The Agile Service Provider: A New Model For Service Delivery

Prepared by

Caroline Chappell
Analyst at Large, *Heavy Reading*

Graham Finnie
Chief Analyst, *Heavy Reading*

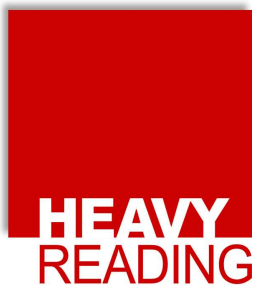

www.heavyreading.com

On behalf of

Cisco Systems, Inc.

May 2009

TABLE OF CONTENTS

LIST OF FIGURES*

SECTION I

* All charts and figures in this report are original to *Heavy Reading*, unless otherwise noted.

Executive Summary

The world's service providers are at an important turning point. Broadband is becoming ubiquitous and commoditized; network operators are replacing legacy switching and routing gear with all-IP NGNs; both networks and computer hardware/software are being relentlessly virtualized; and services are increasingly delivered from the Web and the software that drives it. This massive multi-layered transformation poses one over-riding question: what is the role of service providers in the new world that is emerging?

We believe that network and other service providers need a new strategic narrative that sets out a clear route they can follow in this new world. Their objective must be to become a central link in the value chain for the delivery of a myriad of services to end users. If they fail, they will be little more than providers of physical access infrastructure.

The purpose of this research is to help create a model for the next 3-7 years that network operators and other service providers can use to link changes in the emerging services environment to the services delivery infrastructure that underpins and facilitates those services. It is based on a detailed interview program with progressive providers of web hosting, content delivery and telco services, looking at what their key challenges are and how they are resolving them, and it also draws on an extensive range of survey work and secondary sources.

In Section 1, we look at the new service universe, and what it means for service providers. We show that, while it's impossible to predict what services will be successful 3-7 years' hence, it **is** possible to predict the kinds of attributes these services will have, and to infer from this the kind of infrastructure that will be needed. We argue that the changes required are enormous, and create huge challenges for service providers. They need to create a far more flexible hardware and software infrastructure able to handle any kind of service request much more rapidly and cheaply than today. However, with the correct strategic approach to putting such an infrastructure in place, we show that service providers will ultimately be able to enhance their ability to address new service opportunities and establish a cost base that is appropriate to this new situation.

If service providers create such a service delivery infrastructure (SDI), they may reap rich rewards. They will be able to reach new customers, on new devices, offering new kinds of services, and in particular strike new deals with the third parties that create many of the services, as well as themselves creating the applications and content at the core of the service package. New areas they will be able to colonize include B2C and B2B advertising, consumer IT services, enterprise IT and outsourcing, and the entertainment market. Improving their ability to strike worthwhile deals with third parties is perhaps the single most important advantage of a new SDI.

In order to map a route in this very complex environment, service providers need to identify a set of core Key Performance Indicators (KPIs) and develop some targets in each area. In our model, we identify four areas which taken together will provide strong evidence of a successful transition: service innovation average time to ingest new service; average cost to deliver and support each service; service range & richness (number of services per customer); and enhancement to revenue earned via third party services. Orders of magnitude improvement will be required in most areas.

In Section 2 of the report, we set out the key phases that service providers must go through to achieve the required services delivery infrastructure transformation. We show that silos dividing networks from IT processing platforms and software will be dismantled as service providers go through five distinct but overlapping stages in the transformation of IT, network and data center infrastructure: consolidation; virtualization; convergence; scaling; and federation. The paper shows the different approaches service providers may take to implementing these phases over a five year period, although other timelines could be valid, depending on the operator and its particular circumstances. Each phase can be undertaken separately, but together they can yield the massive KPI improvements that we suggest are necessary.

In the final part of Section 2, we show how each of the key phases contributes to the KPI targets set out in Section 1. Although the actual targets and timing of the phases are likely to be specific to each service provider, the general principles will apply to all.

I The New Service Universe

1.1 The Shape Of Services To Come

Predicting the future of consumer services in the information, communications and entertainment (ICE) sector is fraught with difficulty. Many recent Web and Internet service successes, including P2P file downloading, YouTube and Facebook, were not predicted by anyone, and seemed to emerge from nowhere, creating whole new service categories from scratch.

A major reason for this inherent unpredictability is the Web itself: a hugely creative service development environment, spawning new ideas daily, some of which succeed massively, some of which flare brightly and then die, and many of which make no progress at all. Only the fittest survive, and as in nature, so on the Web, predicting what features and qualities confer survival benefits is a risky exercise.

But if we cannot safely predict which new services will have caught the public imagination in 3-7 years' time, we **can** predict what kinds of **attributes** they are likely to have.

Among these, we believe, are the following probable developments:

- **Even more video content**, including more one-to-one or one-to-many video communications, more live video (e.g. concert and sport feeds), more niche content.
- **Everything on demand,** including all available video, music, gaming, corporate and personal content.
- **Ubiquitous access,** equating to
    - seamless, affordable access on any access network, wherever the customer is located, including locations that are not part of the provider's network
    - seamless, affordable access on any device, including PCs, TV displays, PDAs, mobile devices and so on
- **More (and more varied) "over the top" services** from the Cloud, in its widest sense. In time, the very term "over the top" is likely to lose its meaning as the distinction between telco and third party services blur
- **Targeted advertising and commerce** as the commercial underpinning for more and more new services—replacing or augmenting direct payment for services
- **Shortening shelf-life for certain categories of services,** as the range of services widens and services come in and out of fashion
- **More mash-ups and service blending,** in the widest sense, blurring all existing service boundaries, including:
    - Launching one kind of service from another service
    - Combining elements from one service into another service
    - Switching in session from one service to another, or adding service elements
- **More services developed by end users**, as well as widespread adaptation of existing services to suit end user needs
- **More machine to machine services** sharing information among a network of devices.

1.2 The Shape of the Service Provider Landscape To Come

For the purposes of this report, by "service provider" we mean any owner of infrastructure (broadly defined to include network, data center and software) that is used to deliver services, including network services, hosting services, content services, application services and business process-based services. Some service providers offer all these services to customers while others choose to focus on a particular service niche, such as providing hosting services or content services. Nevertheless, even niche service providers today may put together and operate a complete stack of infrastructure services – network, data center and software – to support their business.

Over the next 3-7 years, a sharper division is likely to emerge between service providers that can scale service delivery infrastructure (SDI) facilities, such as networks, data centers and software, and service providers whose primary business is to create and own content, application and business process-based services. The ability to access infrastructure facilities via the Web (or from the 'cloud', as it is sometimes called) is decreasing the need for some types of service provider to integrate and operate large amounts of infrastructure themselves.

The former category of service provider will seek to drive down the cost of managing and operating service delivery infrastructure so that they can rent out their infrastructure, as a service, to service-creating service providers. To keep development costs low and speed of innovation high, the latter category of service provider may become increasingly 'infrastructure-light', in some cases, maintaining in-house only the data center and network infrastructure that they need to support and secure their core business. Content and application service providers will rely on infrastructure service providers for many of the services they need to deliver their products to customers: at the same time, infrastructure-based service providers must be able to support any content/application/business process service provider requirement.

Some service providers will continue to offer a broad range of services, from infrastructure to content to business process services, to different customer segments. Others will specialize, renting support services from other service providers through the medium of the Cloud. Within 3-7 years, the service provider landscape is likely to become more fragmented and dynamic, reflecting the unpredictability of the Web-driven service landscape. To remain competitive, service providers that intend to remain infrastructure-heavy will need to undertake the phased transformation described in this research.

Figure 1 classifies service providers in the future according to whether they are likely to continue to own facilities (infrastructure-heavy) or move to an infrastructure-light model and whether they are likely to sell a wide range of services or a specific, niche set of services.

Figure 1: The Service Provider Universe: A 3-7 Year View

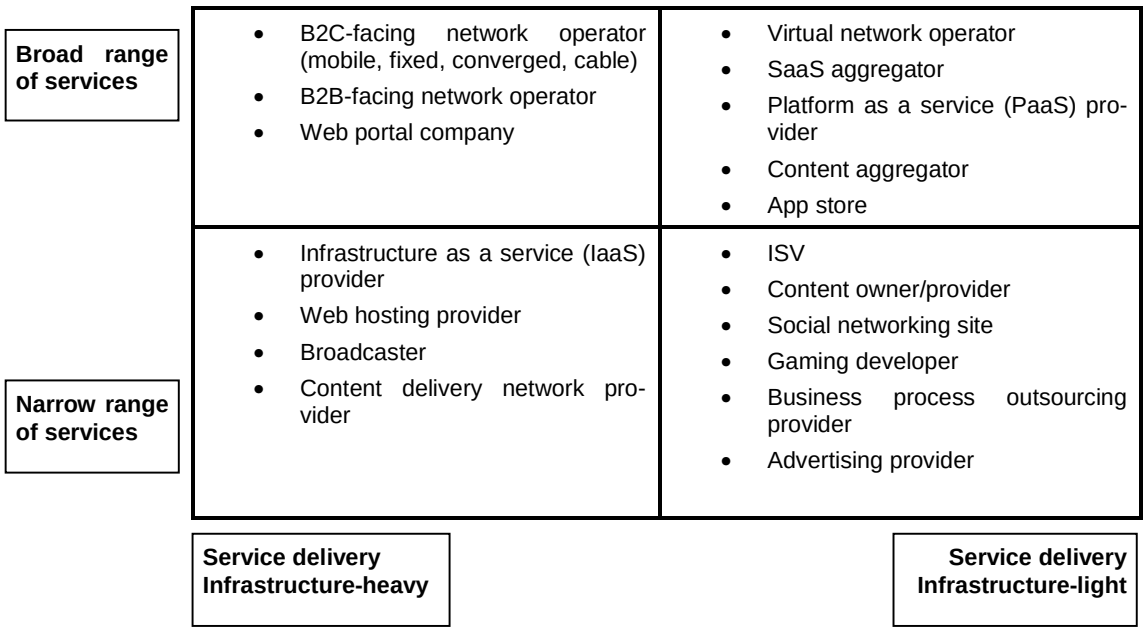

1.3 Implications for Service Providers

These changes create major challenges for infrastructure-based service providers, who must compete across a wide range of dimensions for the business not only of end users but, increasingly, of the third party content/application/business process service providers that will form a more and more important part of the value chain.

The basic problem that faces the average service provider is silos in every aspect of their business—infrastructure, service creation, service deployment, billing and support etc, which hinder their ability to support streamlined, cost-effective, flexible processes between internal stakeholders and external partners.

The above changes imply that infrastructure-based service providers must:

- Deploy a much more flexible infrastructure able to cost effectively handle sudden and unpredictable spikes and troughs in demand, and handle unanticipated applications, without the need to manually add hardware elements
- Minimize infrastructure dedicated to specific services or service tasks, e.g to deliver only specific video services, in order to increase still further services agility and improve the economics of service delivery
- Create the means to reach customers in new geographic locations, as well as existing customers on the move, with new applications and higher bandwidth services
- Ability to create (or the means to create) new applications and add features to existing applications very quickly—in hours/days/weeks, depending on application complexity, rather than months

- Ability to onboard new third party or user-created applications very rapidly—within hours if necessary, rather than weeks or months

- Deploy self-installing, self-adapting service packages—giving consumers not only the ability to add and launch services or service features directly from the portal, but also automated marketing of new services and service features based on analytics

- More flexible organizational structures—in particular, the lowering of barriers between marketing, product, OSS/BSS, IT, and network organizations.

Yet if the new service delivery infrastructure implies a large-scale transformation, it also provides enormous new opportunities to service providers. In fact, in our view it provides the **only credible** long-term way for telcos, for example, to stand any chance of increasing their revenues —short of hoping for unlikely regulatory relief. The new service delivery infrastructure (SDI) can enable them to escape from the saturated services trap that so many currently languish in. At the same time, it can help them to rebalance their cost to revenue ratio. The inflexibility of underlying technology, marketing and organizational paradigms typical in network operators prevent them from realizing their true value, which certainly could be extended beyond the provision of basic infrastructure-based services, such as Internet connectivity and bandwidth, simple communications services, and conventional video entertainment. Such service providers have many underutilized assets, and new and much more flexible models for the development and delivery of services are emerging, and offer hope to service providers that they can play a central role in the emerging service universe.

Among other things, the emerging paradigm for Information Communications and Entertainment (ICE) service creation and deployment will enable service providers to:

- Reach new customers outside their conventional service area through the Cloud
- Serve customers across any device or territory, wherever they are, increasing the perceived value of services
- Offer mainstream computing service support, along with communications and entertainment, to consumers and small businesses, taking advantage of the shift to cloud-based and thin-client computing models
- More fully engage with enterprise customers, right across the enterprise IT value chain, taking advantage of the shift to cloud-based computing and virtualization of enterprise services
- Offer a wide range of capabilities to third parties including mainstream ISVs and Web developers, including subscriber demographic information, billing and charging services, targeted advertising, and so on.

These changes will allow service providers to participate in a much larger revenue pool, including revenue areas that are growing more quickly than their traditional service areas. Figure 1 (next page) shows the opportunities open to network operators, as just one example. Reviewing the Information, Communications & Entertainment (ICE) market as a whole, we can see that this kind of service provider has a strong position in only three areas—fixed telephony, mobile telephony & SMS, and Internet access. All are essentially basic commodity services where operators are vulnerable to market saturation, substitution, and price-led competition. In other areas, although telcos have made some headway in providing basic video entertainment, they have been restricted to a large degree by the inflexibilities described earlier.

Figure 2: Anatomy of the Revenue Opportunity: Network Operator View

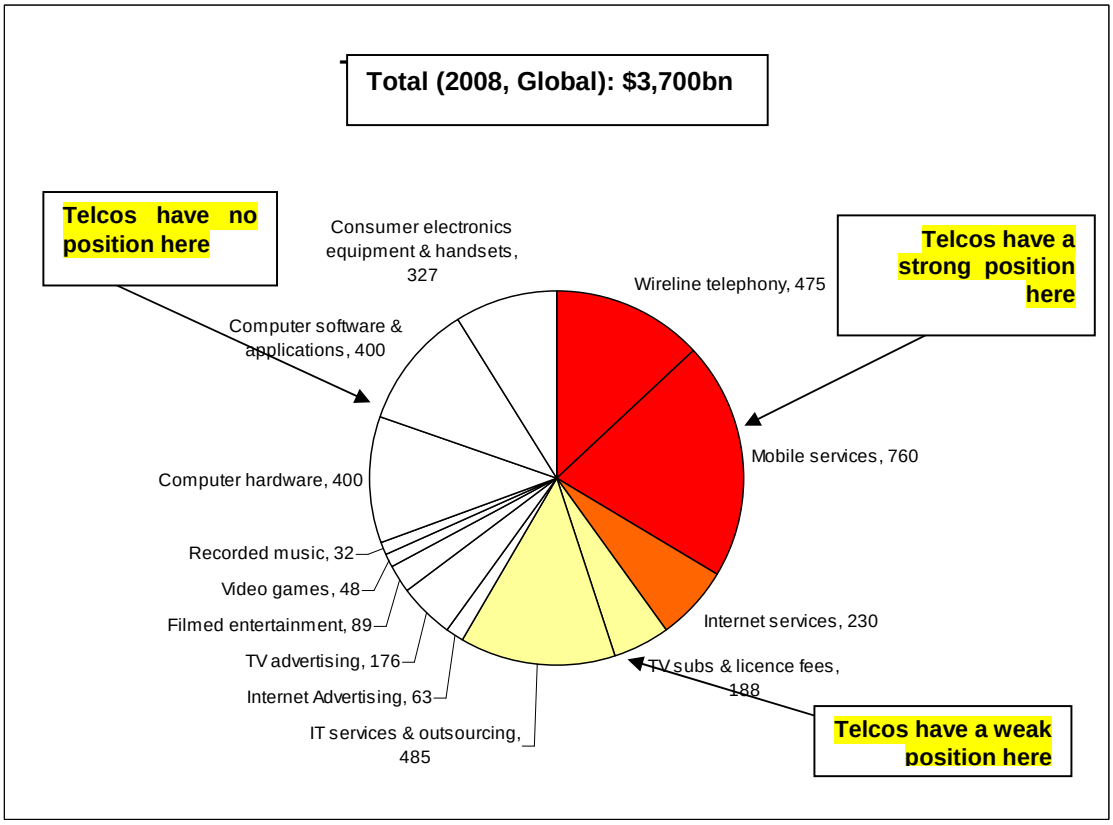

Note: graphic is color-coded to show where network operators currently dominate

In sum, the following new revenue opportunities are potentially opened up to telcos:

- **Revenue from consumer and B2B advertising,** a $500bn global market where telcos earn almost no revenue today. Unless they make it easier to work seamlessly and flexibly with third parties, this huge opportunity will likely remain largely closed to network operators
- **Revenue from consumer IT services of all kinds**, where telcos earn almost no revenue today—enabling network operators not only to share in the existing $180bn market for consumer computer hardware and software, but also to help shift consumer towards cloud-based computing.
- **Revenue from the $500bn enterprise IT & outsourcing market,** where telcos earn some revenue today, but many have struggled to get beyond basic network provision, or struggled to make a reasonable margin in this business. Note that the opportunity here includes potentially very large increases in consumer use of valuable vertical sectoral opportunities in healthcare, education & training, personal & home security, and so on.
- **Revenue from the $350bn entertainment services market**, where telcos are largely restricted at present to trying to develop their own conventional video services, and earn next to nothing from over the top video content owners, aggregators and others likely to become the dominant forces in this area.

In Figure 3, we extend this idea of broadening opportunities to other types of service providers beyond just telcos.

Figure 3: New Opportunities for Other Types of Service Provider

TYPE OF SERVICE PROVIDER	NEW VALUE CHAIN REVENUE OPPORTUNITY
Web portals	Enterprise IT & outsourcing market, leveraging large infrastructure investments to provide infrastructure and software as a service
Managed/hosting services providers	Revenue from B2B advertising, enabling the community of customers they support to advertise to each other
Content delivery network providers	Enterprise IT & outsourcing market, supporting infrastructure as a service and software as a service providers, or providing these services on top of the CDN
Broadcasters	Consumer and enterprise IT services as these can increasingly be blended and bundled with HD video content
Systems integrators	Enterprise IT & outsourcing market, Consumer IT services, entertainment services, where SIs can host infrastructure to support these segments
ISVs	Consumer and enterprise IT & outsourcing, where ISVs can leverage infrastructure investments to provide infrastructure and software as a service.

1.3 Third Parties Are The Key

It's clear from the last section that establishing better working relationships with third parties is the single most important key to success in emerging service markets. As Figure 4 shows, network operators already recognized that this was their biggest opportunity and challenge back in 2007

Figure 4: Partnering With 3rd Parties Is the Number One Telco Objective

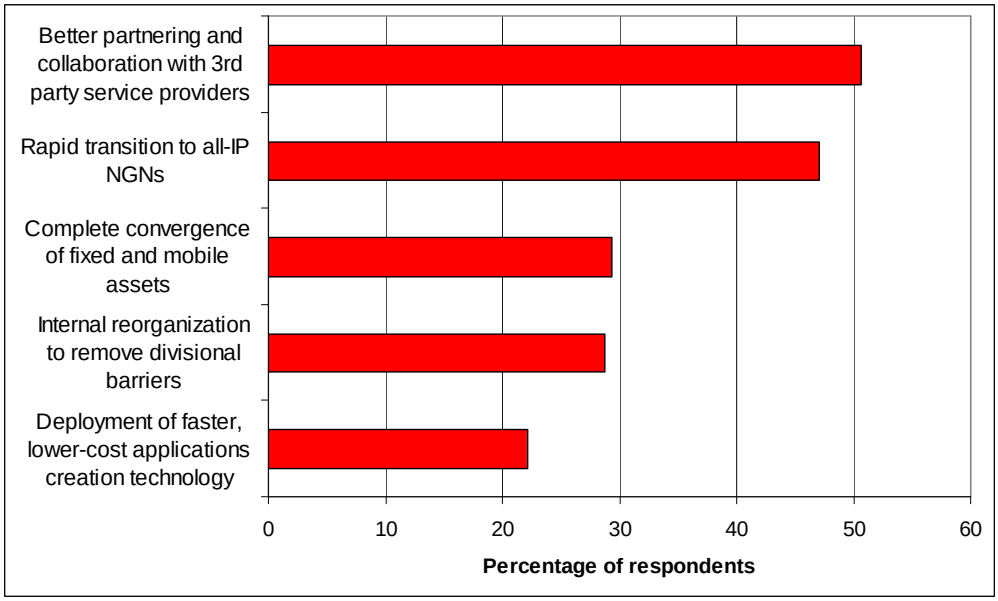

Source: Heavy Reading Survey of Next-Generation Telcos, 2007. n=118 telcos and service providers
Question: *"Which of the following will be the three most important factors to the future success of today's mainstream telcos? (please select three options)"*– **only the top 5 of 11 options offered are shown**

Many telcos have already begun programs in this area as the table on the next page shows. But these programs are really just the start, and much more will be needed by all types of service provider, not just the telcos. We believe that service providers need a roadmap that sets objectives for improvements across a range of relevant KPIs, and these are described in detail in the next section.

Figure 5: Sample Third-Party Service Programs

ORGANIZATION	NAME	PURPOSE & PROGRESS
O2	Litmus	Create very low cost and automated onboarding environment for 3rd party developers; in early stages
Vodafone	Betavine	Create open community for mobile applications developers to create and offer services; in early stages
BT	Ribbit	Initial set of APIs to be released in the first half of 2009; APIs based on other languages are also planned
Orange	Partner	Well-established program for working with 3rd party developers, with many APIs already published; however, onboarding times still lengthy
Sprint	Open Network	Program to enable 3rd parties to write to Sprint-supported handsets and devices, e.g. to facilitate access to major Web sites
Verizon	Open Development Initiative	Program to test and certify new devices and application on Verizon's network

1.4 Key Performance Indicators for Next Generation Services

Service providers can measure their progress in becoming next-generation service providers on four basic metrics, we believe: how long it takes to acquire and deploy new third party services or to enhance existing services with third party assets; the cost of delivering, supporting and managing each service; the relative complexity of services; and the revenue uplift that can be gained from adding new capabilities and serving new markets, especially via third parties.

- **Service innovation average time to market**
  - To ingest and test new third party applications quickly and easily, supporting agile and rapid entry into new markets (from weeks/months to minutes)
  - Turn capacity up/down to meet new application requirement; load, configure and test applications in minutes/hours/days rather than months
  - To take advantage of application framework to innovate new products from reusable internal/external assets (from months/years to hours/days)
  - To provide multi-tenant, on-demand service consumption, cutting time to provision from days/weeks to minutes or real-time
  - ***Objective: 100 to 1000-fold improvement***

- **Average cost per service delivered & supported/managed**
  - Reduce management, maintenance and operational costs by having a standardized environment;
  - Reduce capex and running costs through automated resource allocation and deallocation (ie higher utilization of computing and network assets);
  - Ability to recover from failure (from hours/days to sub-ms to self-healing/proactive prevention;
  - Lower energy costs per service, achieved primarily through higher utlilization of assets;
  - Automated service delivery and a high degree of customer configuration & control when deploying new services;
  - Carrier-grade reliability and QoS characteristics to reduce customer trouble tickets, based on highly adaptive and virtualized service infrastructure
  - ***Objective: 10-100 fold improvement***

- **Service range, richness & adaptability**
  - Number of services offered to customers
  - Number of third party services facilitated
  - Features offered within individual services
  - Degree of personalization possible in customer service packages
  - ***Objective: 10-100 fold improvement***

- **Revenue enhancement**
  - Number of customers/seats/identities served
  - Revenue "universe" available to network operator
  - New direct and indirect revenue earned from delivery of third party services
  - New revenues from new geographic markets
  - New revenue earned from serving customers on all networks/devices
  - New revenue earned from IT-centric consumer & business services
  - ***Objective: up to 10-fold revenue increase***

Figure 6: Translating KPIs Into An Actionable Timeline: A Simplified View

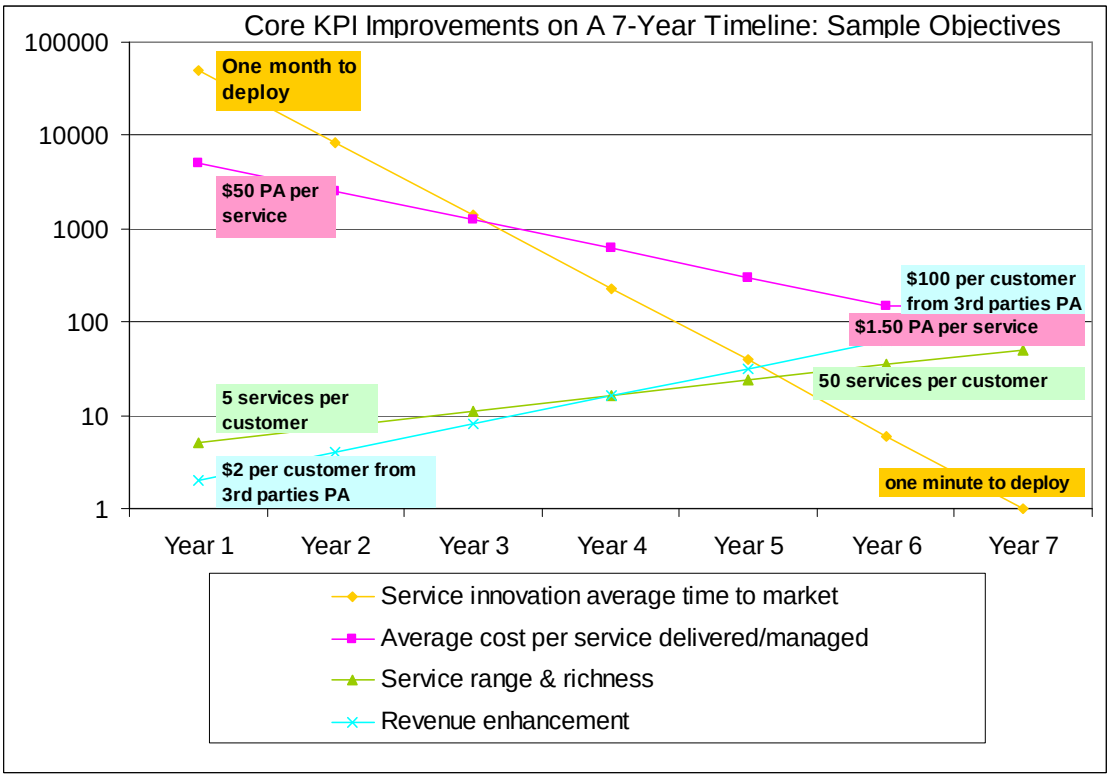

Note: the timeline is illustrative only, and may vary from service provider to service provider, depending on particular circumstances. For a more granular view, with phasing, see Figures 11 and 12 in the final section of this report

Although each network operator will doubtless want to create its own version of this list, the basic objectives, and—more important—the scale of the task ahead is clear. Business as usual, implying small incremental improvements in key indicators, will not carry telcos into the new era, except as providers of basic connectivity. **In Figure 6, we have created a simplified version of the above objectives (for clarity), in order to show that the improvements service providers must seek are exponential, not incremental.**

II Service Delivery Infrastructure Evolution

This section looks at how service providers can evolve their service delivery infrastructure to support the new service universe described in the previous section. It discusses the five phases of next generation infrastructure evolution that service providers will undertake in order to support the new order of rapid and rich service delivery at low cost. It shows how service providers progressing through these phases realise new economies of scale, reduce operating costs and create new revenues from a richer range of services.

The focus of our research is the period 3-7 years' hence, enabling us to take a medium to long-term view of likely developments and their impact.

2.1 The Five Phases of Service Delivery Infrastructure Evolution

Although it is difficult, if not impossible, to forecast the services or types of service that service providers may need to deliver in the next 3-7 years, Section 1 made it clear that the **attributes** of these services are reasonably clear, and require service providers to put in place a next generation service delivery infrastructure capable of supporting them.

However, the infrastructures that most service providers have in place today comprise numerous different and siloed elements that are difficult to manage and to expose in a coherent way to partners and customers. This limits the speed with which service providers can acquire and deliver new services, and therefore their ability to tap into new revenue streams.

Figure 7: The Evolution of Service Delivery Infrastructure

<table>
<tr><th colspan="3">OLD WORLD: NETWORKS + DATA CENTERS</th><th>NEW WORLD: SERVICE DELIVERY INFRASTRUCTURE</th></tr>
<tr><td>NETWORKS

WAN/LAN, storage area networks, network appliances, network application servers</td><td>IT PROCESSING PLATFORMS

Servers, IT appliances, specialized equipment (encoders, storage arrays, ACDs etc)</td><td>SOFTWARE

Utility software, eg operating systems, databases, management tools, and applications</td><td>CONVERGED SOFTWARE LAYER

All IT utility and Layer 4-7 networking software supporting operational and business processes and applications; any other software, eg in-house and third party applications. All software is able to expose itself to any other piece of software as a service</td></tr>
<tr><td colspan="3"></td><td>CONVERGED PROCESSING PLATFORM

Virtualized hardware capable of executing all IT and network software previously associated with dedicated hardware, eg switch/router, IT/network application servers, media gateway, firewall, XML acceleration, cable head end, load balancing, database/storage, content delivery etc, network service brokers, IT middleware etc</td></tr>
<tr><td colspan="3"></td><td>NETWORK TRANSPORT LAYER

Ethernet/VPLS, including DCE, IP/MPLS, FibreChannel over Ethernet etc</td></tr>
</table>

What is required is an evolution to a next generation service delivery infrastructure, in which hardware and software infrastructure components that are discrete entities at the start of this process converge and become more homogenous, flexible and manageable. A single, scalable hardware platform should be able to support any type of software and software itself should be homogenous in its ability to expose itself as a service to any other piece of software. Figure 7

illustrates the goal of this evolution: a layered next generation service delivery infrastructure that operates as a seamless, software-driven platform capable of supporting the delivery of any digital or physical service offered through the Web.

Certain functions have today been embedded into hardware appliances for performance reasons. These functions may continue to be embedded in converged processing platforms, in which case their performance will be enhanced by the all-round capabilities of the service delivery infrastructure. Or they can be extracted into the converged software layer with no loss of, and even a likely increase in, the performance as a result of leveraging the power of the network as a processing platform.

Once next generation service delivery infrastructure (SDI) is established, it will enable service providers to deliver services at high speed, at lowest possible cost, with an experience that matches customer needs and suits their wallets.

Service providers need to take the following five sequential (but probably overlapping) steps to achieve this end goal:

1. ***Consolidate*** the different layers of the infrastructure (network, hardware and software) by standardizing on a single networking technology, a common set of processing platforms and common set of utility software and management tools.

2. ***Virtualize*** the infrastructure by bringing in virtualization technologies at each of the component layers of the infrastructure: network, processing platforms and software. This creates three pools of virtualized resources that can be turned up and down on demand. Service providers already virtualize IP networks using MPLS and/or VPLS. Network virtualization enables service providers to deliver services on demand as a software configuration rather than needing physically to install customer premise infrastructure and equipment which can take weeks or months to provision. Processing platforms can similarly be virtualized using hypervisor technology, allowing a rack of physical servers to be sliced into a much larger number of virtual servers supporting many more applications. Moving applications between virtual servers to balance workloads becomes a software configuration task rather than a physical load and reboot activity. The virtualization of software, including utilities such as operating systems and databases, and applications, can be carried out using technology that translates software built for 'single-user' environments, that is, environments where it is used by one person, one group or one organization, into 'multi-tenant' software, where each user/group is assigned their own, private, virtual copy of the software. Utility and application software is now being re-engineered to naturally be multi-tenant and to understand how to deploy itself in a multi-tenant environment.

3. ***Converge*** the infrastructure, converging and layering all the components in the advanced service delivery infrastructure so that they can be managed seamlessly together as one virtual pool of resources. Since the different layers 'know' about each other, when a virtual instance of an application, and/or virtual server needs to be turned up, 'spun up' or moved to a less busy physical processing platform to handle a spike in user demand, the network can automatically provide the appropriate level of connectivity.

4. ***Scale*** the infrastructure so that it can support the volume and range of services each customer needs, for example to support private virtual copies of the same services for millions of single users and groups belonging to different organizations, distributed across tens of thousands of virtual servers in multiple different geographies

5. ***Federate*** the infrastructure with the SDIs of partner service providers so that service providers can reach and deliver services to roaming customers. This will depend on how quickly a critical mass of service providers adopt SDIs.

2.2 Creating a Timeline for SDI Evolution

Service providers will execute these steps at different rates depending on their market objectives and starting points. The timescales for carrying out each phase will also depend on the scale of the infrastructure to be transformed: a Tier 1 service provider with two hundred network PoPs and 20 data centers may take longer to migrate its infrastructure than a smaller service provider with only a few PoPs and data centers; although a smaller service provider may equally lack the right skill-sets to migrate as quickly as a large service provider with more resources. Nevertheless, an SDI can be implemented by any size of service provider, with smaller service providers and new entrants potentially just as able to adopt a converged infrastructure as larger, more established competitors.

The phases are not necessarily consecutive. A service provider may choose to carry out more than one phase at the same time. Service providers interviewed by Heavy Reading can be divided into three categories in terms of their approach to SDI implementation, as the following figures illustrate. (The six-year timescales are for illustrative purposes only and suggest the possible length of each phase, derived from the plans/actions of those service providers interviewed which are currently transforming their service delivery infrastructures. Year 0 is not intended to equate to a specific year, eg 2010: whichever year a service provider begins this process is designated 'Year 0' in our timescale below).

Asset sweating: These service providers are choosing to evolve their infrastructures progressively, because they have existing legacy services businesses to support and/or infrastructure components that have not yet reached end-of-life. Figure 8 shows how these phases may be implemented progressively.

Figure 8: The Asset-Sweating Strategy

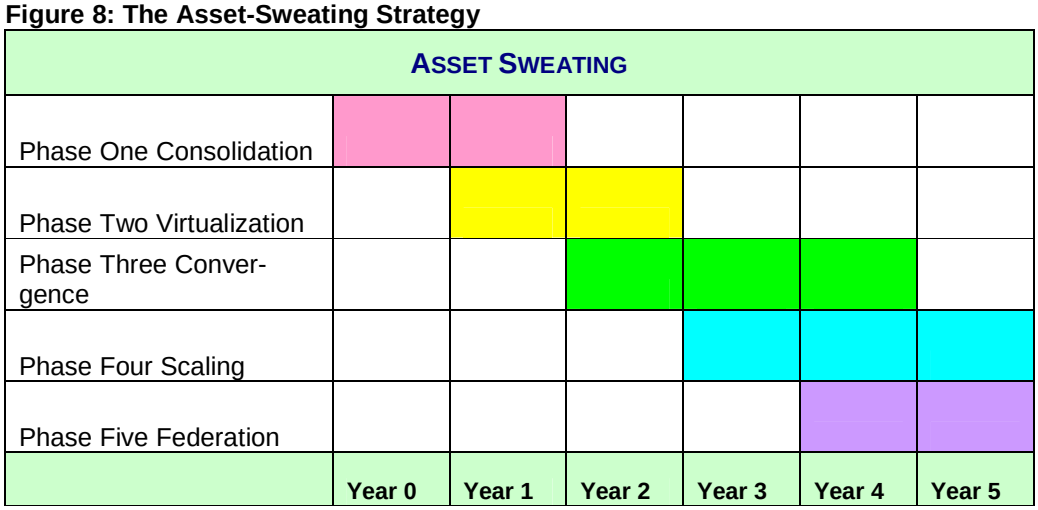

Catch up: These service providers are carrying out virtualization at the same time as consolidating their infrastructures in order to catch up with leading players in the market. However, they want to bed down their chosen technologies before taking the next steps in order to control risk. Figure 9 (next page) shows how the first two phases may be implemented simultaneously: the timescales are for illustrative purposes only.

Figure 9: The Catch-Up Strategy

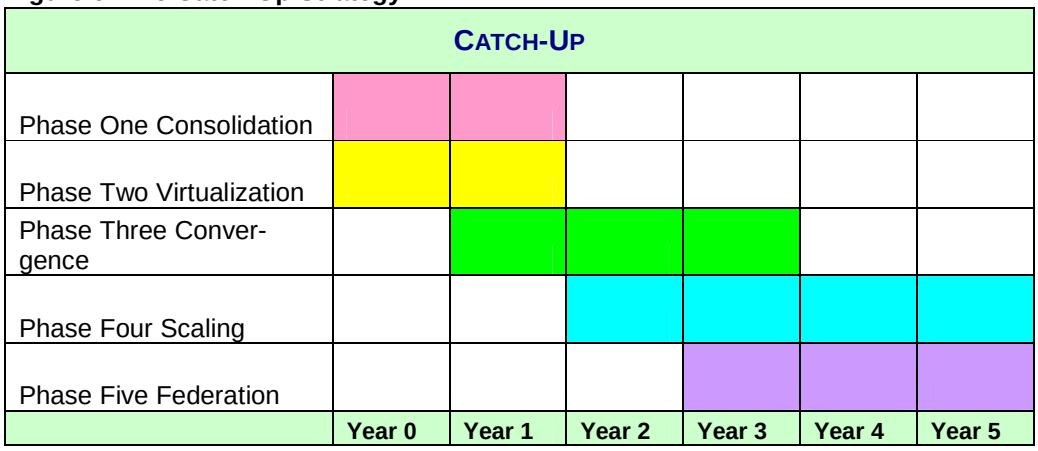

Leapfrog: As the service providers in this category refresh and consolidate their infrastructures, they are choosing to implement virtualized and converged platforms, so that they can establish themselves as new leaders in the market. Such a radical transformation of infrastructure may slow down the process, however, given that it will also require large organizational and skill-set changes. Figure 10 shows how these three phases may be implemented together: the time-scales are for illustrative purposes only.

Figure 10: The Leapfrog Strategy

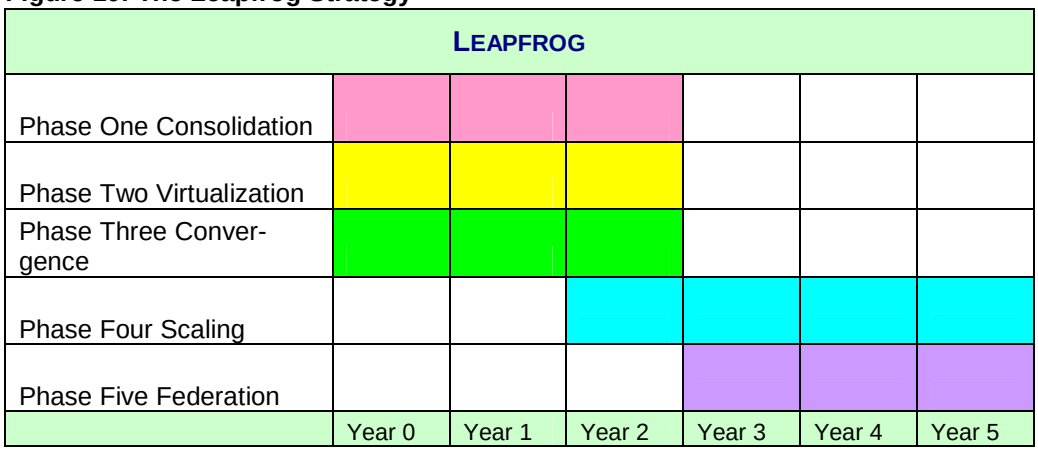

Whatever the approach taken to implementing the first three phases, the scaling and federation phases will be ongoing. A service provider will need to continue to scale as its service business grows and it works towards achieving the improvements in KPIs outlined in Section I. SDI federation, like IP interconnection, wiill take place over time as more service providers adopt SDIs and wish to federate them.

2.3 Critical Success Factors for the Evolving SDI

At each phase of evolution, service providers gain increasing benefit across the four measures (key performance indicators) identified in Section I. This assumes that service providers take into account the critical success factors associated with each phase.

Phase 1. Consolidating the infrastructure

Service providers interviewed by Heavy Reading cite the importance of consolidation and standardisation in:

- Reducing operational costs, due to the need for fewer skill sets and greater ease of managing a more homogenous infrastructure. This translates into lower cost of service delivery
    - **Objective: cost of service delivery reduction of 40%**

- Enhancing service innovation: a common infrastructure and set of interfaces makes it faster to launch new services
    - **Objective: reduce time to launch new services by a factor of 2**

- Improved ability to replicate services across footprint, resulting in new revenues from new geographic markets
    - **Objective: raise revenue by 50%**

- Greater ability to bundle standardised infrastructure components, increasing the richness of service offerings
    - **Objective: increase service portfolio by 20%**

Critical success factors for consolidating the infrastructure

- Define a consolidation strategy with the succeeding phases of evolution in mind
- Mandate dialog and co-ordination between IT and network organizations
- Select technology suppliers based on a roadmap aligned to ours
- Invest in acquiring new skill-sets/retraining staff with legacy skill-sets to support consolidated environment
- Develop/acquire a common, standards-based set of operational processes and management tools in each consolidated domain
- Use open, standards-based interfaces at each level of the consolidated service delivery infrastructure

Phase 2. Virtualizing the infrastructure

By virtualizing these infrastructure components— networks, servers and software— service providers quote the following benefits:

- Reduced operational costs through increased utilization rates. Most service providers have already virtualized their IP/MPLS networks, for example, realizing operational cost savings. Service providers interviewed by Heavy Reading quote server utilization rates in a traditional data centre environment of a mere 5-15%, with processing platforms idling for up to 95% of the time while still consuming power and expensive floor space. Virtualization can increase utilization to nearer the full capacity of the server. However, there is an overhead in virtualizing each component separately and then having to make each virtualized environment aware of the others. There are also overheads involved in re-skilling workforces and making the necessary organizational changes to support virtualization. Service providers point out that it is not always practical to realise cost savings from reduced floor space in multi-functional buildings. As a result of these overheads, we are taking a conservative view of the cost-savings that can be gained from virtualization.
-
    - **Objective: reduce cost of service delivery by at least 20%**

- Accelerated time to market for new third party services, due to the ability to reach and virtualize service owners' services quickly across the infrastructure

    - **Objective: reduce time to market by a factor of 10**
- Increased ability to replicate services across footprint, resulting in an ability to attract a long tail of services.

    - **Objective: raise revenue by 100%**

- Improved ability to handle many more services because of greater capacity and increased process automation, with the ability to deliver services on-demand and with high degree of resilience and availability resulting in a richer user experience

    - **Objective: increase service portfolio (range/features/granularity) by 50%**

Critical success factors for virtualizing the infrastructure

- Make the business case for virtualization with convergence and scaling phases in mind to overcome total cost of ownership issues
- Network and IT organizations should work together to align virtualization technologies
- A repeatable migration strategy from unvirtualized environments should address the large organizational, mindset and operational challenges
- Select virtualization supplier(s) on the basis of their strategy for convergence, including the embedding of virtualization into service delivery infrastructure components
- Ensure operational processes and management tools are exposed to customers so that they are in control of 'their' virtual service delivery environment.

Phase 3. Converging the infrastructure

Advanced service providers Heavy Reading has spoken to want to execute network software alongside utility software and applications on the same physical hardware platforms, running the same virtualization technology. They see the benefits of this approach as:

- Reducing the costs of managing a virtualized infrastructure since there is no management overhead between the different virtualization domains, more automation is possible and such platforms are designed to be cost-effective for the operation and scale of a service provider business

    - **Objective: reduce service delivery costs by 60%**

- Improving time to market for new services by cutting the time to provision to real-time

    - **Objective: reduce time to provision by a factor of 20**

- Enhancing the richness of services, since software function that formerly lived in separate domains, network and IT, live in proximity on the same standardized service delivery infrastructure, served by the same operational processes and management tools. This makes it easier for service providers to create new, cross-domain, bundled and blended products

    - **Objective: increase richness and range of services by up to 100%**

- Enhanced revenue from greater flexibility in the way a larger number of services can be configured and distributed across the service delivery infrastructure, raising ARPU. Ser-

vices that would traditionally be ring-fenced in a service provider's data center(s) or only available through a third party's limited channels, are liberated to run on the service provider's far more widely dispersed converged network platforms, many of which naturally sit very close to large numbers of end-users. This enables service providers to deliver services with improved quality of experience and fine-grained SLA guarantees.

- **Objective: increase revenue by 200%**

Critical success factors for converging the infrastructure

- Converge IT and network organizations to reflect the single, converged nature of the service delivery infrastructure
- Automate operational processes within the converged infrastructure to hide the complexity of interactions taking place between and within the layers of the service delivery infrastructure.
- Create a customer-controlled policy layer so that customers can drive the delivery of 'their' services flexibly, securely and with an appropriate quality of service across the infrastructure
- Design the converged infrastructure to be resilient and self-healing, to lay the foundations for differentiated quality of experience and fine-grained SLAs.

Phase 4. Scaling the infrastructure

As we have pointed out, service providers will scale their infrastructures progressively in response to business demand. The speed with which they will scale the SDI will depend on their individual business growth rates. At scale, service providers will need automated, secure operational processes, proactive management tools, open integration mechanisms and customer self-service capabilities that will allow them to support very large volumes of service transactions across a multi-tenant infrastructure with no degradation of service quality, on-demand responsiveness and the ability to rapidly recruit new services and customers to its platform with minimal delay. Service providers see the benefits of scale as:

- Reduced operational costs due to massive economies of scale and the ability to use less infrastructure to deliver higher numbers of services. Service delivery will involve the automated equivalent of a few keystrokes to reconfigure software for delivery to new customers, plus the running costs of the SDI. Due to the infrastructure's ability to allocate and de-allocate resources dynamically, depending on service demand, these costs will be minimized. At the same time, service providers can amortize these costs across their large customer bases and a long tail of services.

  Objective: reduce average cost of service delivery by 90%

- Eliminating delay in time to market for new services. In a scaled environment, service providers will handle thousands of services with varying attributes: high value, low value, long-lived and of the moment, in many different combinations. The scaled infrastructure will ingest new services in minutes and deliver them on-demand, in real-time.

  Objective: Reduce time to ingest new services by factor of 40

- Supporting an exponential growth in the range of services they deliver across the SDI and an increase in their richness, by enabling customers to mix and match the services they subscribe to in highly individual ways. At scale, service providers can balance the cost-effectiveness of standardized hardware and networking layers against the need for a

highly flexible software layer that lets customers feel they have tailored services, appropriate to their needs.

Objective: increase richness and range of services by up to 150%

- Enhancing revenue both through the growing reach of the SDI, able to support increasing numbers of customers, and through its ability to support more services per customer.

  Objective: increase revenue by 350%

Critical success factors for scaling the infrastructure

- Understand how the costs involved in delivering different types of services in different locations are affected by scaling factors, such as size of addressable market, infrastructure capacity, level of process automation
- Devise metrics that will drive performance optimization at scale
- Automate customer self-service as far as possible to reduce costs in a dynamic, high volume environment
- Ensure that the SDI can track the behavior of potentially millions of customers at high scale: this information is the key to enhanced revenue opportunities and new business models.
- Automate the means by which third parties integrate services into the SDI and its business processes such as billing, third party settlement and fulfillment.

Phase 5. Federating service delivery infrastructures

Even the world's largest service providers do not reach into every corner of the globe and even if they operate at the highest scale, they still have finite service delivery infrastructures which may not be sufficient to cope with spikes in customer demand. All service providers will need to be able to interconnect with partner infrastructure providers if they are to provide their customers with a truly ubiquitous service experience. Service providers Heavy Reading spoke to see the benefits of federation as:

- Managing operational costs by being able to to 'flex and burst' their infrastructures, that is, extend them by temporarily buying infrastructure resources from third parties, on occasion, rather than having to build out their own infrastructure to handle resource spikes and to serve customers roaming in locations beyond their infrastructure footprint.

  Objective: reduce average cost of service delivery by 20%

- Ensuring the range and richness of the services offered to customers is not compromised when the customer roams away from the service provider's geographical footprint

  Objective: improve service richness by 10%

- Gaining additional revenues from keeping their connection to their customers, wherever their customers are, whenever they want access to a service and whatever device they are using. This is critically important as growing numbers of services 'live' on the web rather than on connected devices.

  Objective: revenue uplift of at least 50%

Critical success factors for federating infrastructures

- Play an active role in creating and supporting emerging standards initiatives for service delivery infrastructure interoperability, such as the Cloud Interoperability Forum and IPSphere (now part of the TM Forum), which are defining common management interfaces, policies, SLAs, settlement mechanisms and security features
- Provide open interfaces to operational processes and systems, for example, to transcoding systems to support partner provider devices or to the SLA management system so that partners can synchronize SLAs on behalf of their customers.
- Lead the development of a new business model supporting third party service delivery infrastructure usage.
- Leverage key service provider network interconnect partnerships to align the evolution of service delivery infrastructures. This will make it easier to extend interconnect agreements to support the interconnection of advanced service delivery infrastructures as these evolve.

2.4 Mapping KPIs to Phases

Our final section maps the KPIs we outlined in the first section to the SDI phases set out in this section, in order to illustrate how each phase contributes to the overall objectives

Figure 11 shows how Phase One might be mapped against KPIs, and what its overall impact would be at the end of this phase. Figure 12 maps KPIs against all five phases, showing how each contributes to the overall goal.

Figure 11: Mapping KPIs to the First Phase of SDI

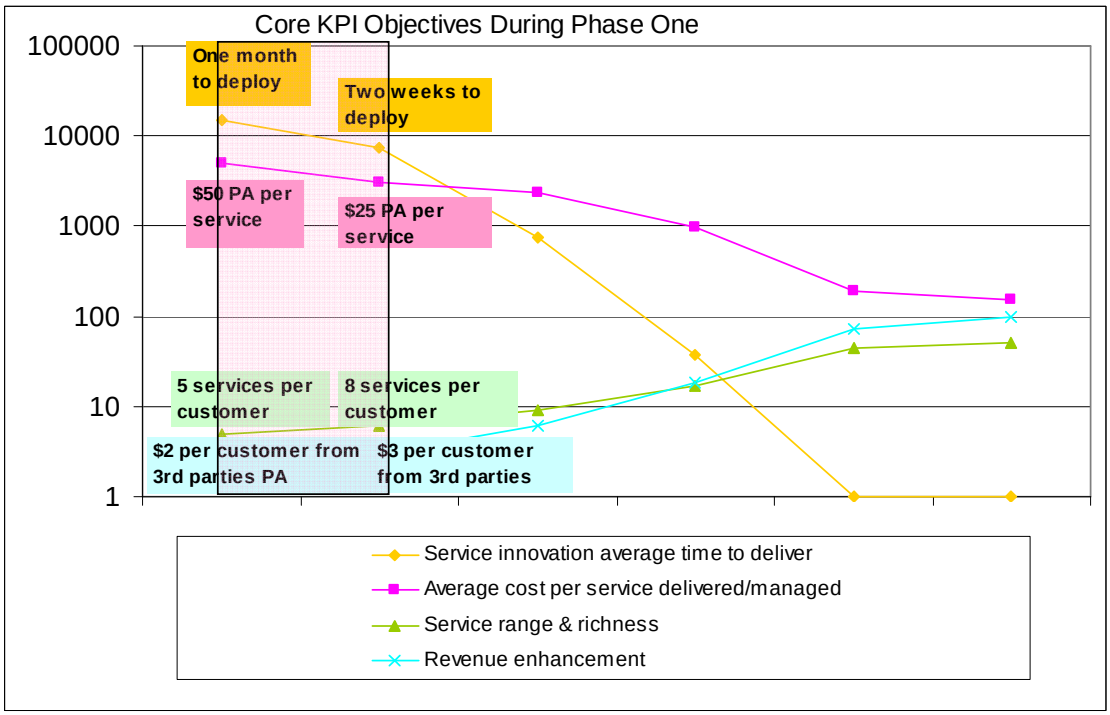

In practice, every telco will need to create its own individual timeline and metrics for this work, but the general principle is clear: an ambitious plan for a transformation of the business, underpinned by a phased deployment of a new service delivery infrastructure, lies at the heart of the creation of next-generation network operators.

Figure 12: Mapping KPIs to all Five Phases of SDI

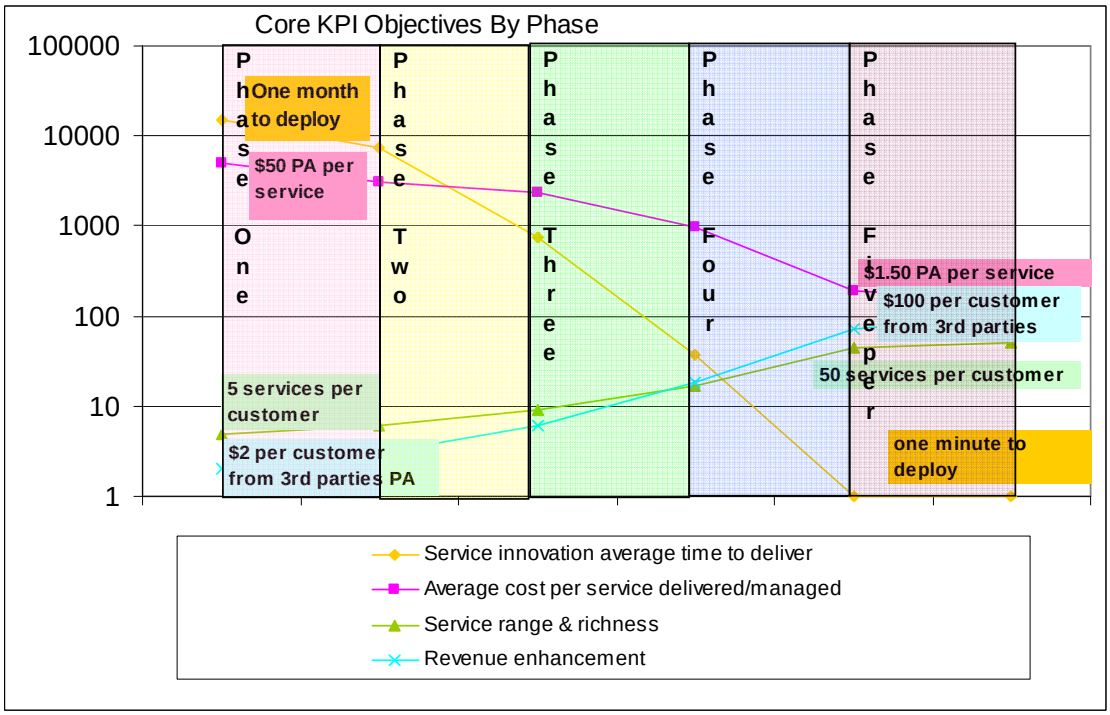

Appendix A: About the Authors

CAROLINE CHAPPELL
ANALYST-AT-LARGE, HEAVY READING

Caroline has spent 20 years analyzing the convergence of IT and network technologies and markets. Her particular interests include service innovation, service delivery platforms, IP-enabled managed services, application service provider and on-demand delivery models. Caroline established Heavy Reading's successful **Services Software Insider** research service in 2005, and has written extensively on service software issues over the past four years. Before establishing Services Software Insider, Caroline worked initially for Ovum where she led large market research projects for clients including IBM and BT. As a freelance consultant, her clients have included the European Commission and a range of UK ITC services companies. Her most recent Services Software Insider report, **"Accelerating Telco Services Through SaaS/PaaS: Strategies for SDPs in the Cloud"** covers platform-as-a-service and cloud-based service creation and delivery developments.

GRAHAM FINNIE
CHIEF ANALYST, HEAVY READING

Graham has 20 years experience in the telecommunications sector as an analyst and consultant. He joined Heavy Reading in September 2004 following a ten-year tenure at the Yankee Group, where he had directed a broadband & media research program. He was appointed Chief Analyst at Heavy Reading in February 2007. As well as setting the overall direction of Heavy Reading's content, Graham has been responsible for a wide range of research, focusing primarily on next-generation broadband services and new applications architectures. His recent publications include **"Re-Inventing the Telco: A Heavy Reading Progress Report"** and **"Policy Control & DPI: The New Broadband Imperative."** Graham also led Heavy Reading's 2008 multi-client study, **"Third Party Service Enablers & The Over The Top Revenue Opportunity,"** and is preparing Heavy Reading's June 2009 report, **"Open Access Telcos & The Third Party Opportunity."**