Beyond IMS

Unconventional Wisdom
Tom Nolle, Founder and President, CIMI Corporation

Introduction

One of the most significant technology shifts that have accompanied "convergence" is the shift from conceptualizing service features as attributes of network devices to conceptualizing them as something hosted "above" the network. Feature hosting is an integral part of the whole convergence concept, in fact.

But what do you host them on? How does feature hosting create a service experience as reliable as the experiences offered by the PSTN, the world standard in communications services? Can the level of stability offered by telephone switch software "generics" of the past ever be replicated in the seemingly disorder space of application development? These are questions that the industry clearly needed to answer once "convergence" was accepted as a goal, and it would be nice to say that there was an explosion of innovation created to provide those answers. Nice, but not true.

The PSTN had a conceptual model called the "Advanced Intelligent Network" or AIN that defined the role of functional components like a Service Switching System, a Service Control Point, and an Intelligent Peripheral. There was no architecture defined for a converged world, and so convergence players apparently had a problem getting a handle on how the whole thing was supposed to come together. One can almost imagine all those vendors, desperate for guidance, floundering about in their labs. Then, when all was dark, came IMS.

The IP Multimedia Subsystem is perhaps the closest thing to transcendentalism that exists in networking. It's a concept that captivates millions and has deployed essentially nothing. We can easily say that it contributes heavily to planning and zero to revenue (and at best next-to-zero in real investment). But what IMS did was create a framework in which people could visualize feature-creating applications. From IMS roots came the central element in feature hosting today, the concept of the service delivery platform or SDP.

SDPs are a kind of server, and of course servers have been increasingly a fixture within network for delivery of content, hosting, etc. One might reasonably ask what the difference between an SDP and a server is, and the answer would be hard to provide. Even if we use the strict IMS definition of SDP, about the only thing you might assert as an SDP-defining characteristic would be NEBS compliance for carrier installation.

There is no question that IMS promotes the notion of SDPs with its notion of signaling-driven feature applications. There is considerable question of whether IMS is a sufficient mission for SDPs. Why not call them "IMSDPs?" There are clearly more services than those defined by IMS, particularly in a time when IMS isn't really defining any services at all in a revenue sense. Linking the concept of SDPs to IMS is like linking the delivery of milk to demand a year down the track. It might show strong planning, but it shows poor revenue realization.

There are already indications that at least some vendors are trying to look beyond IMS in their SDPs, and to create broader architectures ("service delivery frameworks") to embrace more opportunities and justify more sales. Reality always wins, and that means that the current IMS fixation for SDPs will give way to something more rational. What might that be? That will be our topic for this paper.

Let's Start with "Delivered" Services

Why call something a "service delivery platform" if it can't deliver services? The whole media thing about "new product categories" seems to have trapped SDPs in an IMS conceptualization, despite the fact that (ironically) the IEEE Communications Magazine April 2007 edition had a nice article characterizing SDPs in such a way as to make it clear that there was more to them than IMS.

There is, but the IEEE characterization of SDPs is still rooted in technology ("network" or "IT" SDPs, for example). The providers aren't in the technology business; they're in the services business. Thus, it sure seems logical to start our discussion with a kind of taxonomy of the services that an SDP might be involved in. For that, we'll propose this simple outline, which we must make clear we do not propose to be exhaustive in terms of examples:

1. Signaling-Coupled Features of Services
    a) IMS-linked applications/features.
    b) SIP/VoIP features, including PSTN interconnection.
    c) Voicemail and redirection.
    d) Web portal features.
    e) IPTV channel/media selection.

2. Intrinsic Service Components
    a) Identity management.
    b) Digital rights management.
    c) Mobility management.
    d) Presence management.

3. Hosted Content and Experiences
    a) Software as a Service (SaaS).
    b) Content delivery/download.
    c) Storage and backup services (network drive services).
    d) Grid computing services and application hosting.

4. Connection Services
    a) Enterprise "converged" data services.
    b) VPN and VPLS services.
    c) Wholesale trunking services (also perhaps used by larger enterprises).
    d) Traffic management and performance management services.

5. Service, Network, Operations, and Business Management Services
    a) Network operations management, including EMS/NMS and NOC.
    b) Service management, including fulfillment and assurance.
    c) Operations management (OSS/BSS) and eTOM.

From this list, we can safely say that the original mission of SDPs was a part of our 1a category above, and that some players have expanded that mission to include 1b and 1c also. We would also note that some SDP vendors may also indicate that some or even all of our category 2 applications are within their scope. So far, we've seen little vendor support for the inclusion of category 3.

Another way of looking at this is to say that the "classical" conception of SDPs has been the first classification, gradually evolving into the second. The "classical" conception of service provider server missions has been the third category, gradually expanding to embrace the second. Many of our readers will recognize this shifting process as being the SDP analog to the "middleware" emergence we've talked about. The IMS-driven conception of SDPs is more "network-like" and the content/application conception more "IT-like", with the middle zone being a place where both extremes will eventually have to live.

It might be tempting to say that the classifications above are listed in decreasing order of "real-time-ness", but that's an oversimplification. In each of the three classifications, there is a range of performance sensitivity, largely created by a combination of the need to support a lot of interactions and by delay expectations set within the delivered service. Thus, it would not be reasonable to say that there were **any** SDP missions that did not have stringent performance requirements, though it might be possible to reduce their impact on some platforms by distributing the missions between more real-time products (for the really performance-intensive stuff) and more traditional servers (for the rest).

Each of the service types also has an API set associated with it. Obviously the signaling area uses APIs that are linked to the particular signaling framework in use, which could range from Parlay or Diameter for voice to standard HTTP/web access, Web 2.0/Ajax, or web services for data. In the second classification, the general view is that web services would likely be used, and in the third it's likely that a standard TCP/IP API (again, web services, or perhaps a standard sockets interface) would be needed. These specialized APIs are often the primary differentiator of SDPs versus servers, though most people probably realize that all of the SDP APIs are already available on server platforms.

But what about our fourth and fifth categories? Here is the place where we think the conception of "service delivery platform" is most clearly deficient. Despite the fact that every carrier service involves a network is some way or another, that every carrier network demands operations support, and that standardizing network/operations integration is probably critical to most providers, there is little or no acknowledgement of either "networks" or "operations" in the SDP designs of today. Similarly, there is no specific support for either network or operations awareness in servers used in provider missions. **The largest unsolved problem in service delivery is the integration of the delivery application(s) or framework(s) with networks and operations**. We submit that this is therefore the area where SDPs must define and differentiate themselves.

A Slightly Different View

Another way to look at the situation with SDPs is shown in the figure below, which we could title "SDP Worlds in Collision". The shapes on the left represent the service technology distribution of the market, and those on the right the convergence opportunity distribution. The "opportunity percentages" shown represent the percent of the five-year revenue total in each of the categories. The mapping lines show the way that technology and opportunity link.

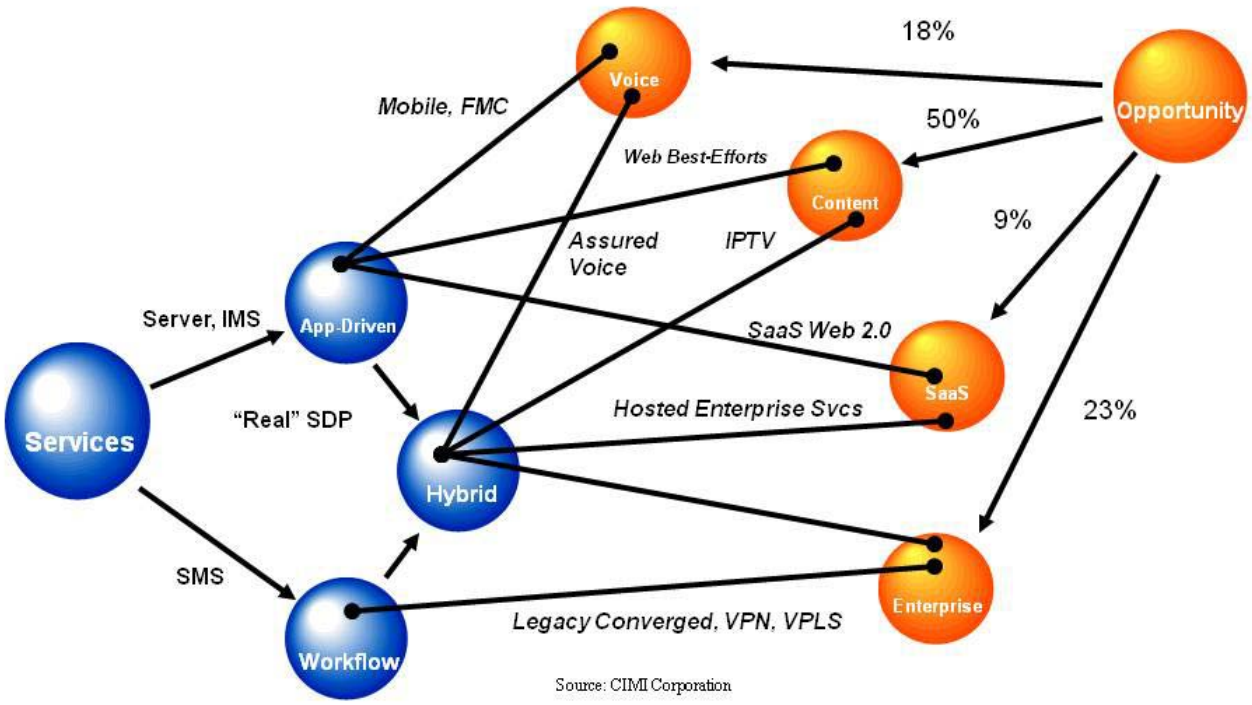

The application-driven component of the services space, which represents both traditional web server and basic IMS capabilities, is in our view largely limited to the best-efforts part of the market owing to the lack of

coordination between the service platforms and resource control. While this still provides a link to a large chunk of opportunity, the linkage is first dependent on suitable best-efforts transport (which appears increasingly unlikely) and second is likely to be the focus of non-carrier or "overlay" providers. Only the mobile space offers the carrier much hope for direct revenue from this path.

The service-management-driven component of the service space is largely focused on the enterprise service set, particularly the long-contract data services such as the converged form of frame relay or ATM and the successor services like VPNs and VPLS. These services are conceptualized around workflow-based automation of service processes, such as those that are defined (to varying degrees) by bodies like the IPsphere Forum and TMF. This represents a full 23% of the addressable opportunity and it is the only way that opportunity can be addressed. This element is therefore absolutely critical in any service strategy.

The hybridization of application-driven and service-management-driven visions of services addresses all of the opportunities, and it is toward this goal of hybridization that we believe service delivery platforms must evolve. To create a hybrid, you have to remake a little of each of the things you start with, and to do that you have to understand a little about what those things are.

Application-driven service experiences really evolved from the notion of the Internet, from a vision of the network of the future as having either all best-efforts service or as having simple grades of service activated by packet tagging at the edge. This vision had developed major stress cracks as early as the mid-90s because the bill-and-keep model of the Internet discouraged QoS cooperation by eliminating settlement among participating players. However dumb the concept might have been at the business level, it had appeal to application types because it made the process of network resource marshalling easy; you just tagged packets.

The other side of the coin is kind of the opposite, as is fitting. Enterprise data and voice/data services have long been based on explicit provisioning and operations management. Connection even within a provider network tends to be explicitly managed, and where multiple providers are involved there is explicit settlement. All of this adds up to a meticulous awareness of the traffic flows the service generates, and increasing concern over how that awareness would translate to converged infrastructure. The particular issue has been controlling operations costs, which already account for more than triple the percentage of total cost that capital equipment offers. Service management systems evolved as a means of automating service provisioning and assurance. This is most easily visualized as a "workflow", simply because the manual processes that have traditionally supported these services have had workflows. However, we must point out that while the seeds of the workflow notion lie in past practices, there have been no alternatives to that approach suggested that have shown any credibility.

Hybridizing two things this different has to start with something, and the logical way to start is by determining what's common with the two extremes. It turns out that's pretty easy—the answer is "the network". The central requirement for hybridizing application- and workflow-driven visions of service/opportunity mapping is to develop a common mechanism for controlling the network, or more generally of controlling service resources. Not only does this minimize the risk of redundant development, it rationalizes the fact that there is already a different vision of at least the network part existing between the application- and workflow-driven models.

The application side of the service framework world (as embodied in the ITU NGN activity) has taken to the notion of resource control as the function of a separate functional layer below the SDP, called the Resource Access Control Facility (RACF). The problems with RACF are that (first) there isn't any, and (second) that nothing practical has been done to try to link this approach to workflow-based provisioning. Thus, there is a potential collision between application-arbitrated resource requests and those created through workflow-based provisioning.

Absent a practical RACF (and, truth be told, we don't really even know if there is a specification for such an animal), even application-based resource requests could create collisions if there are multiple SDPs, which is currently more the rule than the exception. It doesn't take a rocket scientist to realize that no real at-scale IMS implementation could possibly be hosted on a single server. In addition, every new application seems to have its own platform.

How would applications create services that demand real pan-provider, multi-technology and vendor, reserved resources? The IPsphere Forum launched an activity aimed at rationalizing IMS and IPsphere resource usage a year and a half ago, at the insistence of service providers. The Telemanagement Forum's Service Delivery Framework activity seems aimed at some of the issues surrounding IMS, but the scope of the activity and the relevance it will have to all of the service types we've defined can't be assessed at this point. Overall, this is clearly not a problem that providers believe has been solved.

There is also the issue of what could be called "service orchestration", meaning the creation of functional and cooperative service bundles. FMC is an example of this, as are presence-based services. If all of these services are to be coordinated, there has to be an SDP strategy to integrate them. Where is it? There are tools that perhaps could be used by application developers to synchronize, but there's not really even any insurance that they are supported by all the SDPs.

We think all of the problems of hybridization, and therefore all of the problems associated with completely addressing revenue opportunities for converged networks, can be traced to deficiencies in the conceptual role of the SDP. It comes back to what we said at the first; if this was supposed to be a private IMS party, why not call the box an "IMSDP"? Application players have oversimplified the part of the problem that's out of their scope, and that creates an opportunity in the SDP space that cries to be addressed.

When, Oh When, Will the Grief End?

…and perhaps as importantly, who will end it? We've seen and heard the SDP pitches of most of the major players, and frankly they're not particularly insightful. What is driving "SDP marketing" today is the notion that there's an SDP market and not any specific vision of what that market might demand. We've listened patiently to players tell us why their SDP is "best" without ever even mentioning what their SDP did in the way of service delivery! Is that a given, as far as they're concerned? Is IMS such a runaway winner that there aren't even other approaches in the race? If so, how do you explain the fact that IMS is essentially not there in a deployment sense and that elements of the IMS architecture (like RACF) aren't even there in a mature standards sense?

The mobile players like Alcatel-Lucent, Ericsson, Siemens-Nokia, etc. are unlikely in our view to come up with any rational approach. The IP infrastructure players like Cisco and Juniper? Well, Cisco has vacillated between praising and condemning IMS, but its latest announcement of a predictable voice path capability in its 12k routers seems like it might be morphed into a rational data plane strategy. Juniper bought Kagoor, but nothing has come of that investment. Neither Cisco nor Juniper has announced an implementation of committed to support IPsphere's vision of IMS handling. Avici's Soapstone business unit has promised virtual RACF and IPsphere support, but it's not clear at this point whether they see this as the kingpin of an application/workflow hybridization.

As stupid and manipulative the whole idea of "starting a new product category" is, and as totally discredited the very notion of "new category" seems to be based on the outcome of the bubble, we still seem to be locked into the idea. All that matters to the media about SDPs is the name; if you have one it's good enough for all practical purposes. Well, it's good enough for PR purposes, but it may be that it's good for no practical purposes at all…yet.

About the Author

Tom Nolle is the founder and president of CIMI Corporation, a telecommunications, media, and technology consulting and research firm since 1982. Tom is a software architect and developer by background, but currently focuses on strategic market, product, financial, and regulatory issues. He's also the Chief Strategist for ExperiaSphere, an open source initiative that focuses on development of Java-based service logic and service management frameworks for next-generation networking. Tom is a contributor to Network World, the No Jitter CMP publication, Telecommunications Magazine, Tech Target Search Telecom, and Internet Evolution. He's also the author and publisher of Netwatcher, the oldest continuously published networking publication in the industry. He's a 25-year member of the IEEE Communications Society, a member of the Telemanagement Forum, and a speaker at many industry events.

About the Webtorials® Editorial/Analyst Division

The Webtorials® Editorial/Analyst Division, a joint venture of industry veterans Steven Taylor and Jim Metzler, is devoted to performing in-depth analysis and research in focused areas such as Metro Ethernet and MPLS, as well as in areas that cross the traditional functional boundaries of IT, such as Unified Communications and Application Delivery. The Editorial/Analyst Division's focus is on providing actionable insight through custom research with a forward looking viewpoint. Through reports that examine industry dynamics from both a demand and a supply perspective, the firm educates the marketplace both on emerging trends and the role that IT products, services and processes play in responding to those trends.

For more information and for additional Webtorials® Editorial/Analyst Division products, please contact Jim Metzler at jim@webtorials.com or Steven Taylor at taylor@webtorials.com.

Webtorials Briefs
Vol 2, Number 8

Published by Webtorials Editorial/Analyst Division
www.Webtorials.com

Division Cofounders:
Jim Metzler
jim@webtorials.com
Steven Taylor
taylor@webtorials.com

Professional Opinions Disclaimer
All information presented and opinions expressed in this publication represent the current opinions of the author(s) based on professional judgment and best available information at the time of the presentation. Consequently, the information is subject to change, and no liability for advice presented is assumed. Ultimate responsibility for choice of appropriate solutions remains with the reader.