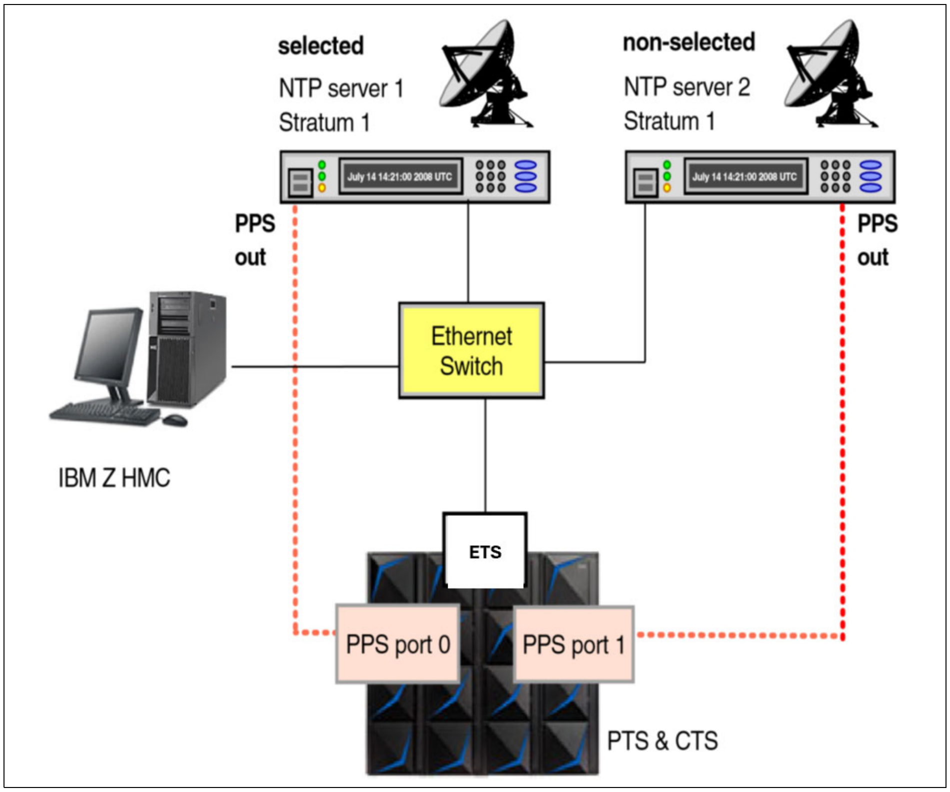

Figure 2-6 IBM Z with NTP PPS

2.3.4 Network Time Security

Network Time Security (NTS) is a modern approach that addresses the specific requirements of time protocol traffic such as NTP or PTP. It uses well-established building blocks such as Transport Layer Security (TLS) and Authenticated Encryption with Associated Data (AEAD) to provide cryptographic security.

NTS for NTP (RFC 8915) authenticates the NTPv4 client/server mode that IBM Z uses. When you use a compatible NTS for NTP, you can authenticate NTP traffic without continuously managing symmetric keys. This approach is more seamless and scalable because it relies on the existing X.509 public key infrastructure in the enterprise.

2.3.5 Comparison of external time sources

IBM Z systems can synchronize their clocks by using standardized external time source technologies. The choice depends on your specific needs around resiliency, accuracy, and security and is often driven by industry regulations, business requirements, or best practices.

Standardized technologies

IBM Z systems can synchronize their clocks by using three main technologies: NTP, PTP, or PPS:

1. Network Time Protocol (NTP)

   NTP is the most widely available time synchronization protocol. It operates in a client/server model within a hierarchical stratum system. It can be sourced from dedicated Stratum 1 time servers, enterprise-wide NTP servers, or public internet servers:

   – Security:

     - Includes integrated security features.
     - The latest addition is Network Time Security (NTS), which uses TLS and AEAD encryption by using existing x509 certificate infrastructure.
     - Using unsecured NTP over public or semi-private networks is discouraged due to potential vulnerabilities.

     For secure setup:

     - Use NTS on IBM z17 or later.
     - Use Hardware Management Console (HMC) security features on IBM z16 and earlier.

   – Accuracy:

     - Software-based timestamping introduces asymmetry because every device along the path, which includes IBM Z, adds delay variations.
     - Transient network conditions can degrade accuracy unexpectedly.
     - Not recommended as the sole method when strict regulatory accuracy standards must be met.

   – Resiliency

     NTP relies on the select and combine algorithms to pick the best NTP sources (z17 and later only). When feasible three NTP sources should be defined for maximum resiliency.

   – Use case:

     Use NTP primarily for general-purpose synchronization with moderate accuracy needs. For higher precision, combine it with PTP or PPS.

2. Precision Time Protocol (PTP)

   PTP provides tighter synchronization than NTP by using hardware timestamping at network interface cards and switches. It requires PTP-compatible time source and network infrastructure:

   – Security

     At the time of writing, integrated security features are still mostly experimental. Use network segmentation where possible. On z17 or later, it is recommended to combine PTP with NTS secured NTP.

   – Accuracy

     Capable of synchronizing devices within nanoseconds, or picoseconds under ideal conditions, over standard Ethernet networks.

   – Resiliency

     Relies on the Best TimeTransmitter Clock Algorithm to provide resiliency in a domain. In addition, resilient domains can be used for an extra layer of redundancy.

– Mechanism

Uses hardware timestamps to reduce latency variability that is caused by software processing delays that are common in NTP setups without the need for extra or special equipment or cabling.

– Use case

Ideal where high precision timing is required such as financial trading platforms, telecommunications networks, or regulated industries requiring strict compliance with timing standards.

3. Pulse Per Second (PPS)

PPS signals provide an electrical pulse precisely once per second to mark exact seconds boundaries directly at hardware level. It requires a PPS generator and coaxial cables wired directly to IBM Z:

– Security

PPS is an analog signal that is delivered over dedicated coaxial cables, excluding an insider attack. Security is not a concern.

– Accuracy

Offers precise timing that is wired directly to the internal clock disciplining mechanisms inside IBM Z systems.

– Resiliency

Requires a corresponding NTP or PTP connection. Dedicated connections to multiple PPS sources can be a challenge.

– Use case

Best suited when ultimate timing precision is needed alongside specialized hardware support, such as GPS receivers providing PPS output connected directly to system clocks.

Summary of time synchronization technologies

Table 2-1 compares the three main time synchronization technologies based on accuracy, security, and usage.

Table 2-1 Comparison of time synchronization technologies

Technology	Accuracy	Security Features	Typical Use Case
NTP	Milliseconds	Supports TLS via NTP	General purpose; moderate accuracy
PTP	Nanoseconds Picoseconds	Depends on implementation; usually secured internally	High precision environments
PPS	Sub-microsecond	Physical signal; inherently secure if isolated	Highest precision with direct hardware

General recommendations:

The following list provides practical recommendations to help you select and implement effective time synchronization tailored to your environment:

- For general use cases without stringent accuracy demands, use secured NTP (preferably with NTS).
- When regulatory compliance requires tight synchronization, supplement or replace with PTP, PPS, or both.

- Always consider your environment's security posture before deploying any external time source over public or semi-private networks.
- Combine multiple methods if necessary to balance resiliency, accuracy, and security requirements effectively.
- Provide fully redundant and isolated timing LANs, which can overlap with existing SE customer-supplied networks.
- Use fully redundant time sources:
  - IBM z16 and earlier: 2 NTP or 2 PTP sources are recommended.
  - IBM z17 and later: 3 NTP sources plus 2 PTP sources are recommended.
- Use all available NTP and PTP sources even if they originate from the same physical device.
- Position/Navigation/Timing (PNT) solutions are typically sourced from GNSS/GPS antennas.
- Consider redundancies or alternatives such as Low Earth Orbit (LEO) subscription services, terrestrial networks, or local atomic references.
- Accuracy requirements:
  - IBM z15 and earlier: PPS is required to meet any accuracy standards.
  - IBM z16 and later: Prefer PTP where available for accuracy enforcement. Otherwise, supplement NTP with PPS.

For more information, see the white paper IBM z16 Server Time Protcol(STP): PTP/NTP Time Direct to CED Accuracy Performance.

2.3.6 ETS Network physical planning

IBM Z connects to External Time Sources by using copper Ethernet connections (CAT5e or better is recommended) to either the SE on IBM z15 and earlier generations or the CPC drawer on IBM z16 and later generations.

A major design principle of ETS is the use of fully redundant paths through customer-provided networks. The key to success in designing these networks is simplicity. Limit the number of clients and supporting network infrastructure to essential components only. This approach provides two advantages:

1. It increases accuracy by reducing the risk of network congestion and queuing delays. The longer the travel time of NTP or PTP traffic, the higher the dispersion. Queuing delays particularly add asymmetry to the link, which invalidates the general assumption that the offset is the center of the dispersion band.
2. This approach increases security by limiting the exposure of unauthenticated NTP or PTP traffic. In many enterprise networks, time protocol traffic is considered an exception to general security checks. This situation creates opportunities for covert channels where attackers can exfiltrate data under the guise of NTP or PTP traffic.

ETS connectivity on IBM z15 and earlier generations

External Time Source connectivity on IBM z15 and earlier generations uses the em3 and em4 connections on the primary and alternate SE. This path provides an accuracy of approximately 100 milliseconds, accounting for NTP or PTP protocol dispersion and internal SE dispersion. Do not consider ETS connectivity through this path without PPS when you must adhere to financial regulations of any kind.

The em3 connection is already connected to the customer-provided network. This connection exists on both the primary and alternate SE. However, only one SE is active at a time. This configuration means that automatic failover is not possible without connectivity through the em4 connection on both SEs. For optional resiliency, provide a secondary ETS-only network that connects to the em4 connection on both SEs when you use this generation of systems. Ensure that the network remains in a separate subnet from the customer-provided network that connects to the em3 connection.

ETS connectivity on IBM z16 and later generations

External Time Source connectivity on IBM z16 and later generations connects to the front of the A10 drawer and, if applicable, the A15 drawer. The first ETS network (ETS1) always connects to A10-H101-J03. On single-drawer systems, the second ETS network (ETS2) connects to A10-H102-J03. For multi-drawer systems, connect ETS2 to A15-H101-J03 and A15-H102-J03 on IBM z16 or IBM z17 and later generations. For more information, see the Installation Manual for Physical Planning in IBM Resource Link® for your IBM Z system generation.

The CPC drawer ETS attachment requires extra tasks, such as physically wiring the ETS network and defining IP addresses, gateways, and static routes (if required) for ETS1 and ETS2 ports in the Customize Network Settings panel on the SE. However, this additional connectivity from the CPC drawer to the customer network provides significantly better accuracy compared to previous IBM Z systems that use the SE interface.

When you design the networks for these connections, consider the routing and failover capabilities of the ETS1 and ETS2 connections. Typically, ETS1 and ETS2 must reside on separate subnets. The only exception to this rule is when you use static routes through multiple dedicated gateways. You can use the existing customer-provided SE network as one of the two ETS networks. However, the second ETS network should remain in a separate subnet.

NTP requirements

The IBM Z platform is compatible with NTP version 4.

NTP communicates over UDP Port 123. This port must remain open and accessible for NTP traffic to flow.

IPv4 and IPv6 transport is supported.

Hostname resolution is supported through DNS. When you use this feature, UDP/TCP Port 53 must remain open and accessible for DNS traffic to flow.

The following integrated security features are supported:

- Symmetric Key Authentication: Supported by HMC NTP clients only.
- Network Time Security (NTS): Supported by IBM z17 and later systems and HMCs. This feature requires TCP Port 4460 to remain open and accessible for NTS traffic to flow.

PTP requirements

The IBM Z platform is compatible with PTPv2 and PTPv2.1. PTP traffic requires UDP Port 319 and UDP Port 320 to remain open and accessible.

On-path network infrastructure should support PTP transparent clocks. IPv4 and IPv6 transport is supported. IPv6 is supported on IBM z17 and later.

End-to-End (E2E) Delay Mechanism is required. One-step or two-step mode is allowed.

Set Sync, Announce, and Delay Requests to once per second. IBM z16 and earlier generations require domain 0.

IBM Z supports two modes of operation:

1. Multicast/Hybrid mode (Enterprise Profile): This mode requires multicast Sync/Follow-Up and Announce messages and unicast Delay Request/Delay Response messages. For security reasons, IBM Z does not transmit IGMP packets. Therefore, you must disable IGMP snooping when you use this mode.
2. Unicast mode on IBM z17 and later: This mode requires support for unicast PTP contracts. All messages use unicast.

PPS requirements

PPS cables should measure a maximum of 150 feet (45 meters) between the last distribution point and the oscillator card PPS port. Use RG6 cabling to reduce noise. Use a right-angle adapter to connect this cable to the oscillator card PPS port.

Before you use PPS, test the functionality by using the Control Pulse Per Second Signal action on the Manage System Time task.

For more information on required signal characteristics, see the *Install Manual for Physical Planning for your specific system* from IBM Z documentation.

2.3.7 ETS steering

IBM Z controls its internal clock by using two functions: adjustment steering and base steering. These functions are described in this section.

Adjustment steering

Adjustment steering occurs automatically through the External Time Source or manually through the Manage System Time panel. Adjustment steering gradually adjusts the local clock. The offset that you specify through adjustment steering is incorporated in small increments or decrements so that operating systems, subsystems, and applications are unaware that time is speeding up or slowing down.

You can use adjustment steering to change the time by a maximum of ±60 seconds. You can implement adjustments greater than 60 seconds in multiple increments of ±60 seconds. This process can take considerable elapsed time to complete.

> **Important:** The adjustment steering rate is approximately 1 second for every 7 hours.

Base steering

Base steering occurs automatically when you use an ETS configuration with NTP Stratums 1, 2, or 3, or better PTP sources.

Each IBM Z system contains multiple oven-controlled crystal oscillators (OCXOs). These oscillators are designed to achieve an accuracy greater than 2 ppm, which means a maximum drift of 2 microseconds per second. Oscillators are not perfect from the factory. Environmental variables such as temperature, pressure, and voltage can affect oscillator frequency. In addition, oscillators degrade slightly as they age.

IBM Z compares multiple OCXO oscillators and uses the long-term adjustment average from external time sources to calculate the base steering rate. When fully calibrated, IBM Z can achieve a skew that is 100 times better, in the range of 10 ppm or tens of nanoseconds of drift per second. During calibrated steady-state operation with highly stable PTP or PPS external time sources, adjustment steering rates often reach the nanosecond range.

Table 2-2 provides a summary of STP clock adjustment functions.

Table 2-2 STP time steering

Adjustment method	NTP/PTP server access adjustment	Manual adjustment	Manual set
Adjustment steering	Yes	Yes	No
Base steering	Yes, after multiple ETS accesses	No	No

Important: After a CTN is configured, the only way to set the date and time is to unconfigure the CTN and then go through the process to initialize the time again. For more information, see Chapter 3, “CTN configuration and operations” on page 41.

2.4 Operating system considerations

Host system time synchronization through STP and ETS provides a centralized reference clock that guest operating systems can use across a hierarchical order of IBM Z systems. Implementing the TOD clock in the hardware layer provides the following advantages:

- Guarantees consistency and coherence of timestamps across different nodes, which is essential for maintaining correctness in distributed algorithms and protocols. This consistency also supports event ordering, data integrity, fault detection, security, and authentication.
- Reduces guest operating system usage for time synchronization because you no longer need to set up external time synchronization for each guest. The host system handles synchronization in one place. In addition, external coupling links pass time synchronization signals for STP. These same coupling links can also exchange timekeeping information and CF messages in a Parallel Sysplex.

Server Time Protocol not enabled

When STP is not enabled, the TOD clock initializes to the TOD value of the SE at Power-on Reset. The operating system that runs in an LPAR can set its own TOD value, which is the only TOD reference it uses. Setting the TOD clock for one logical core in the LPAR sets the TOD clock for all logical cores in that LPAR but does not affect logical cores in other LPARs. The TOD clock value is used during LPAR activation or until a subsequent Set Clock instruction is issued in the LPAR.

Server Time Protocol enabled

When STP is enabled, the TOD clocks for each PU are set to the TOD value from STP. The operating system in each LPAR can independently choose whether to synchronize to the current STP time source, if it is present. Operating systems in LPARs that synchronize to STP run with identical TOD values. Operating systems in LPARs that do not synchronize to STP do not need to be aware of STP and can set their TOD values independently of other LPARs.

2.4.1 z/OS

When you configure a new CTN, you must plan for z/OS considerations. Planning for z/OS includes the following considerations:

- Apply STP-related maintenance to all z/OS systems that require multisystem time synchronization.
- Update SYS1.PARMLIB(CLOCKxx). For more information about available parameters, see CLOCKxx (time of day parameters).

  Pay attention to the ACCURACY keyword. When you set this keyword, the system alerts you when an adjustment from the ETS exceeds a user-specified accuracy threshold.

For a system with STPMODE=YES and PLEXCFG=MONOPLEX, the system switches from LOCAL timing mode to STP timing mode when IBM Z is STP-synchronized. In z/OS, you cannot manually change the value of the TOD setting when synchronized to STP (STPMODE=YES).

You can use the LPAR time offset to support different local time zones in multiple sysplexes by using the STP CTN. Many sysplexes must run with a LOCAL=UTC setting in a sysplex (STPMODE=YES), where the time returned from an STCK instruction yields local time. To fulfill this requirement, the time that you initialize for the STP CTN must be local time. With LPAR time offset support, multiple sysplexes can each have their own local time reported from an STCK instruction. For example, you can set the STP CTN to GMT, specify an LPAR time offset of minus 5 hours for one set of sysplex partitions, and specify an LPAR time offset of minus 6 hours for a second set of sysplex partitions.

When STP is available, use appropriate CLOCKxx statements, such as STPMODE=YES. This approach offers the following advantages over using SIMETRID:

- Changes to DST are handled without manual procedures, such as using the SET TIMEZONE command.
- z/OS posts messages to the z/OS console when it detects certain conditions with STP, such as messages related to ETS.
- z/OS recognizes leap second adjustments that are made in the Manage System Time panel.

Attention: To change from z/OS systems running with SIMETRID to STPMODE=YES, you must plan for an outage of the z/OS systems to activate STP mode. After you configure the CTN, shut down and perform an IPL of the z/OS systems by using the updated CLOCKxx member.

Although running z/OS in SIMETRID simulation mode when STP is available is typically not recommended, it can be helpful in some situations. You can configure a CTN concurrently without stopping production because, in SIMETRID simulation mode, STP interrupts and machine checks do not apply. Continuing to run on a system that is configured for STP exposes problems due to unprocessed sync checks. When a sync check occurs, the LPAR takes an offset equal to the sync check and generates an STP machine check. The LPAR has an offset (an EPOCH) that is not cleared, and the logical TOD for the SIMETRID LPARs drifts away from the STP time of the system as each sync check occurs.

With currently supported z/OS versions, the LPAR offset is not cleared at IPL, so there is no mismatch when systems undergo an IPL because of the LPAR EPOCH. However, when you reactivate an LPAR, the LPAR EPOCH is cleared. When z/OS undergoes IPL in that reactivated LPAR, it cannot join the sysplex because of the LPAR EPOCH mismatch. The running system has a nonzero LPAR EPOCH, and the system undergoing IPL has a zero EPOCH. If you must reactivate one LPAR in the sysplex, you must reactivate all LPARs in the sysplex to ensure that there is no mismatch of LPAR EPOCHs.

For more information on using SIMETR, see Implementing and using SIMETR WITH STP.

z/OS messages

To improve the delivery of important information to the operator and better integrate with system automation tools, z/OS provides STP-related messages in addition to the hardware messages that are posted on the HMC.

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

The IEA395I message informs you that a CTS change has taken place:

IEA395I THE CURRENT TIME SERVER HAS CHANGED TO THE cccccc

Where cccccc is BACKUP or PREFERRED. This informational message does not require any action but ensures that the operational staff that are responsible for STP are aware of the change. For more information, see IEA395I.

The IEA0371 message informs you that an STP alert was received:

IEA037I STP ALERT RECEIVED. STP ALERT CODE = nn, REASON = rr

The alert codes have the following values:

- 01: The alert is related to the ETS.
- 02: The alert is related to the state of the PPS ports.
- 04: The alert is related to the hardware.

For more information, see IEA037I.

There are several IEAxxx and IXCxxx messages that report current and changed timing statuses. For example, the following message reports the result of a successful rename of a CTN:

```
IXC438I COORDINATED TIMING INFORMATION HAS BEEN UPDATED
FOR SYSTEM: SC60
PREVIOUS CTNID: CTN2
CURRENT CTNID: CTLSOMRS
```

In general, no z/OS messages are posted on only the PTS, BTS, or Arbiter. However, you might see the following scenarios:

Certain messages do not appear on the CTS because 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.
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 nnnnnnnnnn is
'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
```

For more information, see IEA messages and IXC messages.

2.4.2 z/VM

z/VM uses the STP facility to generate timestamps for guest and system DASD write I/O operations so that these operations stay synchronized with the I/O operations of other systems. With this support, you can use IBM System Storage z/OS Global Mirror to asynchronously replicate data that z/VM and its guests use over long distances.

z/VM uses the STP facility to generate timestamps for guest and system DASD write I/O operations so that these I/O operations are synchronized with the I/O operations of other systems. With this support, for data that is used by z/VM and its guests, you can use IBM System Storage z/OS Global Mirror to asynchronously replicate data over long distances.

Baseline STP support

Baseline STP support available in z/VM provides the following functions:

- Synchronizes the z/VM TOD clock with the STP server at IPL. With STP support enabled in z/VM, the TOD clock is initialized with the CST and sets the same time zone that is used by the CTN, and displays 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.

To enable this support, the following FEATURES statements for SYSTEM CONFIG have been added:

- STP_Timestamping

  Timestamps are added to write channel programs that are issued to all DASD devices that have the XRC LIC installed.
- STP_TIMEZone/STP_TZ

  The system time zone is derived from the STP server.
- XRC_OPTional

  The 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 second-level z/VM. This statement enables STP_Timestamping without STP availability. A manually specified TOD value is used for timestamping, and is intended for vendor test support.

Example 2-1 shows the FEATURES section of a SYSTEM CONFIG that has STP_TZ enabled only.

Example 2-1 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_Cmds ,             /* 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 information about the FEATURES statements, see *z/VM: CP Planning and Administration*, SC24-6271.

STP state flow in z/VM

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

- During IPL

  If STP_Timestamping, STP_TZ, or both are specified: Perform an activation by querying STP and setting the TOD clock to match the STP TOD value.

  If STP_TZ is specified: Set the system time zone to match the STP time zone.

  If either of these approaches fails, STP enters the SUSPENDED state. Otherwise, STP activation completes successfully and STP is considered ACTIVE.

- STP is ACTIVE

  If STP_Timestamping is specified, the I/O to XRC-capable DASD is timestamped (when required) and I/O to non-XRC-capable DASD is unchanged.

- STP is SUSPENDED:

  If STP_Timestamping is specified and XRC_OPT is not specified: The I/O to XRC-capable

  DASD that must be timestamped is deferred until STP becomes ACTIVE again. The I/O that does not need to be timestamped, and I/O to the non-XRC-capable DASD continues to be issued.

  If STP_Timestamping is specified and XRC_OPT is also specified: I/O to all DASD is issued without a timestamp until STP becomes ACTIVE again.

  Some events in the system might cause STP state changes:

- ACTIVE to SUSPENDED

  Occurs when an STP machine check is received that informs the Control Program (CP) that the TOD value must be synchronized.

- SUSPENDED to ACTIVE

  In response to machine checks or external interrupts that are received, the CP attempts to resynchronize with the STP server. A successful resync does not change the system TOD value, but does update the delta between the system time and the STP TOD value and the system time zone (if STP_TZ is enabled).

- ACTIVE to ACTIVE

  External interrupts might be received that require the CP to query the STP time zone information (for example, the time zone was changed through the HMC). This action does not cause STP to enter the SUSPENDED state, but it causes the system time zone to change.

CP commands

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

- QUERY STP

  This command displays the STP status. Example 2-2 shows a possible response.

  Example 2-2 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 changes the system's time zone ID and time zone offset. It was updated to display the response that is shown in Example 2-6 if z/VM is configured to use STP time zone.

  Example 2-3 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 displays the list of active and inactive time zone definitions on your system. It was updated to show the STP time zones when z/VM is configured to use it.

z/VM IOCP

If z/VM is used to create IOCP decks for the z/OS LPARs (not z/OS guest machines), z/VM can be used to create timing-only links. To define timing-only links, define the peer channel (CFP or CBP) with a control unit of UNIT=STP and no devices.

VM guests

You can establish a Parallel Sysplex among z/OS images that run as guests under z/VM. However, the guests cannot use STP. Only a write I/O to XRC-enabled storage includes STP timestamps when you specify STP_Timestamping. z/OS guests use z/VM timer services and must specify SIMETRID in the parmlib CLOCKxx member to synchronize their virtual TOD clocks.

To change the date and time of a guest system that runs under z/VM, add an OPTION TODENABLE statement in the directory entry for the guest virtual machine. Then, change the system date and time for the guest by using its procedures. Do not specify TODENABLE for guest machines that specify SIMETRID. This support requires that all TOD clocks for the system are synchronized. If you specify TODENABLE, the operator can change the guest TOD clock later.

Do not alter the TOD clock for a VM guest from the value that the LPAR sets during LPAR activation, either by systems that undergo IPL in the partition or by z/VM during its initialization.

For More information about z/VM, see z/VM.

z/TPF

STP is supported in z/TPF with an APAR. For more information, see Server Time Protocol (APAR PJ36831).

Linux on IBM Z

Linux on IBM Z supports STP-based TOD synchronization. Using STP for a Linux-based application can help with more accurate time, as time synchronization is done at the hardware layer. For more information about STP support in Linux, see TOD clock synchronization.

2.5 Management and operations

The Manage System Time task on the HMC is the primary location for all STP and ETS operations. Here you can view, configure, and manage all CTNs. It is recommended that any HMC used to manage a CTN is at the same version level as the latest generation system in a CTN. The following tasks are a sample of the tasks that can be accomplished in the Manage System Time panel:

- Initialize or modify the CTN ID.
- Initialize the time manually or by using an External Time Source (ETS) to keep the Coordinated Server Time (CST) synchronized to the time source that is provided by the ETS.
- Initialize the time zone offset, Daylight Saving Time (DST) offset, and leap second offset.
- Configure connectivity to an NTP, PTP, NTP with PPS, or PTP with PPS External Time Source.

- Manage the CTN membership (add or remove servers, restrict or unrestricted membership, and split and merge CTNs).
- Assign the roles of Preferred, Backup, and Current Time Servers (CTSs), and the Arbiter.
- Adjust the time by up to +/- 60 seconds.

> **Note:** Adjustments take approximately 8 hours to steer out 1 second, so steering out 60 seconds takes a little less than 20 days.

- Schedule changes to offsets that were previously listed (time zone, DST, and leap seconds). STP can automatically schedule DST based on the selected time zone.
- Monitor the status of the CTN.
- Monitor the status of the coupling links that are initialized for STP message exchanges.

HMC web services APIs

The following STP commands are still supported to manage the CTN or to perform automated recovery actions, such as reassigning the PTS, BTS, or Arbiter role in 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 more information, review *Hardware Management Console Web Services API*, SC27-2642.

3

CTN configuration and operations

This chapter describes configuration operations and operational management task for the Coordinated Timing Network (CTN).

The following tasks are covered in this chapter:

- 3.1, “Setup new CTN” on page 42
- 3.2, “Adding systems to CTN” on page 51
- 3.3, “Modify assigned server roles” on page 54
- 3.4, “Remove systems from CTN” on page 57
- 3.5, “CTN membership restriction” on page 61
- 3.6, “Export CTN configuration data to file” on page 64
- 3.7, “Deconfiguring the CTN” on page 64

3.1 Setup new CTN

This wizard guides you through the process of setting up a CTN for the first time.

Important: After you configure a CTN, the only way to set the date and time is to deconfigure the CTN and then create and initialize it again. For more information about how to deconfigure the CTN, see 3.7, “Deconfiguring the CTN” on page 64.

Before you attempt this task, make sure that you complete the planning activities. Planning activities include the following items:

- External time source connectivity
- Server role assignment and type of CTN (restricted or unrestricted)
- Coupling configuration and link validation for STP use on all servers (short-range or long-range coupling, DWDM, number of links, and so on)

This scenario explains how to configure a CTN that consists of IBM Z servers (CPC[1]s) of various generations: one IBM z17, one IBM z16, and one IBM z15.

In preparation for this task, the external time source is configured for two systems that are designated as the primary and backup time servers (PTS and BTS). Configuration of the external time source is described in detail n Chapter 4, “External Time Source Configuration and Management” on page 67.

Coupling links between all systems that are members of the new CTN are also configured and tested.

The initial CTN configuration assumes that none of the existing servers that you add to the new CTN is part of another CTN configuration.

Before you start the configuration, you can check the STP Information in the **System Details** task, as shown in Figure 3-1.

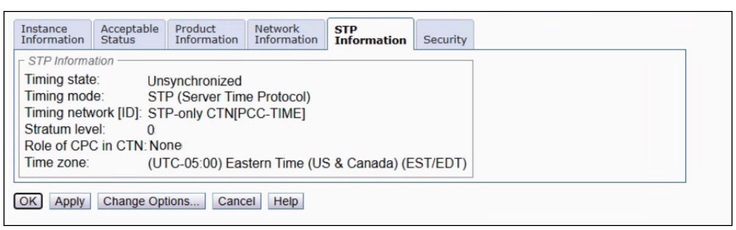

Figure 3-1 STP Information

Use the following steps to configure the CTN:

1. To initiate the configuration, access the **Manage System Time** task on your HMC, as shown in Figure 3-2 on page 43.

[1] CPC - Central Processing Complex (an IBM Z system with CPU and I/O)

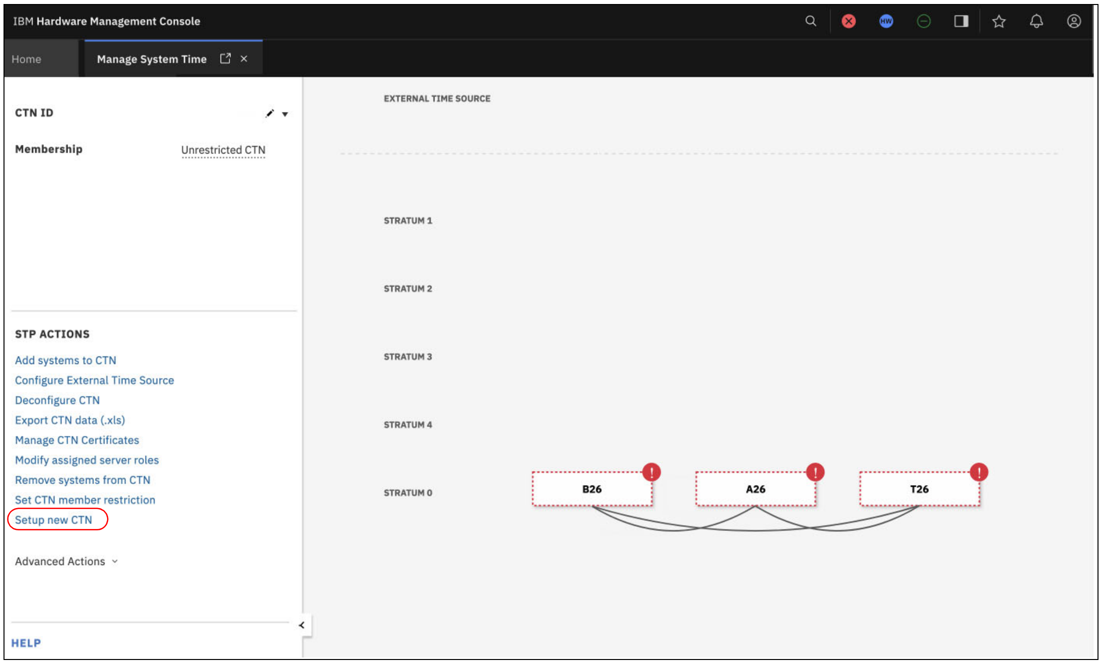

Figure 3-2 Manage System Time task - servers with no CTN membership

2. From the left menu, select “Setup new CTN” to initiate the configuration, as shown in Figure 3-3.

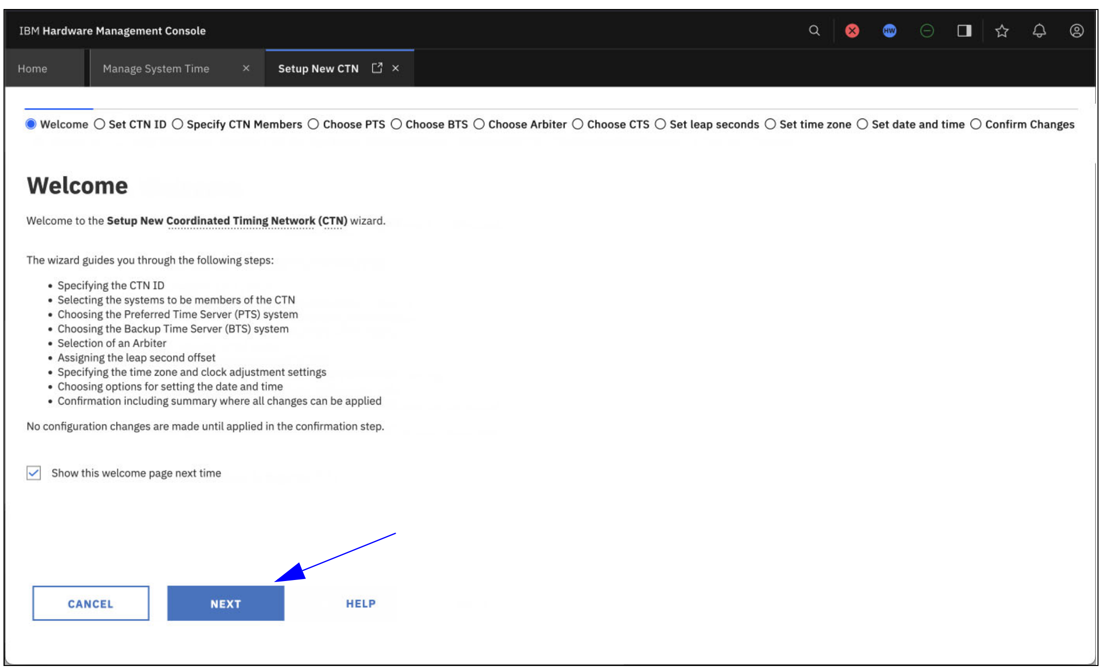

Figure 3-3 Setup new CTN - Welcome window

3. Follow the instructions presented in the window and provide the CTN ID, as in Figure 3-4 on page 44.

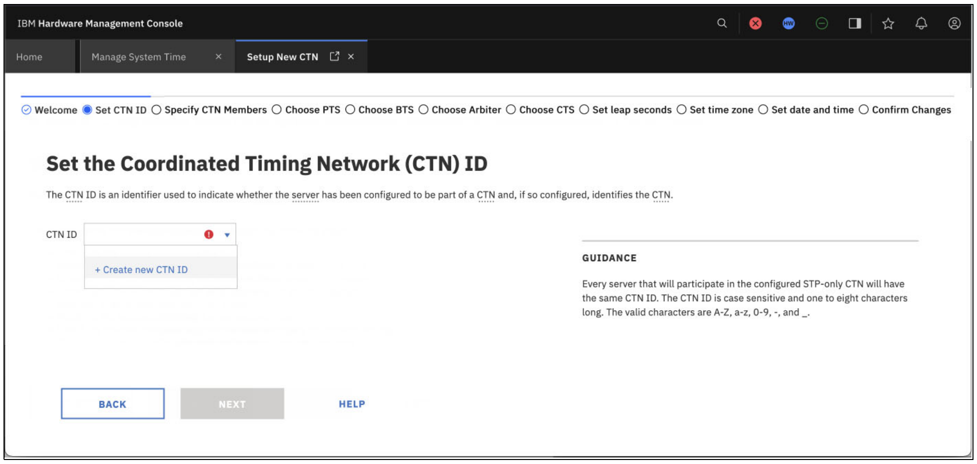

Figure 3-4 Entering the CTN ID

4. The initial CTN ID is empty (unless another CTN has been previously configured and then removed), so you must enter the CTN ID (name) and click **Save**, as shown in Figure 3-5.

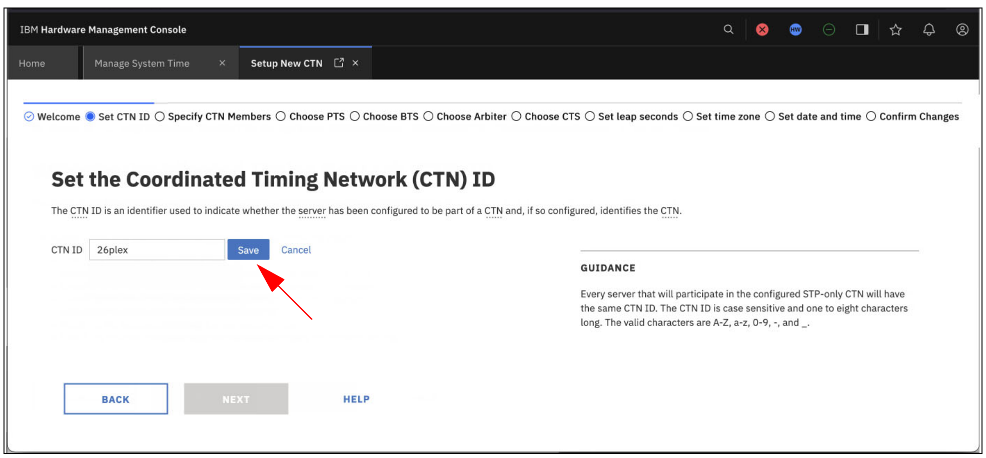

Figure 3-5 Saving the CTN ID

5. After you save the CTN ID (name), click **Next** as in Figure 3-6.

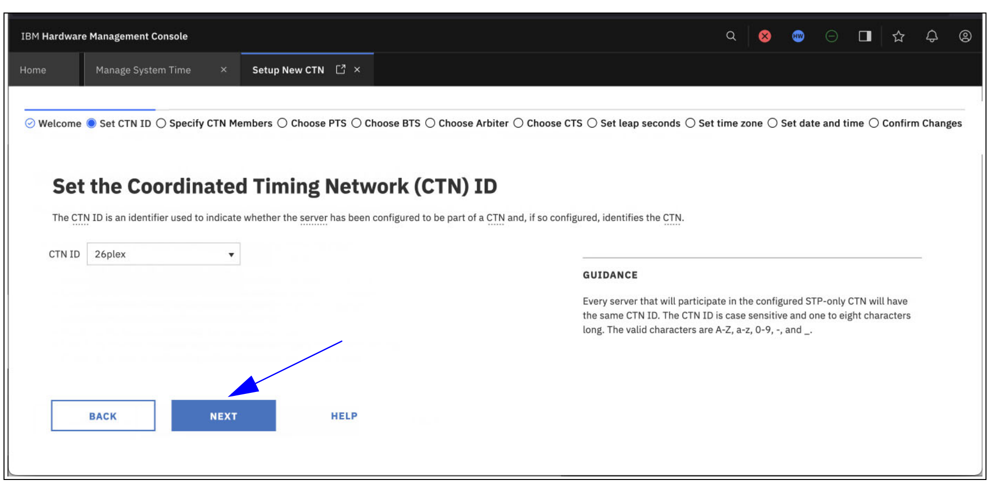

Figure 3-6 CTN ID set

The potential members of the newly defined CTN are shown in Figure 3-7.

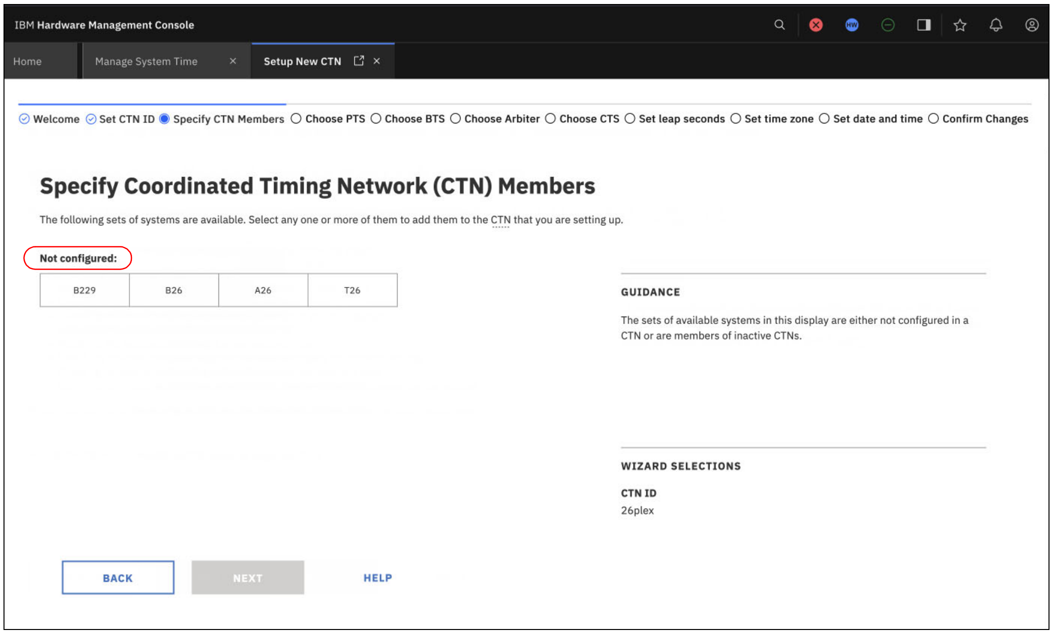

Figure 3-7 Servers available for CTN membership

6. Click each server to add it as a member. The example uses all four servers, as in Figure 3-8. Click **Next**.

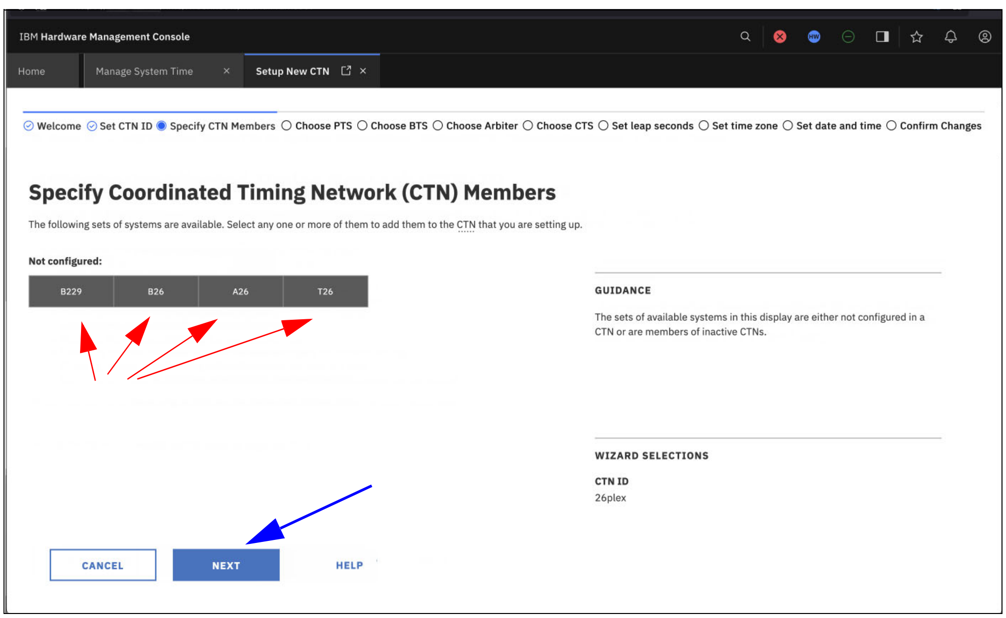

Figure 3-8 Selecting CTN members

7. In the following steps, configure server roles (Preferred Time Server, Backup Time Server, and Arbiter). The first server role that is selected is the Preferred Time Server and is shown in Figure 3-9 on page 46. After selecting the server, click **Next**.

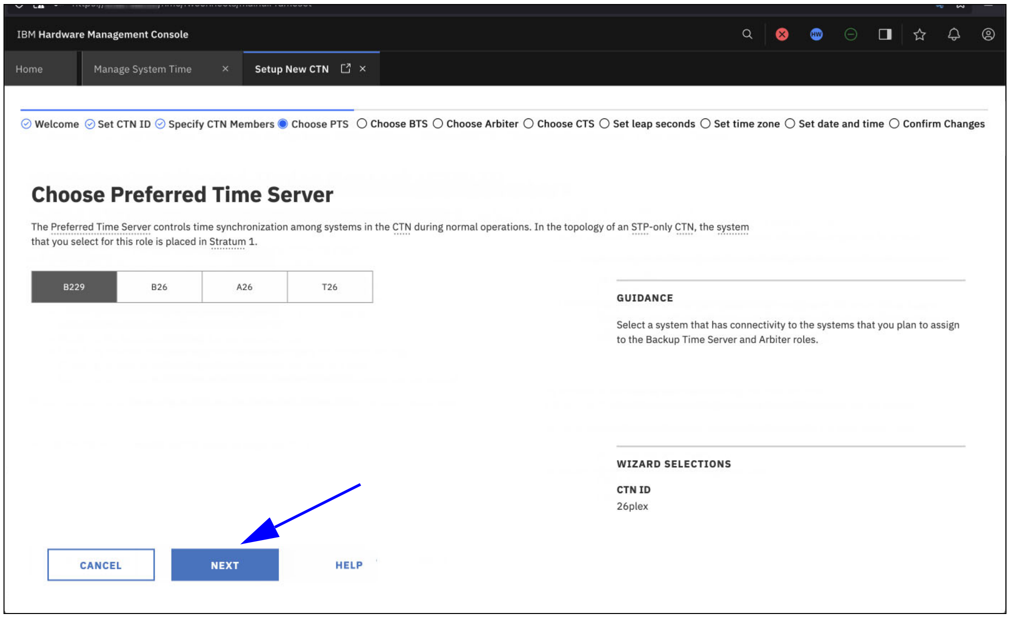

Figure 3-9 Select the Preferred Time Server (PTS)

8. The Backup Time Server and Arbiter are optional. Considerations of when to define these optional roles is provided in 2.2.2, “CTN roles” on page 12. The example uses both the BTS and the Arbiter.
9. Select the Current Time Server, as shown in Figure 3-10.

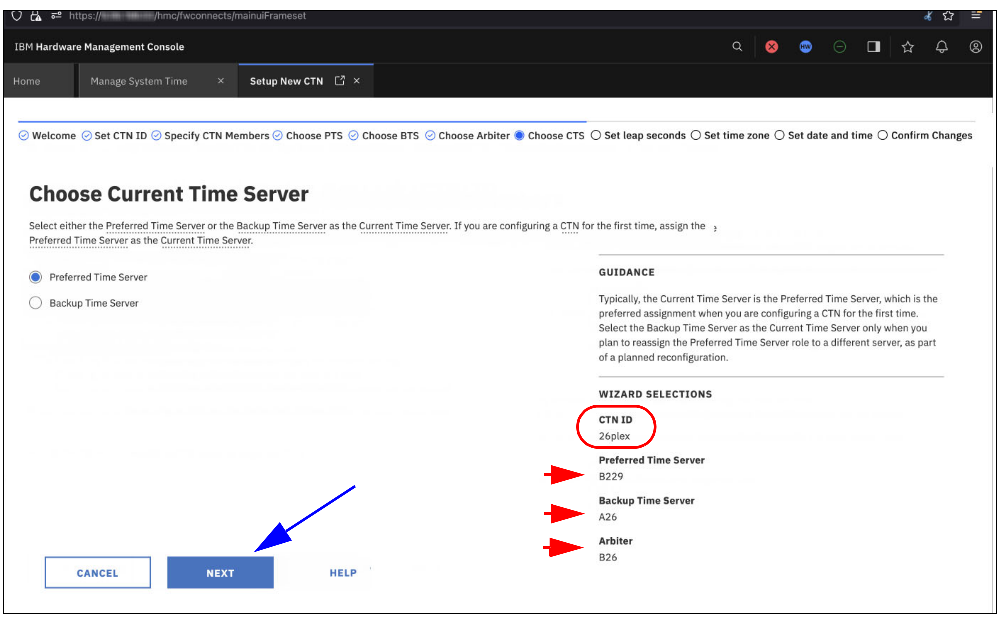

Figure 3-10 Choose the Current Time Server

10. After you select the Current Time Server, you select the leap seconds (if your organization observes leap seconds), as shown in Figure 3-11 on page 47. Click **Next**.

> **Important:** The value that is inputted is the current TAI-UTC leap second value minus 10 seconds. For example. the TAI-UTC at the time of this writing was 37, so use the value ‘**27**’.

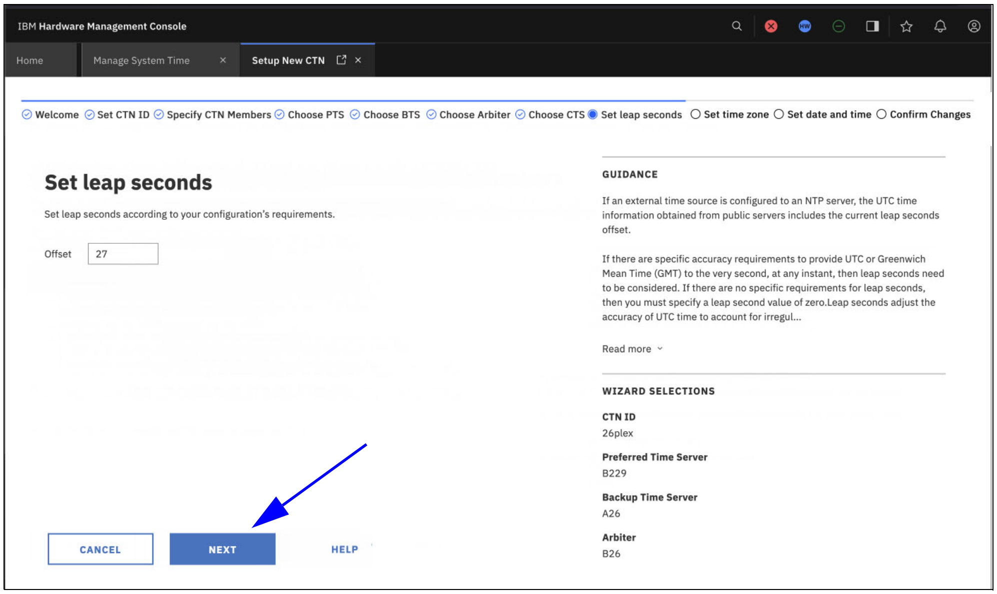

Figure 3-11

11. Configure the **Time Zone** (TZ) information and **Daylight Saving Time** (DST), if applicable, as shown in Figure 3-12. Click **Next**.

> **Note:** It is also possible to configure your own preference for TZ and DST by selecting in the pull-down menu User Defined Time Zones.

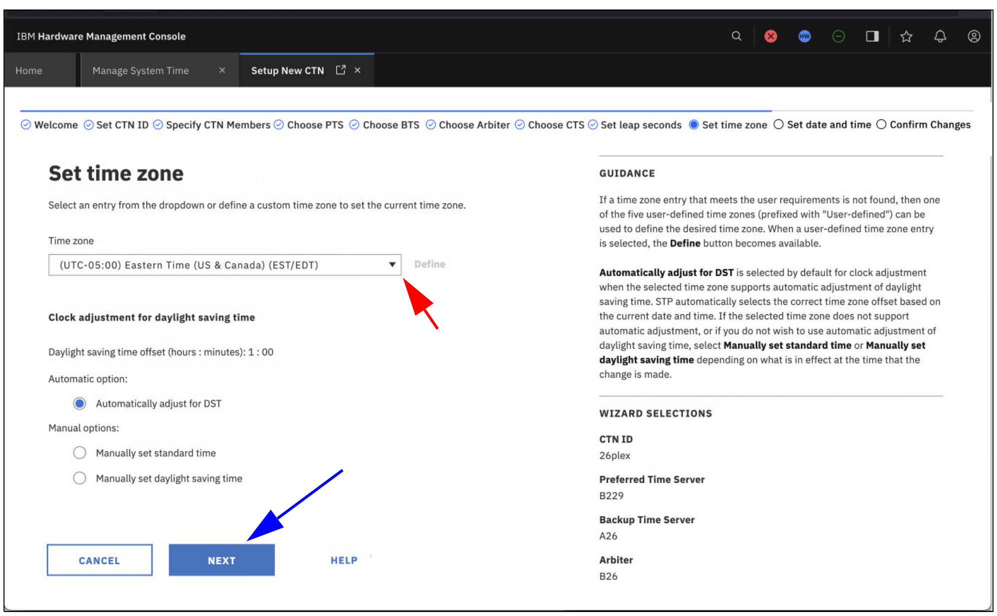

Figure 3-12 Time Zone and Daylight Saving Time

> **Important:** After a CTN is configured, the only way to set the date and time is to deconfigure the CTN and then create the CTN and initialize the time again. For more information on how to deconfigure the CTN, see 3.7, "Deconfiguring the CTN" on page 64.

12.After setting the TZ information, configure the Date and Time, as in Figure 3-13. The recommended option for setting the date and time is to use an External Time Source (ETS) that is either NTP or PTP. Click **Next**.

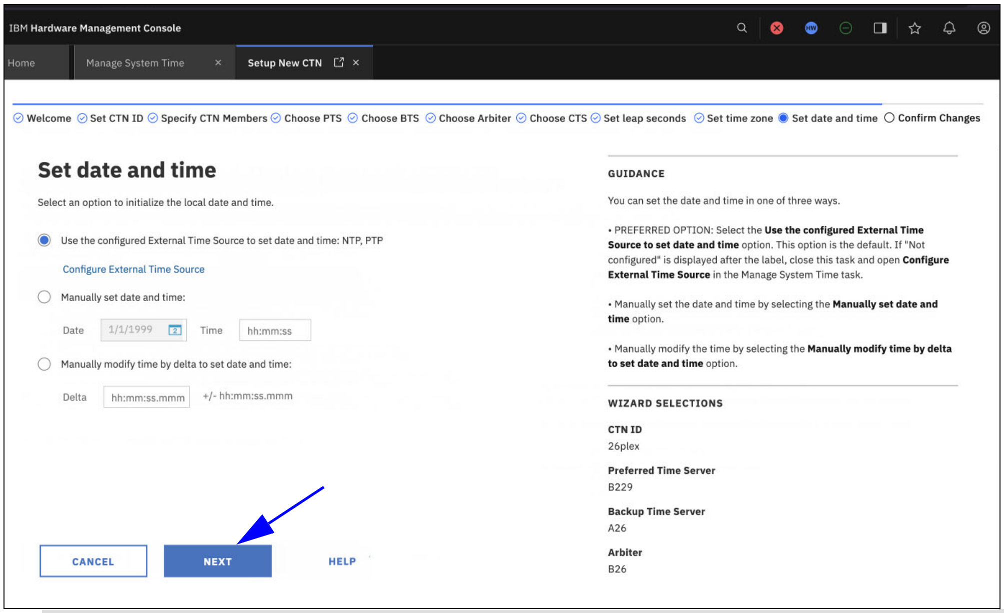

Note:

Figure 3-13 Set Date and Time

A final confirmation window is displayed as in Figure 3-14 on page 49.

Note: The CTN membership restriction is available for single- or dual-server CTNs.

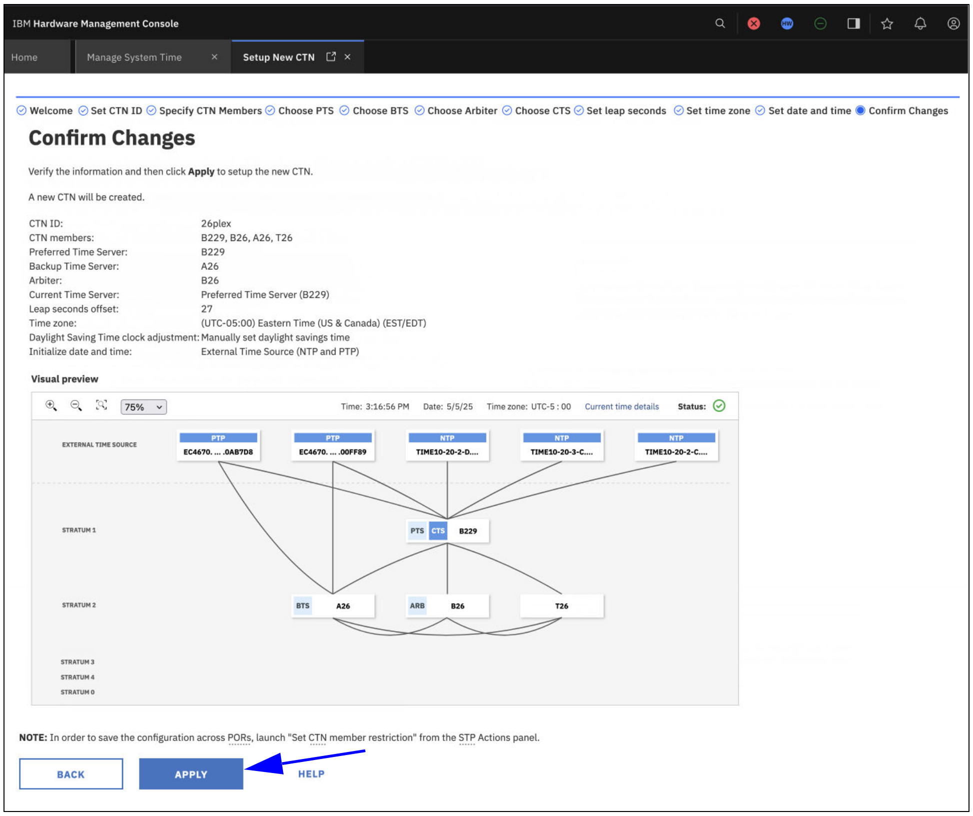

Figure 3-14 CTN confirmation step

13. After reviewing the proposed configuration, click **Apply** and wait for the configuration to finish, as shown in Figure 3-15.

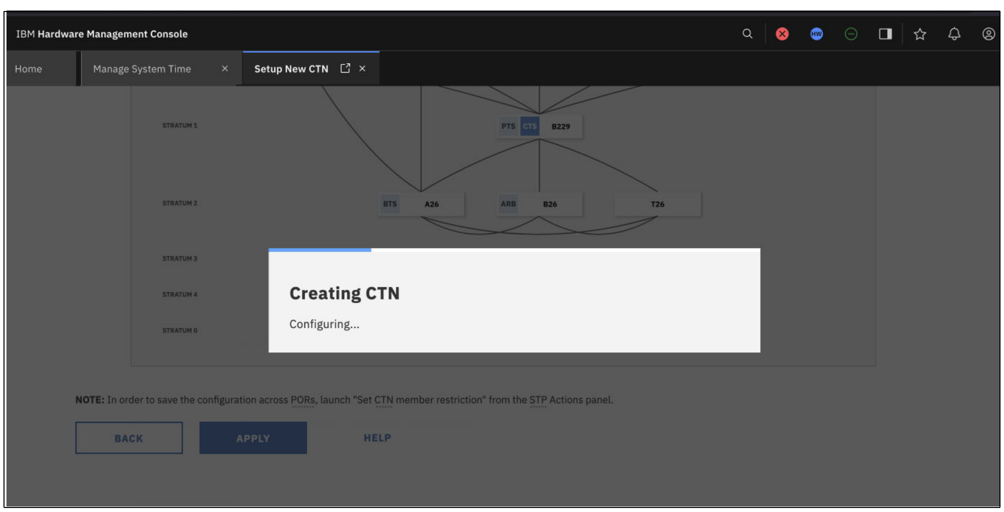

Figure 3-15 CTN configuration progress

14. When the configuration is finalized, the successful CTN setup message is shown as in Figure 3-16 on page 50. Click **Close**.

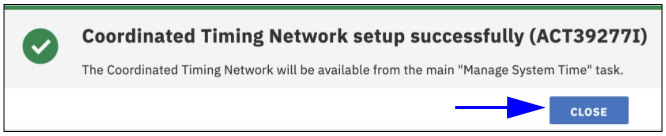

Figure 3-16 Successful CTN configuration message

15. View the final configuration, as shown in Figure 3-17

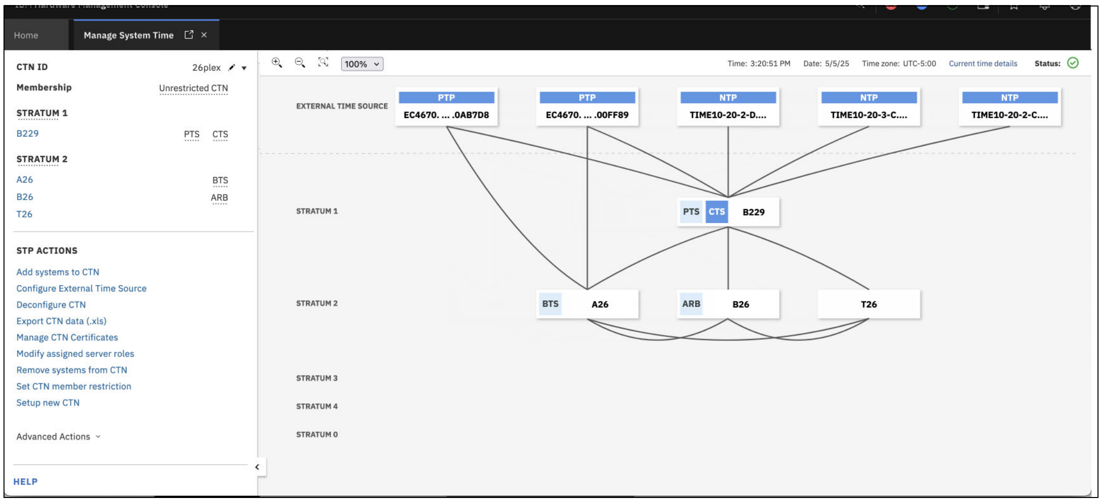

Figure 3-17 CTN configuration shown in the Manage System Time tab

3.1.1 Identify active STP links

Important: Information about active STP links is needed when maintaining the servers in the CTN (that is, servicing a server, such as replacing and adding drawers and parts).

After you configure the CTN, in the Manage Time Server tab you can hover the mouse cursor over the links to check for additional information, as in Figure 3-18.

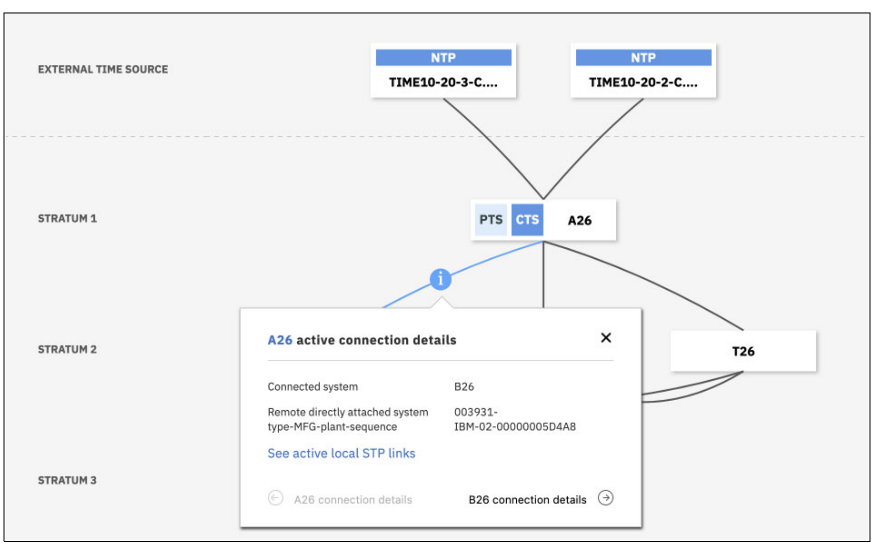

Figure 3-18 Link information

One important aspect of the CTN is to observe which coupling links are actively used for transmitting time. You can view this information by clicking the **See active local STP** link. The Active STP links information is shown in Figure 3-19.

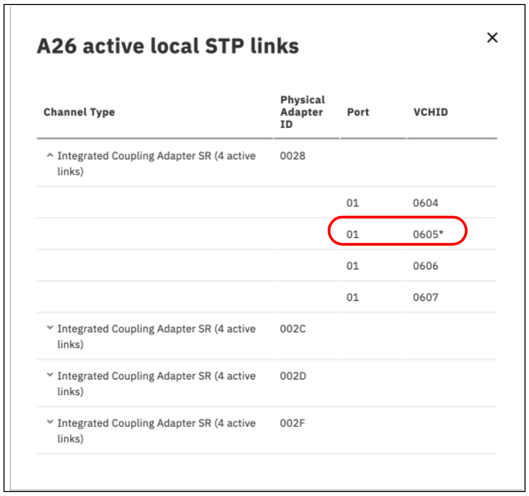

Figure 3-19 Active STP links

3.2 Adding systems to CTN

Use this task to add a new system to an existing CTN. The starting point is a three-server CTN shown in Figure 3-20.

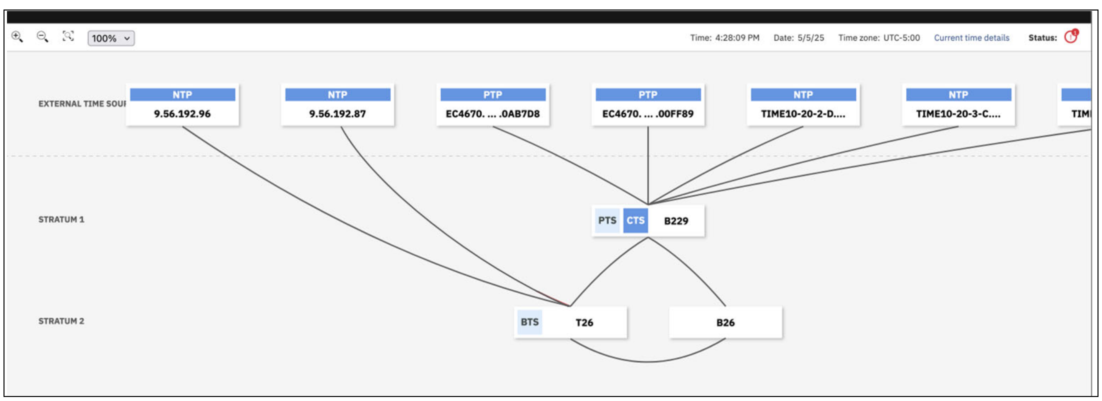

Figure 3-20 Starting pint CTN configuration

1. From the left menu, click **Add systems to CTN** and proceed to the Welcome screen shown in Figure 3-21 on page 52
2. Click **NEXT** to select the systems to add to the CTN, as shown in Figure 3-22 on page 52.

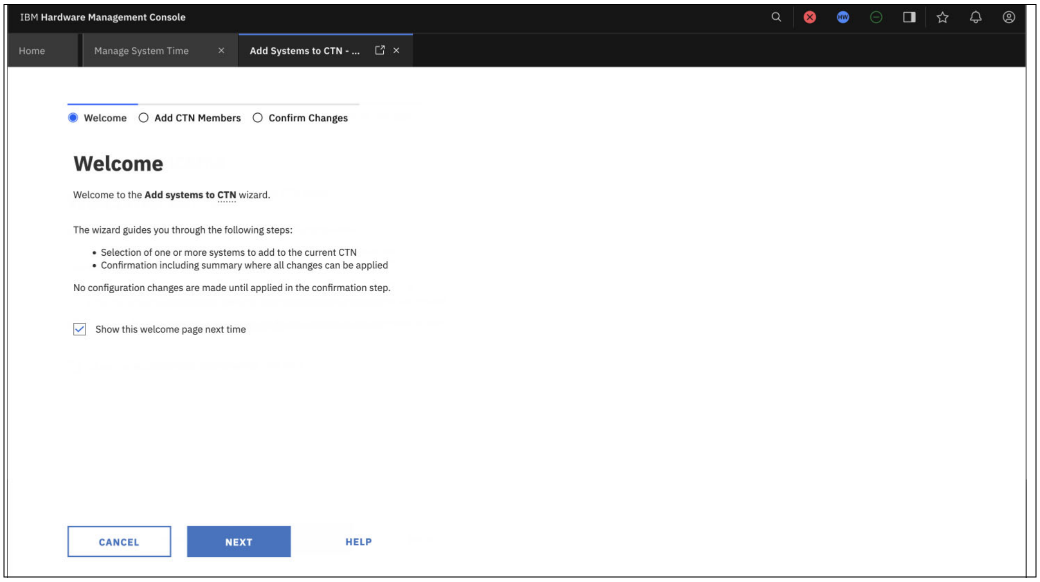

Figure 3-21 Adding systems to CTN

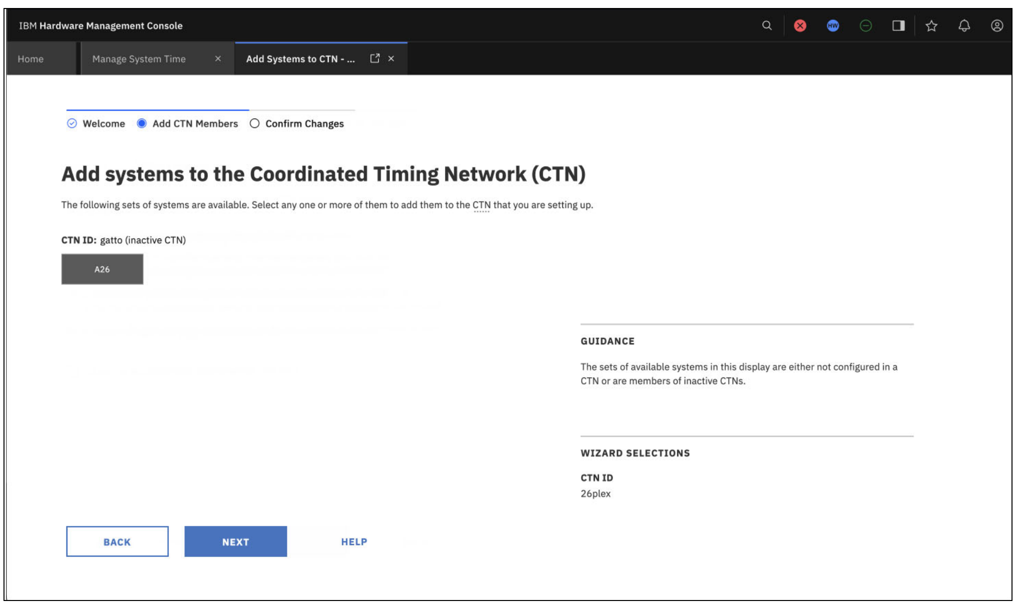

Figure 3-22 Select systems to add to CTN

3. After selecting the system(s), click **APPLY** and the confirmation menu is displayed, as shown in Figure 3-23.

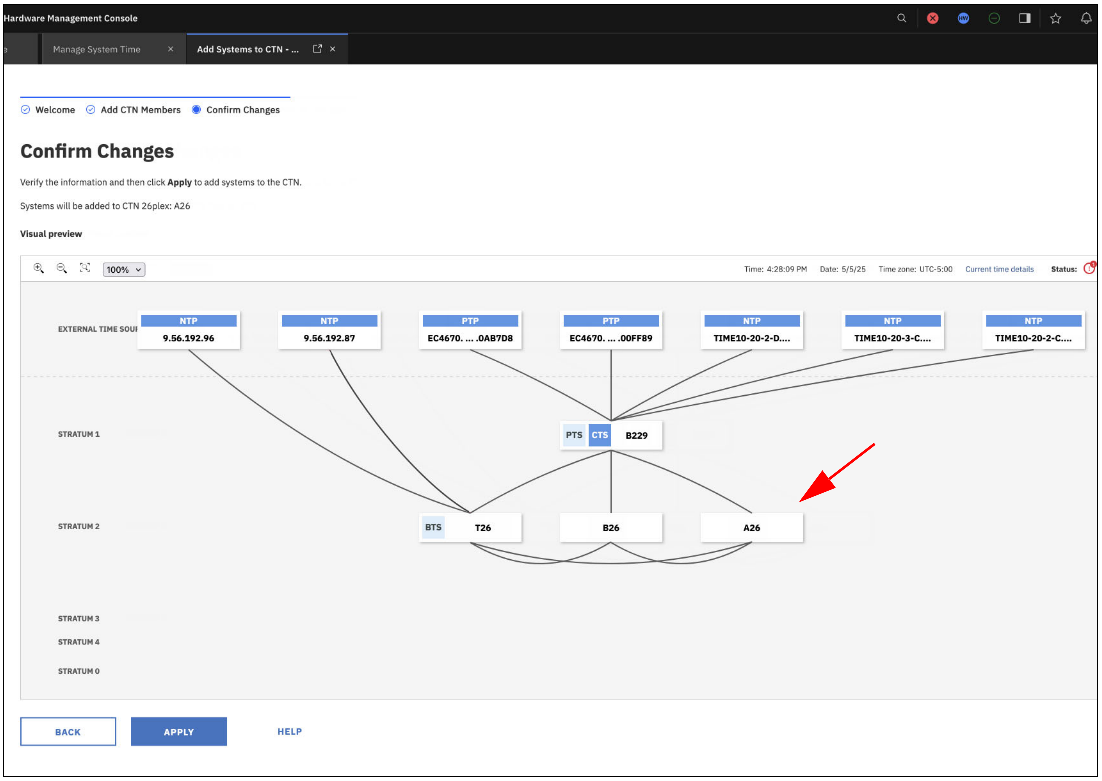

Figure 3-23

4. To confirm changes click **APPLY**. The message shown in Figure 3-24 is displayed because the system to be added to the CTN ID “26plex” was previously configured with a local CTN ID. This is not an error; it is only an informational message.

Figure 3-24 Confirmation message

5. Click **APPLY** to start the operation. The progress window is displayed (Figure 3-25).

Figure 3-25 Operation in progress message

6. The message that is shown in Figure 3-26 is displayed upon successful addition of new system(s) to the CTN.

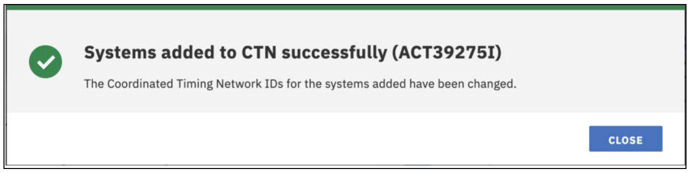

Figure 3-26 Systems successfully added to the CTN

7. At the end of this change, the configuration is shown, as in Figure 3-27.

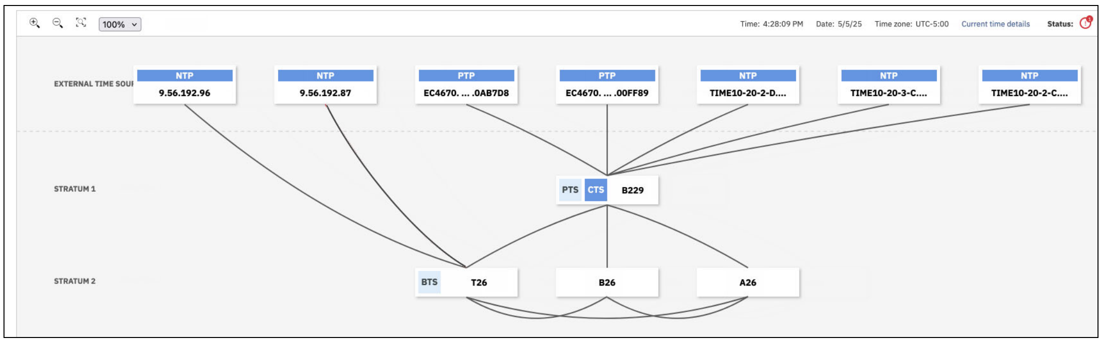

Figure 3-27 CTN configuration after server addition

Observe that the new server does not have a role assigned to it. You can change the server roles by selecting the task from the left menu **Modify assigned server roles**.

3.3 Modify assigned server roles

This task is used to change the Primary Time Server, Backup Time Server, and Arbiter roles.

The starting point of this operation is the CTN displayed in Figure 3-27. Use the following steps to modify the server roles:

1. To modify server roles, select the **Modify assigned server role** task on the left menu of the Manage System Time tab (see Figure 3-17 on page 50).

   The Welcome menu in the Modify assigned server roles tab is displayed, as shown in Figure 3-28 on page 55.

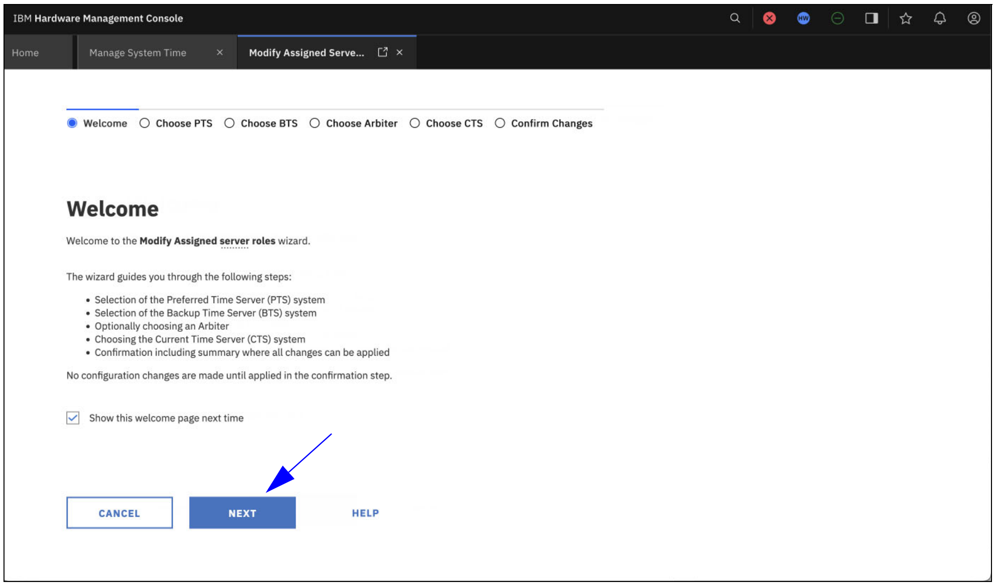

Figure 3-28 Welcome window

2. Click **NEXT** to select the Preferred Time Server, Backup Time Server, and Arbiter assigned roles, as in Figure 3-29.

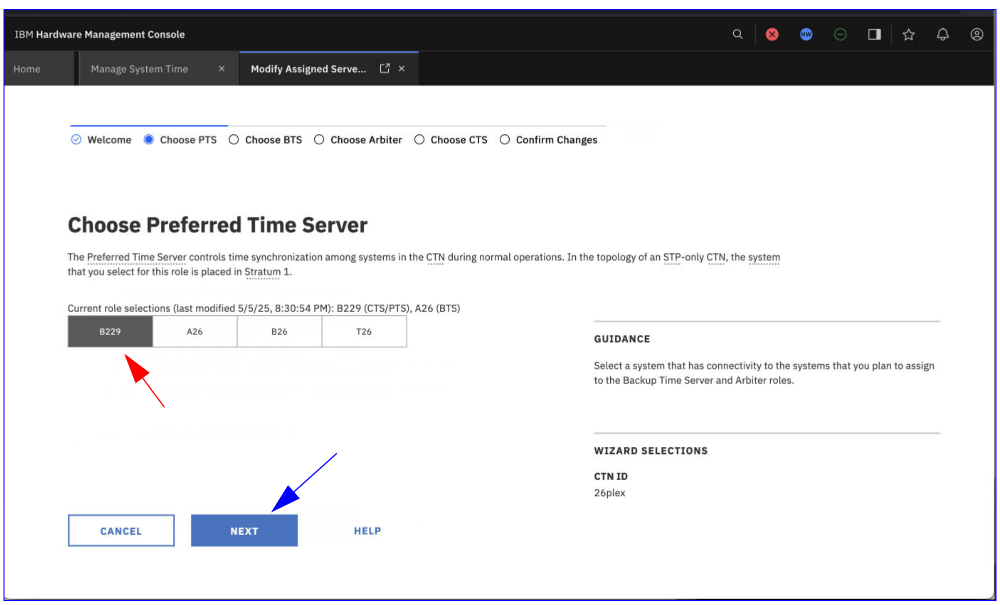

Figure 3-29 Selecting assigned server roles (PTS, BTS, Arbiter)

3. Assigning the Backup Time Server and Arbiter is optional. In the example, both roles are assigned.

   Repeat the step for the Backup Time Server and Arbiter, then click **NEXT**.

4. Select the Current Time Server, as shown in Figure 3-30 on page 56.

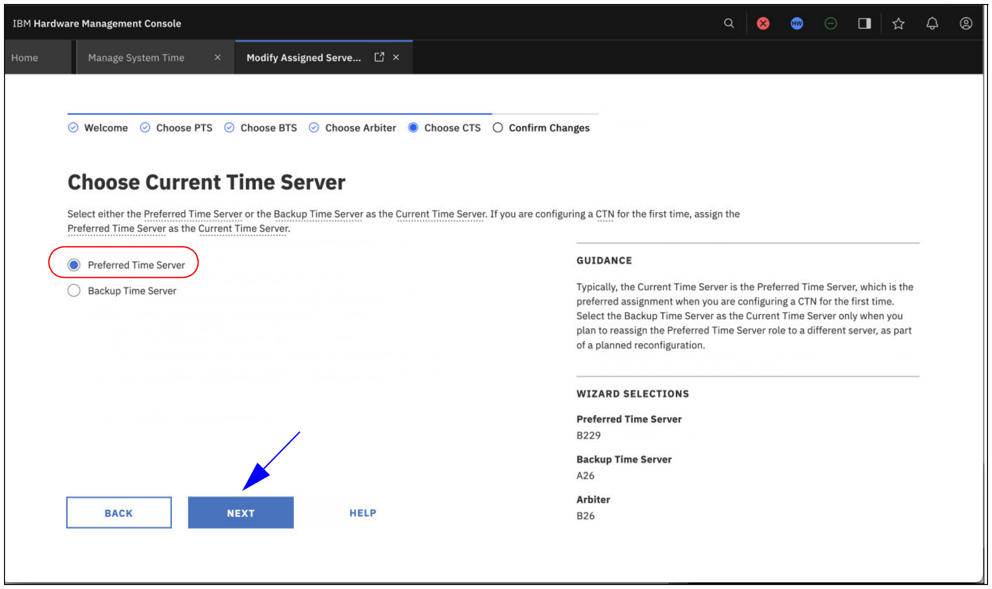

Figure 3-30 Selecting server for the Current Time Server

5. Click **NEXT** to review the changes before applying the new server roles, as shown in Figure 3-31.

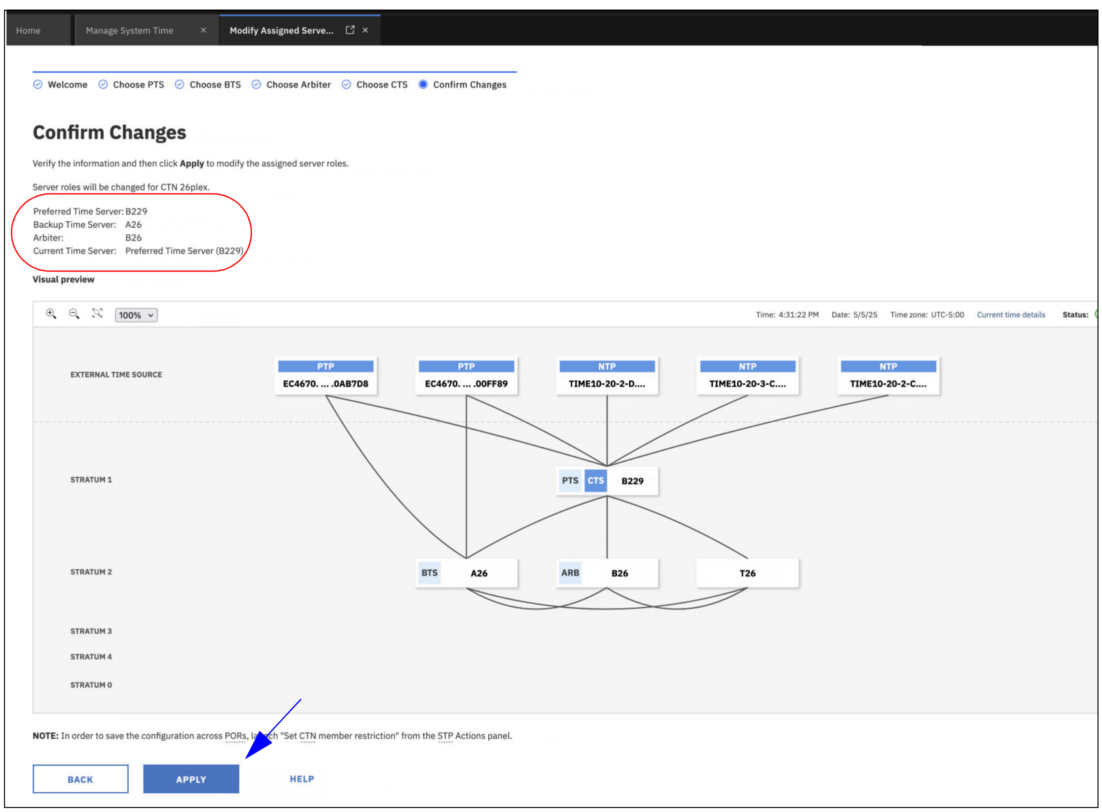

Figure 3-31 Confirm configuration and APPLY

6. Click **APPLY** to launch the task.

   The progress message shown in Figure 3-32 on page 57 is displayed.

Figure 3-32 Modify server roles task progress

7. Upon successful completion, the message shown in Figure 3-33 is displayed. Click **CLOSE** to finalize the task.

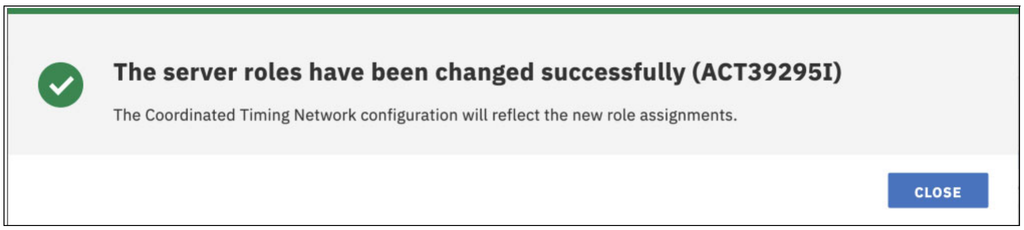

Figure 3-33 Modify server roles task success

The CTN configuration with the modified assigned server roles is shown in Manage System Time tab (see Figure 3-34).

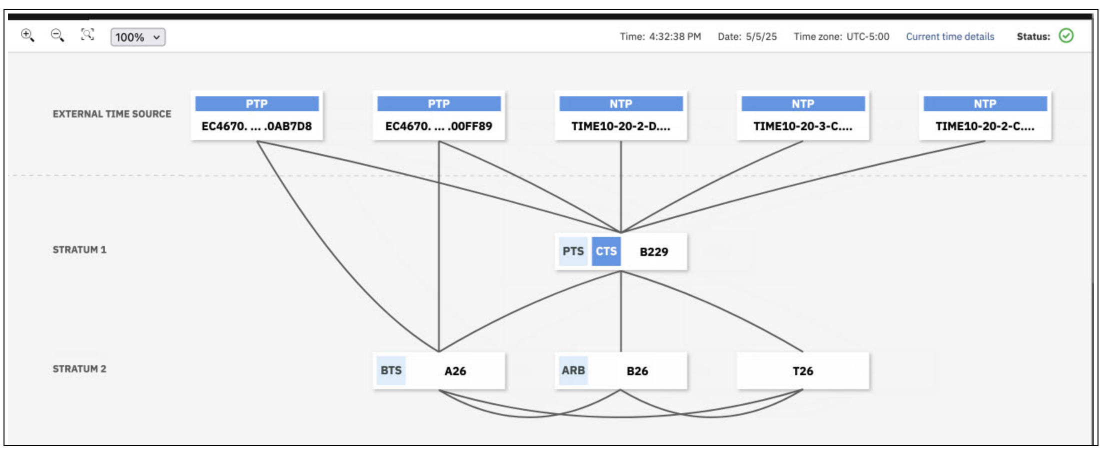

Figure 3-34 Configuration after modifying server roles

3.4 Remove systems from CTN

This task is used to remove one or more systems from a CTN. This task might be required if you plan to decommission servers or re-purpose servers.

Before removing a system from a CTN, you must un-assign any role (PTS/BTS/Arbiter) from the respective system. This can be accomplished through the Modify assigned server role task.

It is not possible to remove systems that have an assigned CTN role by using the Manage System Time task.

Important: If you plan to remove a server from the CTN and re-use it in another CTN, remove it from the current CTN before powering off the system and IML-ing it.

The task is to remove the A26 system from the “26plex” CTN.

The starting point CTN configuration (Figure 3-35) consists of four servers, B229 (PTS/CTS), T26 (BTS) and B26 and A26 (no assigned role). Use the following steps to remove A26:

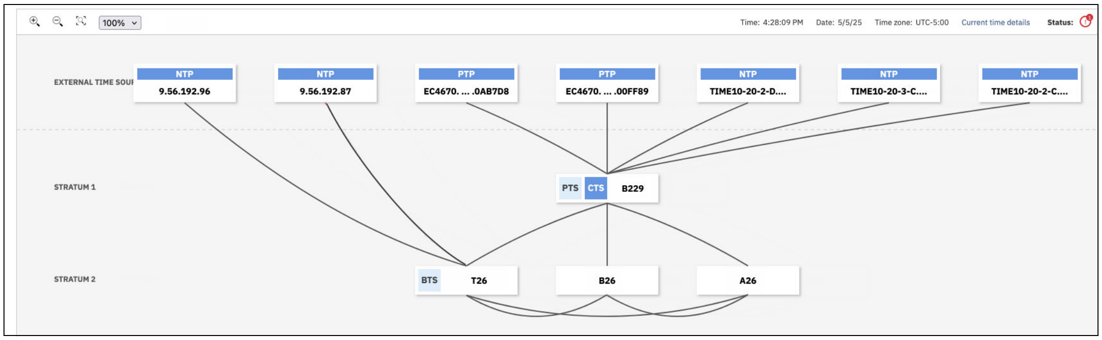

Figure 3-35 Starting point CTN configuration

1. Start the Remove systems from CTN task and proceed to the Welcome menu shown in Figure 3-36.

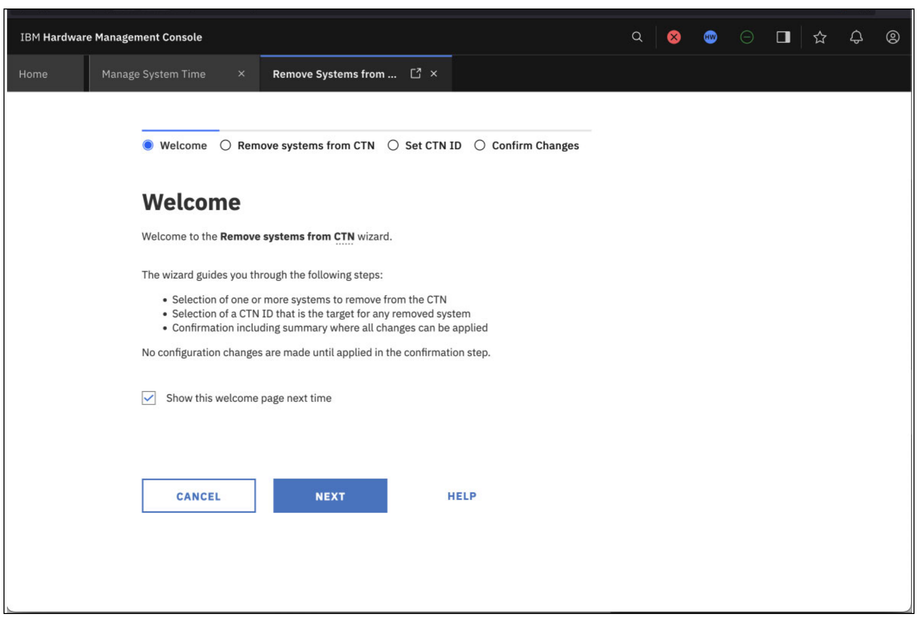

Figure 3-36 Remove systems from CTN task Welcome menu

2. Select the servers to be removed from the CTN as shown in Figure 3-37 on page 59.

   Servers that have an assigned role (PTS, BTS, Arbiter) cannot be removed from the CTN. Before you remove a server, make sure that you un-assign any role from it.

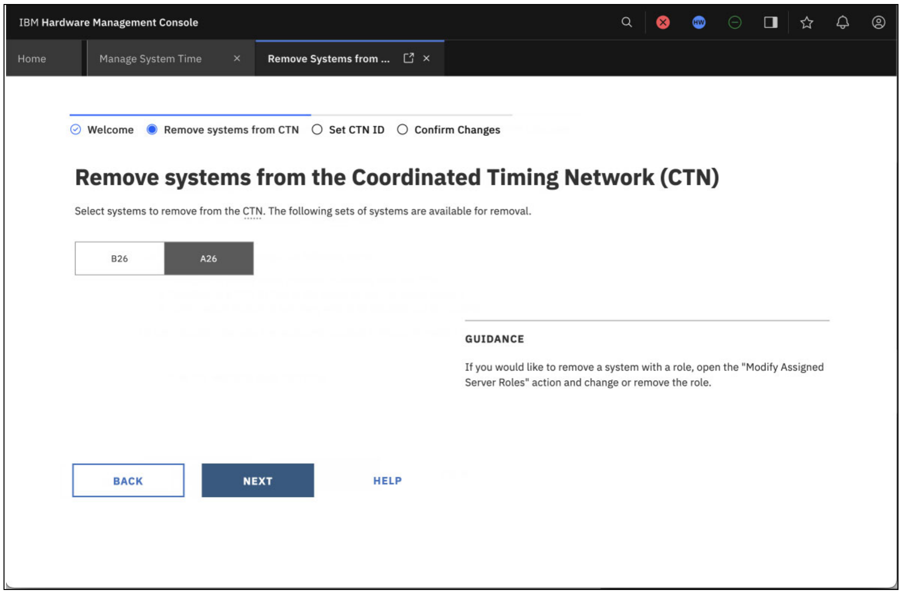

Figure 3-37 Server selection

3. The server to be removed from the current CTN must have a destination CTN, even if that CTN is not configured or does not hold any other server.

   Assign a new CTN ID for the server to be removed, as shown in Figure 3-38. Click **NEXT**.

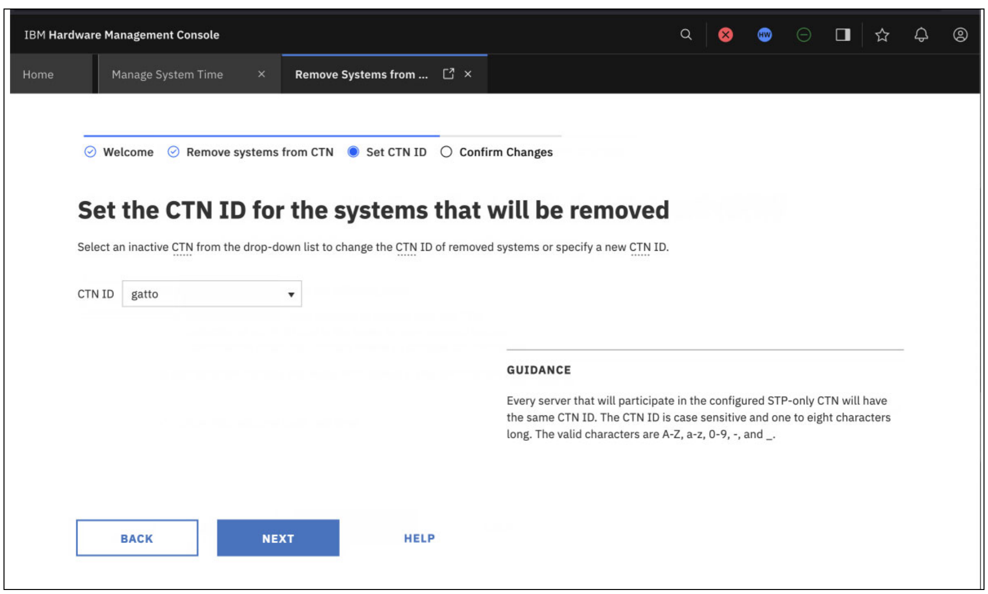

Figure 3-38 Assigning the destination CTN ID for the removed server

4. The “Confirm Changes” menu is shown as in Figure 3-39 on page 60.

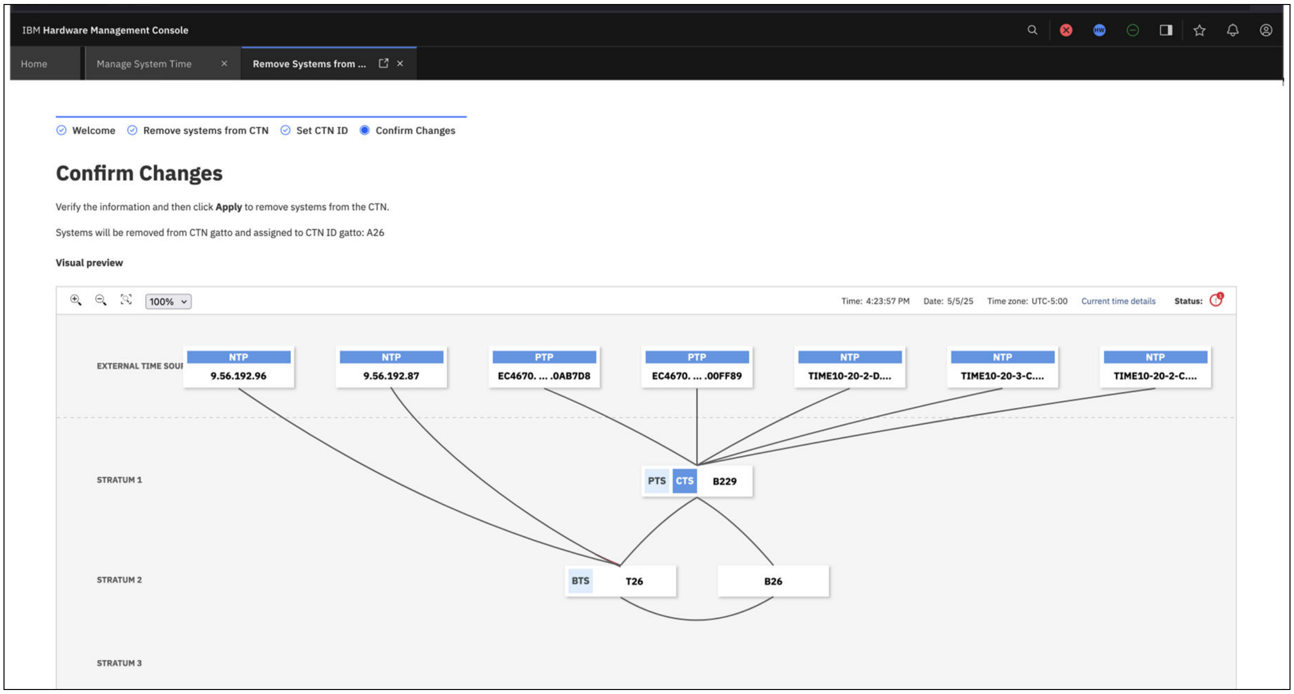

Figure 3-39 Confirm Changes menu

5. Scroll down and click **APPLY** to initiate the operation.
6. Wait for successful task completion message shown in Figure 3-40

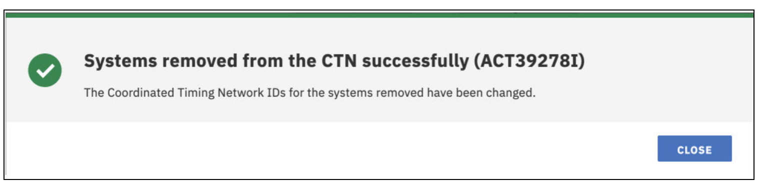

Figure 3-40 Successful removal of server(s) from CTN

The resulting CTN configuration is shown in Figure 3-41.

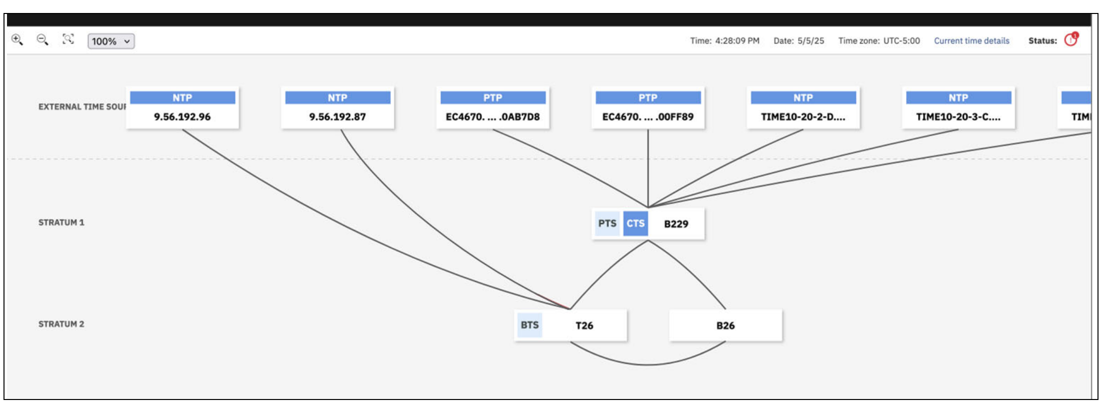

Figure 3-41 Configuration after A26 removal

3.5 CTN membership restriction

This is an optional task that is used to set a CTN with membership restriction. Setting a CTN with membership restriction preserves the configuration during and after POR/IML of all servers in the CTN. The CTN membership restriction task has the following characteristics:

- Saving the CTN configuration (restricting membership for a CTN) is possible only for single- or dual-server CTN configurations.
- In a restricted membership CTN, all members must have a role assigned (PTS for single-server CTN, and PTS and BTS for dual-server CTN).
- For an unrestricted CTN, the configuration of the CTN is maintained if there is one active CTN member that acts as the CTS role.
- If all servers are POR/IML-ed, CTN configuration must be reentered from scratch.

This scenario includes a dual-server CTN that is configured with two servers: A26, which is the PTS/CTS and B26, which has no role. The configuration is shown in Figure 3-42.

When CTN membership is restricted, for any server changes in the CTN, the CTN must be first unrestricted. After it is unrestricted, you can make the required changes, then reapply CTN membership restriction if the resulting configuration change is still valid for restricted CTN membership. This includes the following server changes:

- Adding or removing servers to or from CTN
- Changing server roles

1. Start the CTN membership restriction task by selecting **Set CTN member restriction**.

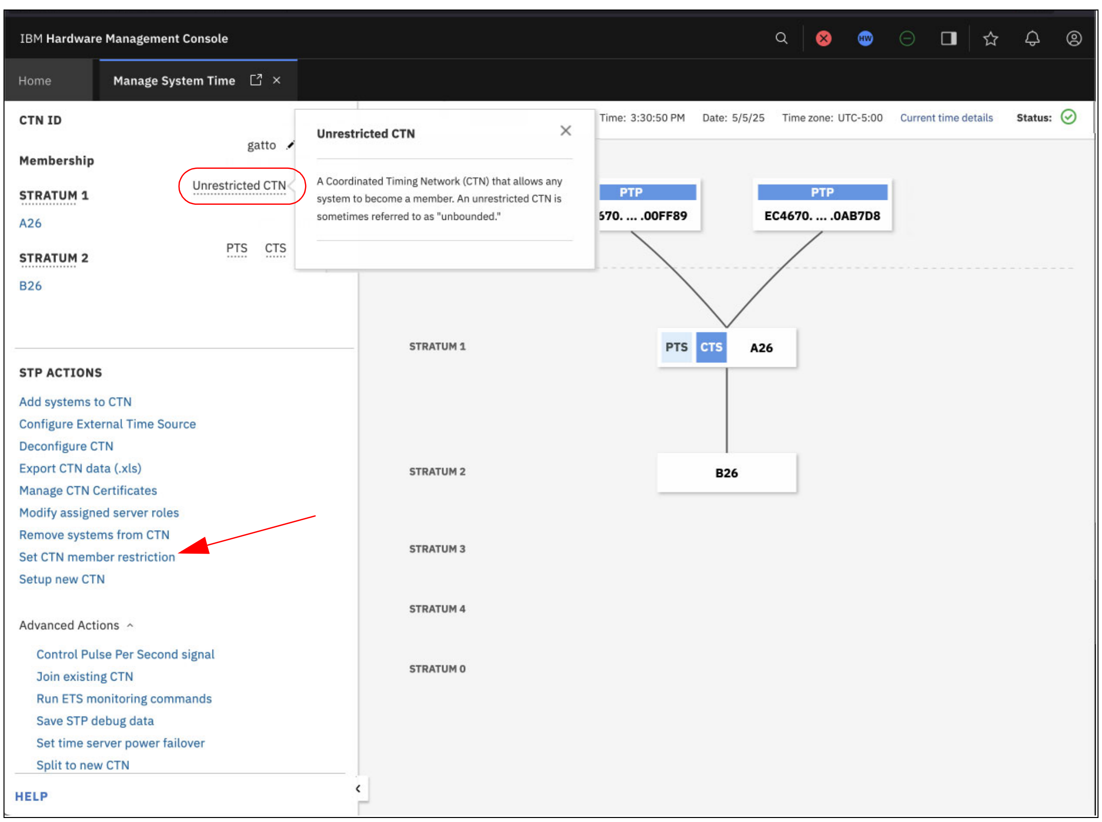

Figure 3-42 Initiating the CTN membership restriction task

2. Because no role is assigned to the second server (B26), it is not possible to restrict CTN membership for this configuration (see Figure 3-43):

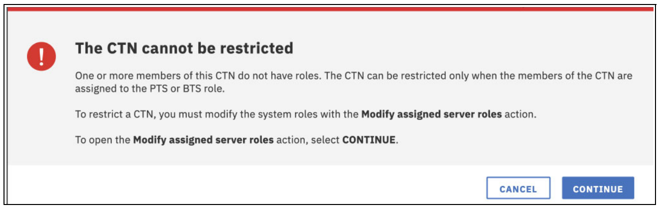

Figure 3-43 Server without an assigned role cannot be restricted

3. To proceed with CTN membership restriction task, assign the BTS role to the second server, B26. (see 3.3, “Modify assigned server roles” on page 54).

   Proceed to Restricting the CTN membership task, as shown in Figure 3-44.

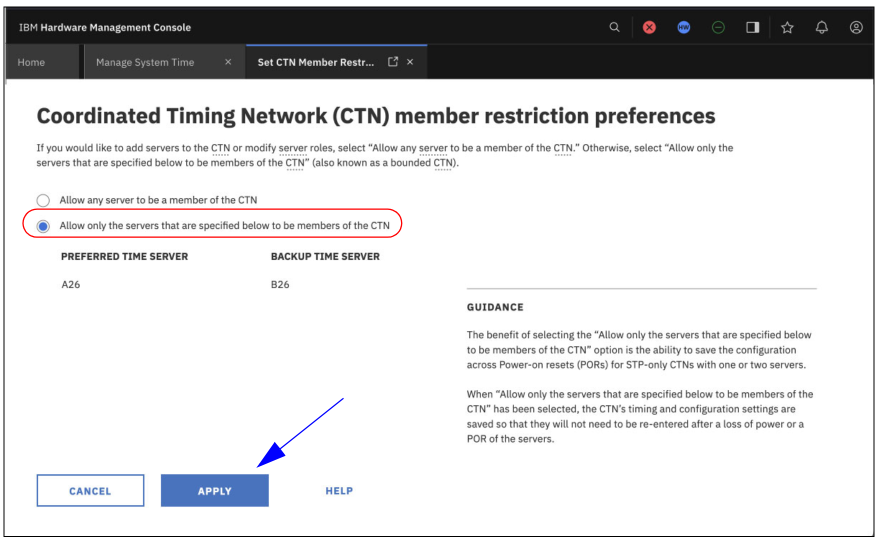

Figure 3-44 Initiating the CTN membership restriction

4. Select **Allow only the servers that are specified below to be members of the CTN** and click **APPLY**. The task progress message is displayed. Wait until configuration is complete and the restricted CTN message is displayed (see Figure 3-45).

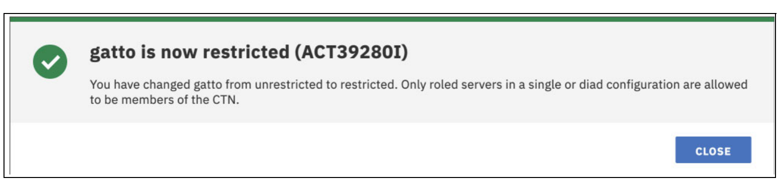

Figure 3-45 Restricting CTN membership task success.

The CTN is now restricted, as shown in Figure 3-46 on page 63.

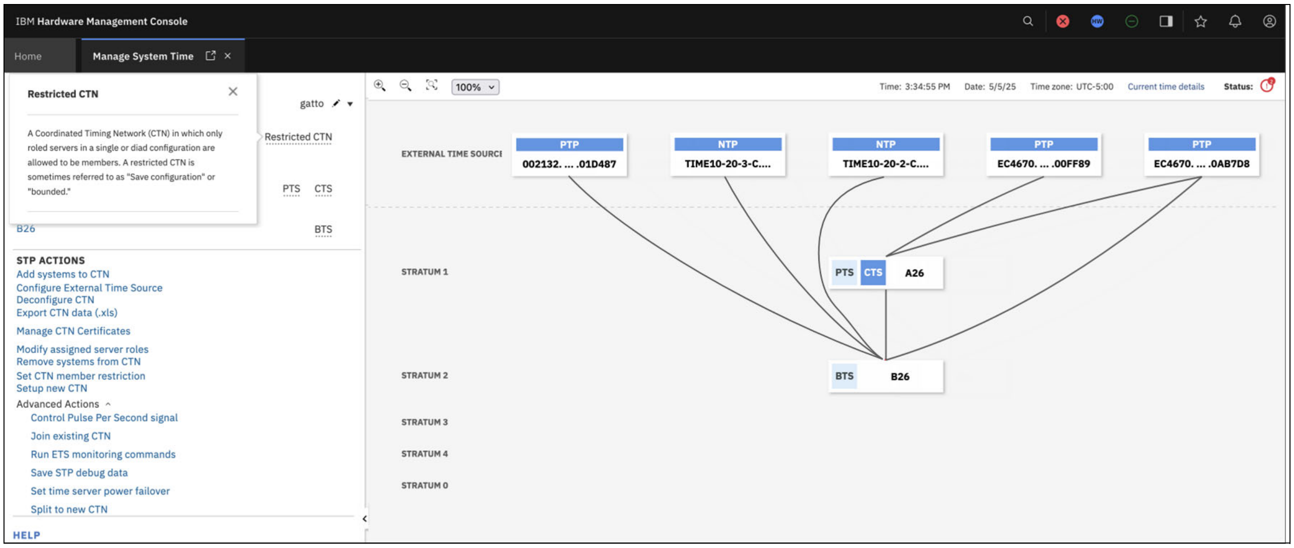

Figure 3-46 Restricted CTN

When it is restricted, a CTN cannot be configured with new members. As such, if you need to add or change/replace servers in a restricted CTN, you need to unrestrict CTN membership first.

To unrestrict a restricted CTN, use the following steps:

1. Select the CTN and then select **Set CTN member restriction** in the Manage Server Time window.
2. Select **Allow any server to be a member of the CTN**.
3. Click **APPLY** as in Figure 3-47.

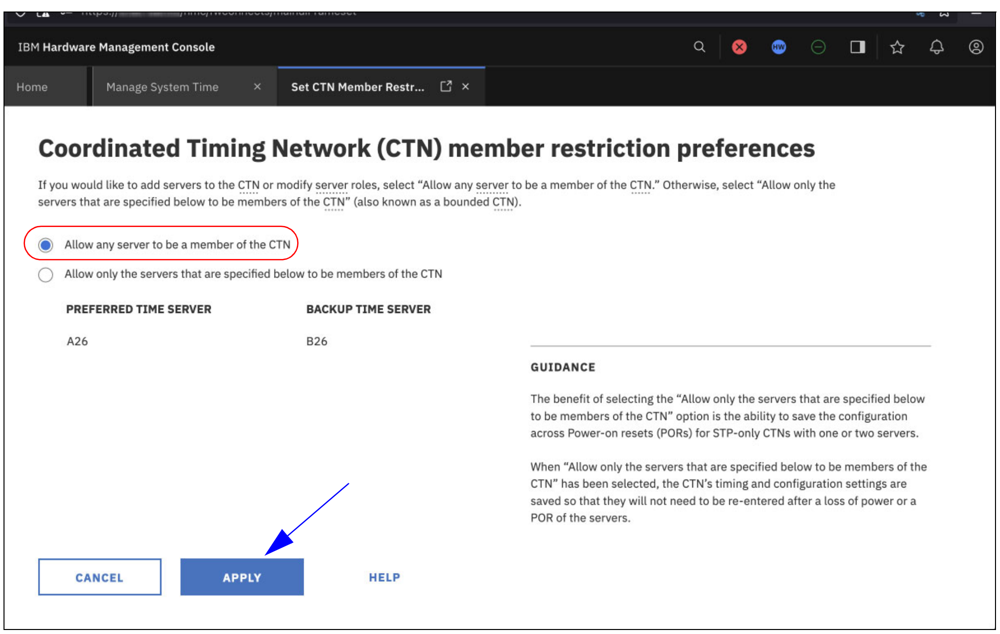

Figure 3-47 Unrestricting CTN membership

3.6 Export CTN configuration data to file

To export CTN configuration data, use the following steps:

1. Click the **Export CTN data (.xls)** menu in the **Manage System Time** task. The message is shown in Figure 3-48.

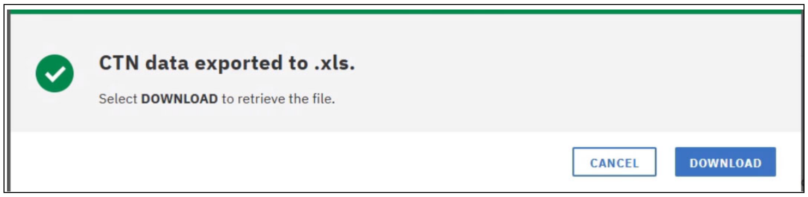

Figure 3-48 Export CTN data

2. Click **DOWNLOAD** to save the file. The export file contains information about the CTN membership, status, ETS information, configured links, and more (see Figure 3-49).

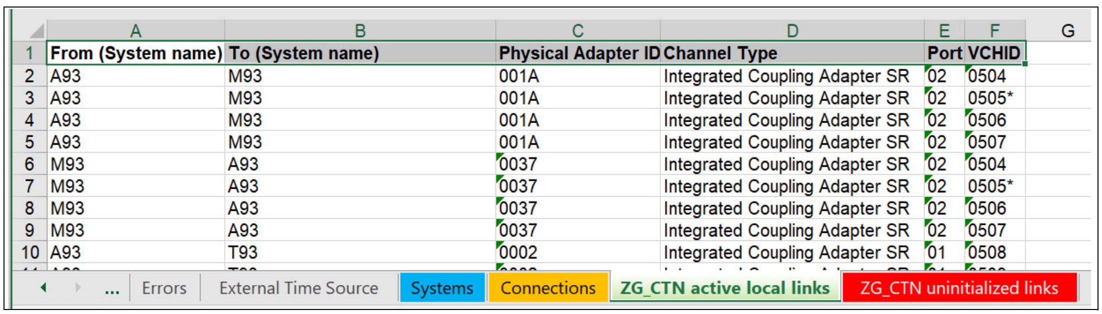

	A	B	C	D	E	F	G
1	From (System name)	To (System name)	Physical Adapter ID	Channel Type	Port	VCHID	
2	A93	M93	001A	Integrated Coupling Adapter SR	02	0504	
3	A93	M93	001A	Integrated Coupling Adapter SR	02	0505*	
4	A93	M93	001A	Integrated Coupling Adapter SR	02	0506	
5	A93	M93	001A	Integrated Coupling Adapter SR	02	0507	
6	M93	A93	0037	Integrated Coupling Adapter SR	02	0504	
7	M93	A93	0037	Integrated Coupling Adapter SR	02	0505*	
8	M93	A93	0037	Integrated Coupling Adapter SR	02	0506	
9	M93	A93	0037	Integrated Coupling Adapter SR	02	0507	
10	A93	T93	0002	Integrated Coupling Adapter SR	01	0508	

Figure 3-49 Sample CTN data export file

3.7 Deconfiguring the CTN

This task is used to deconfigure a CTN, and results in all systems dropping to an unsynchronized state (Stratum 0).

The starting point for this scenario is an unrestricted, dual-server CTN, shown in Figure 3-42 on page 61.

> **Attention - disruptive task:** Deconfiguring the CTN is a disruptive task for any workload that relies on timing information sourced from the Coordinated Server Time. Before attempting this task, make sure that your workloads are not impacted by the servers that change to an unsynchronized (stratum 0) state.

Use the following steps to deconfigure the CTN:

1. Select **Deconfigure CTN** from the Manage System Time menu.

   A warning message is shown (see Figure 3-50 on page 65).

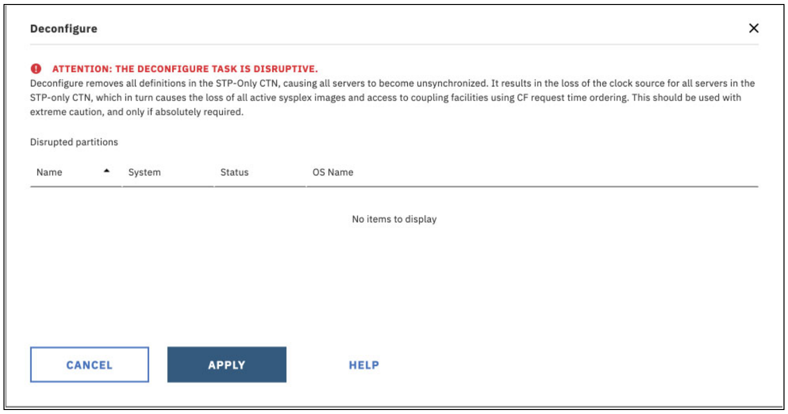

Figure 3-50 Disruptive task warning

2. After you click **APPLY**, the task starts and the task progress is displayed.
3. After the task finishes, you see the message that is shown in Figure 3-51.

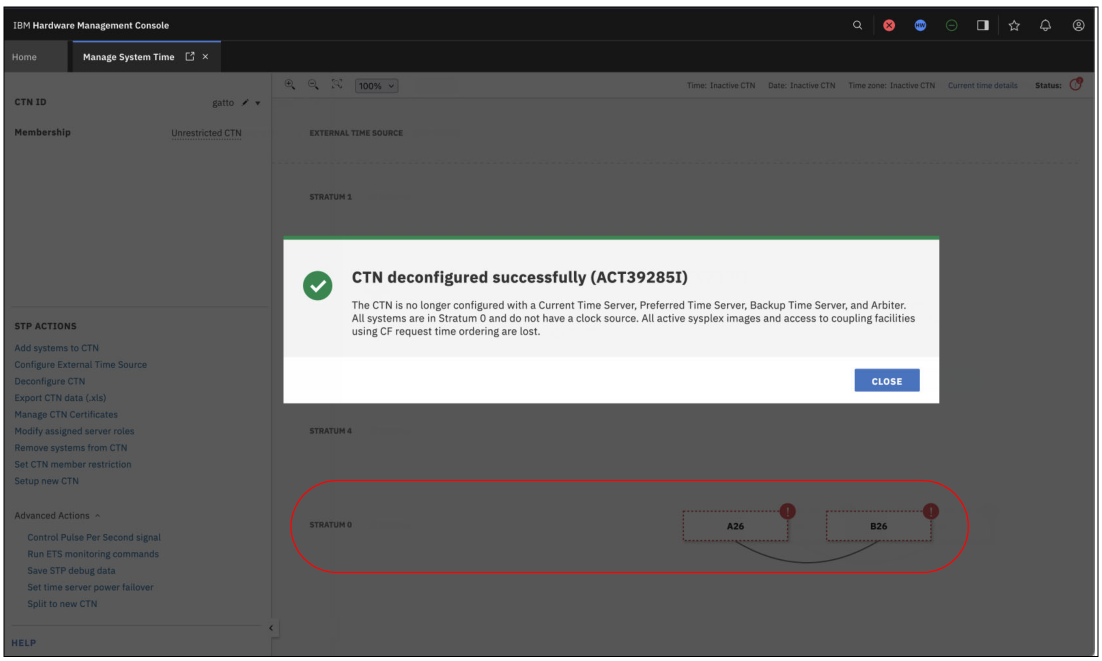

Figure 3-51 CTN deconfiguration success

In Figure 3-51, both servers are now STRATUM 0 and they are not part of the CTN.

4

External Time Source Configuration and Management

This chapter describes how to configure and manage the External Time Source (ETS) configuration for your CTN.

The following tasks are covered in this chapter:

- 4.1, “Configure ETS connectivity” on page 68
- 4.2, “ETS Network Diagnostics” on page 71
- 4.3, “Configure ETS” on page 74
- 4.4, “ETS status” on page 92
- 4.5, “Control Pulse Per Second signal” on page 100
- 4.6, “Manage CTN Certificates” on page 105

Note: Local area network (LAN) configuration for the External Time Source is not covered in this document. In 2.3.6, “ETS Network physical planning” on page 29 you can find planning considerations for providing your IBM Z server with connectivity to the External Time Sources.

4.1 Configure ETS connectivity

Before you can add NTP and/or PTP sources, ensure that ETS connectivity is successfully configured through the Support Element (SE) Customer Network configuration task before attempting this task.

To configure external time source connectivity, access the IBM Support Element. To access the IBM Support Element, use the HMC. From the HMC main menu, select **Single Object Operations** → **SOO**, as shown in Figure 4-1 and then initiate the Customize Network Settings task. The user ID must have administrative privileges, such as ACSADMIN.

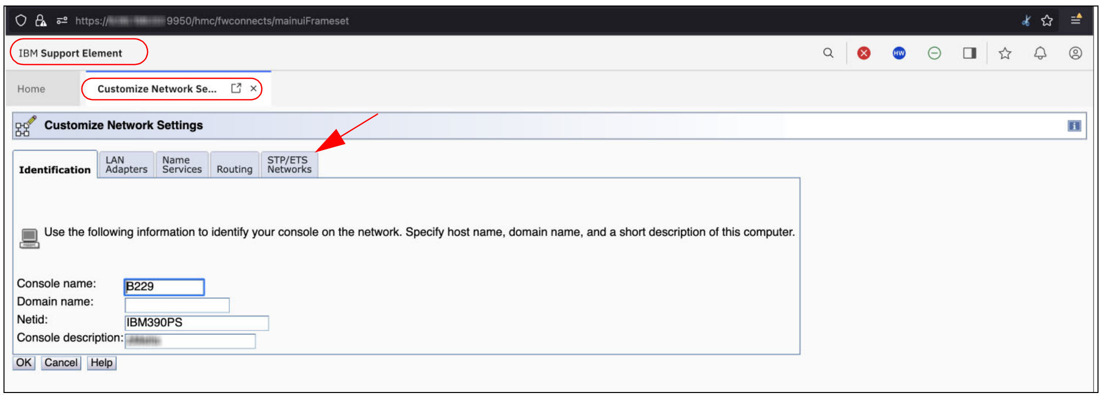

Figure 4-1 Single Object Operations (Support Element Workplace)

Start by selecting the STP/ETS Networks tab to configure the following parameters:

- ETS (Ethernet) interfaces (addresses, netmask)
- IP Routing information
- Name services servers
- Domain naming rules

Use the following steps to configure ETS connectivity:

1. Configure the ETS interfaces IP addresses ("Details"), as shown in Figure 4-2 on page 69.

   You can also configure IPv6 addresses for each interface if your networking infrastructure is configured with IPv6.

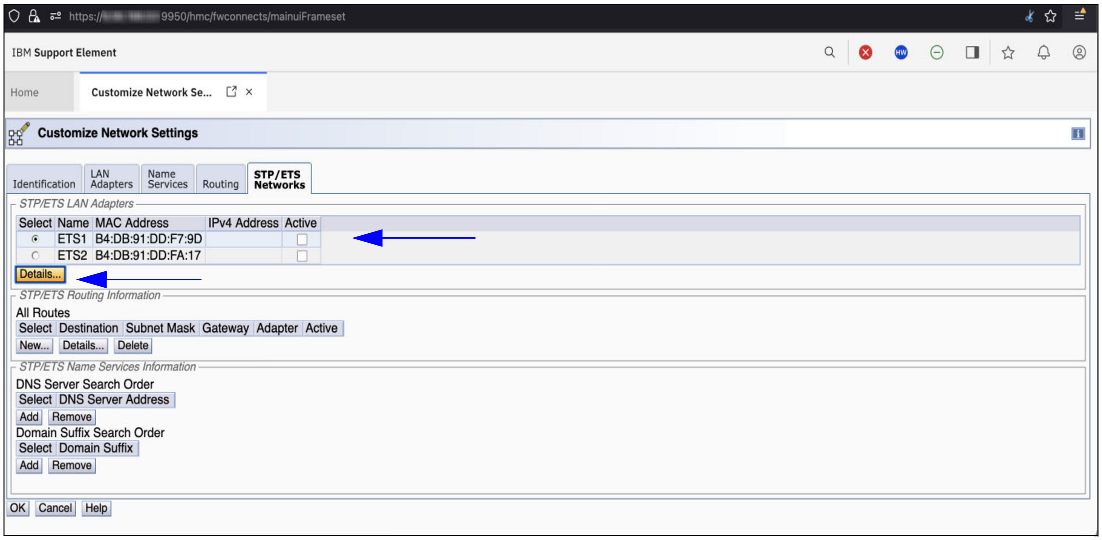

Figure 4-2 ETS interfaces configuration start

2. Enter the address and network mask for the ETS interface as shown in Figure 4-3.
3. Repeat the steps for the second ETS interface.

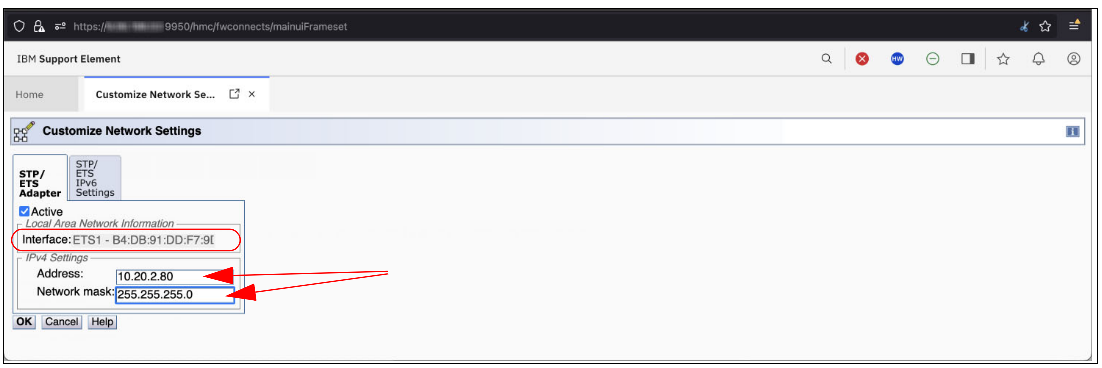

Figure 4-3 IP address and netmask information (IPv4).

4. Click **New** to configure the routes for the newly defined interfaces as shown in Figure 4-4.

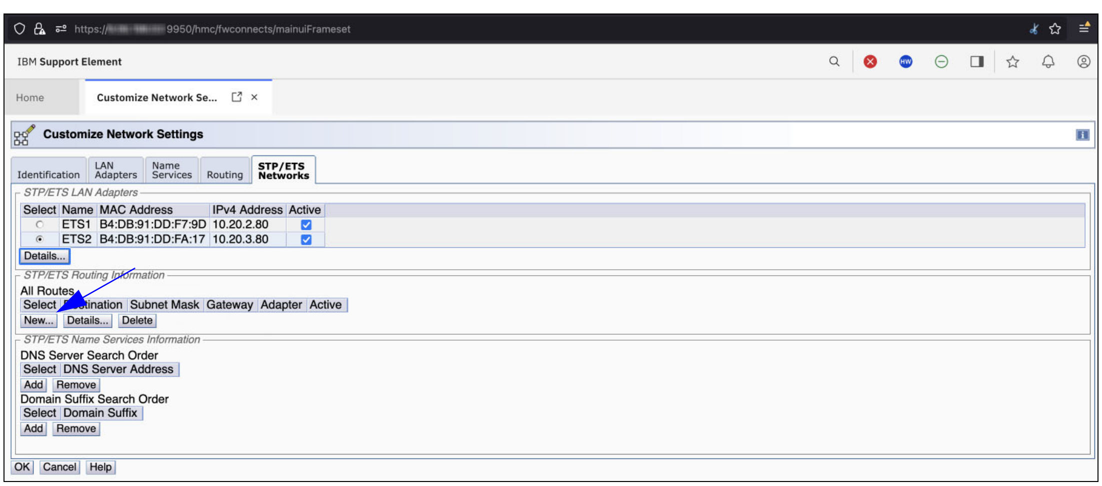

Figure 4-4 Defining IP routes

5. Add routing information for each route that you want to add as shown in Figure 4-5.

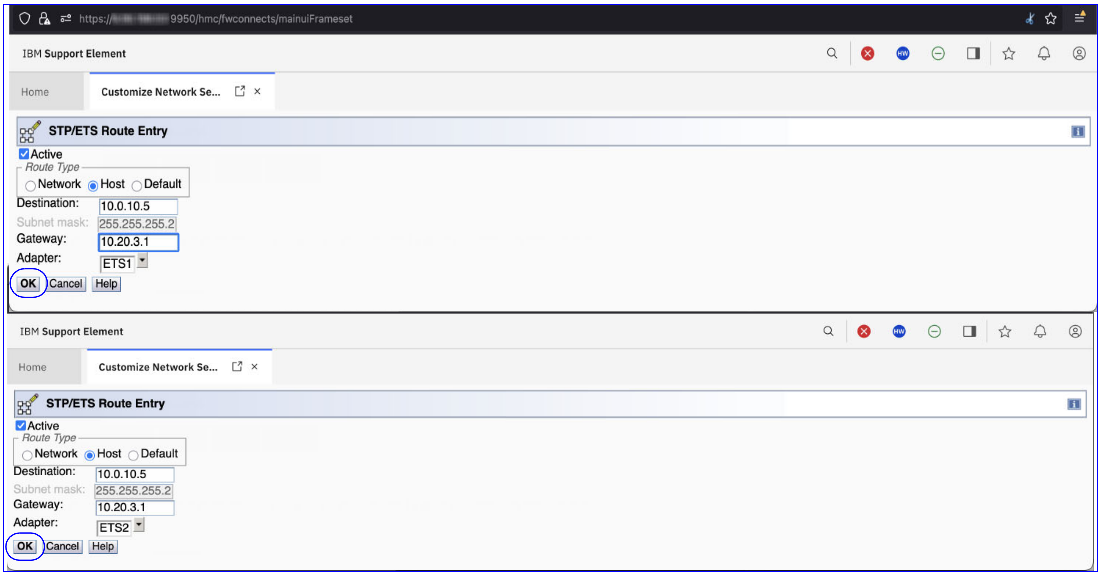

Figure 4-5 Routing information entered

6. Routing information is shown in Figure 4-6.
7. Under DNS Server Search Order, click **Add** to configure DNS servers.
8. Under Domain Suffix Search Order, click **Add** to define the domain naming configuration.

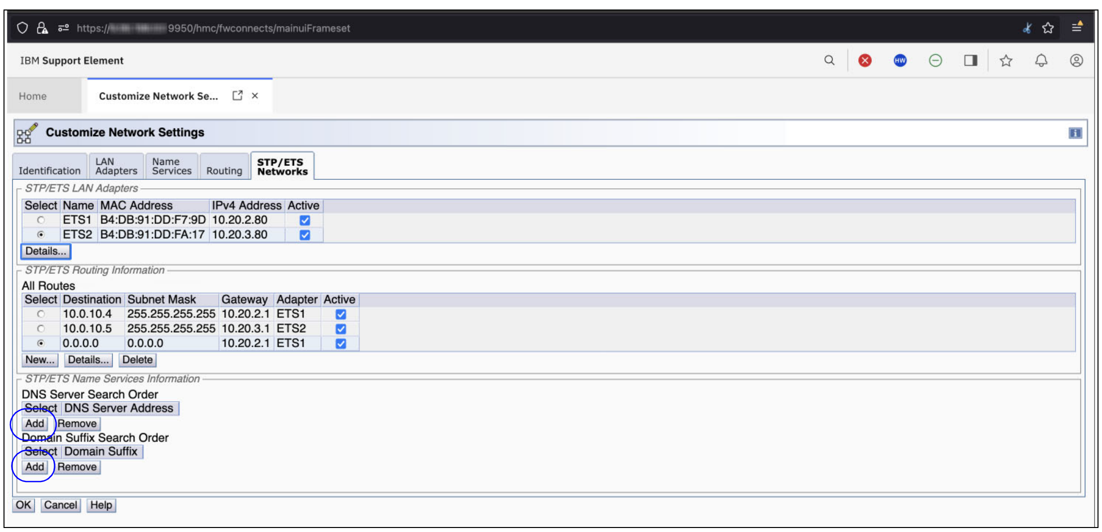

Figure 4-6 Adding DNS and Domain Name suffix information

9. When all the data is entered (Figure 4-7), click **OK** to commit changes.

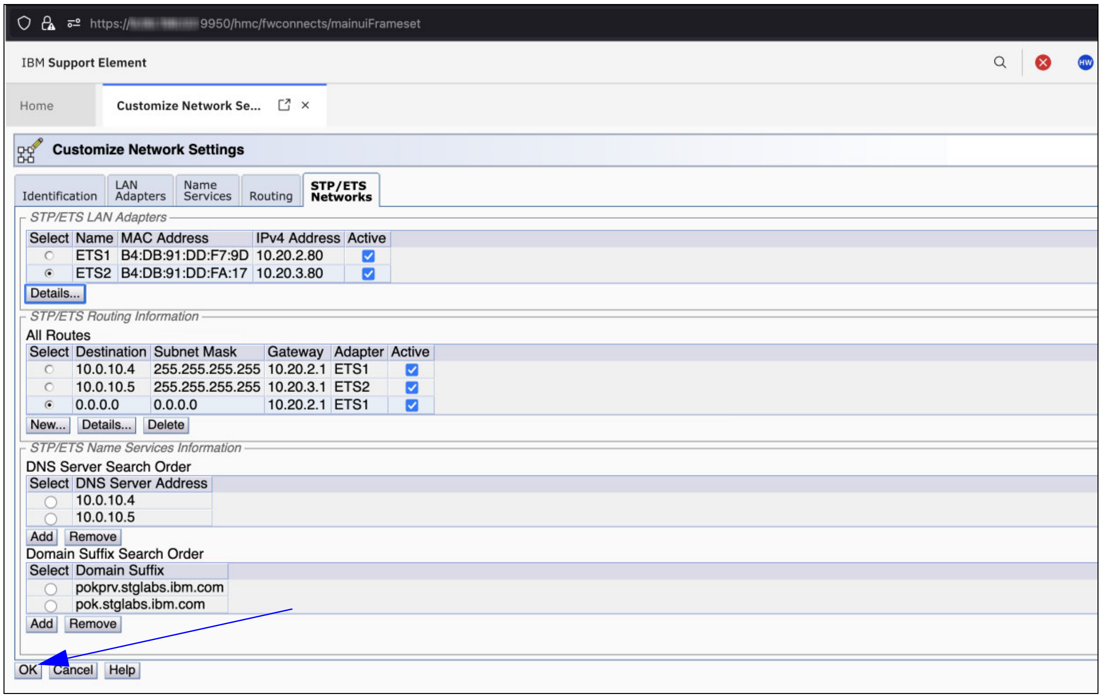

Figure 4-7 ETS Network settings (overview before committing changes)

Click “OK”, and the changes are committed (ETS connectivity is activated - Figure 4-8)

Figure 4-8 ETS connectivity configured and activated

4.2 ETS Network Diagnostics

After the ETS interfaces are configured, test the ETS network connectivity to validate the availability of the external time sources.

To test ETS connectivity, access the Test ETS Connectivity task on the support element workplace (SOO), as shown in Figure 4-9 on page 72, and enter the IP address to ping.

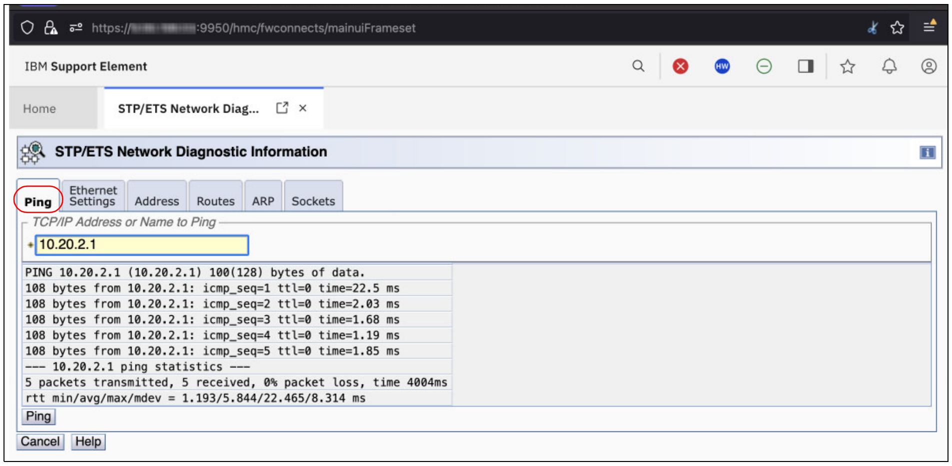

Figure 4-9 ICMP network connectivity testing for ETS network ("ping")

Review the Ethernet configuration by selecting the Ethernet Settings tab as shown in Figure 4-10.

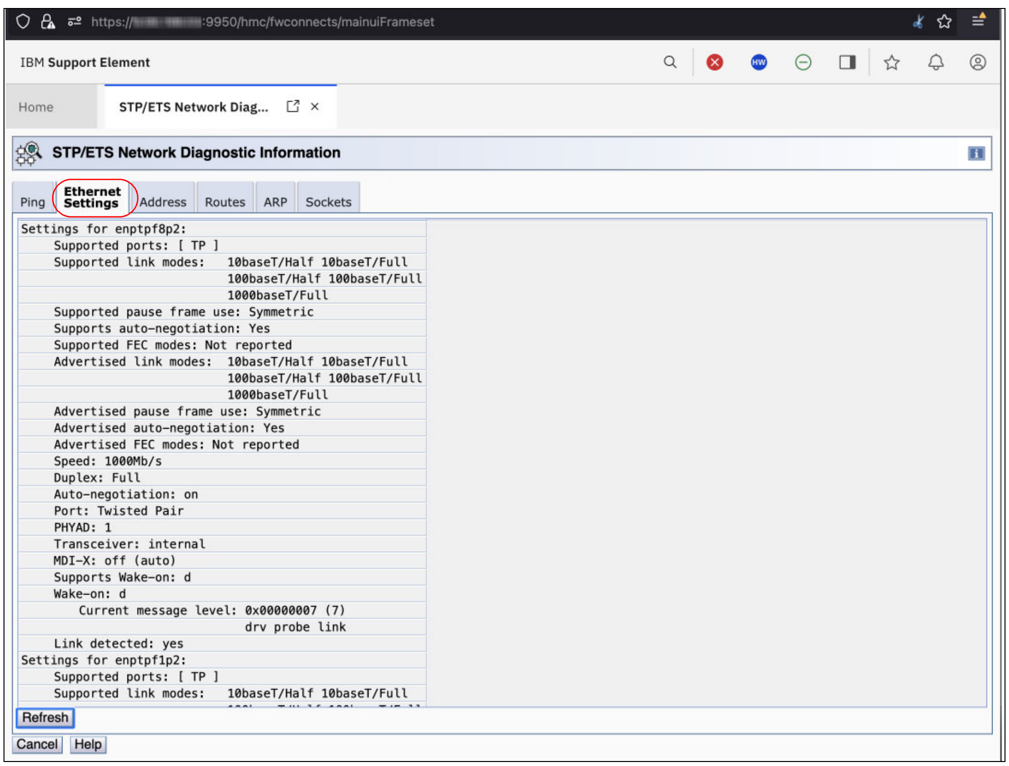

Figure 4-10 ETS interfaces physical connectivity information

Review the IP address configuration for the ETS interfaces by selecting the Address tab, as shown in Figure 4-11.

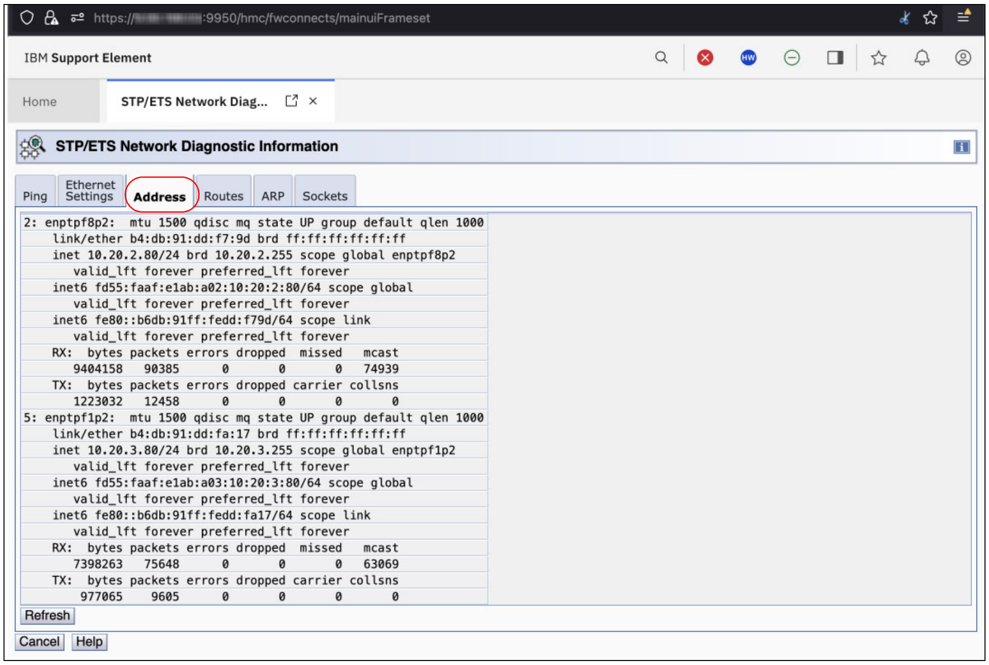

Figure 4-11   Ethernet addresses and statistics

Select the Routes tab to view the routing information for the ETS interfaces as shown in Figure 4-12.

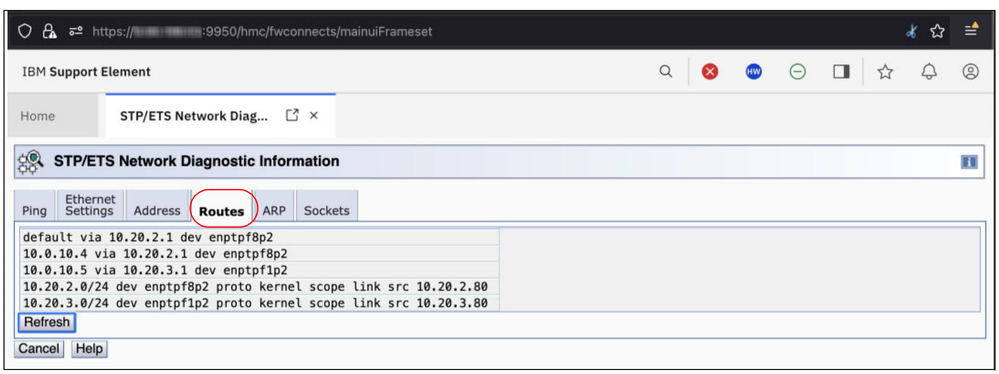

Figure 4-12   Active routing information

Select the ARP[1] tab to view the relationship between MAC addresses and assigned IP addresses as shown in Figure 4-13 on page 74.

1 ARP stands for Address Resolution Protocol.

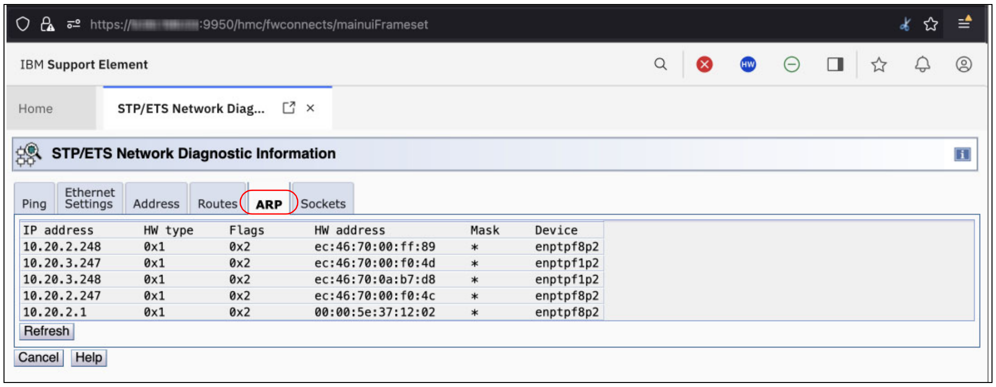

Figure 4-13 Address resolution protocol database

Select the Sockets tab to view the IP connectivity (IP sockets) information as shown in Figure 4-14.

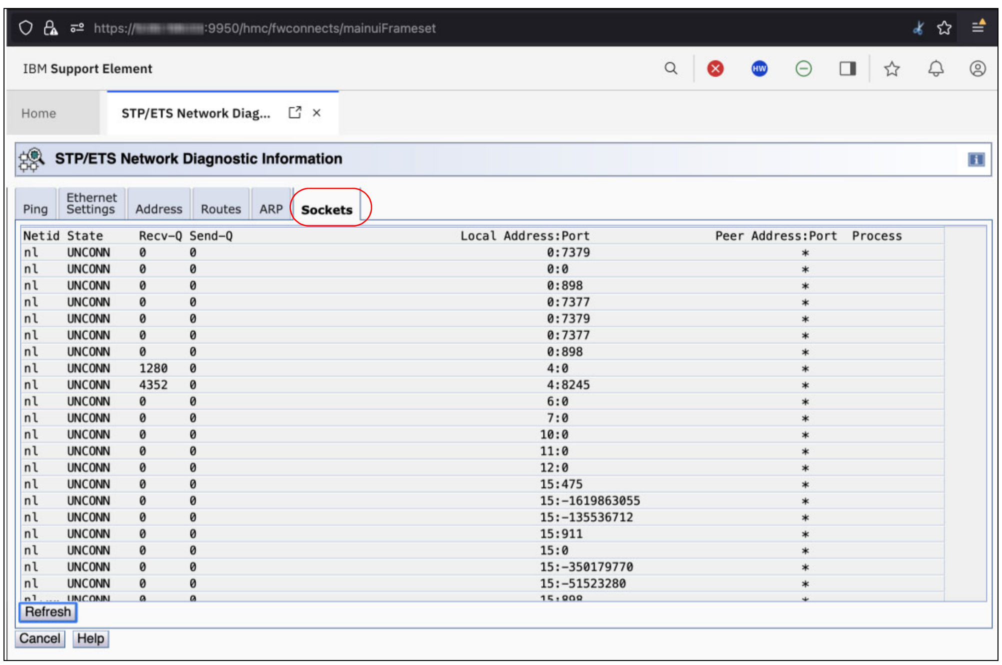

Figure 4-14 Checking (IP) networking sockets

4.3 Configure ETS

The configuration options for the External Time Source (ETS) used in this scenario depend on your organization and site network settings and rules. Ensure that you have the required information to configure ETS connectivity for your location rules.

You established network connectivity for the ETS interfaces on the dedicated ports on the BMC/OSC[2] card on the CPC drawers, and you ensured that the time servers provide the time synchronization data by using NTP or PTP. Therefore, you can proceed to assigning the ETS providers to the CPCs.

---

[2] BMC/OSC card - Baseboard Management Controller and Oscillator combined card - replacing the Flexible Service Processor / Oscillator (FSP/OSC) card from previous IBM Z models (z15 and earlier).

External time source is supported in the following configurations:

- IBM z16 and later generations support both PTP and NTP as an ETS connected directly to the ports of the CPC drawer where ETS is defined.
- IBM z16 supports the definition of two ETS if both are NTP or both are PTP.
- IBM z17 supports the definition of up to five (5) ETS, and combining a maximum of two PTP sources and three NTP sources is supported.
- IBM z16 and z17 support the optional Pulse Per Second signal (direct connection to the CPC drawer) for both PTP and NTP sources.
- IBM z15 supports NTP and PTP as ETS (connected to the ETS port on the Support Element) with optional Pulse Per Second signal (direct connection to the CPC drawer).

4.3.1 ETS for IBM z17

In this scenario, configure two NTP servers as ETS for A26 server, which is an IBM z16 system. Use the following steps to perform the configuration:

1. In the Manage System Time task main menu, select the External Time Source menu (Figure 4-15), and click **NEXT**.

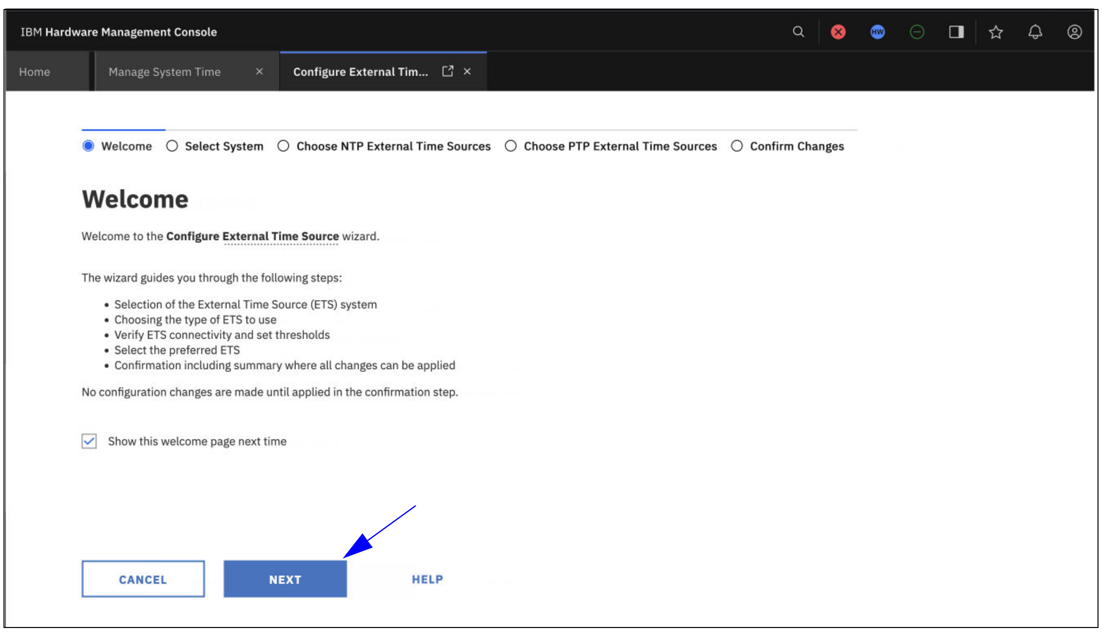

Figure 4-15 Configure External Time Source menu

2. Choose the target server for the configuration and click **NEXT**. In the example, B229 (an IBM z17 server) is selected. (Figure 4-16 on page 76).

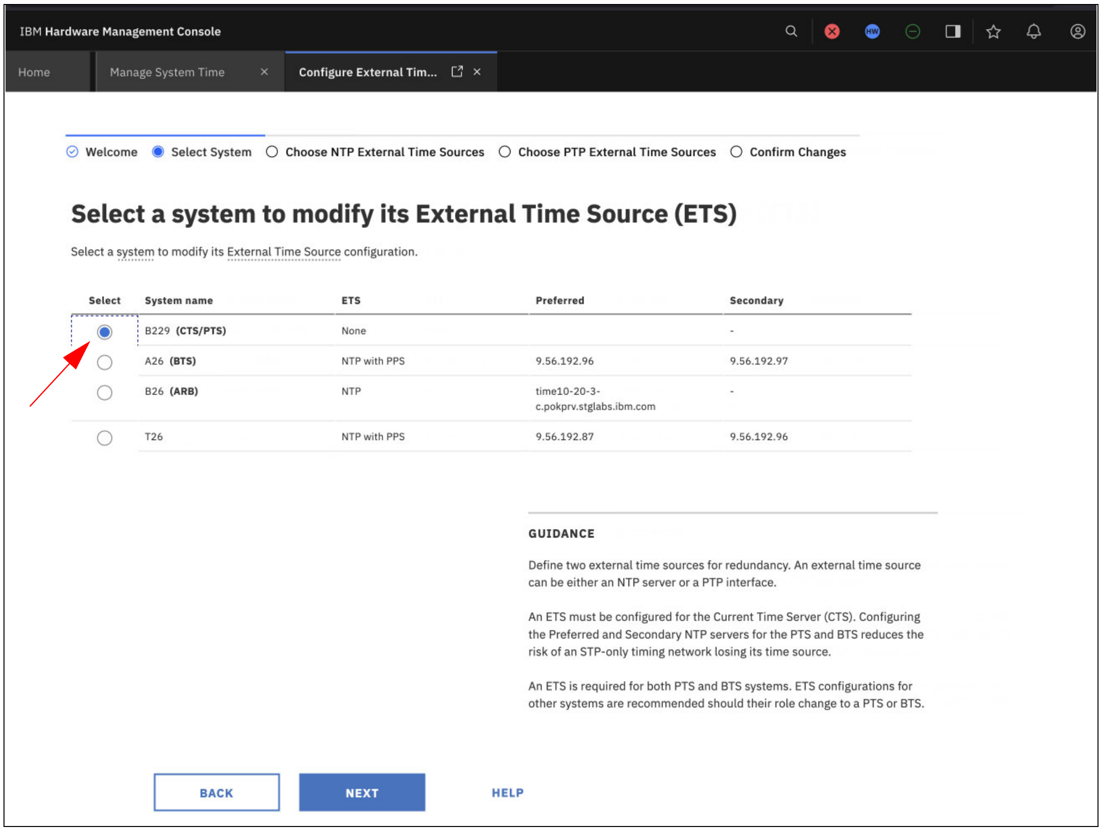

Figure 4-16 Selecting server for ETS configuration

3. Click **ADD NTP SERVER** to configure an NTP server as an ETS for B229 CPC. (Figure 4-17).

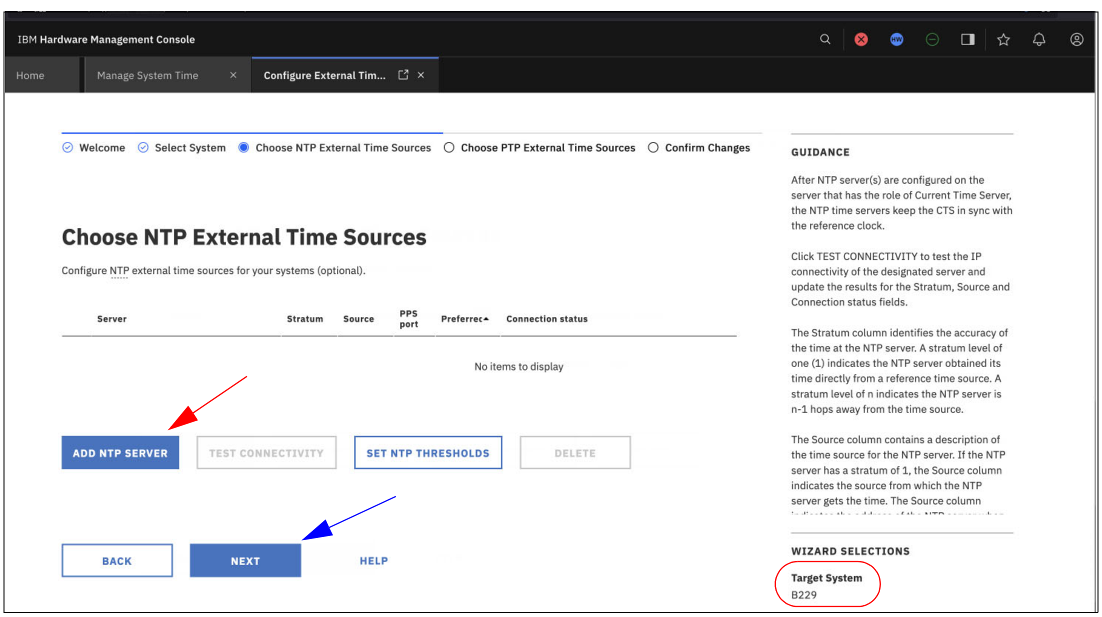

Figure 4-17

4. Enter the following information Figure 4-18 on page 77):
   – Server IP address (or fully qualified domain name, if applicable)

– Select the CPC PPS port associated with this NTP server (optional, if already connected)
– NTP security method (also optional if configured at your location).
– Select the box for the **Preferred server for time adjustments**.

Note: If your site uses NTP security (NTS), then before you enable this feature, you must also add the associated NTP server certificate to your ETS configuration. This task is covered in 4.6, “Manage CTN Certificates” on page 105.

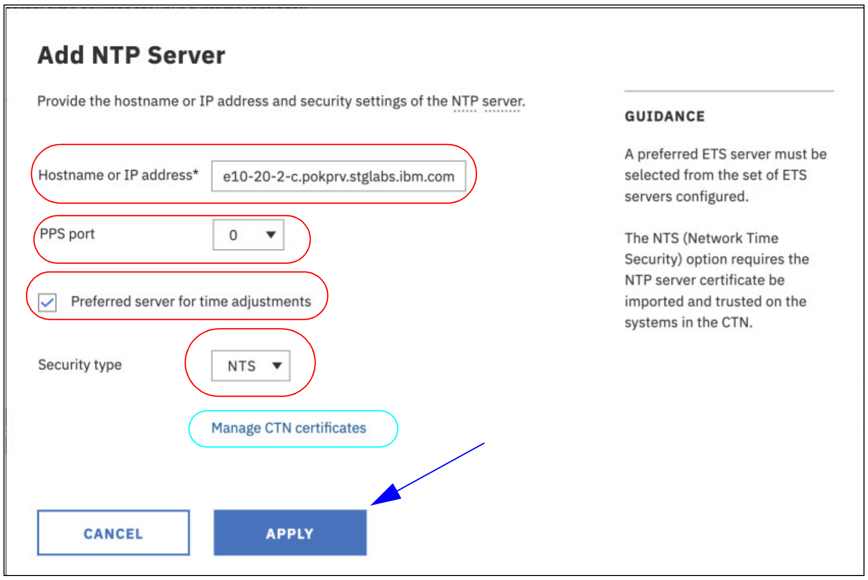

Figure 4-18 NTP server data

5. Apply the information for this server.
6. Add another NTP server as shown in Figure 4-19. For the second NTP server, the Preferred server box is *not* selected. Click **APPLY**.

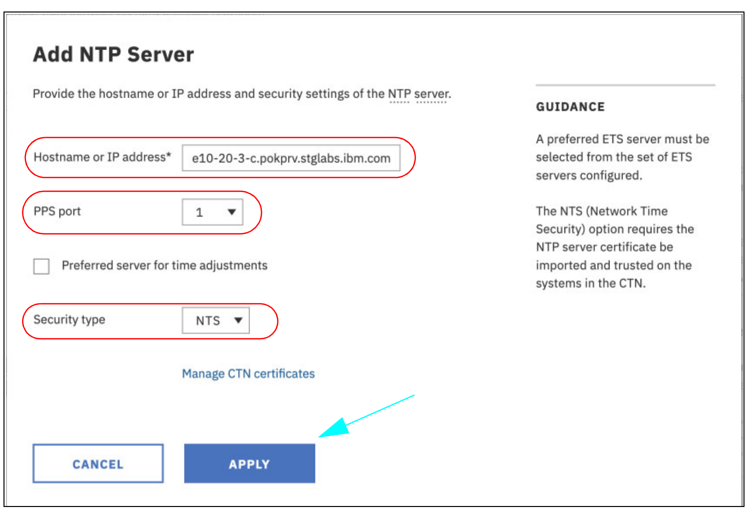

Figure 4-19 Second NTP server data

7. Test the connectivity for the configured NTP servers (Figure 4-20).

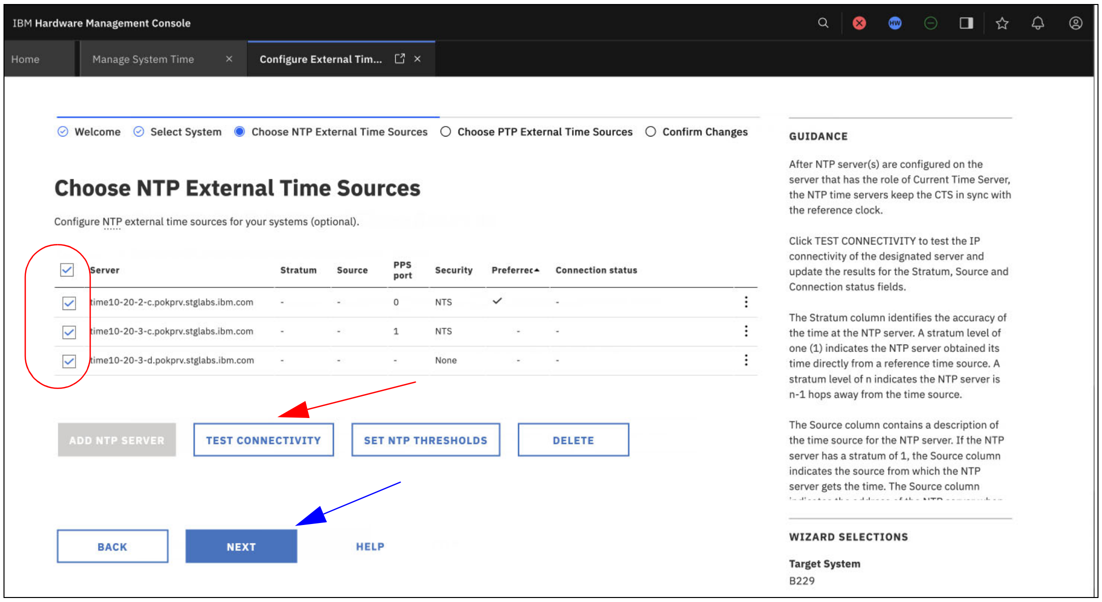

Figure 4-20 NTP servers selected for connectivity test

8. After you test and verify the connection, select **SET NTP THRESHOLDS** (Figure 4-21).

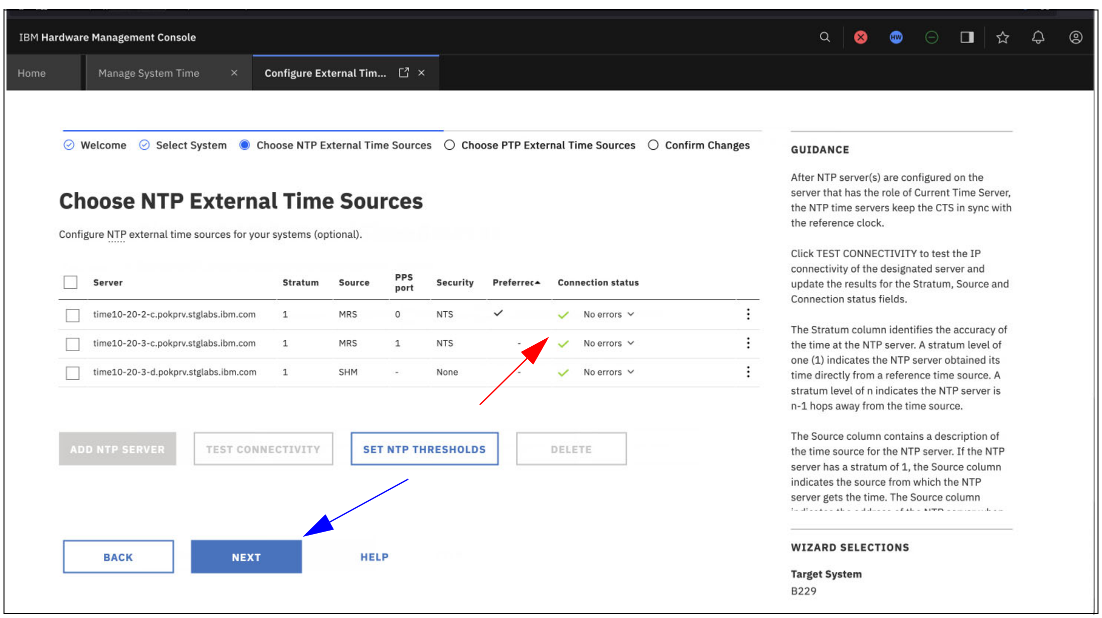

Figure 4-21 NTP servers connectivity successful test

9. Add PTP servers to the B229 CPC (Figure 4-22 on page 79).

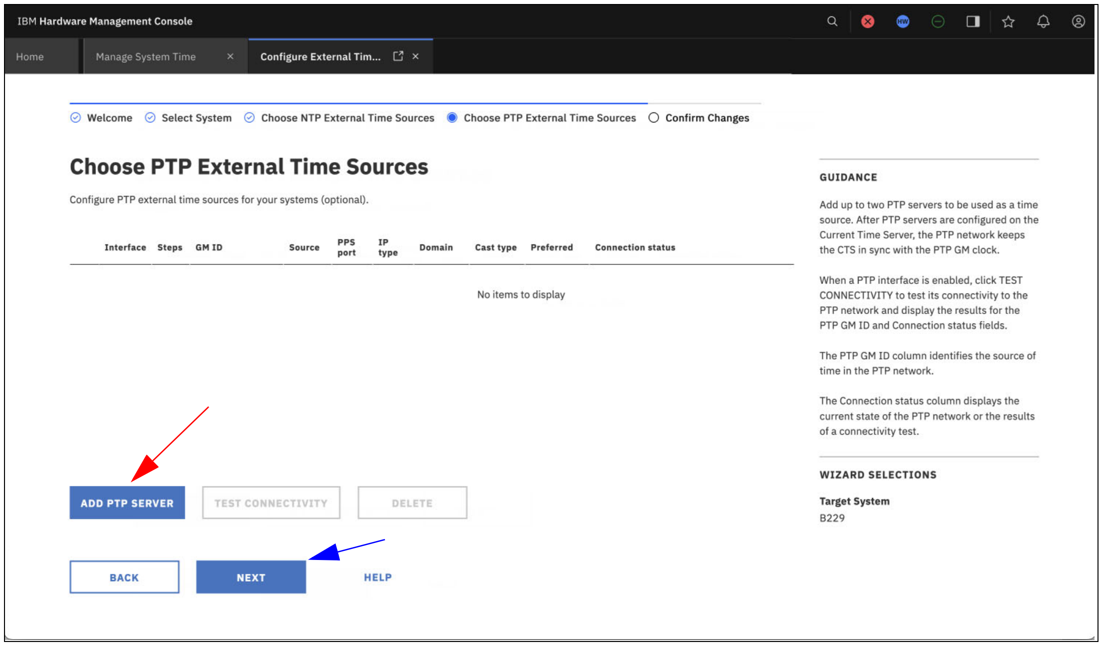

Figure 4-22 Adding PTP sources

10. Select **ADD PTP SERVERS** to add PTP Grandmaster information for ETS1 as shown in Figure 4-23.

    The Domain number must be selected to match your PTP server configuration (verify with your network administrator before configuring this). A wrong domain number prevents communication with the PTP Grandmaster server. For additional information, see 2.3.2, “Precision Time Protocol” on page 23.

11. Click **APPLY**.

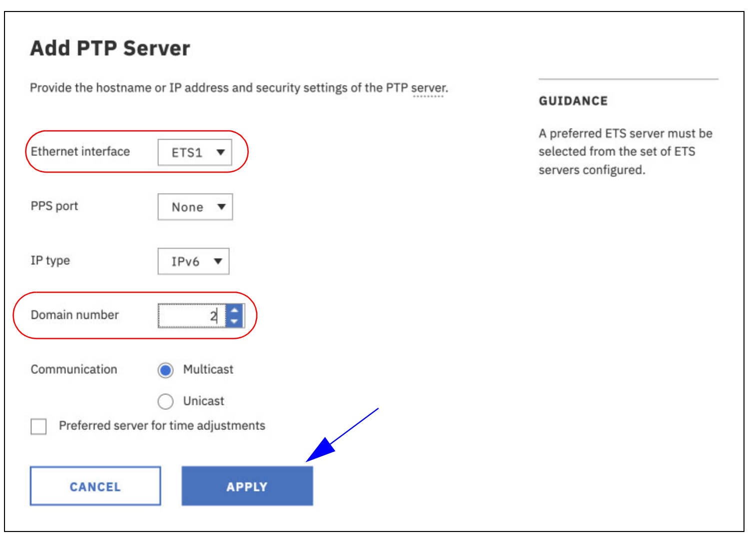

Figure 4-23 PTP server data

12. After you test and verify connectivity, you can view the definition of the PTP source (Figure 4-24 on page 80).

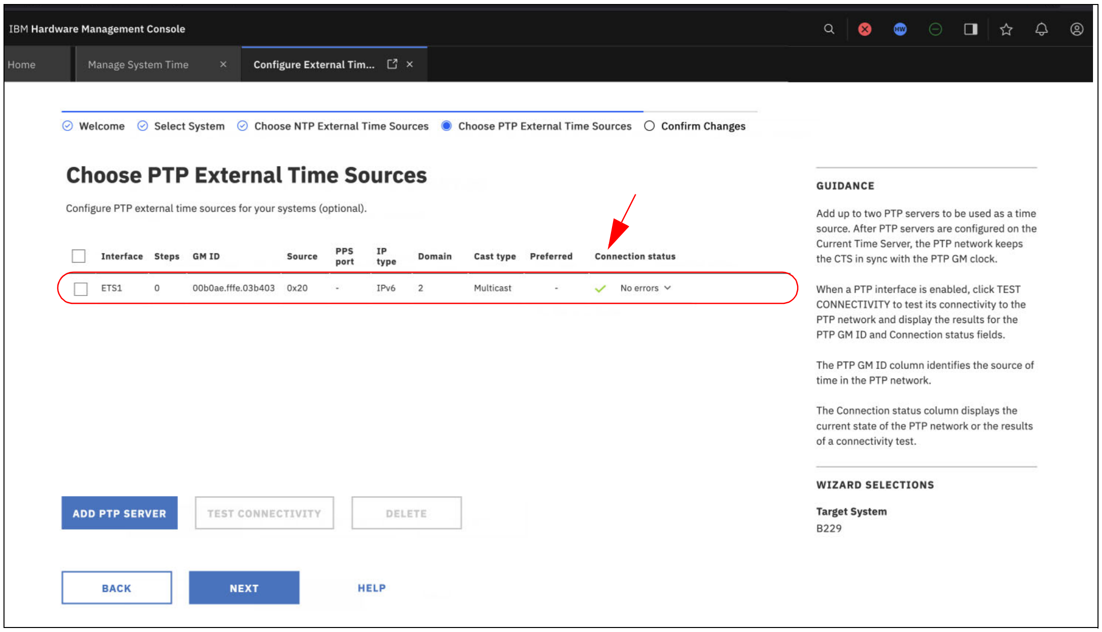

Figure 4-24 First PTP source configured and tested

13. You can change the ETS1 source by clicking the three dots at the end of the server line and selecting **Edit.**
14. Select a different PTP Grandmaster as shown in Figure 4-25.
15. After changing the ETS1 source, test the connectivity.

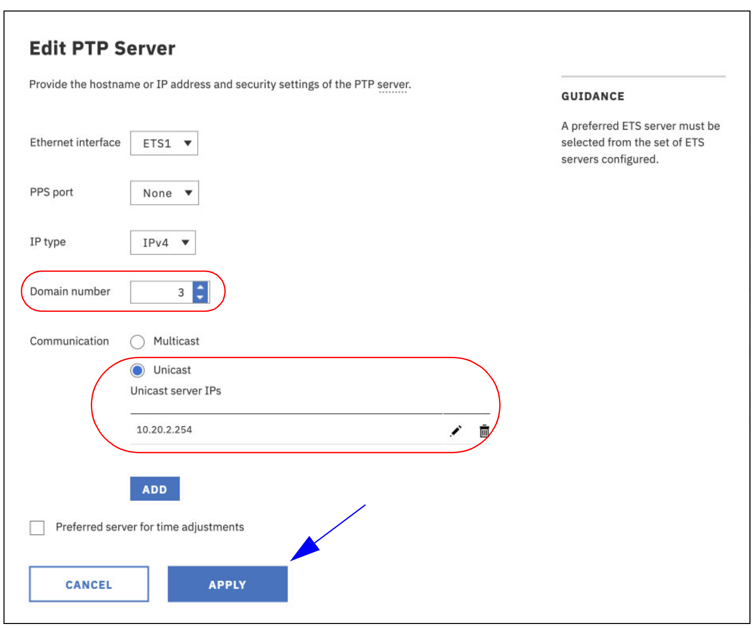

Figure 4-25 Second PTP source

16. Add ETS2 as a second PTP server by selecting **ADD PTP SERVER**.
17. Test and verify connectivity to the second PTP server. The ETS configuration for PTP is shown in Figure 4-26 on page 81.
18. Click **NEXT**.

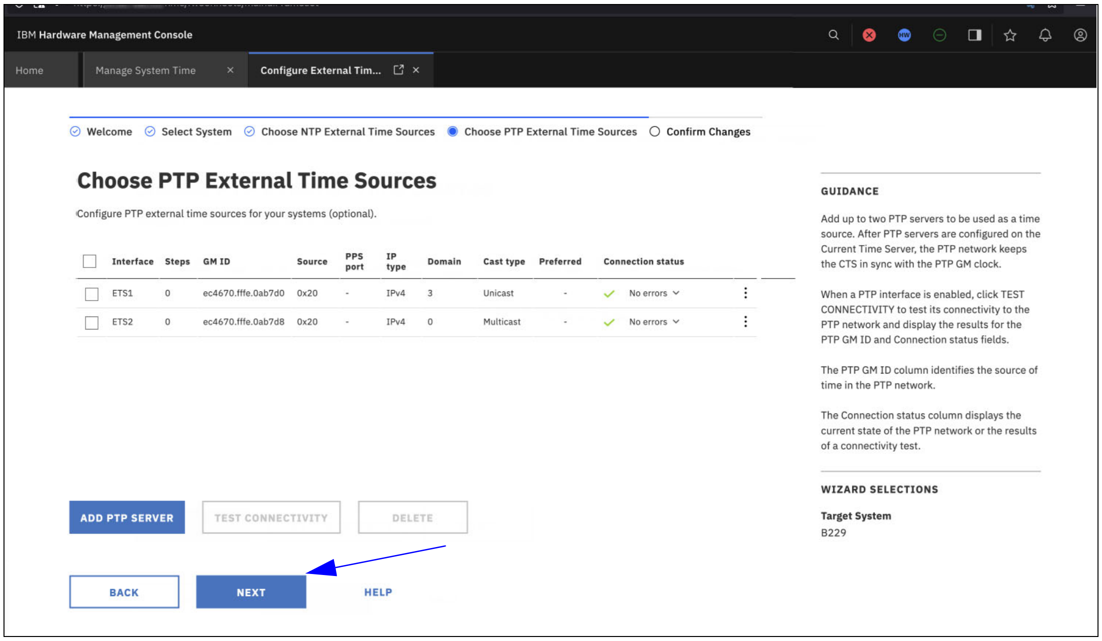

Figure 4-26 ETS1 and ETS2 configured for PTP sources

19. Review the ETS configuration for B229 CPC and click **APPLY**. (Figure 4-27 on page 82).

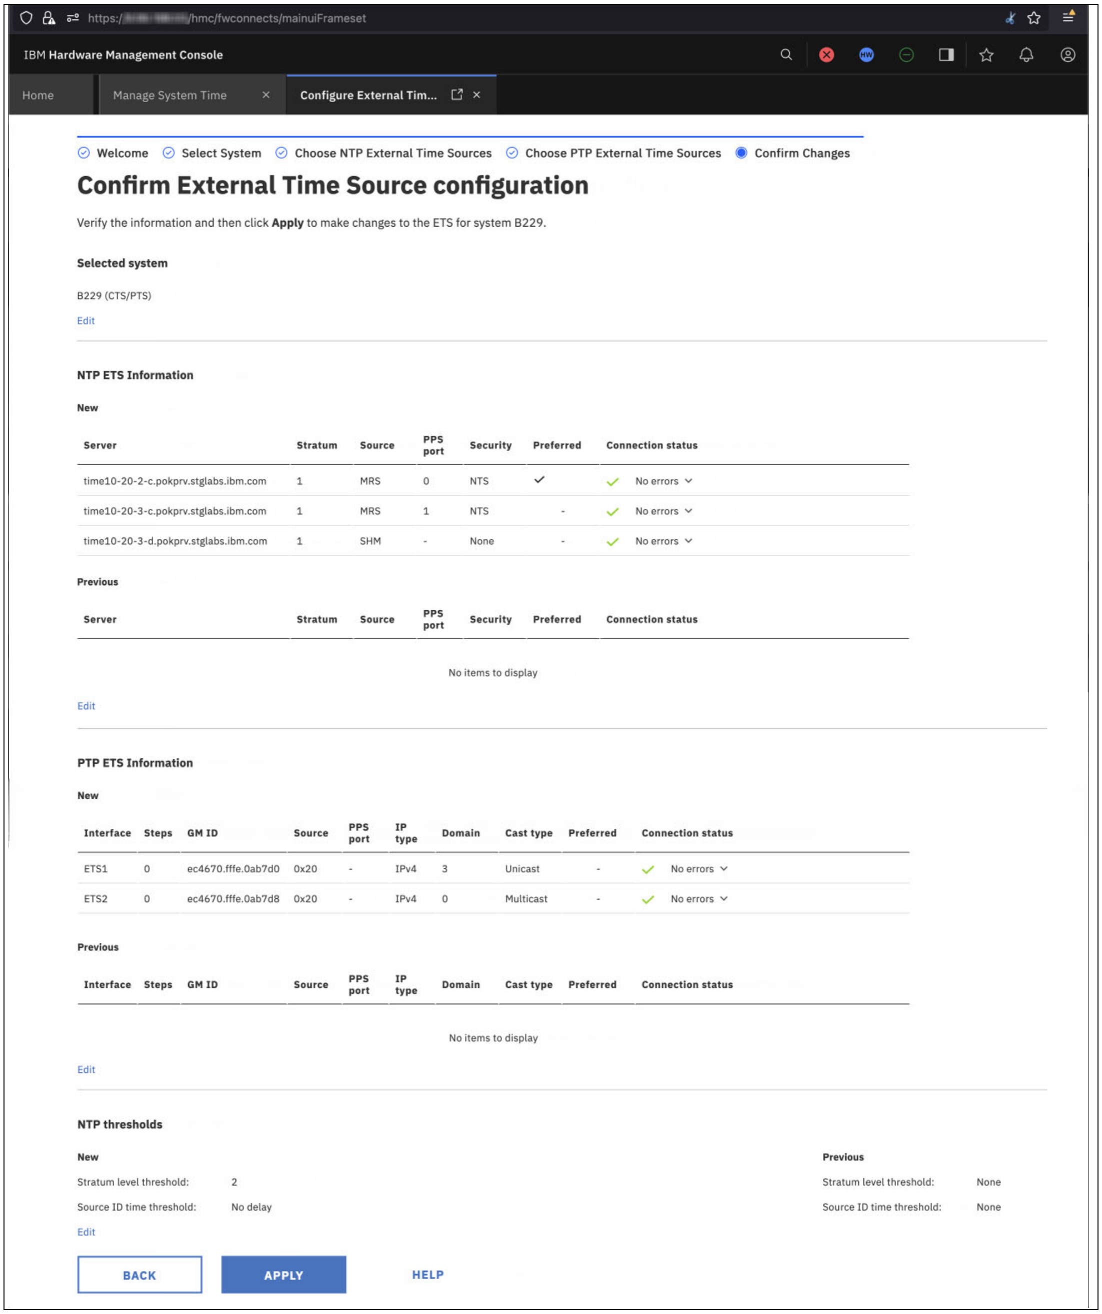

Figure 4-27 ETS configuration overview before commit (B229 server)

After you successfully apply the configuration, select the Manage System Time tab to verify the ETS configuration for your CTN (Figure 4-28 on page 83).

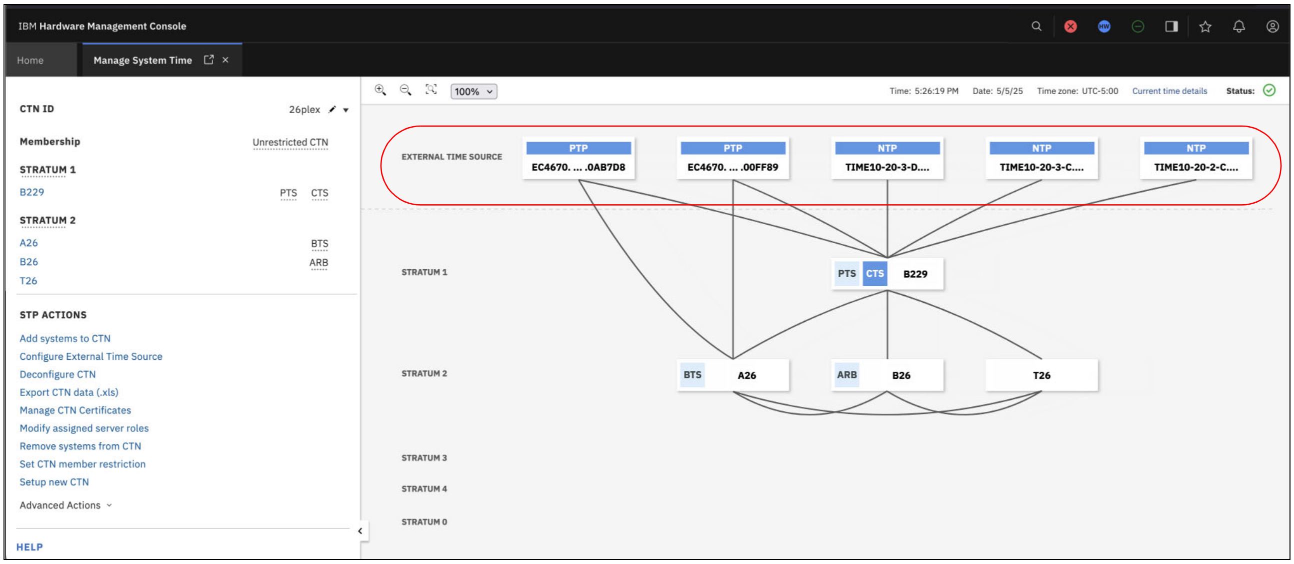

Figure 4-28 ETS configuration in the Manage System Time task

4.3.2 ETS for IBM z16

This section describes how to configure an ETS for an IBM z16 server.

NTP as ETS

In this scenario, configure two NTP servers as an ETS for the IBM z16 server A26. Use the following instructions to add the NTP servers.

1. In the Manage System Time task, select the button **Choose NTP External Time Source** at the top of the screen to view the menu and the target server for the configuration. In this example, A26 is already selected as the target server (Figure 4-29).

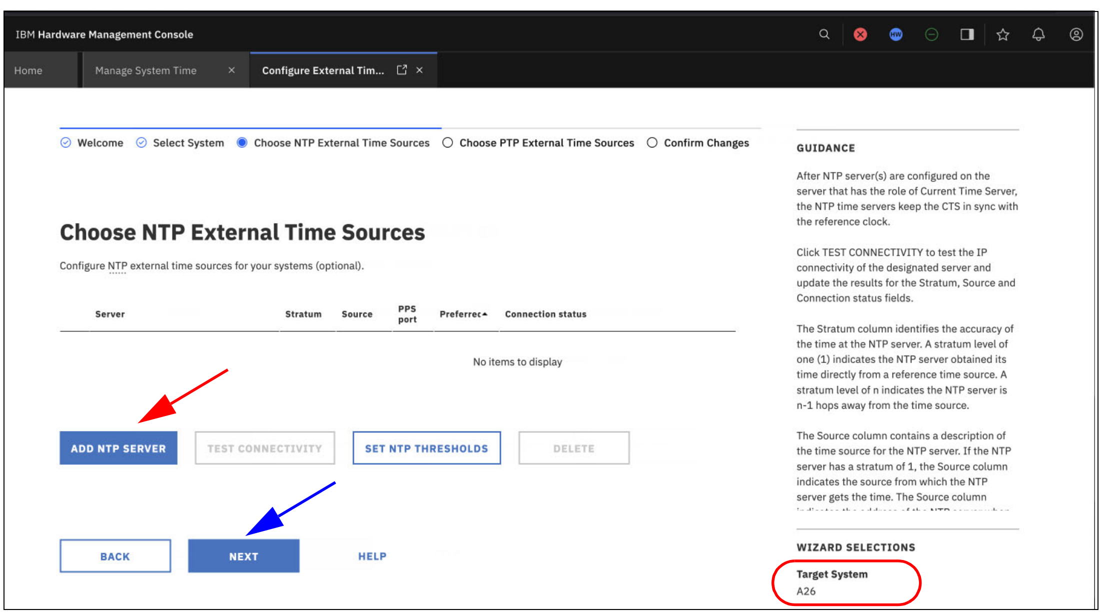

Figure 4-29 Choose NTP External Time Sources

2. Enter the NTP server IP address or fully qualified domain name (FQDN) (Figure 4-30). Click **APPLY**.

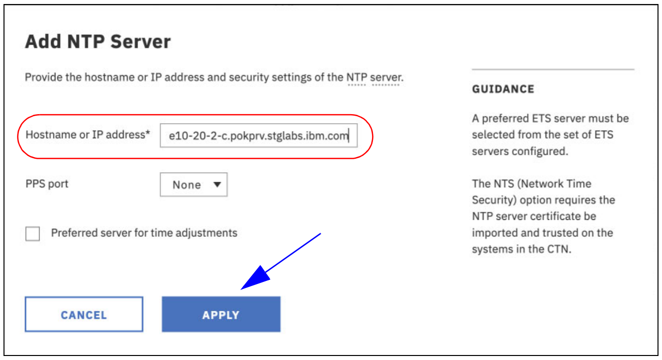

Figure 4-30 NTP server identification (address/FQDN)

3. If the NTP server also has PPS ports, select the CPC PPS port associated with this NTP source server (Figure 4-31).

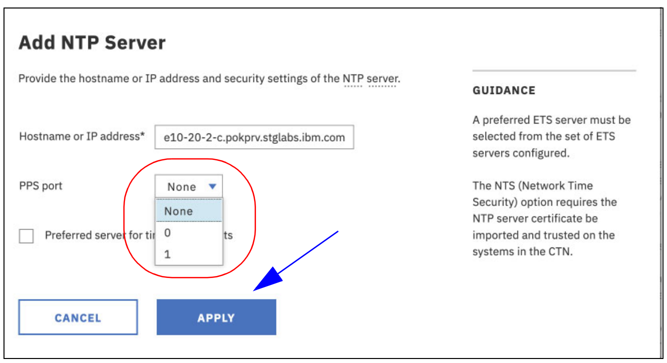

Figure 4-31 PPS port selection

4. Select the box **Preferred server for time adjustments** to designate this NTP server as the preferred source (Figure 4-32). Click **APPLY**.

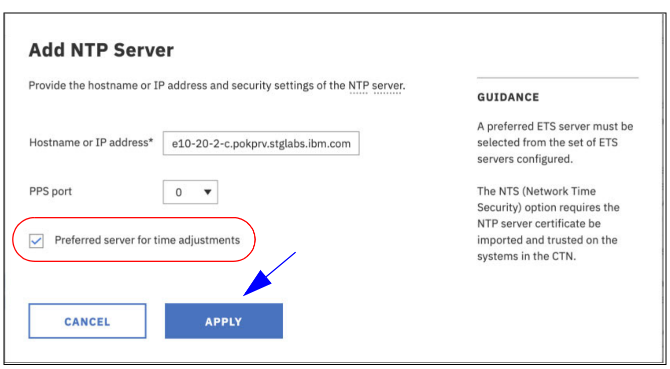

Figure 4-32 Preferred NTP time server for adjustments

5. Add a second NTP source by repeating the operation (Figure 4-33 on page 85).

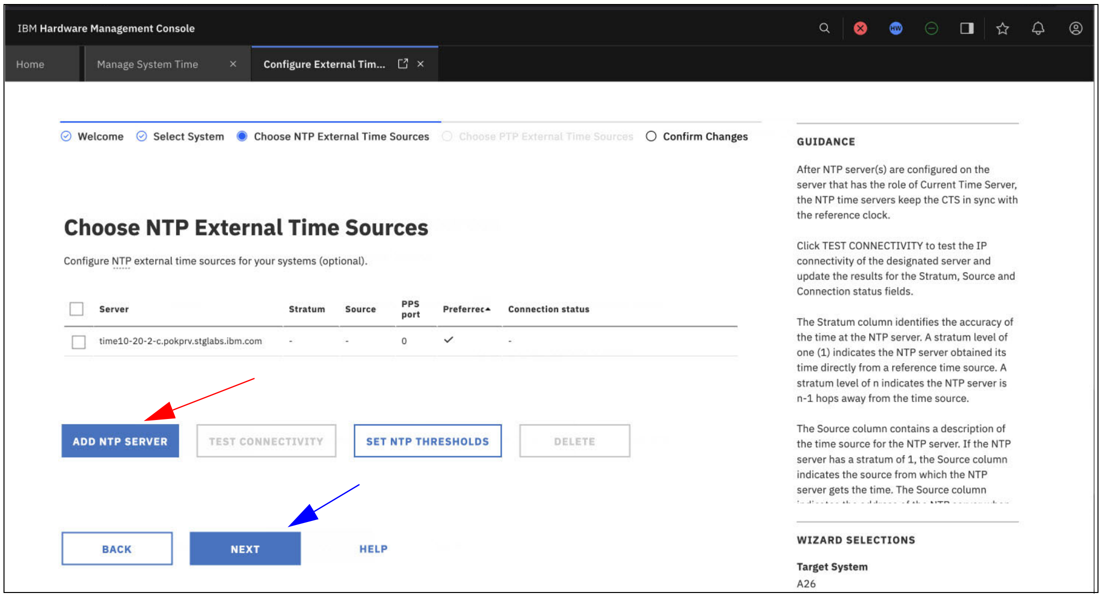

Figure 4-33 Adding another NTP server

6. After you add the second NTP server, click **NEXT (**Figure 4-34).

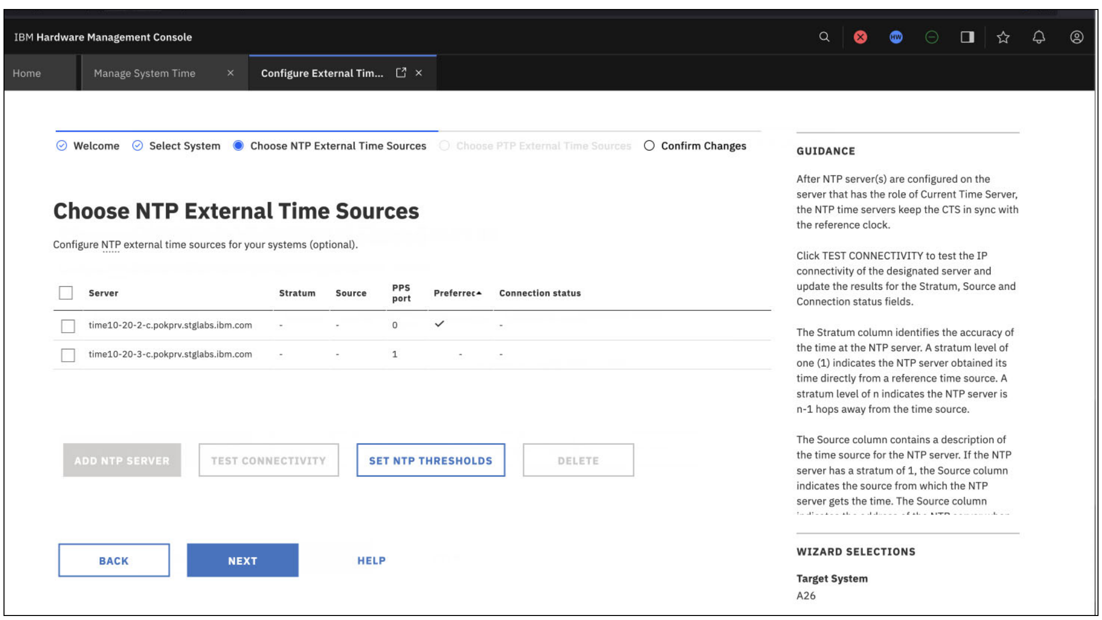

Figure 4-34 NTP servers selected as ETS

7. View the configuration confirmation screen (Figure 4-35 on page 86). Verify the proposed configuration and click **APPLY**.

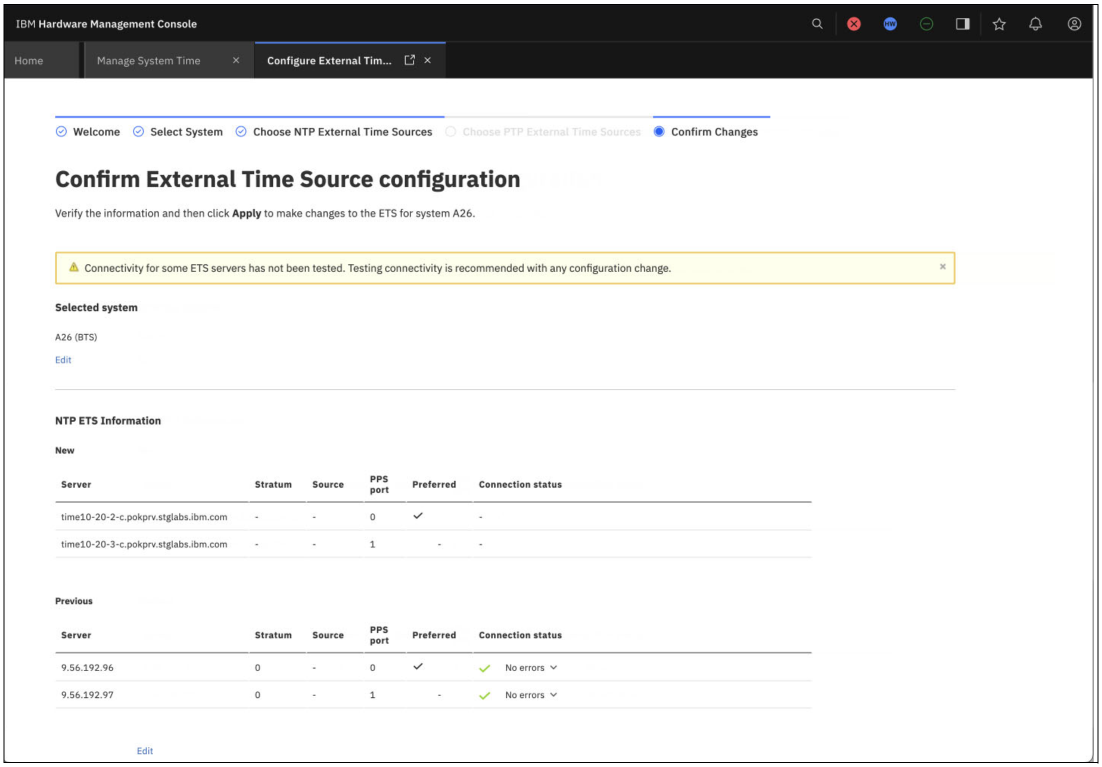

Figure 4-35 ETS changes overview

8. View the ETS configuration success message (Figure 4-36) and click **CLOSE**.

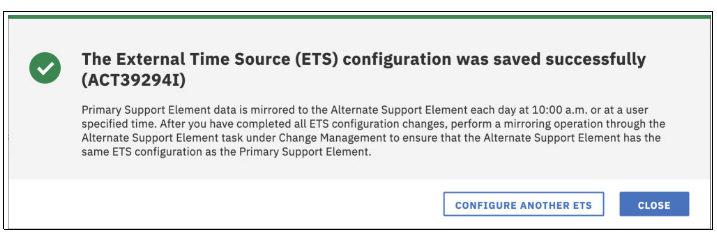

Figure 4-36 ETS successful configuration message

PTP as ETS

In this scenario, configure two PTP servers as ETS for the IBM z16 server A26. For IBM z16, both the ETS must be both NTP or both PTP. PPS is optional for both PTP and NTP. Use the following instructions to configure the PTP servers:

1. In the Manage System Time task, select the button **Choose PTP External Time Source** at the top of the screen to view the menu and the target server for the configuration. (Figure 4-37 on page 87).
2. Select **ADD PTP SERVER**.

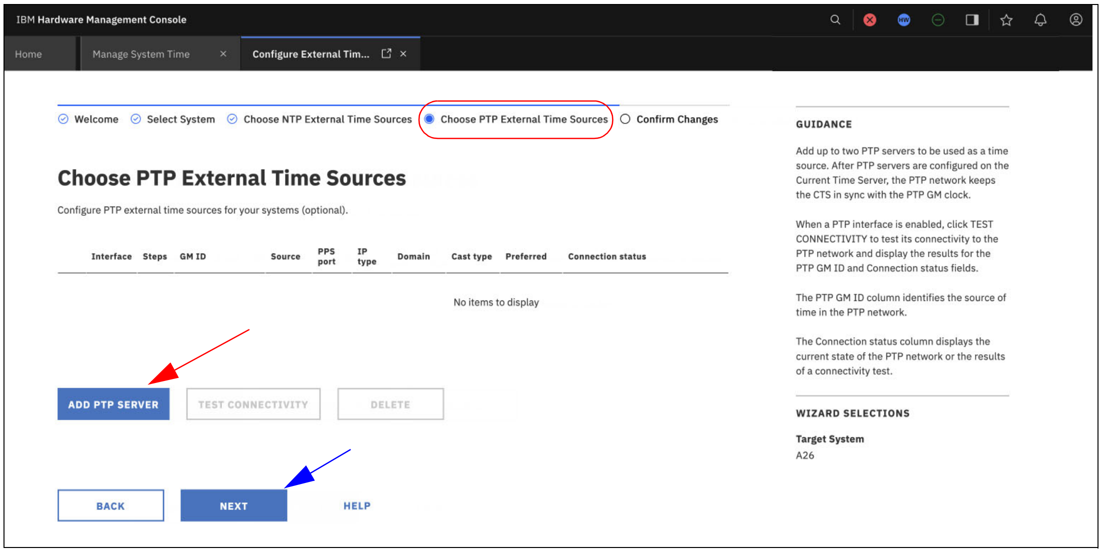

Figure 4-37 Choose PTP as ETS menu

3. Select the Ethernet interface (Figure 4-38).

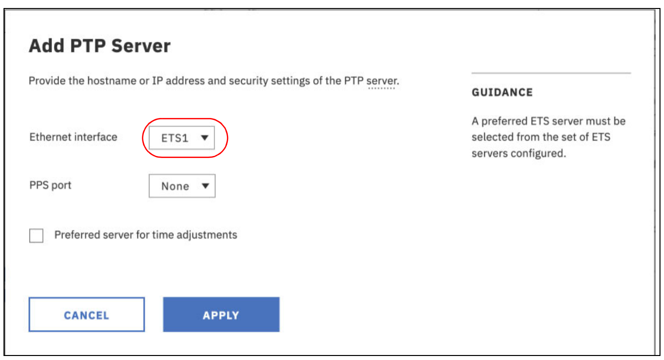

Figure 4-38 Selecting ETS1 port for PTP communication

4. For the PPS port, select **None**.

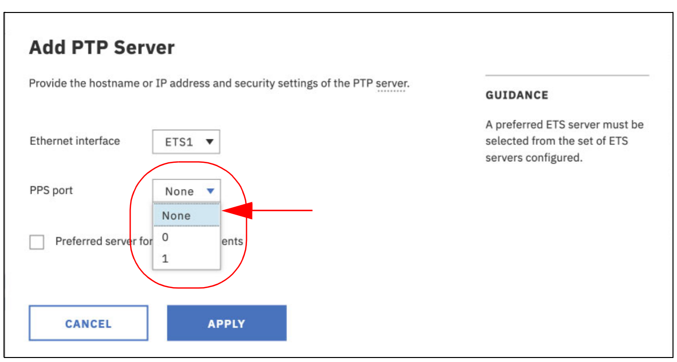

Figure 4-39 Defining the PPS port

5. Click **APPLY**.

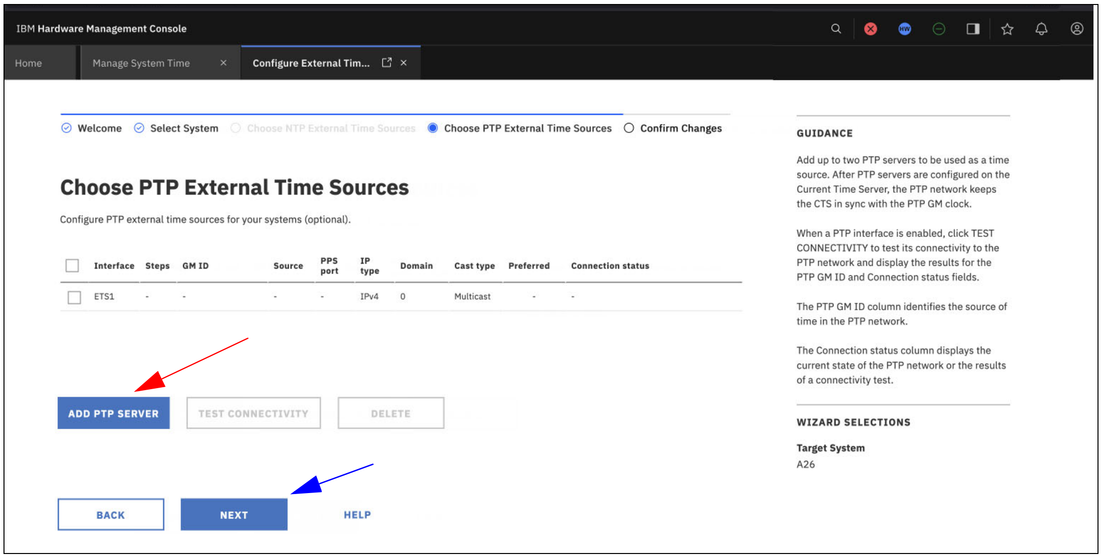

Figure 4-40 First ETS added

6. Add a second PTP server and select **ETS2** as the Ethernet interface server A26.
7. Select **None** as the PPS port.
8. Click **APPLY**.

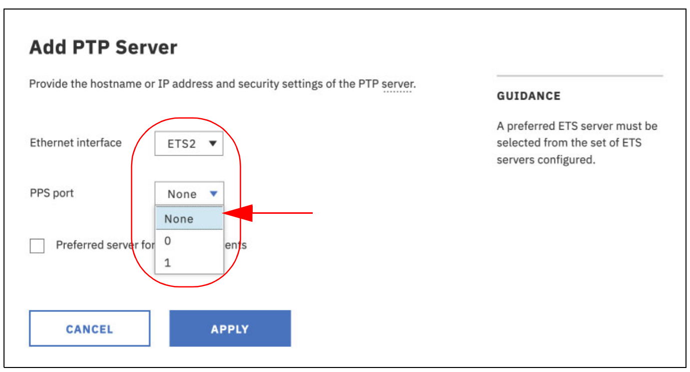

Figure 4-41 PPS port selection for ETS2

9. Test the connectivity for both ETS servers (Figure 4-42 on page 89). To test server connectivity, select the checkbox for the servers.
10. Click **TEST CONNECTIVITY**.

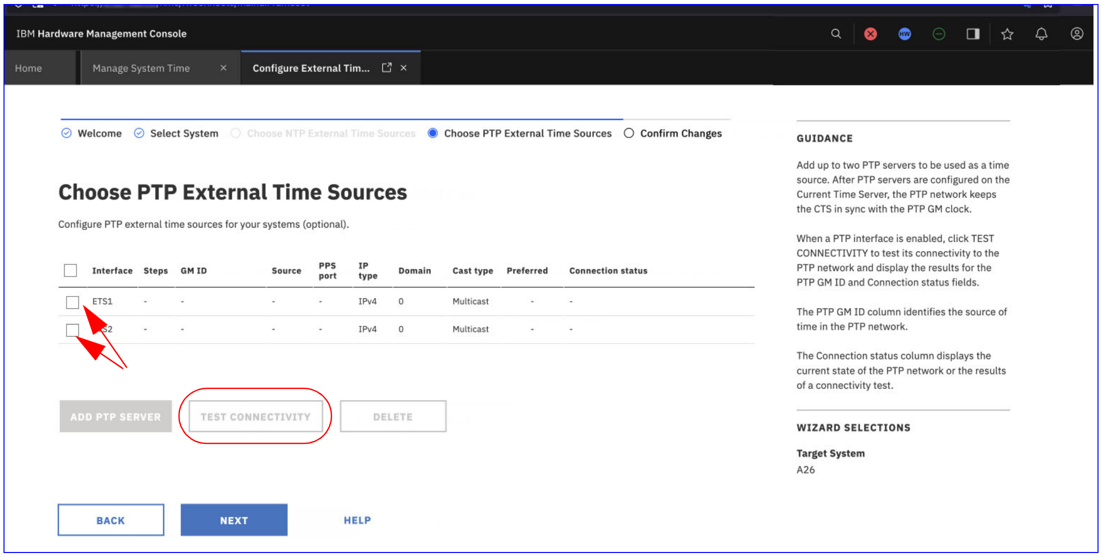

Figure 4-42 Selecting servers to test ETS (PTP) connectivity

11. Verify the connectivity to the ETS servers (Figure 4-43).
12. The preferred time source was not defined during the initial configuration of the ETS servers. Select ETS2 and edit the configuration.

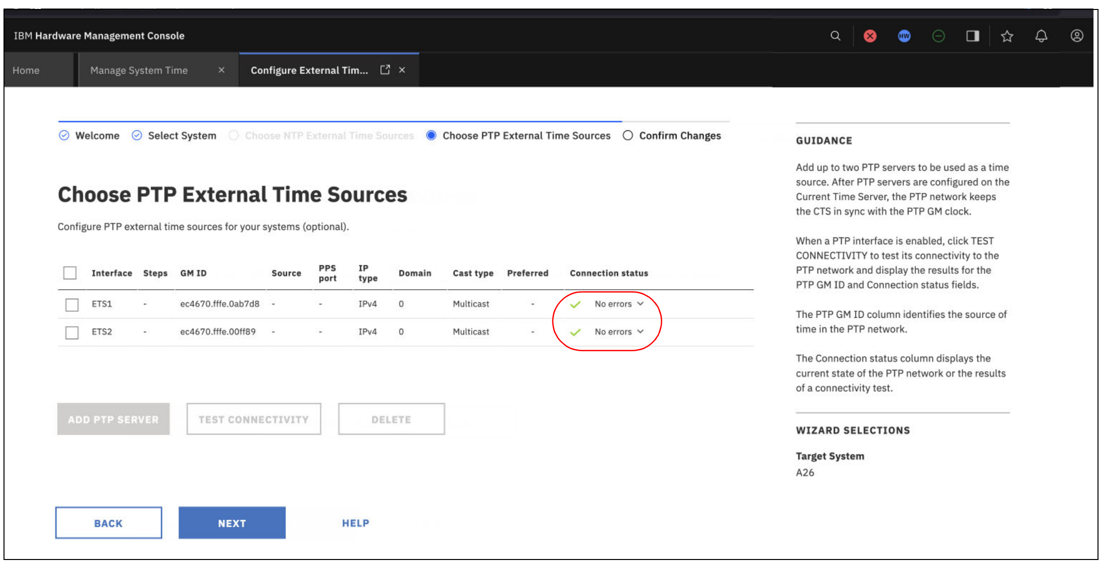

Figure 4-43 test ETS (PTP) connectivity results

13. Select the box **Preferred server for time adjustments (**Figure 4-44 on page 90) and click **APPLY**.

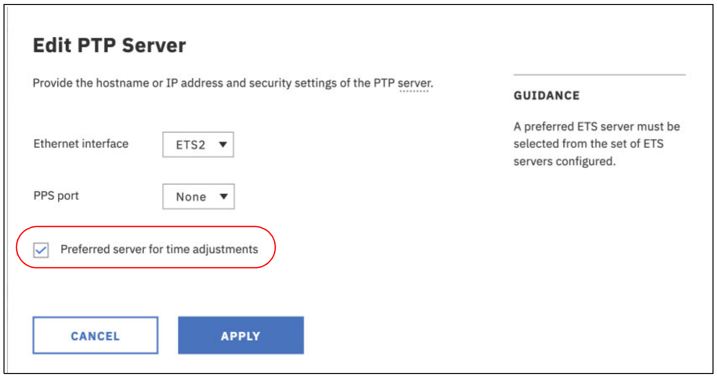

Figure 4-44 Preferred time source selection

14. After you select the preferred time source, click **NEXT** (Figure 4-45) to commit the changes.

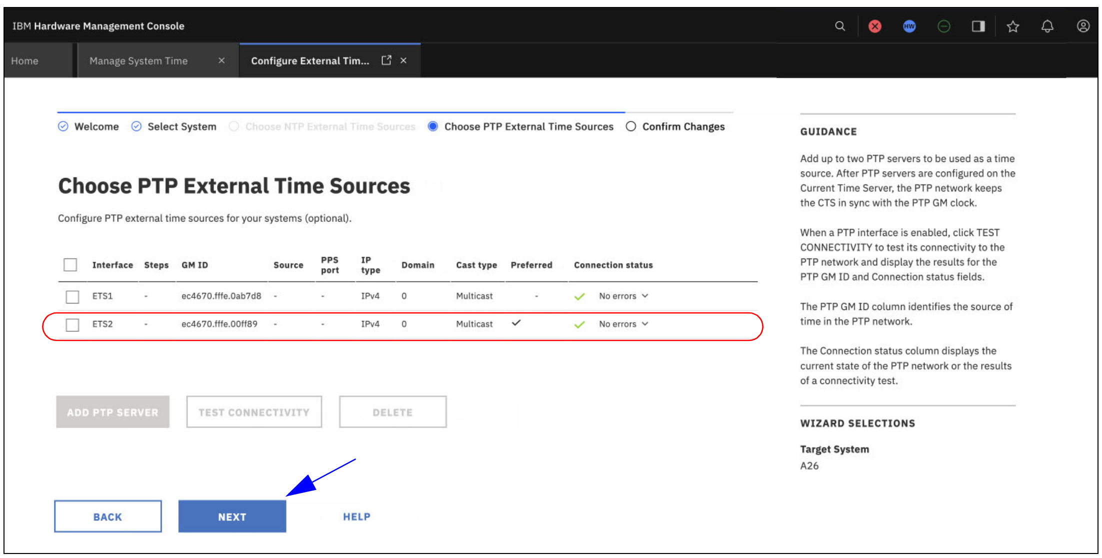

Figure 4-45 ETS using PTP configuration before committing changes

15. Review the changes (Figure 4-46 on page 91) and click **APPLY**. Because the NTP servers were used in a previous configuration, the previous ETS configuration is shown as well.

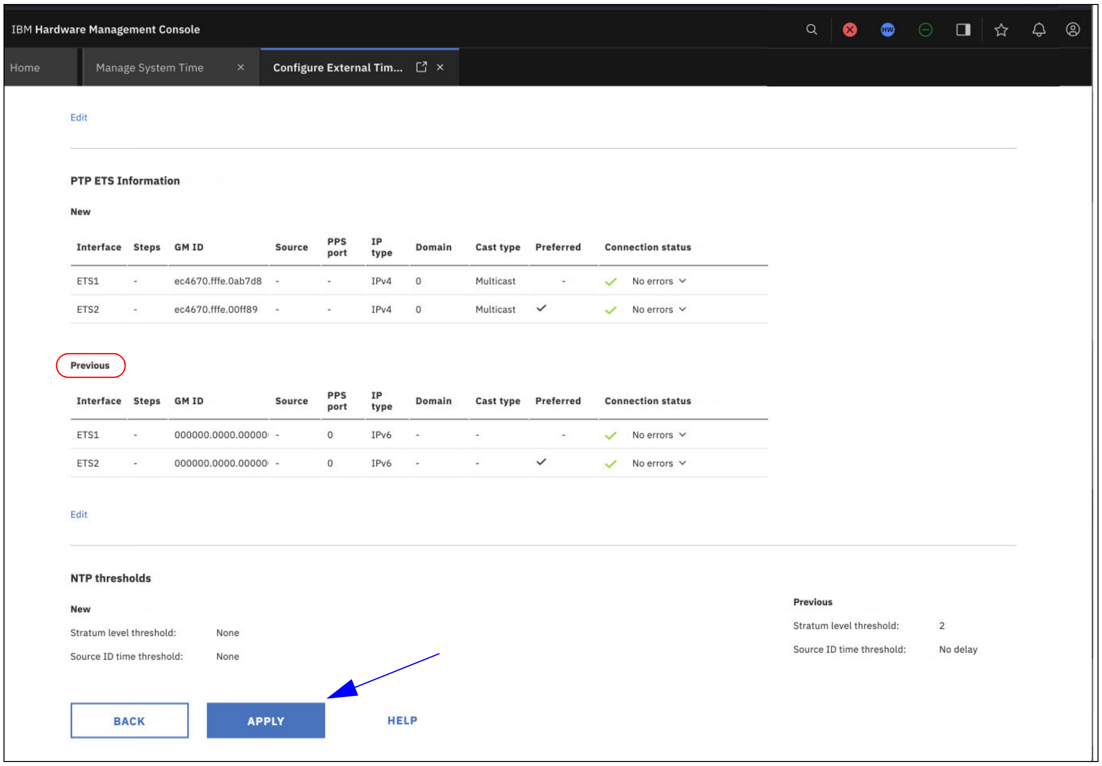

Figure 4-46 ETS changes overview before commit

16. After you click **APPLY**, the Configuring ETS message is displayed (Figure 4-47).

Figure 4-47 Configuring ETS progress

17. Verify that the configuration was successful and click **CLOSE** (Figure 4-48).

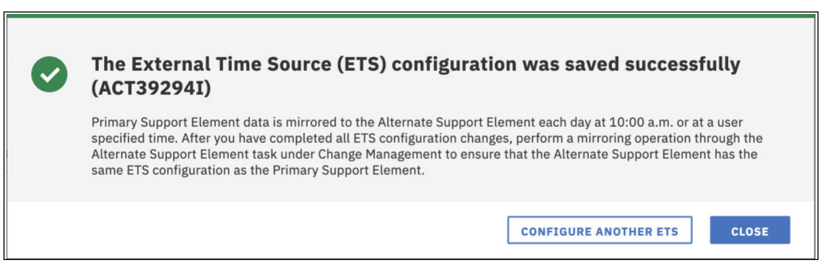

Figure 4-48 Successful ETS configuration

4.3.3 ETS for IBM z15

IBM z15 ETS is provided through the Support Element. Support Element network connectivity is described in the *IBM z15 Installation Manual for Physical Planning*, GC28-7002.

4.4 ETS status

You can view the status of the ETS by using monitoring commands and viewing the audit log. These can be helpful when debugging issues.

4.4.1 ETS monitoring commands

This task is used to obtain detailed ETS status on IBM z17 and later generation systems. This can be helpful when debugging connectivity issues. To run ETS monitoring commands, which are run on the BMC/OSC, use the following steps:

1. Select the task from the Manage Systems Time main menu to view the Run ETS Monitoring Commands screen (Figure 4-49). Each task in the list also includes the associated command for reference.

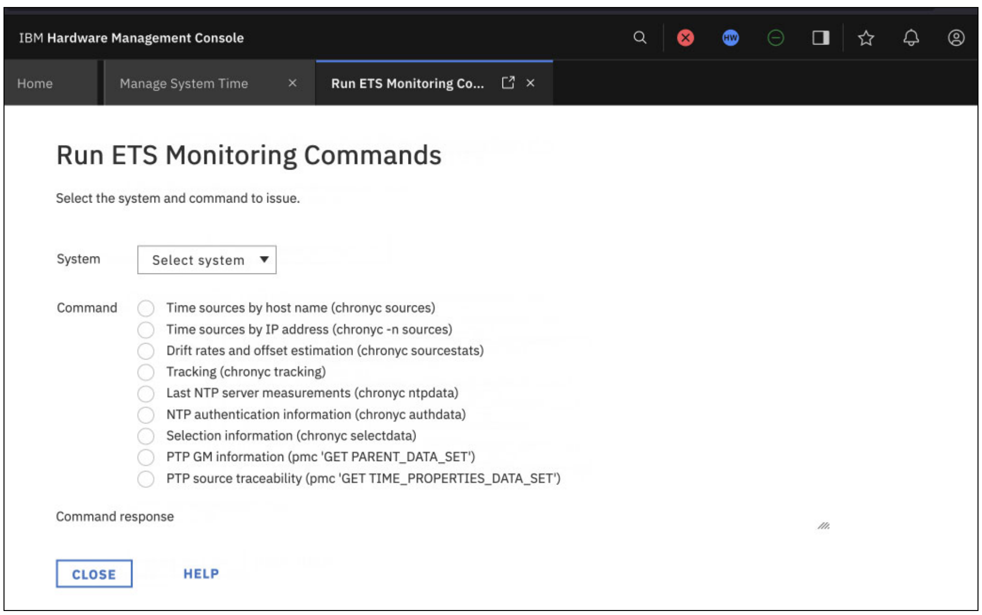

Figure 4-49 Run ETS Monitoring Commands main menu

2. Select the system for which you want to monitor the ETS and the aspect you want to monitor. Figure 4-50 on page 93 shows the time sources configured for the IBM z17 B229 system.

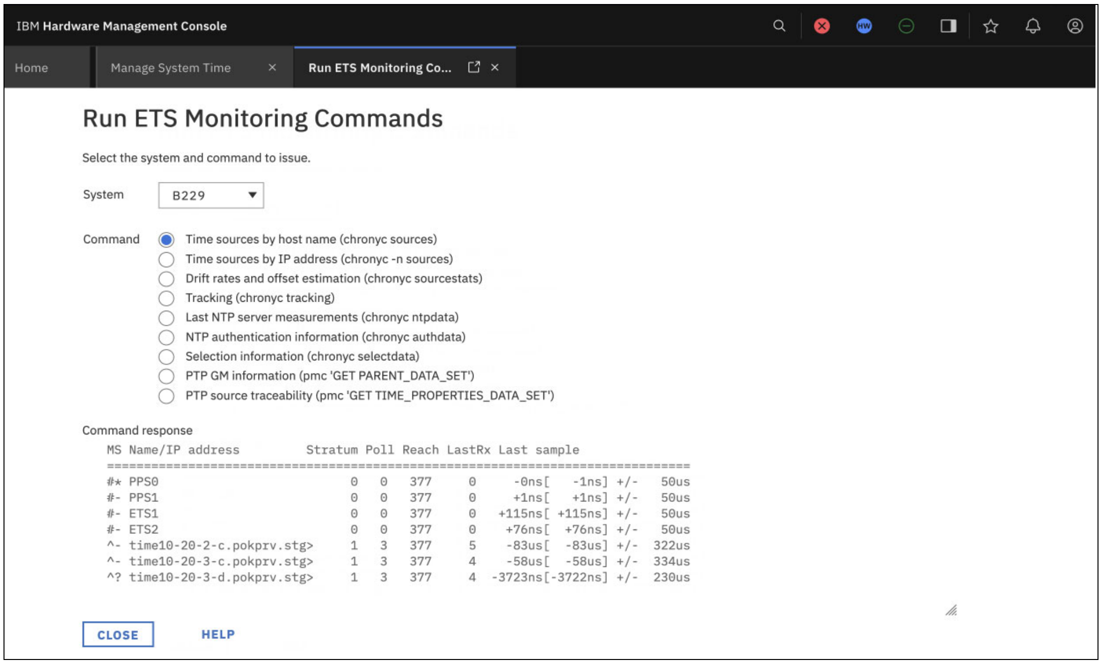

Figure 4-50 System B229 and its time sources by hostname

3. The results of selecting **Time sources by IP address** are shown in Figure 4-51.

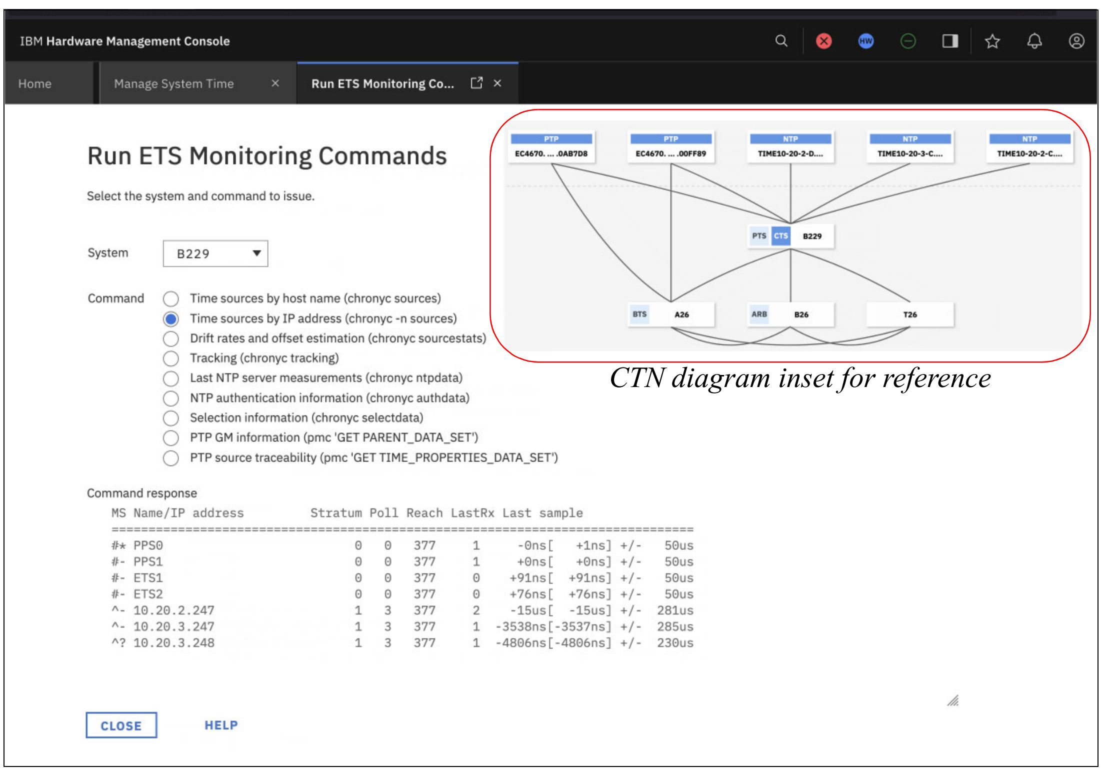

Figure 4-51 Time sources by IP address (CTN diagram inset for reference)

4. Time source drift rates information is shown in Figure 4-52.

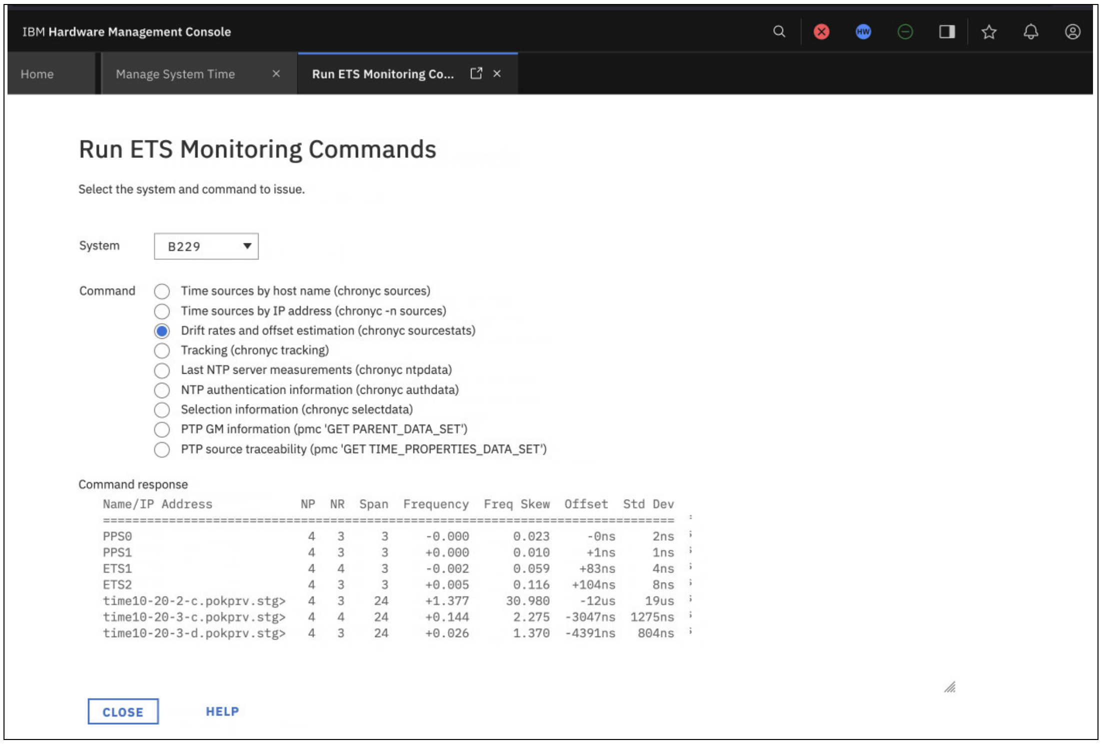

Figure 4-52 Drift rates information for configured time sources

5. Last NTP server measurements are shown in Figure 4-53

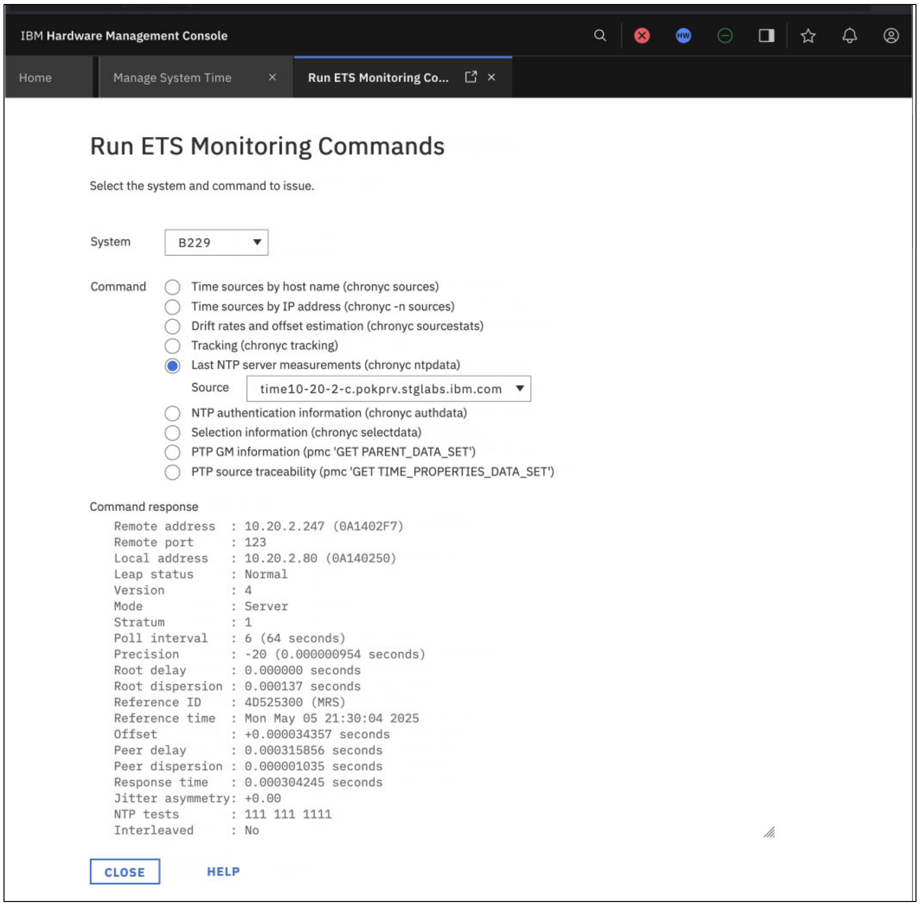

Figure 4-53 Last NTP server measurements

6. NTP authentication information (NTS) data is shown in Figure 4-54 on page 96.

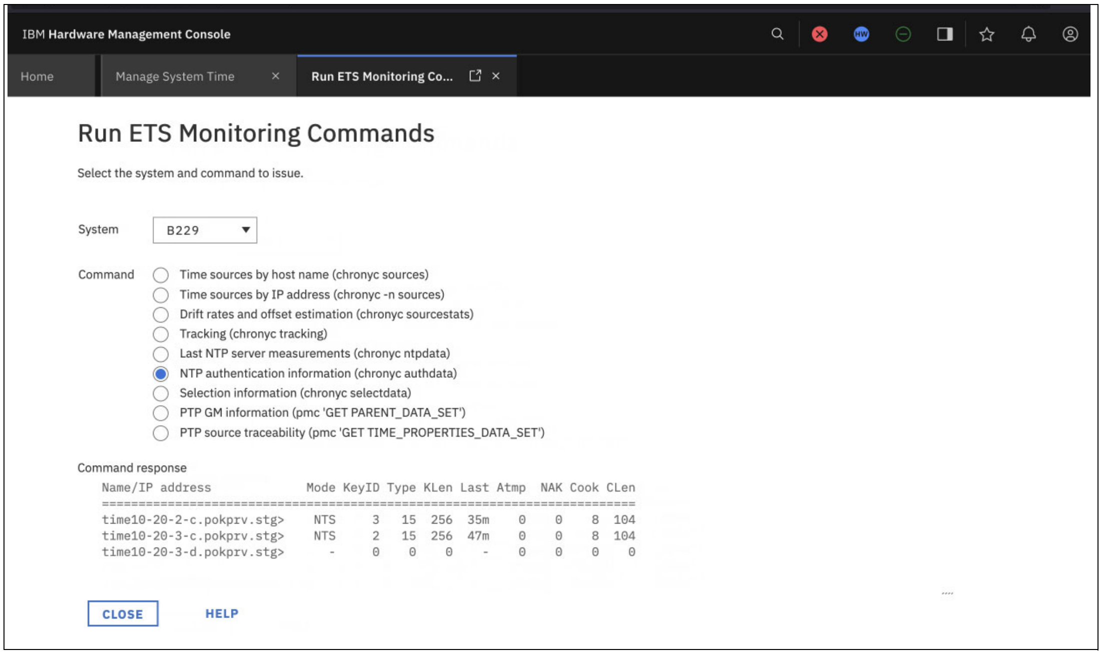

Figure 4-54 NTS information for configured NTP sources

7. Time source selection information is shown in Figure 4-55.

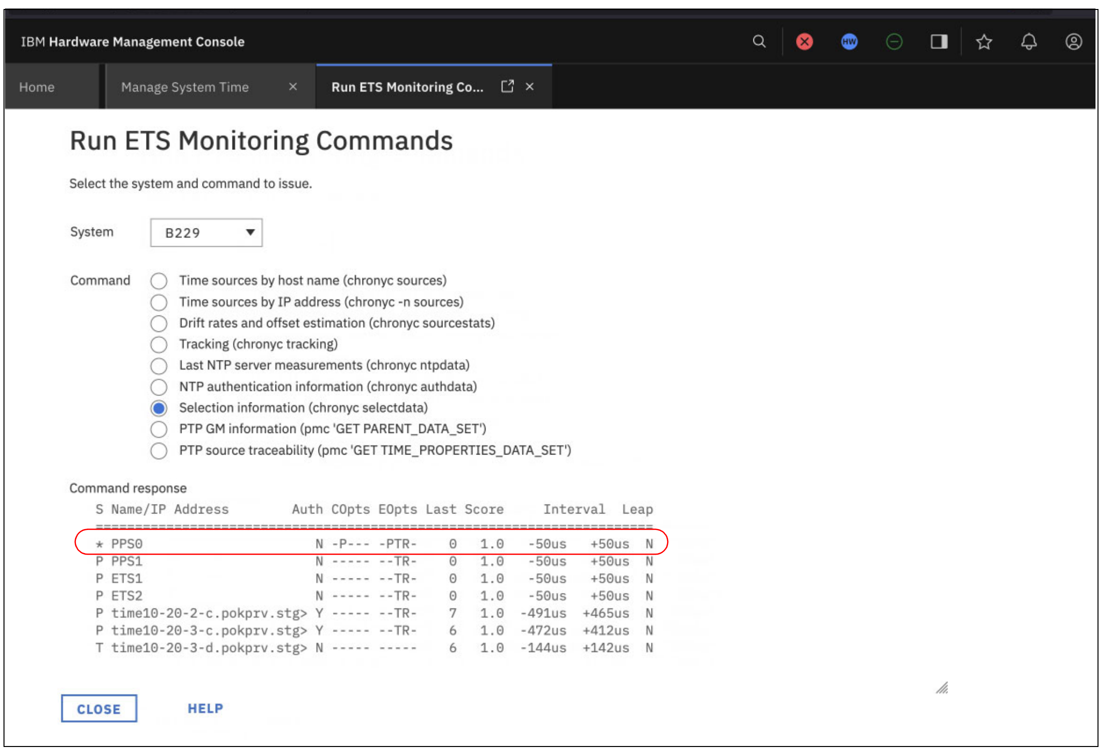

Figure 4-55 Selected time source displayed

8. PTP Grandmaster server information is shown in Figure 4-56.

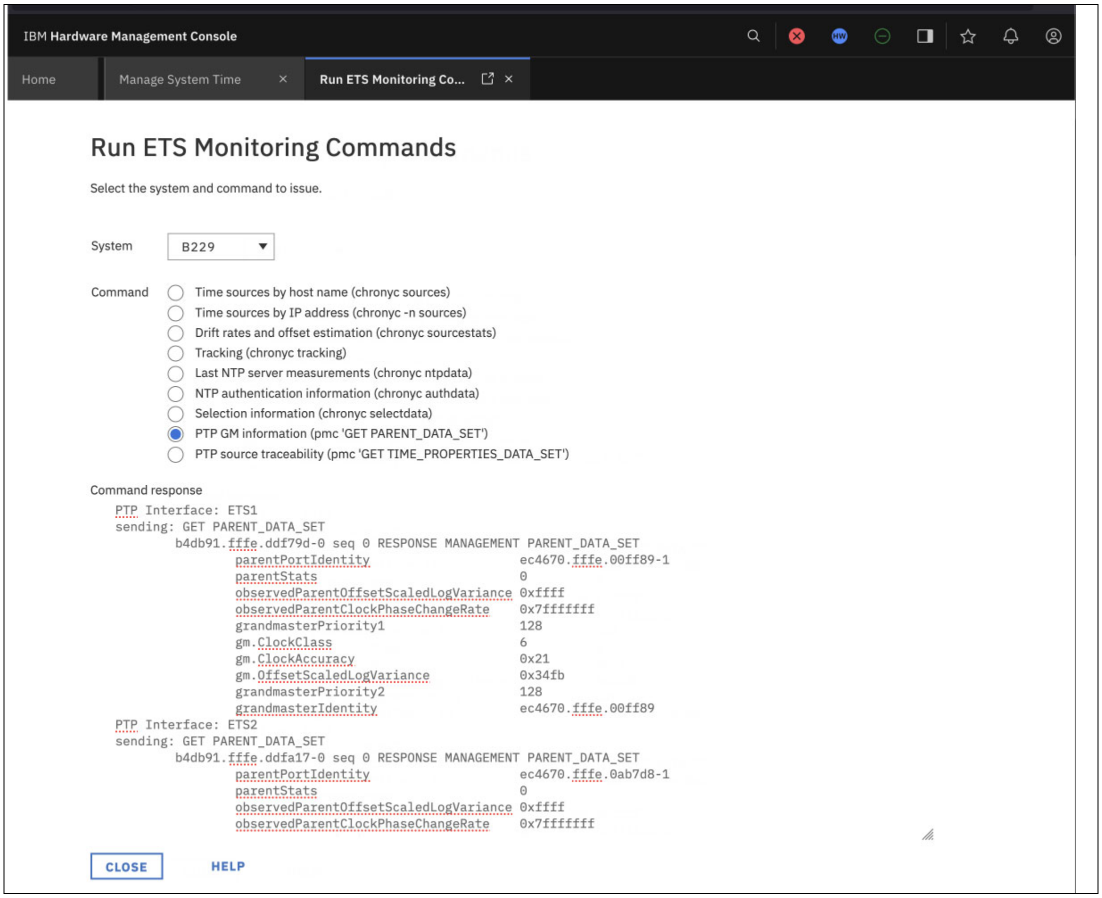

Figure 4-56 PTP Grandmaster information

9. PTP source traceability is shown in Figure 4-57.

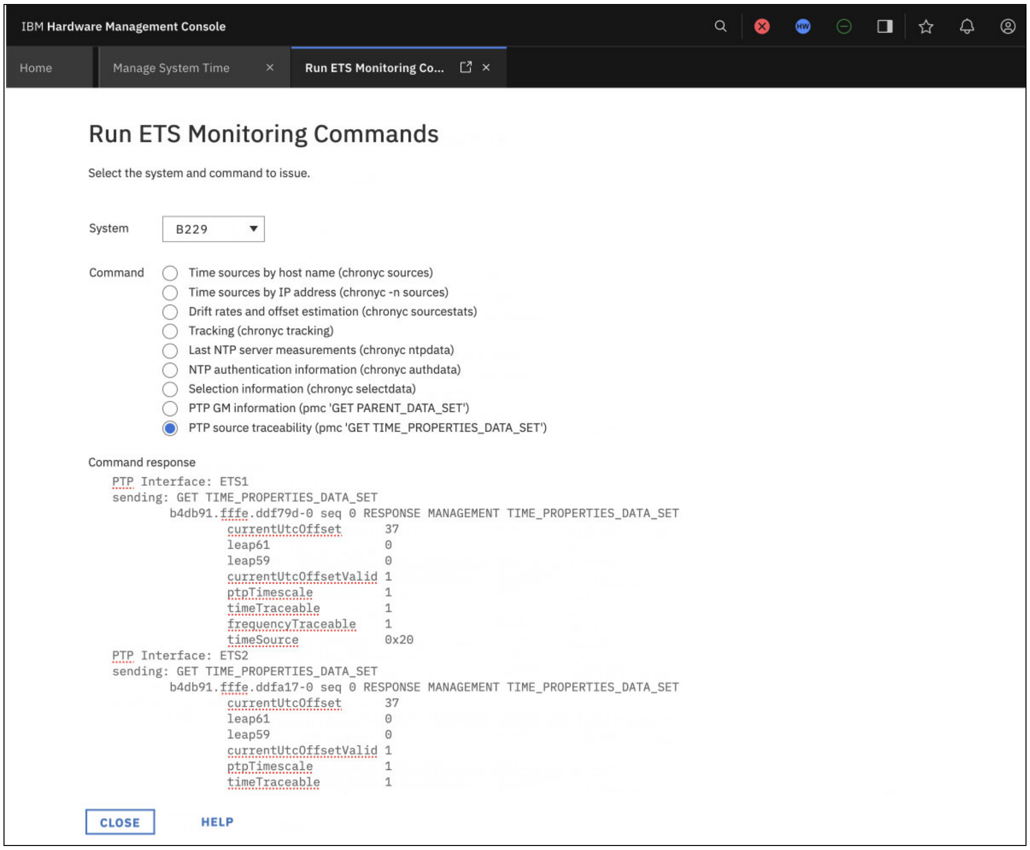

Figure 4-57 PTS source traceability

4.4.2 Audit log

Audit log messages are available through the use of IBM Support Element (Single Object Operations task on the HMC - with appropriate security role).

Use the following steps to access the audit log:

1. In the IBM Support Element, select the Audit Log and Management task (Figure 4-58).

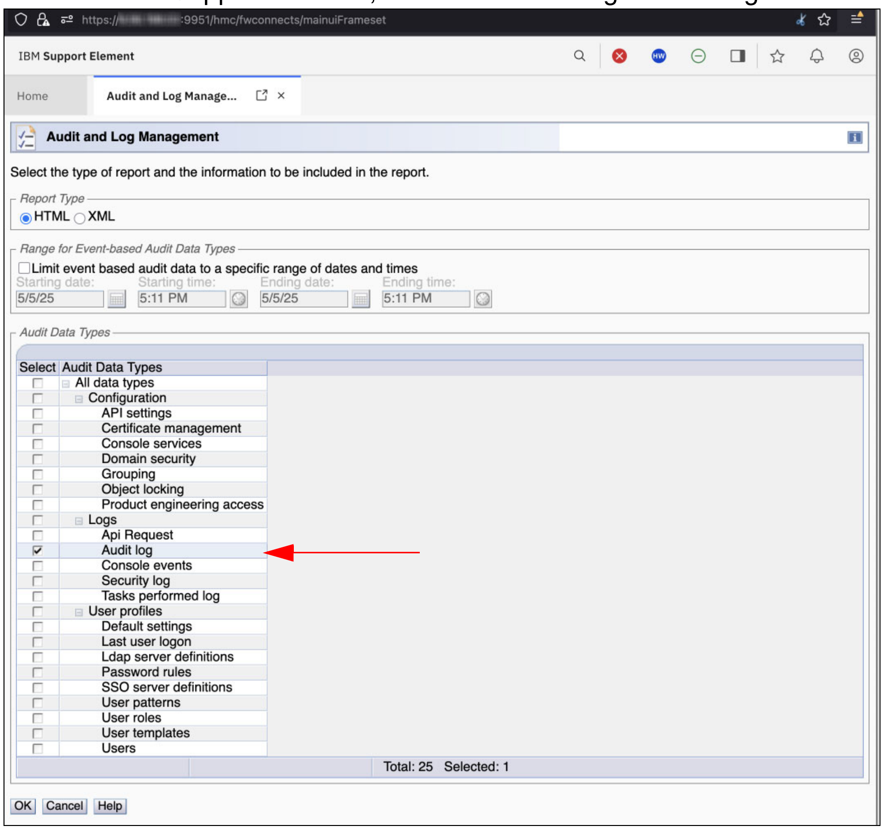

Figure 4-58

2. Look for the messages related to time adjustment such as message ID 1110, as shown in Figure 4-59 on page 100.

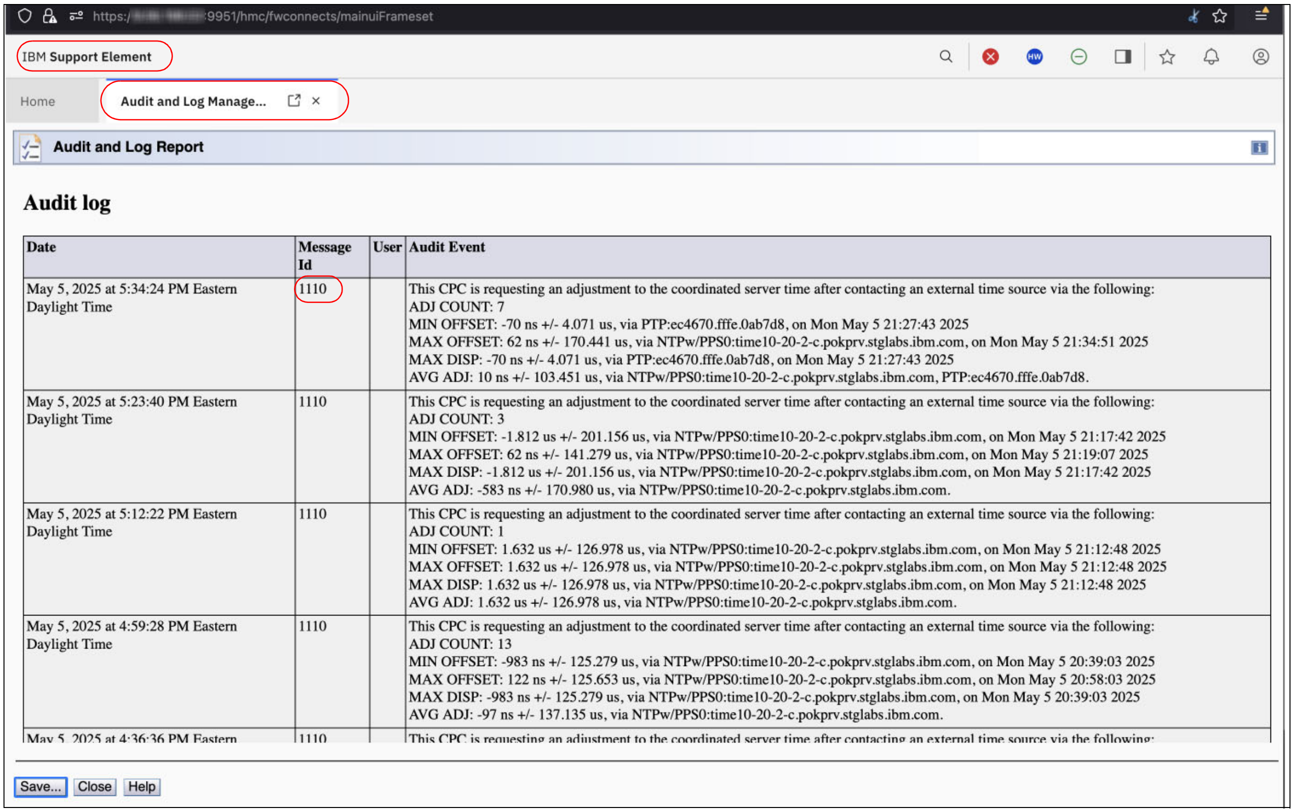

Figure 4-59 Audit Log sample

4.5 Control Pulse Per Second signal

This section provides guidance and procedures to ensure successful PPS connectivity. You can use HMC tasks to verify that the PPS connections are established.

4.5.1 Best practices for implementing PPS connectivity

Use the following guidelines to ensure that your PPS connectivity is working properly:

1. Connect PPS cables

   Coaxial PPS cables are connected point-to-point from your time source appliance to the BMC/OSC card on your CPC. You can run a maximum of two PPS cables to your IBM Z CPC from two time-source appliances.

2. From the Manage System Time task, access the Control Pulse Per Second (PPS) Signal wizard to confirm PPS is operational and detected. If not, do the following to confirm the PPS signal is operational:

   a. Confirm the PPS output is active on time source appliance.

   b. Confirm all coaxial connections, and cable lengths maximums.

   c. Disconnect cables for IBM Z CPC and physically verify the signal at the cable connector by using, for example, an oscilloscope.

   d. Disconnect cables and run a PPS port test in the Control Pulse Per Second (PPS) Signal wizard (see Figure 4-65 on page 103). This confirms that the port is working.

e. Reconnect the cable and repeat Step 2 on page 100.

3. Configure ETS with PPS as described in 4.3, “Configure ETS” on page 74.
4. Use ETS monitoring menus on z17 only to verify whether PPS is functional. For z16 and earlier, check the “i” for the link to the NTP source on the topology map. See 4.4, “ETS status” on page 92.

For Pulse Per Second cabling, see IBM Z documentation and refer to the Installation Manual for Physical Planning (IMPP) for your specific hardware.

4.5.2 Testing and verifying PPS connectivity

To verify the connectivity of the PPS, use the following instructions:

1. From the Manage System Time task, click Control Pulse Per Second (PPS) Signal to access the wizard shown in Figure 4-60.

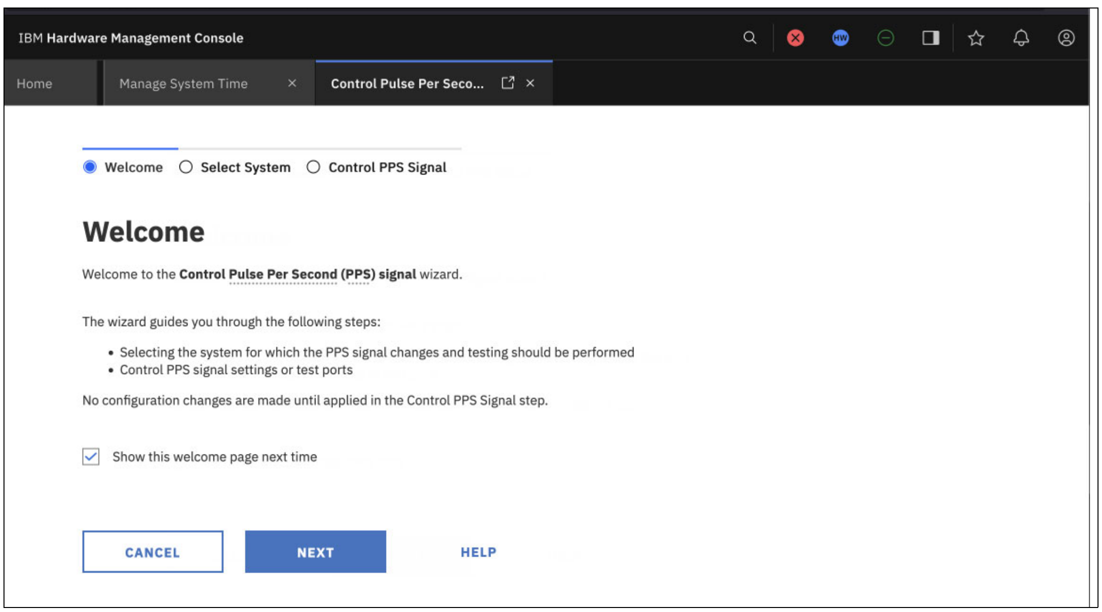

Figure 4-60 Control Pulse Per Second Signal main menu

2. Click the **Select System** button and click **NEXT**.
3. Select the system for which you want to test the PPS signal. The example uses server B229, as shown in Figure 4-61 on page 102.

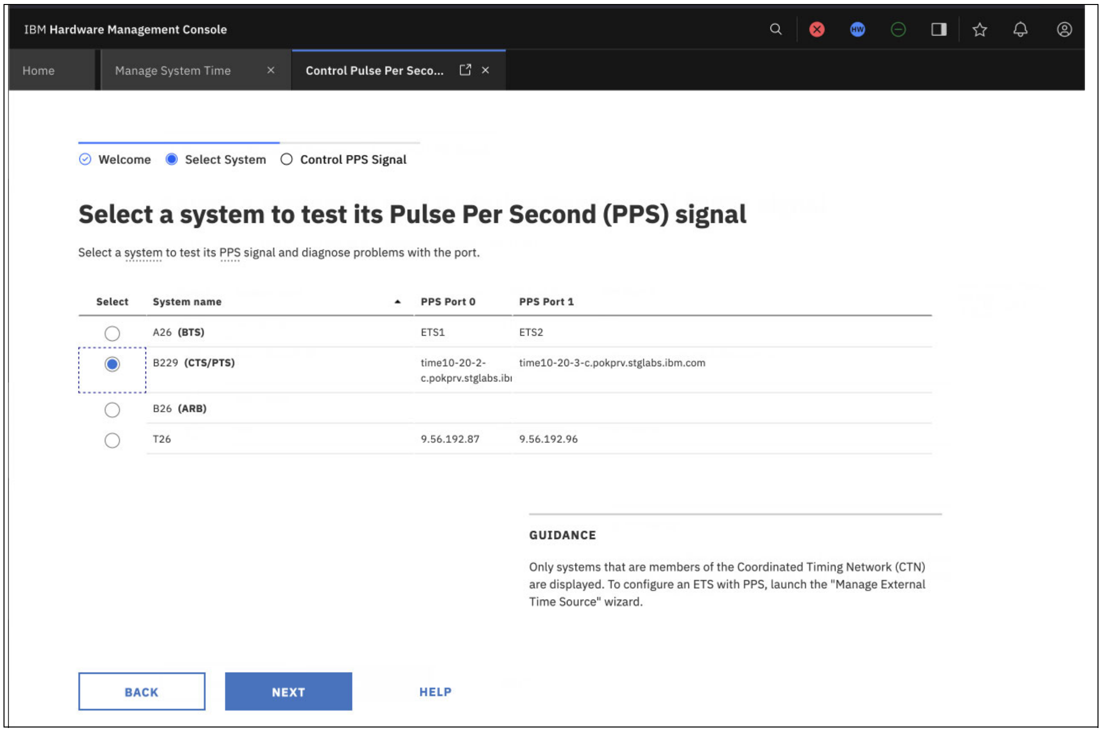

Figure 4-61 Select system to test PPS

4. If a port has a PPS port shows that a PPS port is not connected PPS, you can run a diagnostics test by selecting the button **Perform internal diagnostic test**, as shown in Figure 4-62.
5. Select **INTERNAL TEST**.

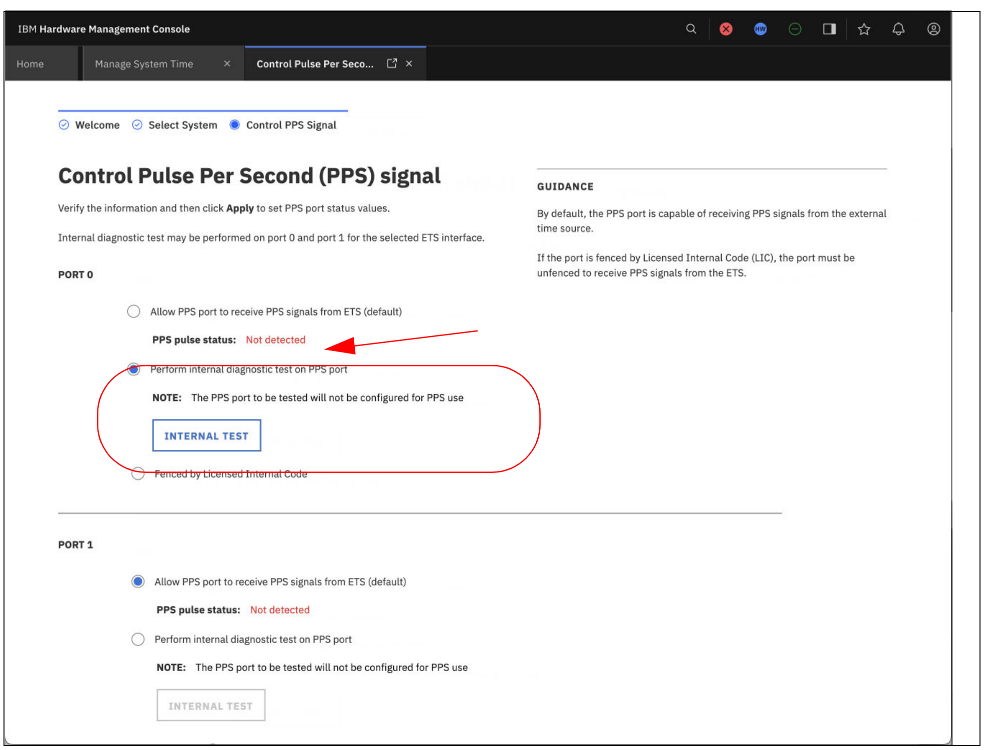

Figure 4-62 Select internal diagnostic test

6. In the confirmation window, click **CONTINUE** (Figure 4-63 and Figure 4-64.

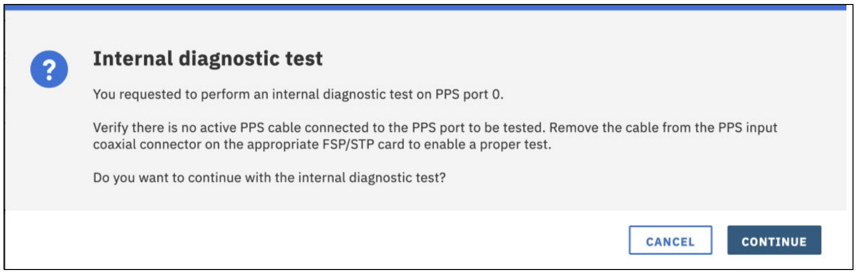

Figure 4-63 Internal diagnostic test - confirm that no cable is connected

Figure 4-64 Running Internal diagnostic progress

7. When the Internal diagnostic test is performed successfully, the message shown in Figure 4-65 is displayed.

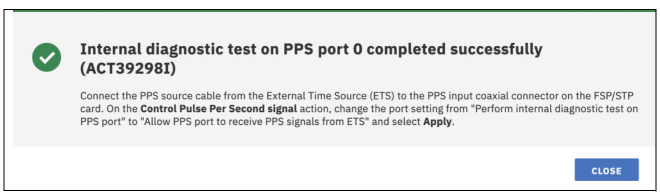

Figure 4-65 Internal diagnostic test successful

8. Reconnect the PPS cables to their ports.
9. To receive a signal, ensure that the default **Allow PPS port to receive PPS signals from ETS** is selected and click **APPLY** to update the port status. (Figure 4-66 on page 104).

   Repeat the step for both PPS designated ports.

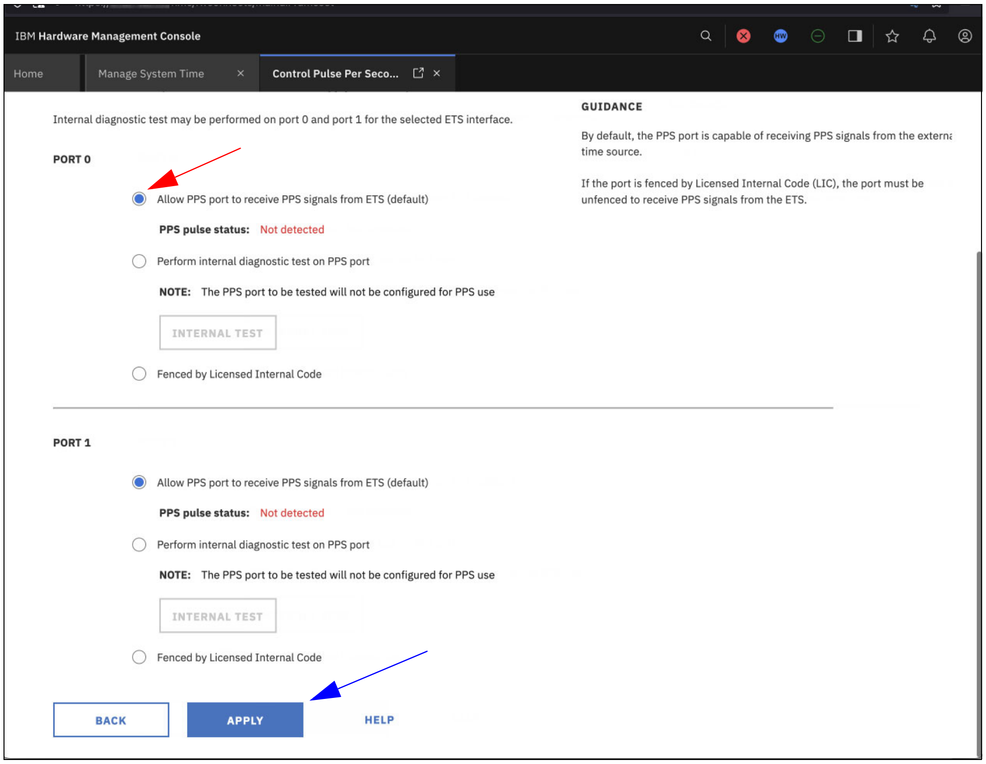

Figure 4-66 Update port status after Internal diagnostic test

10. Verify that cable connectivity and the PPS signal from the time source appliance are functional by verifying that the PPS pulse status is Detected (Figure 4-67).

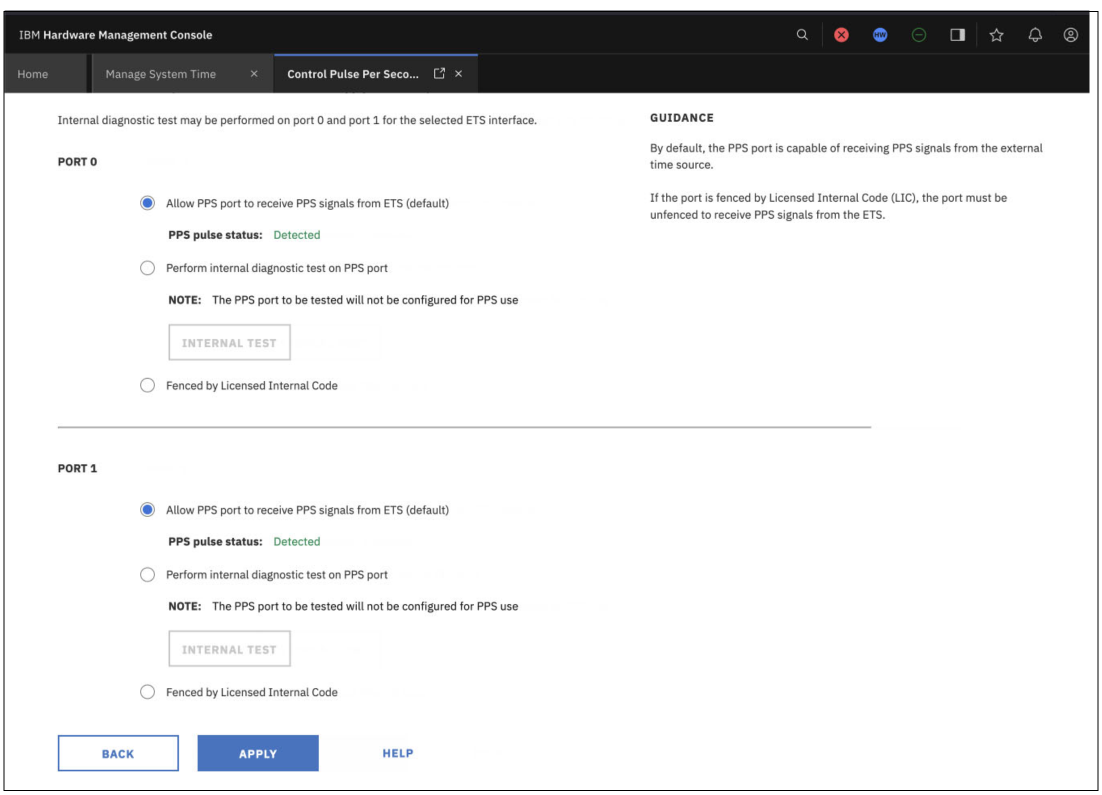

Figure 4-67 PPS cables connected and signal detected

11.Click **APPLY** and the messages in Figure 4-68 are displayed, which confirm that the correct PPS port configuration.

After receiving confirmation, you can configure ETS with PPS. See 4.3, “Configure ETS” on page 74.

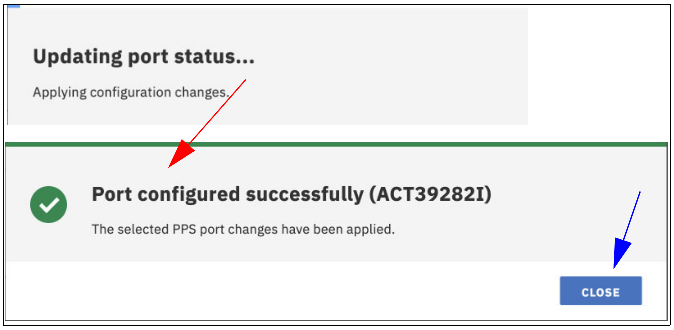

Figure 4-68 PPS connectivity test finalized

4.6 Manage CTN Certificates

This task is used to add x509 certificates used for NTS connectivity. For Additional information, see “NTP security” on page 21.

> **Note:** Manage CTN Certificates (for NTS) is available for IBM z17 and later generations. It is *not* available on IBM z16 and earlier server generations.

To add the x509 certificates, perform the following steps:

1. Access the Manage CTN Certificates menu from the Manage System Time task (Figure 4-69 on page 106).
2. Click **Import**.

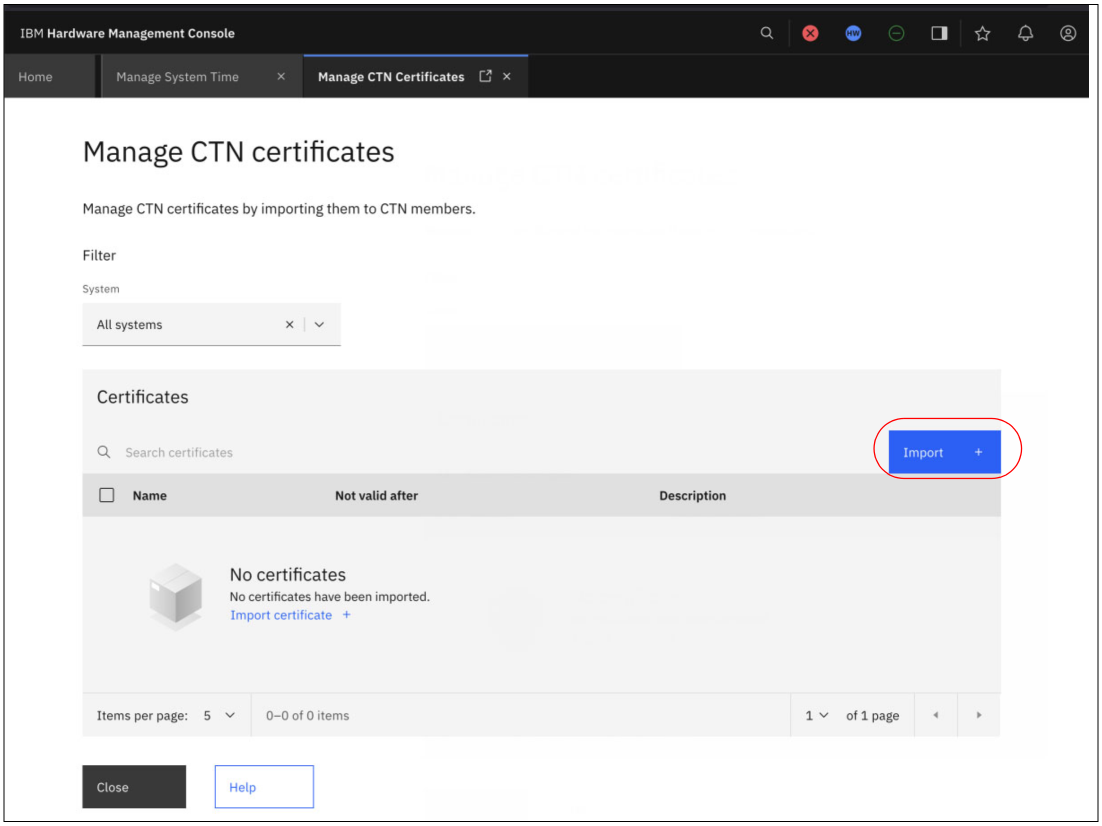

Figure 4-69 Manage CTN Certificates wizard

3. Verify the systems and click **NEXT** to proceed to system selection, as shown in Figure 4-70.

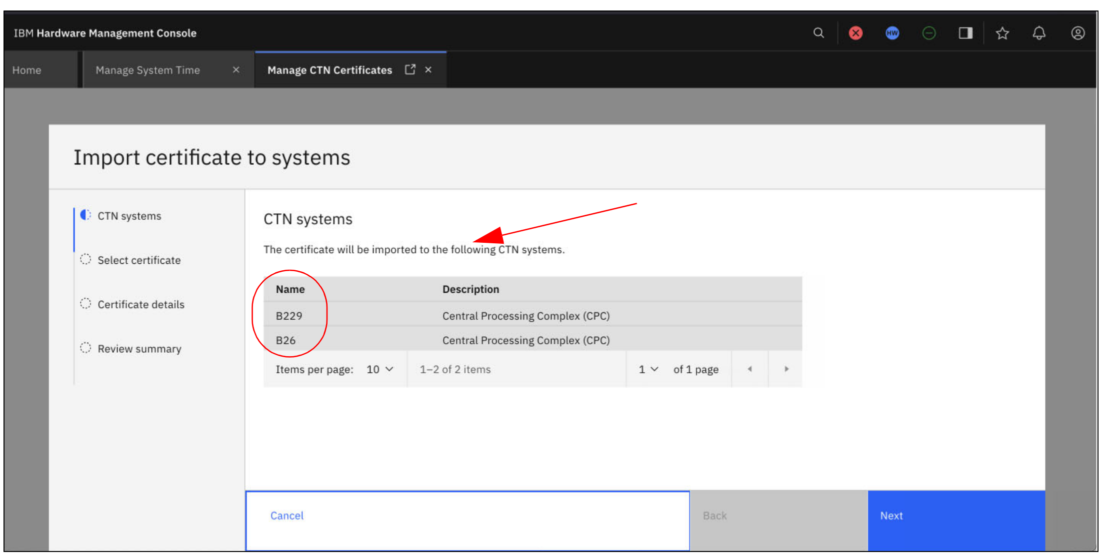

Figure 4-70 Servers that will receive the certificates

4. Verify the selected servers and click **NEXT**.
5. Select a method for uploading the certificates (Figure 4-71 on page 107).

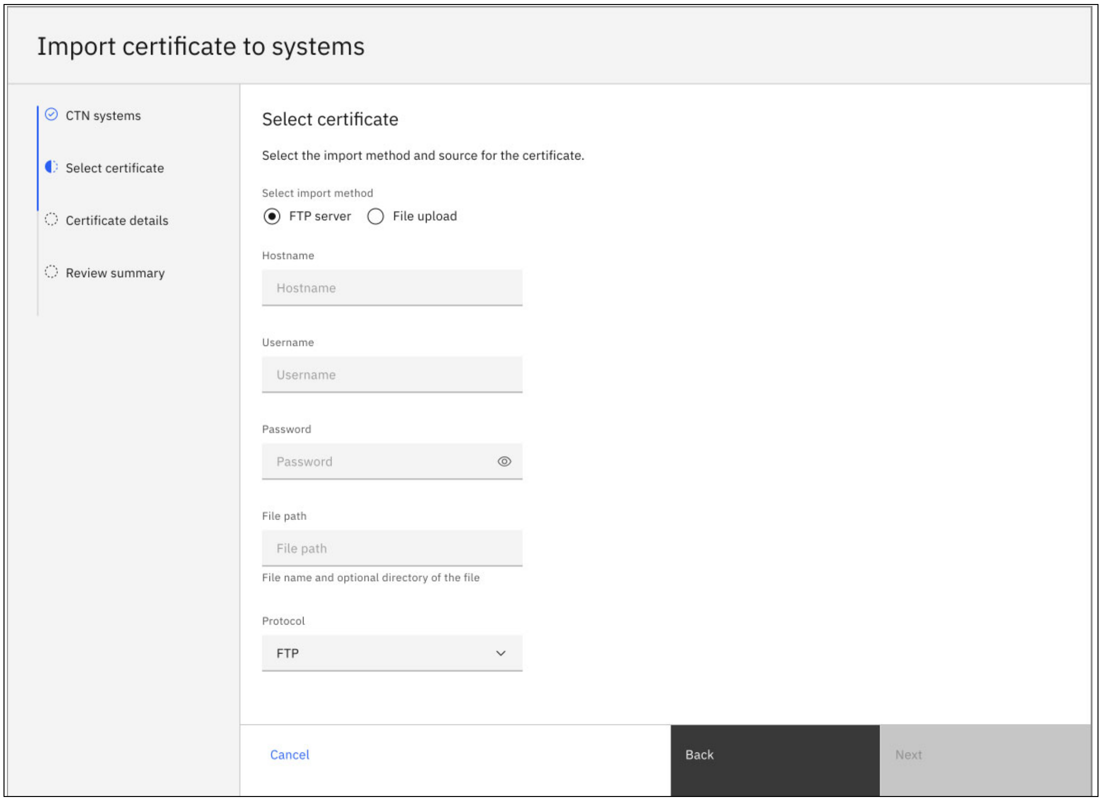

Figure 4-71 Certificate import methods

6. To upload certificates from an existing file, select **File upload** (Figure 4-72) and select **Choose file**.

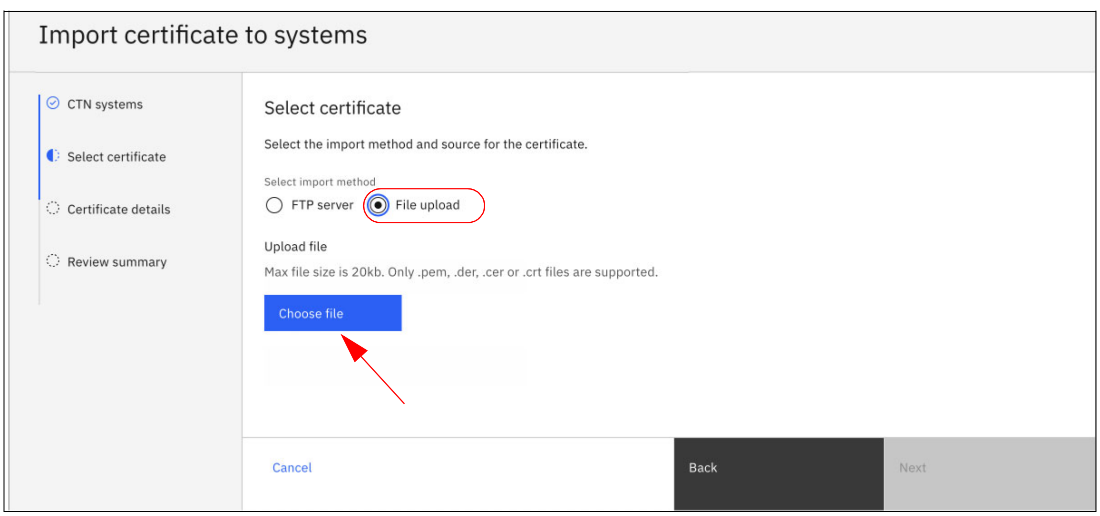

Figure 4-72 Import certificates from file

7. Select the file that contains the certificate (Figure 4-73 on page 108).

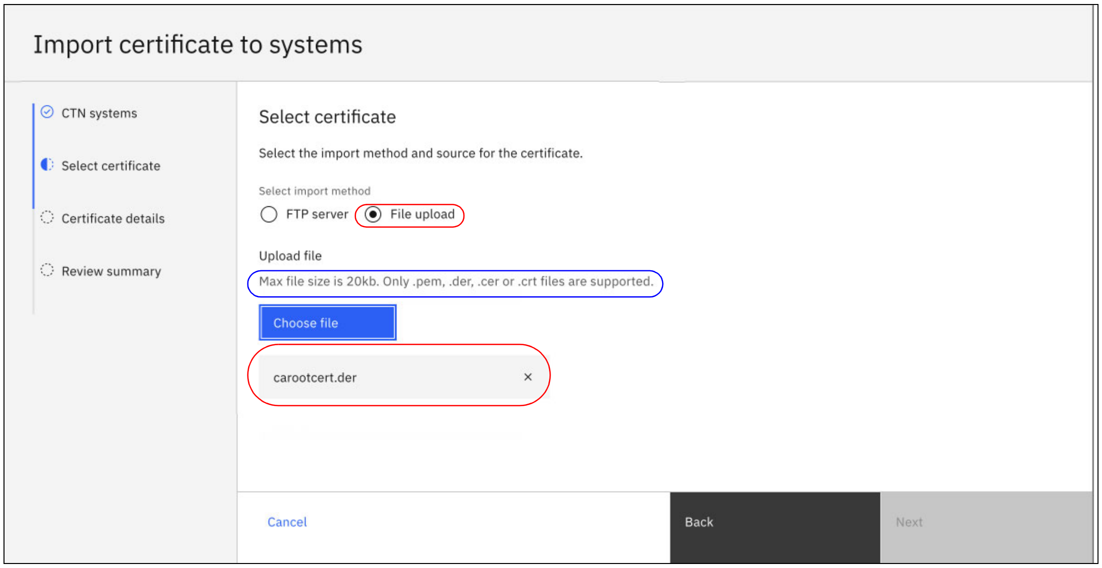

Figure 4-73 File containg certificate

8. Enter a name for the certificate. The name is used for local certificate management (Figure 4-74).

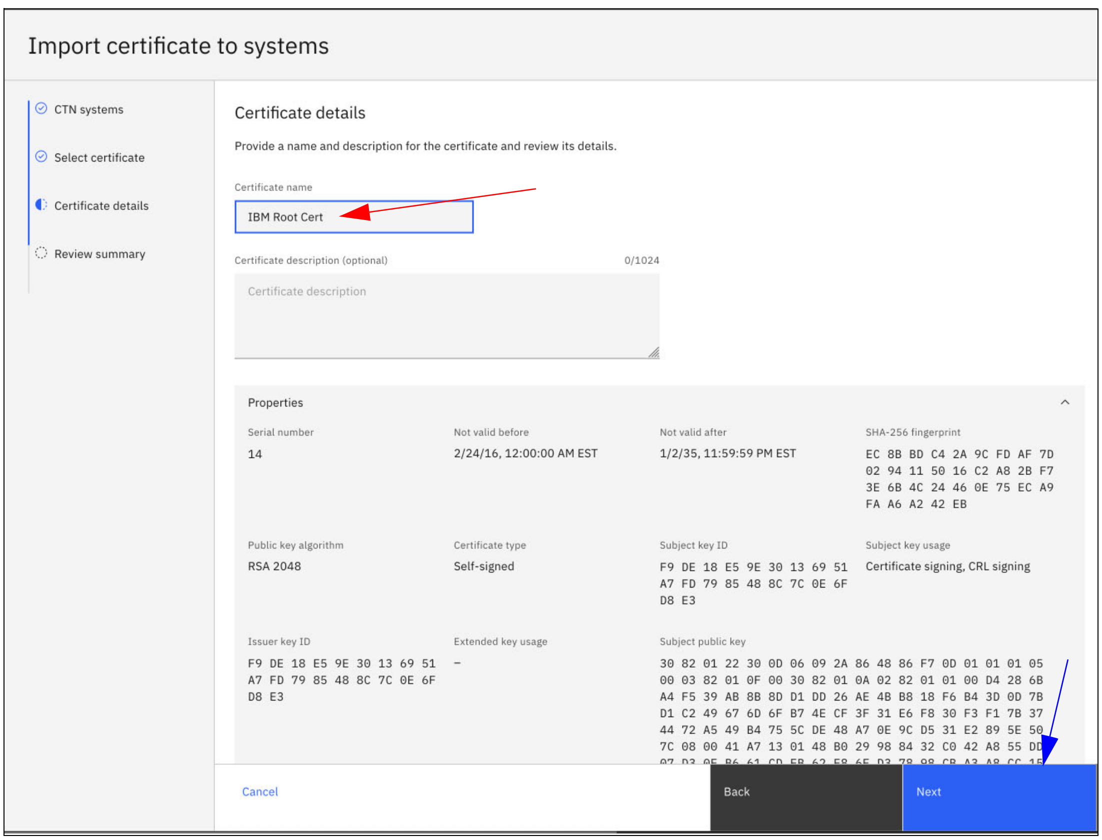

Figure 4-74 View certificate to be imported

9. Click **Next** to upload the certificate (Figure 4-75 on page 109).

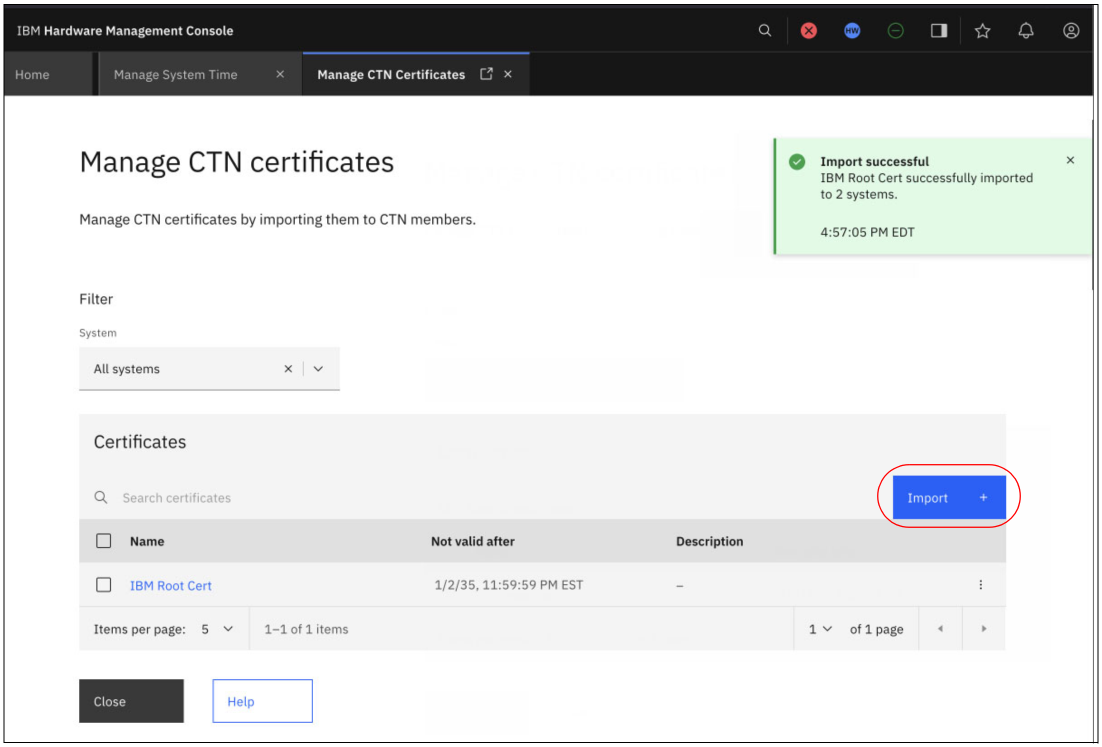

Figure 4-75 Certificate imported

10.If you want to upload multiple certificates, select **Import** and choose another certificate.

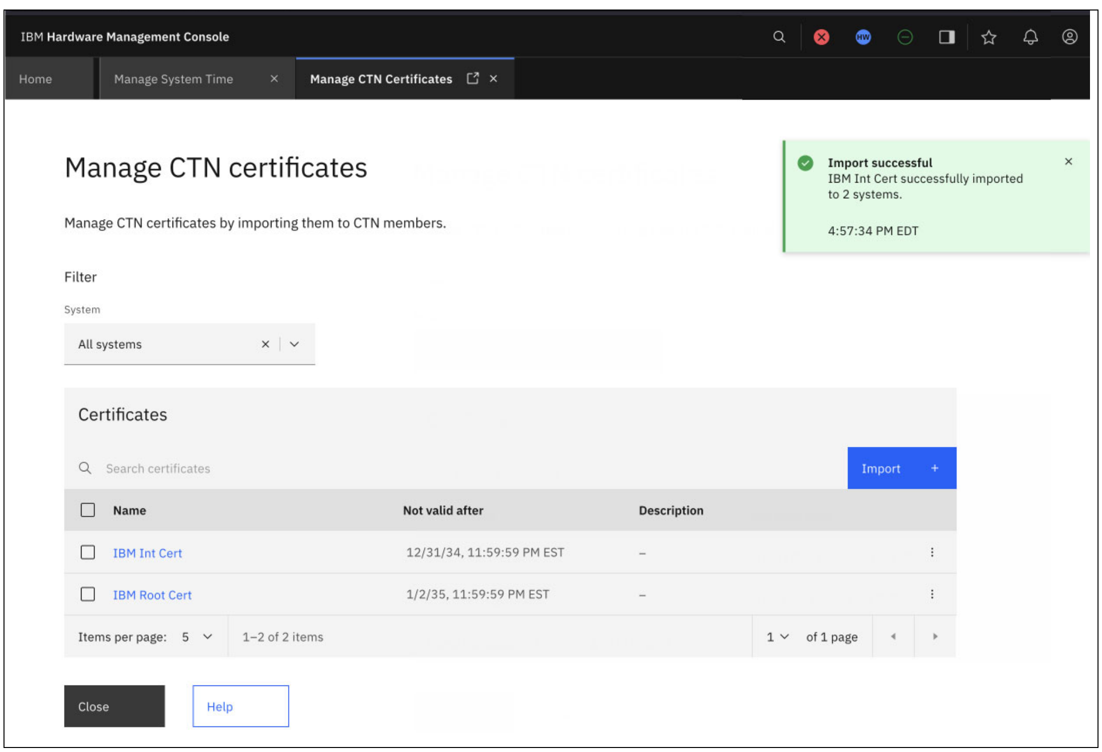

Figure 4-76 Uploaded certificates

11.The certificates uploaded to the example systems are shown in Figure 4-76 on page 109. With all certificates uploaded, select each certificate and review (Figure 4-77).

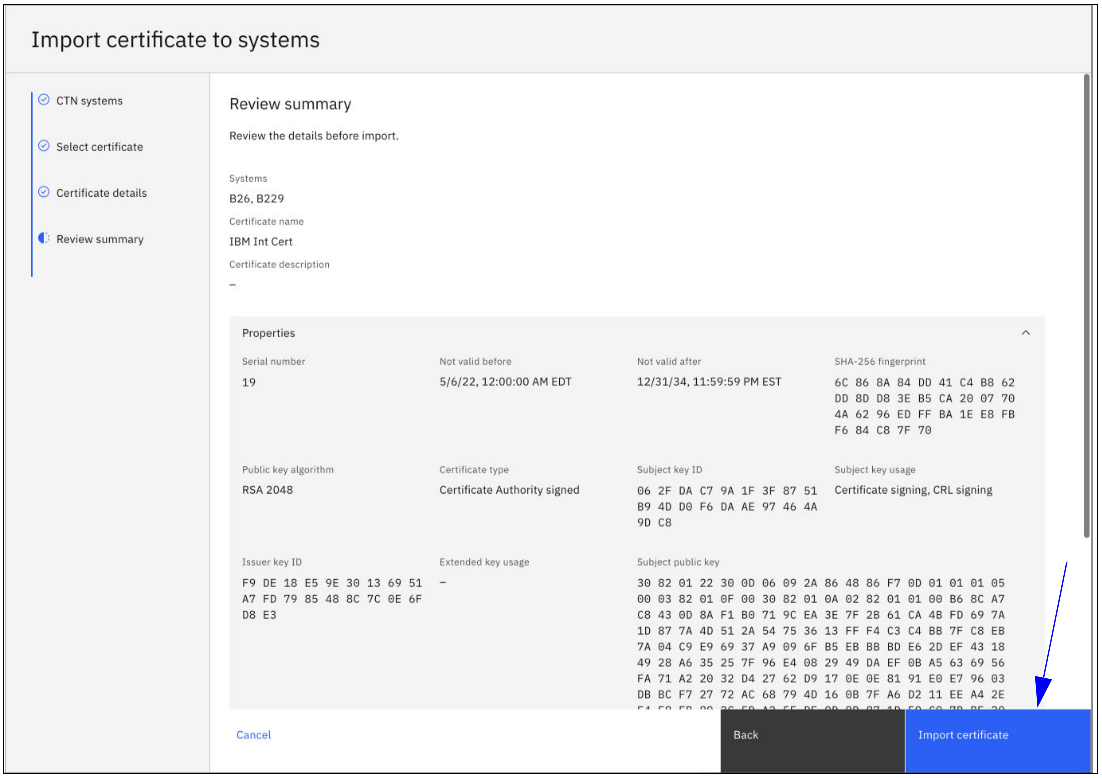

Figure 4-77 Import certificates

12.Click **Import certificate** to import the certificates to the BMS/OSC card.

5

Advanced CTN configurations and time management operations

This chapter describes advanced CTN configuration operations and time management.

This chapter includes the following topics:

- 5.1, “Advanced CTN Actions” on page 112
- 5.2, “Managing the Coordinated Server Time (CST)” on page 123

5.1 Advanced CTN Actions

This section describes the Advanced Actions in the Manage System Time task.

5.1.1 Join existing CTN

Use the join existing task to merge two existing CTNs into one single CTN. The two CTNs to be joined are “26plex” (target CTN ID) and “gatto” (joining CTN), as shown in Figure 5-1.

Before joining systems, consider the following factors:

- The two CTNs to be joined must have their Coordinated System Time values as close to each other as possible to avoid disruption of the workload on the joining CTN, gatto.

  If the CST values for the two CTNs are apart, then for each second of offset, it takes approximately 7 hours of adjustment steering to synchronize the joining servers to the same CST value.

  For additional details, see 2.3.7, “ETS steering” on page 31.
- Ensure that the coupling connectivity exists between the CTS of the target CTN and the CTS of the joining CTN.

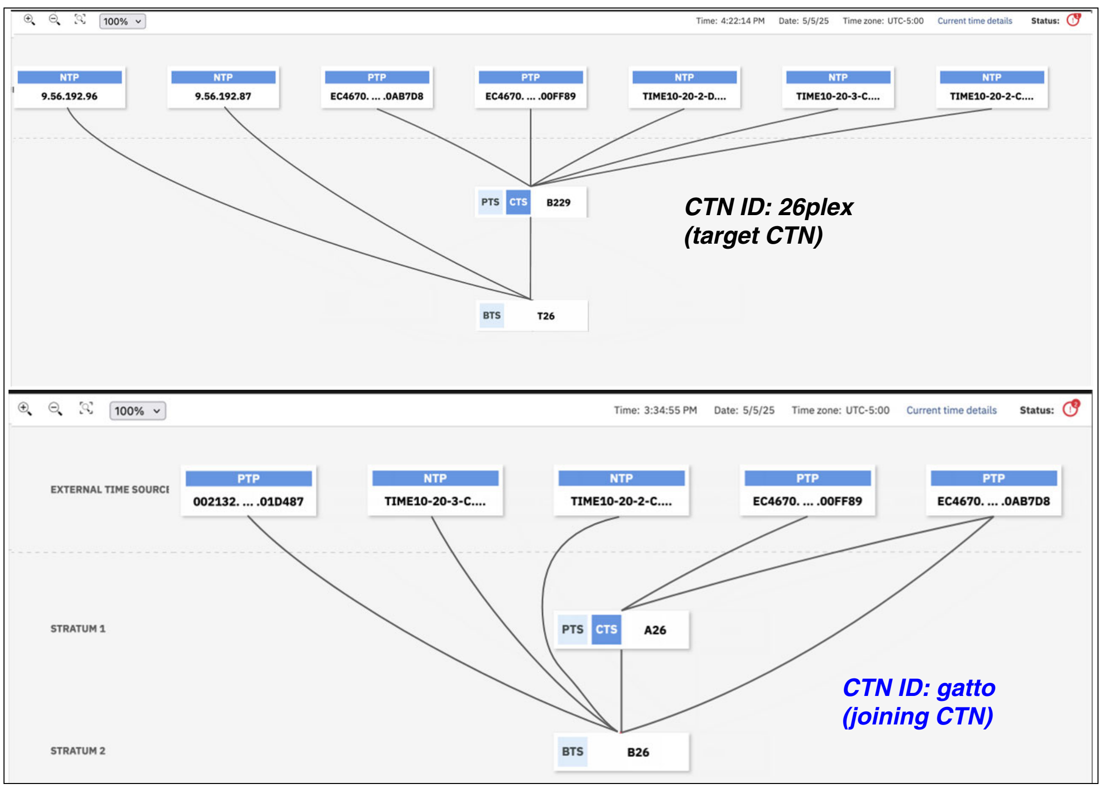

Figure 5-1 The two CTNs to be joined

Use the following steps to join the CTNs:

1. To initiate the process, on the Manage System Time task of one of the servers in the joining CTN gatto, click the **Join Existing CTN** menu.
2. From the **Welcome** menu (Figure 5-2 on page 113), click **NEXT**.

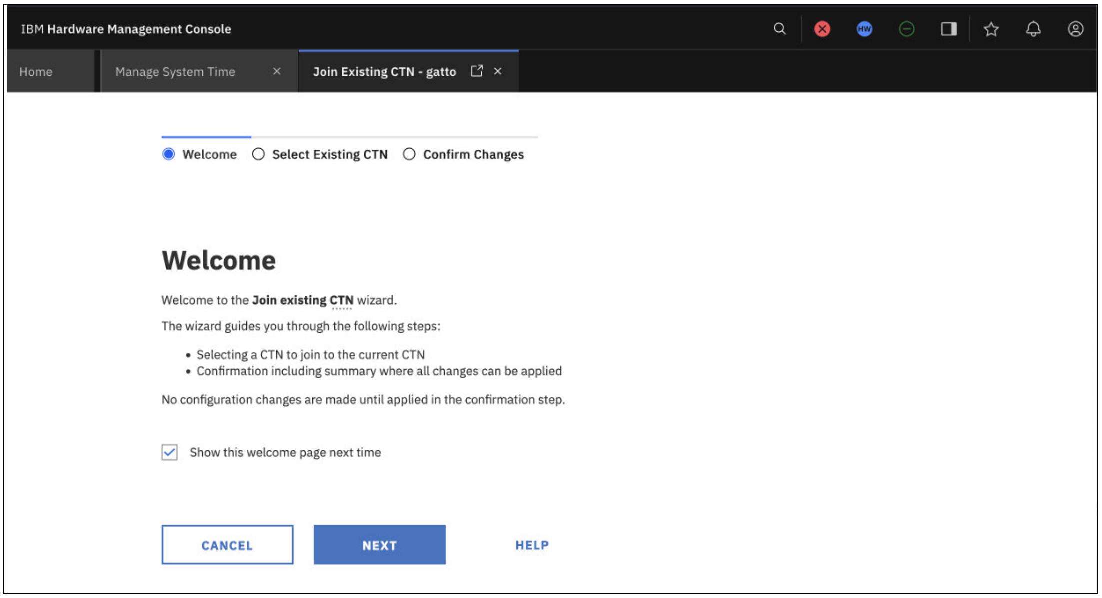

Figure 5-2 Join Existing CTN menu

3. Select the target CTN, as shown in Figure 5-3. Observe the GUIDANCE messages before continuing the process.
4. Click **NEXT**.

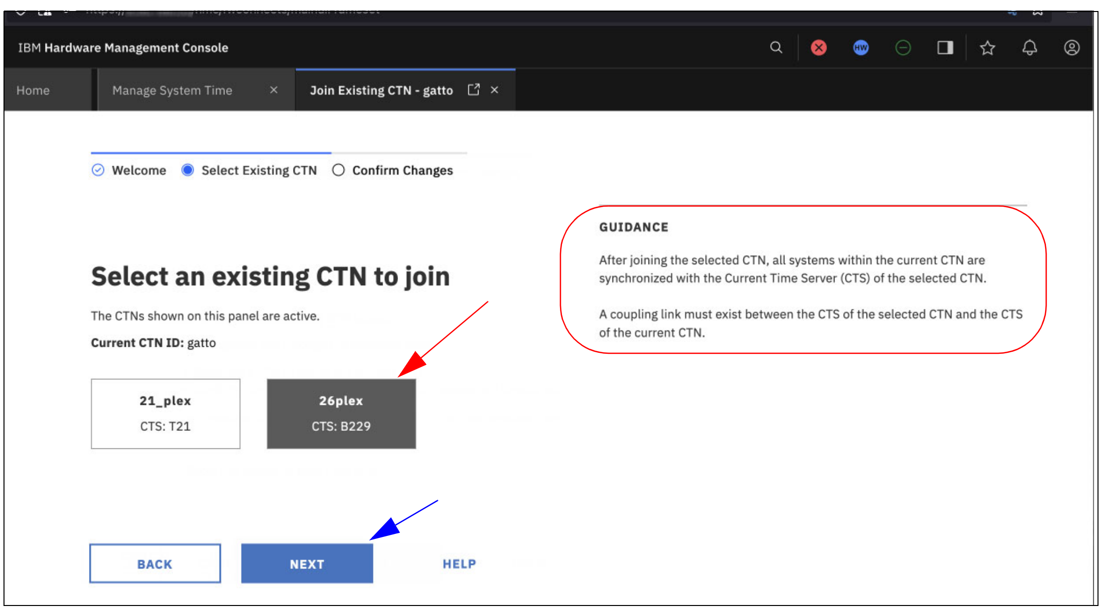

Figure 5-3 target CTN ID selected for servers in "gatto" CTN

5. A visual preview of the target configuration is displayed in the Confirm Changes menu (Figure 5-4 on page 114).

   Verify the changes, then scroll down and click **APPLY**.

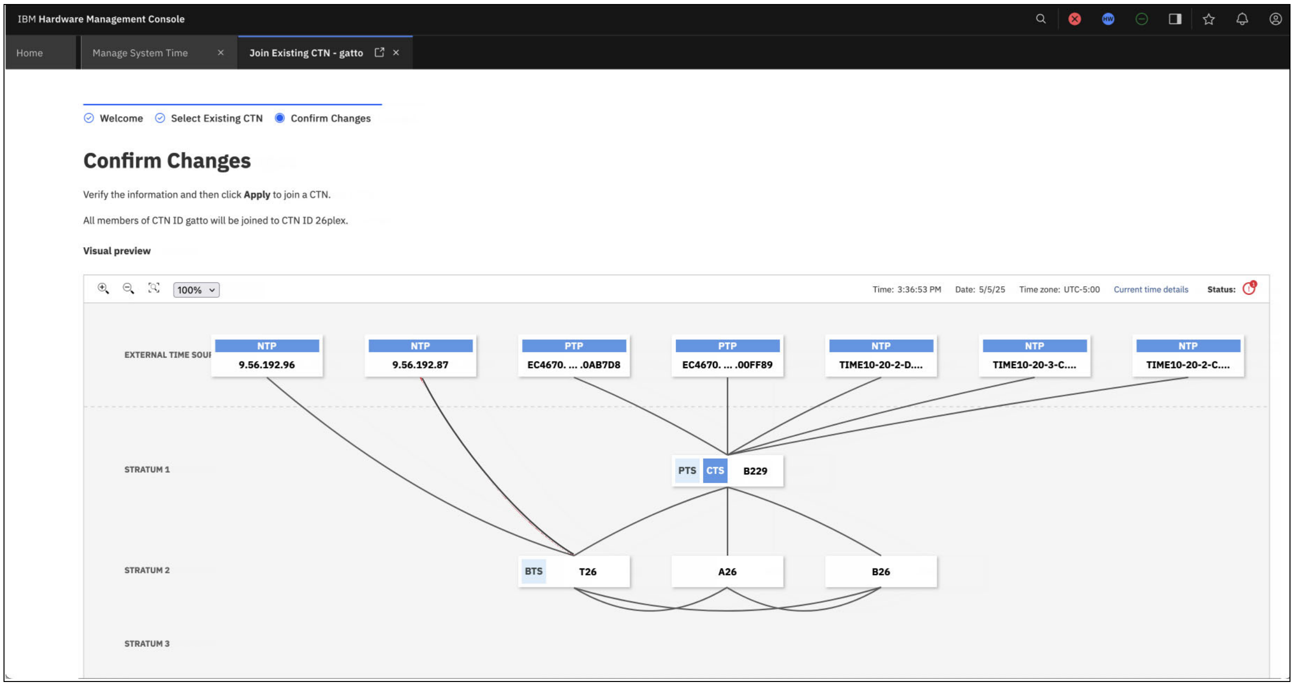

Figure 5-4 Confirm Changes menu

6. The process starts. When the process finishes, verify that it was successful (Figure 5-5).

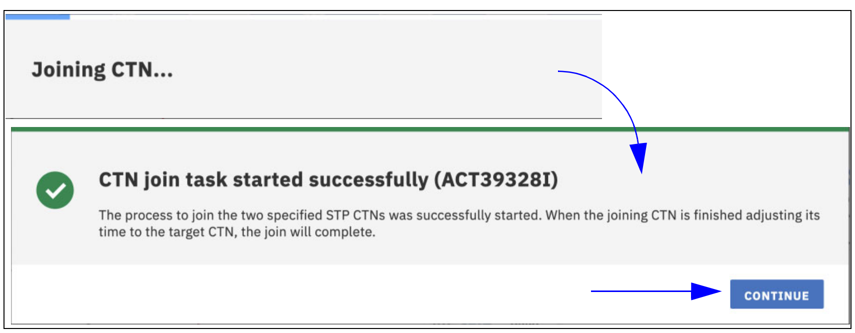

Figure 5-5 Joining CTN progress and successful starting messages

The result is displayed in the Manage System Time task as shown in Figure 5-6 on page 115.

Note: The servers from the joining CTN do not have a role assigned at the time of joining. You can assign roles through the Modify Assigned Server Roles menu.

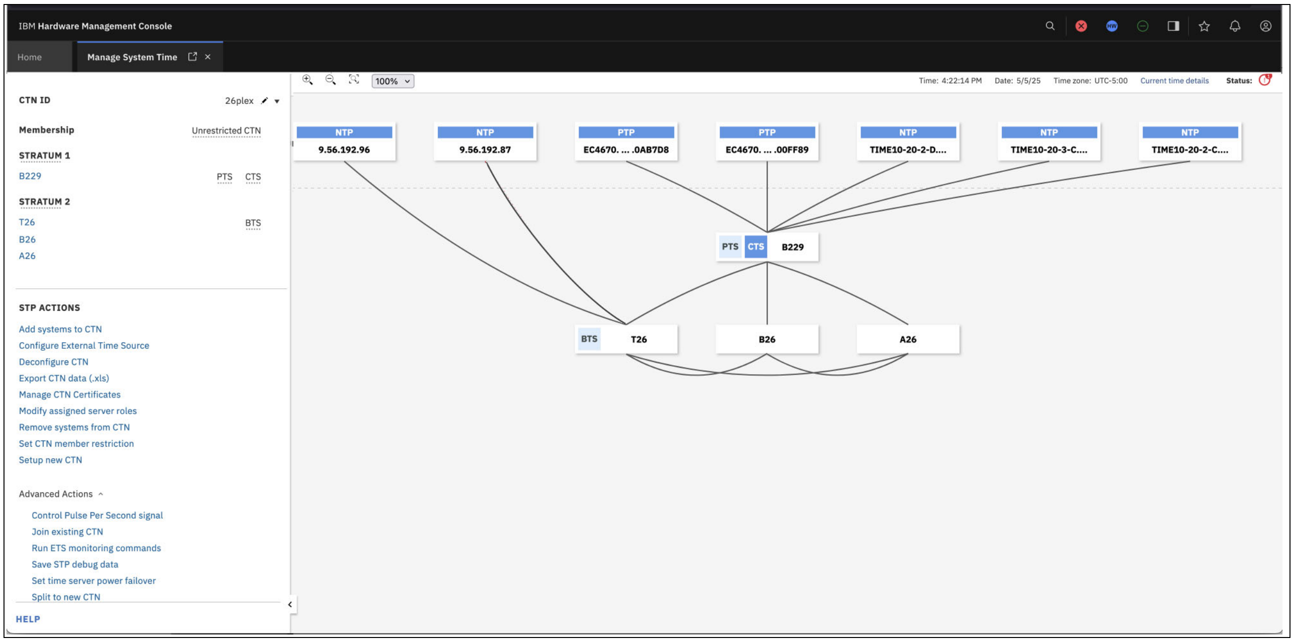

Figure 5-6 Final configuration

5.1.2 Split to new CTN

This task is used to split one CTN into two distinct CTNs. The starting point of this scenario is the configuration depicted in Figure 5-6.

> **Note:** Before attempting this task, verify that workloads can be divided and distributed across the two resulting CTNs.

To split the CTN, use the following steps:

1. To initiate the task, click **Split to New CTN** menu in the Manage System Time task.

   The system posts the message that is shown in Figure 5-7, and asks you to verify that you want to split the CTN.

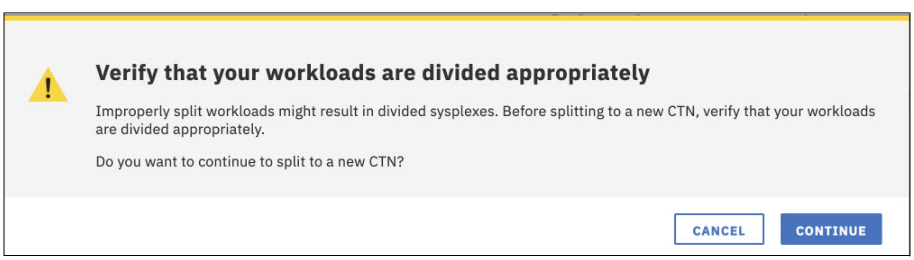

Figure 5-7 CTN split warning for divided workloads

2. Click **CONTINUE** and the Welcome menu is displayed (Figure 5-8 on page 116).
3. Click **NEXT** to enter a new CTN ID.

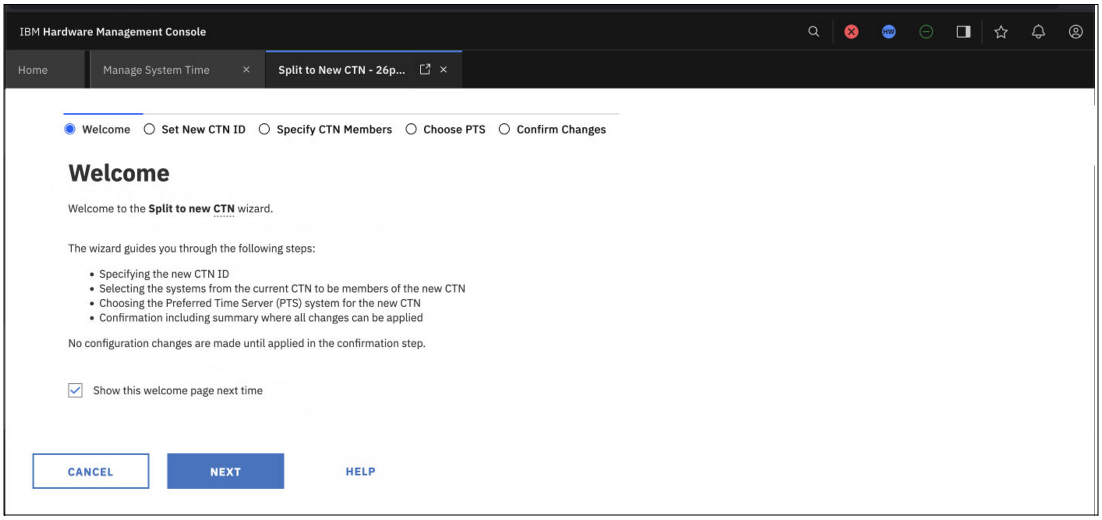

Figure 5-8 Split CTN Welcome menu

4. Enter the new CTN ID, as shown in Figure 5-9.

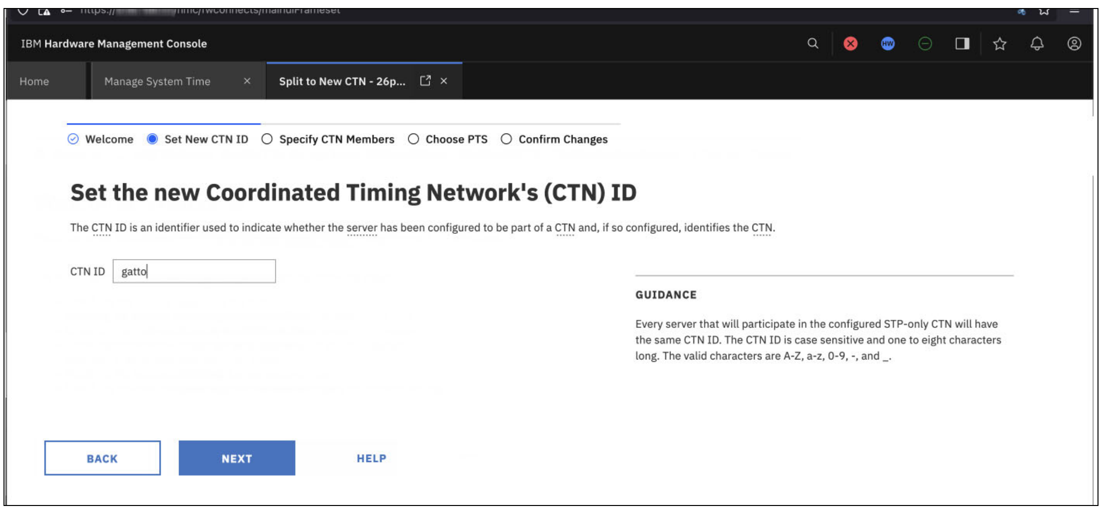

Figure 5-9 Set the new CTN ID

5. Select the servers that you want to become members in the new CTN ID (“gatto”), as shown in Figure 5-10 on page 117.

   Only the servers not having an assigned role are eligible for a new CTN.

6. After you select the servers, click **NEXT**.

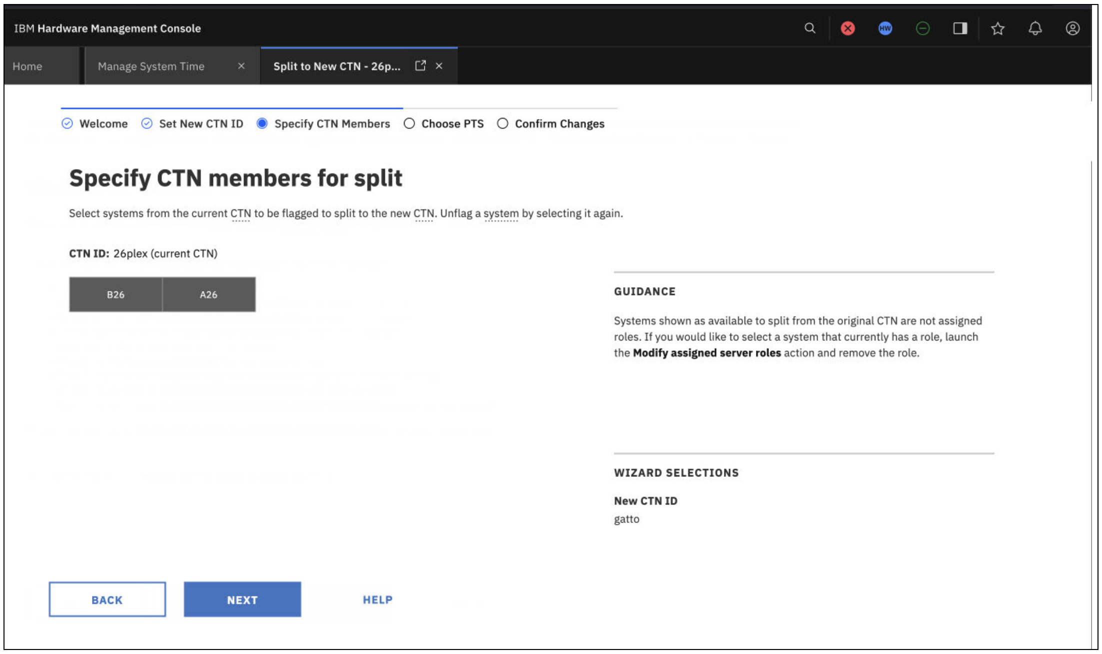

Figure 5-10   Servers for new CTN selected

7. Select the Preferred Time Server for the new CTN, as shown in Figure 5-11. Other assigned roles are optional and can be changed later as needed.
8. Click **NEXT**.

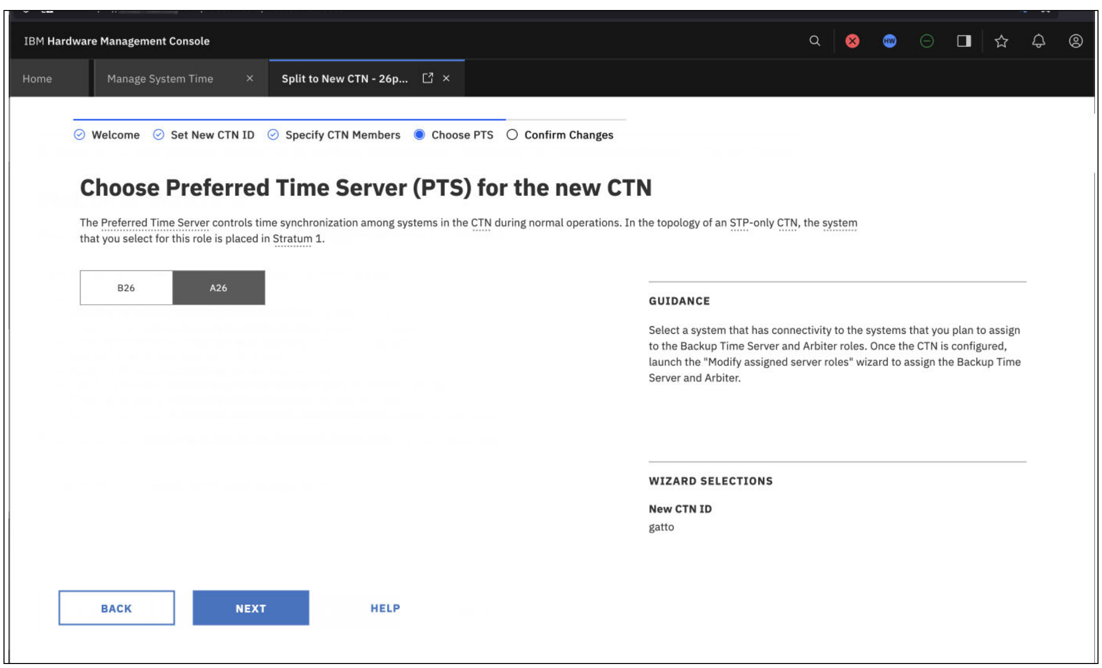

Figure 5-11   Assigning the PTS role

Configuration (diagram) of the new CTN is shown in Figure 5-12 on page 118.

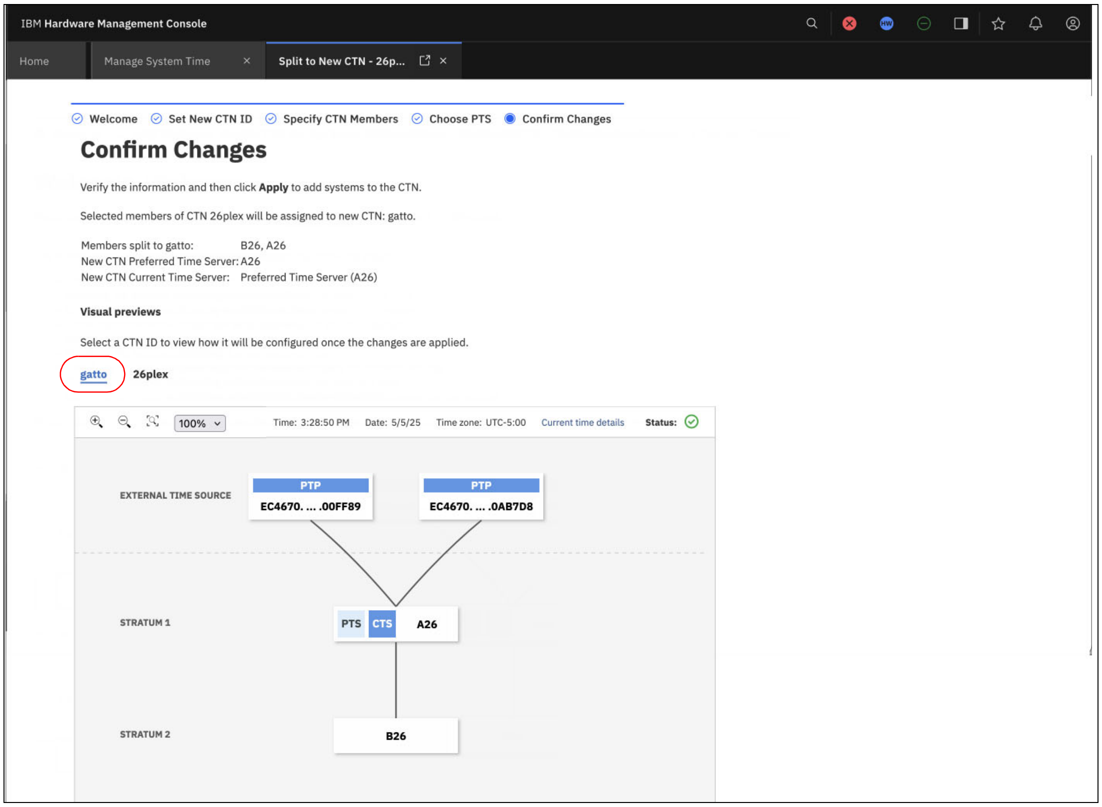

Figure 5-12 New CTN configuration

The configuration of the existing CTN is shown in Figure 5-13.

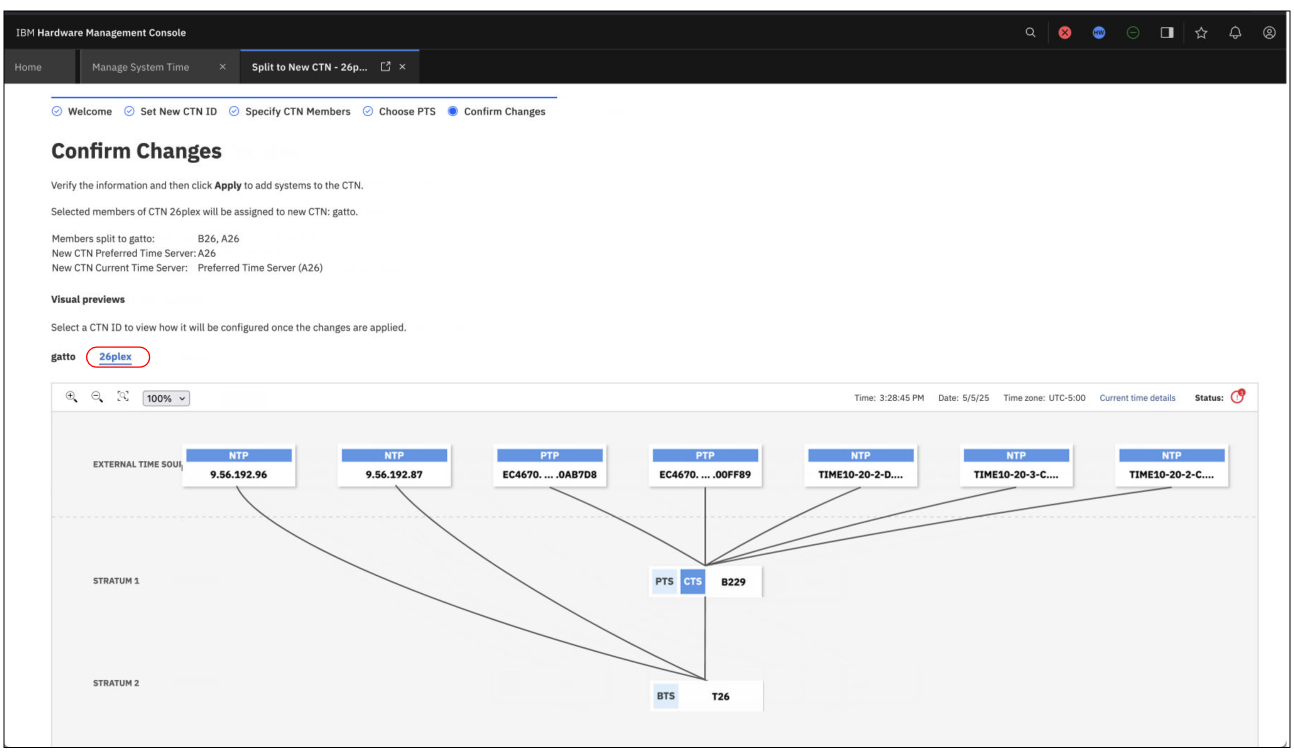

Figure 5-13 Existing CTN configuration after split

9. Review the changes, then scroll down and click **APPLY**. The progress messages followed by the Split CTN task success message are displayed (see Figure 5-14)

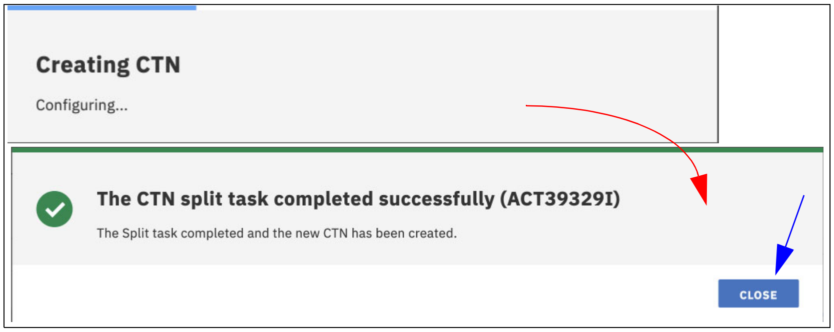

Figure 5-14 Task completed successfully

10. After the CTN split is successful, you can view the new CTN in the Manage System Time task window. See Figure 5-12 on page 118.

5.1.3 Save STP debug data

Use this task to send necessary STP information to IBM after finding a problem. Run this task with the following steps:

1. From the main menu, if necessary, expand Advanced Actions and select **Save STP Debug Data**.
2. From the Save STP debug data panel, select all systems with a target system or a single system (Figure 5-15).

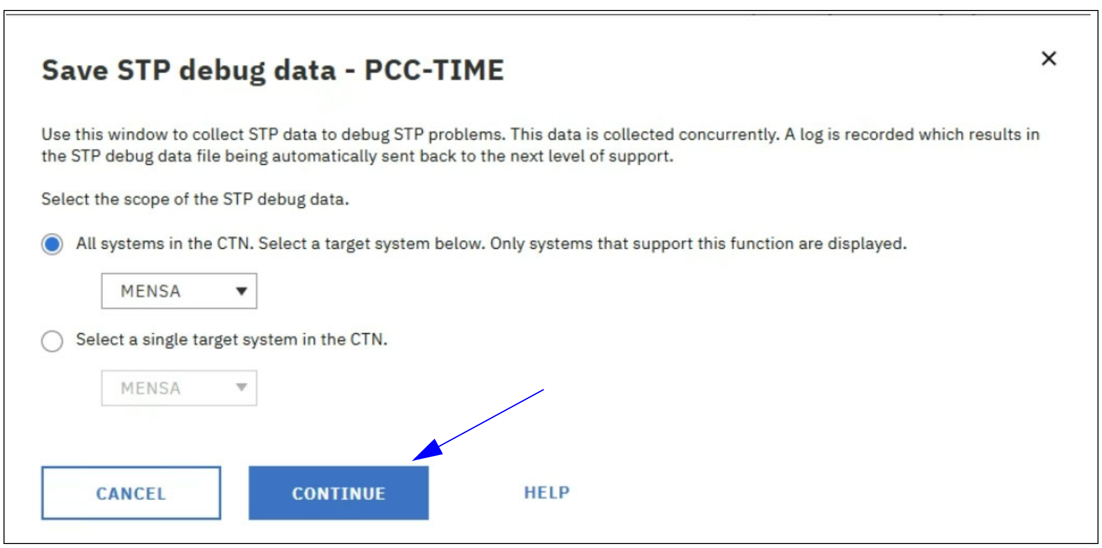

Figure 5-15 Start Save STP debug data menu

3. Click **CONTINUE** and check the progress and finalizing of the task, as shown in Figure 5-16 on page 120.

   When the data is collected, it is transmitted directly to IBM support.

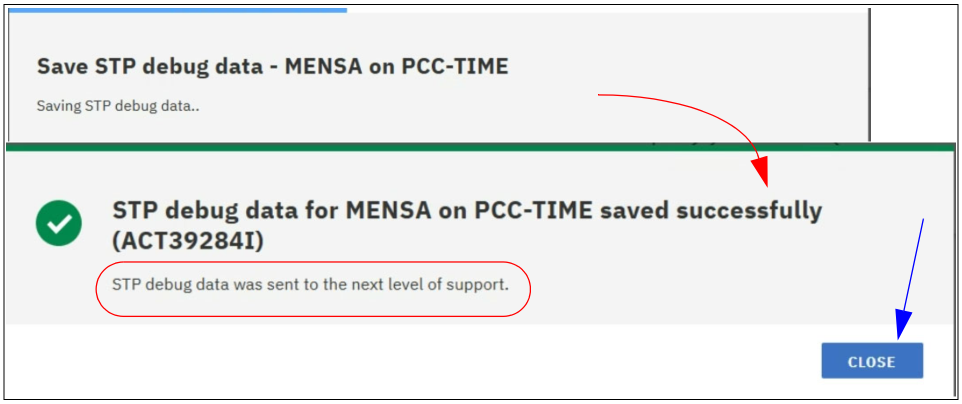

Figure 5-16 STP debug data sent to IBM support

5.1.4 Set time server power failover

Use this task to configure assigned server role failover if power is lost at the CTS site.

The Current Time Server (CTS - functional role) provides time information to the CTN members. In a dual- or triple- server CTN, the CTS role can be active on either the Preferred Time Server (default configuration) or the Backup Time Server (assigned CTN roles).

For high availability, another server can take the role of CTS if the current CTS server fails. For planning information about CTS failover methods, see “CTN recovery” on page 15.

To enable time server power failover, in the Manage System Time task, use the following steps:

1. Under the Advanced Actions section, select **Set time server power failover**, as shown in Figure 5-17 on page 121.

> **Note:** The **Set time server power failover** facility is available for IBM z16 and later systems. When you plan the CTN server roles, ensure that both the systems with the assigned roles (PTS, BTS) support this feature. For additional information, see “Power failure handling (n-mode power)” on page 15.

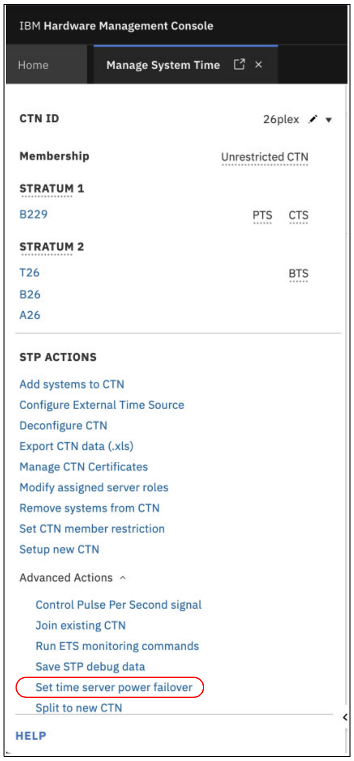

Figure 5-17 Set time server power failover menu

2. When you see the confirmation message that is shown in Figure 5-18, select the checkbox to enable the time server power failover option.
3. Set the timer, which is the interval after which the CTS fails over to the BTS.
4. Click **APPLY**.

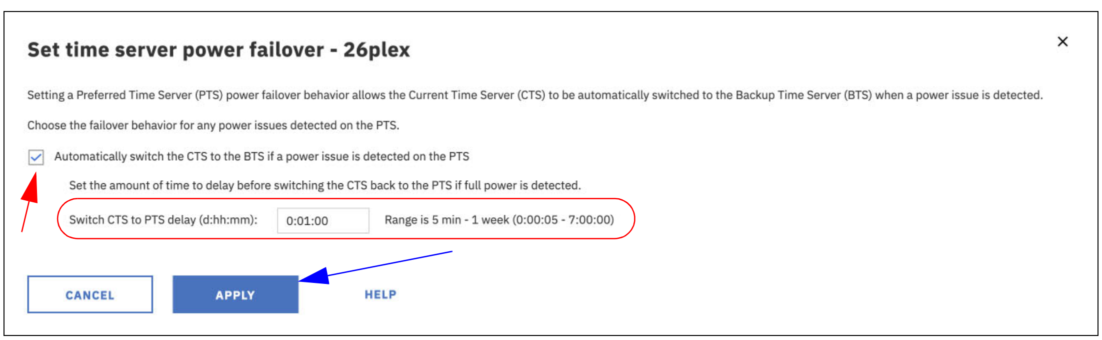

Figure 5-18 Set time server power failover menu

If one of the assigned server roles is not at a supported server generation (IBM z16 or newer servers), a warning message is displayed (Figure 5-19 on page 122).

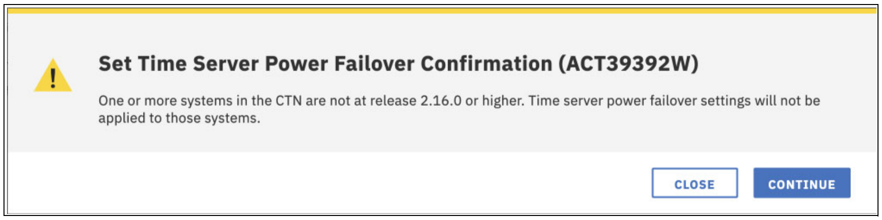

Figure 5-19 Time server power fallover warning message

5. You can click **CONTINUE** to finalize the operation. However, if the servers are not at a supported level, the CTS does not fail over to the BTS, as shown in Figure 5-20.

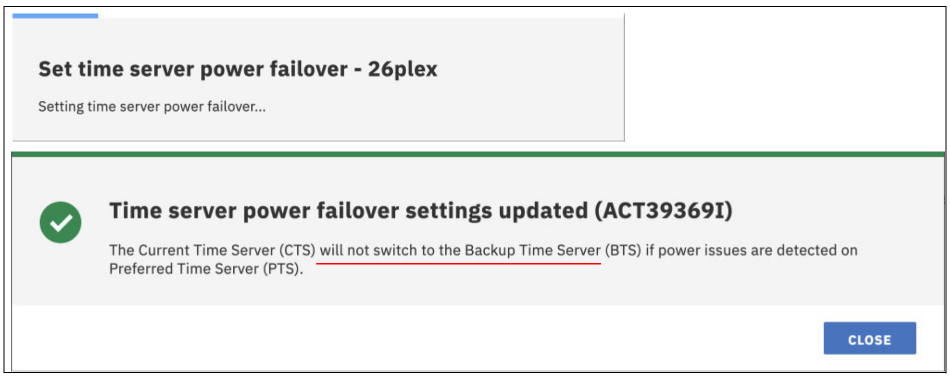

Figure 5-20 Message if servers with roles assigned are not IBM z16 or newer

However, if all servers are at the correct level, the feature is enabled, as shown in Figure 5-21

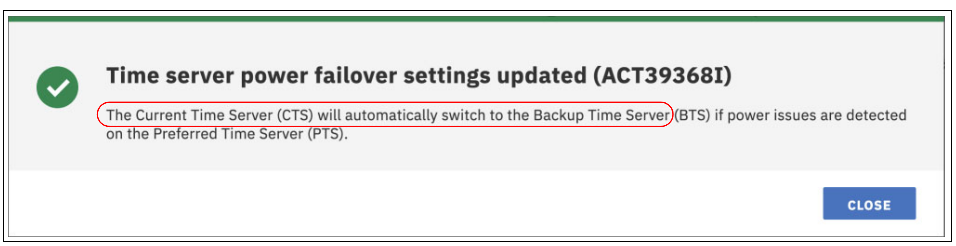

Figure 5-21 Time server power failover successfully enabled

Further CTN role changes, such as the Modify server roles task can impact the Time server power failover if one of the assigned roles is not at a supported level. Therefore, when you modify server roles, you see the message shown in Figure 5-22.

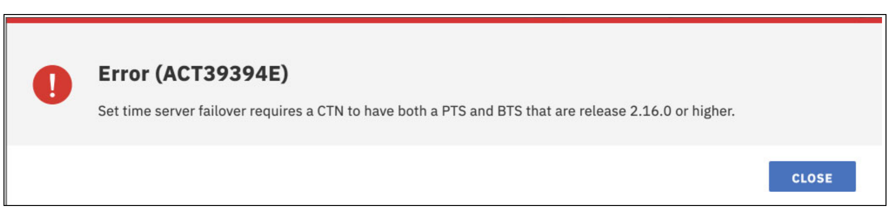

Figure 5-22 Error message: PTS and BTS not at 2.16.0 level (or higher)

5.2 Managing the Coordinated Server Time (CST)

This section describes the following time modifications within a CTN and the operational requirements:

- STP time adjustment
- STP offset adjustments
- Changes in local time
- Logical partition (LPAR) time offset
- Parallel Sysplex and multiple time zones

> **Important:** To set the date and time after a CTN is configured, you must deconfigure the CTN and then create the CTN and initialize the time again. For more information about how to deconfigure the CTN, see 3.7, “Deconfiguring the CTN” on page 64.

5.2.1 STP time adjustment

Without regular adjustment, the time on the CTS drifts from its initial setting because of the frequency drift of the oscillator that is used to step the TOD clock. This situation might not meet time accuracy requirements. Therefore, adjustments might need to be made regularly, either manually or by referencing an ETS.

To do a manual adjustment of time, use the following steps:

1. In the Manage System Time tab, access the Current time details panel (over the graphical topology view) and click **Adjust time** (Figure 5-23).

   The **Adjust Time** button is not available when the CTN time source is NTP with PPS. Because the STP uses PPS input every second, manual adjustment is disabled.

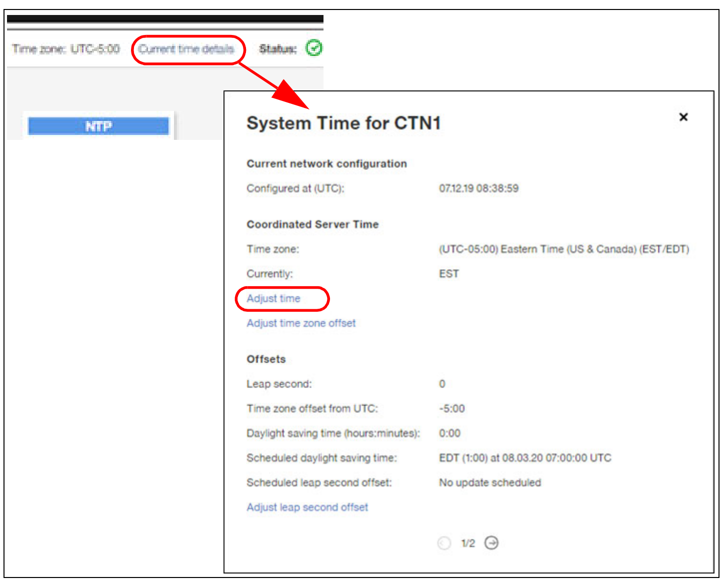

Figure 5-23 Adjusting the time

2. In the Adjust Time window, if you click **Access ETS**, the deviation of the CTN time from the ETS time is calculated and displayed.
3. Click **APPLY** to adjust the CTN time to the ETS time (Figure 5-24 on page 124).

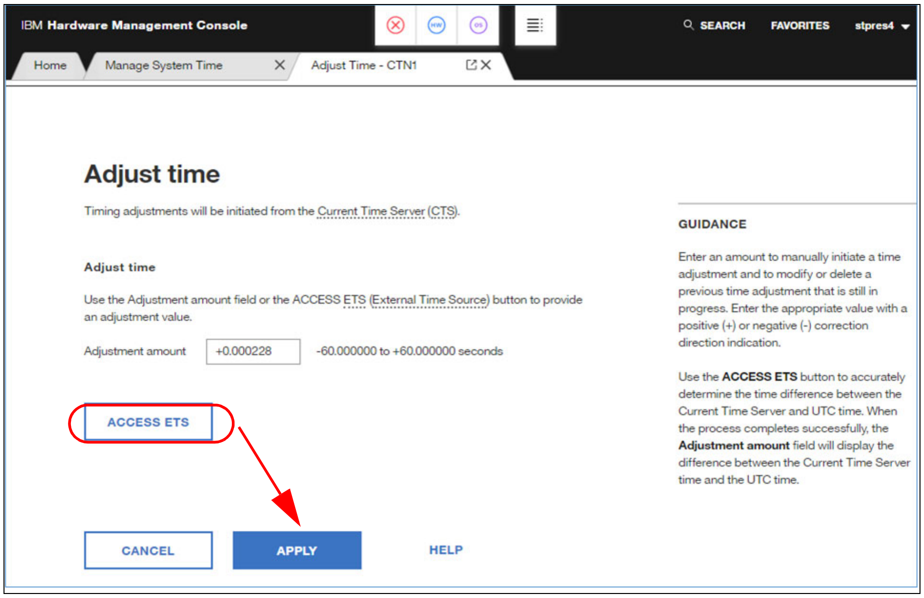

Figure 5-24 Applying the time adjustment

The following adjustments are possible:

- Adjustment steering
- Base steering

Adjustment steering

STP supports adjustment steering, which allows the time at the CTS to be changed by up to ±60 seconds. Adjustments greater than 60 seconds can be implemented in multiple increments of ±60 seconds. This change can take considerable elapsed time to achieve.

The offset that is specified is gradually incorporated into the standard timing messages in small enough increments or decrements that the operating systems (OSs), 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.

> **Important:** In a CTN, the adjustment steering rate is approximately 1 second for every 7 hours.

Base steering

There is another time-steering method that is built into the STP facility that works in a similar way to adjustment steering. This method is an automatic function that does not require user intervention. This method is known as base steering, which is performed at the CTS and requires an ETS.

> **Note:** Base steering occurs automatically when you use an ETS configuration with NTP Stratums 1, 2, or 3, or better PTP sources.

The stratum level of the NTP server can be seen in the Manage System Time window by clicking the NTP server (Figure 5-25 on page 125).

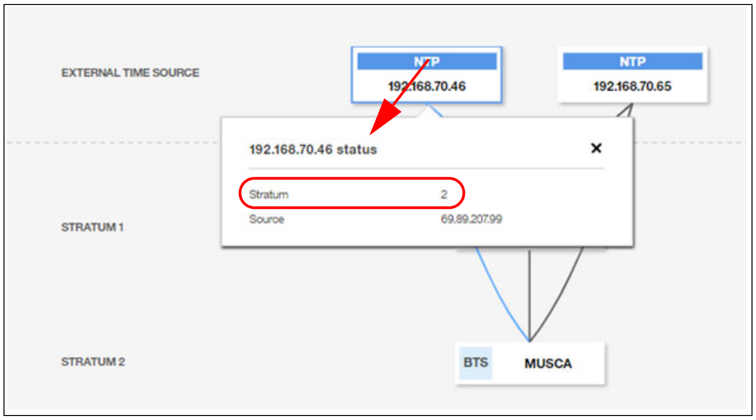

Figure 5-25 NTP stratum information

By comparing the time that is obtained from subsequent accesses of a configured ETS with the corresponding Coordinated Time Server (CST) values, STP can compute the amount of drift that occurred between the events. This amount represents the inherent inaccuracy of the CTS oscillator over time. With this information, STP can automatically introduce a compensation offset into the STP timing messages by adding extra steering to counter the drift. As a result, the servers in the CTN self-correct over time so that the offset that is returned from future ETS accesses approaches zero as greater accuracy is achieved.

Table 5-1 provides a summary of STP clock adjustment functions.

Table 5-1 STP time steering

Adjustment method	NTP server access adjustment	Manual adjustment	Manual set
Adjustment steering	Yes	Yes	No
Base steering	Yes, after multiple ETS accesses	No	No

Important: To set the date and time after a CTN is configured, you must deconfigure the CTN and then create the CTN and initialize the time again. For more information, see 5.1, “Advanced CTN Actions” on page 112.

5.2.2 STP offset adjustments

The STP timing message includes the following items:

- CST
- Leap second offset
- Time zone offset
- DST offset

The values in this list are transmitted from the CTS to all the servers in the CTN. How the z/OS system image uses these values depends on options that are specified in the TIME macro with options that are specified in CLOCKxx at IPL.

Different time results can be received depending on the options that are specified in the TIME macro, as shown in Table 5-2 on page 126.

Table 5-2 TIME macro options

Option	TIME macro with ZONE=LT	TIME macro with ZONE=UTC	TIME macro with STORE CLOCK
Include TOD in the result.	Yes	Yes	Yes
Include leap second offset.	Yes	Yes	No
Include a time zone offset.	Yes	No	No

In addition, the parameters that are specified in the CLOCKxx member at IPL determine where these values are obtained.

Table 5-3 shows comparable details for z/OS systems on a server in STP timing mode.

Table 5-3 z/OS system on a server in STP timing mode

Option	STPMODE=NO STPZONE=NO	STPMODE=YES STPZONE=NO	STPMODE=YES STPZONE=YES
Step TOD to CTS.	No[a]	Yes	Yes
Include a time zone offset from the CTS.	No	No	Yes
Include time zone offset from CLOCKxx.	Yes	Yes	No
Allow local time adjustment through z/OS SET commands.	Yes	Yes	No

a. The base TOD steering affects all LPARs whether they operate in STP mode or not because it changes the hardware TOD gradually without notifying the OSs. The STP mode corrects the difference when the local hardware TOD drifts and differs from Stratum 1 (the CTS or CST), which results in an STP sync check. The STP sync check is presented only to LPARs running in STP mode. Several other STP machine checks and external interruptions are also presented only to those LPARs running in STP mode.

Leap second offset adjustment

Adjustments are necessary (if leap seconds are used) only if the applications require precise synchronization accuracy to Coordinated Universal Time (UTC). Examples of such specific requirements might be legal or contractual requirements for timestamps to be within tolerance of UTC time or if timestamps are used for time-dependent banking, scientific, or navigational purposes.

Periodically, the International Earth Rotation and Reference Systems Service (IERS) in Paris advises to introduce a positive or negative leap second adjustment into civil time. For the most current leap seconds information, go to IERS and search for “leap seconds”.

Although most leap second adjustments are positive, it is possible to have a negative leap second adjustment.

Important: Leap seconds are automatically built into the UTC time that is obtained from an ETS. Any leap second offset that is defined is included when calculating the delta between the CTS and the time that is received from the ETS to prevent double accounting.

If the ETS is used to incorporate an extra leap second into the CTS TOD, then the requirement that leap second adjustments should occur simultaneously worldwide has been breached. This situation occurs because TOD adjustments through the ETS are implemented over a period through adjustment steering rather than immediately.

If leap seconds are used, use the adjust leap seconds facility 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 30 June or 31 December. However, the offset adjustment might be scheduled for a particular date, depending on your requirements.

A positive leap second is inserted after UTC 23:59:59, as UTC 23:59:60, and then UTC 00:00:00 occurs as normal. Note that 23:59:60 is never made visible to z/OS users of the TOD clock. Although at the time of writing, a negative leap second has not occurred, it is implemented by excluding UTC 23:59:59.

When the leap second offset occurs, z/OS is interrupted 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.

To apply the leap second offsets, use the Adjust Leap Second window (see Figure 5-26). To apply the leap second offset, click **Adjust Leap Second offset** from the Current time details window (over the graphical topology view) in the Manage System Time tab.

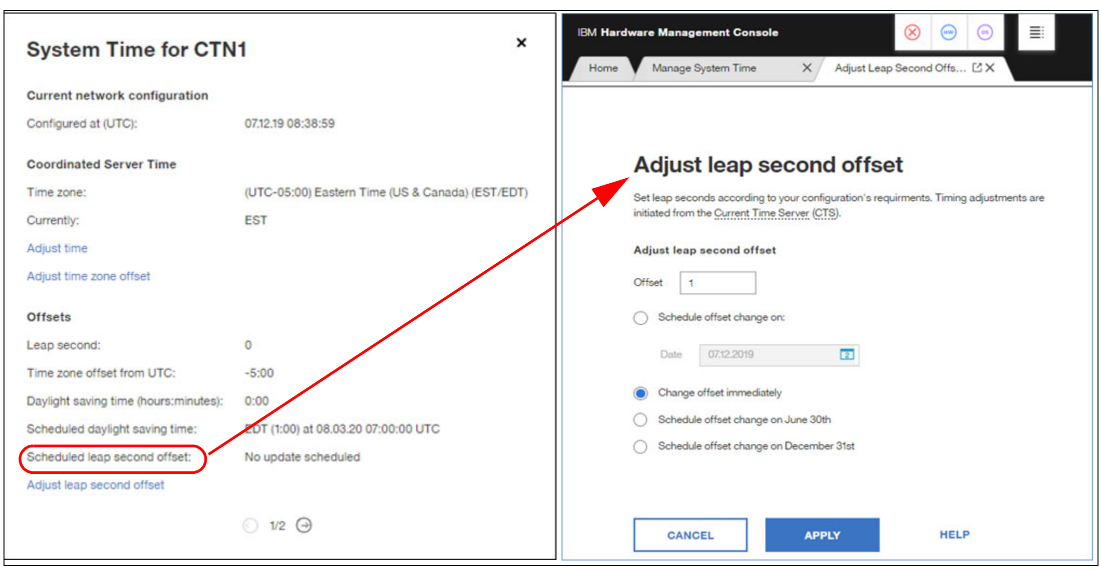

Figure 5-26 Leap seconds adjustment

Attention: The value of "1" for the leap seconds is only for display purposes of this scenario. Use a correct value for the leap second offset. At the time of writing, the value for leap seconds is 27.

When the Adjust Leap Second Offset window opens, the current leap second offset in effect for the STP 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.

If you want to add a leap second (negative or positive), you must add or subtract this leap second from the value that is displayed in the window.

Although z/OS spins for positive leap seconds, there is no need for a spin during a negative leap second.

A negative leap second appears to any application that 1 second jumped past without them seeing it. A large change in leap seconds can cause a system outage.

You can also set the time to implement the leap second adjustment in the window.

Important: Leap second adjustments occur simultaneously worldwide (regardless of the time zone) when UTC 23:59:59 occurs. STP and the OS take care of the adjustment. No user intervention or interaction is needed.

Time zone offset adjustment

In the Current time details window (over the graphical topology view) in the Manage System Time tab, you can see the set time zone offset.

You can change the time zone by using the **Adjust Time Zone Offset** tab. Select one of the supported time zones that are provided by default or use a user-defined time zone to meet the requirements (see Figure 5-27 on page 129).

Supported time zone offsets

Several supported time zone entries are provided by default. Each of these entries has a defined offset from UTC.

Note: The time zone on the SE is independent of the time zone that set for the CTN. The same is also true for the HMC. It is possible to have a different time zone at each SE, the HMC, and for the CTN.

In addition, each entry can optionally have a time zone algorithm that is 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** button can be used to activate automatic adjustment for the site.

Alternatively, the time zone might not support automatic adjustment, or it is handled manually. In this case, select either **Set standard time** or **Set Daylight Saving Time** to indicate whether daylight saving is active when the selected time zone is activated.

The Scheduled Clock Adjustment for daylight saving section is displayed if an adjustment is scheduled either automatically as the result of an Adjust Time Zone offset algorithm or manually through the schedule change on facility.

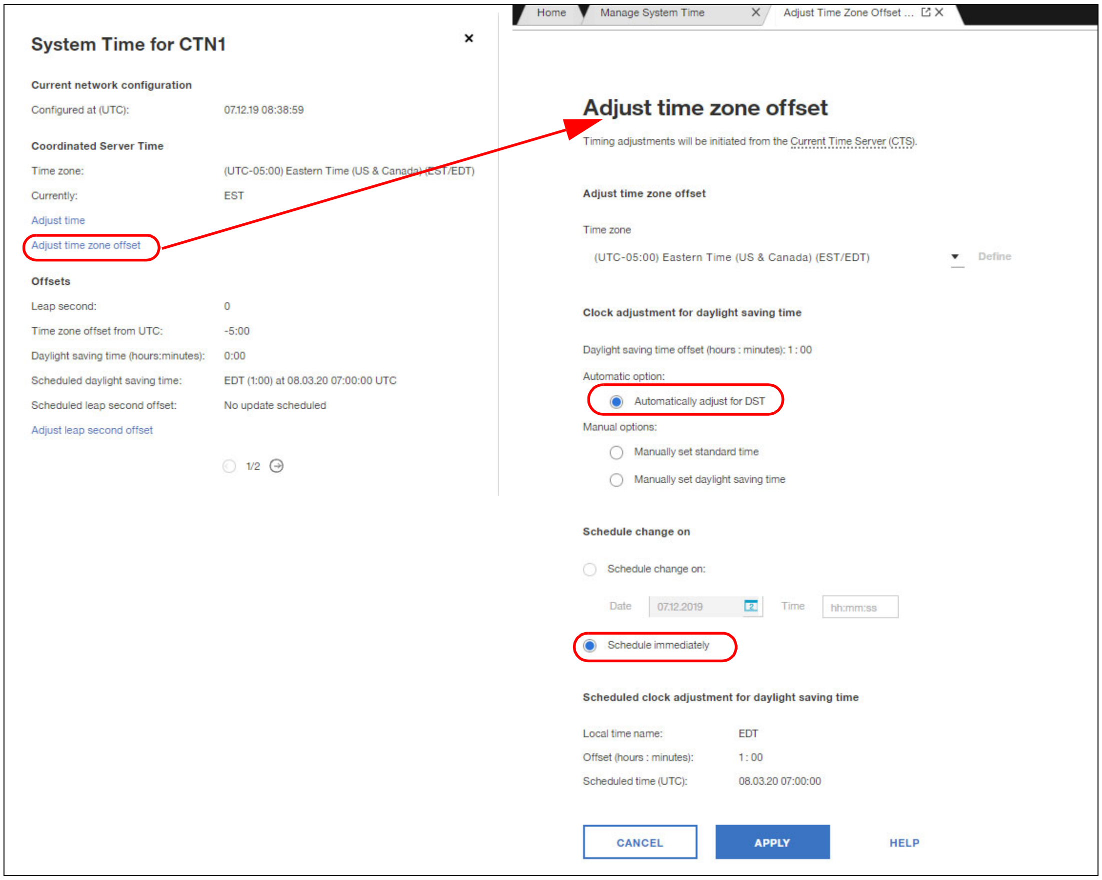

Figure 5-27 Time zone adjustment

Reviewing a time zone algorithm

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 offset entry for the time zone meets the requirements. However, you can check this situation indirectly by applying the time zone offset entry for the time zone with *automatically adjust*, and then review fields in the Schedule Clock Adjustment for daylight saving section for correctness in relation to the next time zone adjustment.

Daylight saving changes that are implemented by a government authority

An example of daylight saving changes that are implemented by a government authority is the US Energy Policy Act of 2005 that changed DST in the US in 2007.

For an STP CTN, which has the option of automatically scheduling DST, support for the new DST start and end dates is available if the server has the correct level of microcode.

Future DST changes that are implemented by any governmental authority might require MCL updates. Update the server MCLs as needed.

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, UD1–UD5, can be used to define the time zone. If a user-defined time zone entry is selected, the **Define** button is enabled (Figure 5-28 on page 130).

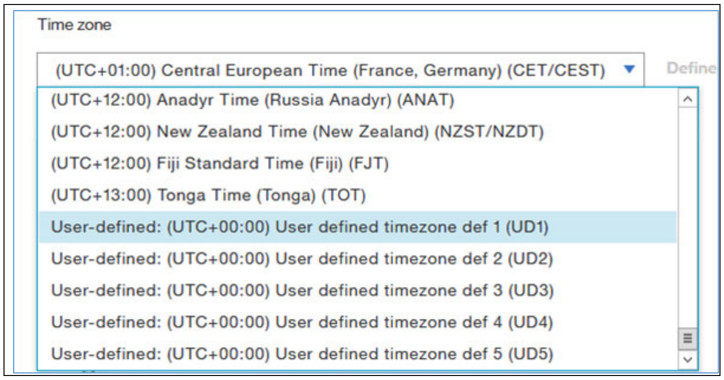

Figure 5-28 User-defined time zone entries

The Description (maximum of 80 characters) and Standard time name fields (maximum of 4 characters) must be entered, or an error message appears when you click **OK**. The standard time name is an abbreviation that is displayed on various windows to differentiate standard time from daylight saving.

The UTC offset must be entered in ±hours and minutes, and the range is -14 to +14 hours.

Also, if the time zone is subject to daylight saving adjustments, then the daylight saving name and daylight saving offset must be specified. Optionally, algorithms for DST start and DST end can be defined to support automatic clock adjustment by selecting the “Define adjustment of clock for daylight saving” option. The algorithm is saved when you click **OK**, but it is not sent to STP until you click **4$ZSeeeeeeeee** on the Adjust Time Zone Offset window.

5.2.3 Changes in local time

You can use z/OS to obtain STORE CLOCK (STCK) time, UTC time, or local time, depending on your requirements. The difference between UTC time and local time is usually the time zone offset under normal circumstances. You can manage the time zone offset at the z/OS level by specifying the STPZONE=NO option in the CLOCKxx member.

In z/OS, you can set ETRMODE in the CLOCKxx member, which requires a connection to a Sysplex Timer. In terms of hardware, connecting the mainframe server to a Sysplex Timer is no longer possible because IBM z196 does not support it. Therefore, do *not* set ETRMODE. If the server runs in STP timing mode and both STPMODE and ETRMODE are specified as YES, z/OS uses STPMODE YES.

The TIMEZONE parameter in the CLOCKxx member sets the time zone offset at IPL, and you can use several z/OS SET commands to dynamically adjust the offset when required. Similarly, the CF supports the concept of time zone offset and allows dynamic modification of the time zone offset through a command.

z/OS commands

On a z/OS system, the local date and time can be modified dynamically. The ability to do this task depends on what options are specified in the CLOCKxx member at IPL, as shown in Table 5-4 on page 131.

UTC time cannot be changed by a z/OS command. It can be modified during the IPL processing response to the following message:

```
IEA888A   [UTC DATE=yyyy.ddd,CLOCK=hh.mm.ss] LOCAL DATE=yyyy.ddd,CLOCK=hh.mm.ss
```

REPLY U, or **UTC/LOCAL TIME** is issued at IPL when **OPERATOR PROMPT** is specified in CLOCKxx with STPMODE NO and SIMETRID not specified.

Any attempt to change the local time or date when the server is operating in STP mode generates the following message:

```
EA279I: IEA279I ALL CLOCK RELATED SET COMMANDS ARE IGNORED WHEN IN STP MODE.
```

Table 5-4 z/OS time adjustment through command cross-reference

	Adjust time through z/OS command		
	Local time ZONE=LT	**UTC time ZONE=UTC**	**STCK time STCK**
`STPMODE=NO, STPZONE=NO`	Yes	Yes[a]	Yes[a]
`STPMODE=YES, STPZONE=NO`	Yes	No	No
`STPMODE=YES, STPZONE=YES`	No[b]	No	No

a. UTC time cannot be changed by a z/OS command.
b. Attempts to change the local time or date in STP mode generates message `IEA279I`.

DISPLAY TIME

This command can be used to display the local time and date, and the UTC time and date, as shown in Example 5-1.

Example 5-1 The 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 the DST offset, if any) that is applicable to the time zone.

SET DATE

This following command is used to change the local date:

```
SET DATE=yyyy.ddd
```

The command has the following restrictions:

- *yyyy* is the year. It must be in the range 1900– 2042. The value that is specified must consist of four digits and must be within 70 years of the UTC date, or 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 following command is used to change the local time:

```
SET CLOCK=hh.mm.ss
```

This command is used with the SET DATE command to set a maximum value of 23.53.47 on 2042.260. This command does not update the server's TOD clock or the logical TOD of the LPAR on which this z/OS image operates. The change that is made by this command is effective until the next IPL.

Also, z/OS does not change the date when the new time implies a change of date, so either use the **SET DATE** command or enter the command after midnight if the new time is for tomorrow.

SET RESET

This command resets the time zone offset to the value that was read from the CLOCKxx member during IPL, therefore changing the local date and time with the following syntax:

```
SET RESET
```

This command annuls all previous SET DATE, SET CLOCK, and SET TIMEZONE commands, and reestablishes the relationship between the local date and time and the UTC date and time plus time zone offset.

SET TIMEZONE

This command changes the time zone offset to a different value from what was specified at IPL through the TIMEZONE parameter in CLOCKxx. This command automatically adjusts the local date and time with the following syntax:

```
SET TIMEZONE={W|E}.hh[.mm]
```

The time zone offset direction is West (W) or East (E). West is the default. The value for *hh* must be 00–15, and the value for *mm* must be 00–59.

You can manually change the DST by using the SET CLOCK command instead of STP changing it automatically. When you use the manual method, there is always a degree of error because the difference between the local time and UTC time does not exactly match the time zone offset that would have been achieved by updating the TIMEZONE statement in CLOCKxx and then running an IPL.

The z/OS SET TIMEZONE command overcomes this problem by applying the correct offset value to the CVTTZ field, which causes an exact time zone offset to be applied.

Tip: If the SET CLOCK command is used to change local time for DST offset purposes, then use the SET TIMEZONE command instead for better accuracy.

If the time zone offset for DST is changed dynamically by 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 does not regress at the next IPL.

Coupling Facility commands

In a Parallel Sysplex environment, Coupling Facility commands (CFs) require time awareness to support CF request time ordering. The server TOD is used for this purpose.

CFs also support the concept of time zone offset, which is used only for timestamping messages that are displayed on the console.

There is no Coupling Facility Control Code (CFCC) command to display the time. However, all messages that appear on the CF console include a timestamp 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 timestamp 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 with the following syntax to display the current time zone offset that is used by the CF:

```
DISPLAY TIMEZone
```

This command produces a single line indicating how many hours and minutes the current time zone is east or west of Greenwich Mean Time (GMT) (Example 5-2).

Example 5-2 The 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 to change the time zone offset, if needed, and has the following syntax:

```
TIMEZone {0|hh|hh:mm|:mm} {East|West}
```

Use the command that is shown in Example 5-3 to adjust the local time as shown in messages on the CF console for the onset and removal of DST.

Example 5-3 The CFCC command to adjust the time zone

```
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 with Daylight Saving Time changes

When a CF image partition is activated and the server is using an STP source, the CFCC uses only one of the following time offset options:

- The LPAR time offset that is specified in the image profile. The TIMEZ offset overrides the LPAR time offset.
- Use the TIMEZ command for DST changes, as described in Coupling Facility - TIMEZ Command during Daylight Saving Time Changes.

5.2.4 Logical partition time offset

IBM Z servers support a function called *LPAR time offset*. LPAR time offset supports the optional specification of a fixed time offset (specified in days, hours, and quarter hours) for each LPAR activation profile. If you specify an offset, z/OS applies it to the time that an LPAR receives from STP.

Sometimes, you must run multiple Parallel Sysplexes with different local times and set the time to TOD Clock LOCAL. This situation causes the results that the STCK instruction returns to reflect local time. With LPAR time offset support, LPARs on each server in a Parallel Sysplex can specify an identical time offset that shifts time in the LPAR sysplex members to the local time. If you specify an LPAR time offset value, z/OS applies it to the time value that an LPAR receives from the time source. The time zone offset and DST offset are independent of this parameter.

The remaining LPARs on the servers can continue to participate in current date production Parallel Sysplexes that use the same STP messages with the time that STP provides. All supported releases of z/OS support this function.

To set the LPAR time offset, complete the following steps:

1. 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 5-29).

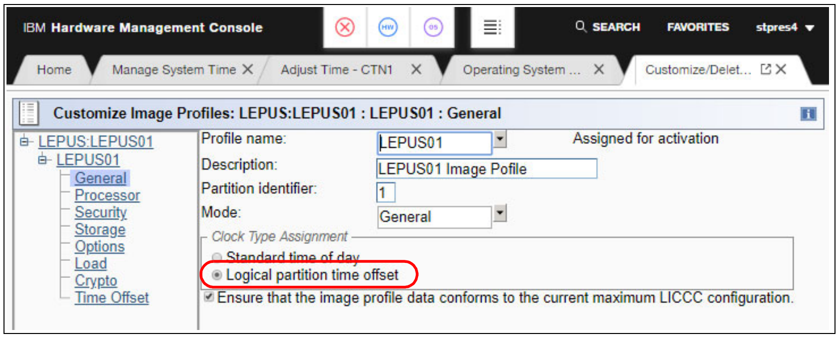

Figure 5-29 Customize Image Profiles: General

2. A new selection that is called **Time Offset** becomes available at the left. Use the Time Offset window to set the offset, as shown in Figure 5-30.

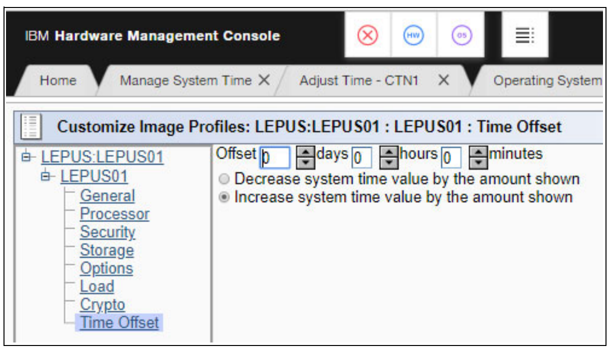

Figure 5-30 Customize Image Profiles: Time offset

The settings that are 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 that is supplied by its time source. The offset can be in the following range:

- 0–999 days
- 0–23 hours
- 0, 15, 30, or 45 minutes

– **Decrease system time value by the amount shown**

Select this choice to set the LPAR's clock back from the time of day that is 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 LPAR's clock ahead of the time of day that is 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.

5.2.5 Parallel Sysplex and multiple time zones

Some Parallel Sysplex installations have a requirement where different z/OS sysplex members undergo IPL in different time zones. You can meet this requirement by using the following means:

- Having all sysplex members on a common time source with STPMODE YES.
- Not using STP to obtain time zone information by specifying STPZONE NO and TIMEZONE *d.hh.mm.ss*.
- Alternatively, one group of sysplex members in a common time zone can use STP as the time zone source by using STPZONE YES while other sysplex members do not.

This section describes how to create and manage a multiple time zone environment.

Creating the multiple time zone environment

The example in this section reflects a consolidation of multiple sysplexes into a single sysplex. Although this action is not remarkable, what is unusual is that the sysplex supports applications that operate in images that use different time zones.

You must address the following issues when you create the environment:

- IBM software considerations
- Vendor software considerations
- Configuration considerations

IBM software considerations

Creating a sysplex with staggered time zone offsets introduces many complexities into the environment that might not be acceptable to everyone. However, business reasons such as sysplex consolidation or enterprise mergers might favor this configuration. The following examples illustrate some of these complexities:

- SMF considerations

  SMF records contain local time. If you use SMF records from two systems for accounting purposes, the data is inaccurate. For example, assume that a job is submitted on LPAR 1A, obtains a job reader start time in one time zone, and runs on another image with job start and end times in a different time zone.

- RMF considerations

  RMF displays and reports are unreliable for the same reasons that apply to SMF.

- OPERLOG considerations

  Although the records are ordered by UTC (GMT) timestamps, they display local time offsets. For example, in a merged OPERLOG, messages from the two systems appear to be out of sequence.

- JES2 time-sensitive command considerations

  Some JES2 commands, such as $TA and $PJQ, have time-sensitive parameters that cause actions to occur at different times. For example, the commands run N hours earlier on one system than on the other.

The system functions properly, but the external view complicates the role of some people in the enterprise.

Vendor software considerations

Because a sysplex allows work to be routed to any machine in the configuration, programs that request local time can receive different answers depending on the server where the work is sent. To help in this situation, some software provides an altered date and time to users, user applications, subsystem software, IBM CICS® regions, and even specific transactions or terminal IDs.

Configuration considerations

To set up the sysplex with different time zones, you must consider z/OS and middleware requirements. In-depth details are beyond the scope of this book, but the following examples illustrate some considerations:

- CICS complex (IBM CICSplex®)

  Although a full-function CICS originally ran in a single address space (region) in the z/OS environment, most CICS users now run multiple, interconnected CICS regions. By using the CICS Multi-region Operation (MRO) Intercommunication Facility, you can combine CICS regions into a complex of subsystems.

  With MRO, you can separate CICS functions into individual regions, with the different types of CICS regions classified as resource managers. With the latest enhancements to MRO, these CICS resource manager regions can run in one or more z/OS images.

- CICSPlex System Manager

  If servers are in different geographical locations, determine whether connections exist between those processors or whether they are managed as separate entities, each with its own workload. If these separate units exist in the enterprise, you likely must define multiple CICSplexes and manage the enterprise CICS systems as though they belong to more than one enterprise.

Much of the CICSPlex System Manager activity is time-dependent. For example, you can specify that a monitor definition or an analysis definition is active during a particular period. CICSPlex System Manager does not require every instance in a single CICSplex to run in the same time zone, so it must accommodate any time zone differences between entities:

- When you create a time period definition by using the CICSPlex System Manager PERIODEF view, you must specify a time zone in the definition. For example, you can create a time period definition called MORNING for the hours 08:00–11:59 Eastern Standard Time.
- You must specify a time zone for each CICSPlex System Manager address space (CMAS) in its data repository initialization job, which is called EYU9XDUT. You can make a permanent change to the CMAS time zone value, even while the CMAS is running, by using the CICSPlex System Manager user interface.
- You must establish a time zone for each managed CICS system. When you define a CICS system in a CICSPlex System Manager, you can specify the time zone in which the system runs. Alternatively, if you do not specify a time zone in the CICS system definition, the CICS system is assumed to run in the time zone that is the default for the CMAS to which it connects. Allow the time zone of a managed CICS system to default to that of its CMAS. You can alter the time zone of a managed CICS system later while the CICS system is running. Any change that you make in this way lasts for the lifetime of the CICS system or until you change it again, whichever occurs first.
- You must specify a time zone for every CICSplex when you first define it. This time zone is used by the CICSPlex System Manager monitor function to determine the actual time at which the monitor interval for the CICSplex expires. You can alter the CICSplex time zone by using the CICSPlex System Manager user interface.

Time zones are specified by using single-character codes B–Z. For example, code S represents Mountain Standard Time, code T represents Central Standard Time, and code C represents Eastern Europe Time. For a complete list of the codes, see CICS Transaction Server for z/OS 5.5: CICSPlex SM Administration

. CICSPlex System Manager allows offsets (known as time zone adjustments) in the range of 0–59 minutes to accommodate regions that do not run in standard time zones. You can also specify DST.

Because multiple CICSPlex System Manager entities require a time zone to be specified, there is potential for conflicting time zones. For example, a CMAS and a CICS instance in the same CICSplex might be in different time zones. CICSPlex System Manager always acknowledges the time zone of the CMAS. Suppose that the following conditions exist:

- The time period definition time zone is S.
- The CMAS time zone is B.
- The CMAS time zone is C.

Time zone C is used by the CMAS, and the CMAS makes any necessary adjustments between time zones B, C, and S to ensure that the time zone is acknowledged.

Managing a multiple time zone environment

Note: In this section, it might be preferable to use the z/OS SET TIMEZONE command instead of the z/OS SET CLOCK command. For more information, see “SET TIMEZONE” on page 132.

Because each sysplex member in this example is in a different time zone, you must stagger DST changes. You can accomplish DST changes individually in each sysplex member by using the SET CLOCK command from the z/OS console.

The cities in the following examples 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 5-31, each of the three cities accurately reflects the number of hours from the Greenwich meridian when DST is in effect in the southern hemisphere. Sydney's TIMEZONE parameter is shown while DST is in effect. The London and Atlanta TIMEZONE parameters are shown while standard time is in effect. Each hemisphere changes the offset in different directions. When the southern hemisphere changes from DST to standard time (back), the northern hemisphere changes from standard time to DST (forward).

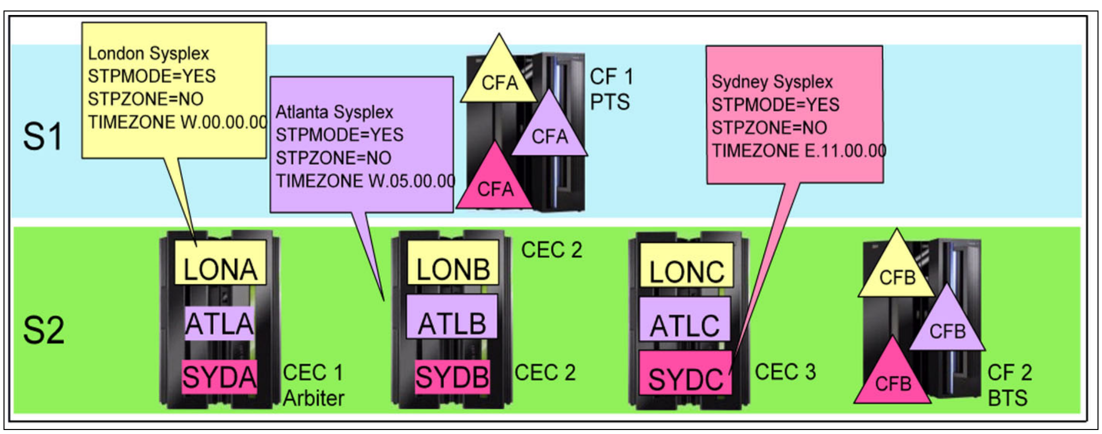

Figure 5-31 Staggered offset changes: one sysplex supporting multiple time zones

When you plan for DST changes in different countries, not all countries agree on the dates when the switch occurs. In fact, each country's government might change the DST switch dates from year to year. A country that has multiple time zones might have different change dates in each zone. In this example, assume that all countries switch either to DST (London and Atlanta) or to standard time (Sydney) on the same arbitrary day.

Also, assume that this sysplex is in Los Angeles, California. Any z/OS images representing Pacific Standard Time in the United States are not shown in Figure 5-31 on page 137. However, assume that those images contain STPZONE YES and that STP is being used to obtain DST offsets.

The world time zones that are shown in Table 5-5 represent only the local standard time offset from the Greenwich meridian and do *not* include DST offsets in either hemisphere.

Table 5-5 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

Determine which sysplex members change their offset first. The sysplex members in Sydney, SYDA, SYDB, and SYDC, are the first to change because a new day begins in Sydney before London and Atlanta. The operators in Los Angeles must be at Sydney's z/OS console to enter the SET TIMEZONE command at the correct time.

Determine the local time in Los Angeles when Sydney changes its DST offset. 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 changes from 02:00 a.m. to 01:00 a.m. (back one hour).

Several hours later, you must do the same for the London sysplex members. Again, assuming that London changes 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 changes the offset in the opposite direction from Sydney. London moves from standard time to DST (forward).

Atlanta is the next location to change the time. Atlanta also moves from standard time to DST. The Los Angeles operator must enter the SET CLOCK command on Atlanta's z/OS console at 11:00 p.m.

You can change the offset for Los Angeles through STP in the Adjust time zone offset window. These images have STPZONE YES in their CLOCKxx members.

Note: Each customer situation is different. Certain ones are more complicated, and others less so. Planning and excellent communications with the user community are essential when DST schedules are 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 you enter a SET CLOCK command, 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.

Abbreviations and acronyms

AAR	Arbiter Assisted Recovery
AEAD	Authenticated Encryption with Associated Data
AID	adapter ID
ARB	Arbiter
BIPM	Bureau International des Poids et Mesures
BMCA	Best Master Clock Algorithm
BPA	Bulk Power Assembly
BTCA	Best TimeTransmitter Clock Algorithm
BTS	Backup Time Server
CAR	Console Assisted Recovery
CAT	consolidated audit trail
CE LR	Coupling Express Long Reach
CF	Coupling Facility
CFCC	Coupling Facility Control Code
CFs	Coupling Facility commands
CMAS	CICSPlex System Manager address space
CMG	Computer Measurement Group
CP	Control Program
CST	Coordinated Server Time
CTN	Coordinated Timing Network
CTS	Current Time Server
DR	disaster recovery
DST	Daylight Saving Time
E2E	End-to-End
EDA	Enhanced Drawer Availability
EPSPT	Enhanced Preventive Service Planning Tool
ESMA	European Securities and Markets Authority
ETR	External Time Reference
ETS	External Time Source
EU	European Union
FINRA	Financial Industry Regulatory Authority
FQDN	fully qualified domain name
FSP	flexible support processor
GAS	Going Away Signal
GDPS®	Geographically Dispersed Parallel Sysplex
GMT	Greenwich Mean Time
HA	high availability
HADR	high availability and disaster recovery
HCA	Host Channel Adapter
HCD	Hardware Configuration Dialog
HCM	Hardware Configuration Manager
HFT	High Frequency Trading
HMC	Hardware Management Console
IBF	Internal Battery Feature
IBM SSR	IBM System Services Representative
IBM	International Business Machines Corporation
IC	Internal Coupling
ICA SR	Integrated Coupling Adapter Short Reach
ICB	Integrated Cluster Bus
IERS	International Earth Rotation and Reference Systems Service
IMPP	Installation Manual for Physical Planning
IPv4	IP address and netmask information
ISC	Inter-System Channel
ITU	International Telecommunication Union
LAN	Local area network
LEO	Low Earth Orbit
LIC	Licensed Internal Code
LPAR	Logical partition
MCL	maintenance change level
MIF	Multiple Image Facility
MIS	Management Information Systems
MRO	Multi-region Operation
MTOF	Message Time Ordering Facility
MiFID	Market in Financial Instruments Directive
NASD	National Association of Securities Dealers
NAT	network address translation

NIST National Institute of Standards and Technology

NTP Network Time Protocol

NTS Network Time Security

OADM optical add-drop module

OATS Order Audit Trail System

OCXOs Oven-controlled crystal oscillators

OS operating system

OSC Oscillator Card

OSs operating systems

OTN optical transport network

PAR Project Authorization Request

PCHID physical channel identifier

PDU Power Distribution Unit

PNT Position/Navigation/Timing

POR Power on Reset

PPS Pulse per Second

PR/SM Processor Resource/Systems Manager

PSP Preventive Service Planning

PTP Precision Time Protocol

PTS Preferred Time Server or Primary Time Server

RII Redundant I/O Interconnect

RTS Regulatory Technical Standard

SDO standards development organization

SE Service Elements

SE Support Element

SEC Securities and Exchange Commission

SFP small form-factor pluggable

SI System of Units

SNTP Simple Network Time Protocol

SOO support element workplace

SPE small programming enhancement

STCK STORE CLOCK

STCKE STORE CLOCK EXTENDED

STCKF STORE CLOCK FAST

STP Server Time Protocol

TAI Temps Atomique International also known as International Atomic Time

TLS Transport Layer Security

TOD Time of Day

TZ Time Zone

UDP User Datagram Protocol

UPS uninterruptible power supply

UT1 Universal Time 1

UTC Coordinated Universal Time also known as Universal Time Coordinated also known as temps universel coordonné

V2 Version 2

VCHID Virtual Channel Path Identifier

WANs wide area networks

WDM Wavelength Division Multiplexer

XRC Extended Remote Copy

iPDU intelligent Power Distribution Unit

z/TPF PUT z/Transaction Processing Facility Product Update

Related publications

The publications that are listed in this section are considered suitable for a more detailed description of the topics that are covered in this book.

IBM Redbooks

The following IBM Redbooks publications provide more information about the topics in this document. Some publications that are referenced in this list might be available in softcopy only.

- *IBM z15 (8561) Technical Guide*, SG24-8851
- *IBM z15 (8562) Technical Guide*, SG24-8852
- *IBM z15 Technical Introduction*, SG24-8850
- *IBM z16 Technical Introduction*, SG24-8950
- *IBM z16 (3931) Technical Guide*, SG24-8951
- *IBM z16 A02 and IBM z16 AGZ Technical Guide*, SG24-8952
- *IBM z17 Technical Introduction*, SG24-8580
- *IBM z17 (9175) Technical Guide*. SG24-8579
- *IBM Z Connectivity Handbook*, SG24-5444

You can search for, view, download, or order these documents and other Redbooks, Redpapers, web docs, drafts, and additional materials, at the following website:

ibm.com/redbooks

Other publications

The following publication is also relevant as a further information source:

Hardware Management Console (HMC) Operations Guide Version 2.15.0:

https://www.ibm.com/support/pages/hardware-management-console-operations-guide-version-2150 (requires IBMid authentication).

Online resources

The IBM Resource Link for documentation and tools website is also relevant as another information source:

http://www.ibm.com/servers/resourcelink (requires IBMid authentication)

Help from IBM

IBM Support and downloads:

ibm.com/support

IBM Global Services:

ibm.com/services

Redbooks IBM Z Time Synchronization Implementation Guide

(0.2”spine)
0.17”<->0.473”
90<->249 pages

SG24-8480-02

ISBN 073846225X

Printed in U.S.A.