Server Time Protocol Role Assignment Considerations

December 09, 2011 (update)

Gregory Hutchison (gkhutch@us.ibm.com)

Agenda

- Server roles in an STP-only CTN
- Server Role Planning
- Additional Information

Assumption:

Typically STP role assignment considerations apply to CTNs with 3 or more servers.

For a single CEC CTN or a two CEC CTN, see…

http://www.ibm.com/support/techdocs/atsmastr.nsf/WebIndex/TD105103

Server roles in an STP-only CTN

- All servers/CF's in STP-only CTN have to be STP configured
  - 9037s no longer required and not part of STP-only CTN
- Server/CF roles
  - **Preferred Time Server/CF (PTS)**
    - Server that is preferred to be the S1 server
  - **Backup Time Server/CF (BTS)**
    - Role is to take over as the S1 under planned or unplanned outages, without disrupting synchronization capability of STP-only CTN
  - **Current Time Server/CF(CTS)**
    - S1 Server/CF
      - Only one S1 allowed
      - Only the PTS or BTS can be assigned as the CTS
      - Normally the PTS is assigned the role of CTS
  - **Arbiter**
    - Provides additional means to determine if BTS should take over as the CTS under unplanned outages

Stratum levels

- STP transmits timekeeping information in layers or Stratums
  - Stratum 1 (S1)
    - Highest level in the hierarchy of timing network that uses STP to synchronize to Coordinated Server Time (CST)
  - Stratum 2 (S2)
    - Server/Coupling Facility (CF) that uses STP messages to synchronize to Stratum 1
  - Stratum 3 (S3)
    - Server/Coupling Facility (CF) that uses STP messages to synchronize to Stratum 2
- STP supports configurations up to Stratum 3

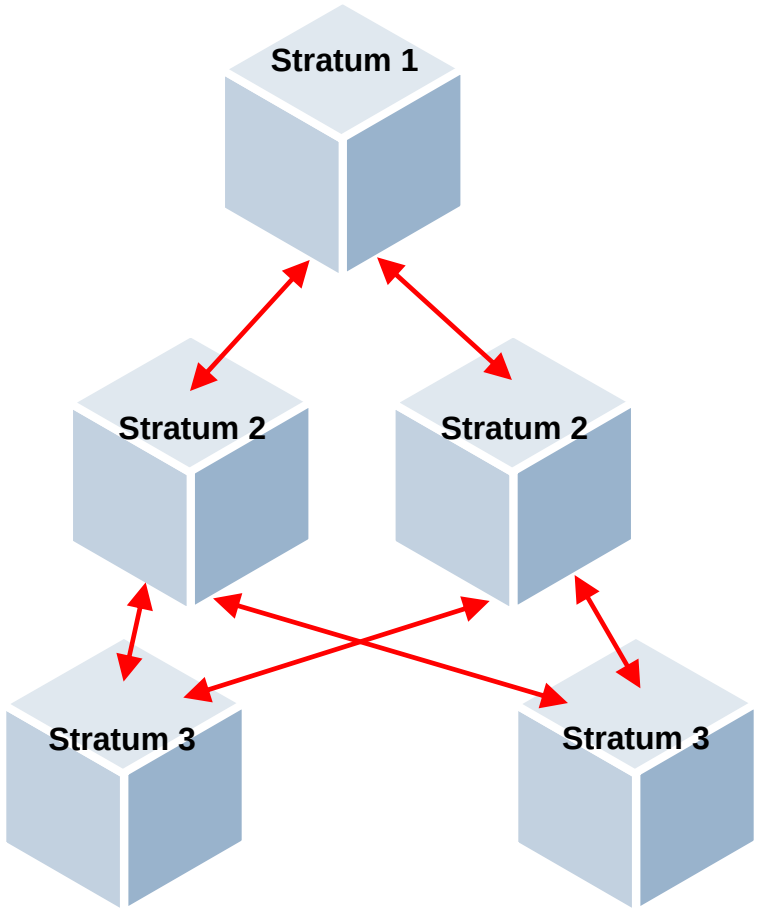

Connectivity rules and best practices recommendations

- Recommend PTS be assigned as the CTS during initial STP role assignments
  - Only the PTS can automatically re-takeover CTS role after recovery
  - PTS can automatically re-takeover CTS role only if:
    - None of the roles have been re-assigned in the interim
    - PTS has at least same connectivity as BTS
- The PTS, BTS, Arbiter must connect to each other
  - At least two coupling links between servers for redundancy.
  - To **temporarily** bypass STP connectivity check between these 3 servers, “Force” option can be used

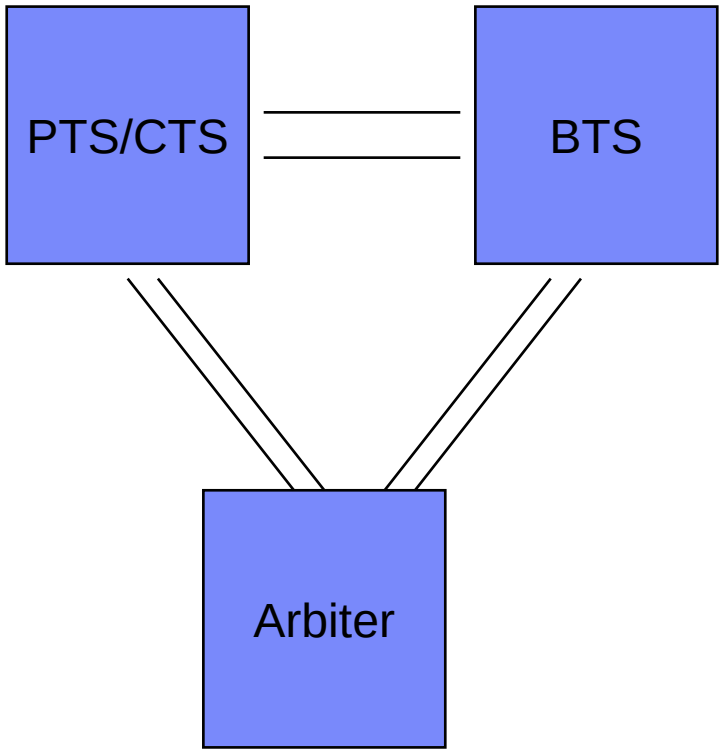

Connectivity rules and best practices recommendations –(contd)

- The BTS connectivity to other servers in the CTN should be such that when the BTS becomes the CTS, all the other servers in the CTN remain synchronized.
  - Servers may change Stratum levels from S2 to S3
- Servers stay synchronized after recovery, but recommendation for S3 server to have connectivity to at least two S2 servers not met.

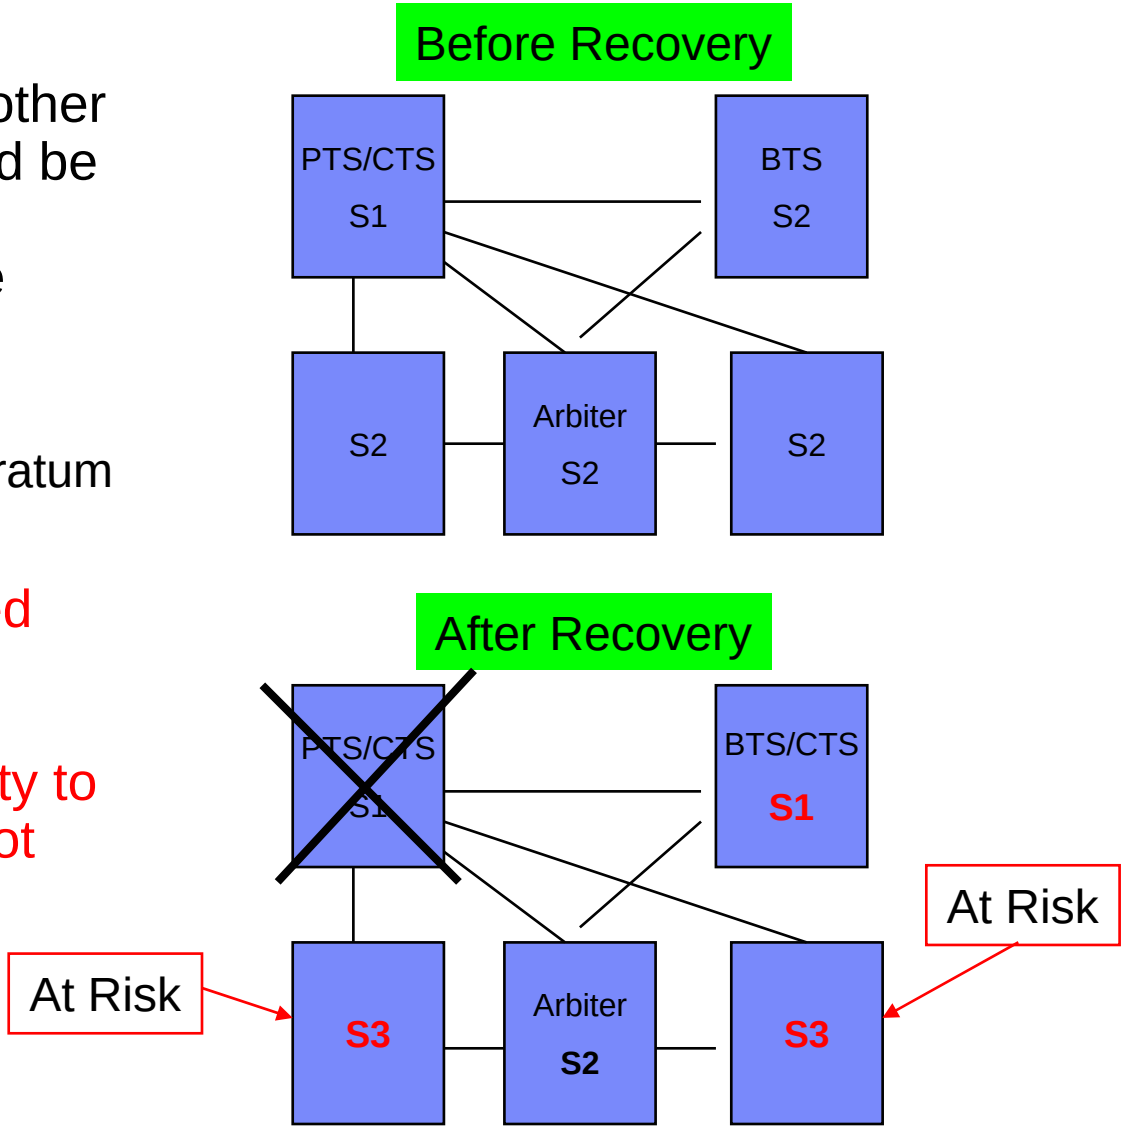

Best Practices recommendations – role reassignments

- Plan for which servers/CFs will be reassigned to be the PTS, BTS, Arbiter in the case when the **PTS** has an extended planned or unplanned outage.
  - The reassignment is to provide a robust fault tolerant configuration during the extended outage.
  - Without the reassignment, no other server has the capability to do a takeover, if a second failure now impacts the new CTS.
  - The reassignment can now be automated if the new API on the HMC is used.

Additional connectivity may be needed to follow best practices recommendations

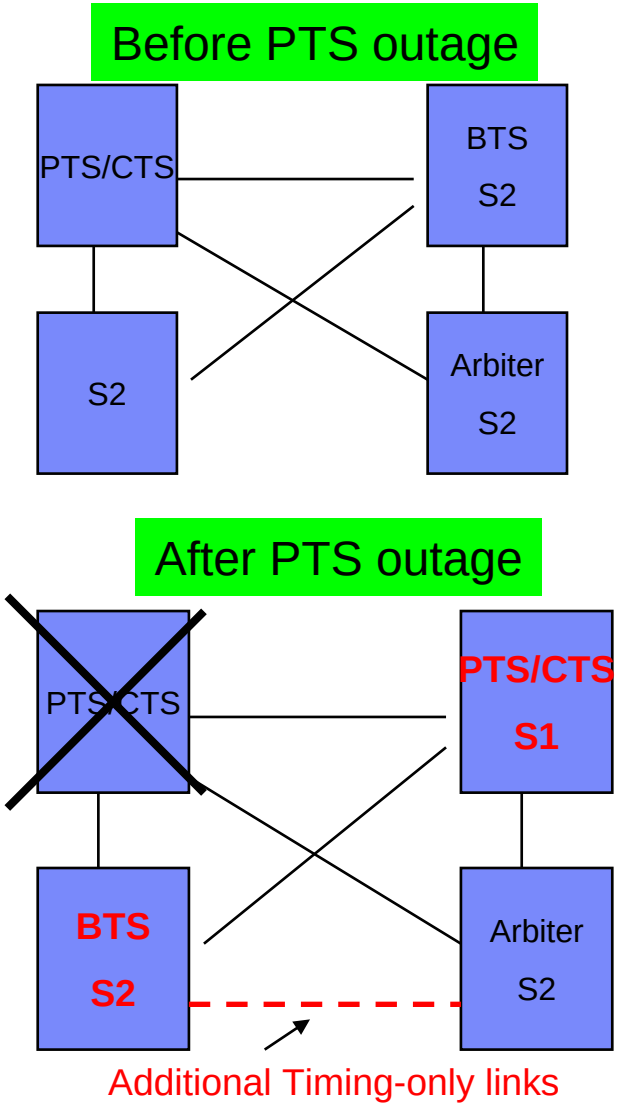

Best Practices recommendations – role reassignments (continued)

- Plan for which servers/CFs will be reassigned to be the PTS, BTS, Arbiter in the case when the **BTS** has an extended planned or unplanned outage.
- Plan for which servers/CFs will be reassigned to be the PTS, BTS, Arbiter in the case when the **Arbiter** has an extended planned or unplanned outage.
- If all servers with roles can not be reassigned, assign the affected server as “Not Configured”, however a CTS is always required.

Additional connectivity may be needed to follow best practices recommendations

PTS/CTS role planning

- **Location**
  - Most important site if multi-site configuration
  - Recovery considerations as to which site or server survives after outage
    - Note: For a two site CTN, recommendation is to locate Arbiter in same site as PTS/CTS
- **Technology**
  - Enhancements may be available only on latest server technology
  - Examples of functions not available on z990, z890
    - NTP Client support
    - Restore STP configuration for two server CTN
- **Maintenance**
  - Server requiring frequent maintenance may not be best choice
  - STP will block disruptive actions on server assigned CTS

PTS/CTS role planning (continued)

- **Connectivity**
  - Standalone CFs will typically provide best connectivity to other servers requiring time synchronization
  - Existing coupling link connectivity can be used to exchange STP messages, but ....
    - What additional connectivity do you need to follow connectivity rules and best practices recommendations for reassigning roles?
  - Standalone CFs will not produce z/OS messages for operations or interception by automation routines
    - Information messages at IPL or interrupt time may not be displayed
      - IEA380I THIS SYSTEM IS NOW OPERATING IN STP TIMING MODE.
      - IEA381I THE STP FACILITY IS NOT USABLE. SYSTEM CONTINUES IN LOCAL TIMING MODE
    - Warning messages identifying a condition being associated for a given server.
      - If the condition being raised relates to connectivity between two servers, the information might be available to a z/OS system image at the other end of the link.
      - However, if both ends of the link are CF partitions, no warning message is available to the user.

Examples of z/OS warning messages

- If the condition being raised relates to connectivity between two servers, the information might be available to a z/OS system image at the other end of the link.
  - Example 1: If one of the two coupling links between PTS/CTS and BTS fails,
    - z/OS image P1 will still report
      - **IEA382I** THIS SERVER HAS ONLY A SINGLE LINK AVAILABLE FOR TIMING PURPOSES
  - Example 2: If both ends of the link are CF partitions, no warning message is available to the user on the z/OS console
    - HMC will post a hardware message related to the link failure, but not a warning message that you have a potential SPOF

P1, P2, P3 – z/OS images in a Sysplex

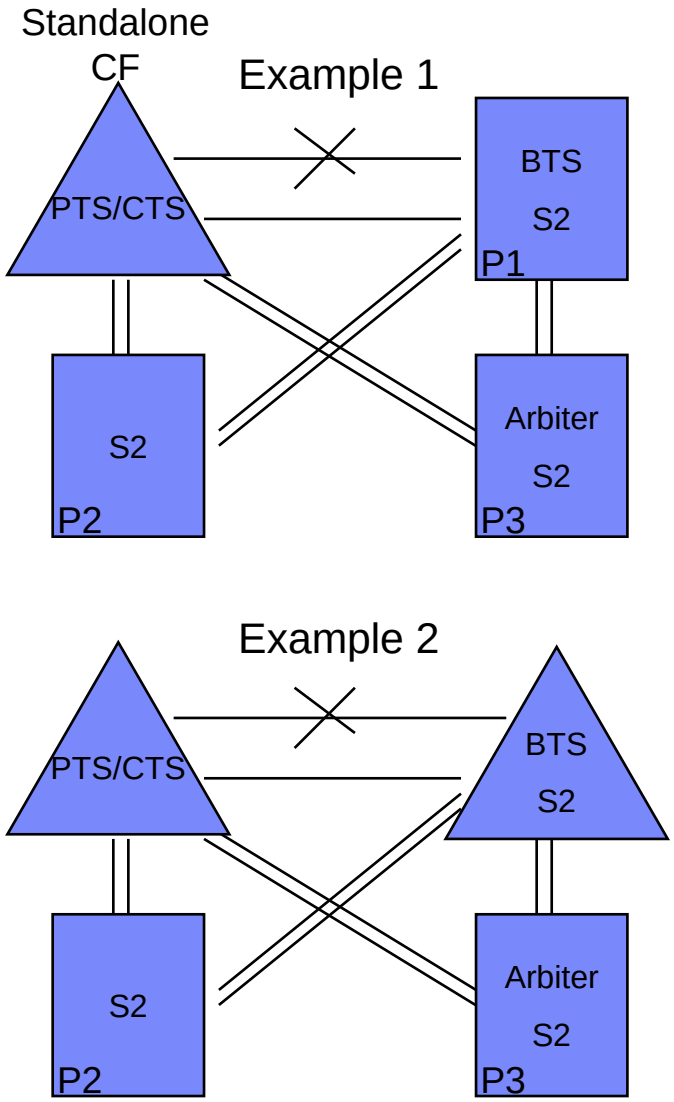

PTS/CTS role planning - z/OS Messages

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

  IXC438I COORDINATED TIMING INFORMATION HAS BEEN UPDATED

  FOR SYSTEM: sysname

  PREVIOUS CTNID: IBM_NET1-15

  CURRENT CTNID: IBM_NET1

- Question: Is it significant that these IEAxxx and IXCxxx messages are not displayed if the Standalone CF is assigned the role of PTS/CTS?

For a list of Supervisor, XCF/XES messages related to STP see Backup Slides

PTS/CTS role planning - z/OS Messages (contd)

- **In general, there are no z/OS messages that are posted *only* on the PTS, BTS, or Arbiter.**
- **Certain messages will not appear on the CTS since it is the time source**
  - IEA382I THIS SERVER HAS ONLY A SINGLE LINK AVAILABLE FOR TIMING PURPOSES.
  - IEA383I THIS SERVER RECEIVES TIMING SIGNALS FROM ONLY ONE OTHER NETWORK NODE.
  - IEA281I STP SYNC CHECK THRESHOLD EXCEEDED. CPC CONTINUES IN LOCAL MODE.
  - IEA390I TOD CLOCKS DYNAMICALLY ADJUSTED TO MAINTAIN STP SYNCHRONISM.
- **Following message may not appear on some special role servers**
  - IEA388I THIS SERVER HAS NO CONNECTION TO THE nnnnnnnnn
    - Where nnnnnnnnnnn = 'PREFERRED '| 'BACKUP '| 'ARBITER '
  - For example,
    - IEA388I THIS SERVER HAS NO CONNECTION TO THE BACKUP will never appear on a z/OS system running on the BTS.

Need to examine complete list of STP z/OS messages before deciding significance of Standalone CF not displaying STP messages

Summary – Server role planning

- After deciding which server assigned PTS/CTS, similar review should be made for the BTS
  - If multi-site sysplex, Arbiter should be placed in same site as PTS/CTS
- Many considerations to evaluate before deciding which servers assigned special roles
  - Location, Technology, Maintenance
  - Connectivity
  - Stand alone CF
  - z/OS message automation
- Consider following best practices recommendations for reassignment of roles
- Every customer needs to do their own assessment
  - No single recommendation fits all configurations

Backup Slides

STP Messages - Supervisor

- **IEA031I** STP ALERT RECEIVED. STP ALERT CODE = nn
  - Alert code is x'01' to x'84‘
- **IEA380I** THIS SYSTEM IS NOW OPERATING IN STP MODE.
- **IEA381I** THE STP FACILITY IS NOT USABLE. SYSTEM CONTINUES IN LOCAL MODE.
- **IEA382I** THIS SERVER HAS ONLY A SINGLE LINK AVAILABLE FOR TIMING PURPOSES.
- **IEA383I** THIS SERVER RECEIVES TIMING SIGNALS FROM ONLY ONE OTHER NETWORK NODE.
- **IEA384I** SETETR COMMAND IS NOT VALID IN STP TIMING MODE.
- **IEA385I** CLOCKxx ETRDELTA & TIMEDELTA BOTH SPECIFIED. yyyyyyy IGNORED.
- **IEA387I** STP DATA CANNOT BE ACCESSED. SYSTEM CONTINUES IN yyyyy TIMING MODE.
- **IEA388I** THIS SERVER HAS NO CONNECTION TO THE nnnnnnnnnn
- **IEA389I** THIS STP NETWORK HAS NO SERVER TO ACT AS nnnnnnnnnn
- **IEA392I** STP TIME OFFSET CHANGES HAVE OCCURRED.
- **IEA393I** ETR PORT n IS NOT OPERATIONAL. THIS MAY BE A CTN CONFIGURATION CHANGE.
- **IEA394A** THIS SERVER HAS LOST CONNECTION TO ITS SOURCE OF TIME.
- **IEA395I** THE CURRENT TIME SERVER HAS CHANGED TO THE ccccccccccc
  - (where cccccc is BACKUP or PREFERRED)

STP ETR Messages - XCF

- **IXC434I** SYSTEM sysname HAS TIMING DEFINITIONS THAT ARE NOT CONSISTENT WITH THE OTHER ACTIVE SYSTEMS IN SYSPLEX sysplex-name.
  - EFFECTIVE CLOCK VALUES ARE NOT CONSISTENT
  - SYSTEM: sysname IS RUNNING IN LOCAL MODE
  - SYSTEM: sysname IS USING ETR NET ID: xx
  - SYSTEM: sysname IS USING CTN ID: yyyyyyyy-xx
  - SYSTEM: sysname IS USING CTN ID: yyyyyyyy
  - where:
    - xx = ETR NET ID in use by system sysname
    - yyyyyyyy = CTN ID in use by system sysname
    - yyyyyyyy-xx = CTN ID in use by system sysname

- **IXC435I** ALL SYSTEMS IN SYSPLEX sysplexname ARE NOW SYNCHRONIZED TO THE SAME TIME REFERENCE.
  - SYSTEM: sysname IS USING ETR NETID: ee
  - SYSTEM: sysname IS USING CTNID sssssss-ee
  - SYSTEM: sysname IS USING CTNID sssssss

STP ETR Messages – XCF (continued)

- **IXC437I** - SYSTEMS CAN NOW ENTER THE SYSPLEX USING
  - COORDINATED SERVER TIME CTNID=ssssssss-ee
  - COORDINATED SERVER TIME CTNID=ssssssss

- **IXC438I** COORDINATED TIMING INFORMATION HAS BEEN UPDATED FOR SYSTEM: sysname
  - PREVIOUS ETR NETID: ee
  - PREVIOUS CTNID: ssssssss-ee
  - PREVIOUS CTNID: ssssssss

  - CURRENT ETR NETID: ee
  - CURRENT CTNID: ssssssss
  - CURRENT CTNID: ssssssss-ee
  - CURRENT TIMING: LOCAL

- **IXC439E** ALL SYSTEMS IN SYSPLEX sysplexname ARE NOT SYNCHRONIZED TO THE SAME TIME REFERENCE.
  - SYSTEM: sysname IS USING ETR NETID: ee
  - SYSTEM: sysname IS USING CTNID sssssss-ee
  - SYSTEM: sysname IS USING CTNID sssssss
  - SYSTEM: sysname IS RUNNING IN LOCAL MODE

- **IXC468W** XCF IS UNABLE TO ACCESS THE CTN AND HAS PLACED THIS SYSTEM INTO NON-RESTARTABLE
  - WAIT STATE CODE: 0A2 REASON CODE: 158

STP ETR Messages – XES

- **IXL160E** CF REQUEST TIME ORDERING: REQUIRED AND NOT-ENABLED
  - COUPLING FACILITY SIMDEV.IBM.EN.CF0100000000
    - PARTITION: 00 CPCID: 00
  - REASON: ETR NETID MISMATCH - CF ETR NETID: 0F
  - REASON: ETR NOT CONNECTED TO COUPLING FACILITY
  - REASON: MESSAGE TIME ORDERING FUNCTION FAILURE
  - REASON: TIME ORDERING FUNCTION NOT INSTALLED ON THE CF
  - REASON: TIME ORDERING FUNCTION NOT INSTALLED ON THIS SYSTEM
  - REASON: CF IS OUT OF SYNC WITH TIMING NETWORK
  - REASON: CTNID MISMATCH - CF CTNID: CTNNET01 0F
- **IXL162E** CF REQUEST TIME ORDERING: REQUIRED AND WILL NOT BE ENABLED
  - COUPLING FACILITY SIMDEV.IBM.EN.CF0100000000
    - PARTITION: 00 CPCID: 00
  - REASON: CTNID MISMATCH - CF CTNID: CTNNET01 0F

Additional Information

- **Redbooks®**
  - Server Time Protocol Planning Guide SG24-7280
  - Server Time Protocol Implementation Guide SG24-7281
- **Education**
  - Introduction to Server Time Protocol (STP)
    - Available on Resource Link™
    - www.ibm.com/servers/resourcelink/hom03010.nsf?OpenDatabase
- **STP Web site**
  - www.ibm.com/systems/z/pso/stp.html
- **Systems Assurance**
  - The IBM team required to complete a Systems Assurance Review (SAPR Guide SA06-012) and complete Systems Assurance Confirmation Form via Resource Link
    - http://w3-03.ibm.com/support/assure/assur30i.nsf/WebIndex/SA779
- **Techdocs and WSC Flashes**
  - http://www-03.ibm.com/support/techdocs/atsmastr.nsf/Web/Techdocs
    - Search on “STP”

IBM Implementation Services for System z – Server Time Protocol (6948-J56)

Offering Description	• **This offering is designed to assist clients to quickly and safely implement Server Time Protocol within their existing environments. STP provides clients with the capability to efficiently manage time synchronization within their multi-server infrastructure. Following best practices and using detailed planning services, IBM helps clients identify various implementation models and engage in the appropriate configuration required to effectively support STP for driving a more responsive business and IT infrastructure.**
Program, Play, Industry Alignment	• **Infrastructure Improvement; Energy Efficiency; Better performance and lower operational cost**
Client Value ***(enables customers to...)***	• **Swift and secure implementation of STP for improved availability, integrity and performance** • **Improves multi-server time synchronization without interrupting operations** • **Enables integration with next generation of System z infrastructure**
Target Audience	• **Primarily core, Large Enterprise customers.** • **Existing z midrange clients**
Key Competitors	• **In house staff**
Competitive Differentiation	• **Leverages best practices with secure implementation** • **Short implementation time – lower risk** • **Provides support and facilitates knowledge sharing through IBM's mainframe expertise**
Proof Points & Claims for Client Value / Differentiation	• **Need to safely implement a reliable replacement for Sysplex Timer® while maintaining continuous operations** • **Cost of providing and maintaining hardware, floor space and solution support for additional Sysplex Timer intermediate site** • **Lack of in-house expertise, skills and resources for implementing Server Time Protocol**
Engagement Portfolio	• **http://spimweb1.boulder.ibm.com/services/sosf/dyno.wss?oid=50423&loc=All&langcd=en-US#1**
Offering Manager	• **Anna Lee/Southbury/IBM, 512-590-8914, T/L: 268-9318**

IBM Announces –
IBM Implementation Services for System z – Server Time Protocol

Implementation of STP for improved availability and performance

Offering

Assist clients to quickly and safely implement Server Time Protocol within their existing environments. IBM helps clients identify various implementation models and engage in the appropriate configuration required to effectively support STP for driving a more responsive business and IT infrastructure

Customer Value:

- Improves multi-server time synchronization without interrupting operations
- Enables integration with next generation of System z infrastructure
- Swift and secure implementation of STP for improved availability, integrity, and performance
- Reduces hardware maintenance and power costs while eliminating intermediate site requirements for Sysplex Timer

Leverages IBM's knowledge and best practices to help implementation of Server Time Protocol

Reference Material - Terminology

APAR	Authorized Program Analysis Report
ARB	Arbiter
BTS	Backup Time Server
CF	Coupling Facility
CTS	Current Time Server
CTN	Coordinated Timing Network
DWDM	Dense Wave Division Multiplexer
ETR	External Time Reference
ETS	External Time Source
FC	Feature Code
HMC	Hardware Management Console
HCA	Host Channel Adapter
ICB	Integrated Cluster Bus
IPL	Initial Program Load
ISC	InterSystem Coupling Channel
LAN	Local Area Network
LIC	Licensed Internal Code
LPAR	Logically Partition
NTP	Network Time Protocol
PR/SM	Processor Resource / Systems Manager
PSIFB Infiniband	Parallel Sysplex
PTF	Temporary Program Fix
PTS	Preferred Time Server
SW and	Software (programs operating systems)
SE	Support Element
TPF	Operating System
UTC	Coordinated Universal Time
zVM	Operating System
zVSE	Operating System
z/OS	Operating System
z/VM	Operating System

Questions?

- **What are your questions?**

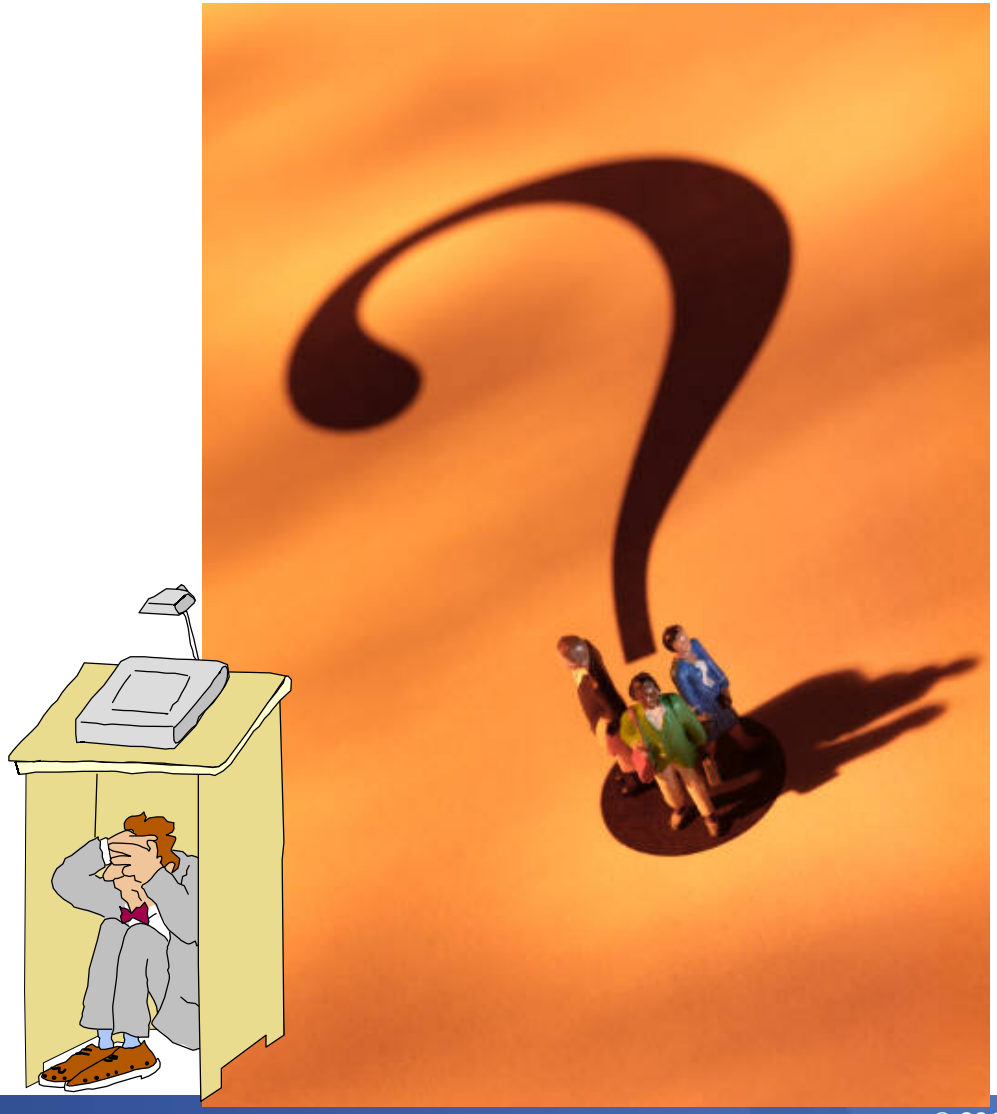

Trademarks

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

Not all common law marks used by IBM are listed on this page. Failure of a mark to appear does not mean that IBM does not use the mark nor does it mean that the product is not actively marketed or is not significant within its relevant market.

Those trademarks followed by ® are registered trademarks of IBM in the United States; all others are trademarks or common law marks of IBM in the United States.

For a complete list of IBM Trademarks, see www.ibm.com/legal/copytrade.shtml:

*, AS/400®, e business(logo)®, DBE, ESCO, eServer, FICON, IBM®, IBM (logo)®, iSeries®, MVS, OS/390®, pSeries®, RS/6000®, S/30, VM/ESA®, VSE/ESA,

WebSphere®, xSeries®, z/OS®, zSeries®, z/VM®, System i, System i5, System p, System p5, System x, System z, System z9®, System z10®, BladeCenter®

The following are trademarks or registered trademarks of other companies.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.
Cell Broadband Engine is a trademark of Sony Computer Entertainment, Inc. in the United States, other countries, or both and is used under license therefrom.
Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.
Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.
Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
UNIX is a registered trademark of The Open Group in the United States and other countries.
Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.
ITIL is a registered trademark, and a registered community trademark of the Office of Government Commerce, and is registered in the U.S. Patent and Trademark Office.
IT Infrastructure Library is a registered trademark of the Central Computer and Telecommunications Agency, which is now part of the Office of Government Commerce.

* All other products may be trademarks or registered trademarks of their respective companies.

Notes:
Performance is in Internal Throughput Rate (ITR) ratio based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput that any user will experience will vary depending upon considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve throughput improvements equivalent to the performance ratios stated here.
IBM hardware products are manufactured from new parts, or new and serviceable used parts. Regardless, our warranty terms apply.
All customer examples cited or described in this presentation are presented as illustrations of the manner in which some customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics will vary depending on individual customer configurations and conditions.
This publication was produced in the United States. IBM may not offer the products, services or features discussed in this document in other countries, and the information may be subject to change without notice. Consult your local IBM business contact for information on the product or services available in your area.
All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.
Information about non-IBM products is obtained from the manufacturers of those products or their published announcements. IBM has not tested those products and cannot confirm the 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.
Prices subject to change without notice. Contact your IBM representative or Business Partner for the most current pricing in your geography.