

HP VoIP solution for service providers

white paper

- Introduction2
- Target voice.....4
- VoIP networks4
- OCSC components6
- Deployment in operation....9
- Edge domain12
- Core Class IV domain13
- Class V domain14
- Conclusion16
- Appendix.....17


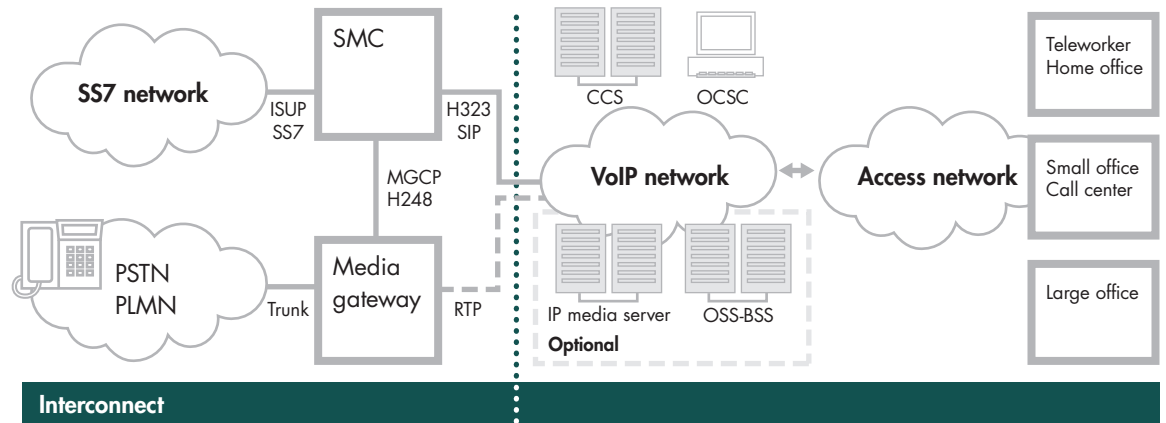


Figure 1
Network architecture



Introduction

This document describes the solution that HP, through a successful partnership with NetCentrex, offers for the deployment of telephony services over a broadband network using Internet Protocol (IP) technology.

The components of the solution are

- CCS (Call Control Server) softswitch—performs advanced Call Control and Switching functions and provides enhanced Voice over IP (VoIP) revenue-generating services
- OCSC (OpenCall Service Controller)—allows the introduction of value-added services on top of the softswitch. OCSC serves as a central point for service deployment and hosts the central repository and the service provisioning
- SMC (Signaling and Media Controller)—interconnects the existing Public Network (i.e., PSTN) and the VoIP network at a Call Control–Signaling level and controls a set of Media Gateways responsible for the interconnection at a media and voice trunk level

Optionally HP IP Media Server can be offered to perform IVR (interactive voice response) or IP voice mail functionalities.

The scope of the document is to provide service providers with technical information on the network architecture.

Requirements and long-term architecture

Implementing and deploying a complete VoIP network generally requires the following actions.

- Building the core infrastructure:
 - Enable voice on the IP infrastructure by deploying voice-enabled gateways, ensuring voice quality and security
 - Interface to the PSTN (Public Switched Telephone Network) or partner VoIP networks at various locations; each PSTN location may have its own signaling flavor, and VoIP networks provide fewer standardization issues but more security risks
 - Decide on a hierarchical routing plan and numbering convention inside the network
 - Predefine automatic recovery call flows to resist internal network failures
 - Predefine call rerouting rules to recover from call completion failures in the PSTN or other partner's networks
 - Decide on a disaster recovery strategy
 - link the network accounting records to a rating and billing engine, in batch mode or real time, and develop recovery strategies in case the billing system fails
 - Comply with the legal requirements that have an impact at the infrastructure level, such as ensuring caller ID is managed properly and Local Number Portability (LNP) has an impact on the routing plans
 - Provide network announcements

- Complying with national regulations:
 - Calling-line identification restriction
 - Malicious call identification
 - Legal interception
 - Carrier selection
 - Local Number Portability
 - Emergency services (i.e., 911 in the US; 112 in Europe)

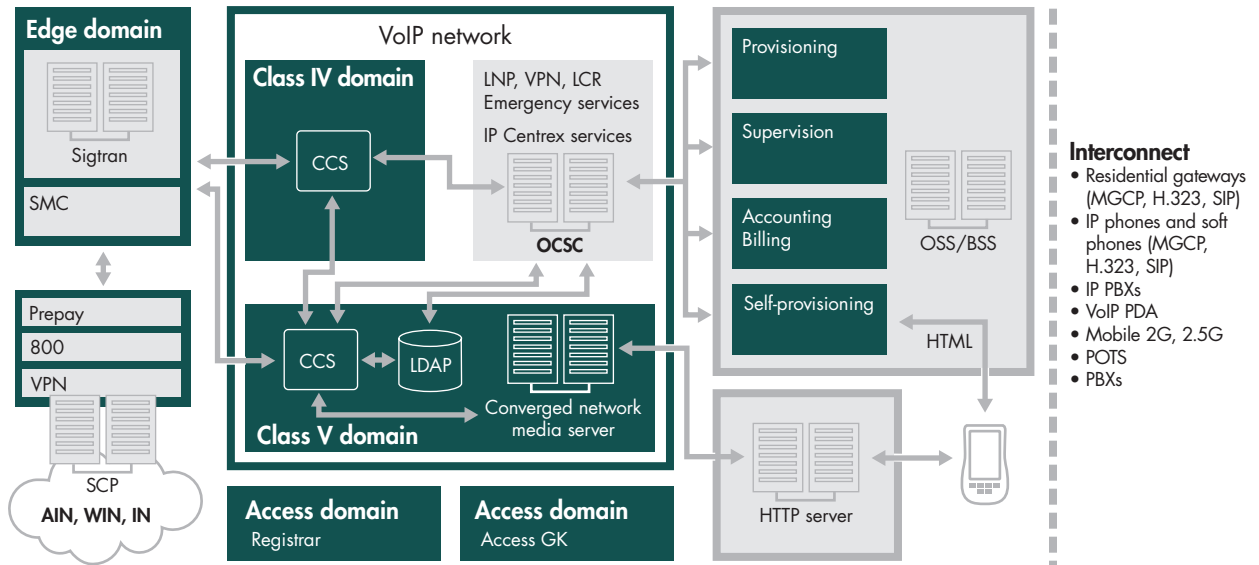
The scope of the document is to provide service providers with technical information on the network architecture.

- Class IV services:
 - On-net-to-On-net calls (calls originate and terminate on the IP voice network)
 - On-net-to-Off-net calls (calls originate on the IP voice network and terminate on PSTN)
 - Off-net-to-On-net calls (calls originate on PSTN and terminate on the IP voice network)
 - Management of dial plans (inter-area code, international, and inter-carrier)
 - Call routing, load balancing, and fault recovery
 - Network announcements
- Residential Class V services:
 - Call Forward on Busy
 - Call Forward on No Answer
 - Call Forward Unconditional
 - Call Hold
 - Call Waiting
 - CLIP (Calling Line Identification Presentation)
 - Code Restriction services
 - Denied Origination
 - Denied Termination
 - Feature activation
 - Find Me and Follow Me services
 - Three-way calling

- Applications interface:
 - Class V softswitches for residential telephony applications or Centrex
 - Applications for the corporate market, such as network Interactive Voice Response (IVR) units (Voice Extensible Markup Language [Voice XML]), Virtual Call Centers in the network, voice VPN, and videoconferencing
 - Third-party application service providers
- VoIP solution management:
 - Integrate VoIP components in the existing OSS and BSS or provide standalone network operation capabilities: FCAPS, billing
 - Enable user access and control to some data (i.e., self-provisioning of user profiles)

In order to ensure unlimited scalability and to cover any geographic region, the functions listed above need to be distributed over several servers in the network. This also allows the service provider to delegate regional network management to several regional administrators.

Figure 2
IP telephony network architecture



Target voice architecture

The target network is built on a layered architecture, including:

- An edge domain—responsible for peering relationships with all third-party networks, including PSTN and other VoIP networks that may be used to terminate calls
- A core Class IV domain—responsible for routing calls between all other domains and Local Number Portability (LNP) and emergency calls, as required
- A set of application domains—responsible for more advanced features, such as Class V call control
- A set of access domains—responsible for the direct management of Customer Provided Equipment (CPE), including IP address and alias registrations; this domain is most useful in the context of a Class V telephony service and is not needed for other, simpler services, such as voice VPNs

The core product of the network is the CCS VoIP softswitch as illustrated in figure 2 above.

An IN approach to VoIP networks

When services are deployed within a legacy telephony network, they can take advantage of the network architecture. The latter is typically an Intelligent Network (IN) architecture in which there are basically two ways of implementing services:

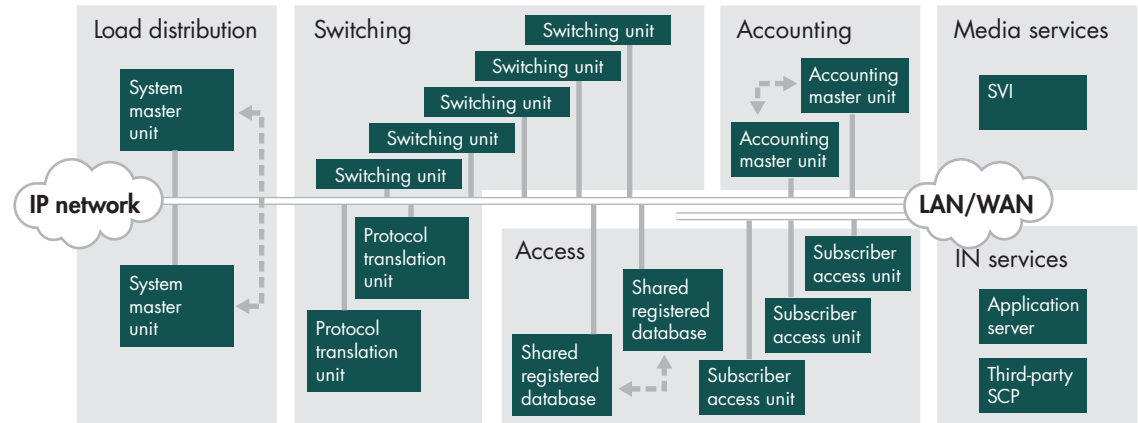
- A network service platform called the “Service Node.” This element has a few interactions with the rest of the IN. The IN routes the call to the Service Node, which takes the call, processes it, and typically issues an out-bound call to the final call destination. The service is completely hosted by the Service Node.
- A kind of “Service Logic” of the Service Control Point (SCP). This service is hosted partly by the SCP and partly by equipment called the “Intelligent Peripheral,” which is mainly in charge of voice processing.

Both options have their assets and drawbacks. They illustrate how regular telephony networks provide frameworks for deploying value-added services.

The world of IP telephony is on the move and is described by protocols H.323, SIP (Session Initiation Protocol), MGCP (Media Gateway Control Protocol), upcoming Megaco (Media Gateway Control), and so on, but there is no architecture on top to implement services, since they exist in the PSTN world.

It would be highly valuable to take all the assets of IN architecture and combine them with the openness, flexibility and mobility features of the IP world. This is the objective with the proposed architecture and its combination of the CCS and OCSC (OpenCall Service Controller) products—to provide the VoIP networks with a true, carrier-grade call-control architecture, making it possible to implement value-added services.

Figure 3
CCS distributed design



Without going into too much detail about CCS and OCSC, the following comparisons can be made:

- Call Control Server (CCS), which is composed of an H.323v2 gatekeeper, and a SIP Proxy Server, to which an Intelligent Network state machine is implemented (Q.1224), carry out the equivalent of the Service Switching Function (SSF).
- The equivalent of the Service Control Function (SCF) is carried out by OpenCall Service Controller (OCSC), which services the creation environment (call handling, media, contact centers, and so on).

CCS has Detection Points (DPs), which makes it possible to trigger specific processing events during the different steps of a call (connecting, busy, no answer, and so on) or upon the reception of H.323- and SIP-specific events. Services running on OCSC control the CCS Detection Points by activating and deactivating specific events either statically or dynamically during the progression of a call.

Thanks to the OCSC and CCS interface, services are capable of controlling and modifying the IP telephony signaling “on the fly,” making it possible to implement a number of advanced features.

CCS capabilities are critical since they implement call-switching functions, as well as sort and sequence the call progression phases.

With this carrier-grade VoIP architecture implementing an IN call model and providing Class IV and Class V capabilities, it’s possible to achieve a “happy marriage” by combining the best of the Intelligent Network philosophy with the best of the IP world.

CCS components

CCS can support multiple VoIP protocols, mediate between VoIP protocols, and enable protocol-agnostic applications that leverage the ITU standard Intelligent Network triggers.

The CCS softswitch

- Is scalable to system demands and has the capacity to support a linear increase in simultaneous calls and calls per second by adding Switching Units (SUs)
- Supports all the carrier-class redundancy and fault-tolerance features required by service providers

System Master Unit

The System Master Unit (SMU) is instrumental in achieving true carrier-class reliability and fault tolerance. All calls are directed to the SMU and then distributed to available CCS Switching Units (SUs). Each SMU centralizes the SNMP capabilities, makes sophisticated call-load distribution algorithms available to the SUs, and performs proactive SU monitoring to ensure uninterrupted call completion. Automatic failover to a backup SMU assures seamless carrier-class, fault-tolerant operation.

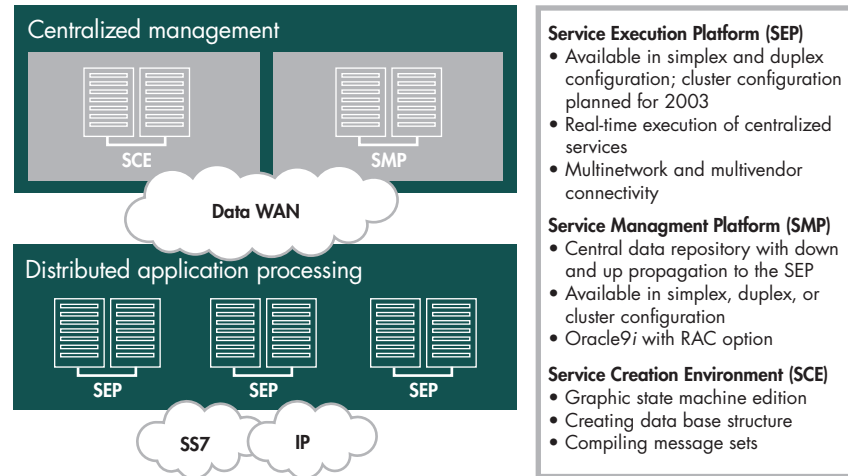
Switching Unit

The CCS Softswitch Switching Unit (SU) handles core call control and switching technology. The Intelligent Network (IN)-based design enables the creation of VoIP-independent services that use an IN-state machine and Q.1224 triggers. Thus, both the core softswitch services and applications are available to all supported VoIP protocols (H.323, SIP, MGCP), depending on selected protocol options, endpoints, and networks.

The SU also includes a comprehensive Class IV call routing engine, configurable through a telnet interface. The routing engine enables a two-level dial plan translation:

- From the source format to the “pivot” format, which will be used by the accounting system
- From the “pivot” format to the destination format

Figure 4
OpenCall Service Controller Components Overview



In addition, the SU also provides integration with OpenCall Service Controller.

In small trial deployment, the SU also includes the Subscriber Policy Engine, which implements the IETF (Internet Engineering Task Force) Call Processing Language (CPL).

Shared Registration Database Unit

The Shared Registration Database (SRD) Unit handles CPE registration by storing all registered endpoints in a memory-resident database. Automatic failover to a backup SRD ensures seamless carrier-class, fault-tolerant operation. The SRD is used only for limited deployments. For large-scale deployments, a distributed architecture with an external access domain, as described above, is recommended.

Subscriber Access Unit

The Subscriber Access Unit (SAU) provides multiline features (i.e., call hold, three-way conferencing, and so on) to residential gateways or IP phones that use the Media Gateway Control Protocol (MGCP). While H.323 or SIP endpoints handle multiline features internally, MGCP/L endpoints are much simpler and require the use of SAU, which provides the following features to MGCP/L endpoints: call hold, call retrieve, call transfer, call waiting, dial tone, call return, last number redial, call hold voice prompt, and three-way conferencing. The SAU can be used in residential services, as well as IP Centrex offerings, to provide class V multiline features.

Note: The Accounting Master Unit is not used in this architecture since CDRs (call data records) are being generated by OCSC (OpenCall) and not by CCS.

Value-added services can be implemented in two non-exclusive ways:

1. In the Media Server—using the Service Creation Environment (SCE) that makes it easy to implement and operate Media Services, thanks to a completely graphic environment
2. In an Intelligent Network—triggering Service Logic by specific CCS detection points; the Service Logic is located on a third-party SCP and may involve media resources that act as a VoIP Intelligent Peripheral, featuring INAP CS1, CAMEL, and WIN-IS41 protocols

OCSC components

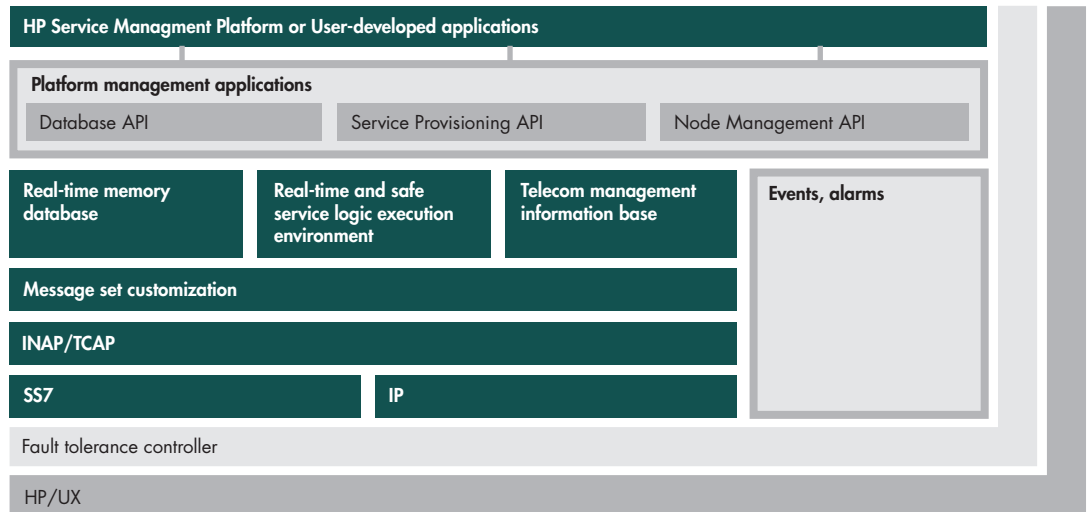
As shown in figure 4 above, the HP OpenCall Service Controller is composed of three components:

1. OpenCall SEP—Service Execution Platform
2. OpenCall SMP—Service Management Platform
3. OpenCall SCE—Service Creation Platform (optional)

HP OpenCall SEP

OpenCall Service Execution Platform (SEP) provides the capability to introduce, execute, control, and manage customer-built services quickly, efficiently, and economically within the converging telephony and IP networks. The SEP is an environment for running service applications and an in-memory database, providing real-time access to the customer information used by the services. The SEP architecture is based on active and standby mechanisms. The active SEP handles all the calls, and when it fails, the standby SEP is activated.

Figure 5
OpenCall SEP software architecture



The OpenCall SEP serves as a central point for service execution and is invoked systematically by the CCS through an API (Application Programming Interface) whenever a call is presented to the softswitch.

The services running on SEP are responsible for call routing within the VoIP network, management of numbering plans, and CDR generation.

OpenCall SEP software components

OpenCall SEP software includes the following components:

- Service Logic Execution Environment (SLEE)—provides the execution environment for Service Logic Programs (SLPs). The SLEE also incorporates the SEP database and associated management functions. SLEE provides the environment and resources to the run advanced call-routing functions necessary for VPN-IP Centrex services.
- Fault Tolerance Controller (FTC)—monitors and controls the duplicated processes running on each platform host and controls automatic switchover
- Event Handler—controls the display and logging of messages and alarms
- The CCS-Proxy—implementing INAP over IP—controls and manages the communication with the CCS Softswitch

OpenCall SEP transaction processing components perform the following actions:

- The SEP real-time database stores service information used by the services.

- For High Availability (HA) of the platform, the SEP transaction processing components are duplicated. SEP contains a high-performance, in-memory relational database for storing subscriber and service-related information. The information can be read and updated by the services running on SEP. The database is replicated in the HA SEP configuration.

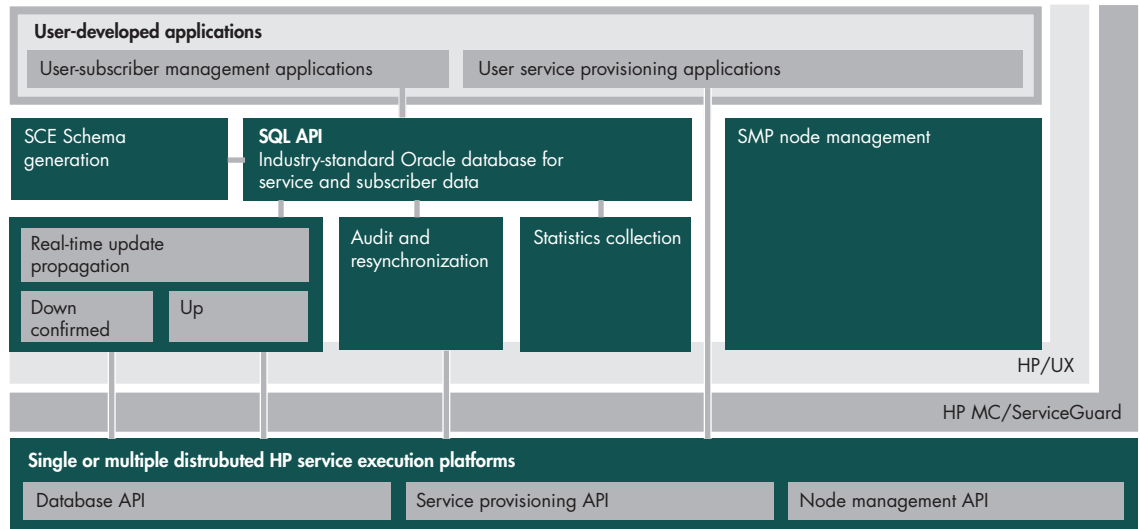
The database is optimized for real-time read/write access, including writes replicated to the standby processor for the HA configuration. The database consists of a series of linked tables that can be freely defined by the Service Designer. It is possible to configure which data is local to the active processor, which data is replicated to the standby processor, and which data is replicated to the optional Service Management Platform (SMP), if applicable. Applications running in SLEE have complete access to the database for read, write, navigate, and other operations.

SMP accesses the SEP database through a published management API. For the HA configuration, the database in the active SLEE is replicated in the standby SLEE. The associated files on disk in the active SEP host are also replicated on disk in the standby SEP host.

SEP-replicated hardware

In an HA SEP, the hardware consists of two computers, active and standby, which are connected by a duplicated LAN. The standby computer is a hardware replicate of the active computer. For CCS softswitch connectivity, each host contains two dedicated LAN cards, ready to take over.

Figure 6
HP OpenCall SMP software architecture



In normal mode, the active system controls all network connections. The standby system takes control of all connections during a failure of any critical hardware or software in the active system. This process is called switchover, where the standby system becomes the new active system and is ready to process network messages.

HP OpenCall Service Controller SMP

The HP OpenCall Service Controller Service Management Platform (OCSC SMP) can synchronize and manage subscriber data through an industry-standard SQL interface and an Oracle® Relational Database Management System. This is achieved even among geographically distributed OCSC SEPs, CCSs, and SAUs.

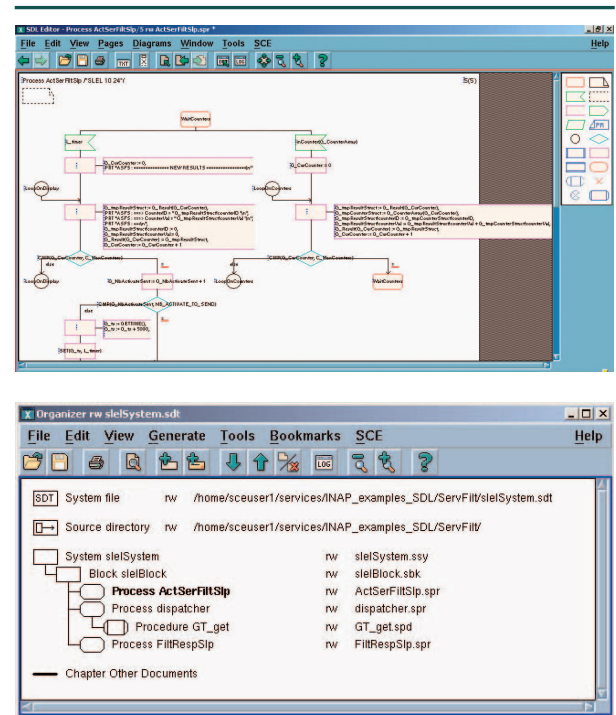
Centralized SMP enables customer care applications to access and utilize a rich set of subscriber information more easily and cost-effectively, supporting the development of differentiated customer care services.

OCSC SMP provides a central repository for network topology, service, and customer data, serving as a focal point for service deployment and data management across the network. The SMP is built around an Oracle database, providing open access to service and customer data through multilevel APIs (e.g., Java™ data access objects and Oracle stored procedures).

The HP OCSC platform provides sophisticated propagation mechanisms between the SMP Oracle database and the SEP real-time database, allowing instant and seamless availability of provisioned data to the service. As soon as the provisioning manager submits updates, service configuration changes are taken into account by the service.

HP OCSC also stores CDRs before their collection by reporting and billing systems.

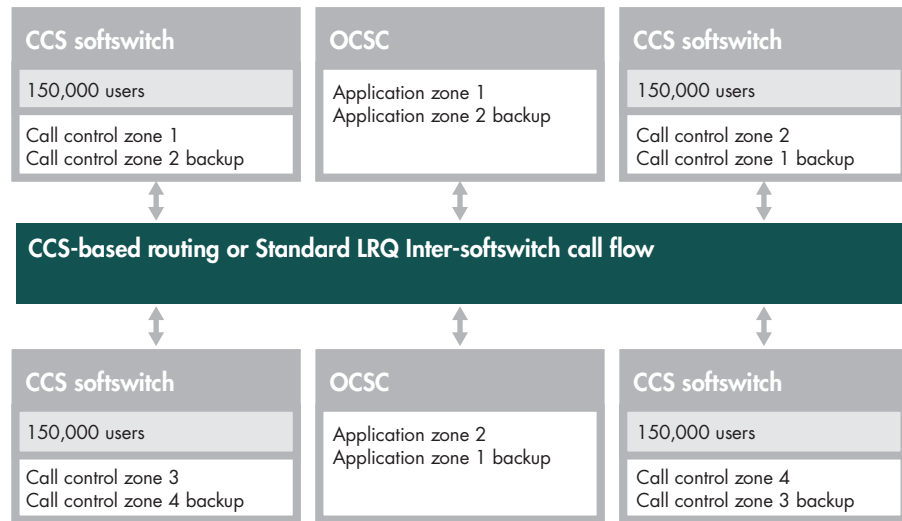
Figure 7
HP OpenCall service logic creation



HP OCSC SCE

HP OpenCall provides a powerful Service Creation Environment (SCE) for future service expansions and service customization with a GUI-based environment that automatically generates state machines and execution code that will be running in the SEP with the real-time database.

Figure 8
Distributed Architecture



Deployment in operation

High availability

When going into operation, it is critical for any telecommunications company to have a High Availability (HA) telephony system in order to provide the carrier-grade level of service that these clients expect.

With the High Availability option:

- There is no single point of failure in the system.
- It becomes possible to modify system hardware and software without stopping operations.
- Network wiring and connectivity is fully redundant (also requires redundant routers in the service provider network, using HSRP [Hot Standby Routing Protocol] or a similar health-check real-time protocol, and appropriate routing protocols).

CCS high availability

The CCS platform features an N+1 redundant policy. Making CCS high availability requires:

- One additional SMU
- One additional SU
- One additional SRD (optional—trial configuration)
- One additional SAU (optional—residential or IP Centrex)
- One additional server for the Media Server (optional)

OCSC high availability

The OCSC features an active/standby redundant policy for both the Service Execution Platform and the Service

Management Platform. To move from a simplex to duplex configuration, thus making OCSC high availability, requires:

- One additional SEP server
- One additional SMP server (optional)

High availability can only be deployed on the SEP, which is the real-time service engine, while SMP remains in simplex. For a full HA, both platforms require an additional server.

Disaster recovery

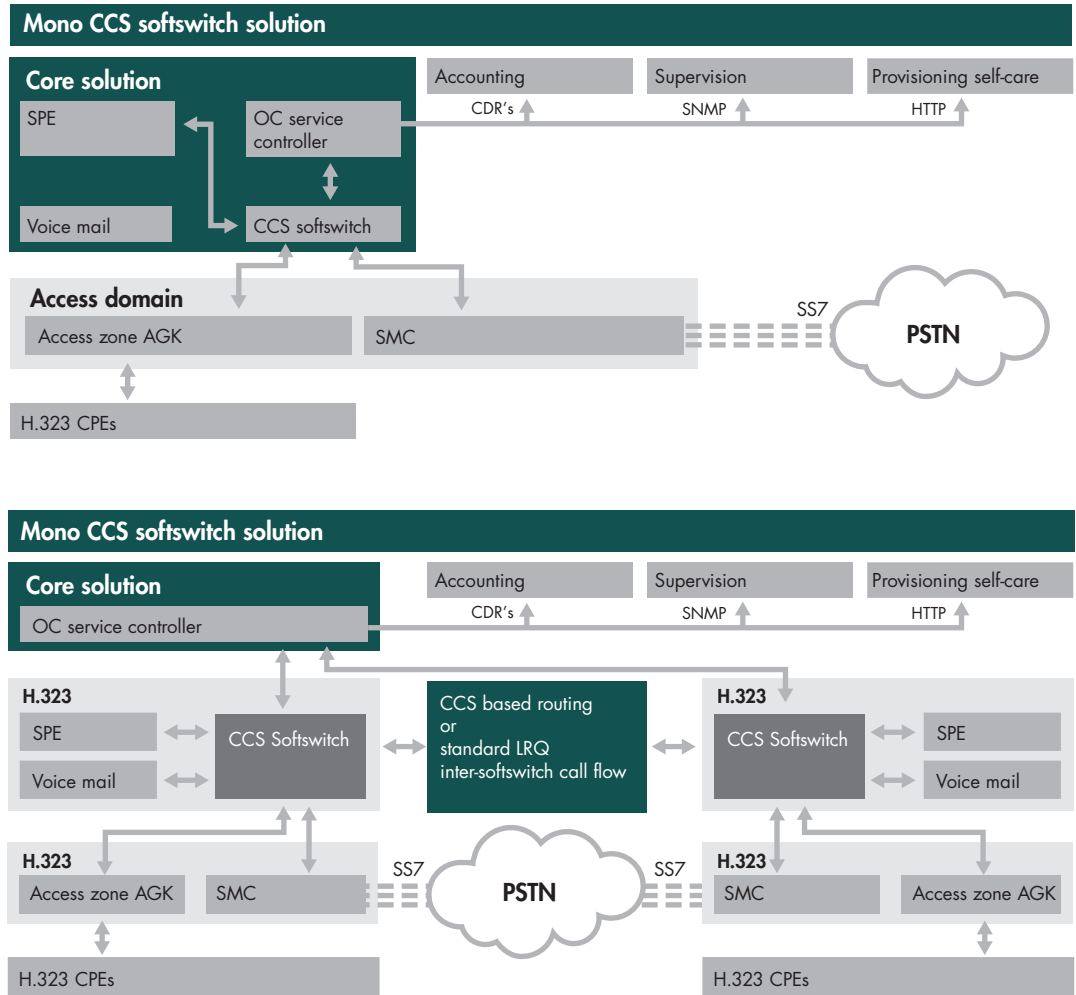
A network of CCS softswitches and OCSC application servers can resist the complete destruction of a given system (or complete network isolation of a given system)—i.e., a CCS, an OCSC, or both. Each edge device needs to be configured with a primary and backup CCS softswitch, and the routing tables of CCS softswitches in the network need to have two routes or more defined for each destination. Similarly, each CCS softswitch needs to be configured with a primary and backup OCSC application server.

Call processing capacity increase

There are two models for scaling a software-based function.

- Operating system–based multiprocessor scaling—writing the code in a multithreaded form allows the operating system to use multiple processors to increase system capacity. This approach works well for medium-scale configurations, but its cost increases exponentially with the number of processors, as the additional processing power made available to the application for each new processor decreases steadily with the number of processors already in place. There are also some operating systems with this approach.

Figure 9
Distributed Architecture hardware impact



- Distributed programming—this is the approach used by the CCS softswitch. The CCS code is not only multi-threaded but also distributed on several physical machines, each with its own operating system. The system is linearly scalable, i.e., each additional machine allows the system to gain a fixed capacity, regardless of the number of systems already in place.

A single CCS configuration can grow by adding switching units to an aggregate capacity of 1.8 million Busy Hour Call Attempts (BHCA) and 15,000 simultaneous calls.

Call processing capacity increase means two things:

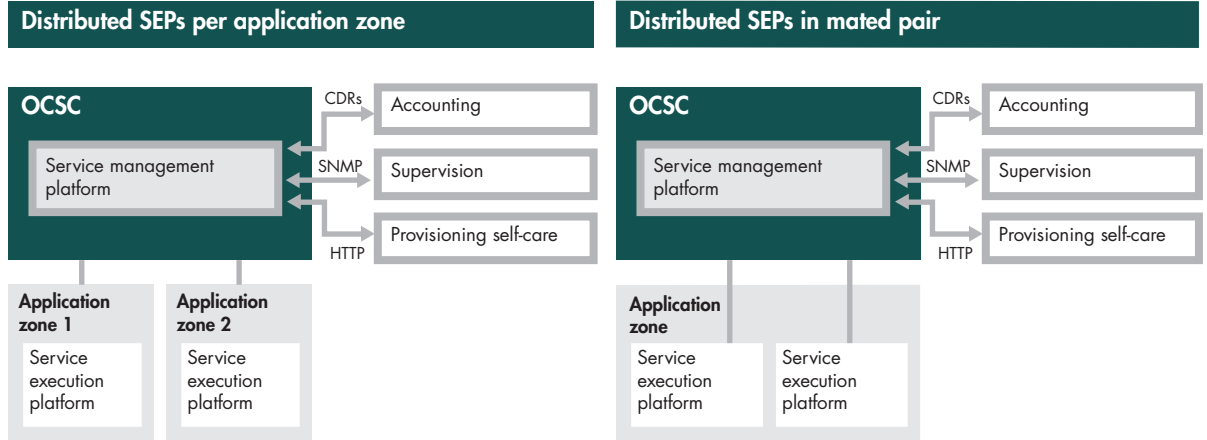
1. Increasing the capacity of the CCS, which means increasing the number of SUs (Switching Units) and the number of corresponding servers for the Media Server

(network announcement servers and “simplified voice mail” for which the required capacity is a function of the number of total calls likely to be handled by the system—1 to 10 typically) to reach the maximum call processing capacity of 15,000 calls (the maximum per CCS)

2. Increasing the capacity beyond 15,000 calls and/or taking into account geographic constraints (see the “Distributed Architecture” discussion below)

The solution architecture presented describes a centralized system using only a single softswitch and application server resource. This solution has some limitations and drawbacks that can be solved by using a distributed architecture using multiple CCS softswitches and OCSC application servers.

Figure 10
OCSC Distributed Configuration



Distributed Architecture

Using Distributed Architecture offers the following advantages:

- Pass-by CCS softswitch limits: One CCS softswitch can manage a maximum of 15,000 simultaneous calls. One SU can manage a maximum of
 - 10 to 20 calls per second (depending on options)
 - 300 to 1,000 simultaneous calls (depending on options)
- One CCS softswitch can handle a maximum of
 - 300 calls per second
 - 15,000 simultaneous calls
 - 150,000 users (average usage)

The Distributed Architecture shares the system load of multiple CCS softswitches

- Consequently, the platform capacity can be raised.
- Similarly, it can pass by the OCSC application server limits, adding more application servers when more capacity is required at the application layer.
- From a geographic point of view, Distributed Architecture prevents a single point of failure. In fact, if a location loses IP connectivity, calls can be routed to an alternate location. By taking full advantage of the call-routing module, multiple backup locations can be defined in order to provide full telephony service continuity.

It is possible to decide to migrate from a single CCS softswitch architecture to a multi-CCS softswitch architecture. The impact on the architecture is sketched in figure 9 on page 10.

When adding a new zone to a VoIP telephony network, additional hardware must be introduced for the following components:

- Front-end computers to load-balance calls on the additional zone
- Media servers to duplicate voice resources (optional)

Introducing a new zone with a Distributed Architecture is not necessarily linked to an increase of the telephony system's capacity. In fact, when considering a mono-CCS softswitch platform managing 60,000 users, the platform requires 20 SUs. When migrating to a multi-CCS softswitch architecture, the SUs can be shared between the two different zones (10 SUs for each zone). The OCSC components can remain centralized, since they can manage multiple zones from a single site.

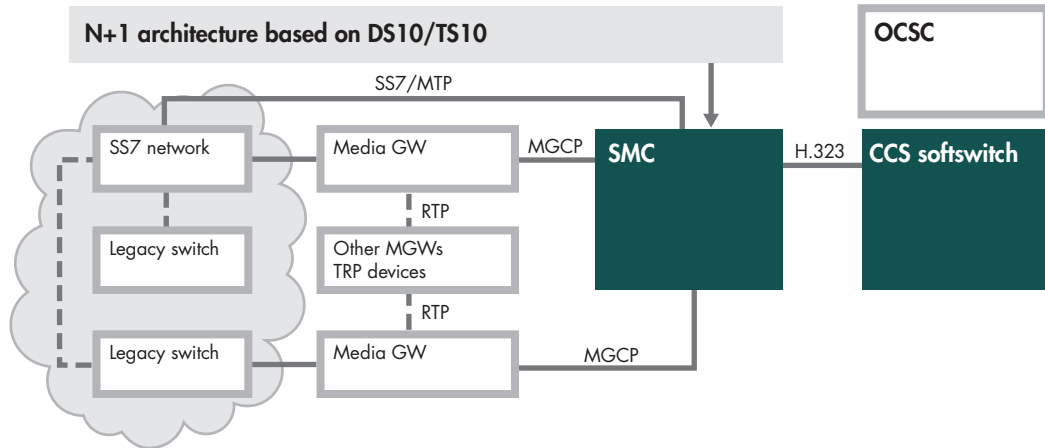
OCSC-Distributed Configuration

The OCSC Service Execution and Service Management platforms can be configured in various distributed ways to address specific requirements, such as capacity extensions, geographic constraints, and/or disaster recovery policies. OCSC SEP can be distributed as application servers dedicated to a given application zone (see figure 9 on page 10) or as mated pairs handling the same application zone in load sharing for disaster recovery purposes (see figure 10 above).

When configured as mated pairs, several SEPs can share the same subscriber and application data. Any update will be synchronized automatically on the various distributed SEPs using the SMP data propagation mechanisms.

A single OCSC SMP, keeping a centralized provisioning and management system, can manage distributed OCSC SEPs. A backup SMP can also be provided to address disaster recovery requirements for provisioning and management applications.

Figure 11
SMC interface with PSTN networks



Edge domain: SMC (Signaling and Media Gateway Controller)

SMC: interface with the PSTN networks

The interface to the PSTN networks uses the Signaling and Media Controller.

The SMC allows legacy PSTN networks and VoIP networks to be interconnected.

SMC consists of:

- An SS7 Signaling Control Function that allows connection to legacy PSTN networks based on ISUP (Integrated Services Digital Network User Part)
- A Media Gateway Control Function (MGCF) that:
 - Controls Media Gateways (the equipment that interconnects the voice trunks and the IP-based media)
 - Interconnects to IP-based Call Control Networks (e.g., H.323 in the current version of the SMC; SIP in the next version)

The MGCF maps ISUP messages to H.323 messages (or to SIP messages in the next version of the SMC) and controls Media Gateways using MCGP (or H.248 in future version of the SMC), thus forming the interworking in the convergent network.

Rely on the HP IN7 ISUP—a proven product

The SS7 ISUP connectivity and control are based on the HP OpenCall IN7 stack implementation, which benefits from a long and proven experience that:

- Conforms to the ITU-T, ANSI, and China standards
- Is deployed in 180 networks worldwide (wireless and wireline)
- Is certified by major telecommunications operators

SMC supports both E1 and T1 connectivity to SS7 Networks. SMC is also available on Linux with OpenCall SS7. (Contact your HP representative for details on availability.)

Distributed implementation

Scalability

The distributed processing structure of the HP OpenCall IN7 ISUP stack allows dimensioning of the SMC, according to a given requirement, and incrementally increases this capacity in line with growing network demand.

High availability

An N+1 architecture is implemented to ensure the global availability of the SMC. This architecture relies on HP IN7, which allows the SMC to be seen as a single point code. When a machine fails or is taken down for upgrade or maintenance, SS7 messages are distributed to the remaining machines in order to maintain the global call-handling capacity of the product. The distributed implementation of IN7 also allows applications to be insulated from changes and modifications in the SS7 network.

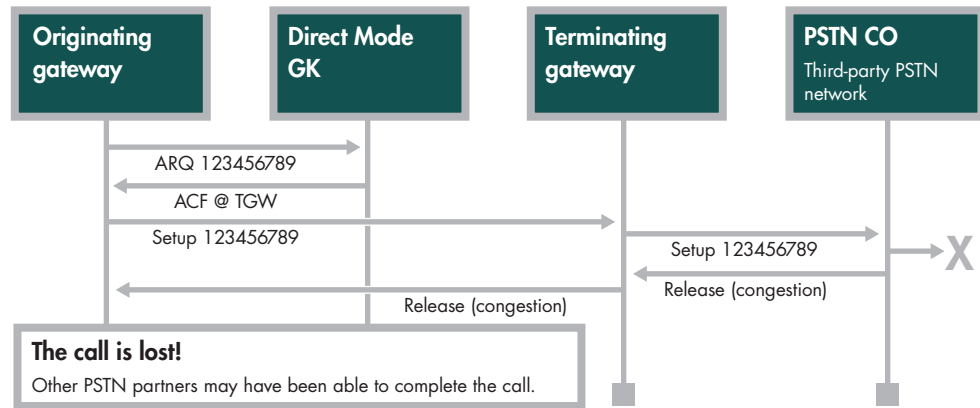
IP standards-based implementation

A major benefit of the distributed implementation is the standards-based implementation for the Call Control and Media Gateways Control, which allow the deployment of the Signaling and Media Controller with virtually any Gatekeeper/Proxy server or Media Gateway server in the marketplace.

The SMC runs on low-end HP servers (e.g., AlphaServer DS10 or TS10) as illustrated in figure 11 above.

Typically, all calls arriving on an Edge device are sent to the Class IV domain, except for some specific cases described below.

Figure 12
Call failure using “light” Class IV



Use of third-party infrastructure softswitches within the Edge domain

The Edge domain can implement its own call-routing or load-balancing policies, using for instance, infrastructure gatekeepers that monitor the state and load of several PSTN gateways. This monitoring is transparent to the Class IV domain as long as the infrastructure softswitch offers an LRQ-based interface for H.323 or uses SIP.

The Edge domain may also be used for direct trunking applications that do not require any service, e.g., leased-line emulation or pure SS7 trunk emulation. In this case, some calls may not be relayed through the Class IV domain.

Support of media gateways

The SMC supports Cisco AS53xx media gateway, as well as AudioCodes’ Median Gateway. Support of Cisco AS58xx is being validated. Other MGCP gateways can be supported after validation. Each validation generally requires several weeks to complete.

Core Class IV domain

Description

The role of the Class IV domain is to route calls between application domains, between Edge components, and between Edge components and application domains. It contains the central routing configuration of the network, including least-cost routing features, and it also implements the local number portability feature when required. It does not prevent the Edge components from doing their own call routing within the Edge domain, if relevant.

The core Class IV domain is composed of one or more Call Control Servers, configured as Class IV switches and controlled by one or more Service Controllers. The core Class IV domain supports SIP and H.323 call control protocols.

“Light” Class IV versus “true” Class IV

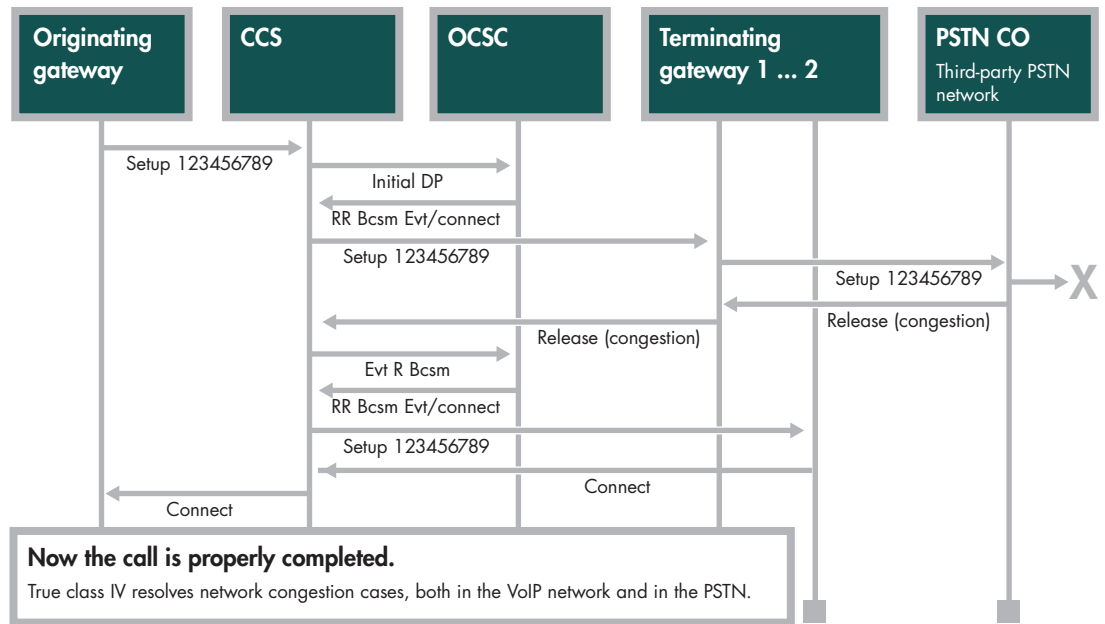
Both “light” and “true” models are fully supported by the proposed architecture. However, we recommend using the “true” Class IV model.

“Light” Class IV is the simplest way of routing calls in a VoIP network. It corresponds to the use of direct mode H.323 gatekeepers supporting the Location Request message, or SIP redirect servers. In such a network, any device willing to make a VoIP call asks where to send the call set-up message to the gatekeeper or redirect server. The Class IV service that is offered to the network is “light” because it does not provide these useful features:

- A single termination is offered for each call; the softswitch doesn’t take responsibility for rerouting the call to alternate destinations if the primary destination isn’t responding. In light Class IV networks, the call termination rate of the network cannot be better than the termination rate of the preferred partner on any given destination.
- Alias translation is not possible, requiring a homogeneous dial plan inside the network, and smart gateways that can convert from the PSTN format to the VoIP intra-network format and vice versa.
- CDR generation is gateway-centric, often requiring a single vendor network.
- Since edge devices have to communicate directly, a network with N vendors requires $N(N-1)/2$ interoperability tests.

“True” Class IV Domain mimics the current architecture of PSTN networks. The softswitch combined with the service controller can take full responsibility of terminating all calls it receives and may try multiple termination devices to terminate a call. As a result, the termination rate of a network using a true Class IV network is much better than the individual termination rates of each termination device or partner. In order to do this, the softswitch is on the path

Figure 13
Rerouted call using “true” class IV



of all call set-up messages (H.323 routed mode; SIP stateful proxy). Because it is aware of all call-related events, a true Class IV switch can create centralized CDRs that were independent of the gateways. True Class IV networks inherently facilitate multivendor deployments; each vendor requires only an interop test with the Class IV softswitch.

The call flow illustrations in figures 12 and 13 show two examples: The first is a call that fails in the PSTN because of congestion (a very common situation when doing least-cost routing). The congestion is reported using a Q850 error code. A light Class IV network cannot recover from this situation, while a true Class IV network can, as noted in figure 13.

Interfaces with other domains

Calls received from the Edge or Application domains

The Class IV domain receives calls from the Edge domain on IP address IPcore, optionally using an H.323 LRQ message in which case, the IP address used to forward the Q931 signaling is specified in the LCF. For SIP, a regular INVITE message is used, indicating a user or service is being invited to participate in a SIP-based call session.

Calls sent to the Edge or Application domains

Depending on the Edge domain device, the Class IV domain can either send an LRQ message or send a direct Setup message to the proper IP address. If an LRQ is used, the IP address used to forward the Q931 signaling is expected in the LCF. For SIP, a regular INVITE message is used.

Class V domain

The Class V domain is responsible for all subscriber-specific features, including:

- Unconditional call forwarding
- Time-dependent call forwarding
- Source-dependent call forwarding
- Redirect on busy or no-answer calls
- CLIP/CLIR (Calling Line Identification Presentation/Calling Line Identification Restriction)
- Call blocking
- Legal interception

In order to improve scalability, the Class V domain does not maintain CPE (Customer Provided Equipment) registrations directly, but communicates using LRQs with one or several Access domains that maintain this information.

This distributed design allows operators to resolve issues associated with centralized designs when the network starts. It also facilitates the use of multiple CPE vendors, by allowing the deployment of CPE-specific access servers as required. (See the “Access domain” section below for details.)

Access domain

The Access domain is composed of direct mode gatekeepers for H.323 (or Registrars for SIP), usually from the same vendor as the CPE vendor to ensure interoperability of security mechanisms.

The role of access gatekeepers is to:

- Maintain registrations of CPE, as well as keep-alive registrations

For a complete list of “supported” (tested) endpoints, contact your HP representative.

- Respond to the ARQs of devices and point to the Class V CCS for all authorized calls
- Block call attempts from CPEs to destination numbers not authorized by the local dial plan (for instance, a ported prefix)

In addition, the Access domain includes all necessary servers to allow CPE to download their firmware and configuration, including the registration alias, security parameters, and software keys that limit features of the CPE (e.g., multiple call handling or three-way conferencing).

If a CPE authentication scheme is required, the Access domain will also require authentication servers, such as RADIUS servers.

Using a separate Access domain facilitates dimensioning the network. While the Class V and Class IV domains scale according to the number of simultaneous calls in the network, the Access domain scales proportionally to the number of CPE lines. The relative sizes of the two domains may vary widely according to average line usage. A very low line usage (typical of software endpoints) will require a generous sizing of the Access domain, but relatively less capacity in the Class V domain. A higher line usage (e.g., hardware CPE or high-usage lines, such as call center agents) requires the exact opposite dimensioning.

Redundancy in the access zone with such gatekeepers is achieved using the “alternate gatekeeper” feature or clustering pairs of Access gatekeepers. Current Access gatekeeper is supported by the Cisco IOS gatekeeper. Other H323 gatekeepers could be supported after validation.

For endpoints using the MGCP protocol, Access domain management is performed by the SAU. This unit manages MGCP endpoint registrations, as well as multiline features, which are not handled by MGCP devices because of the protocol’s design.

IP Media Server

The HP Media Server solution set offers a robust and scalable platform for the development of value-added services applications, supporting thousands of simultaneous calls in a wide variety of telecom configurations. Media Server, in its three major configurations for Service providers (Intelligent Peripheral, IVR, and Service Node), is a tool for the implementation of multiple customer applications, such as network announcements, unified messaging access, voice portals, and IP voice mail.

The Media Server comes with the following applications:

- Network announcements—allows call mapping of certain numbers to prerecorded announcements.
- Simple voice mail—allows recorded voice messages—when the end-user CPL call management rules instruct the voice platform (softswitch + OpenCall)—to forward to e-mail. The e-mail addresses of end users are stored in their profiles in the LDAP directory, and the Media Server sends an SMTP mail with the voice message as an attachment to the address of the user. The voice mail then appears as a regular voice mail message in the user mailbox.
- Advanced voice mail (optional)—is included if the mail server also supports IMAP 4 mail retrieval, and it allows users to consult and manage voice mail by phone. The application needs to be configured with the user mailbox password in the user LDAP profile (for security reasons). Note: HP does not provide the SMTP mail server as part of the basic offering, and the provisioning tool only configures the LDAP user profile. User mailbox creation and management are done using the tools provided by the mail server vendor.
- Other applications can be developed on top of the VXML environment.

Endpoints

The HP solution supports multiple endpoints: H323, SIP, MGCP and either soft phones, IP hard phones, or analog phones. The solution can also support H323 “video” phones. For a complete list of “supported” (tested) endpoints, contact your HP representative. Other endpoints can be supported after validation.

HP focuses more than 25 years of telecommunications expertise into a powerful integrated team, the Network Service Provider Business Unit (NSPBU).

Management: FACPS (fault, accounting, configuration, performance, security)

- Fault management—the CCS and OpenCall platforms provide SNMP traps that can be integrated in HP OpenView.
- Accounting—HP OpenCall generates CDRs (call data records) that can be collected by billing platforms.
- Configuration-service configuration—HP OpenCall provides the Service Management Platform, which provides all configuration data for the VoIP value-added services. A Web-based application is provided using these data and enabling the operator, as well as end users, to configure their profiles.
- Network configuration—network configuration is being taken care by the network infrastructure solution, generally using an HP OpenView system-based solution.
- Performance—the HP platform provides MIB information, configurable to generate performance data. HP platform can also provide online statistics on calls.
- Security—the HP platform is a fully multitenant carrier-grade platform where each customer’s data is totally secure and not accessible by other customers. In terms of user authentication, a login hierarchy is available with different profile configurations, from administrator and basic operator to different categories of end users. Internet security related to address plan—public, private addressing—is also being taken care of with adjacent components (optional) that enable proxy functionality.

Conclusion

The HP VoIP solution offers a unique and optimized combination of open/scalable architecture and a set of features enabling service providers to:

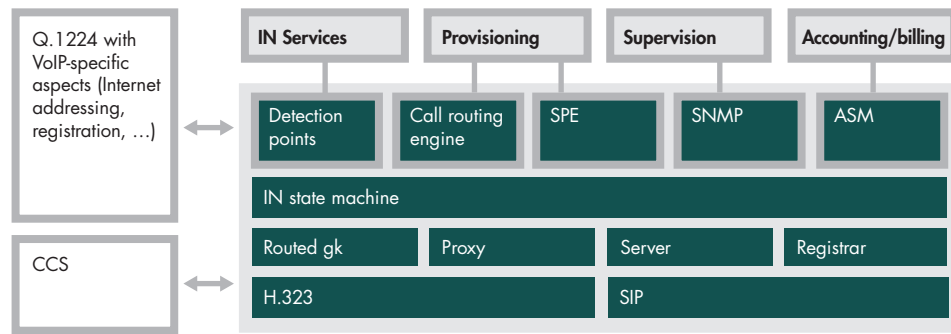
- Deploy rapidly an off-the-shelf, ready-to-operate-and-bill service
- Integrate a solution with current legacy technology, thanks to a complete set of traditional TDM or IP capabilities
- Enhance applications and their functionality by customizing them to adapt to specific requirements
- Develop an architecture and a set of products that guarantee medium- to long-term evolution—in the pure IP space, as well as in the converged 3G and multimedia IP communication space

Communications solutions are highly complex, and service providers must deliver even more innovative services to the market while keeping customers loyal and insulated from the complexities behind the services. In order to achieve this, service providers need strategic partners that can do more. HP offers a range of targeted, seamless solutions that are integrated with partners and delivered quickly and efficiently. HP systems and solutions are open and flexible, empowering customers to customize or create value-added services. Our service capabilities provide the expertise to develop, integrate, test, install, and support the most complex service launches. This one-stop shopping approach lets service providers focus on their customers—not their suppliers.

HP focuses more than 25 years of telecommunications expertise into a powerful integrated team, the Network Service Provider Business Unit (NSPBU). The NSPBU, along with 500 valued solutions partners, assists the world’s top 200 service and equipment providers and meets the voice and data needs of hundreds of millions of wireline and wireless subscribers.

With solutions, technologies, and services—including HP OpenCall and HP OpenView telecommunications capabilities arrayed across network infrastructure, network services, operations and business support, mobile and rich media solutions, and end-user access—the HP NSPBU is a major player that is leading the network and service provider industry.

Figure A-1
CCS logical diagram



Appendix A: Interfacing VoIP infrastructure with existing IN services

Introduction

The objective of this appendix is to describe how IN services already deployed in the PSTN network can be accessed from subscribers connected to a service provider's Voice over IP (VoIP) network.

The proposed approach consists of delivering a trial solution to the service provider to validate the target solution and further identify any specific requirements in this area.

In the sections below, it is assumed that the target VoIP network is based on HP's standard VoIP solutions. In particular, the CCS Softswitch offers the VoIP network core-switching function.

The proposed interface between existing IN services and the VoIP network will be implemented by an Intelligent Network Application Part (INAP) gateway supporting a standard INAP over SS7 interface toward the PSTN network and an extended INAP over IP interface toward the VoIP network core-switching function.

- The PSTN network interface (i.e., INAP/SS7) of the gateway is based on proven HP OpenCall SS7 and IN software currently deployed in almost all operator networks worldwide.
- The VoIP network interface is based on the state-of-the-art extended INAP interface delivered by the CCS Softswitch, which is also deployed in large VoIP networks worldwide.

The sections below describe the core switching functions of the CCS softswitch support-standard INAP call model, the proposed gateway architecture, and the IN call flows supported for the trial.

IN call model supported by the CCS

As already mentioned in an earlier section of this document, the CCS product is a VoIP Softswitch providing advanced call control architecture and implementing an IN call model (Q.1224) with additional detection points that take into account specific aspects of the VoIP framework (e.g., registration, a mobility function).

These VoIP-specific additions are further referred to in this appendix as INAP CCS extensions, and they are marked with an asterisk (*) in some of the diagrams and call flows on the following pages.

A Detection Point (DP) is associated with each Q1224 state. A default action that will be the "next step by default" can be defined for each DP.

Options exist for filters to be positioned on every DP: If a filter is positioned on a DP, the action is passed on to the Service Control Point; if there is no filter associated to a DP, the process is continued on the CCS softswitch.

Filters are either static or dynamic, are defined at the system initialization, and can be modified "on the fly."

There are two types of Detection Points for which actions have to be triggered on the SCP.

1. "For information"—the DP and its associated parameters are sent to the SCP, but CCS doesn't wait for any feedback from the SCP. Next state is the default one.
2. "Blocking"—The DP and its associated parameters are sent to the SCP, and CCS waits for a return from the SCP or for a Timeout that puts the process in the next state.

Returns from the SCP are

- DP parameters that will be changed and the next state is the default one.
- The next step is not the default one, but the Service Logic, which has been activated, defines it.

Figure A-2
CCS O-BCSM and T-BCSM

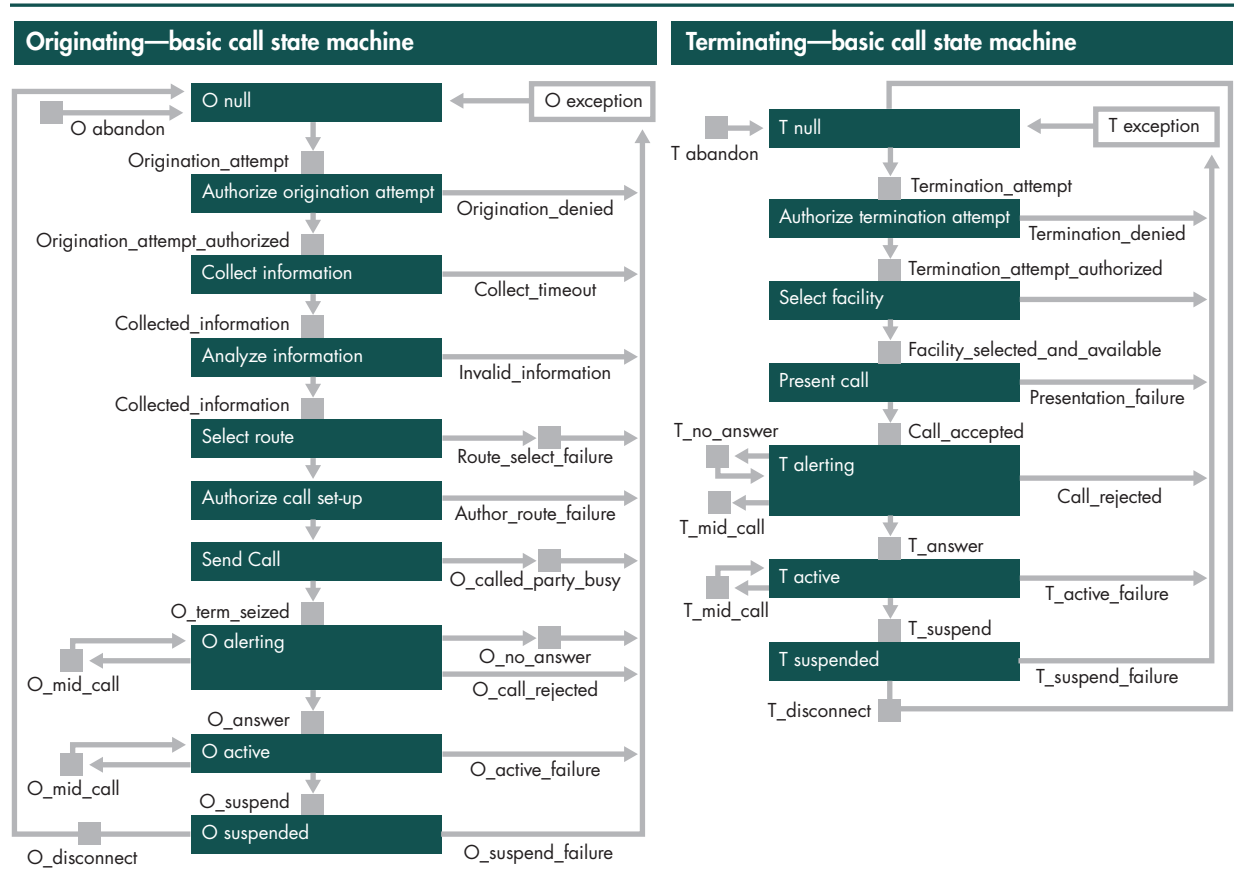


Figure A-3
INAP gateway overview

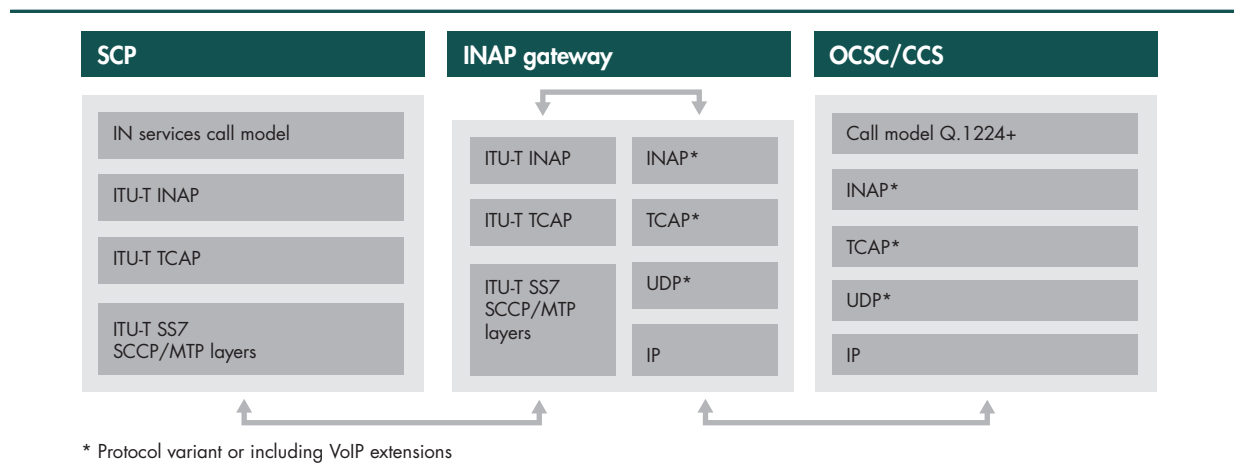


Figure A-4
Scenario 1—basic call

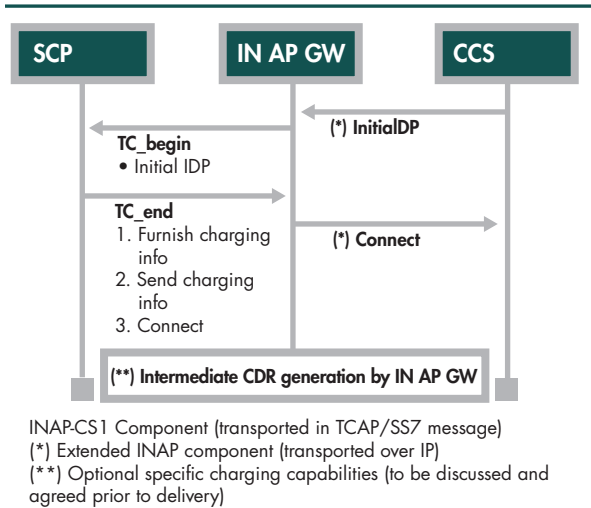


Figure A-5
Scenario 2—call aborted (time out)

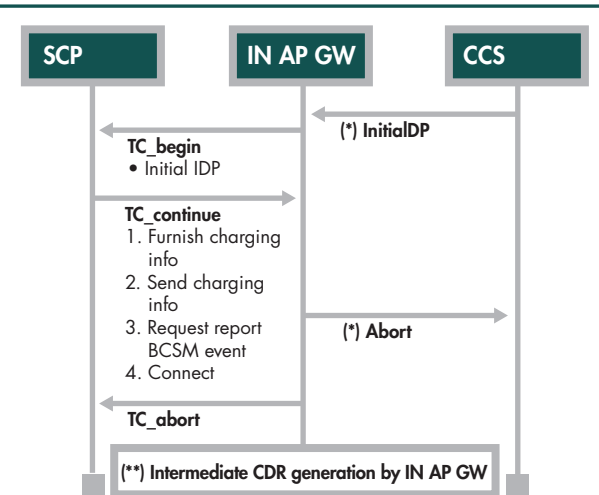


Figure A-6
Scenario 3—called party busy or no answer
 Note: This scenario could continue recursively or end as per scenarios 2 and 4.

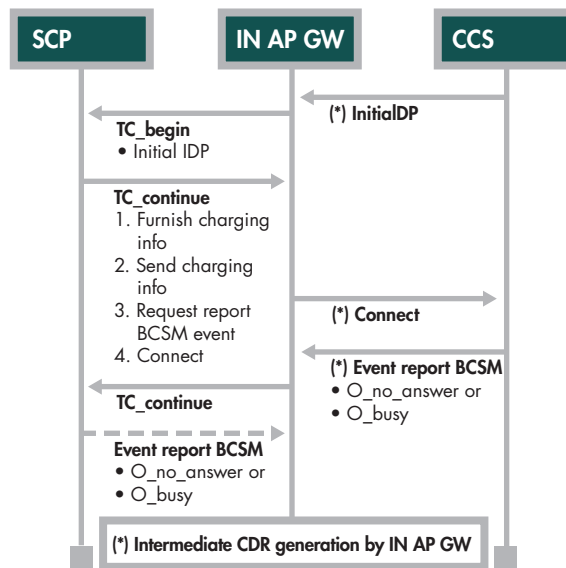
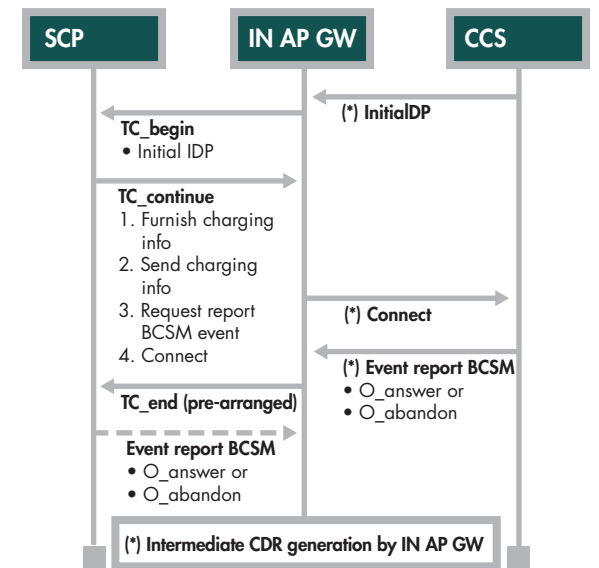


Figure A-7
Scenario 4—call answered or abandoned



The CCS undertakes Gatekeeper/Registrar operations. (For instance, in H.323, RAS and call routing are based on H.225 and H.245.) Therefore, the pure IN Detection Points (DP-alerting, DP-connect) are enriched by H.323/SIP-derived Detection Points (DP-register, DP-admission-pending).

CCS Detection Points make it possible to trigger specific processing during the different steps of a call (e.g., connecting, busy, no answer, and so on) or upon reception of H.323/SIP-specific events.

CCS is also capable of modifying the IP telephony signaling “on the fly,” making the implementation of a number of services possible. CCS capabilities are critical in the sense that they implement call-switching functions, as well as sorting and sequencing the call progression phases.

INAP gateway overview

As shown in the illustration below, the INAP gateway provides the VoIP interface core-switching function (i.e., the CCS softswitch) to the existing IN services running on the SCPs. The gateway relays two types of INAP messages:

- From the SCPs (INAP over SS7 layers) to the CCS (INAP over IP)
- From the CCS (INAP over IP) to the SCPs (INAP over SS7 layers)

This INAP gateway proposal brings the following major benefits:

- The gateway will implement the specific INAP adaptation required to interoperate with the existing SCPs in the PSTN network. The mapping between INAP and CCS-extended INAP (and vice versa) will be managed at the gateway level. Therefore, the SCPs will not have to handle any of the CCS INAP VoIP extensions, and the CCS, as well as the whole VoIP network infrastructure, will not have to handle any of the specific INAP parameters.
- Keeping network specificities at the INAP gateway, which is a centralized and highly scalable piece of equipment, will enable more flexibility and cost-effectiveness in managing future network evolutions.

As IN services are subjected to operator specificities on one side (e.g., country-specific and/or network equipment vendor-specific parameters), and as all PSTN/VoIP interoperability issues are not currently addressed by standard protocols on the other side, an inventory of operator specificities on IN services, INAP protocol parameters, and network implementation is needed. This is especially the case regarding all charging information and its associated handling. Also, for subscribers accessing IN services from the VoIP network, there may be some limitations due to current PSTN/VoIP convergence protocol limitations (e.g., advice of charging, metering pulse indication, and so on).

HP expects to discuss and agree with service providers regarding the impact of the above-mentioned, specific requirements on the current proposal.

The trial platform includes one gateway made of two servers (active and standby configuration types), enabling the transport adaptation between two SCPs and one CCS. We have sized the INAP Gateway, taking into account a maximum of four SS7 links per SCP.

The trial platform will support the sample INAP Call Flows described in the following section between the SCPs and the CCS.

Supported IN call flows

Figures A-1 through A-7 on pages 17–19 are the scenarios supported by the INAP gateway for the trial period. HP is ready to discuss and investigate the support of additional Call Flows to support a wider range of IN services as required.

For more information, go to www.hp.com/communications

© Copyright 2003 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Java is a U.S. trademark of Sun Microsystems, Inc.
Oracle is a registered U.S. trademark of Oracle Corporation, Redwood City, California.

5981-8273EN, 05/2003

