CONSULTATIVE EDUCATION IN GLOBAL TELECOMMUNICATIONS
Technical Education | Management Consulting | Media Development

The Femtocell Revolution

Steven Shepard
Shepard Communications Group, LLC
+1-802-878-0486 *Office*
+1-802-238-1007 *Mobile*
Steve@ShepardComm.com

The Femtocell Evolution

When data networks first came into prominence they shared a handful of defining characteristics. They were wire-based, not wireless; they assumed that all signaling functions would be performed by SS7 deep in the core of the network; they took it for granted that end-user devices had no innate intelligence and that any required networking intelligence would be centralized in the core; the network, by *everyone's* definition, terminated within the central office or, in some radical definitions, at the protector box on the side of the customer's house or at a business premise; and they assumed that the overwhelming majority of the data transmitted would be transactional in nature, low-speed and shared between fixed devices.

What a difference a couple of decades makes. From those early (and by today's standards, charmingly naïve) days of networking, much has changed – in fact, every one of those benchmarks has been cast aside. While the wireline infrastructure still plays and will continue to play a central role in data networking, access is increasingly mobile and therefore wireless. Signaling, once the exclusive domain of SS7 and executed in the network core on behalf of "dumb devices," is increasingly performed at the edge – in fact, within the user's mobile device – and is often SIP-based rather than SS7-dependent. The network no longer ends at the walls of the central office or at the box on the side of the house, but today extends all the way into the customer's home or office, usually at the request of the customer in exchange for better and more granular service quality. And the transmitted data? More and more it is richly contextual, multimedia-based, bandwidth-hungry and two-way in nature, and more often than not transmitted between mobile devices.

This evolution has taken place for all of the right reasons, some of them painfully wrenching for service providers, all of them necessary. The driving force that is most germane to the subject of this paper is the inexorable and entirely predictable shift in the nature of customer demand.

When network-based services were delivered from the core of the network, the ability to customize them was extraordinarily difficult and prohibitively expensive. The ability to create customer-specific content was simply not within the budgets of most carriers, so "mass customization" was the norm – a small number of core-based features from which customers could pick and choose.

As time went on and the sophistication level of the modern network improved, network designers began to realize that they were leaving money on the table because of the centralized nature of their network architectures. IP became a key design consideration and with it came what is for all intents and purposes a "functional inversion" of the network. Content, once exclusively housed on the customer's device, began to creep into the network, taking up residence on servers hosted there; and SS7 signaling functions began to leak out of the network on their way to the end user device in the form of a SIP client. Network-based content severed the user's dependency on a particular type of device – or on any device for that matter; and device-based signaling made it possible for the user to identify herself or himself to the network along with content, device, access modality and billing preferences. Furthermore, it became clear to the telcos that service customization could be done cost-effectively if three changes were to take place in the network: First, that all customization be done at the edge of the network and *only* at the edge; second, that customization be done exclusively in software; and finally, that customers be given the ability to self-provision customized services.

This transformation is well underway and is being embraced by both traditional carriers and non-traditional providers as well – Google, Apple, Yahoo and AOL to name a few. Helped along by a mass adoption of mobility thanks to low-cost multimedia devices and the strong proliferation of broadband wireless, networking has become a broadband, multimedia, two-way mobile game.

This evolution is critical to content and network providers alike. For content players it represents a new, as yet untapped market. For network providers it represents an opportunity for them to improve a flat or declining revenue stream and to increase their marketplace relevance by offering increasingly customized services to customers on an individual basis. To do this, however, providers must establish a service monitoring beachhead as close to the customer as possible, and today that means within the residence or small business location. And as network access has become increasingly wireless, it makes sense to create a service provider-manageable in-building wireless access point – not a WiFi hotspot, because WiFi does not facilitate mobility; rather, an access point that serves as a functional extension of the existing mobile network that works effectively within the confines of the home or small office.

Of course, this evolution comes with a number of challenges. For years the wireless industry has searched for cost-effective in-building solutions, but for both engineering and economic reasons have not (until now) found the magic wireless bullet.

In today's network, in-building coverage is generally viewed as an extension of the existing cellular network, with propagation patterns and coverage scenarios designed to provide a statistically meaningful level of wireless connectivity inside of buildings. Traditional cellular technology, however, doesn't lend itself to 100% in-building coverage; the cost to do so is far too prohibitive.

There is increasing interest in the development of indoor coverage solutions, and the result is the femtocell – a cost effective connectivity option for indoor scenarios that meet the economic criteria for its deployment. The economics, however, have to date been tricky. Given that a large percentage of the intended market is the home, the ability to prove cost effectiveness and manageability has been elusive.

There is good news, however. Thanks to advances in semiconductor products for the second and third generation wireless space, it is now possible to create femtocell solutions whose economics are easy to prove in and whose manageability is assured. This is by no means a guaranteed market, but its attractiveness is increasingly high.

Steps for Success

Mobile operators have one primary goal: To attract and keep customers by making them so irretrievably sticky that they have no desire to go to an alternate carrier. The rollout of a femtocell solution will measurably add to their relevance factor provided the manufacturers of femtocell technology meet three critical requirements. First, because of the home-based nature of the service, the femtocell solution must be priced below the economic "pain point" of the typical household – around $250. Second, because the femtocell concept is rolled out as a functional extension of the mobile carrier's network, it must be transparently scalable, must integrate seamlessly with the existing network (zero to one truck roll max), must be cost-effective for the carrier to deploy, and must operate within RF interference parameters for devices of this type.

Cost Management

The opportunities in this space for silicon foundries like LSI are immense. Femtocells are to the wireless industry what the set-top box is to the cable industry: The single most expensive cost element in the deployment strategy for in-building coverage. To meet the $250 cost target, femtocell designs must be stringently managed, a concept that is far from unfamiliar to established design houses.

Scalability and Legacy Integration

Today's networks rely on centralized Operations Support Systems (OSS), Billing/Business Support Systems (BSS), and Operations, Administration, Maintenance and Provisioning (OAM&P) systems to perform the behind-the-scenes functions required to cost-effectively operate a large network. The bulk of these systems were designed around the concept of a centralized, hierarchical, intelligent network model – not for a distributed, peer-to-peer network working closely with intelligent devices that originate as much content as they consume.

Integration of the femtocell model, therefore, comes with its own set of unique challenges. In traditional 3G networks, Radio Network Controllers (the 3G network entity responsible for control of all Node Bs) communicate with the Base Transceiver Station (Node B) over dedicated high-speed facilities that rely on the protocol known as Iu-b, described shortly. However, there is a caveat: For reasons of cost control, femtocell designs typically use IP connections, often the Internet, for access to the network rather than the dedicated facilities typically deployed between base stations and remote nodes. This creates a network design (and potentially, engineering) challenge for the service provider: how to cost-effectively design, build, operate and monetize the distributed, peer-to-peer network described earlier? Clearly this concept is far removed from traditional network design considerations. First, because it assumes the deployment of thousands of femtocell nodes, existing back room systems must be modified to accept and integrate them with full functionality. Failure to do so is a show-stopper of the highest order. Second, because femtocells “live” on customers’ premises, their management processes must be built around worst case scenarios in terms of device availability. Femtocells don’t necessarily enjoy access to uninterruptible power or to central office-quality maintenance as traditional base stations do; this assumption must be designed into the service logic. Finally, security considerations must be taken into account to ensure that femtocells are properly authenticated before being granted network access and to guard against the potential for network intrusion.

Finally, given the potential size of the market and the economics that will derive from it, standardized interfaces must be put into place to facilitate the involvement of third party manufacturers and to embrace what is rapidly becoming an open development environment.

Compliance with Interference Mandates

The deployment of femtocells, for all the right reasons and best intentions, can result in unintended consequences in terms of network performance. Unless dedicated spectrum is made available for the femtocells to avoid the potential for interference, RF engineers must go through the often painful process of propagation pattern design and spectrum management to avoid what could be a serious service-affecting problem. Consider something as simple as a call handoff from a 3G network to a femtocell. Femtocell standards limit the number of cell sites that the femtocell can scan for to no more than 16, a potentially serious limiting factor. If a CDMA system and a femtocell system are operating within the same spectrum, interference can occur due to power differences between the two systems. And perhaps more common than any other is the challenge of adjacent femtocells in multi-tenant dwellings that interfere with one another.

There are, of course, solutions in the works for these impairments, including using dedicated spectrum for femtocell systems and stringently controlled power on femtocells. In traditional operating environments (standard 3G networks), frequency deployment is carefully prescribed to avoid such concerns, a procedure that is manageable given the relatively small number of sites that must be managed. In a femtocell environment, however, where the number of remote nodes can easily climb into the hundreds of thousands, the process becomes untenable.

Manufacturers of femtocell devices, aware of the potential problems associated with interference-related phenomena, have done considerable research in the area of intersystem interference and have proffered a variety of viable solutions that will ultimately become part of global femtocell design and management schemes.

Approaches to Femtocell Deployment

Today there are three basic design models for femtocell deployment and integration. The first is the IP-based Iu-b interface (3GPP Rel.5), the second, a SIP-based approach (Iu/A Interface), and the third, the use of unlicensed spectrum in a technique known as Unlicensed Mobile Access (UMA). A fourth approach, which takes into account emerging IMS standards (IMS VCC), will be briefly described as well.

Iu-b BSC/RNC-Based Femtocell Deployments

In the Iu-b model, femtocells are fully integrated into the wireless carrier's network and treated like any other remote node in the network. The Iu-b protocol has a number of responsibilities including management of common channels, common resources, and radio links; configuration management, including cell configuration management; measurement handling and control; TDD synchronization; and error reporting. In Iu-b configurations, mobile devices access the greater network and its services via the Node B link, and femtocells are treated as traditional base stations.

SIP-Based Femtocell Deployments

As a central element in IP-based access and transport networks for signaling functions, SIP also stands to play a role in femtocell deployment. A SIP client, embedded in the femtocell, uses SIP to communicate with the SIP-enabled Mobile Switching Center (MSC). The MSC performs the operational translation between the IP SIP network and the traditional mobile network.

UMA-Based Femtocell Deployments

Many believe that the use of Unlicensed Mobile Access (UMA) may represent the most logical approach to femtocell deployment. In light of the ongoing evolution to an all-IP network, it certainly has its advantages. The UMA model, which has now been named Generic Access Network (GAN), offers an alternative way to access GSM and GPRS core network services over broadband. To support this approach, the standards define a new network element (the UMA Network Controller, UNC) and protocols that guarantee secure transport of signaling and user traffic over IP. The UNC interfaces into the core network via existing 3GPP interfaces including those described earlier. Its primary advantage? It supports core network integration of femtocell-based services by delivering a standards-based, scalable, IP interface for mobile core networks.

IMS VCC-Based Femtocell Deployments

Voice Call Continuity is a standard that is currently making its way through the 3GPP development process. 3GPP R7, the standard that will support VCC functionality, provides for a network design that extends an IMS network to include cellular coverage and addresses the very important handoff process. It is designed to provide seamless call continuity between cellular networks and any networks that support VoIP. VCC is a remarkably comprehensive standard that provides for interoperability between GSM, UMTS, and CDMA cellular networks and any IP-capable wireless access. IMS-VCC also supports the use of a single phone number or SIP identity and offers a broad collection of functional advantages, including support for multiple markets and market segments; provisioning of enhanced IMS multimedia services, including greater service personalization and control; seamless handoff between circuit-switched and IMS networks; and access to services from any IP device.

The attractiveness of the IMS VCC option should be fairly obvious. IMS, or something resembling IMS, is clearly the model for the future IP network. An access solution that provides seamless connectivity between existing network elements, the emerging IMS architecture and the femtocell is clearly a forward-looking solution, the attractiveness of which will continue to climb as IMS becomes more and more common.

Marketplace Potential

The potential of femtocell deployment in terms of market potential is attention-grabbing. Based on the deployment of large numbers of femtocells, service providers can expect to save between $12 and $15 annually per customer, which may not sound like much until it is aggregated. Furthermore, manufacturers of femtocell products should sit up and take notice: once begun, the deployment of femtocell architectures will only accelerate.

This concept of femtocell-based access, like the widespread rollout of IP, is a 'when,' not an 'if' question. As the point of service delivery migrates from the network's core to the farthest reaches of the network's edge, the breadth and depth of potential opportunities becomes apparent. As long as back room systems evolve to support the special requirements of this very different architecture, and as long as service providers put into place data collection methodologies to ensure their ability to respond effectively to requests for service, the femtocell model will yield enormous upward revenues. It is therefore critical that manufacturers not only pursue the best possible design for femtocell products, but that they create awareness among their service provider clients about the potential that femtocell deployment offers.

And what of the future? Standards will continue to evolve, service providers will continue to seek deployment opportunities, and manufacturers will respond in kind. Most important of all, however, are the demands that customers will continue to place on the network as mobile content, delivered to a seemingly endless array of device types, displaces voice as the largest component of mobile data.

* * *

Steven Shepard is the president and founder of the Shepard Communications Group, a Vermont-based firm that provides industry analysis, technical education, management consulting, and media development to companies throughout the world. He can be reached at Steve@ShepardComm.com.