



System Configuration Guidelines

This chapter describes the issues to consider when putting together a new system or upgrading an existing system.

Ethernet Connections

A dedicated Ethernet connection for your host-to-switch connection and for SS7-to-switch installations eliminates contention delays imposed by other traffic. Sites where the connection is shared with other business traffic may function adequately until either the host-to-switch traffic or the other traffic increases to the point where delays are incurred. The number of open sockets affects the available call processing time on the VCO's CPU.



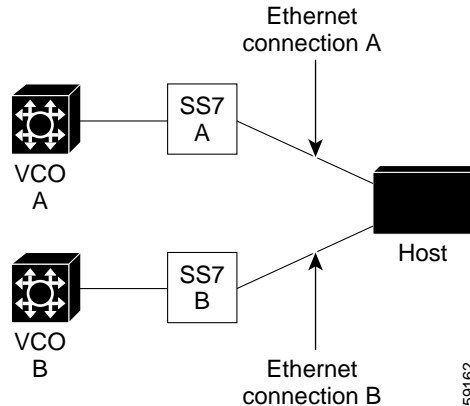
Note

Multiple VCO hosts may not necessarily mean multiple physical hosts. For example, in an SS7 configuration the circuit interworking box acts as a gateway to the VCO system and supports up to eight physical host connections with the VCO.

Ethernet Connections and Redundancy in SS7 Systems

Cisco cannot guarantee that systems configured with a single Ethernet connection between the host and the SS7/VCO pairs (see Figure 6-1) will successfully do a switchover from the active to the standby sides if only Ethernet connection A (or B) is installed. If the feature flag to allow switchover upon host link failure is set, switchover does not occur unless both Ethernet connections A and B are installed.

Figure 6-1 Ethernet Connections and SS7/VCO Redundancy Switchover



SS7 Stacks for Country Variants

The following country SS7 stacks variants are available:

- ANSI—Used in North America.
- ITU—Used outside North America.
- China TUP—Used in China.



Note

Contact your Cisco sales representative for the list of available country variants.

ARP Cache Entries



Note

The following information applies only to system software prior to V5.1(2).

The VCO does not correctly respond to an Address Resolution Protocol (ARP) message requesting a Media Access Control (MAC) address. Instead of sending a reply to the requesting device, it responds with a broadcast message.

For example, host A sends a broadcast ARP message to a VCO switch with IP address x.x.x.x requesting the MAC address of the switch. The switch responds with a broadcast ARP reply. Because host A did not receive the needed MAC address, it must always send a broadcast ARP request whenever it wants to communicate with the switch. Repeated broadcast messages cause congestion on the Ethernet network.

The workaround (prior to V5.1(2)) is to statically define the VCO's MAC address (both sides, if redundant) in the devices that are likely to send an ARP message for the VCO's MAC address. This includes hosts, routers, SS7 stations, VRUs, and so forth.



Note

Refer to the product manual of each device for instructions on how to statically define (force) a MAC address into its ARP cache.

TeleRouter as Host Backup

The TeleRouter functionality is built into the VCO/4K as a virtual host. It can be configured to do automatic call handling and routing without host intervention. It can also be configured to do backup call handling. This means that the TeleRouter virtual host can take over in the event of loss of communication between the VCO and the host. The following are given as some of the configuration changes needed to use the TeleRouter as a host backup:

- The host application must be designed to work with TeleRouter
- Use impulse rules to accept the incoming call
- Use REP EACH to capture all digit fields
- Use WAIT TIME to pause the call
- Set up the host to abort the impulse rule and send a message to route the call under normal circumstances
- If the host does not abort the rule, the TeleRouter ROUTE token takes over after the WAIT TIME expires

Another example could use the Host setup timer and a NOHOST token in the impulse rule.



Note

This backup mechanism must be carefully tested before being placed in service. This mechanism does not work with SS7 systems.

For more details on using the TeleRouter, refer to the *Cisco VCO/4K TeleRouter Reference Guide*.

Calculating Resource Needs

When you are designing a new system to meet specific needs, determine the hardware requirements of the switch. Alternately, you may want to determine if a defined system will satisfy the needs of a specific site. You can work from a base of the available time slots that are a function of physical system constraints. From that base, and a given card population, you can determine what can be allocated.

For either approach, you must also consider what effect the host application might have.



Note

This section provides guidelines for resource calculation. Final system design requires input and guidance from a Cisco Systems Engineer (SE).

Calculating Loading

Use the following steps to determine the resources required to support a given call loading:

-
- Step 1** Determine the call loading on the switch, based on the following:
- Incoming traffic load
 - Time required to complete the transaction; for DTMF ports, the digit collection time
 - Grade of service—usually P.01 (1% blocking)

Traffic load can be based on the following:

- Busy hour call (BHC), if known.
- Million minutes per month. Calculate the calls per day and then use the percentage of calls in the busy hour. With this information you can calculate backward to obtain BHC rate.

In addition to BHC or the calculated rate (above), you need the average call hold-time.

Assume the calls arrive uniformly for the hour. (This has been proven reliable in actual practice.) Simultaneous seizures are usually not an issue in this calculation.

Step 2 Determine the call flow.

- Calculate the following, based on the call rate and expressed in erlangs. Erlangs are based on a statistical distribution of call arrivals and average call times. Erlang-B is best for voice traffic:
 - How many network ports are needed?
 - How many connections are needed for VRU and other application-specific peripherals?
 - How many ports are needed for resources (DTMF, for example)?

The above calculations result in a total port count.

- From the port count, calculate the number of E1/T1 spans needed.

The difficulty in assigning a number to the resource port needs is because it is heavily call model dependent. You should work with the application developer to define the needed resources—where, when, and for how long. This may be mostly DTMFs used to tie to external IVR resources.

- From the port count, calculate the number of SPC/SRMs needed.

Step 3 Determine signaling requirements.

- Note the following:
 - With SS7 or PRI there is no in-band signaling.
 - If in-band signaling is used, it is usually MF or MFCR2.
 - With MF you need a port for a short time—6 seconds or so.
- Perform the same erlang analysis as for the voice traffic (Step 2 a.), but for a time of 4 to 8 seconds (typical) rather than the average call time.

Note the following:

- MFCR2 is an international requirement and is the same as MF. The only exception is that it is two-way. Use 10 to 12 seconds in your calculations.
- CPA usage is typically 20 seconds or so. Generally, the CPA determines a busy condition or an answer before it gives back the port.
- A special (and relatively unusual) situation occurs when the interface is not to a switch or is to a connection that does not have answer supervision. This requires that the CPA listen for voice (actual answer by a person) or a hang up. This could result in the CPA being connected for the entire call. This occurs occasionally on international calls. The erlang analysis for this component is therefore based on the entire call time.

Step 4 Determine conferencing requirements.

This involves passing ports to external IVRs for messages or other interactive functions. The IVR passes the port back, but the IVR requires the port for a while. This means that there is a three-way call during this time (in/out/IVR). Typically this is for 20 seconds or so. By the very nature of conferencing, you must assume that three ports are in use for a portion of the call.

Step 5 Determine IPRC usage requirements.

IPRC prompts need a port for about 6 seconds. On rare occasions, such as when a user supplies music, the time requirement can extend over a much longer time.

Step 6 Calculate resource requirements.

As an example:

- a. Calculate MF ports when no call loading is given, by doing the following:

Convert ports to erlangs. This is the total number of incoming network ports. With a 1% grade of service, and a trial and error approach, consult erlang tables to identify the erlang value that represents the number of ports you have. (Start with a number 3% to 6% less than the number of ports.)

Unless you have been given the Average Call Hold Time (ACHT) for the application, assume 3 minutes.

- b. Use the following equation to convert erlangs to BHCs:

$$\text{BHC} = \text{erlangs} \times 3600 / \text{ACHT (in seconds)}$$

With this incoming BHC you can calculate the MF ports required. For MF ports you can assume that the average hold time is 6 to 10 seconds to deal with the ANI and DNS. The port is normally released back to the pool after this time.

- c. Use the following equation to calculate the number of MF erlangs:

$$\text{MF erlangs} = \text{BHC} / 3600 / \text{MF AHT (in seconds)}$$

- d. Use the erlang tables to convert erlangs to a number of MF ports.



Note

Erlang B values can be calculated using formulas found in various telecommunications sources. Tables of calculated values also exist that can help you determine loading for a specific customer.

Other Considerations

The longer the average call, the less that is required of the CPU. Conversely, shorter calls tend to place more load on the CPU. Longer calls are very likely to require more system physical resources.

When the system is heavily loaded, the system CPU is taxed to the point where the system console may respond slowly.

Port/Time Slot Allocation

Port/time slot allocation does not affect system performance. However, an understanding of the allocation algorithm is useful when you perform system configuration. You can achieve more efficient port assignments if you know how the system is designed to allocate physical resources. Allocation varies by system type and by span type.

Time Slot Availability

The fundamental difference between 2K and 4K systems is time slot availability as a result of the C-Bus. There are also a small number of additional time slots used by the 2K system that are freed up to become a resource in 4K systems.

To achieve maximum efficiency when a system is heavily loaded, it is useful to know how allocation actually happens. Within a span, time slots are contiguous, but the spans/service do not necessarily have to be adjacent to other spans on the same card.

The 0 to 7 bit section at the beginning of the A-bus is unavailable for either 2K or 4K systems. It is a VCO/4K system software limitation and results, in a 4K system, in the total time slot availability actually being 4087 rather than the binary 4K of 4096.

Time Slot Availability in a Fully Licensed VCO/4K

Use the following guidelines when you configure time slots and logical interfaces in your fully licensed VCO/4K system:

- Listen and transmit time slots—4088.
- Listen-only time slots—4000 (these are separate from the listen and transmit time slots).
- Logical number of time slots—240. A logical number is defined as one interface or card number. For example, a fully loaded ICC constitutes 16 logical time slots. An SPC can have 32 logical time slots. It is not possible to have more than 240 logical time slots.



Note

If you exceed these maximum numbers as you configure your system, error messages are generated in the system log file. Refer to *Cisco VCO/4K System Messages* for more information.

Span and Slot Availability—C-Bus Disabled

Figure 6-2 shows the available time slots for a 2K system (C-bus disabled).

Figure 6-2 Time Slot Availability by Bus—C-Bus Disabled

C-Bus	Not available			
B-Bus	8	120	128	48
	HW	Available	DTG	HW
	720 Available			
A-Bus	8	968		48
	API/HW	Available		HW

59163

Table 6-1 lists the number of spans supported.

Table 6-1 Spans Supported—No C-Bus

Bus / Total Spans	T1 Spans	E1 Spans
C-bus	0	0
B-bus	35	25

Table 6-1 Spans Supported—No C-Bus (continued)

Bus / Total Spans	T1 Spans	E1 Spans
A-bus	40	30
Total spans supported	75	55

See Table 5-2 on page 5-2 for a list of the number of time slots available.

Span and Slot Availability—C-Bus Enabled

In a 4K system with a C-bus, all the space above 2K can be used. Figure 6-3 shows the available time slots for a 4K system.

Figure 6-3 Time Slot Availability by Bus—C-Bus Enabled

Table 6-2 shows the number of time slots available.

Table 6-2 Time Slots Available—With C-Bus

Time Slots	T1 Spans	E1 Spans
Total	4096	
Reserved	8	
DTG	128	
Available for use	3960	
Used by network interfaces	3936	
Available for services	24	

See Table 5-3 on page 5-3 for a list of the number of spans supported.

Time Slot Considerations for Older Cards in a VCO/4K System

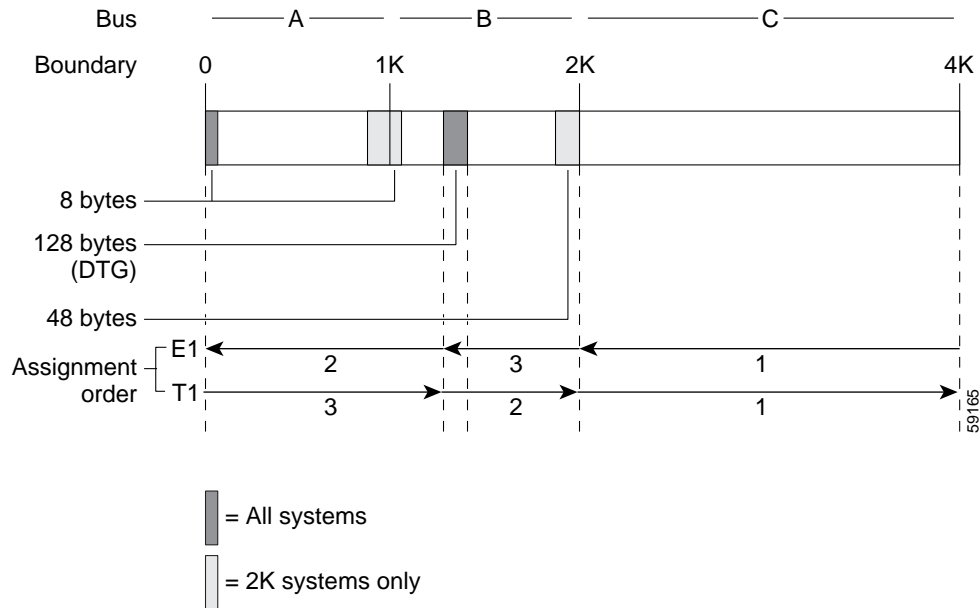
Older cards (such as IPRC, D&I, LTC-8) installed in VCO/4K systems must obey the following rules:

- An IPRC-64 (for example) requires 64 consecutive time slots on the A- or B-bus (The C-bus cannot be used, the port addresses must be less than 08 00).
- The time slots must not straddle bus boundaries (the A- to B-bus boundary is 03 FF to 04 00).
- The unavailable ports on a VCO/4K system are: 00 00 to 00 07 and 04 80 to 04 FF.

Time Slot Mapping by Bus Type

The manner in which time slots are assigned depends on the nature of the span (E1 or T1). E1 spans are assigned in multiples of 32 and T1 spans in multiples of 24. The order in which the three possible segments (from 0 up: 1st K, 2nd K, or 3rd 2K) is also a function of span type (see Figure 6-4).

Figure 6-4 Time Slot Mapping



When allocating E1 spans to a 4K system, the allocation order proceeds from the top down (in general): 4095 to 2048, 1151 to 8, and then 2047 to 1152 (no DTG). If you are adding E1 spans, the first span would go in 4064 to 4095, then 4032 to 4063, and so on; then 1120 to 1151, 1088 to 1119, and so on; and then 2016 to 2047, and so on.

When allocating E1 spans to a 2K system, the allocation order also proceeds from the top down but has to skip the holes created by hardware limits (see Figure 6-2, Figure 6-4, and Table 5-2 on page 5-2).

When allocating T1 spans to a 4K system, the allocation order proceeds from the bottom up: 2048 to 4097, 1280 to 2047, 8 to 1279 (no DTG). The first span would be 2048 to 2071, and so on.

When allocating T1 spans to a 2K system, the allocation order also proceeds from the bottom up but has to skip the holes created by hardware limits (see Figure 6-3, Figure 6-4, and Table 5-2 on page 5-2 and Spans Supported—With C-Bus, page 5-3).

Time Slot Assignment

Table 6-3 describes how several time slot assignment parameters are handled by 2K and 4K systems.

Table 6-3 Time Slot Assignment Parameters by Bus Type

Parameter	2K (Non-C-bus)	4K (C-bus)
When assigned to a card	When added to the database	When added to the database

Table 6-3 Time Slot Assignment Parameters by Bus Type (continued)

Parameter	2K (Non-C-bus)	4K (C-bus)
Assignment by card port capacity	Full capacity (contiguous)	Partial capacity (noncontiguous)
Assignment of unused slots	Time slots must be contiguous All time slots assigned to a card must be on the same bus	Unused slots flexibly assignable (span-by-span) to any Type 2 card ¹

1. The network type 2 card is the ICC card.

VCO/4K Companding Algorithms

A VCO/4K is configured to use only one companding law internally. The following terms are synonymous and refer to the companding law:

- Backplane law
- DTG/DTG2 law
- Internal law
- System law

Backplane law was previously the most popular choice, system law has become the terminology of choice with VCO/4K. All tones are generated (for example, DTG) and interpreted (for example, SPC and CPA) under whichever law (A or mu) is set for the system law. In other words, all service cards in a given switch operate under the law that is set as the system law.



Note

The DTG law for each DTG/DTG2 card is specified in the *Cisco VCO/4K Tone Plan Release Notes*.

For older cards (for example, PRI/N, CPA, and 4xT1), the system law is set individually on each card by a hardware jumper. New cards (for example, SPC and ICC), get the system law via software from the system-law feature flag from within the System Features screen. This feature flag is Set system to A-law, where Y sets the system law to A and N sets it to mu. All of these settings (jumpers on the cards and the feature flag) must match, either all mu or all A.

For connections to the network, the companding law is set via software to match the particular network connection. For older cards, it is set for each span, in the upper section of the span configuration screen. For the ICC, the law is set for each port/channel in the lower portion of the configuration screen, so it can be set differently for different ports in the same span. (An exception is an ICC span configured for ISDN. Then, as with the older cards, the law setting is done for all channels on the span at once.)



Note

Companding law is specified from the Database Administration menu. See the *Cisco VCO/4K Tone Plan Release Notes* for more information.

If the law setting for a network connection matches the system law, no conversion is done. If the law setting doesn't match, then the PCM stream is converted before further processing is done by the switch. The conversion is done on the network card.

In the worst case, if for example the system law is set to A, and a call comes in via an ICC configured as mu and goes out on a PRI/N configured as mu (but jumpered to A), then the PCM stream will be converted twice, mu-to-A by the ICC then A-to-mu by the PRI/N, even if there is no law-dependent internal processing (such as tone generation) done on the call.

The VCO can process A law and mu law calls simultaneously because they are all converted to whatever the system law is.

The ICC has a third choice, SYS, for port law settings. You should not do any conversion—assume the network connection uses the same law as is configured for the system. The SYS choice is typically used for data channels such as an SS7 stream.

**Note**

The older cards do not have the SYS choice in the configuration screen. For any spans carrying data, you must ensure that you set the span law the same as the backplane law.

You cannot set the law for ISDN D-channels and E1 framing channels; the network card knows not to convert those.

Association Between Tones and the Companding Law

The DTG/DTG2, which does only tone generation, has the encoded waveforms stored directly as tables of bits in the firmware. For each tone plan, you have to install a separate set of DTG firmware, with its own set of proms. The bit pattern in each table is the A- or mu-law encoded form of the particular waveforms, as determined by the tone plan's requirements.

The SPC has the A- or mu-law conversion embedded in the DSP software. The basic DSP algorithms do their work internally with linear 16-bit representations; the outgoing stream is then compressed into 8 bits (in the case of tone generation), or the incoming stream is first expanded from 8 bits (in the case of tone interpretation), by a layer of DSP software that performs the compression/expansion for either A or mu law, whichever is set as the system law.