An Introduction to Circuit Emulation Services

Contents

Introduction
Before You Begin
Conventions
Prerequisites
Components Used
CES Concepts
Types of Signaling
Frame and Cell Processing
Types of CES
Clocking in Circuit Emulation
Clocking Modes
Clock Distribution
Configuring CES
Before You Configure
Sample Configurations
Verifying Configurations
Basic Troubleshooting
The Circuit is not Coming up or the CPE is in Alarm
The Circuit is Experiencing Clocking Slips
PBX Reporting Framing Errors or Carrier Drops
Users Hear Static or Clicking on Phone Calls
You Suspect a Bad Reference Clock
There are Clocking Problems in a Network with PA-A2
Related Information

Introduction

Circuit Emulation Service (CES) allows DS-n and E-n circuits to be transparently extended across an ATM network using constant bit rate (CBR) ATM permanent virtual circuits (PVCs) or soft PVCs. CES is based on the ATM Forum standard af-vtoa-0078.0000 (PDF). This standard defines the CES Interworking Function (CES-IWF), which allows communication between non-ATM CBR circuits (such as T1, E1, E3, and T3) and ATM UNI interfaces. CES is typically implemented on ATM switches, but it can be implemented on ATM edge devices (such as routers) as well. CES is mostly used for communication between non-ATM telephony devices (like PBXs, TDM, and channel banks) or video devices (such as CODEC) and ATM devices (such as the Cisco LS1010 and Catalyst 8540-MSR ATM switch), or via an ATM uplinks (such asd the PA-A2 on the Cisco 7200 router).

Before You Begin

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

Prerequisites

There are no specific prerequisites for this document.

Components Used

This document is not restricted to specific software and hardware versions.

The information presented in this document was created from devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If you are working in a live network, ensure that you understand the potential impact of any command before using it.

CES Concepts

This section introduces some basic CES terminology. Refer to the sub-topics within this section for more details.

Note: This document focuses more on T1 examples, but you can apply the theory to E1 as well.

CES is typically used to transfer voice or video traffic across an ATM network. Voice and video, unlike data traffic, is very sensitive to delay and delay variance. CES uses virtual circuits (VCs) of the CBR ATM service category, which guarantees acceptable delay and delay variation. Therefore, it satisfies both voice and video traffic requirements. ATM adaptation layer one (AAL1) specified by ITU-T.I.363.1 is used at CES-IWF.

Some typical applications of CES are listed below:

- Extension of a private telephone network across multiple campuses, as illustrated below. For example, there are two campuses with a private branch exchange (PBX) on each. You can use an ATM network to connect two the PBXs without having ATM capabilities on the PBX itself. By doing so, voice traffic between two campuses uses your private ATM backbone instead of leased lines, thereby using the same ATM network for your voice and data needs.

- Video conferencing between multiple sites, as illustrated below:

The ATM Forum defined CES-IWF for many types of Telco circuits (such as DS-1, DS-3, E-1, E-3, J-1 and J-3), but for CES-IWF, the most common types are DS-1 service and E-1 service. In the Enterprise arena, Cisco provides T-1 and E-1 CES on the 8510-MSR, Catalyst 8540-MSR, and PA-A2 port adapter for the 7200 series router. Cisco also supports CES on some of its service provider products like the MGX 8220. However, this document concentrates on Enterprise products.

CES-IWF converts the entire DS-n or E-n frame arriving from the customer premises equipment (CPE) (such as a PBX) into AAL1 ATM cells and transmits them across the ATM network using a single VC. The ATM switch or router on the remote end converts the AAL1 ATM cells into a DS-n or E-n frame, which is then transmitted to a Ds-n or E-n CPE device. This type of CES is called *unstructured CES*, which extends clear channel T1 (all 24 channels) across an ATM network (on a single VC).

In addition to this basic functionality, CES supports channelized T1 services by splitting T1 into multiple Nx64k circuits and transmitting those channelized T1 circuits across different ATM VCs with one or multiple destinations. This allows, for example, a single PBX to communicate with multiple remote PBXs using a single T1 port on a hub PBX. This type of hub-and-spoke example, known as *structured CES*, is shown below.

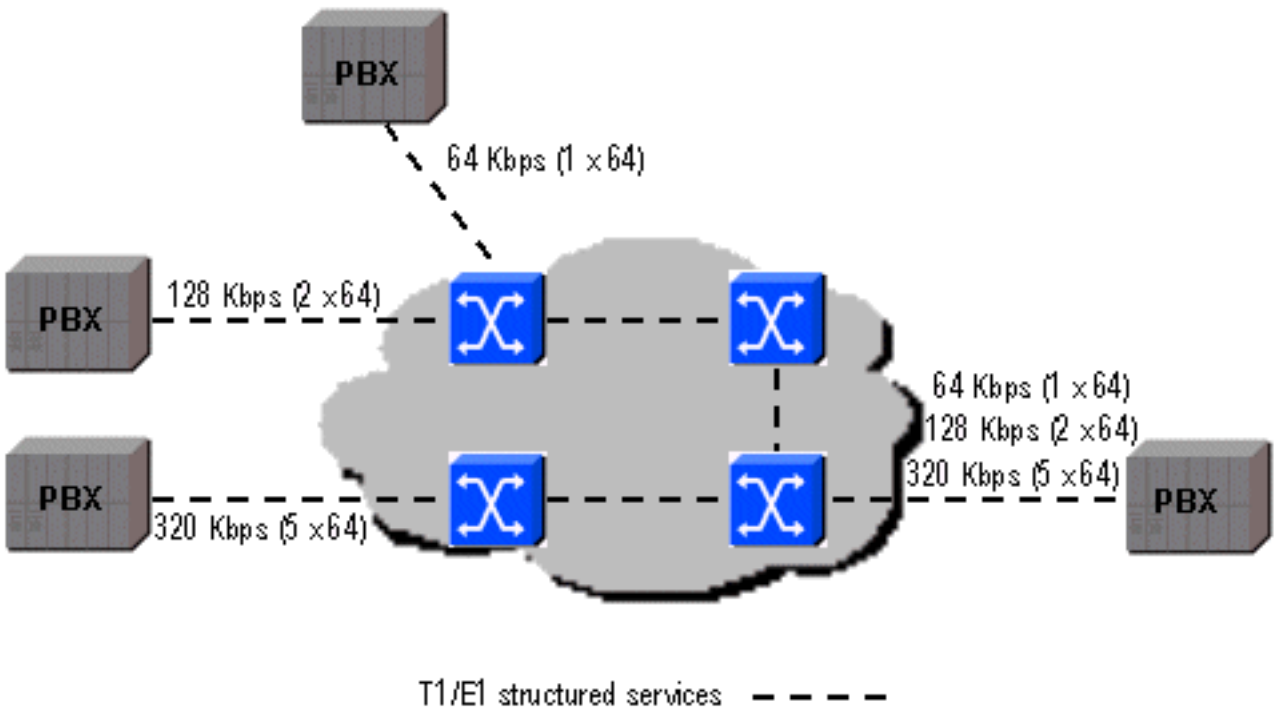

Types of Signaling

There are two types of signaling associated with T1 and T1 circuit emulation: channel associated signaling (CAS) and common channel signaling (CCS). CAS is in-band signaling and CCS is out-of-band signaling.

You can typically use CAS to transparently transmit proprietary signaling protocols that use the ABCD bits of a T1 frame. On Cisco ATM switches configured for CAS, the ABCD bits will not be changed or acted upon, which provides the extension of proprietary signaling across the ATM network.

Note: You need to use structured CES if you are providing CAS.

You can also use CAS for on-hook detection on Cisco enterprise ATM switches. CAS with on-hook detection is supported for (56k/64k) DS0 circuits only. CES-IWF mandates that voice be transmitted as CBR ATM traffic, a method that forces the ATM switch to reserve bandwidth for the voice circuit even when there is no user traffic (voice) to be sent. So when there is no voice communication, AAL1 cells are still using bandwidth on the ATM link sending "NULL" data. The solution for minimizing the "NULL" cells on ATM links is to not send "NULL" cells if there is no voice communication.

8510-MSR implements on-hook detection as follows:

- Detect on-hook/off-hook. This requires the ABCD pattern to be configured in a way that indicates the on-hook signal that the CPE is using. In other words, the CPE dictates how this must be configured on the 8510-MSR; the CPE and the 8510-MSR must be configured the same.
- Stop sending AAL1 cells when on-hook is detected.
- Indicate to the ATM switch that has the destination CBR circuit that it is in on-hook mode. This prevents the remote switch from declaring a loss of cell delineation (LCD) if no cells (data or "NULL") are received.
- Start sending AAL1 cells when on-hook is no longer detected (that is, when the ABCD pattern coming from the CPE equipment no longer matches the configured pattern).

Note: CAS with on-hook detection on the 8510-MSR can be used only if the CPE equipment supports CAS and can detect on-hook state.

Robbed bit signaling on Cisco Enterprise switches and routers is configured using the **ces dsx1 signalmode robbedbit** command. CAS and on-hook detection is configured using the **ces circuit** command.

CES ports on Cisco Enterprise switches support CAS, which "robs" one bit out of every channel in the sixth T1 frame in order to transmit signaling messages. CAS is also referred as "robbed bit signaling"; robbed bits are referred to as AB (in SF) or ABCD (in ESF) bits. CAS can be used for on-hook detection, which allows for better utilization of network resources at times when there is no user traffic.

CCS uses the entire channel of each basic T1 frame for signaling. An example of CCS is ISDN PRI, where entire the 64k D channel is used for signaling. CCS is not supported natively on Cisco LightStream and Catalyst ATM switches; however, the 8510-MSR (or 8540-MSR, LS1010) together with the Cisco VSC2700 signaling controller can provide similar function using Simple Gateway Control Protocol (SGCP). This solution is implemented by 8510-MSR propagating the signaling DS0 channel to the VSC2700 gateway, which is capable of understanding multiple signaling protocols and signaling back to the 8510-MSR the ATM address to which the 64k soft PVC needs to be set up. Once an end-to-end circuit is established, the 8510-MSR is responsible for transfer of user traffic. By doing bandwidth-on-demand in such a way, the total number of required interfaces is reduced and the need for tandem PBX can be eliminated.

CES can be implemented using PVCs or soft PVCs. PVC requires manual configuration on each ATM switch in the ATM cloud; soft PVC relies on ATM signaling to establish the VC, and VC configuration is required on only one ATM switch. Another advantage of soft PVC is that the VC can be re-routed in a case of link failure.

On the other hand, PVCs are more stable because they don't depend on any dynamic components, like ATM signaling. If an ATM network has ATM switches that do not support ATM signaling, PVCs are the only option. It is very important to note that clocking is of significant

importance to CES. The receiving T1 stream on a remote CPE must have the same clocking characteristics as the transmitting T1 stream. To ensure this, the ATM network must not significantly change clock characteristics. To achieve this, you can use one of several clocking schemes discussed in Clocking in Circuit Emulation.

Frame and Cell Processing

As previously mentioned, CES-IWF converts T1 frames to AAL1 ATM cells. The CES-IWF function is implemented on the CES port adapter module (PAM) of an ATM switch. In simpler terms, the T1 frame comes in CES PAM, where it gets buffered and segmented into 47-byte cells. One byte of AAL1 header is added to each 47-byte cell, forming a 48-byte cell. Five bytes of ATM cell header is added and the 53-byte cell is switched to the outgoing ATM interface. Depending on the type of CES service, additional steps might also occur. At the receiving end, the process is reversed.

Types of CES

CES services can be differentiated in two ways: synchronous versus asynchronous, and structured versus unstructured.

Synchronous versus Asynchronous

- Synchronous service assumes that synchronized clocks are available on each end. Therefore, no clocking information is transported in the ATM cell. Propagation of the clock source throughout the network is required.
- Asynchronous service sends clocking information in ATM cells to the remote end of the circuit. Clocking information sent in the ATM cell is called Synchronous Residual Time Stamp (SRTS).

The SRTS value is specified using four bits and is sent per eight cells using one bit in the AAL1 header for every odd sequence numbered cell. The reference clock must still be propogated throughout the network.

Structured versus Unstructured

- Unstructured service (also called "clear channel") utilizes the entire T1 bandwidth, (meaning there is one single channel). The ATM switch doesn't look into the T1, but simply reproduces a stream of bits with clocking from the receiving port to the target port.
- Structured service (also called *channelized T1* or *cross-connect*) is designed to emulate point-to-point Fractional T1 (Nx64k) connections. This allows the T1 to break into multiple DS-0 channels towards different destinations. More than one circuit (AAL1) entity will share the same physical T1 interface. To provide this service, AAL1 is capable of delineating repetitive fixed-size blocks of data (block size is the integral number of octets, where an octet represents a 64k channel).

For a block size larger than one octet, AAL1 uses a pointer mechanism to indicate the start of the structure block. A convergence sub-layer (CS) indicator (CSI) bit in the AAL1 header set to 1 indicates structured service, while a CSI bit of 0 indicates unstructured service. So, if CSI = 1, the pointer identifying the beginning of the structure is inserted in the CSI field of even numbered cells. Using this pointer, the receiving switch will know how to convert the AAL1 cells into the appropriate fractional T1.

On Cisco Enterprise switches and routers, this type of circuit emulation service is configured using the **ces aal1 service** command.

Clocking in Circuit Emulation

Clocking is so important for CES. This section focuses on two clocking concepts:

- clocking modes
- clock distribution

Clocking modes define multiple ways to achieve the same clock in the transmitting and receiving ends of a T1 circuit end-to-end. This means that the T1 stream that PBX1 transmits has same the clocking characteristics as the T1 stream that PBX2 receives, and vice versa.

Some clocking modes (like synchronous and SRTS) rely on a reference clock source that needs to be the same throughout the network. For those clocking modes, clock distribution of the reference clock source is required.

The following sections discuss various clocking modes and methods of clock distribution. We will also list advantages and disadvantages of each clocking mode.

Clocking Modes

There are three major clocking modes:

- Synchronous clocking
- SRTS
- Adaptive clocking

It is of significant importance to remark that accurate clocking distribution can be done with hardware support. The Phased Lock Loop (PLL) chip used to do this is only present in the ASP-PFQ card on the LS1010 and RP equipped with network clock modules on the 8540-MSR. Using those modules is *highly* recommended when designing ATM networks that use CES. Refer to Clocking Requirements for the LightStream 1010, Catalyst 8510-MSR and Catalyst 8540-MSR for more information.

Synchronous Clocking

Transmit clock frequency is produced by an external source (also called the primary reference signal [PRS]). PRS is distributed throughout the ATM network so that all devices can synchronize to the same clock.

Advantages	Disadvantages
Supports both structured and unstructured CES services	Requires network clock synchronization.

.	
Exhibits superior wander and jitter characteristics.	Ties CES interface to the PRS; in case of PRS failure, the circuit might be degraded unless redundant PRS is available.
	Other interfaces (besides the CBR or ATM interface used to derive the network clock on the ATM switch) might be affected in case of PRS failure because Cisco ATM switches use that derived clock as a system clock for all interfaces in the switch, not just interfaces involved with CES.

SRTS

SRTS is an asynchronous clocking method. SRTS measures the difference between the service clock (received on the CBR interface) and the network-wide reference clock. This difference is the Residual Time Stamp (RTS). The RTS is propagated to the remote end of the circuit in the AAL1 header. The receiving end reconstructs the clock by adjusting the reference clock by the RTS value. Keep in mind that the reference clock needs to be propagated throughout the network; in other words, the switch needs to be capable of distributing the clock.

Advantages	**Disadvantages**
Conveys an externally-generated user (such as PBX, MUX, or CODEC) clocking signal throughout the ATM network, providing an independent clocking signal for each CES circuit.	Requires network clock synchronization services.
Useful in networks that have multiple external clock sources.	Supports only unstructured CES services.
	Exhibits moderate wander and jitter.

Adaptive Clocking

In adaptive clocking, the source CES IWF simply sends the data to the destination CES IWF. The destination CES IWF writes data to the segmentation and re-assembly (SAR) buffer and reads it with the local T1 service clock. The local (interface) service clock is determined from the actual CBR data received.

The level of the SAR buffer controls the local clock frequency by continuously measuring the fill

level around the median position and feeding this measurement to drive the Phase Lock Loop (PLL), which in turn drives the local clock (transmit clock). So the transmit clock frequency is modified to keep the re-assembly buffer depth constant. When CES IWF senses that its SAR buffer is filling up, it increases the transmit clock rate. When CES IWF senses that the SAR buffer is emptying, it decreases the transmit clock rate.

The proper choice of buffer length can prevent buffer overflow and underflow and, at the same time, control delay (greater buffer size implies greater delay). The buffer length is proportional to the maximum cell delay variation (CDV), which the user can configure on Cisco ATM switches. The network administrator can estimate what the maximum CDV should be by summing the CDV of each network device in the circuit path. The sum of the measured CDVs that each piece of equipment introduces must be smaller than the maximum CDV configured. If not, underflows and overflows will occur. On Cisco equipment, you can view the actual CDV with the **show ces circuit interface cbr *x/y/z* 0** command if you are using unstructured service.

Advantages	**Disadvantages**
Does not require network clock synchronization.	Supports only unstructured CES.
	Exhibits poorest wander characteristics.

On Cisco Enterprise products, this clocking mode is configured using the **ces aal1 clock CBR interface** command.

Clock Distribution

Synchronous and SRTS clocking modes require distribution of the PRS throughout the network. If you use one of these two clocking modes, you will first have to choose which clock source will serve the role of PRS and design a network-level clock distribution topology.

Things to consider when deciding on the PRS is clock accuracy and position of PRS in network:

- Clock accuracy is determined by the stratum level. Typically, the service provider will provide better accuracy clock (stratum1 or 2) than local oscillators on equipment (ATM switches or CPE equipment). In absence of the service provider clock (which is often the case with video applications), choose the device with the most accurate local oscillator as the PRS.
- Another thing to consider when deciding on PRS is the position of the devices that will be the PRS in the network. This is typically the case if you have multiple potential clock sources with the same level of accuracy, or if you have a very large ATM network. You need to choose the position of the PRS so that it minimizes the number of network devices that the clock needs to traverse from the PRS to the edge devices because the clock gets degraded as it traverses network nodes.

Once you choose the PRS, the next decision is to find the best way to propagate the reference clock. The network distribution topology must be loop-free; in other words, it needs to be a tree structure or a set of trees. The clock distribution topology should also impose a strict hierarchical ordering of the active components of the topology based on the stratum level of the various network equipment. That is, if there are two equal-hop paths to choose from, choose the one that goes through the more accurate equipment (lower stratum).

See the network clock distribution tree in the following illustration:

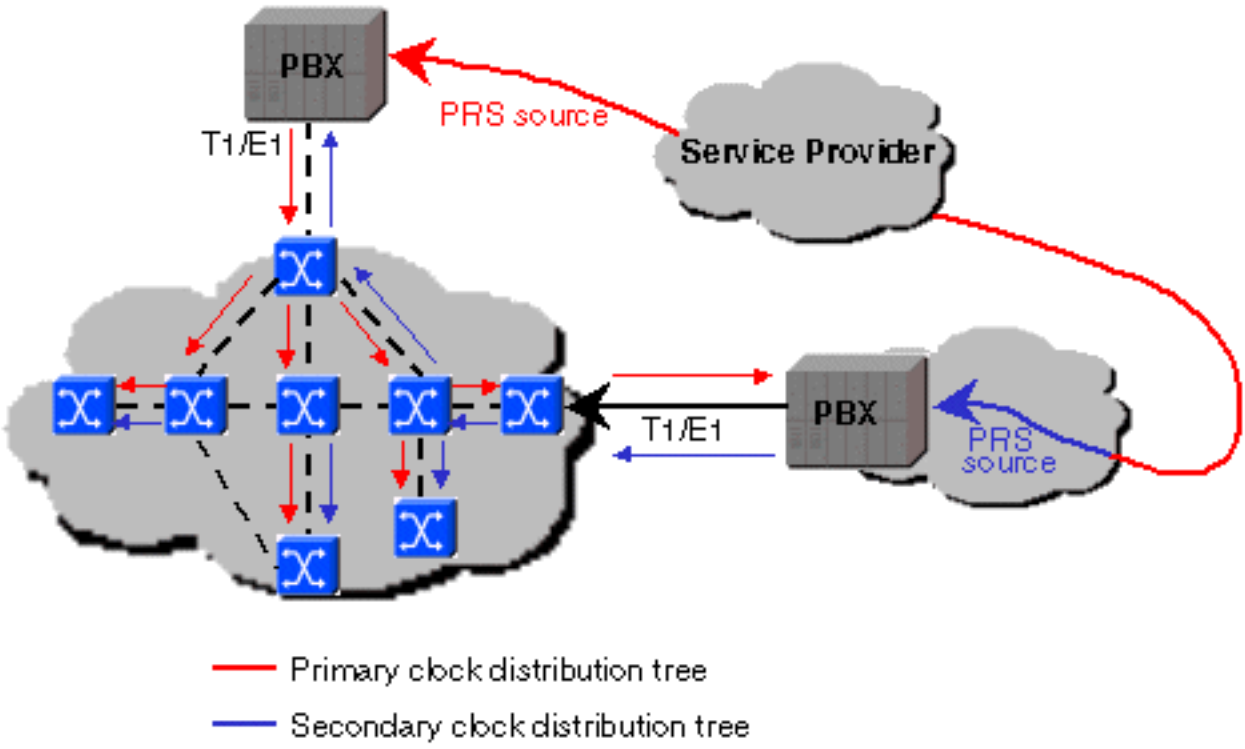

Oscillators on the 8510-MSR and the PA-A2 on the Cisco 7200 can provide stratum 4 clock. The Catalyst 8540-MSR with the optional network clock module can provide stratum 3 clock source. Without the optional network clock module, the Catalyst 8540-MSR provides stratum 4 clock. If the Catalyst 8540-MSR is equipped with the optional network clock module, the T1/E1 Building Integrated Timing Supply (BITS) port can also be used as a clock source.

Once you decide how the clock distribution tree will look for the entire network, you need to implement it on each device, including Cisco ATM switches (that is, internal clock distribution within the ATM switch needs to be configured). Internal clock distribution on Cisco Enterprise ATM switches and routers can be configured using these two commands: **ces dsx1 clock source** and **network-clock-select**.

Use the **network-clock-select** command to specify which clock source (interface or internal oscillator) to use as the system clock on the ATM switch. On Cisco products supporting CES, you can specify multiple network clock sources and their priority for redundancy purposes. If nothing is configured, the 8510-MSR and Catalyst 8540-MSR use the local oscillator on the ATM switch processor (ASP) or route processor (RP) as the system clock by default. All interfaces that are configured to use network-derived clock use the clock source specified in the **network-clock-select** statement as a transmit clock on that interface. All ATM and CBR interfaces on the 8510-MSR and Catalyst 8540-MSR are configured to be network-derived by default. So too are the ATM and CBR interfaces on the PA-A2 port adapter. The **ces dsx1 clock source** statement specifies for each individual interface which clock source to use as a transmit clock on that interface. The following options are available:

- Network-derived: As previously mentioned, if the interface is configured to be network-derived, the clock source specified by the **network-clock-select** statement is used as the transmit clock on that interface (that is, the transmit clock is derived from the source provided by the ATM switch internal clock distribution mechanism). Use the **show network-clock** command to find out which clock source is being used. Network-derived is the default setting on all Cisco ATM switch interfaces.
- Loop-timed: Transmit clock on the interface is derived from the clock source received on the

same interface. This mode can be used when connecting to a device with a very accurate clock source.
- Free-running: Transmit clock on the interface is derived from the port adapter's local oscillator, if one exists. If the port adapter does not have a local oscillator, the oscillator from the processor board is used. In this mode, the transmit clock is not synchronized with any receive clocks in the system. This mode should be used only if synchronization is not required, like some LAN environments.

Configuring CES

Before You Configure

Before implementing and configuring CES, you should make the following decisions based on the information discussed in this document so far:

1. Which type of service do you need (unstructured or structured)?
2. Which clocking mode will you use (synchronous, SRTS, or adaptive)?
3. If you decide to use synchronous or SRTS clocking mode, which device in your network will provide clock source to the rest of the network? Do you have devices equipped with PLLs? Do you plan to derive clock from interfaces that do not support it? Refer to Clocking Requirements for the LightStream 1010, Catalyst 8510-MSR and Catalyst 8540-MSR for more information.
4. How do you plan to distribute clock source throughout the network so that you have a loop-free clock tree while still preserving the clock characteristics of the PRS as much as possible?
5. Determine the T1/E1 characteristics (such as linecode and framing) defined at the CPE or line provided by the service provider.
6. Determine the distance between the CES PAM and the closest device that regenerates the T1/E1 signal (this can be CPE or CSU/DSU, for example). If the distance is greater than 110 feet, you need to change the lbo configuration on the CES PAM.

Sample Configurations

Here are some examples of configurations with:

- T1 Unstructured CES using Synchronous Clocking and PVCs
- T1 Unstructured CES with SRTS Clocking and Soft PVCs
- T1 Unstructured CES with Adaptive Clocking and Soft PVCs

See also Configuring Circuit Emulation Services.

Verifying Configurations

You can use the **show** commands explained below to verify the configuration. The output from these **show** commands from all involved devices is also helpful to Cisco Technical Assistance Center (TAC) engineers if you have to open a case.

Comm	Description

and	
show version	Displays the current version of Cisco IOS. You'll need to know the IOS version when verifying supported features or searching for bugs on CCO.
show run	Displays the current running configuration.
show int cbr *x/y/z*	Displays interface status.
show ces int cbr *x/y/z*	Displays line state and all T1/E1 error counters (the definition of all counters are in RFC 1406 (). It also shows port and service configuration. Make sure that the line-code and framing configured on the switch is the same as configured on the CPE device.
show ces circuit int cbr *x/y/z n*	Where ***n*** is the channel ID (0 = unstructured; 1-24 = structured). Displays information about underflows and overflows. **Note:** There will always be some underflows/overflows as a circuit is coming up, so be sure to look at the relative increase and not the absolute number. Underflows and overflows indicate clocking slips.
show ces address	Displays the address and the VPI/VCI pair to be used if you want to terminate the soft PVC on this CBR port. You must first configure the CES circuit to view this information. If you have structured service with multiple channels, there will be multiple addresses and VPI/VCI pairs.
show ces stat	Displays the status of all circuits.
show network-clock	Displays the configuration of network clock source preferences and indicates if the active clock source is indeed the one that is configured to be preferred.
show log	Displays any past clock switch-over events or interface events. To benefit from the log, you should configure timestamps on your switch and enable logging. You can configure this in the global configuration mode using the following commands: • **logging buffered** • **service timestamps log date msec** • **service timestamps debug date msec**

Basic Troubleshooting

Some of the most common problems encountered with CES are listed below, along with troubleshooting tips.

The Circuit is not Coming up or the CPE is in Alarm

1. Make sure that you are using the correct cable. Refer to PA-A2 ATM CES Cables, Connectors, and Pinouts for pinouts of all CES ports for the PA-A2.
2. Make sure that the framing and line-code are the same on the CPE and the switch. Use the **show ces interface *x/y/z*** command to see how the switch is configured. To change the framing and line-code, use the **ces dsx1 framing** and **ces dsx1 linecode** commands.
3. Make sure that all hardware is in working condition, such as the port on the CPE, and the cable and port on the switch. You can troubleshoot hardware problems by replacing one component at a time or by using loopbacks to localize the problem. You can use user-configurable loopbacks to do this by using the **ces dsx1 loopback** command for CBR interfaces and the **loopback** command for ATM interfaces. It may be necessary to make an external loopback plug on the CBR T1 interface or to externally loop the transmit cable to the receive cable on the ATM interface. Loopback tests are useful in general when troubleshooting CES problems.
4. Check the alarm indicators:A red alarm indicates failure on a local device.A yellow alarm indicates remote-end failure.A blue alarm is declared when all one pattern is detected (AIS). CPE equipment connected to the port in blue alarm should see this condition as loss of signal (LOS). A blue alarm often indicates that there is a problem in the ATM network and/or the connection has possibly gone down.On the 8510-MSR, LEDs indicate different alarms.
5. Measure the distance between the CPE (or the closest signal regenerating device, such as CSU/DSU) and the CBR port on the CES PAM. The default line build out is 0 -110 feet. If your distance is longer, use the **ces dsx1 lbo** command to increase the default value. The maximum supported distance is about 700 feet.

The Circuit is Experiencing Clocking Slips

To determine if there are clocking slips on a circuit, check for underflows and overflow using the **show ces circuit interface cbr *x/y/z n*** command, where *n* is circuit ID (always 0 for unstructured CES).

As AAL1 cells are received on an ATM interface, they are stored in the SAR buffer, which resides on the CES PAM. Then, the *framer* will take the AAL1 data from that buffer, strip all the headers out, form a T1 frame, and transmit it on the CBR interface. The size of this buffer is implementation-dependant and is chosen to accommodate specific end-to-end maximum CDV while avoiding excessive delay. If there is a slight clocking difference between the device doing segmentation (conversion from T1 frames to ATM cells) and the device doing re-assembly (conversion from ATM cells to T1 frames), the SAR buffer will get either underflows or overflows.

- *Overflows:* The segmentation side is faster than the re-assembly side, resulting in dropped frames.
- *Underflows:* The segmentation side is slower than the re-assembly side, resulting in repeated frames.

PBX Reporting Framing Errors or Carrier Drops

Check all ATM links for cyclic redundancy check (CRC) or other errors. Use the **show controller atm** and **show interface** commands.

Users Hear Static or Clicking on Phone Calls

Check the clocking of all ATM and CES devices. Try adaptive clocking and see if the problem ceases.

You Suspect a Bad Reference Clock

1. Reference clock can be degraded if the original clock source provided by the service provider has problems, if the ATM network degrades the clock, or if clock distribution throughout the network is mis-configured.
2. Try adaptive clocking. If that solves the problem (while SRTS and synchronous were experiencing the problem), you can conclude that your suspicion was accurate.

There are Clocking Problems in a Network with PA-A2

The ATM interface on the PA-A2 also uses network-derived clocking by default on the ATM uplink port. By default, the clock source is atm clock internal, which is the equivalent of network-derived. By network-derived, we mean that we use the highest-priority active clock source, as displayed in the output of the **show network-clock** command.

Use the **no atm clock internal** command to set the transmit clock to the line. This configuration is equivalent to a loop-timed transmit clock source, in which the transmit clock source is derived from the clock source received on the same interface.

Related Information

- **Asynchronous Transfer Mode (ATM)**
- **Technical Support & Documentation - Cisco Systems**