Cisco Systems

Deploying Cisco Voice over IP Solutions

Learn real-world voice-over-IP deployment solutions and strategies from the Cisco experts

Edited by **Jonathan Davidson**, CCIE® No. 2560
and **Tina Fox**

ciscopress.com

Deploying Cisco Voice over IP Solutions

Edited by: Jonathan Davidson and Tina Fox

Contributing Authors:

Phil Bailey
Rommel Bajamundi
Wayne Cheung
Thu Dao
Jonathan Davidson
Tina Fox
Sachin Gupta
Christina Hattingh
Ted Huff
Stephen Liu
Curt Mah
Greg Mercurio
Jeremy Pollock
Jim Rushton
Ravindar Shankar

Cisco Press

Cisco Press
800 East 96th Street
Indianapolis, IN 46240 USA

This makes it possible to build hybrid topologies where RSVP is used around the edges of the network to protect slower WAN access links from oversubscription, while the high-speed campus and WAN backbone links do not use RSVP. Of course, this topology compromises the true end-to-end reservation and guaranteed QoS promise of RSVP, but it might be a workable compromise. The backbone links can receive a measure of protection from over-engineering or from one of the other CAC mechanisms discussed earlier, while the highest contention links (typically the WAN edge) can make use of RSVP.

Figure 4-23 shows a hypothetical network that is configured for DiffServ in the backbone and campus, but uses RSVP reservations across the WAN edge links.

Figure 4-23 *Hybrid DiffServ/RSVP network topology.*

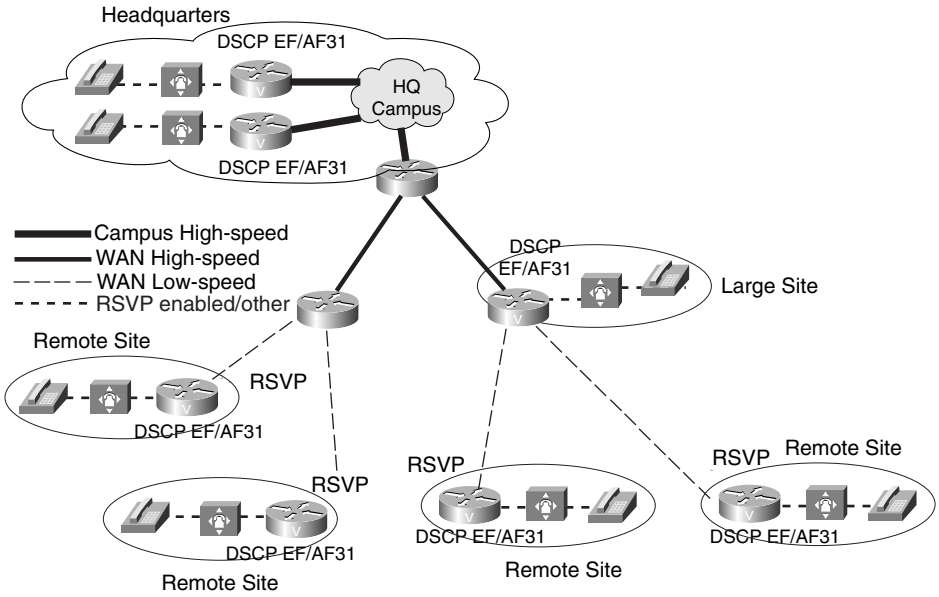

RSVP CAC Summary

Keep these factors in mind regarding the use of RSVP as a CAC mechanism.

In current Cisco IOS Software, H.323 calls are initiated by default using FastConnect when RSVP is configured:

- RSVP packets (PATH and RESV) travel as Best Effort traffic.
- WFQ must be enabled on an interface/PVC as a basis for LLQ.

RSVP is a true end-to-end CAC mechanism only if configured on every interface that a call traverses.

For the unique ability to serve as an end-to-end CAC mechanism, and guarantee the QoS for the entire duration of the call, RSVP does incur some "costs" on the network:

- Signaling (messaging and processing).
- Per flow state (memory).
- Post-dial delays.
- RSVP doesn't provide for call redirection after call setup if a link in the network should fail.
- RSVP is not yet supported on the Cisco IP phones.

Table 4-18 evaluates the RSVP mechanism against the CAC evaluation criteria described earlier in this chapter.

Table 4-18 *Summary of RSVP*

	Evaluation Criteria	Value
1	VoX supported	VoIP/H.323 only
2	Trunking/IP telephony	Currently trunking only
3	Platform/Release	Cisco IOS gateways Cisco IOS Releases 12.1(5)T and 12.2
4	PBX trunk types supported	All
5	End-to-end/Local/IP cloud	End to end between outgoing gateway and terminating gatekeeper (provided all intermediate nodes are RSVP configured) Could be used at WAN edge with DiffServ backbone
6	Per call/ interface/endpoint	Per call
7	Topology awareness	Yes
8	Guarantees QoS for duration of call	Yes
9	Post-dial delay	Yes
10	Messaging network overhead	PATH/RESV and periodic keepalives

Feature Combinations, Interactions, and Sequencing

Although some overlap exists between the functionality they provide, several of these solve different aspects of the CAC problem and make sense to use together in a network design. Two questions often arise:

1. Can two CAC methods be used together on the same gateway at the same time for the same calls?
2. If the answer to the preceding question is yes, in what sequence is the CAC decision reached?

Figure 4-24 summarizes the sequencing of CAC features that can be active on an outgoing gateway, based on Cisco IOS Releases 12.1(5)T and 12.2. As features and software releases change, and as bugs are fixed, this information might change without notice. As you can see from the flow diagram in Figure 4-24, the only features that are mutually exclusive are RSVP and PSTN Fallback.

Figure 4-24 *Sequence of CAC feature utilization on an outgoing gateway.*

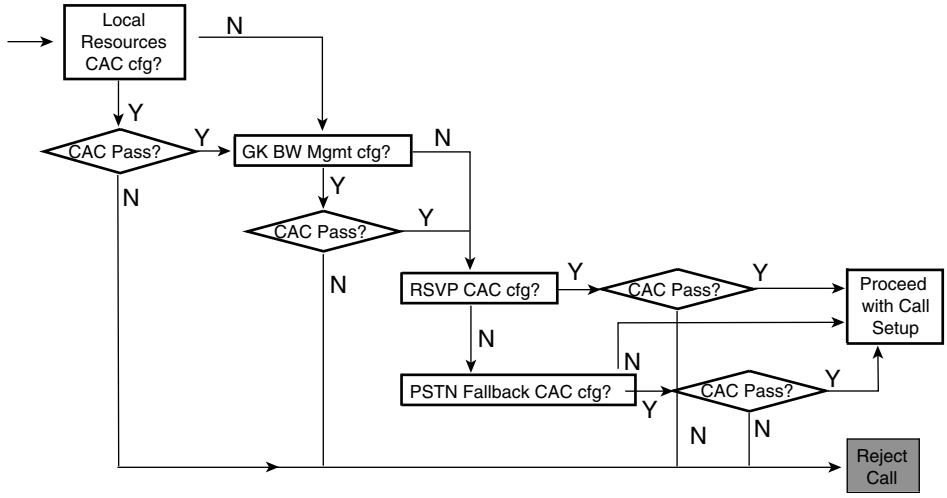

When Should I Use Which CAC Mechanism?

With a plethora of CAC mechanisms available, the immediate design question is, "When should I use which CAC feature?" As has been pointed out during the discussions of the individual features, and through the comparisons and summaries that have been drawn throughout the text, the various features often do different things and solve different aspects of a CAC problem. Some of these aspects can be more important design criteria for your network than others. Thus, there is no single recipe prescribing exactly when to use which mechanism. Like all other software features, you have to make the decision while considering your network design goals.

This section attempts to provide some guidance concerning design criteria that might exist for your network, and if so, which features might fit the solution. Before proceeding, it should be noted that the first feature selection criteria that should be used are the Evaluation Criteria listed at the end of each feature section. For example, if a SIP-based VoIP network is being designed, there is no point in considering an H.323 CAC feature. Provided you have already accomplished that level of screening, use the suggestions in this section to further narrow your choice of features.

CAC in Connection Trunk Networks

Unlike switched networks, where each call is set up individually across the packet network after a user dials, "connection trunk" networks consist of nailed-up connections across the packet network. The PBX might perceive that it makes each call individually, but the packet network has a permanent trunk in place (a point-to-point link—similar in concept to a leased line) that is always present, always ready, and always terminates to a fixed and predetermined destination. These nailed-up packet network configurations are typically used when some signaling is present between the PBXs that must pass transparently and unchanged through the packet network. The gateways cannot interpret the signaling; they merely tunnel it through the packet network.

The following are two major applications for this type of network:

- Networks in which signaling such as flash-hook and Message Waiting Indications (MWI) must be passed through the packet network to a PBX to be activated for OPX (Off Premise Extension) phones—phones that are separated by the packet network from the PBX from which they draw their features.
- Networks in which proprietary signaling is used between PBXs to enable private PBX networking features. (Examples include Lucent DCS, Siemens CorNet, NEC CCIS, and others.)

Cisco IOS gateway connection trunk configurations use the same basic tools (such as dial-peers) as switched networks to set up connections. The difference is that these "calls" are set up only once, when the gateway boots up or when the configuration is inserted, and remain in place indefinitely. If a link in the network should fail and bring the call down, the router will reestablish it at its earliest opportunity. Whether or not there is actually a real call active (with people talking) over this connection is transparent to the gateways. For this reason, the standard CAC mechanisms, in most cases, do not apply. Connection trunk configurations will not come up properly if there is not enough bandwidth for the connection, so once the configuration is in place, it is assumed that there is sufficient bandwidth available for the calls.

The following call-by-call CAC mechanisms apply only to switched networks and should not be used with connection trunk configurations:

- Max connections
- PSTN Fallback
- Gatekeeper bandwidth
- Gatekeeper RAI

Connection trunk configurations can, however, benefit from the PBX busyout CAC features. When something in the network is down and the nailed-up connections fail, or the

interfaces they use fail, it would certainly be useful to busyout the trunk to the PBX. These features include:

- LVBO
- AVBO
- Trunk conditioning

In concept, RSVP could be used to guarantee (reserve) bandwidth for the nailed-up calls in order to protect the voice quality from fluctuating network conditions. However, because connection trunk networks are fixed, point-to-point connections, the number of calls active across any network segment (from the router's perspective) is fixed and relatively easily designed by manually engineering the bandwidth and by using standard LLQ configurations to ensure bandwidth. The value-add that RSVP can offer here should be carefully considered.

Areas of the Network to Protect

CAC methods are most useful and most needed in switched networks where it is often impossible to predict exactly how many calls might want to use a particular network leg at a given point in time. Statistical methods for engineering voice networks have existed for decades; nevertheless, there is no mechanism by which to know exactly who will call whom across the network at any given time. Unless the topology of the network is very simple, it is possible that bandwidth, at some point in the network, might be oversubscribed by too many calls. In the PSTN, this condition results in reorder tone or an intercept message indicating "all circuits are busy."

When considering CAC methods to trigger a comparable "all circuits are busy" condition when a packet network is too congested to carry a call, the goals of the network design must be considered. All the aspects of CAC shown in Figure 4-25 exist in every network, but some attributes will almost always be more important to a particular customer than others. The aspects of the network that might need protection with CAC features have been divided into four areas, as shown in Figure 4-25.

Figure 4-25 *Division of areas of the network.*

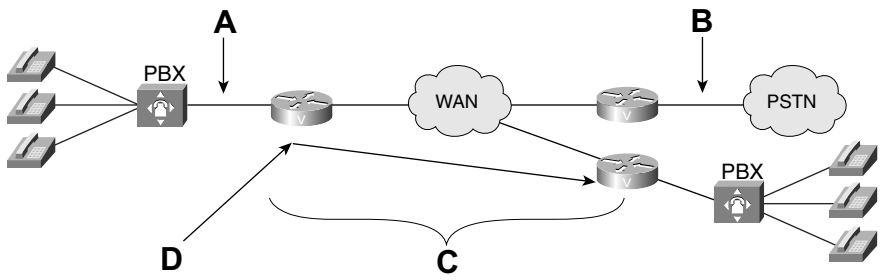

The area labeled A is the originating POTS connection. If it is important to keep the originating PBX from attempting to place a call onto the packet network when the network is incapable of completing the call, then the busyout CAC features should be considered. This might be important if hairpinning is an unacceptable call reject recovery method, or if the PBX/Key System does not have the ability to choose another route for a rejected or hairpinned call.

Area B is the terminating POTS side of the connection. If it is likely because of specific traffic patterns that the terminating POTS side is the part of network most susceptible to oversubscription, then gatekeeper RAI should be used. In enterprise networks, this is seldom of overarching importance, but in service provider networks, this is often an extremely important section of the network to protect.

Area C is the IP backbone part of the network. This is the most typical area of the packet network that enterprise customers (including Service Provider Managed Services networks) wish to protect their calls against, because this infrastructure is not dedicated to voice, but shared by many types of traffic. The CAC features protecting the network "cloud" include:

- PSTN Fallback
- Gatekeeper zone bandwidth
- RSVP

These CAC methods are all IP-based methods, which means that more CAC methods are available for VoIP networks than for VoFR and VoATM networks. VoIP also needs it more, because the Layer 2 technologies like FR and ATM cannot intrinsically protect against VoIP packet loss, as they can with VoFR and VoATM traffic.

Area D is a logical section of the network between sites. Regardless of the actual infrastructure connecting sites together, you might desire not to limit traffic within a site, or to limit it based on very different criteria than the traffic limitations between sites. For example, if the Headquarters location has the capability to handle 24 active calls at once, you might want to make sure that all 24 calls cannot be used by any one other site at any one time, but that there is a certain amount of capacity available to different remote sites so that the low-traffic sites don't get locked out by the high-traffic sites.

The CAC features you would use in this situation include max connections and gatekeeper zone bandwidth.

Network Topology Considerations

At a general level, two network topologies exist:

- Hub-and-spoke
- Multilayer hierarchical network with distribution layers