Making IMS Work: Current Realities, Challenges And Successes

Jeff Fried and Duane Sword

IMS is real, and it is proceeding. Enterprises should be prepared.

The IP Multimedia Subsystem (IMS) is steadily progressing from concept to reality. But significant real-world issues remain.

All major equipment providers have committed to the IMS architecture and are implementing parts of it—as are a selection of major IT vendors and a variety of specialists. Many vendors and carriers now tout IMS as the single most significant technology change of the decade. New product announcements, reports, newsletters and articles are continually appearing. So what is the current reality, and how can you prepare for the inevitable challenges?

This article will survey the current state of IMS readiness and touch on some current issues and alternatives.

The Premise And Promise Of IMS

IP Multimedia Subsystem (IMS) is an emerging reference architecture for carrier-class networks that is evolving through several standards groups, research labs, vendors and carriers. IMS allows wireless and wireline carriers to use a common IP application core to deliver a host of new, content-rich multimedia services combined with legacy services across a variety of access technologies. (For more background on IMS, see *BCR*, June 2005, pp. 18–23). The basic idea behind IMS is to shift the network architecture from "silos" to "layers" in delivering services across different access networks (Figure 1).

Jeff Fried (jfried@empirix.com) is chief technology officer for Empirix Inc. Duane Sword (dsword@empirix.com) is vice president of product management for Empirix Inc.

FIGURE 1 The Premise Of IMS

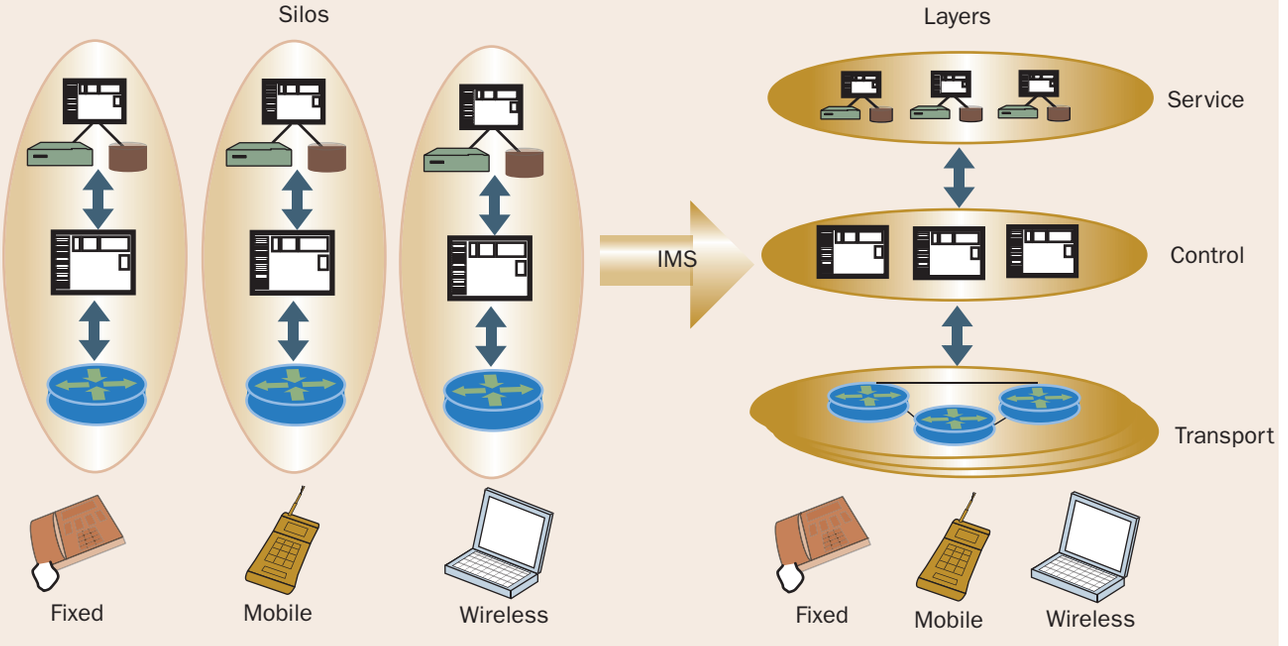

IMS's primary advantage for enterprises will be in fixed-mobile convergence

IMS promises to accelerate convergence in many dimensions (technical, business-model, vendor and access network) and make "anything over IP and IP over everything" a reality. However, the stratification of the transport layer, control/session plane and applications—a key change to the network with IMS—creates unique challenges from a service quality-assurance perspective.

To some, IMS is "simply" a flavor of Service Oriented Architecture (SOA–also see this issue, pp. 58–64), and is fueled by the same forces that drive SOA in enterprises and service provider networks. Many also hearken back to the Intelligent Network (IN) initiative of the 1980s that delivered many real benefits but fell short of the dream of rapid application delivery.

Some call IMS the "last gasp of the smart network" and predict that it will be circumvented even as it collapses of its own weight (see *BCR*, October 2005, pp. 54–58). Others truly regard it as the culmination of major forces and the realization of convergence on a global scale. But no matter how you look at it, a few premises in IMS are now very well accepted:

- Decouple access from applications.
- Provide functions as modules.
- Use standardized interfaces, reusing as much as possible (e.g., SIP).

These may sound simple, but there are wide-ranging implications, many of which are already evident in early IMS products and trials.

Currently, the move to IMS is primarily being spearheaded by Tier-1 operators in the U.S. and Europe. Many of these have already announced IMS contracts, including Cingular, Sprint, Bell-South, AT&T, France Telecom, Telecom Italia Mobile, Vodaphone, O2, Swisscom Mobile, and more. These carriers are faced with flat or decreasing average revenue per unit (ARPU), increased competition and a desire to roll out a large variety of compelling services to increase customer "stickiness."

For enterprises, IMS could provide advantages in fixed-mobile convergence (FMC), enabling increased mobility at lower costs, as well as the potential for access and control of services in the network.

Of course, for this promise to be realized, progress has to happen on many fronts: Standards, products from vendors, infrastructure rollouts, new applications, business models, etc. IMS is a complex architectural framework, and itself is an enabler and part of a larger picture as network services converge and evolve.

This article touches on some of the current realities and lessons, to shed some light on how things are unfolding and what can be done to prepare for the inevitable challenges.

IMS Standards—Surprisingly Fast Progress, But Still Unfolding

IMS began in the 3GPP group as an effort by wireless carriers to standardize service delivery. It has gained remarkable steam (and a remarkable alpha-

FIGURE 2 IMS Standards Players

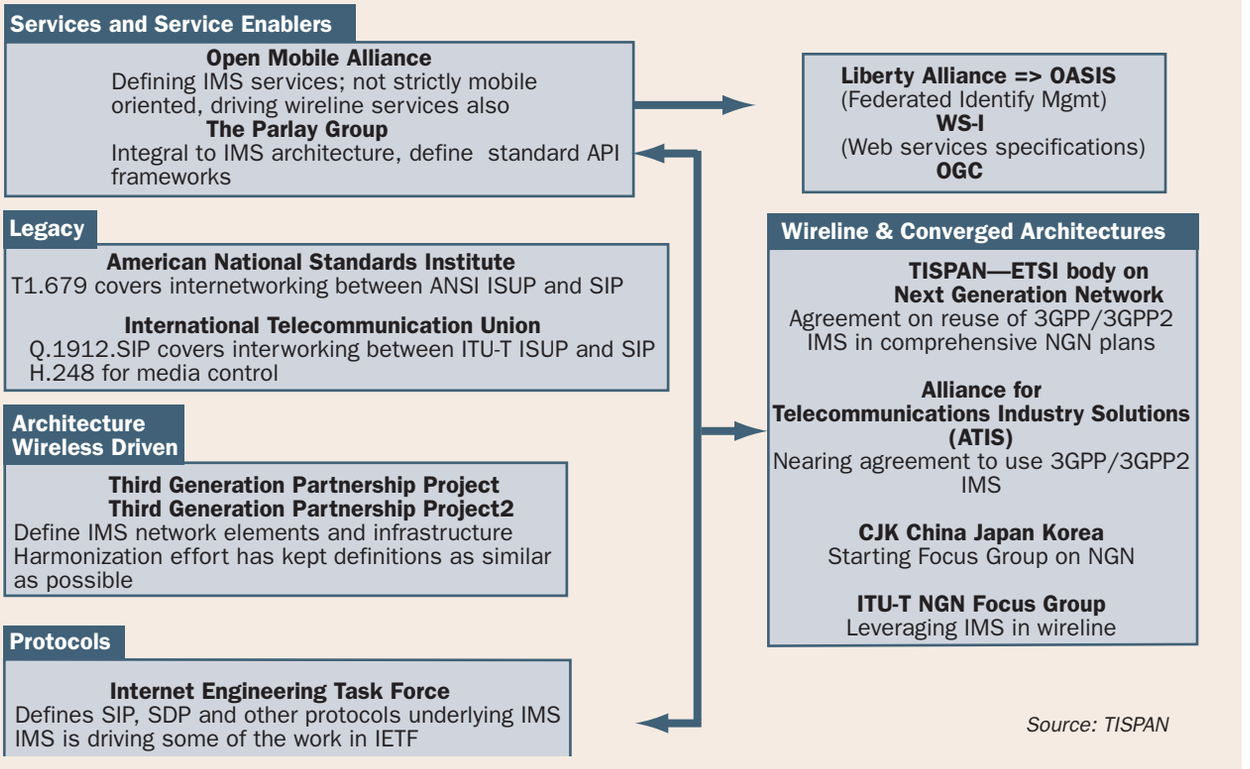

Source: TISPAN

bet soup of acronyms), and is now active in at least 14 different standards forums in a major way (Figure 2), with more than 40 standards groups in the picture. We won't delve into the intricacies of these groups and the various versions of IMS and "pre-IMS" standards here. It is, however, worth commenting that the complexity of the standardization effort currently under way is predictably represented in the complexity of the IMS standards framework.

IMS is not a standard; it is reference architecture that defines functions within a three-layer architecture consisting of

1.) The access/transport/device layer.
2.) The control layer.
3.) The applications layer.

Figure 3 shows a simplified view of the reference architecture. As much as possible, existing standards are incorporated into the architecture to specify functions. This re-use of existing standards such as SIP and Parlay, plus a general sense of urgency in these standards groups, has resulted in remarkable progress. For an industry that took 12 years to standardize T1 and E1 and 15 years to standardize AIN, getting to the current IMS framework in 5 years is remarkable.

However, many issues remain, and standards haven't yet addressed several big areas. These include the definition of common security elements and security aspects such as handling denial of service, topology hiding, and overload protection. Legal requirements for "lawful intercept" aren't dealt with yet, nor are the "transcoding" approaches that will be necessary when users roam between wireline and wireless networks.

It's already clear that the conformance and interoperability challenges with IMS will be huge. To date, there haven't been any "interoperability" events or conformance test suites beyond those in place with the existing protocols IMS adopted.

The first such events are coming in mid-2006. As a result, most of the early trials have focused on single vendors' product lines. But since no vendor covers everything and carriers are pushing for standardization in order to have interoperability, you should expect much more effort on interoperability testing.

IMS Products—Early Implementation Challenges

All major traditional telecom vendors are committed to the development of a comprehensive IMS solution. Figure 4, p. 46, provides an overview of the various types of products and vendors involved: There are at least 11 major "suite" vendors, 6 softswitch players, and 45–50 specialists. Although IMS specifies functions, it does not specify packaging, so many products encompass more than one function. Different vendors package the functions differently, which complicates the construction of a true multi-vendor system (see *BCR*, April 2006, pp. 40–45).

What is particularly interesting about IMS is that there are more than 12 "traditional IT vendors" involved, and many products and concepts come from the IT world rather than the telecom world. Application servers, Web services interfaces, and several other IMS elements are new to telecom but well established in the IT world. The equivalence is not always direct; for example an "application server" in IMS runs the gamut from a SIP proxy to a J2EE app server.

Because of the sheer complexity of IMS deployments, integration is crucial. A breed of "Network Integrator" (NI) is emerging with the expertise to make complex multi-vendor implementations work in demanding carrier environments—something many carriers have traditionally done by themselves.

To date, all the major network integrators are part of network equipment manufacturers: Lucent, Alcatel, Sonus and Siemens, for example. There is overlap between "Network Integrators" and traditional "System Integrators" like IBM, Accenture,

The carriers' big vendors are serving as their integrators

FIGURE 3 IMS Reference Architecture

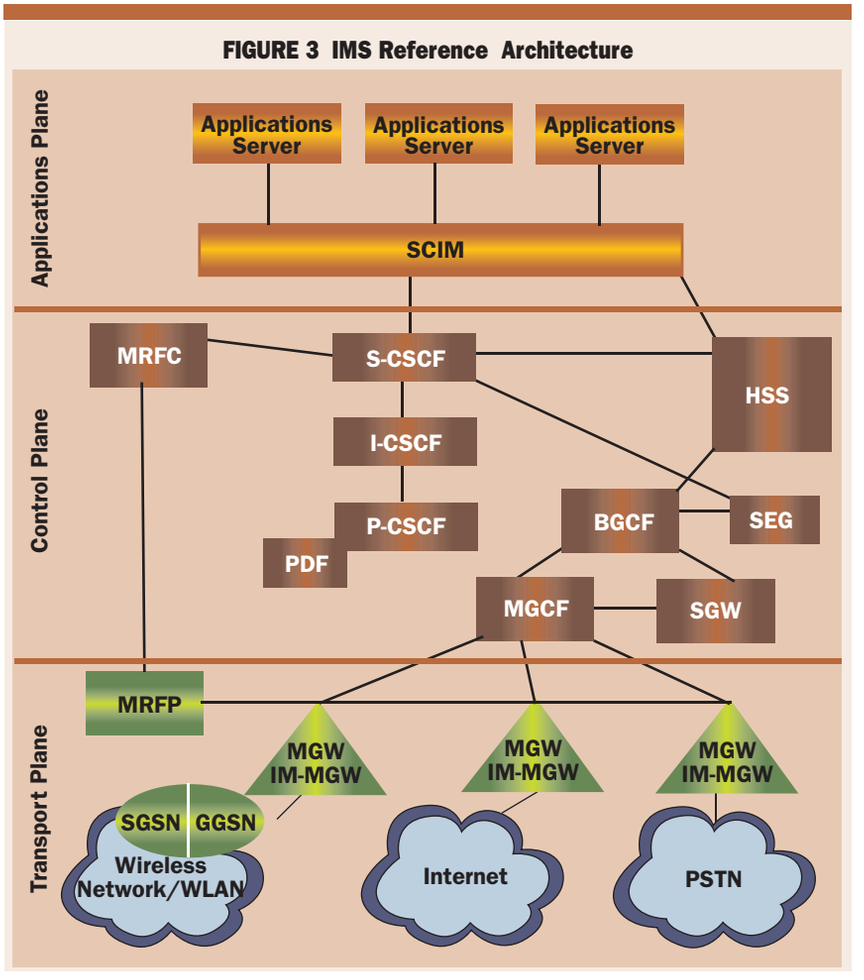

Device interoperability and service interworking are—not surprisingly—the big challenges

FIGURE 4 IMS Vendor Landscape

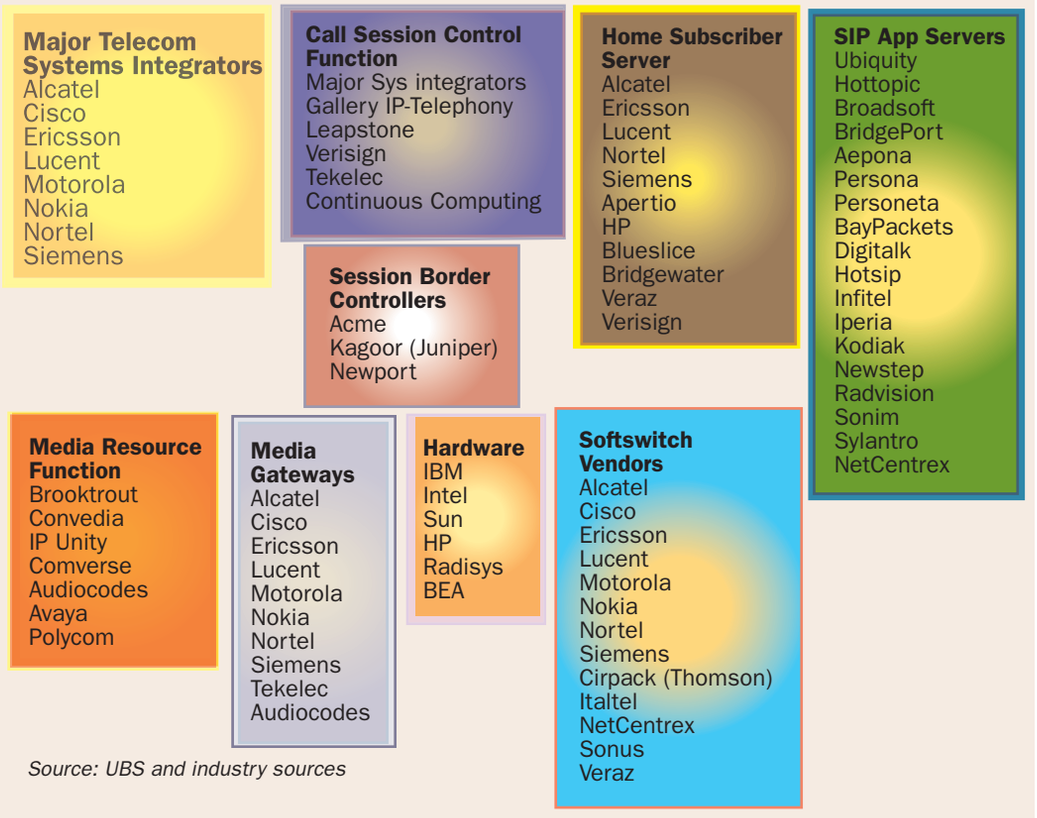

Source: UBS and industry sources

HP, etc., but so far the NIs have been the prime contractors in all IMS contracts.

Lucent, Ericsson and Sonus in particular have some notable successes and have won many high-profile IMS contracts to date. All these vendors invested in IMS development starting in 2003–2004, and are now fielding product lines that cover most of the core IMS functions (including Call Session Control Function—CSCF; Home Subscriber Server—HSS; Media Gateway Control Function—MGCF; and Media Gateway—MGW). They have also built up their Network Integration capabilities and service organizations.

Vendors have engaged in much marketeering around nuances of interpretation of the IMS architecture and future product road maps; this in turn has slowed time to deployment and acceptance. Naturally, vendors are eager to show that their IMS offerings are stable and distinctive; in many cases this means casting existing products in IMS language and providing distinctive (i.e., non-interoperable) functions.

Though this "standardize while jockeying" phenomenon isn't unique to IMS, it is very visible and confusing just now. IMS's sheer scale and complexity makes this particularly troublesome, to say nothing of the carriers' challenge in migrating smoothly from existing network cores to IMS. Rigor in compliance, test coverage, device interoperability and service interworking will be sorely needed as IMS rolls out.

In our work performing IMS readiness tests in labs and operational networks, we've observed two main kinds of problems that delay deployment of IMS-enabled services: device interoperability and service interworking. These aren't surprising: SIP interoperability alone is a big issue, and service interactions have troubled the telecom industry for decades.

The good news is that there is reality to interoperability, and service composition has been demonstrated for real in a number of labs and pilots. Service providers continue to rely on their vendors to provide them with pre-integrated, reliable solutions, but integration work is brought to a whole new level of complexity with IMS. IMS promises plug-and-play services, but this is far from the current reality; achieving this will require new development and integration practices.

It's About The Applications And How They Perform

No major architecture or infrastructure change is complete without some controversy about the driving applications. IMS is classic in this regard; like "triple-play," there is a focus on bundling, and many carriers regard the killer app as whatever becomes "the second app,"—i.e., the add-on

service that the customer purchases—since multiple services are well-proven to enhance customer stickiness.

Some carriers, notably Verizon, envision custom packages of services for enterprise customers, and specialized applications for users such as real-estate brokerages. Others predict that opening the network to third-party application developers will generate a tsunami of new yet-to-be-imagined applications, akin to the application explosion happening on the Internet.

There's no shortage of new application ideas for IMS, yet the initial products and deployments largely focus on providing solid PSTN basic features and building a migration path. Voice and data messaging—shared across different access methods—is a prosaic but powerful next step. Video and presence are common new applications being built out even within some "pre-IMS" products. Two particular applications on opposite ends of the service-complexity spectrum—video on demand and ringback tones—are appearing on the maps of multiple vendors and service providers.

The application categories covered by IMS aren't new, but their sheer breadth is: push-to-*X*, instant conferencing, messaging, gaming, personal information management, presence, location-based services, IP Centrex, voice self-service and many more. There is even a new term called "combinational services," which actually has two meanings, depending on the person using the term. "Combinational services" can either be:

- Services that combine circuit-switched voice and IMS core services as a transition step to "pure" IMS; or
- Services that combine capabilities such as presence and location-based services into specific offerings.

Each application category has its own nuances. In many cases, information services compete with Internet-based services delivered to the phones' browser. Applications that require responsiveness (low customer-perceived latency, or CPL) are of particular value—and also are particularly difficult to implement.

Push-to-talk features, which allow a mobile phone to function as a two-way radio, are an interesting case study in IMS apps. They are being standardized under the term PoC (Push-to-talk over Cellular), and have been generalized into "push-to-*X*", where *X* includes talking, sending a snapshot, messaging, conferencing, video, etc. Nextel (now merged with Sprint) introduced a PoC feature in 1997 with its proprietary iDen architecture. This feature became remarkably popular, and other carriers hope to break into this area and provide cross-carrier interoperability.

PoC became an early target application for IMS in late 2003 as part of IMS release 6—after a bit of a false start around IMS-based instant messaging and presence (IMPRES). However, the initial PoC approach suffered from poor performance due to the complexity of its signaling (requiring as many as 14 round trips to set up a call). Users expect very quick connections from this kind of service, and the system didn't provide this responsiveness (also see *BCR*, January 2006, pp. 19–20).

The Open Mobile Alliance is defining PoC as part of IMS, and a first version of the OMA PoC standard was finalized in first half of 2005. "Pre-standard" versions of PoC are rolling out now, including the first implementations that perform well. Compatibility between these versions is still lagging, and different phones work with different services.

Even the "IMS compliant" versions of PoC from different vendors haven't yet been demonstrated to work together. This lack of interoperability is exactly what IMS is aiming to resolve, and there is definite progress—but we aren't there quite yet.

Will there be a large field of independent application providers? To date, a handful have appeared. OpenEra, one of the startup specialists in IMS mobile client applications, was recently acquired by NMS (formerly Natural Microsystems). Their services include PoC, P2P video sharing, IM/chat, and an "active phonebook." Vendors such as Qualphone and Flextronics that have been providing client frameworks and applications for 3GPP are competing alongside IBM and BEA, who offer applications together with their application servers.

Today, most applications are still quite vendor-specific. The "best of breed interoperability" vision is a long ways off, as evidenced by the number of interoperability issues we've witnessed. Most vendors still link their applications to a specific application server. But the existence of independent application providers is a promising sign.

Where Are The System Bottlenecks?

There are several bottlenecks in the IMS architecture

In a sense, IMS is a "swing of the pendulum." Traditional networks were intelligent in the core and dumb at the access, with the "bottlenecks" in the core—bottlenecks in many senses, including capacity, interoperability and service development and delivery. Some networks have evolved to intelligence in the endpoint and a dumb, yet fast, core network where the bottlenecks of policy, security and provisioning are seen at the edge.

IMS has "intelligence at the access," and attempts to have a single application network. This is complex: In order to deliver a "seamless user experience," the network senses the intelligence and capabilities of the endpoints (handsets, terminals, set-top-box, gaming console, etc). The theory is that there are no centralized bottlenecks in this decoupled arrangement. A disaggregated network can optimize the main function of each service, including broadcast IPTV, wireless voice, video telephony, instant messaging, etc.

The IMS business case remains unproven

Today, there are numerous bottlenecks in IMS. For example, the Home Subscriber Server (HSS) is involved in every registration, and is consulted on location changes; HSS databases are likely to be huge and at risk for being slow, and the coordination of stateful (presence) and non-stateful (find-me/follow-me) services is complex.

Another area with performance challenges is the border functions (in IMS-speak, I-BGF, A-BGF, and P-CSCF/PDF functions, implemented in session border controller products). Encryption and transcoding, both very compute-intensive, are handled in the SBC, and many details of this area in IMS have yet to be worked out.

A third example is the service capability interaction manager (SCIM), which acts as a broker between various application servers and the control plane functions; the message load on the SCIM for some applications is likely to be quite high, and latency is important for this function.

Overall, IMS is expected to provide more multimedia real-time services, resulting in performance demands on the IMS elements as well as functional, interoperability and security requirements. Services done in IMS tend to have higher complexity than their conventional counterparts, as a result of having more functional elements and delivering access independence. Performance of IMS applications, especially latency for signaling and media, will continue to be a focus area for quite some time.

Service Provider Activities And Experiences

Plans for IMS deployment have been announced by approximately 35 carriers to date. A few examples of announced IMS service plans are shown in Table 1.

Most of these plans are multiyear rollouts, starting with small-scale trials. Technical trials have been under way at several service providers for up to a year. However, there are no "real life" carrier examples of IMS yet deployed.

In general, Tier 1 service providers are moving toward centralized control approaches. Tier 2/Tier 3 providers are largely waiting for IMS to become clearer and more affordable, but those that are rolling out services are moving towards a peer-to-peer (intelligent gateway) control architecture. Many providers want to carry their branding and identity into the devices and services (so SIP phones would have carriers' names on them, as cell phones do today).

To date, nearly all the trials have been primarily single-vendor, primed by one of the major network equipment manufacturer/integrators, such as Lucent, Siemens, Ericsson or Nokia. The level of experience and the satisfaction with lab testing and early trials varies. However, there seems to be a consensus on the need for some key capabilities: Inter-realm access (ability to connect to multiple networks with the same phone), real-time billing, security, transport QOS, and database-level integration for operations support systems. The "converged core and divergent access" approach—a basic premise of IMS—is being embraced at nearly all these service providers.

Unlike other communications architectures which have remained localized (e.g., GSM, CDMA, HSDPA, packet cable, ePON, AIN, ISDN, etc), IMS is proving to have global appeal to traditional service providers.

While the IMS technical trials and work are ongoing, the business challenges in IMS are also being tested and explored. Fundamentally, the IMS business case still is unproven—there isn't yet a solid business case for savings in opex and capex. Many providers are working toward developing and proving out this case, but most are planning IMS deployments on an "architectural" rationale. Some service providers feel that this is their "last chance" to make a major architectural change that keeps them competitive with the new breed of providers (such as Vonage, AOL, Google, etc). The strategy is to get into an architecture that lets the incumbents be competitive in rolling out new services—before erosion of their revenue stream makes an investment of this scale impossible.

The migration costs to move millions of subscribers and their services from legacy networks to the new infrastructure will be huge and must be made up in substantial new revenues from new services and savings in the cost of rolling out these new services. Finally, existing business relationships between carriers and their traditional vendors may limit the commercial deployment of third-party services.

Carriers are very interested in ways to expand their service bundles and create stronger interdependencies among the parts—for example moving from "triple" to "quadruple play" (i.e., voice, video, data plus wireless). While IMS hype centers on open architectures, most of this openness is still directed toward the carriers' benefit rather than customers'. Non-IMS based VOIP solutions as provided by Yahoo, Skype, Google, etc., can pose a severe threat to IMS's success.

TABLE 1 Announced IMS Service Plans (selected)

Bell South	VOIP
Cingular	Voice, video, data and multimedia
O2	Push-to-*X*, messaging, content sharing
SBC	Caller ID on TV, dual-mode telephony
TDK	VOIP, IP Centrex
Telefonica	VOIP, IP Centrex
Vodafone	Peer-to-peer services
AT&T	Multimedia, seamless wireless/wireline offering
France Telecom	Fixed-mobile convergence
SK Telecom	IM, video phone, push-to-*X*, VOIP

IMS's advantages in security, QOS, etc., must outweigh the relative speed and innovation of these "disruptive" services.

IMS In The Enterprise

Most early IMS implementations target consumer applications. The interests of carriers and enterprises are often at odds, and IMS is no exception to this. But IMS has near-term implications for enterprises as well, most notably around enterprise access to IMS functions and dual-mode handsets.

Some carriers, most notably Sprint, focus on providing converged services to enterprises via IMS. These services—IP-VPN, hosted PBX and IP Centrex—will be an integral part of the convergence plan for Tier 1 operators over time—to the benefit of enterprises.

Support for dual-mode handsets—letting people use a single phone both in and out of the office, with handoff between an in-building Wi-Fi network and a cellular network—is probably the nearest-term IMS application most enterprises will encounter. This is an active area with "pre-IMS" approaches, and handoff remains the major technical issue. Many people expect 2006 to be the year when dual-mode phones and services appear and customers begin signing up. Both proprietary and standard methods now exist for voice session hand-off—in particular, 3GPP TS 23.806 specifies Voice Call Continuity.

A looming difficulty with IMS in the enterprise is SIP interoperability. IMS is based solidly on SIP, but the current version of SIP used in IMS does not support the many advanced features found in business phones. Even outside of IMS, this is an issue, and an initiative to develop SIP-B (SIP for business) has run aground amid the complexity and vested interests involved. Today, a SIP phone which works with IMS networks often doesn't even register correctly when plugged into an enterprise IP-PBX.

Summary

IMS promises delivery of new applications across multiple devices, media types and locations. Service providers of all types (especially wireless carriers) are actively progressing down the IMS path, in search of a more flexible, modular, access-independent application delivery platform. Initial focus is on consumer services, but implications for enterprises are showing up sooner than expected, notably with dual-mode handsets and enterprise access into IMS functions.

Initial products are largely evolutions of current IP-oriented products to encompass one or more of the IMS functions, as would be expected. Early products and pilots have shown real functionality and benefit, even amid device interoperability and service interworking issues.

Done right, IMS should enable faster deployment of new, content-rich services for both consumer and enterprise customers. Done wrong, it can prevent vendors and service providers from launching new products and services on time, and wreak havoc on service facets such as dial tone, voice quality, security integrity and transaction reliability. The transition from today's networking architecture to IMS is sure to be bumpy for some. Thorough interoperability, security and performance testing should be a hallmark of good service rollouts.

Whether IMS will be a technical or commercial success remains to be seen, but there is life: real products and real pilots with real problems□

The transition will be bumpy for some. But it is under way

Companies Mentioned In This Article

AT&T (www.att.com)
AOL (www.aol.com)
Accenture (www.accenture.com)
Alcatel (www.alcatel.com)
BEA (www.bea.com)
BellSouth (www.bellsouth.com)
Cablevision (www.cablevision.com)
Cingular (www.cingular.com)
Comcast (www.comcast.com
Cox Communications (www.cox.com)
Ericsson (www.ericsson.com)
Flextronics (www.flextronics.com)
France Telecom (www.francetelecom.com)
Google (www.google.com)
HP (www.hp.com)
IBM (www.ibm.com)
Lucent (www.lucent.com)
NMS (www.nmscommunications.com)
Nokia (www.nokia.com)
O2 (www.o2.com)
Qualphone (www.qualphone.com)
Siemens (www.siemenscommunications.com)
Skype (www.skype.com)
Sonus (www.sonusnet.com)
Sprint (www.sprint.com)
Swisscom Mobile (www.swisscom-mobile.ch/scm/scm_home-en.aspx)
Telecom Italia Mobile (www.company.tim.it)
Time Warner Cable (www.timewarnercable.com)
Verizon (www.verizon.com)
Vodaphone (www.vodafone.com)
Vonage (www.vonage.com)
Yahoo (www.yahoo.com)