



Emergency Services

Revised: June 14, 2016

Emergency services are of great importance in the proper deployment of a communications system. This chapter presents a summary of the following major design considerations essential to planning for emergency calls:

- [911 Emergency Services Architecture, page 15-2](#)
- [Cisco Emergency Responder, page 15-10](#)
- [High Availability for Emergency Services, page 15-12](#)
- [Capacity Planning for Cisco Emergency Responder Clustering, page 15-13](#)
- [Design Considerations for 911 Emergency Services, page 15-13](#)
- [Cisco Emergency Responder Deployment Models, page 15-22](#)
- [ALI Formats, page 15-29](#)

This chapter presents some information specific to the 911 emergency networks as deployed in Canada and the United States. Many of the concepts discussed here are adaptable to other locales. Please consult with your local telephony network provider for appropriate implementation of emergency call functionality.

In the United States, some states have already enacted legislation covering the 911 functionality required for users in a multi-line telephone system (MLTS). The National Emergency Number Association (NENA) has also produced the *NENA Technical Requirements Document on Model Legislation E9-1-1 for Multi-Line Telephone Systems*, available online at

<http://www.nena.org/>

This chapter assumes that you are familiar with the generic 911 functionality available to residential PSTN users in North America.



Note

The topics discussed in this chapter apply to Cisco Emergency Responder only when it is used in conjunction with Cisco Unified Communications Manager (Unified CM). Cisco TelePresence Video Communication Server (VCS) currently does not support emergency services.

What's New in This Chapter

Table 15-1 lists the topics that are new in this chapter or that have changed significantly from previous releases of this document.

Table 15-1 *New or Changed Information Since the Previous Release of This Document*

New or Revised Topic	Described in	Revision Date
Service Provider ALI (SP-ALI)	Service Provider ALI, page 15-3	June 14, 2016
SIP trunk support	Dynamic ANI (Trunk Connection), page 15-7	June 14, 2016
Access point tracking	Device Location Discovery Methods in Cisco Emergency Responder, page 15-10	June 14, 2016
Location awareness for wireless clients	Cisco Emergency Responder and Location Awareness for Wireless Clients, page 15-19	June 14, 2016
Other minor updates	Various sections of this chapter	June 14, 2016
Cisco Unified Communications Manager Native Emergency Call Routing	Emergency Call Routing Using Unified CM Native Emergency Call Routing, page 15-27	June 15, 2015

911 Emergency Services Architecture

This section highlights some of the functionality requirements for emergency calls in multi-line telephone systems (MLTS). In the context of this section, emergency calls are 911 calls serviced by the North American public switched telephone network (PSTN).

Any emergency services architecture usually consists of the following elements:

- A distressed caller should be able to dial the emergency services from a fixed line, a mobile phone, a public phone, or any device capable of making the voice call.
- An emergency services call handler must be available to respond to the emergency request and dispatch the needed services such as police, fire, and medical.
- In order to provide help, the call handler should be able to identify the location of the distressed caller as precisely as possible.
- An emergency services network is needed to route the call to the nearest emergency services call handler with jurisdiction for the location of the caller.

The following sections explain some of the important architectural components of 911 emergency services architecture.

Public Safety Answering Point (PSAP)

The public safety answering point (PSAP) is the party responsible for answering the 911 call and arranging the appropriate emergency response, such as sending police, fire, or ambulance teams. The physical location of the phone making the 911 call is the primary factor in determining the appropriate PSAP for answering that call. Generally, each building is serviced by one local PSAP.

To determine the responsible PSAP for a given location, contact a local public safety information service such as the local fire marshal or police department. Also, the phone directory of the local exchange carrier usually lists the agency responsible for servicing 911 calls in a given area.

Typical Situation

- For a given street address, there is only one designated PSAP.
- For a given street address, all 911 calls are routed to the same PSAP.

Exceptional Situation

- The physical size of the campus puts some of the buildings in different PSAP jurisdictions.
- Some of the 911 calls need to be routed to an on-net location (campus security, building security).

Selective Router

The selective router is a node in the emergency services network that determines the appropriate PSAP for call delivery, based on caller's geographic area and the automatic number identification (ANI). The Local Exchange Carrier (LEC) usually operates the selective router. Hence, it is imperative to ensure that the enterprise IP communications network is designed in such a way that the caller is routed to the appropriate selective router based on its location.

Automatic Location Identifier Database


Location information of the caller is an important part of the 911 services infrastructure. The Automatic Location Identifier (ALI) database maintains the location information for the particular geographical location served by the LEC. For every 911 call, the PSAP searches the ALI database to retrieve the caller's location based on the ANI of the calling number. The addresses are stored in the Master Street Address Guide (MSAG) format in the ALI database. The ALI database is maintained on behalf of the local emergency services administration by a contracted third party, generally the incumbent Local Exchange Carrier (LEC).

Service Provider ALI

Service Provider ALI (SP-ALI) refers to a configuration in which the service provider is responsible for defining and maintaining the ALI information for all emergency calls over the connection. SP-ALI service uses the physical interconnection at the LEC to determine the source location of the call. For residential customers, the ALI information is associated with the address of the subscriber and the directory number of that resident. Because the ALI information is determined by the service provider based upon the physical interconnection in the LEC, the subscriber does not have the ability to change or set the ALI information.

The setting of the ALI information based on the physical point of interconnection of the line or trunk applies to PRI trunk connections also. By default, an MLTS operator that uses PRI trunks for PSTN access will have SP-ALI service. The LEC defines the calling party number (CPN) and ALI address for emergency calls. Typically, the calling party number used for emergency calls is the customer's bill-to number (BTN) or the MLTS operator's main number. The physical address associated with the emergency calling number is the address of the demark of the PRI at the customer's facility. If a PRI trunk is set for SP-ALI service, all calls to 911 have the calling party number replaced by the LEC to match the ALI record for the customer. (See [Figure 15-1](#).)

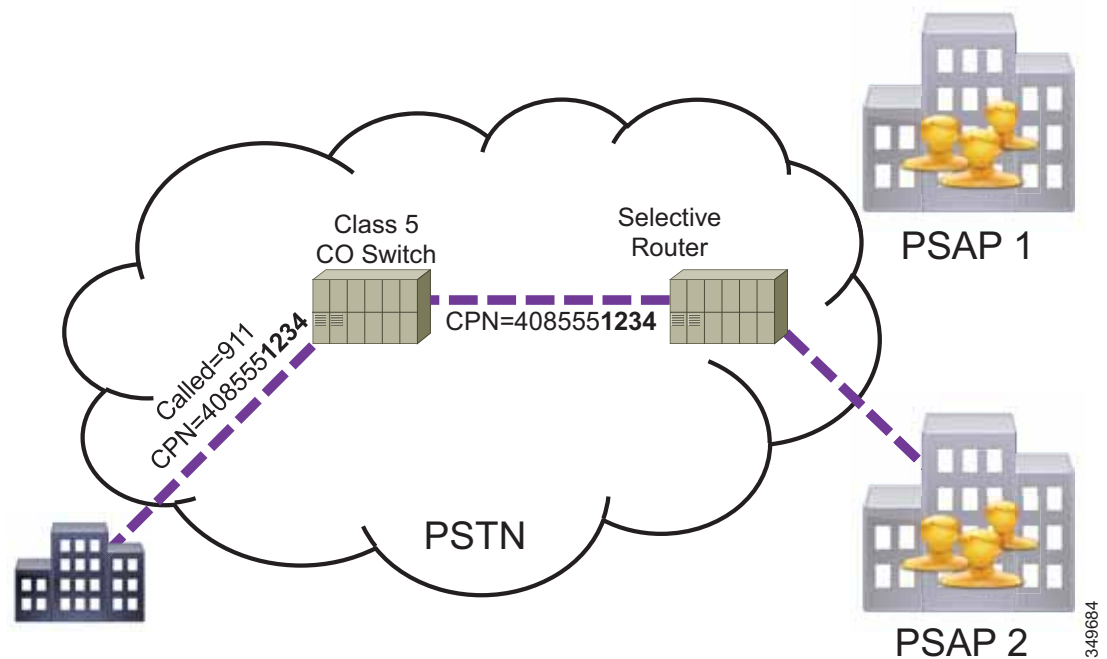
Figure 15-1 Service Provider ALI



Private Switch ALI

Private Switch ALI (PS-ALI) is an enhancement to 911 emergency response systems that enables MLTS operators to provide more specific address and location information for each endpoint. The service allows a customer-generated address table to be loaded into the ALI database so that each station of an MLTS system can be uniquely identified if a call is placed to 911 from that telephone number. The station-specific or location-specific automatic number identification (ANI) generated by the communications system can be passed directly to the E911 system to pinpoint the precise location of the caller. (See [Figure 15-2](#).) The PSAP operator can then direct emergency response personnel to the correct address, building, floor, room, or even cubicle, thereby streamlining operations and increasing accuracy.

Figure 15-2 Private Switch ALI



911 Network Service Provider

After identifying the responsible PSAPs, you must also identify the 911 network service providers to which each PSAP is connected. It is commonly assumed that PSAPs receive 911 phone calls from the PSTN, but that is not the case. Instead, 911 calls are carried over dedicated, regionally significant networks, and each PSAP is connected to one or more such regional networks. In the majority of cases, the incumbent Local Exchange Carrier (LEC) is the 911 network service provider for a PSAP. Some exceptions include military installations, university campuses, federal or state parks, or other locations where the public safety responsibility falls outside the jurisdiction of the local authorities and/or where a private network is operated by an entity other than a public local exchange carrier.

If you are in doubt about the 911 network service provider for a given PSAP, contact the PSAP directly to verify the information.

Typical Situation

- For a given street address, the 911 network service provider is the incumbent Local Exchange Carrier (LEC). For a location served by Phone Company X, the corresponding PSAP is also served by Phone Company X.
- All 911 calls are routed directly to an off-net location, or all 911 calls are routed directly to an on-net location.

Exceptional Situation

- The local exchange carrier (LEC) through which the MLTS interfaces to the PSTN is *not* the same LEC that serves as 911 network service provider to the PSAP. (For example, the communications system is served by Phone Company X, but the PSAP is connected to Phone Company Y.) This situation might require either a special arrangement between the LECs or special, dedicated trunks between the phone system and the PSAP's 911 network service provider.
- Some LECs may not accept 911 calls on their networks. If this is the case, the only two options are to change LECs or to establish trunks (dedicated to 911 call routing) connected to a LEC that can route 911 calls to the appropriate PSAPs.
- Some (or all) of the 911 calls have to be routed to an on-net location such as campus security or building security. This situation can easily be accommodated during the design and implementation phases, but only if the destination of 911 calls for each phone has been properly planned and documented.

Interface Points into the Appropriate 911 Networks

For larger communications systems, 911 connectivity might require many interface points. Typically, more than one E911 selective router is used within a LEC's territory, and these routers usually are *not* interconnected.

For example, an enterprise with a large campus could have the following situation:

- Building A located in San Francisco
- Building B located in San Jose
- San Francisco Police Department and San Jose Police Department are the appropriate PSAPs
- San Francisco Police Department and San Jose Police Department are served by the same 911 network service provider
- However, San Francisco Police Department and San Jose Police Department are served by different E911 selective routers operated by that same 911 network service provider!

This type of situation would require two separate interface points, one per E911 selective router. The information pertaining to the E911 selective router territories is generally kept by the incumbent LEC, and the local account representative for that LEC should be able to provide an enterprise customer with the pertinent information. Many LECs also provide the services of 911 subject matter experts who can consult with their own account representatives on the proper mapping of 911 access services.

Typical Situation

- For single-site deployments or campus deployments, there is usually only one PSAP for 911 calls.
- If access to only one PSAP is required, then only one interface point is required. Even if access to more than one PSAP is required, they might be reachable from the same E911 selective router, through the same centralized interface. If the enterprise's branch sites are linked via a WAN (centralized call processing), it is desirable to give each location its own local (that is, located inside each branch office) access to 911 to prevent 911 isolation during WAN failure conditions where Survivable Remote Site Telephony (SRST) operation is activated.

Exceptional Situation

- The physical size of the campus puts some of the buildings in different PSAP jurisdictions, *and*
- Some of the 911 calls have to be routed to different E911 selective routers, through different interface points.

**Note**

Some of the information required to establish the geographical territories of PSAPs and E911 selective routers is available online or from various competitive local exchange carrier (CLEC) information web sites. (For example, <https://clec.att.com/clec/hb/shell.cfm?section=782> provides some valuable data about the territory covered by AT&T in California and Nevada.) However, Cisco strongly recommends that you obtain proper confirmation of the appropriate interface points from the LEC prior to the design and implementation phases of 911 call routing.

Interface Type

In addition to providing voice communications, the interfaces used to present 911 calls to the network must also provide identification data about the calling party.

Automatic Number Identification (ANI) refers to the North American Numbering Plan number of the calling party, which is used by networks to route a 911 call to the proper destination. This number is also used by the PSAP to look up the Automatic Location Identification (ALI) associated with a call.

911 calls are source-routed, which means that they are routed according to the calling number. Even though different locations are all dialing the same number (911), they will reach different PSAPs based on their location of origin, which is represented by the ANI (calling number).

You can implement 911 call functionality with either of the following interface types:

- Dynamic ANI assignment
- Static ANI assignment

While dynamic ANI assignment scales better (because it supports multiple ANIs) and lends itself to all but the smallest of applications, static ANI assignment can be used in a wider variety of environments, from the smallest to the largest systems.

Dynamic ANI (Trunk Connection)

The dynamic aspect of ANI refers to the fact that a communications system has many endpoints sharing access to the 911 network across the same interface, and the ANI transmitted to the network might need to be different for each call.

There are three main types of dynamic ANI interfaces:

- Integrated Services Digital Network Primary Rate Interface (ISDN-PRI, or simply PRI)
- Session Initiation Protocol (SIP) trunk
- Centralized Automatic Message Accounting (CAMA).

PRI

This type of interface usually connects a communications system to a PSTN Class 5 switch. The calling party number (CPN) is used at call setup time to identify the E.164 number of the calling party.

Most LECs treat the CPN differently when a call is made to 911. Depending upon the functionality available in the Class 5 switch and/or upon LEC or government policy, the CPN may not be used as the ANI for 911 call routing. Instead, the network may be programmed to use the listed directory number (LDN) or the bill-to number (BTN) for ANI purposes.

If the CPN is not used for ANI, then 911 calls coming from a PRI interface all look the same to the 911 network because they all have the same ANI, and they are all routed to the same destination (which might not be the appropriate one). The replacement of the CPN by the LEC is typically called Service Provider ALI (SP-ALI), because the service provider specifies the CPN for ALI lookup.

Some LECs offer a feature to provide CPN transparency through a PRI interface for 911 calls. With this feature, the CPN presented to the Class 5 switch at call setup is used as ANI to route the call. The feature name for this functionality varies, depending on the LEC. (For example, SBC calls it Inform 911 in California.)

**Note**

When SP-ALI service is used, the CPN *must* be a routable North American Numbering Plan number, which means that the CPN must be entered in the routing database of the associated E911 selective router.

**Note**

For Direct Inward Dial (DID) phones, the DID number could be used as the ANI for 911 purposes, but only if it is properly associated with an Emergency Service Number in the 911 service provider's network. For non-DID phones, use another number. (See [Emergency Location Identification Number Mapping, page 15-14](#), for more information.)

Many Class 5 switches are connected to E911 selective routers through trunks that do not support more than one area code. In such cases, if PRI is used to carry 911 calls, then the only 911 calls that will be routed properly are those whose CPN (or ANI) have the same Numbering Plan Area (NPA) as the Class 5 switch.

Example

An MLTS is connected to a Class 5 switch in area code 514 (NPA = 514). If the MLTS were to send a 911 call on the PRI trunk, with a CPN of **450.555.1212**, the Class 5 switch would send the call to the E911 selective router with an ANI of **514.555.1212** (instead of the correct **450.555.1212**), yielding inappropriate routing and ALI lookup.

To use PRI properly as a 911 interface, the system planner must ensure that the CPN will be used for ANI and must properly identify the range of numbers (in the format NPA NXX TNTN) acceptable on the link. For example, if a PRI link is defined to accept ANI numbers within the range 514 XXX XXXX, then only calls that have a Calling Party Number with NPA = 514 will be routed appropriately.

SIP Trunk

SIP trunking is an IP-only interface that connects a communications system to a service provider, typically through a Session Border Controller (SBC). SIP trunks allow for the same dynamic calling party number delivery to the carrier as PRI trunks; but unlike PRI trunks, SIP trunks do not have a physical limit on the number of calls that can be established concurrently.

When emergency services are called over a SIP trunk, delivery of the call to the correct selective router must be verified with the provider. Unlike PRI circuits that terminate at the local LEC, SIP trunks might not have a physical connection with the local LEC and as a result will not automatically route 911 calls to the selective router in the municipality of the calling party. Additionally, each SIP trunk provider might have different E911 routing capability; for example, one service provider may be able to deliver calls to selective routers across the US based upon the calling party number (even outside the local area), while another service provider may allow E911 calls into only one customer-specified selective router. A Cisco Unified CM administrator should always confirm the 911 call delivery capabilities with the carrier, especially when a SIP trunk is providing centralized call routing.

SIP service providers are required to route 911 calls to the appropriate rate center or PSAP for any DID number that they service over a SIP trunk. For example, assume that a deployment has a SIP trunk that physically terminates in a data center in Dallas Texas that services DIDs for a San Francisco office with the range of 415-555-1xxx and for a New York office with the range of 212-448-2xxx. If a call to 911 is placed from 415-555-1800, then the SIP provider must route the call to the San Francisco selective router for PSAP delivery. If a user at extension 212-448-2840 in the New York City office dials 911, the call can be routed on the same SIP trunk to the appropriate selective router in the New York City area to reach the PSAP appropriate for the caller.

CAMA

Centralized Automatic Message Accounting (CAMA) trunks also allow the MLTS to send calls to the 911 network, with the following differences from the PRI approach:

- CAMA trunks are connected directly into the E911 selective router. Extra mileage charges may apply to cover the distance between the E911 selective router and the MLTS gateway point.
- CAMA trunks support 911 calls only. The capital and operational expenses associated with the installation and operation of CAMA trunks support 911 traffic only.
- CAMA trunks for the MLTS market may be limited to a fixed area code, and the area code is typically implied (that is, not explicitly sent) in the link protocol. The connection assumes that all calls share the same deterministic area code, therefore only 7 or 8 digits are sent as ANI.

Static ANI (Line Connection)

Static ANI provides a line (rather than a trunk) connection to the PSTN, and the ANI of the line is associated with all 911 calls made on that line, regardless to the CPN of the calling phone. Static ANI is based on the physical interconnection point in the LEC. Because the Static ANI is defined by the carrier on the interconnection point in the LEC, Static ANI emergency call routing is also referred to as Service Provider ALI (SP-ALI). A plain old telephone service (POTS) line is the most common type of connection used for this purpose.

POTS lines are one of the simplest and most widely supported PSTN interfaces. A POTS line usually comes fully configured to accept 911 calls. In addition, the existing E911 infrastructure supports 911 calls from POTS lines very well.

The POTS approach has the following attributes:

- The operational costs associated with a POTS line are low.
- The POTS line can even serve as a backup line in case of power failure.
- The POTS line number can be used as the callback number entered into the ALI database.
- POTS lines represent the lowest cost 911 support for locations where user density does not justify local PRI or CAMA access into the PSTN.
- POTS lines are ubiquitous in PSTN installations.

All outgoing 911 calls through this type of interface are treated the same by the E911 network, and any tools that enable ANI manipulation presented to the E911 network (such as translations or transformations) are irrelevant because the ANI can be only the POTS line's number.

Cisco Emergency Responder

Ease of administration for moves, adds, and changes is one of the key advantages of IP communications technology. To provide for moves, adds, and changes that automatically update 911 information without user intervention, Cisco has developed a product called the Cisco Emergency Responder (Emergency Responder).

Cisco Emergency Responder provides the following primary functionality:

- Dynamic association of a phone to an Emergency Response Location (ERL), based on the detected physical location of the phone.
- Dynamic association of the Emergency Location Identification Number (ELIN) to the calling phone, for callback purposes. In contrast to the general emergency services scenarios outlined in preceding sections, Cisco Emergency Responder enables the callback to ring the exact phone that initiated the 911 call.
- On-site notification to designated parties (by pager, web page, email, or phone call) to inform them that there is an emergency call in progress. Email, pager, and web page notifications include the calling party name and number, the ERL, and the date and time details associated with the call. Phone notification provides the information about the calling number from which the emergency call was placed.

For more information on ERLs and ELINs, see [Emergency Response Location Mapping, page 15-13](#), and [Emergency Location Identification Number Mapping, page 15-14](#). For more information on Cisco Emergency Responder, see [Cisco Emergency Responder Design Considerations, page 15-19](#), and refer to the Cisco Emergency Responder product documentation available online at

http://www.cisco.com/en/US/products/sw/voicesw/ps842/tsd_products_support_series_home.html

Device Location Discovery Methods in Cisco Emergency Responder

Cisco Emergency Responder uses multiple methods to determine the physical location of a device. Because more specific location discovery results in a shorter time to locate the emergency and administer emergency services, Emergency Responder uses the following methods (listed in priority order) to identify an emergency caller's location:

1. Switch port discovery
2. Access point association
3. IP subnet
4. Static DN assignment
5. Default route

Switch Port Discovery

The primary method for location identification in Cisco Emergency Responder is the detection of an endpoint via Layer 2 discovery at the switch port level. Discovering an endpoint through Layer 2 Cisco Discovery Protocol (CDP) discovery enables Emergency Responder to determine the exact physical location of the calling device based on the physical termination of the network cable to a network jack in a cubicle or office. Although the discovery mechanism of the connected device is reliable, the accuracy of the physical location relies on two main assumptions:

- The wired infrastructure of the enterprise is well established and does not change sporadically, and any wiring closet changes trigger notification to the Emergency Responder administrator indicating what changed.
- The infrastructure is available for Cisco Emergency Responder to browse; that is, Cisco Emergency Responder can establish Simple Network Management Protocol (SNMP) sessions to the underlying network infrastructure and can scan the network ports for the discovery of connected phones.

Once Cisco Emergency Responder discovers the originating port for the call, it associates the call with the pre-established ERL for the location of that port. This process also yields an association with a pre-established ELIN for the location and the selection of the appropriate egress point to the E911 infrastructure, based on the originating ERL.

Access Point Association

Because wireless devices do not have the same discovery capability and tracking characteristics as a wired endpoint, Cisco Emergency Responder tracks wireless clients by using the Location Awareness feature available in Unified CM 11.5 and later releases. The Location Awareness feature allows Emergency Responder to synchronize all deployed access points in Unified CM and to assign the APs to the appropriate ERL. The Location Awareness feature also allows for the updating of mobile device movement between APs.

Emergency Responder is able to track wireless clients across the enterprise through the Location Awareness feature in Unified CM. When a mobile client associates with an AP in the enterprise, the device sends the Basic Service Set Identifier (BSSID) of the AP to Unified CM through call control. Unified CM then updates the database with the new AP association. Periodically, Emergency Responder requests device updates from Unified CM for any device that has updated its AP association since the last request. Emergency Responder receives only the devices that have moved since the last request. In Unified CM 11.5, the request interval is 2 minutes.

IP Subnet

Cisco Emergency Responder also provides the capability to configure ERLs for IP subnets and to assign IP endpoint location by IP address. This capability may be used to locate wireless IP phones, IP softphones, collaboration endpoints that do not support Cisco Discovery Protocol (CDP), and third-party SIP endpoints registered to Cisco Unified CM, which Cisco Emergency Responder cannot locate by connected switch port. It may also be used instead of, or in addition to, connected switch port locations for wired Cisco Collaboration endpoints. If both connected switch port and IP subnet locations are available for a Cisco Collaboration endpoint, Cisco Emergency Responder will prefer the connected switch port location because it is usually more specific than the IP subnet location. Using both connected switch port and IP subnet locations is a best practice because it provides assurance that an appropriate ERL will be assigned, even in case of any delay or error in detecting the connected switch port.

Cisco Emergency Responder allows for the use of two or more ELINs per ERL. The purpose of this enhancement is to cover the specific case of more than one 911 call originating from a given ERL within the same general time period, as illustrated by the following examples.

Example 1

- Phone A and phone B are both located within ERL X, and ERL X is associated with ELIN X.
- Phone A makes a 911 call at 13:00 hours. ELIN X is used to route the call to PSAP X, and PSAP X answers and releases the call. Then, at 13:15 hours, phone B makes a 911 call. ELIN X is again used to route the call to PSAP X.
- PSAP X, after releasing the call from phone B, decides to call back phone A for further details pertaining to phone A's original call. The PSAP dials ELIN X, and gets phone B (instead of the desired phone A).

To work around this situation, Cisco Emergency Responder allows you to define a pool of ELINs for each ERL. This pool provides for the use, in a round-robin fashion, of a distinct ELIN for each successive call. With the definition of two ELINs for ERL X in our example, we now have the situation described in Example 2.

Example 2

- Phone A and phone B are both located within ERL X. ERL X is associated with both ELIN X1 and ELIN X2.
- Phone A makes a 911 call at 13:00 hours. ELIN X1 is used to route the call to PSAP X, and PSAP X answers and releases the call. Then, at 13:15 hours, phone B makes a 911 call, and ELIN X2 is used to route this call to PSAP X.
- PSAP X, after releasing the call from phone B, decides to call back phone A for further details pertaining to phone A's original call. The PSAP dials ELIN X1 and gets phone A.

Of course, if a third 911 call were made but there were only two ELINs for the ERL, the situation would allow for callback functionality to properly reach only the last two callers in the sequence.

High Availability for Emergency Services

It is very important for emergency services to always be available to the user even under the most critical conditions. Therefore, high availability planning must be done carefully when deploying emergency services in an enterprise.

Cisco Emergency Responder supports clustering with a maximum of two servers in active/standby mode. The data is synchronized between the primary and the secondary Cisco Emergency Responder servers. To ensure that calls are routed to the secondary server if the primary server is unavailable, the system administrator must follow certain provisioning guidelines for configuring CTI route points and the directory numbers (DNs) associated to those CTI route points in Cisco Unified CM. For more details on configuration, refer to the *Cisco Emergency Responder Administration Guide*, available at

http://www.cisco.com/en/US/products/sw/voicesw/ps842/prod_maintenance_guides_list.html

If both of the Cisco Emergency Responder servers are unavailable, a local route group (LRG) may be used to route the call to the appropriate PSAP with an appropriate ELIN/ERL (which might be less specific than what Cisco Emergency Responder could have provided). Alternatively, the call may be routed to an internal security office to determine the caller's location. In either case, this provisioning must be done in Cisco Unified CM.

Apart from Cisco Emergency Responder redundancy, Cisco Unified CM redundancy and gateway/trunk redundancy should also be considered to route the 911 emergency calls and to avoid any single point of failure.

Capacity Planning for Cisco Emergency Responder Clustering

In a Cisco Emergency Responder cluster, the quantity of endpoints roaming outside the tracking domain of their home Cisco Emergency Responder group is a scalability factor that must be kept within the limits set forth in the section on *Network Hardware and Software Requirements* in the *Cisco Emergency Responder Administration Guide*, available at:

http://www.cisco.com/en/US/products/sw/voicesw/ps842/prod_maintenance_guides_list.html

For deployments that exceed the Emergency Responder maximum roaming capacity limit (for instance, large campus deployments with multiple Unified CM clusters), phone movement can be tracked by IP subnets. By defining the IP subnets in each of the Cisco Emergency Responder groups and by assigning each ERL with one ELIN per Cisco Emergency Responder group, you can virtually eliminate roaming phones because all phones in the campus will be part of the tracking domain of their respective Cisco Emergency Responder group.

To ensure proper sizing, use the Cisco Unified Communications Sizing Tool (Unified CST). This tool is available only to Cisco partners and employees, with appropriate login required, at <http://cucst.cloudapps.cisco.com/landing>. If you do not have access to this sizing tool, work with your Cisco account team or partner integrator to size your system appropriately.

Design Considerations for 911 Emergency Services

When planning 911 emergency services for multi-line telephone system (MLTS) deployments, first establish all of the physical locations where phone services are needed. The locations can be classified as follows:

- Single building deployments, where all users are located in the same building
- Single campus deployments, where the users are located in a group of buildings situated in close proximity
- Multisite deployments, where users are distributed over a wide geographical area and linked to the call processing site through WAN connectivity

The locations, or type of deployment, affect the criteria used to design and implement 911 services. The following sections describe the key criteria, along with typical and exceptional situations for each. When analyzing and applying these criteria, consider how they are affected by the phone locations in your network.

Emergency Response Location Mapping

The National Emergency Number Association (NENA) has proposed model legislation to be used by state and federal agencies in enacting the rules that govern 911 in enterprise communications systems. One of the concepts in the NENA proposal is that of the emergency response location (ERL), which is defined as:

A location to which a 911 emergency response team may be dispatched. The location should be specific enough to provide a reasonable opportunity for the emergency response team to quickly locate a caller anywhere within it.

Rather than having to identify each endpoint's location individually, the requirement allows for the grouping of endpoints into a "zone," the ERL. The maximum size of the ERL may vary, depending upon local implementation of the legislation, but we will use 7000 square feet (sq ft) as a basis for discussion in this section. (The concepts discussed here are independent of the maximum ERL size that may be allowed in any given state or region.)

An emergency location identification number (ELIN) is associated with each ERL. The ELIN is a fully qualified E.164 number, used to route the call within the E911 network. The ELIN is sent to the E911 network for any 911 call originating from the associated ERL. This process allows more than one phone to be associated with the same fully qualified E.164 number for 911 purposes, and it can be applied to DID and non-DID phones alike.


Note

This document does not attempt to present the actual requirements of any legislation. Rather, the information and examples presented here are for the purposes of discussion only. The system planner is responsible for verifying the applicable local requirements.

For example, assume a building has a work area of 70,000 sq ft and 100 endpoints. In planning for 911 functionality, the building can be divided into 10 zones (ERLs) of 7000 sq ft each, and each endpoint can be associated with the ERL where it is located. When a 911 call is made, the ERL (which could be the same for multiple endpoints) is identified by sending the associated ELIN to the PSAP. If the endpoints were evenly distributed in this example, each group of 10 endpoints would have the same ERL and, therefore, the same ELIN.

The various legislations define a minimum number of endpoints (for example, 49) and a minimum work area (for example, 40,000 sq ft) below which the requirements for MLTS 911 are not applicable. But even if the legislation does not require 911 functionality for a given enterprise, it is always best practice to provision for it.

Emergency Location Identification Number Mapping

In general, you must associate a single fully qualified E.164 number, known as the emergency location identification number (ELIN), with each ERL. (However, if using Cisco Emergency Responder, you can configure more than one ELIN per ERL.) The ELIN is used to route the call across the E911 infrastructure and is used by the PSAP as the index into the ALI database.

ELINs must meet the following requirements:

- The ELIN must be routable across the E911 infrastructure. (See the examples in the section on [Interface Type, page 15-7](#).) If an ELIN is not routable, 911 calls from the associated ERL will, at best, be handled according to the default routing programmed in the E911 selective router.
- Once the ERL-to-ELIN mapping of an enterprise is defined, the corresponding ALI records must be established with the LEC so that the ANI and ALI database records serving the PSAP can be updated accurately.
- The ELIN must be reachable from the PSAP for callback purposes.

The ELIN mapping process can be one of the following, depending on the type of interface to the E911 infrastructure for a given ERL:

- Dynamic ANI interface

With this type of interface, the calling party number identification passed to the network is controlled by the MLTS. The telephony routing table of the MLTS is responsible for associating the correct ELIN with the call, based on the calling endpoint's ERL. In scenarios where Cisco Emergency Responder is not deployed, the calling party number for calls made to 911 can be

modified by Unified CM using transformation masks. For example, all endpoints located in a given ERL can share the same calling search space that lists a partition containing a translation pattern (911) and a calling party transformation mask that would replace the endpoint's CPN with the ELIN for that location. On the other hand, if Cisco Emergency Responder is deployed, calling party number modification should be done on the Emergency Responder system.

- **Static ANI interface**

With this type of interface, the calling party number identification passed to the network is controlled by the PSTN. This is the case if the interface is a POTS line. The ELIN is the phone number of the POTS line, and no further manipulation of the phone's calling party identification number is possible.

PSAP Callback

The PSAP might have to reach the caller after completion of the initial conversation or if the caller hangs up before the PSAP operator answers the call. The PSAP's ability to call back relies on the information that it receives with the original incoming call.

The delivery of this information to the PSAP is a two-part process:

1. The Automatic Number Identification (ANI) is first sent to the PSAP. The ANI is the E.164 number used to route the call. In our context, the ANI received at the PSAP is the ELIN that the MLTS sent.
2. The PSAP then uses the ANI to query a database and retrieve the Automatic Location Identification (ALI). The ALI provides the PSAP attendant with information such as:
 - Calling company name
 - Physical address
 - Applicable public safety agency
 - Other optional information, which could include callback information. For example, the phone number of the enterprise's security service could be listed, to aid in the coordination of rescue efforts.

Typical Situation

- The ANI information is used for PSAP callback, which assumes that the ELINs are PSTN dialable numbers.
- The ELINs are PSTN numbers associated with the MLTS. If someone calls the ELIN from the PSTN, the call will terminate on an interface controlled by the MLTS.
- It is the responsibility of the MLTS system administrator to program the call routing so that calls made to any ELIN in the system will ring a phone (or multiple phones) in the immediate vicinity of the associated ERL.
- Once the ERL-to-ELIN mapping is established, it needs to be modified only when there are changes to the physical situation of the enterprise. If phones are simply added, moved, or deleted from the system, the ERL-to-ELIN mapping and its associated ANI/ALI database records need not be changed.

Exceptional Situation

- Callback to the immediate vicinity of the originating ERL may be combined with (or even superseded by) routing the callback to an on-site emergency desk, which will assist the PSAP in reaching the original caller and/or provide additional assistance with the emergency situation at hand.
- The situation of the enterprise could change, for example, due to area code splits, city or county service changes requiring a new distribution of the public safety responsibilities, new buildings being added, or any other change that would affect the desired routing of a call for 911 purposes. Any of these events could require changes in the ERL-to-ELIN mapping and the ANI/ALI database records for the enterprise.

Dial Plan Considerations

It is highly desirable to configure a dial plan so that the system easily recognizes emergency calls, irrespective of whether an access code (for example, 9) is used or not. The emergency string for North America is generally 911. Cisco strongly recommends that you configure the system to recognize both the strings 911 and 9911.

Cisco also strongly recommends that you explicitly mark the emergency route patterns with Urgent Priority so that Unified CM does *not* wait for the inter-digit timeout (Timer T.302) before routing the call.

Other emergency call strings may be supported concurrently on your system. Cisco highly recommends that you provide your system users with training on the selected emergency call strings.

Also, it is highly desirable that users be trained to react appropriately if they dial the emergency string by mistake. In North America, 911 may be dialed in error by users trying to access a long distance number through the use of 9 as an access code. In such a case, the user should remain on the line to confirm that there is no emergency, and therefore no need to dispatch emergency personnel. Cisco Emergency Responder's on-site notification capabilities can help in identifying the phone at the origin of such spurious 911 calls by providing detailed accounts of all calls made to 911, including calls made by mistake. If the emergency dispatch center cannot confirm that a call to 911 was accidental, then emergency services must be dispatched to the calling location. More than three emergency services dispatches to a single customer in a month often times will result in a fine to the company.

In a multisite deployment, the dial plan configuration should ensure that the emergency calls are always routed through the PSTN gateway local to the site, thereby making sure that the emergency call is routed to the nearest PSAP within the jurisdiction. One of the mechanism to achieve this could be to use the Local Route Group feature of Cisco Unified CM. In the case of multisite deployments with centralized PSTN access, local call routing to the PSAP is not possible. For deployments with centralized PSTN access, the Unified CM administrator must verify that the PSTN provider will route emergency calls to the proper PSAP based on ANI or ELIN. If the service provider cannot provide emergency call routing services for multiple sites, then any site not included in E911 coverage must have a location connection (an analog line) or the centralized PSTN access must support 911 call delivery for remote sites (a SIP trunk). (See the examples in the section on [Interface Type](#), page 15-7.)

Also, in a multisite deployment it is very important to make sure that the emergency number is always reachable and routed through the local PSTN gateway for the mobility users (extension mobility and device mobility) independent of the implemented Class of Service (CoS). If the site/device approach is being used, the device calling search space (CSS) could be used to route the emergency calls.

Cisco recommends enabling Calling Party Modification on Cisco Emergency Responder. When this feature is enabled, the calling party number is replaced with the ELIN by Cisco Emergency Responder for the emergency call. If Calling Party Modification is not enabled, either the DID will be sent to the PSAP or Cisco Unified CM must be configured to replace the calling party with the ELIN defined on the route pattern or the gateway.

Gateway Considerations

Consider the following factors when selecting the gateways to handle emergency calls for your system:

- [Gateway Placement, page 15-17](#)
- [Gateway Blocking, page 15-17](#)
- [Answer Supervision, page 15-18](#)
- [Answer Supervision, page 15-18](#)

Gateway Placement

Within the local exchange carrier (LEC) networks, 911 calls are routed over a locally significant infrastructure based on the origin of the call. The serving Class 5 switches are connected either directly to the relevant PSAP for their location or to an E911 selective router, which itself is connected to a group of PSAPs significant for its region.

With Cisco's IP-based enterprise communications architecture, it is possible to route calls on-net to gateways that are remotely situated. As an example, an endpoint located in San Francisco could have its calls carried over an IP network to a gateway situated in San Jose, and then sent to the LEC's network.

For 911 calls, it is critical to choose the egress point to the LEC network so that emergency calls are routed to the appropriate local PSAP. In the example above, a 911 call from the San Francisco endpoint, if routed to a San Jose gateway, could not reach the San Francisco PSAP because the San Jose LEC switch receiving the call does not have a link to the E911 selective router serving the San Francisco PSAP. Furthermore, the San Jose area 911 infrastructure would not be able to route the call based on a San Francisco calling party number.

As a general rule, route 911 calls to a gateway physically located with the originating endpoint. Contact the LEC to explore the possibility of using a common gateway to aggregate the 911 calls from multiple locations. Be aware that, even if the 911 network in a given region lends itself to using a centralized gateway for 911 calls, it might be preferable to rely on gateways located with the calling phones to prevent 911 call routing from being impacted during WAN failures.

Gateway Blocking

It is highly desirable to protect 911 calls from "all trunks busy" situations. If a 911 call needs to be connected, it should be allowed to proceed even if other types of calls are blocked due to lack of trunking resources. To provide for such situations, you can dedicate an explicit trunk group just for 911 calls.

It is acceptable to route emergency calls exclusively to an emergency trunk group. Another approach is to send emergency calls to the same trunk group as the regular PSTN calls (if the interface permits it), with an alternative path to a dedicated emergency trunk group. The latter approach allows for the most flexibility.

As an example, we can point emergency calls to a PRI trunk group, with an alternate path (reserved exclusively for emergency calls) to POTS lines for overflow conditions. If we put 2 POTS lines in the alternate trunk group, we are guaranteeing that a minimum of two simultaneous 911 calls can be routed, in addition to any calls that were allowed in the main trunk group.

If the preferred gateway becomes unavailable, it may be acceptable to overflow emergency calls to an alternate number so that an alternate gateway is used. For example, in North America calls dialed as 911 could overflow to an E.164 (non-911) local emergency number. This approach does not take advantage of the North American 911 network infrastructure (that is, there is no selective routing, ANI, or ALI services), and it should be used only if it is acceptable to the applicable public safety authorities and only as a last resort to avoid rejecting the emergency call due to a lack of network resources.

Answer Supervision

Under normal conditions, calls made to an emergency number should return answer supervision upon connection to the PSAP. The answer supervision may, as with any other call, trigger the full-duplex audio connection between the on-net caller and the egress interface to the LEC's network.

With some North American LECs, answer supervision might not be returned when a "free" call is placed. This may be the case for some toll-free numbers (for example, 800 numbers). In exceptional situations, because emergency calls are considered "free" calls, answer supervision might not be returned upon connection to the PSAP. You can detect this situation simply by making a 911 test call. Upon connection to the PSAP, if audio is present, the call timer should record the duration of the ongoing call; if the call timer is absent, it is very likely that answer supervision was not returned. If answer supervision is not returned, Cisco highly recommends that you contact the LEC and report this situation because it is most likely not the desired functionality.

If this situation cannot be rectified by the Local Exchange Carrier, it would be advisable to configure the egress gateway *not* to require answer supervision when calls are placed to the LEC's network, and to cut through the audio in both directions so that progress indicator tones, intercept messages, and communications with the PSAP are possible even if answer supervision is not returned.

By default, Cisco IOS-based H.323 gateways must receive answer supervision in order to connect audio in both directions. To forego the need for answer supervision on these gateways, use the following commands:

- **progress_ind alert enable 8**

This command provides the equivalent of receiving a progress indicator value of 8 (in-band information now available) when alerting is received. This command allows the POTS side of the gateway to connect audio toward the origin of the call.

- **voice rtp send-recv**

This command allows audio cut-through in both backward and forward directions before a connect message is received from the destination switch. This command affects all Voice over IP (VoIP) calls when it is enabled.

Be advised that, in situations where answer supervision is not provided, the call detail records (CDRs) will not accurately reflect the connect time or duration of 911 calls. This inaccuracy can impede the ability of a call reporting system to document the relevant statistics properly for 911 calls.

In all cases, Cisco highly recommends that you test 911 call functionality from all call paths and verify that answer supervision is returned upon connection to the PSAP.

Cisco Emergency Responder Design Considerations

Device mobility brings about special design considerations for emergency calls. Cisco Emergency Responder (Emergency Responder) can be used to track device mobility and to adapt the system's routing of emergency calls based on a device's dynamic physical location.

Device Mobility Across Call Admission Control Locations

In a centralized call processing deployment, Cisco Emergency Responder can detect Cisco endpoint relocation and reassign relocated endpoints to appropriate ERLs automatically. However, Cisco Unified CM location-based call admission control for a relocated endpoint might not properly account for the WAN bandwidth usage of the phone in the new location, yielding possible over-subscription or under-subscription of WAN bandwidth resources. For example, if you physically move a phone from Branch A to Branch B, the endpoint's call admission control location remains the same (Location_A), and it is possible that calls made to 911 from that endpoint would be blocked due to call admission control denial if all available bandwidth to Location_A is in use for other calls. To avoid such blocking of calls, manual intervention might be required to adapt the device's location and region parameters.

Cisco Unified CM device mobility provides a way to update the endpoint's configuration automatically (including its calling search space and location information) in Unified CM to reflect its new physical location. If device mobility is not used, manual configuration changes may be necessary in Cisco Unified CM.

For more details on the Device Mobility feature, refer to the section on [Device Mobility](#), page 21-15.

Default Emergency Response Location

If Cisco Emergency Responder cannot directly determine the physical location of an endpoint, it assigns a default emergency response location (ERL) to the call. The default ERL points all such calls to a specific PSAP. Although there is no universal recommendation as to where calls should be sent when this situation occurs, it is usually desirable to choose a PSAP that is centrally located and that offers the largest public safety jurisdiction. It is also advisable to populate the ALI records of the default ERL's emergency location identification numbers (ELINs) with contact information for the enterprise's emergency numbers and to offer information about the uncertainty of the caller's location. In addition, it is advisable to mark those ALI records with a note that a default routing of the emergency call has occurred. Alternatively, the call may be routed to an internal security office to determine the caller's location.

Cisco Emergency Responder and Location Awareness for Wireless Clients

Cisco Emergency Responder 11.5 and later releases can track wireless endpoints and clients to an access point in the enterprise. To minimize configuration changes in Cisco Emergency Responder, all access points must be synchronized from Cisco Unified Communications Manager. The synchronization process also handles any access point additions, updates, or removals that occur in Cisco Unified CM. Any access point changes in Unified CM are seen in Emergency Responder within 2 minutes of the change. Access points cannot be defined within Emergency Responder, and all access points that are to be used for location identification in Emergency Responder must be defined in Unified CM. For access point management, Unified CM uses the Cisco Wireless LAN Controller Synchronization Service to automatically synchronize access points into the Unified CM database. The Cisco Wireless LAN Controller Synchronization Service integrates with Cisco Wireless LAN Controllers (WLCs) for access point information. If another vendor is used for WLC services, then the access points must be bulk imported into the Cisco Unified CM database using the Bulk Administration Tool (BAT).

For a mobile client or wireless device to be associated to a wireless access point, the client or device must send the Basic Service Set Identifier (BSSID) of the associated access point to Cisco Unified CM. Due to the frequency of updates that a mobile client can generate, Unified CM limits the rate of location updates from mobile devices and wireless clients to 90 updates per second per node. If location updates exceed this rate for a sustained period of time, Unified CM defers further updates with a 480 "Busy Here" message. The client responds by waiting a period of time before sending the location update again. The amount of delay before sending the update again depends on the client and not on Cisco Emergency Responder or Unified CM.

When a mobile client or wireless device updates its location in Cisco Unified CM, the update is reflected in Cisco Emergency Responder in less than 2 minutes.

Cisco Emergency Responder and Extension Mobility

Cisco Emergency Responder supports Extension Mobility within a Cisco Unified CM cluster. It can also support Extension Mobility Cross-Cluster (EMCC), provided that both Cisco Unified CM clusters are supported either by a common Cisco Emergency Responder server or group, or by two Cisco Emergency Responder servers or groups configured as a Cisco Emergency Responder cluster. In either case, the Cisco Unified CM clusters must not be configured to use the Adjunct Calling Search Space (CSS) associated with EMCC for 911 calls, but must be configured to use Cisco Emergency Responder for all 911 calls in both Cisco Unified CM clusters.

Cisco Emergency Responder and Video

Cisco Emergency Responder can discover Cisco Video Collaboration endpoints in the following ways, depending on their capabilities:

- [Video Collaboration Endpoints that Support CDP, page 15-20](#)
- [Video Collaboration Endpoints that Do Not Support CDP, page 15-21](#)

Regardless of which way the video endpoints are discovered, it is important to note that video is not supported as media for emergency calling to the PSAP.



Note

The topics discussed in this chapter apply to Cisco Emergency Responder only when it is used in conjunction with Cisco Unified Communications Manager (Unified CM). Cisco TelePresence Video Communication Server (VCS) currently does not support emergency services.

Video Collaboration Endpoints that Support CDP

For video collaboration endpoints that support Cisco Discovery Protocol (CDP) and that are within the corporate premises, Cisco recommends treating them like any other collaboration endpoints tracked by Cisco Emergency Responder through CDP, as described by the Emergency Responder switch configuration information in the latest version of the *Cisco Emergency Responder Administration Guide*, available at

http://www.cisco.com/en/US/partner/products/sw/voicesw/ps842/prod_maintenance_guides_list.html

For video collaboration endpoints with CDP support that are outside the corporate premises, Cisco recommends treating them like voice collaboration endpoints as described in the information for off-premises support of IP phones in the latest version of the *Off-Premise Location Management User Guide for Cisco Emergency Responder*, available at

http://www.cisco.com/en/US/partner/products/sw/voicesw/ps842/products_user_guide_list.html

Video Collaboration Endpoints that Do Not Support CDP

For video collaboration endpoints that do not support Cisco Discovery Protocol (CDP), Cisco recommends using a dedicated line for a voice collaboration endpoint. If you require tracking of the video collaboration endpoint, then Cisco recommends configuring an IP subnet ERL as described in the information about setting up IP subnet-based ERLs found in the latest version of the *Cisco Emergency Responder Administration Guide*, available at

http://www.cisco.com/en/US/partner/products/sw/voicesw/ps842/prod_maintenance_guides_list.html

Cisco Emergency Responder and Off-Premises Endpoints

In cases where endpoints are located outside of the enterprise boundary but connect back to the enterprise using VPN or VPN-less solutions (for example, Cisco Expressway mobile and remote access from a home office or hotel), Cisco Emergency Responder will not be able to determine the location of the caller. Furthermore, it is unlikely that the system would have a gateway properly situated to allow sending the call to the appropriate PSAP for the caller's location.

It is a matter of enterprise policy to allow or not to allow the use of off-premises endpoints for 911 calls through the enterprise. It might be advisable to disallow 911 calls by policy for those endpoints that connect over the Internet through VPN or Cisco Expressway. Nevertheless, if such a user were to call 911, the best-effort system response would be to route the call to either an on-site security force or a large PSAP close to the system's main site.

The following paragraph is an example notice that you could issue to users to warn them that emergency call functionality is not guaranteed for off-premises endpoints and users:

Emergency calls should be placed from devices that are located at the site for which they are configured (for example, your office). A local safety authority might not answer an emergency call placed from a device that has been removed from its configured site. If you must use this device for emergency calls while away from your configured site, be prepared to provide the answering public safety authority with specific information regarding your location. Use a device that is locally configured to the site (for example, your hotel phone or your home phone) for emergency calls when traveling or telecommuting.

Cisco Emergency Responder also supports integration with Intrado V9-1-1, an emergency call delivery service that can reach almost any PSAP in the United States. With the combination of Cisco Emergency Responder and Intrado V9-1-1, users of IP phones and softphones outside the enterprise can update their locations by using the display screen on most Cisco IP Phones and Cisco IP Communicator or by using a web page provided by Cisco Emergency Responder. Emergency calls from an off-premises location will then be delivered through Cisco Emergency Responder to Intrado and then to the appropriate PSAP for the caller's location.

Test Calls

For any enterprise telephony system, it is a good idea to test 911 call functionality, not only after the initial installation, but regularly, as a preventive measure.

The following suggestions can help you carry out the testing:

- Contact the PSAP to ask for permission before doing any tests, and provide them with the contact information of the individuals making the tests.
- During each call, indicate that it is *not* an actual emergency, just a test.
- Confirm the ANI and ALI that the call taker has on their screen.

- Confirm the PSAP to which the call was routed.
- Confirm that answer supervision was received by looking at the call duration timer on the endpoint. An active call timer is an indication that answer supervision is working properly.

PSAP Callback to Shared Directory Numbers

Cisco Emergency Responder handles the routing of inbound calls made to emergency location identification numbers (ELINs). In cases where the line from which a 911 call was made is a shared directory number, the PSAP callback will cause all shared directory number appearances to ring. Any of the shared appearances can then answer the call, which means that it may not be the phone from which the 911 call originated.

In Cisco Unified CM 11.5 and later releases, a PSAP callback to a shared DN will ring only the device that placed the call to the PSAP. Unified CM will override device and line settings (such as Call Forward All and Do Not Disturb) to deliver the callback from emergency services.

Cisco Emergency Responder Deployment Models

Enterprise communications systems based on multiple Unified CM clusters can benefit from the functionality of Cisco Emergency Responder (Emergency Responder).

The *Cisco Emergency Responder Administration Guide* provides detailed descriptions of the terms used herein, as well as the background information required to support the following discussion. Of specific interest is the chapter on *Planning for Cisco Emergency Responder*. This documentation is available at

http://www.cisco.com/en/US/products/sw/voicesw/ps842/prod_maintenance_guides_list.html



Note


Cisco Emergency Responder does not support Cisco Unified Communications Manager Express (Unified CME) or Survivable Remote Site Telephony (SRST). In case of SRST deployment, configure the appropriate dial-peer to route the 911 calls to the PSTN with the published site number. Unified CME natively supports E911.

Single Cisco Emergency Responder Group

A single Emergency Responder group can be deployed to handle emergency calls from two or more Unified CM clusters. The design goal is to ensure that an emergency call from any phone is routed to the Cisco Emergency Responder group, which will assign an ELIN and route the call to the appropriate gateway based on the endpoint's location.

One advantage of using a single Cisco Emergency Responder group is that all ERLs and ELINs are configured into a single system. An endpoint registered on any cluster will be located by the single Cisco Emergency Responder group because that group is responsible for polling all of the system's access switches. [Figure 15-3](#) illustrates a single Cisco Emergency Responder group interfaced with two Unified CM clusters.

Figure 15-3 A Single Cisco Emergency Responder Group Connected to Two Unified CM Clusters



132073

The single Cisco Emergency Responder group in [Figure 15-3](#) interfaces with the following components:

- Each Unified CM cluster, via SNMP, to collect information about their respective configured endpoints.
- Enterprise access switches, via SNMP, where IP telephony endpoints are connected. This connection is not required if the endpoint locations are being identified based on IP subnets. For details on configuring IP subnet-based ERLs, refer to the *Cisco Emergency Responder Configuration* chapter in the *Cisco Emergency Responder Administration Guide*, available at http://www.cisco.com/en/US/products/sw/voicesw/ps842/prod_maintenance_guides_list.html
- Each Unified CM cluster, via JTAPI, to allow for the call processing required by any endpoint that dials 911 – for example, identification of the calling endpoint's ERL, assignment of the ELIN, redirection of the call to the proper gateway (based on the calling endpoint's location), and the handling of the PSAP callback functionality.
- Each Unified CM cluster, via SNMP, to collect access point information from a Cisco Wireless LAN Controller (WLC).

The version of the JTAPI interface used by Cisco Emergency Responder is determined by the version of the Unified CM software to which it is connected. At system initialization, Cisco Emergency Responder interrogates the Unified CM cluster and loads the appropriate JTAPI Telephony Service Provider (TSP). Because there can be only one version of JTAPI TSP on the Cisco Emergency Responder server, all Unified CM clusters to which a single Cisco Emergency Responder group is interfaced *must* run the same version of Unified CM software.

For some deployments, this software version requirement might present some difficulties. For instance, during a Unified CM upgrade, different clusters will be running different versions of software, and some of the clusters will be running a version of JTAPI that is not compatible with the version running on the Cisco Emergency Responder servers. When this situation occurs, emergency calls from the cluster running a version of JTAPI different than that of the Cisco Emergency Responder group might receive the call treatment provided by the call forward settings of the emergency number's CTI Route Point.

When considering if a single Cisco Emergency Responder group is appropriate for multiple Unified CM clusters, apply the following guidelines:

- Make Unified CM upgrades during an acceptable maintenance window when emergency call volumes are as low as possible (for example, after hours, when system use is at a minimum).
- Use a single Cisco Emergency Responder group only if the quantity and size of the clusters allow for minimizing the amount of time when dissimilar versions of JTAPI are in use during software upgrades.

For example, a deployment with one large eight-server cluster in parallel with a small two-server cluster could be considered for use with a single Cisco Emergency Responder group. In this case, it would be best to upgrade the large cluster first, thus minimizing the number of users (those served by the small cluster) that might be without Cisco Emergency Responder service during the maintenance window of the upgrade. Furthermore, the small cluster's users can more appropriately be served by the temporary static routing of emergency calls in effect while Cisco Emergency Responder is not reachable because they can be identified by the single ERL/ELIN assigned to all non-ER calls made during that time.

Multiple Cisco Emergency Responder Groups

Multiple Cisco Emergency Responder groups can also be deployed to support multi-cluster systems. In this case, each ER group interfaces with the following components:

- A Unified CM cluster via the following methods:
 - SNMP, to collect information about its configured endpoints
 - JTAPI, to allow for the call processing associated with redirection of the call to the proper gateway or, in the case of roaming endpoints, the proper Unified CM cluster
- The access switches (via SNMP) to which most of the endpoints associated with the Unified CM of the Cisco Emergency Responder group are most likely to be connected
- Each Unified CM cluster (via SNMP) to collect Access Point information from a Cisco Wireless LAN Controller (WLC)


This approach allows Unified CM clusters to run different versions of software because each is interfaced to a separate Cisco Emergency Responder group.

To allow endpoints to roam between various parts of the network and still be tracked by Cisco Emergency Responder, you might have to configure the Cisco Emergency Responder groups into a Cisco Emergency Responder cluster. For details on Cisco Emergency Responder clusters and groups, refer to the chapter on *Planning for Cisco Emergency Responder* in the *Cisco Emergency Responder Administration Guide*, available at

http://www.cisco.com/en/US/products/sw/voicesw/ps842/prod_maintenance_guides_list.html

Figure 15-4 presents a sample topology illustrating some of the basic concepts behind Cisco Emergency Responder clustering.

Figure 15-4 Multiple Cisco Emergency Responder Groups



132074

Figure 15-4 illustrates the following topology:

- Cisco Emergency Responder group A is interfaced to Unified CM cluster A to access switches A1 and A2, and it is deemed to be the home Cisco Emergency Responder group of all endpoints registered to Unified CM cluster A.
- Likewise, Cisco Emergency Responder group B is interfaced to Unified CM cluster B to access switches B1 and B2, and it is deemed to be the home Cisco Emergency Responder group of all endpoints registered to Unified CM cluster B.

Endpoint Movements Within the Tracking Domain of a Cisco Emergency Responder Group

The emergency call processing for endpoints moving between access switches controlled by the same home Cisco Emergency Responder group is the same as the processing done for a deployment with a single Unified CM cluster. For example, an endpoint moving between access switches A1 and A2 remains registered with Unified CM cluster A, and its location is determined by Cisco Emergency Responder group A both before and after the move. The endpoint is still under full control of Cisco Emergency Responder group A, for both the discovery of the endpoint by Unified CM cluster A and the determination of the endpoint's location on switch A2 by Cisco Emergency Responder. The endpoint is therefore not considered to be an unlocated phone.

Endpoint Movements Between the Various Tracking Domains of a Cisco Emergency Responder Cluster

A Cisco Emergency Responder cluster is essentially a collection of Cisco Emergency Responder groups that share location information. Each group shares the location of any endpoint it finds on an access switch or in an IP subnet.

Cisco Emergency Responder groups also share information about endpoints that cannot be located within a Cisco Emergency Responder group's tracking domain (in switches or IP subnets) but which are known to be registered in the group's associated Unified CM cluster. Such endpoints are deemed *unlocated*.

If an endpoint is roaming between access switches monitored by different Cisco Emergency Responder groups, those groups must be configured in a Cisco Emergency Responder cluster so they can exchange information about the endpoint's location. For example, endpoint A3 is registered with Unified CM cluster A, but it is connected to an access switch controlled by Cisco Emergency Responder group B. Cisco Emergency Responder group A is aware that endpoint A3 is registered with Unified CM cluster A, but group A cannot locate endpoint A3 in any of the site A switches. Therefore, endpoint A3 is deemed *unlocated* by Cisco Emergency Responder group A.

Cisco Emergency Responder group B, on the other hand, has detected the presence of endpoint A3 in one of the switches that it monitors. Because the endpoint is not registered with Unified CM cluster B, endpoint A3 is advertised through the Cisco Emergency Responder database as an *unknown* endpoint.

Because the two Cisco Emergency Responder groups are communicating through a replicated database table, they can determine that Cisco Emergency Responder group B's *unknown* endpoint A3 is the same as Cisco Emergency Responder group A's *unlocated* endpoint A3.

The Unlocated Phone page in Cisco Emergency Responder group A will display the endpoint's MAC address along with the remote Cisco Emergency Responder group (in this, case Cisco Emergency Responder group B).

Emergency Call Routing within a Cisco Emergency Responder Cluster

Cisco Emergency Responder clustering also relies on route patterns that allow emergency calls to be redirected between pairs consisting of a Unified CM cluster and a Cisco Emergency Responder. For more details, refer to the section on *Creating Route Patterns for Inter-Cisco Emergency Responder Group Communications* in the *Cisco Emergency Responder Administration Guide*, available at

http://www.cisco.com/en/US/products/sw/voicesw/ps842/prod_maintenance_guides_list.html

If endpoint A3 places an emergency call, the call signaling flow will be as follows:

1. Endpoint A3 sends the emergency call string to Unified CM cluster A for processing.
2. Unified CM cluster A sends the call to Cisco Emergency Responder group A for redirection.
3. Cisco Emergency Responder group A determines that endpoint A3 is located in Cisco Emergency Responder group B's tracking domain, so it redirects the call to a route pattern that points to Unified CM cluster B.
4. Unified CM cluster A sends the call to Unified CM cluster B over a SIP trunk or an intercluster trunk.
5. Unified CM cluster B sends the call to Cisco Emergency Responder group B for redirection.
6. Cisco Emergency Responder group B identifies the ERL and ELIN associated with endpoint A3's location (based on calling party number) and redirects the call to Unified CM cluster B. The calling number is transformed into the ELIN associated with the ERL of endpoint A3, and the called number is modified to route the call to the proper gateway.

7. Unified CM cluster B routes the call according to the new called number information obtained from Cisco Emergency Responder group B.
8. Unified CM cluster B sends the call out the gateway toward the Emergency PSTN network.

**Note**

The ERL and ELIN match in step 6 is based on the calling party number of the endpoint placing the call to 911. If the SIP trunk or intercluster trunk modifies the calling party number (perhaps to a full +E.164 number), then the Emergency Responder for Group B will not be able to match the calling party number over the trunk with the directory number learned from the access switch Cisco Discovery Protocol (CDP) neighbor. Therefore, emergency calls that traverse a SIP trunk or intercluster trunk must not undergo any calling party transformations.

WAN Deployment of Cisco Emergency Responder

Cisco Emergency Responder supports two main sites using clustering over the WAN. Install one Emergency Responder server in each site, and configure one server as the publisher and the other server as a subscriber. The Emergency Responder publisher should be located with the primary Unified CM CTI Manager, and the Emergency Responder subscriber should be located with the secondary Unified CM CTI Manager. Any Unified CM server remote from either Emergency Responder server must be within 80 ms round-trip time (RTT) of both Emergency Responder servers. The Emergency Responder publisher and subscriber must also be within 80 ms RTT of each other. The minimum bandwidth required between the Cisco Emergency Responder servers is 1.544 Mbps.

Emergency Call Routing Using Unified CM Native Emergency Call Routing

Customers that require accurate location identification but have a single site or small number of locations that need to be identified, can use the Cisco Unified Communications Manager Native Emergency Call Routing feature. The Native Emergency Call Routing feature allows an administrator to define Emergency Location Identification Numbers (ELINs) at the device pool level or device level so that a device's location can be determined and identified at the public safety answering point (PSAP).

Cisco Unified CM Native Emergency Call Routing provides the following functionality:

- ELIN association based on a static device assignment or device pool assignment
- Dynamic association of the ELIN to the calling phone for callback purposes
- For mobile devices, Device Mobility Groups used to track mobile devices with Native Emergency Call Routing
- Automatic replacement of the calling party number with the appropriate ELIN
- Routing emergency calls to the appropriate gateway for emergency call completion

Design Considerations for 911 Native Emergency Call Routing Services

When designing an emergency call routing plan using Cisco Unified CM Native Emergency Call Routing services, give special consideration to the boundaries of an emergency location inside a building. An emergency location should be an identifiable location with physical or logical boundaries to reduce the amount of time for emergency services to locate an individual in an emergency situation. Examples of physical or logical boundaries can include: a single floor of a building, a lab, an office, or a directional floor indicator (for example, West side of first floor).

The design for Native Emergency Call Routing requires an ELIN to be defined and assigned to devices or device pools, but the Native Emergency Call Routing feature does not allow the administrator to define the ERL information to be associated with the ELIN. The ERL definition for a given ELIN must be done outside of Cisco Unified CM and uploaded to the local PSAP per the instructions provided by the local exchange carrier when establishing E911 services.

Similar to a Cisco Emergency Responder deployment, Native Emergency Call Routing can support multiple unique and concurrent calls to emergency services from the same location. Native Emergency Call Routing allows the creation of a pool of ELINs that are associated with an emergency location. The number of locations that can be defined is based on the number of ELINs assigned to an individual Emergency Location (ELIN) Group. Native Emergency Call Routing supports a maximum of 100 ELINs. If the deployment requires only one concurrent call per a location, then the system can support 100 unique Emergency Location Groups. If the deployment requires the ability to track 2 concurrent callers from the same location, then the administrator must define 2 ELINs for a single Emergency Location (ELIN) Group. If 2 ELINs are required for a single location, Unified CM will be able to support 50 locations ($2 \text{ ELINs} * 50 \text{ ERLs} = 100 \text{ ELINs}$). Using more ELINs to support concurrent and uniquely identified callers from a location will reduce the total number of locations that can be defined. The following formula can be used to determine the maximum number of locations that can be defined based on the number of concurrent and unique callers from an ERL:

$$100/(\text{Number of unique and concurrent callers per ERL}) = \text{Max ERLs}$$

ELINs are not required to be the same for each Emergency Location (ELIN) Group. If one ERL covers a high-density user population, the Emergency Location (ELIN) Group may contain 4 ELINs to support 4 concurrent and unique emergency callers. But if the same building has a large lab floor or warehouse that has a small number of regular employees, then that location might have only one ELIN assigned to the Emergency Location (ELIN) Group.

If the PSAP needs to call back and get additional information from the caller, the call will return to Unified CM using the ELIN that originated the call. To route the return call correctly, the dial plan must be configured so that the inbound called number matches the ELIN defined in Unified CM. If the inbound trunk delivers only the last 5 digits of the called party, then the administrator must include a translation pattern to expand the collected digits to match the ELIN. For proper return call operation, the called number must match exactly the ELIN number as defined in Unified CM. Although ELINs can be any number in a customer's DID range, Cisco recommends keeping the ELIN numbers contiguous to use as few call translation patterns as possible.

ALI Formats

In multi-cluster configurations, there might be instances where the physical locations of ERLs and ELINs defined in a single Cisco Emergency Responder group span the territory of more than one phone company. This condition can lead to situations where records destined for different phone companies have to be extracted from a common file that contains records for multiple LECs.

Cisco Emergency Responder exports this information in ALI records that conform to National Emergency Number Association (NENA) 2.0, 2.1, and 3.0 formats. However, many service providers do not use NENA standards. In such cases, you can use the ALI Formatting Tool (AFT) to modify the ALI records generated by Cisco Emergency Responder so that they conform to the formats specified by the service provider. The service provider can then use the reformatted file to update their ALI database.

The ALI Formatting Tool (AFT) enables you to perform the following functions:

- Select a record and update the values of the ALI fields. AFT allows you to edit the ALI fields to customize them to meet the requirements of various service providers. The service provider can then read the reformatted ALI files and use them to update their ELIN records.
- Perform bulk updates on multiple ALI records. Using the bulk update feature, you can apply common changes to all the records that you have selected.
- Selectively export ALI records based on area code, city code, or a four-digit directory number. By selecting to export all the ALI records in an area code, for example, you can quickly access all the ELIN records for each service provider, thereby easily supporting multiple service providers.

Given the flexibility of the AFT, a single Cisco Emergency Responder group can export ALI records in multiple ALI database formats. For a Cisco Emergency Responder group serving a Unified CM cluster with sites in the territories of two LECs, the basic approach is as follows:

1. Obtain an ALI record file output from Cisco Emergency Responder in standard NENA format. This file contains the records destined for multiple LECs.
2. Make a copy of the original file for each required ALI format (one copy per LEC).
3. Using the AFT of the first LEC (for example, LEC-A), load a copy of the NENA-formatted file and delete the records of all the ELINs associated with the other LECs. The information to delete can usually be identified by NPA (or area code).
4. Save the resulting file in the required ALI format for LEC-A, and name the file accordingly.
5. Repeat steps 3 and 4 for each LEC.

For more information about the ALI formatting tools, refer to the online documentation available at

http://www.cisco.com/en/US/products/sw/voicesw/ps842/prod_maintenance_guides_list.html

For LECs not listed at this URL, the output from Emergency Responder can be formatted using standard text file editing tools, such as spreadsheet programs and standard text editors.

