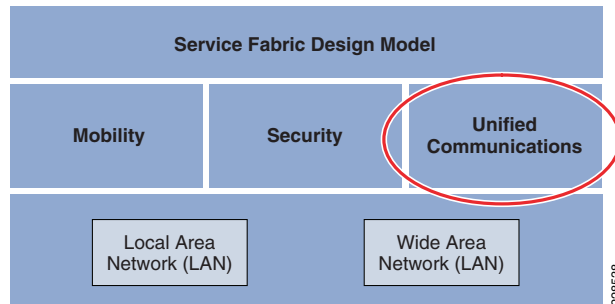


Unified Communications Design

As mentioned earlier in this document, the service fabric is made up of 4 distinct components: local and wide area network (LAN/WAN), security, mobility, and unified communications. This chapter focuses on the unified communications area within the service fabric. See [Figure 1](#).

Figure 1 Unified Communication Design



In today's environment, we observe community college and vocational education campuses moving towards a model that makes it convenient for the students to attend by not traveling to a central campus. This requires many smaller locations that still require the network facilities and voice communications features of the main campus. Instructors in this environment may teach classes at several locations within the same day or week. If these instructors had to contend with different phone systems at each location with no unified dial plan, the frustration level would be high and missed information would ultimately impact the customer, in this case the students. For this reason, it is important to have a unified communication architecture that allows uniformity in features and dialing as well as providing a high level of availability.

The SRA for CCVE solution is designed with unified communications considerations for the following:

- LAN design
- Call processing
- PSTN trunk sizing
- Dial plan
- High availability

LAN Design

Proper LAN design is critical to the performance of all IP telephony components. Telephony is expected to be always available. In order for this to be accomplished, a well designed LAN with considerations end-to-end QoS is required.

Layer 2 Access

The following are the access layer recommendations:

- Single IP subnet per virtual LAN (VLAN). Typically, a VLAN should not span multiple wiring closet switches; that is, a VLAN should have presence in one and only one access layer switch. Confining a VLAN to a single access layer switch serves to limit the size of the broadcast domain. There is the potential for large numbers of devices within a single VLAN or broadcast domain to generate large amounts of broadcast traffic periodically, which can be problematic.
- Typical access layer switches Power-over-Ethernet (PoE) capable include the stackable Cisco Catalyst 2950, 3500XL, 3550, and 3750, as well as the Cisco 3560 and the larger, higher-density Catalyst 4000 and 6000 switches fixed configuration Cisco Catalyst 2960-E, 3560-E and 3750-E series switches. Cisco Catalyst modular switching platform like 4500-E or 6500-E Series switches can be deployed in large size access-layer network domain that requires higher port-density and control-plane redundancy for resilient network services to critical endpoints (i.e., Cisco TelePresence, fixed configuration Cisco Catalyst 2960-E, 3560-E and 3750-E Series switches. Cisco Catalyst modular switching platform like 4500-E or 6500-E Series switches can be deployed in large size access-layer network domain that requires higher port-density and control-plane redundancy for resilient network services to critical endpoints (i.e., Cisco TelePresence).

Cisco recommends enabling two VLANs at the access layer: a native VLAN for data traffic and a voice VLAN under Cisco IOS.

Separate voice and data VLANs are recommended for the following reasons:

- Address space conservation and voice device protection from external networks private addressing of phones on the voice or auxiliary VLAN ensures address conservation and ensures that phones are not accessible directly via public networks. PCs and servers are typically addressed with publicly routed subnet addresses; however, voice endpoints should be addressed using RFC 1918 private subnet addresses.
- QoS trust boundary extension to voice devices QoS trust boundaries can be extended to voice devices without extending these trust boundaries and, in turn, QoS features to PCs and other data devices.
- Protection from malicious network attacks VLAN access control, 802.1Q, and 802.1p tagging can provide protection for voice devices from malicious internal and external network attacks such as worms, denial-of-service (DoS) attacks, and attempts by data devices to gain access to priority queues via packet tagging.
- Ease of management and configuration. Separate VLANs for voice and data devices at the access layer provide ease of management and simplified QoS configuration.

To provide high-quality voice and to take advantage of the full voice feature set, access-layer switches should provide support for the following:

- 802.1Q trunking and 802.1p for proper treatment of Layer-2 CoS packet marking on ports with phones connected.

- Multiple egress queues to provide priority queuing of RTP voice packet streams.
- The ability to classify or reclassify traffic and establish a network trust boundary.
- Inline power capability (Although inline power capability is not mandatory, it is highly recommended for the access layer switches.)
- Layer 3 awareness and the ability to implement QoS access control lists. These features are required if you are using certain IP telephony endpoints, such as a PC running a softphone application, that cannot benefit from an extended trust boundary.)

Spanning Tree Protocol (STP)

To minimize convergence times and maximize fault tolerance at Layer 2, enable the following STP features:

- *PortFast*—Enable PortFast on all access ports. The phones, PCs, or servers connected to these ports do not forward bridge protocol data units (BPDUs) that could affect STP operation. PortFast ensures that the phone or PC, when connected to the port, is able to begin receiving and transmitting traffic immediately without having to wait for STP to converge.
- *Root guard or BPDU guard*—Enable root guard or BPDU guard on all access ports to prevent the introduction of a rogue switch that might attempt to become the Spanning Tree root, thereby causing STP re-convergence events and potentially interrupting network traffic flows. Ports that are set to **errdisable** state by BPDU guard must either be re-enabled manually or the switch must be configured to re-enable ports automatically from the errdisable state after a configured period of time.
- *UniDirectional Link Detection (UDLD)*—Enable this feature to reduce convergence and downtime on the network when link failures or misbehaviors occur, thus ensuring minimal interruption of network service. UDLD detects and takes out-of-service links where traffic is flowing in only one direction. This feature prevents defective links from being mistakenly considered as part of the network topology by the spanning tree and routing protocols.

Note With the introduction of RSTP 802.1w, features such as PortFast and UplinkFast are not required because these mechanisms are built into this standard. If RSTP has been enabled on the Catalyst switch, these commands are not necessary.

Routed Access

For designs requiring simplified configuration, common end-to-end troubleshooting tools, and the fastest convergence, a distribution block design using Layer-3 switching in the access layer (routed access) in combination with Layer-3 switching at the distribution layer provides the fastest restoration of voice and data traffic flows.

In both the traditional Layer 2 and the Layer 3 routed access designs, each access switch is configured with unique voice and data VLANs. In the Layer 3 design, the default gateway and root bridge for these VLANs is simply moved from the distribution switch to the access switch. Addressing for all end stations and for the default gateway remains the same. VLAN and specific port configurations remain unchanged on the access switch. Router interface configuration, access lists, “**ip helper**,” and any other configuration for each VLAN remain identical but are configured on the VLAN Switched Virtual Interface (SVI) defined on the access switch instead of on the distribution switches. There are

several notable configuration changes associated with the move of the Layer 3 interface down to the access switch. It is no longer necessary to configure an HSRP or GLBP virtual gateway address as the “router” interfaces because all the VLANs are now local. Similarly, with a single multicast router, for each VLAN it is not necessary to perform any of the traditional multicast tuning such as tuning PIM.

query intervals or ensuring that the designated router is synchronized with the active HSRP gateway.

The many potential advantages of using a Layer-3 access design include the following:

- Improved convergence
- Simplified multicast configuration
- Dynamic traffic load balancing
- Single control plane
- Single set of troubleshooting tools (for example, ping and traceroute)

Of these advantages, perhaps the most significant is the improvement in network convergence times possible when using a routed access design configured with EIGRP or OSPF as the routing protocol. Comparing the convergence times for an optimal Layer 2 access design (either with a spanning tree loop or without a loop) against that of the Layer 3 access design, you can obtain a four-fold improvement in convergence times, from 800 to 900 msec for the Layer 2 design to less than 200 msec for the Layer 3 access design.

Call Processing

When considering a design for unified communications, many factors come into play in making an overall call processing decision such as number of sites, number of users, functionality requirements (Call Center, IVR, etc.), network capabilities (WAN), etc.

If a Community College has one large campus that houses the primary data center and many small remote sites then a Centralized Call Processing model with Survivable Remote Site Telephony (SRST) might be the best choice. In this design, the remote sites depend on the primary campus Unified CM components for normal operation. If the WAN fails, the remote sites would resort to their local SRST router for primary call control until the WAN recovers. This model requires less administration since all call control is accomplished within one cluster.

If a very large Community College has multiple large campuses, it might require Distributed Call Processing. Within this model, you will find a Unified CM cluster at each large site for call control. Each cluster would communicate with the other clusters but administration is accomplished on a per-cluster basis. This design does not allow for some features that might be important in a Community College environment such as shared lines across clusters. Extension Mobility is not easily accomplished across clusters as well. In a campus environment Extension Mobility might be a very important feature for those instructors that might teach at multiple locations.

For this document, it is assumed that the main campus and large remote campus each have a data center and several hundred handsets. Instead of administering two clusters, clustering over the IP WAN will be implemented. Each of the smaller sites will then depend on this cluster for primary call control and SRST in case of WAN failure.

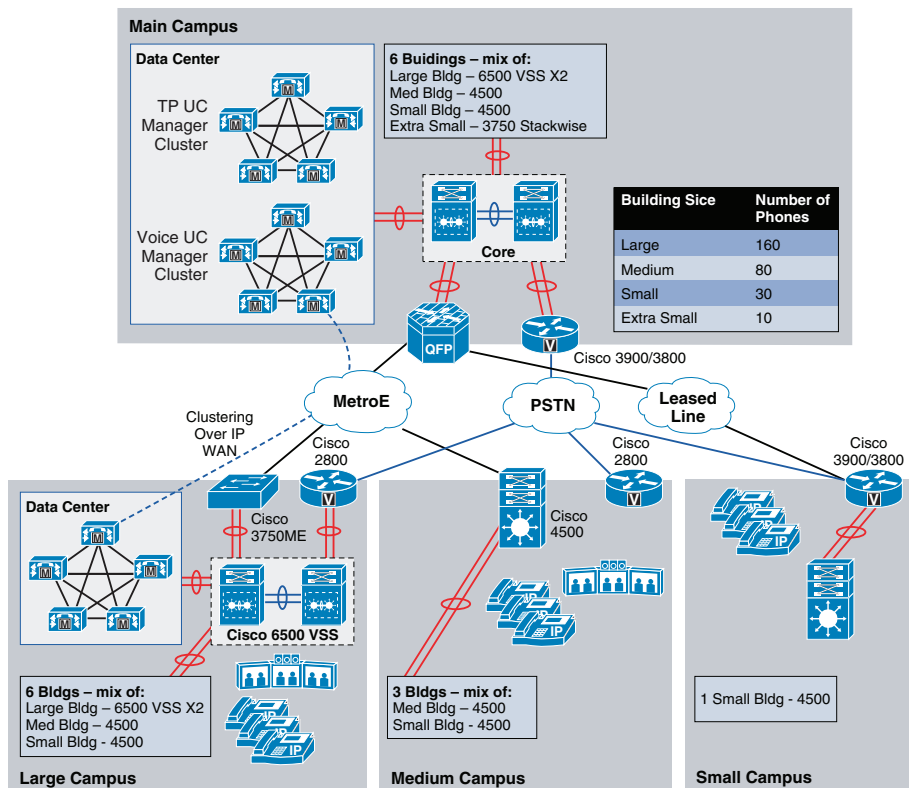
Follow these guidelines and best practices when implementing multisite centralized call processing deployments:

- Minimize delay between Unified CM and remote locations to reduce voice cut-through delays (also known as clipping).

- Configure locations (static or RSVP-enabled) in Unified CM to provide call admission control into and out of remote branches.
- The number of IP phones and line appearances supported in Survivable Remote Site Telephony (SRST) mode at each remote site depends on the branch router platform, the amount of memory installed, and the Cisco IOS release. SRST on a Cisco IOS gateway supports up to 720 phones, while Unified CME running in SRST mode supports 240 phones. (For the latest SRST or Unified CME platform and code specifications, refer to the SRST and Unified CME documentation available at <http://www.cisco.com>.) Generally speaking, however, the choice of whether to adopt a centralized call processing or distributed call processing approach for a given site depends on a number of factors such as the following:
 - IP WAN bandwidth or delay limitations
 - Criticality of the voice network
 - Feature set needs
 - Scalability
 - Ease of management
 - Cost

In [Figure 2](#), each location is indicated with number of phones anticipated.

Figure 2 Main Campus



For more information on the different call processing models, refer to the *Cisco Unified Communications SRND Based on Cisco Unified Communications Manager 7.x* at the following URL:

http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/srnd/7x/uc7_0.html

Clustering over the IP WAN

Clustering over the IP WAN Advantages

Clustering over the IP WAN best fits the needs of this Community College design for the following reasons:

- Single point of administration
- Feature transparency
- Shared line appearances (even across locations)
- Extension Mobility
- Unified Dial Plan

In the education vertical this model allows educators to move from campus to campus and retain their phone identity with Extension Mobility. It also allows for features such as shared lines that might not be possible in other deployment models. Clustering over the IP WAN does however have some very strict requirements in regard to bandwidth and latency.

Before providing details on clustering it might be good to have some background on what components exist in a cluster.

In every Communications Manager Cluster there is only one publisher but there can be several subscribers. The publisher server replicates a partial read-only copy of the master database to all other servers in the cluster. Most of the database modifications are done on the publisher. If changes such as administration updates are made in the publisher's master database during a period when another server in the cluster is unreachable, the publisher will replicate the updated database when communications are re-established. Database modifications for user-facing call processing features are made on the subscriber servers to which the IP phones are registered. These features include:

- Call Forward All (CFA)
- Message Waiting Indication (MWI)
- Privacy Enable/Disable
- Do Not Disturb (DND) Enable/Disable
- Extension Mobility Login (EM)
- Hunt Group Logout
- Device Mobility
- CTI Certificate Authority Proxy Function (CAPF) status for end users and application users
- Credential hacking and authentication

Each subscriber replicates these changes to every other server in the cluster. Any other configuration changes cannot be made on the database during the period when the publisher is unreachable or offline.

Most normal operations of the cluster, including the following, will *not* be affected during the period of publisher failure:

- Call processing
- Failover
- Registration of previously configured devices

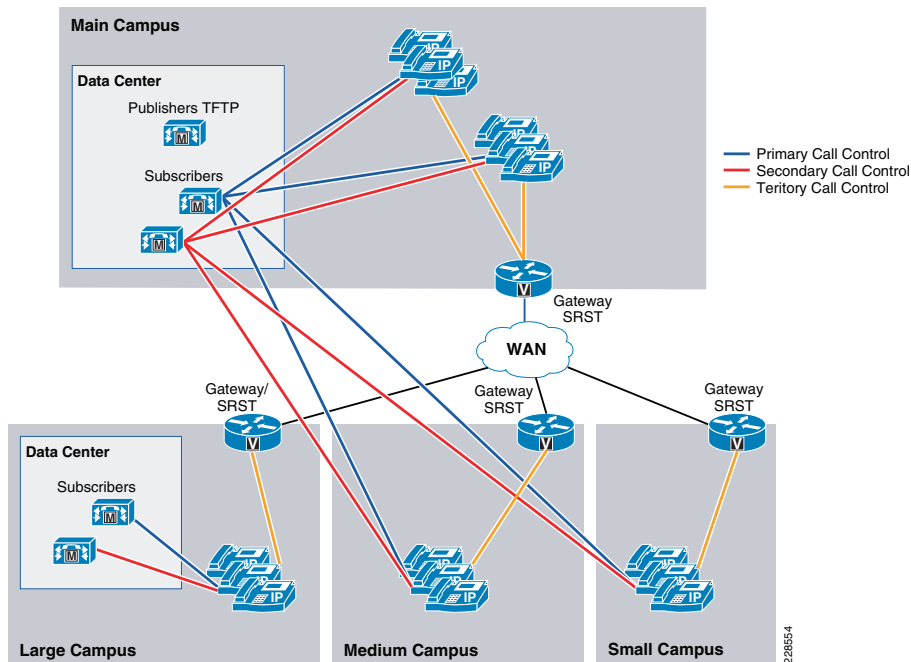
Other services or applications might also be affected, and their ability to function without the publisher should be verified when deployed.

Within the clustering over the IP WAN deployment scenario there are two possible failover models:

- *Option 1—Local Failover Deployment Model*
 - This deployment scenario provides the highest level of redundancy since each site that contains cluster servers houses a primary subscriber/s and a backup subscriber. Users within the site would fail from their assigned subscriber at the site to the backup subscriber at the same site. If by chance all subscribers at the site fail SRST is the final failover call control target.
- *Option 2—Remote Failover Deployment Model*
 - In this model, each site that contains a cluster server does not contain a backup server target. The backup servers might be within the primary campus for all users or users might not have a backup target at all and uses SRST during server failures. If a backup server is used at primary campus then users would register with that server over the WAN. This additional traffic must be considered when designing for failover.

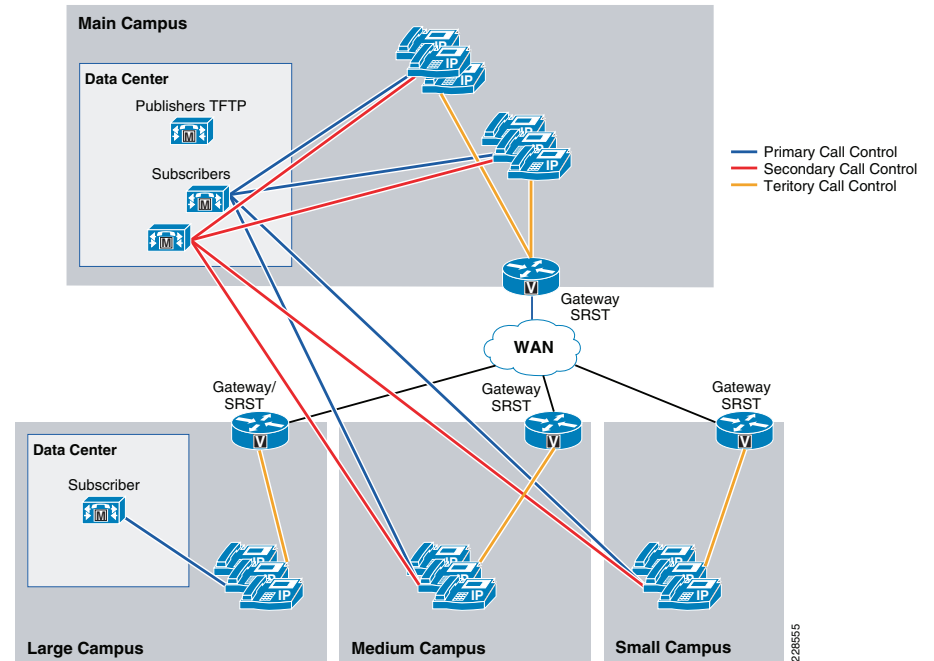
In **Figure 3**, option 1 is depicted with a backup server at the main campus and also at the remote large campus.

Figure 3 Option 1—Local Failover Deployment Model



In **Figure 4**, Option 2 is depicted with a backup server at the main campus only. Users at the remote large campus would failover to the main campus subscriber over the WAN and then to their local SRST target.

Figure 4 Option 2—Remote Failover Deployment Model



WAN Requirements

For clustering over the WAN to be successful, you must carefully plan, design, and implement various characteristics of the WAN itself. The Intra-Cluster Communication Signaling (ICCS) between Unified CM servers consists of many traffic types. The ICCS traffic types are classified as either priority or best-effort. Priority ICCS traffic is marked with IP Precedence 3 (DSCP 24 or PHB CS3). Best-effort ICCS traffic is marked with IP Precedence 0 (DSCP 0 or PHB BE).

ICCS consists of the following:

- *Database traffic*
 - Configuration information
- *Firewall management traffic*
 - Used to authenticate the subscribers to the publisher to access the publisher's database
- *ICCS real-time traffic*
 - Signaling, call admission control, and info about calls as they are initiated or completed
 - Uses TCP connection to all servers with Cisco CallManager Service-enabled

- Full mesh between all servers (maximum of 8 servers in a cluster running CM Service so there may be up to seven connections on each server)
- *CTI Manager*
- Used for CTI devices involved in calls or for controlling or monitoring other third-party devices

In order for the ICCS traffic to traverse the WAN there are very specific WAN requirements:

- *Delay*
 - Maximum one-way delay between any 2 servers can not exceed 40 msec (or 80msec round-trip). If clustering over the WAN is being used between data centers, any additional security that is applied both within and between those data centers has to fit within the maximum round-trip time that is allowed between nodes in a cluster.
- *Jitter*
 - Jitter for the IP Precedence 3 ICCS traffic must be minimized using Quality of Service (QoS)
- *Packet loss and errors*
 - The network should be engineered to provide sufficient prioritized bandwidth for all ICCS traffic, especially the priority ICCS traffic. Standard QoS mechanisms must be implemented to avoid congestion and packet loss. If packets are lost due to line errors or other “real-world” conditions, the ICCS packet will be retransmitted because it uses the TCP protocol for reliable transmission. The retransmission might result in a call being delayed during setup, disconnect (teardown), or other supplementary services during the call. Some packet loss conditions could result in a lost call, but this scenario should be no more likely than errors occurring on a T1 or E1, which affect calls via a trunk to the PSTN/ISDN.
- *Bandwidth*
 - Provision the correct amount of bandwidth between each server for the expected call volume, type of devices, and number of devices. This bandwidth is in addition to any other bandwidth for other applications sharing the network, including voice and video traffic between the sites. The bandwidth provisioned must have QoS enabled to provide the prioritization and scheduling for the different classes of traffic.
 - A minimum of 1.544 Mbps (T1) is required for ICCS for 10,000 busy hour call attempts (BHCA) between the sites that are clustered over the WAN. The *Cisco Unified Communications SRND Based on Cisco Unified Communications Manager 7.x* contains the calculations for establishing how much bandwidth is required.
 - In addition to the bandwidth required for Intra-Cluster Communication Signaling (ICCS) traffic, a minimum of 1.544 Mbps (T1) bandwidth is required for database and other inter-server traffic for every subscriber server remote to the publisher.
 - When directory numbers are shared between sites that are clustered over the WAN, additional bandwidth must be reserved.

- *QoS*
 - The network infrastructure relies on QoS engineering to provide consistent and predictable end-to-end levels of service for traffic. Neither QoS nor bandwidth alone is the solution; rather, QoS-enabled bandwidth must be engineered into the network infrastructure.

Intra-Cluster Communications

Intra-cluster communications represents all traffic between the servers in a cluster. The real-time protocol called Intra-Cluster Communication Signaling (ICCS) provides the communications with the Cisco CallManager Service process.

The intra-cluster traffic between the servers consists of the following:

- Database traffic from the IBM Informix Dynamic Server (IDS) database that provides the main configuration information. The IDS traffic may be reprioritized in line with Cisco QoS recommendations to a higher priority data service (for example, IP Precedence 1 if required by the particular business needs).
- Firewall management traffic, which is used to authenticate the subscribers to the publisher to access the publisher's database. The management traffic flows between all servers in a cluster. The management traffic may be prioritized in line with Cisco QoS recommendations to a higher priority data service (for example, IP Precedence 1 if required by the particular business needs).
- ICCS real-time traffic, which consists of signaling, call admission control, and other information regarding calls as they are initiated and completed. ICCS uses a Transmission Control Protocol (TCP) connection between all servers that have the Cisco CallManager Service-enabled. The connections are a full mesh between these servers. Because only eight servers may have the Cisco CallManager Service enabled in a cluster, there may be up to seven connections on each server. This traffic is priority ICCS traffic and is marked dependant on release and service parameter configuration.
- CTI Manager real-time traffic is used for CTI devices involved in calls or for controlling or monitoring other third-party devices on the Unified CM servers. This traffic is marked as priority ICCS traffic and exists between the Unified CM server with the CTI Manager and the Unified CM server with the CTI device.

For detailed information on various types of traffic between Unified CM servers, refer to the port usage document at http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/port/7_0/CCM_70PortList.pdf.

Gateways

In this design, gateways will be configured at every location. If a location has cluster servers within that location, gateways at that location will register with the local servers via device pools. Partitions and Calling Search Spaces will also be configured so local gateways are the first choice for PSTN calls. In some cases, overflow to other location gateways is allowed but WAN bandwidth must be considered so as to not oversubscribe it. Another consideration for overflow is in the event of local gateway failure causing all PSTN calls to overflow to another locations gateway. This could suddenly oversubscribe not only the WAN but the remote locations PSTN call handling ability (number of available

trunks). Having PSTN access at every location is necessary for emergency services and also as a means of making and receiving calls when in SRST mode. Every location will also support SRST as a last resort call control mechanism.

PSTN Trunk Sizing

Calculating Traffic

In voice communications, it is important to understand how much traffic is being carried or is expected to be carried by the public switched telephone network (PSTN) circuits. This traffic is measured in Erlangs. 1 Erlang can be considered as 1 hour of traffic. As an example, a single call lasting 1 hour would be 1 Erlang of carried traffic.

There are several ways to calculate how much traffic is currently being carried over a trunk group. A traffic study on the existing telephony system or a traffic study from the carrier. In a new installation, this is hard to project. The following subsections describe this.

Traffic in an Existing Location

Existing Phone System Traffic Reports

If a location already has a telephony system installed, traffic reports can usually be run from the perspective of that telephony system. In general terms, you are looking for the worst case scenario: busy day and busy hour. These traffic reports come in many varieties and measure in many different ways. The goal is to calculate how many seconds of use per phone user is consumed during the busy hour. Some reports break traffic into the following three parts:

- In (inbound from an external trunk)
- Out (outbound to an external trunk)
- Station to station (calls between two phones on the same system).

In calculating trunk traffic, we are not concerned with station to station call traffic so it should be subtracted from the total. Station to station traffic might be a consideration for calls over the WAN to another station but not for local PSTN traffic.

Typically, traffic from the station perspective is measured in Centum Call Seconds (CCS). A CCS is equal to 100 seconds so 400 seconds would be expressed as four CCS. There are 3600 seconds (36 CCS) in an hour so the most traffic a user could theoretically consume is 36 CCS. In a call center, each user that will be on-line is usually considered to be able to use trunk facilities for a full 36 CCS.

For example, a report that might indicate that the average phone usage (In + Out) per user is 400 seconds during the busy hour (09:00 to 10:00). In telephony, this would be called 4 Centum Call Seconds (4 CCS) per user.

After deriving the CCS/user, the total trunk traffic can be calculated. Trunk traffic is usually measured in Erlangs or hours of traffic. If a location has 125 users at 4 CCS/user, then the total busy hour traffic for that location would be 500 CCS or 50,000 seconds. $50,000/3600$ (seconds in an hour) = 13.9 Erlangs (hours) of traffic. Another way to calculate is $500 \text{ CCS}/36 \text{ CCS}$ (36 CCS in an hour) = 13.9 Erlangs.

After calculating the total busy hour traffic (13.9 Erlangs), it is now possible to estimate how many trunks are required. Before doing this, a couple decisions need to be made. The calculators that suggest the number of trunks required will ask for a Grade-of-Service

(GoS) or level of acceptable blockage. In most cases, a GoS of .01 (1 call in 100 will be blocked) is used. The other question that usually asked is what calculation method should be used. The following are the typical methods:

- Erlang B
 - Unsuccessful calls are blocked (typical non-call center)
- Erlang C
 - Unsuccessful calls are queued (call center based)

There are several other methods however, these are the most widely used.

To continue the example, in this instance we will be using 13.9 (round up to 14 Erlangs), a blockage factor (GoS) of .01, and the Erlang B calculation method. If these values are used in one of the many on-line traffic calculators, it will provide a value of 23 lines (trunks).

One on-line calculator can be found at the following URL:

<http://www.erlang.com/calculator/erlb/>

Carrier Traffic Reports

If reports are not available from an existing telephony system, the carrier can usually provide this information. They do not usually store this data, so a traffic report is usually requested for a specific date and time. In some cases, the carrier will provide a "Peg" report that is only a count of calls without duration. This is not very helpful unless you have some other way of deriving the average call length. If users are complaining of not being able to make an external call or customers are complaining of getting a busy on inbound calls, these reports will also usually show call blockage patterns so more trunks can be added.

After receiving the data from the carrier, a busy hour calculation can be made to see what the total traffic is and the same methods as above can be used to arrive at the total number of trunks used. In many cases companies are over-trunked so this can be used to reduce the financial burden associated with too many trunks.

Traffic in a New Location

If traffic reports are available for existing locations, this can be used to project the traffic for a new location. If no traffic is available some assumptions can be made. The rule of thumb is that most hospitals and universities are in the 3 CCS to 4 CCS/user range and most Enterprise customers in the 5 CCS-7 CCS range. Hospitals and universities are usually on the lower end because they have a large number of rarely used phones (patient rooms, classrooms, etc.). The typical Enterprise has fewer phones, therefore the traffic per phone goes up. The nature of the business can also drive the average CCS/user higher as well. It is always better to over-trunk, monitor and fine tune than to under-trunk and cause employee and customer frustration.

Dial Plan Considerations

In many cases, the dial plan is conceived when a system is very small and there is no expectation that multiple locations will be added in the future. For a community college that might be in the initial phases of designing a system, it would be very important to design a dial plan that can be expanded without users having to change their DN. This not only causes user frustration, but ultimately causes customer confusion and added expense since stationary, business cards, Web Sites, etc. must all be updated.

For this design a variable-length on-net dial plan will be used. This allows users within a site to dial a 4 or 5 digit number. To reach users in other sites, an access code (usually 8) followed by the location code and then extension would be dialed. This dial plan allows for future growth and more flexibility within each site.

For more information on dial plans, refer to the following documents:

- *Cisco Service Ready Architecture for Schools Design Guide*
http://www.cisco.com/en/US/docs/solutions/Enterprise/Education/SchoolsSRA_DG/SchoolsSRA-DG.html
- *Cisco Unified Communications SRND Based on Cisco Unified Communications Manager 7.x*
http://www.cisco.com/en/US/docs/voice_ip_comm/cucm/srnd/7x/uc7_0.html

SRST

Cisco Unified SRST provides Cisco Unified Communications Manager with fallback support for Cisco Unified IP phones that are attached to a Cisco router on the local network. Cisco Unified SRST enables routers to provide call-handling support for Cisco Unified IP phones when they lose connection to remote primary, secondary, or tertiary Cisco Unified Communications Manager installations or when the WAN connection is down.

When Cisco Unified IP phones lose contact with primary, secondary, and tertiary Cisco Unified Communications Managers, they must establish a connection to a local Cisco Unified SRST router to sustain the call-processing capability necessary to place and receive calls. The Cisco Unified IP phone retains the IP address of the local Cisco Unified SRST router as a default router in the Network Configuration area of the Settings menu. The Settings menu supports a maximum of five default router entries; however, Cisco Unified Communications Manager accommodates a maximum of three entries. When a secondary Cisco Unified Communications Manager is not available on the network, the local Cisco Unified SRST Router's IP address is retained as the standby connection for Cisco Unified Communications Manager during normal operation.

Configuration examples, compatibility, etc. can be at the following URL:
http://www.cisco.com/en/US/partner/products/sw/voicesw/ps2169/tsd_products_support_series_home.html

Community College Mission Relevancy

This chapter discusses the possibilities of a comprehensive Unified Communications architecture for community colleges. The following subsections provide a non-technical discussion of how the UC considerations might fit into the foundation services or introduced earlier in this document. These foundation services are as follows:

- Virtual Learning Environment
- Secure Connected Classrooms
- Safety and Security
- Operational Efficiencies

Virtual Learning Environment

This UC architecture provides a base call processing platform that allows remote student or instructor softphone or hardphone connectivity incorporating all of the features and capabilities of an on-site handset. This design also provides the call processing and media ready network for high bandwidth video applications such as Telepresence

Secure Connected Classrooms

Secure Connected Classrooms typically implies network connectivity. From a UC perspective this could also include secure handsets within the classroom. With the use of Extension Mobility the classroom phones would have very basic capabilities like 911 service and internal station to station access. When an instructor enters the room and logs on to the phone it now takes on the characteristics assigned to that instructor that include voicemail indicator, calling capabilities like long distance or international. The instructor can also then be reached regardless of campus location.

Safety and Security

Safety and security within the Unified Communications infrastructure can pose many questions:

- Is the system secure from network intrusion?
 - Is this covered in security section?
- Does the UC platform offer any features or design considerations to assist campus security or emergency services?
 - Although not covered in this document the UC platform can interface with a Cisco Product called Emergency Responder. Cisco Emergency Responder enhances the existing emergency 9-1-1 functionality offered by Cisco Unified Communications Manager. It assures that Cisco Unified Communications Manager will send emergency calls to the appropriate Public Safety Answering Point (PSAP) for the caller's location, and that the PSAP can identify the caller's location and return the call if necessary. In addition, the system automatically tracks and updates equipment moves and changes.

Cisco Emergency Responder includes the following features:

- Real-time location-tracking database and enhanced routing capabilities
- Supports automatic notification of customer security personnel when an emergency call is in progress and provides the caller's location
- Requires no administrative support for moving phones or staff from one location to another
- The design does include PSTN trunks at every location. This will provide an address for each location to the PSAP but not offer the granularity provided by CER.
- Does the system offer high availability and survivability so communication can continue even when the WAN is down?
 - At the main campus each device has a primary, secondary, and tertiary target for call control.

- All other locations have a primary and secondary target for call control. The key point is that call control for these locations is typically the main campus under normal conditions. If a WAN component fails then each site has a local SRST router that can act as call control and still offer station to station as well as PSTN access.

Operational Efficiencies

The UC design set forth in this document provides for a single system with one system administration location. Regardless of the location of the phone all system administration is handled within a single administration access.