Technology White Paper

IMS inside the Enterprise

IP Multimedia Subsystem (IMS) is a technology that grew out of the carrier world, was standardized by 3GPP, and subsequently adopted and endorsed by 3GPP2 and ETSI (TISPAN). But how is it applicable to the Enterprise?
The rigid definition of protocol combinations and procedures that IMS brings to the world of IETF protocols make it an attractive proposition, in a world where seamless interoperability between VoIP solutions from different vendors has been slow to materialize.
But what will really drive enterprises towards participation in IMS is the desire to better incorporate the mobile phone into the enterprise telecommunications environment.
That means participating in an IMS applications environment. Through its User-Centric Broadband Architecture and its commitment to IMS, Alcatel will ensure that those applications can be located wherever is most appropriate for that enterprise and its SPs.

ALCATEL

P. Carden, E. Darmois, P. Tournassoud

IMS INSIDE THE ENTERPRISE

How will this technology, born in the world of service providers, infiltrate the enterprise? And who should be paying attention?

SIP (Session Initiation Protocol) and RTP (Real-Time Protocol) provide the major building blocks for advanced conversational services over IP, but it is the IP Multimedia Subsystem (IMS) that provides a standard architecture able to offer global scale, support for roaming, negotiation of QoS, application invocation and sequencing based on user profile, charging mechanisms, and interoperation with circuit-switched networks.

IMS is a set of standards born in the world of the Service Provider (SP). It takes standard IETF (Internet Engineering Taskforce) protocols for conversational services over IP, and articulates how to use them, in combination with other protocols and procedures, to build an end-to-end architecture for SPs.

What relevance does IMS have for enterprise users, what are the participation options, and what do those options mean for involved SPs?

The building blocks

The building blocks of an IP conversation are RTP, which provides transport of media streams, and SIP, which sets up those media streams.

In the simplest case, SIP session establishment can be a relatively straightforward process. First, the caller sends a SIP Invite message to the IP address of the called party. This Invite message includes the proposed RTP technical parameters (e.g. Codec type). The endpoints exchange additional SIP messages (including confirmation of negotiated RTP parameters), and the RTP session begins directly between the endpoints.

However, since IP addresses are both unfriendly and may change often, a mechanism is required to support persistent user-friendly addresses suitable for listing in a directory. For SIP, such addresses take an email-like form, for example sip://p.carden@domain.com or sip://+1212222453@domain.com. One or more intermediate systems between the two endpoints must be introduced to provide this address resolution and other session control functions.

By combining this session control with a directory, we can provide all the necessary support for a limited population of users, such as an enterprise, to communicate with each other. If we then introduce a PSTN gateway, we can also enable that enterprise to place calls to phones on the PSTN or, via the PSTN, to other VoIP (Voice over Internet Protocol) enabled enterprises with their own PSTN gateways.

But what happens when the role of PSTN is diminished? In a world where enterprises connect to each other directly via IP, and end devices such as mobile phones primarily use IP, we need a consistent global mechanism for SIP address resolution and session control. IMS provides that mechanism.

What is IMS?

IMS is a technology standardized by 3GPP (3rd Generation Partnership Project), and subsequently adopted and endorsed by 3GPP2 and ETSI (European Telecommunications Standards Institute) TISPAN (Telecoms & Internet converged Services & Protocols for Advanced Networks). It describes how standard Internet protocols (SIP, RTP) are implemented in a holistic architecture for conversational services (voice or video) over IP. These capabilities are delivered via a number of architectural building blocks *(Figure 1)*.

Figure 1: IMS components

Call Handling

- **P-CSCF:** entry point for a terminal in the visited network
- **I-CSCF:** entry point in the home network
- **S-CSCF:** basic call handling and application invocation
- **AS:** Application Server(s)
- **HSS:** directory of subscribers
- **PDF:** tunes the access when a call is initiated.

Announcements

- **MRFC:** controls the media system.
- **MRFP:** media system (e.g. announcements, codec translation)

Media Pocessing

- **MGCF:** gateway for call control
- **MGW:** voice gateway

In this article we focus primarily on the Call Handling elements and their relevance to an Enterprise context.

At the heart of IMS is the Call/Session Control Function (CSCF), which is split into three parts – the proxy-

CSCF (P-CSCF), interrogating-CSCF (I-CSCF) and serving-CSCF (S-CSCF). The CSCF components of the originating and terminating networks work together to resolve the called party address into an IP address.

At the same time, IMS addresses several other carrier requirements, including roaming support, QoS negotiation, application invocation and sequencing based on user profile, charging mechanisms, and interoperation with circuit-switched networks.

Figure 2: IMS call flow, mobile to mobile

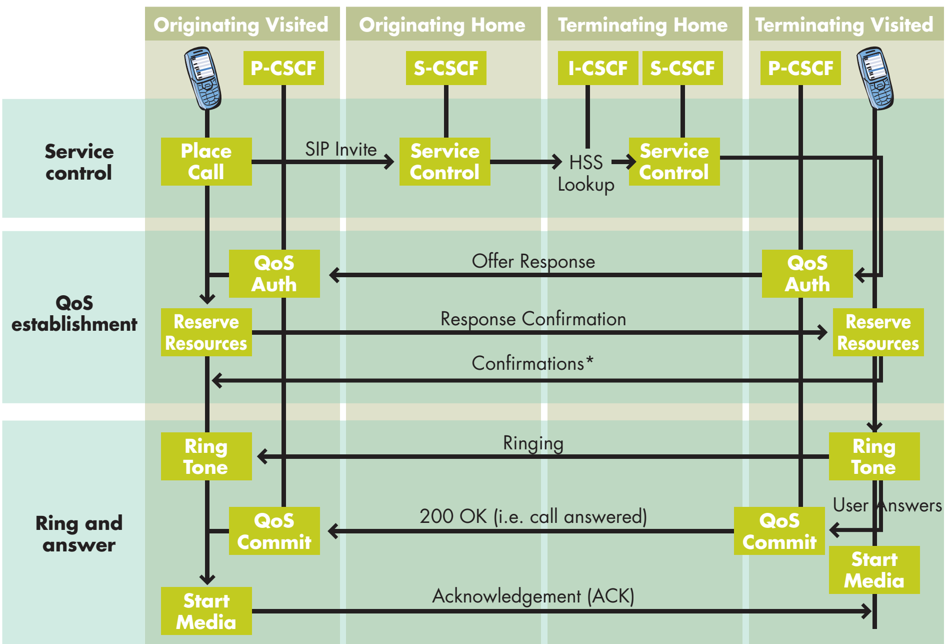

* *Simplified, Refer text*

A typical IMS call flow

To illustrate the CSCF component roles, consider the establishment of a conversation (audio or video) from one mobile phone to another *(Figure 2)*.

Since the mobile phones involved may be roaming (i.e. in a "visited" network), we can have four SP networks involved (originating visited/home networks, plus terminating visited/home networks).

For ease of understanding, it is useful to think of call establishment in three phases:
- Service Control provides address resolution so that the SIP Invite message gets to the right destination user terminal;
- QoS Establishment authorizes and reserves the necessary network resources;
- Ring and Answer, i.e. completion of the call, including commitment of resources.

Within the Service Control phase, the SIP Invite message, having traversed the originating networks (visited and home), comes to the I-CSCF and S-CSCF of the terminating home network. The S-CSCF has to apply the Service Control logic for the termination of the call to the right end device. For that to happen, it needs to know about the destination terminal and its associated P-CSCF, which occurs by a process called registration.

Becoming known – the registration process.

Figure 3 illustrates the mechanism by which a terminal in a visited network locates the S-CSCF it should register with. Such a process occurs when the terminal is switched on or enters network coverage.

First, the terminal must discover the address of a P-CSCF, which it typically does in conjunction with obtaining its IP address e.g. via DHCP (1). The P-CSCF is always the first point of contact in the network for the IMS client. After registration, it will forward SIP Invites received from the client directly to the correct S-CSCF. However, during registration, it does not yet know which S-CSCF is the right one, so first it uses the domain name associated with the mobile phone to look up an I-CSCF for that domain via DNS (2). The primary function of the I-CSCF is to refer a SIP message to the right S-CSCF based on user profile.

User profile information in IMS is always stored in a Home Subscriber Server (HSS). While in the current version of IMS there can only be one HSS per IMS user, users from different domains may be stored on different HSSs. In that case, the I-CSCF refers to a Subscriber Location Function (SLF) to determine which HSS is associated with the domain of the user registering (3). It then looks up the right HSS to determine the right S-CSCF (4), and forwards the registration message to the correct S-CSCF (5).

After completion of the registration process, both the S-CSCF and the P-CSCF maintain the association with that user so that later, when the terminal originates a call, the P-CSCF forwards the SIP Invite directly to the correct S-CSCF (6). In other words, for outgoing calls, the I-CSCF is not involved after the registration process.

Figure 3: The registration process

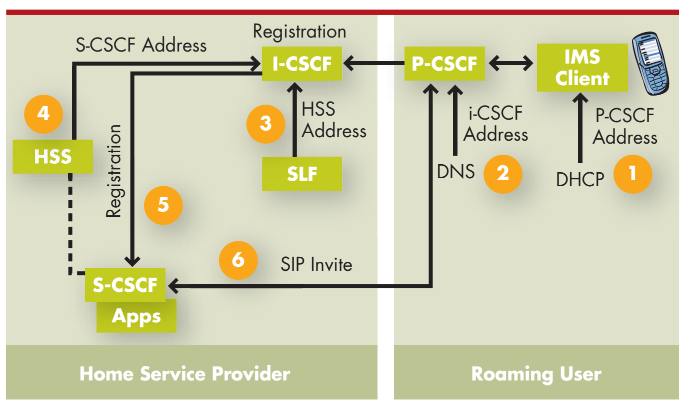

Someone's calling . . .

The same is not true for incoming calls. Because it is the I-CSCF that is listed in DNS, it is to the I-CSCF that an originating S-CSCF will send SIP messages for a particular domain. As for the registration process, the I-CSCF looks up the HSS to find the right S-CSCF to handle termination of the call (*Figure 4*).

The I-CSCF is therefore used for both registration and incoming calls. These two functions are similar but distinct. As we investigate the implications for enterprises, it is particularly useful to consider the two different DNS lookups involved. During registration, the P-CSCF looks up the I-CSCF(s) for the domain associated with the phone, while during call establishment the originating S-CSCF looks up the I-CSCF(s) for the domain in the SIP address (URL). Those domains might be the same, or different. For example, the domain of the user identity stored on the phone's SIM card would typically be the domain of the mobile SP, while the domain in the SIP URL might be that of the enterprise (in the same way that the domain name in an email address belongs to the enterprise). That in turn means that the I-CSCF for incoming calls may be different to the I-CSCF for registration – i.e. the I-CSCF for the enterprise need not be the same as the I-CSCF of the mobile SP.

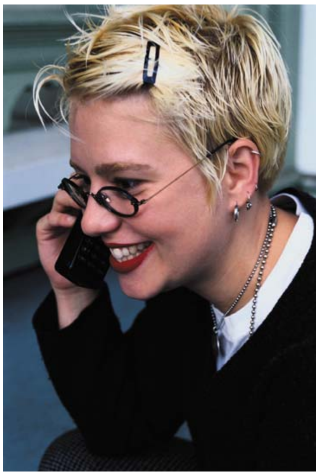

therefore on something that everyone can agree with: that the mobile phone will continue to be an important tool for the enterprise user. As mobile phones evolve to use VoIP, the mechanism they will use is IMS. We may therefore conclude in the simplest of terms that IMS is important to the enterprise because a significant proportion of end devices will use it.

However, at the same time, enterprises are beginning to demand much tighter coordination between their mobile phones and the rest of the corporate telephony environment. The features available to end users, as well as the controls available to network managers,

Figure 4: Incoming call

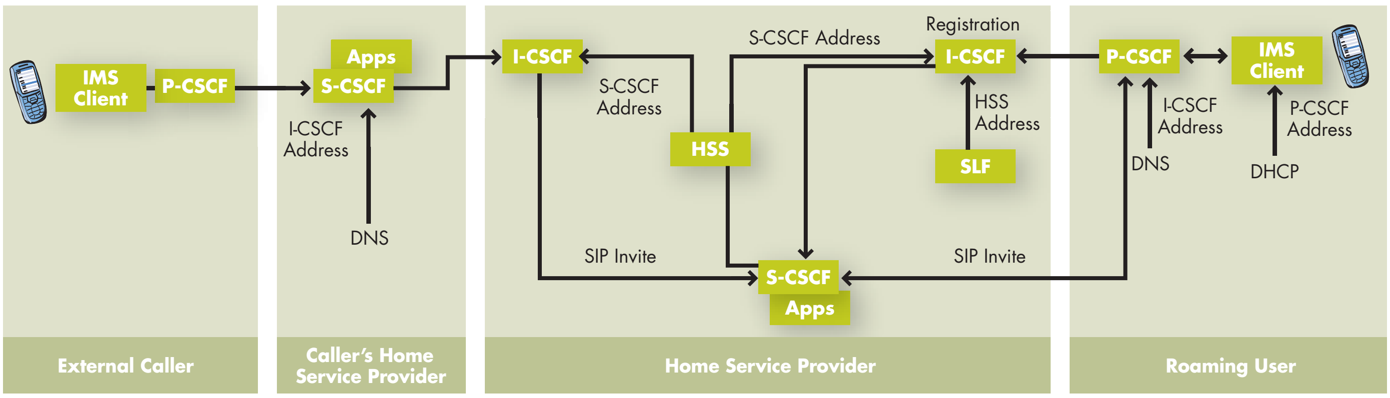

The **P-CSCF** is the SIP proxy which is the entry point for SIP terminals. The SIP terminals receives the address of the P-CSCF of the visited network at start time. Then, the terminals send their SIP messages to the P-CSCF. The P-CSCF analyzes the request. Extracting the terminal IMS address (and its domain) from the SIP message, it forwards the request to the home I-CSCF. The P-CSCF also sends the terminal description to the PDF to tune the network QoS parameters if needed.

The **I-CSCF** is the SIP proxy at the entry point of the IMS domain. The I-CSCF receives the SIP messages from subscriber terminals through the P-CSCF of the domain or through the P-CSCF of other domains (in the case of roaming), e.g., a call by a subscriber. The I-CSCF also receives the SIP messages from the CSCF of other domains (e.g., reception of a call from outside). Then, the I-CSCF gets the user status from the HSS and forwards the SIP message to an S-CSCF. Note that it is possible to use an I-CSCF in a domain for outgoing messages in order to hide the internal IP network topology.

IMS in the enterprise

There are many reasons why IMS could be important to the enterprise, but because the technology set is relatively new, many of those reasons are still open to debate. Let us focus

should be similar, regardless of whether the end device is a mobile phone, a desk phone or a soft phone (using the term phone in the broadest multimedia sense).

Enterprise deployment scenarios

For that to happen, the various devices must use a common set of applications – and since at least a proportion of end-user devices use IMS clients, the application environment must itself be IMS-compatible. That certainly doesn't preclude telephony applications from providing parallel support for other VoIP schemes. Indeed there are some features that could be envisaged within the enterprise that are not supported by IMS today. For example, IMS today doesn't address splitting audio and video for a particular session across different devices, which might be desirable in conferencing applications. Nor does it address some of the additional QoS complexities that exist today within enterprise networks. And whether or not IMS addresses such features in the future, there will nevertheless be the need to support parallel VoIP schemes for migration purposes. The important point to note is that even if the application environment is not purely IMS, it must support IMS, and incorporate IMS end devices.

By considering the preceding analysis of IMS, three mechanisms to incorporate IMS end devices into enterprise applications become apparent.

The **Home Subscriber** Server (HSS) describes the many database functions that are required in next-generation mobile networks. These functions will include the HLR (Home Location Register), DNS (Domain Name Server) and security and network access databases.

Option 1. Pure "net-sourced" model

For some enterprises, it will become increasingly attractive for all telecommunications applications to reside entirely in the network. This is especially true for small enterprises with highly mobile workforces. In such cases, all the IMS CSCF infrastructure, HSS and applications would reside in the network. User devices could consist of both IMS-compliant mobile phones and premises-based, IMS-compliant desk phones, soft phones, or user agents in iPBXs. The mobile phones would connect through the Home SP and its roaming partners, while the premises-based users would connect directly to a P-CSCF located in the Home SP's network (*Figure 5*).

Within Alcatel's User-Centric Broadband strategy, such "net-sourced" enterprise applications are a key component of Alcatel's Managed Communications Services (MCS) portfolio.

Option 2. Enterprise-based application server(s)

However, it is also recognized in the MCS portfolio that many enterprises will continue to maintain premises-based solutions (either fully in-house, or managed in an out-sourcing arrangement). Where such enterprises rely on a single mobile SP (typically medium-sized enterprises or larger, single-country organizations like government), the enterprise may maintain its own application server(s) while still relying on the SP for providing the CSCF and HSS infrastructure. Here, the SP's S-CSCF would hand off to an application (e.g. the enterprise's iPBX) via an IMS Service Control Interface (ISC) (*Figure 6*).

Since the S-CSCF refers SIP messages to a sequence of application servers (based on HSS profile), the SP may easily apply enterprise-independent service logic before or after the enterprise application is referenced. Not only does this provide a straightforward mechanism to accommodate "operational" requirements such as number portability, it also enables the SP to provide added-value applications in the network that work in conjunction with the premises-based applications. Such applications could range from virtual private networking (across fixed and mobile), to advanced mobility, collaboration, and contact center applications. In short, the enterprise has the flexibility to mix and match between premises and network-based applications, while the SP maximizes its addressable market because it can promote new services to enterprise customers, even when those enterprises are not considering replacement of their premises applications.

Option 3. Enterprise-based CSCF and HSS

For large multinationals, an additional level of complexity is introduced by the need to work with multiple "home" SPs in different countries. At the same time, it is desirable to provide consistency of user experience, not only from mobile to desk phone to soft phone, but also between regions. Such enterprises may wish to deploy an entire IMS infrastructure within the enterprise.

Here, the enterprise itself could maintain the HSS and CSCF infrastructure. Most likely the S-CSCFs would be geographically dispersed (in a similar footprint to today's mail servers), while more advanced applications and HSSs would reside in a very small number of data centers.

Its mobile SP(s) would use their I-CSCFs (with the SLF) to hand off registration of mobile users (including roamers) to the S-CSCF specified in an enterprise HSS.

Conversely, since the I-CSCF address is determined by a DNS lookup of the domain name (typically granular only to the level of the enterprise), the signaling for incoming calls (ter-

Figure 5: pure "net-sourced " model

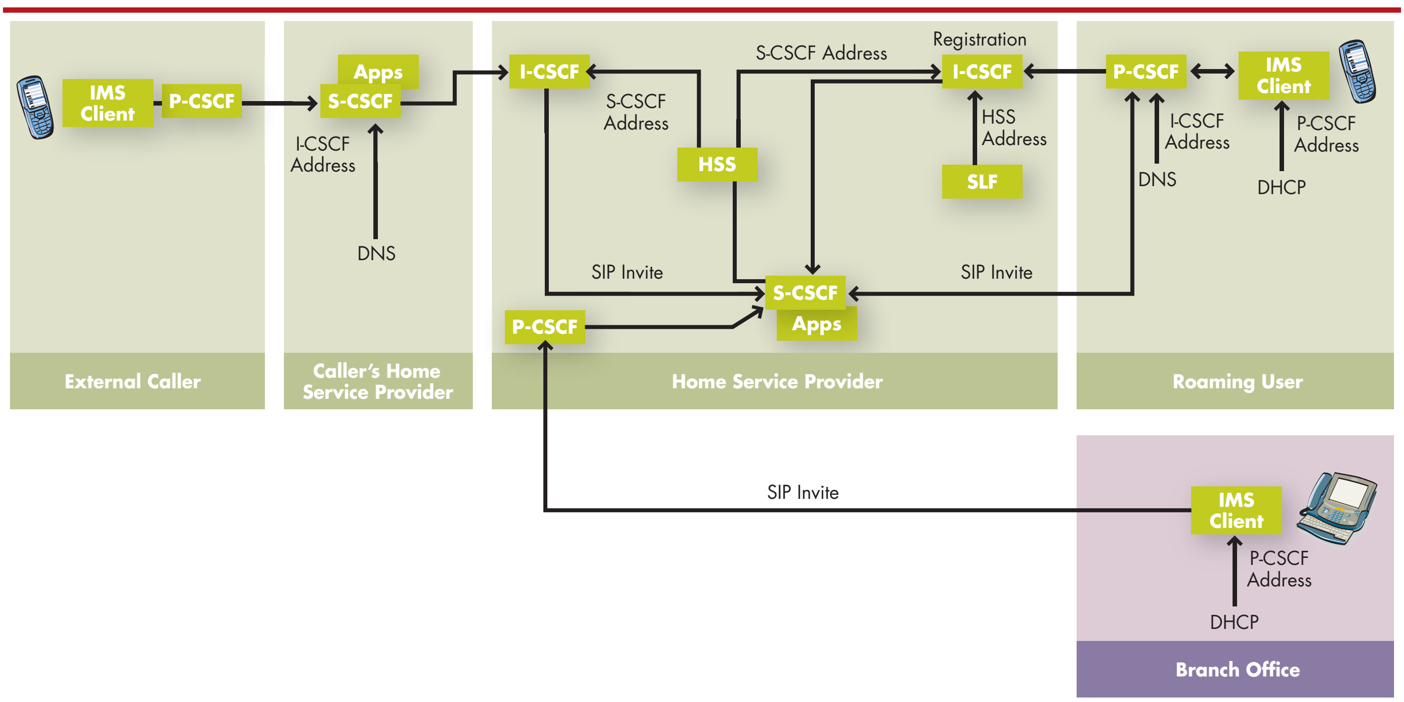

Figure 6: Premises-based solutions

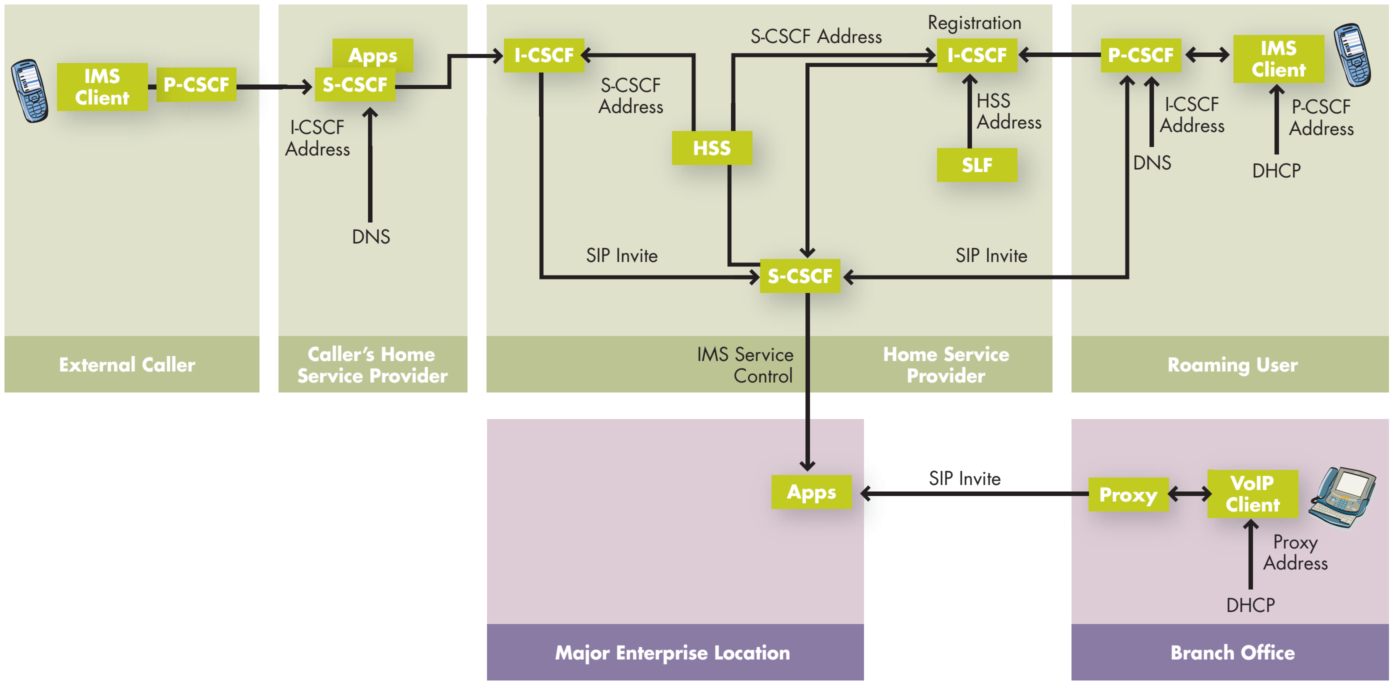

minated to the enterprise) would come from the originating S-CSCF to an enterprise-maintained I-CSCF, also in the data center (*Figure 7*).

The advantage of such an approach is that it allows a large enterprise to maintain an application environment that is independent of its mobile SPs. While still relying on mobile SPs to manage the complexities of macro-mobility and roaming, it enables the delivery of a consistent user experience from region to region and from mobile phone to desk phone to soft phone.

Such consistency of experience becomes increasingly important as the richness of advanced telecommunications applications increases. For example, if the collaboration client on my laptop allows me to invite a colleague from the Shanghai office to a conference, it should also let me invite my New York-based colleague who is currently visiting Tokyo in the same way – incorporating presence, location, and reachability information consistently. If I then transfer my leg of the conference call from the meeting room's conference phone to

Figure 7: "Do-it-yourself" IMS for multinationals

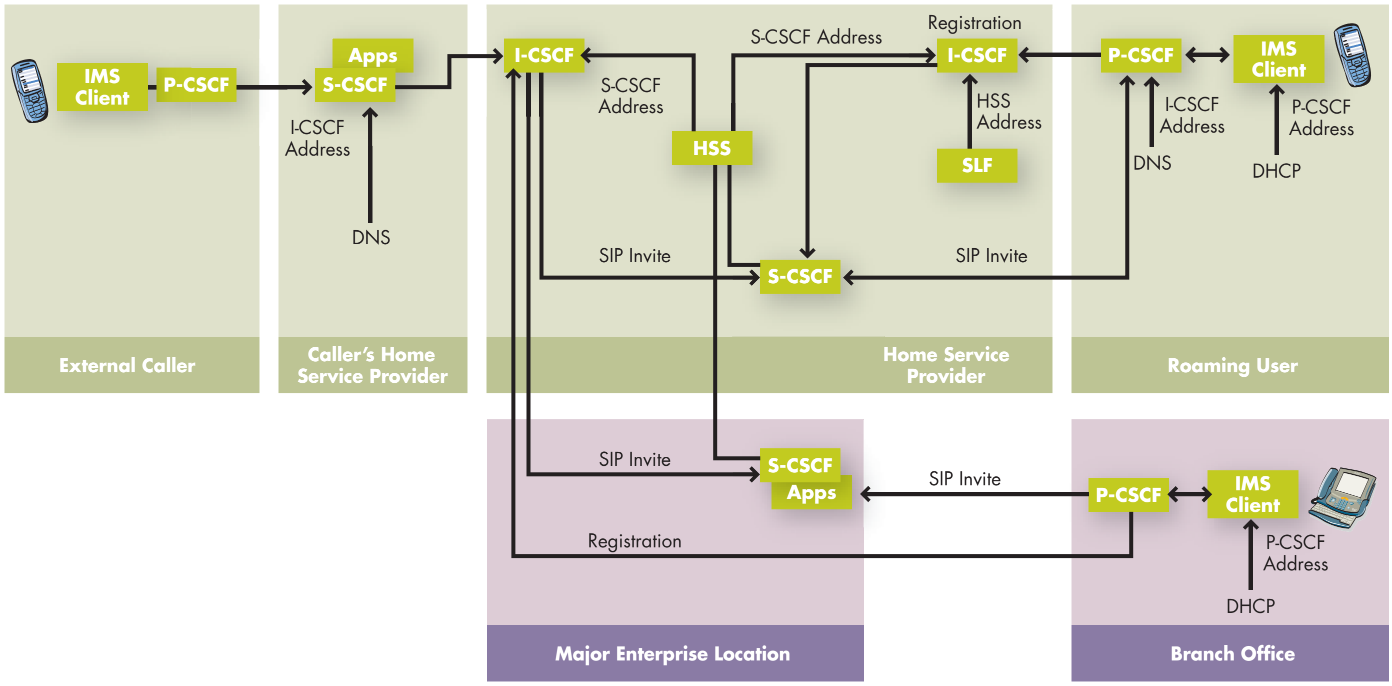

my mobile (so I can continue participating while in a taxi to the airport), I should do that in a standard way despite the fact that I am in Sydney with a home SP based in France. As applications get more sophisticated (picture the same scenario with video and shared documents), the need for consistency becomes more important.

However, providing consistency of end-user applications doesn't mean that all applications need to be provided in the enterprise. As in the previous option, the user experience can be delivered through a combination of applications located in the enterprise and in the network. Application location becomes unimportant – but what is important is that specific aspects of the user experience are delivered through the same application, regardless of his location, or which home SP he uses. In other words, where there are multiple mobile SPs involved, it is not practical that each is responsible separately for the user experience associated with that phone.

The required application consistency can be achieved in a number of ways – either through global consortia of mobile SPs that deliver a global solution, through a particular SP handling S-CSCF, HSS and applications on behalf of an enterprise, or through an enterprise maintaining much of the IMS infrastructure and applications itself.

An increasing level of complexity

None of these models is without challenges, and adding to those challenges are other industry developments including one-phone solutions, and the near-term need in many markets to support emergency calling services (ECS), commonly known as E911 in the U.S. and Canada.

By *one-phone* solutions, we refer to converged fixed/mobile solutions and/or services that employ dual-mode, short-range and cellular radio – e.g. WiFi and GSM/UMTS/CDMA – to switch between "fixed" networks (accessed via WiFi) and mobile PLMNs. Since such solutions must track the state of each phone, they are tightly linked to the notion of profile, which for IMS is managed by the HSS.

ECS solutions have a similarly tight dependency on profile. A primary requirement is the ability to determine location – not only to assist emergency personnel in locating an incident, but also to ensure routing of an emergency call to the right public response centre. There are many ways to determine location, nearly all of which rely on correlation between real-time information and static profile data.

For enterprises that rely on a single mobile SP, the conclusions are straightforward – despite an increasingly elegant user experience, and the simplifying effects of the standardized architecture that IMS provides for VoIP, the complexity of the telecommunications applications environment will continue to increase. This complexity, combined with the increasing participation of the mobile phone in enterprise applications, will make it increasingly attractive for such enterprises to adopt a net-sourced model for some or all of their telecommunications needs.

For multinational enterprises and their SPs, the impact of IMS is more complex, and requires careful consideration, not only in terms of technology, but also in terms of SP relationships. Because IMS drives the enterprise towards a consistent set of application services, enterprises must resolve who will be responsible for those applications. Will they do it themselves (in which case their SPs will need to provide a hand-off mechanism to the enterprise HSS for registration), or will they rely on a single SP for applications, with others providing mobile access (in which case the mobile access SPs will have to hand off to the application SP)? And a combination of the two approaches is certainly practical (at least for applications, if not for the CSCF/HSS).

Some large enterprises will conclude that the new technical and commercial challenge of getting different mobile SPs to work together towards a separate, consistent application environment suddenly makes a fully outsourced or net-sourced approach much more attractive. To meet that demand for a global managed communications service, SPs will need to consider another set of relationships – with each other.

In all likelihood, different deployment models will prove attractive to different enterprises in different industries in different regions. Such deployment flexibility has always been a cornerstone of Alcatel's Managed Communication Services philosophy.

Conclusion

One-phone and Emergency Calling Services are both areas where technologies and standards continue to develop, and indeed they are both likely to influence the future development of IMS itself. What they serve to illustrate in this context is that the application environment will continue to increase in its complexity, reinforcing the need to manage that complexity through a structured, consistent IMS architecture. By incorporating a variety of different deployment scenarios in its Managed Communications Services framework, Alcatel not only accommodates sophisticated future IMS scenarios, but also, equally importantly, provides an elegant migration towards IMS from today's state-of-the-art enterprise solutions.

Bibliography

3GPP TS 23.228 V7.1.0 (2005-09), 3rd Generation Partnership Project, Technical Specification Group Services and System Aspects, IP Multimedia Subsystem (IMS) Stage 2 (Release 7)

Philip Carden is Strategy Director for Managed Communication Services (MCS) in Alcatel's Corporate Network Strategy Group, Paris, France. (philip.carden@alcatel.com)

Emmanuel Darmois is Vice President, Unified Interaction Management, Enterprise Solutions Division, based in Colombes, France. (Emmanuel.Darmois@alcatel.fr.)

Pierre Tournassoud is Vice President, Network Strategy Group, in Alcatel Headquarters, Paris, France. (Pierre.tournassoud@alcatel.com)

Glossary

3GPP 3rd Generation Partnership Project
3GPP2 3rd Generation Partnership Project 2
AAC Advanced Audio Codec
AS Application Server
DNS Domain Name System
ETSI European Telecommunications Standards Institute
G.711 A codec
HSS Home Subscriber Server
HTTP Hypertext Transfer Protocol
I-CSCF Interrogating-Call/Session Control Function
IETF Internet Engineering Taskforce
IMS IP Multimedia Subsystem
MGCF Media Gateway Control Function
MGW Media Gateway
MP3 Motion Picture expert group audio layer 3
MPEG Motion Picture Expert Group
MX Mail Record in DNS
NAPTR Naming Authority Pointer (A DNS Record)
P-CSCF Interrogating-Call/Session Control Function
PDF Policy Decision Function
QoS Quality of Service
RTP Real-Time Protocol
S-CSCF Interrogating-Call/Session Control Function
SIP Session Initiation Protocol
SLF Subscriber Location Function
SMTP Simple Mail Transfer Protocol
SRV Service Record in DNS
THIG Topology Hiding Internetworking Gateway
TISPAN Telecoms & Internet converged Services & Protocols for Advanced Networks (TISPAN), a standardization body of ETSI
UA User Agent