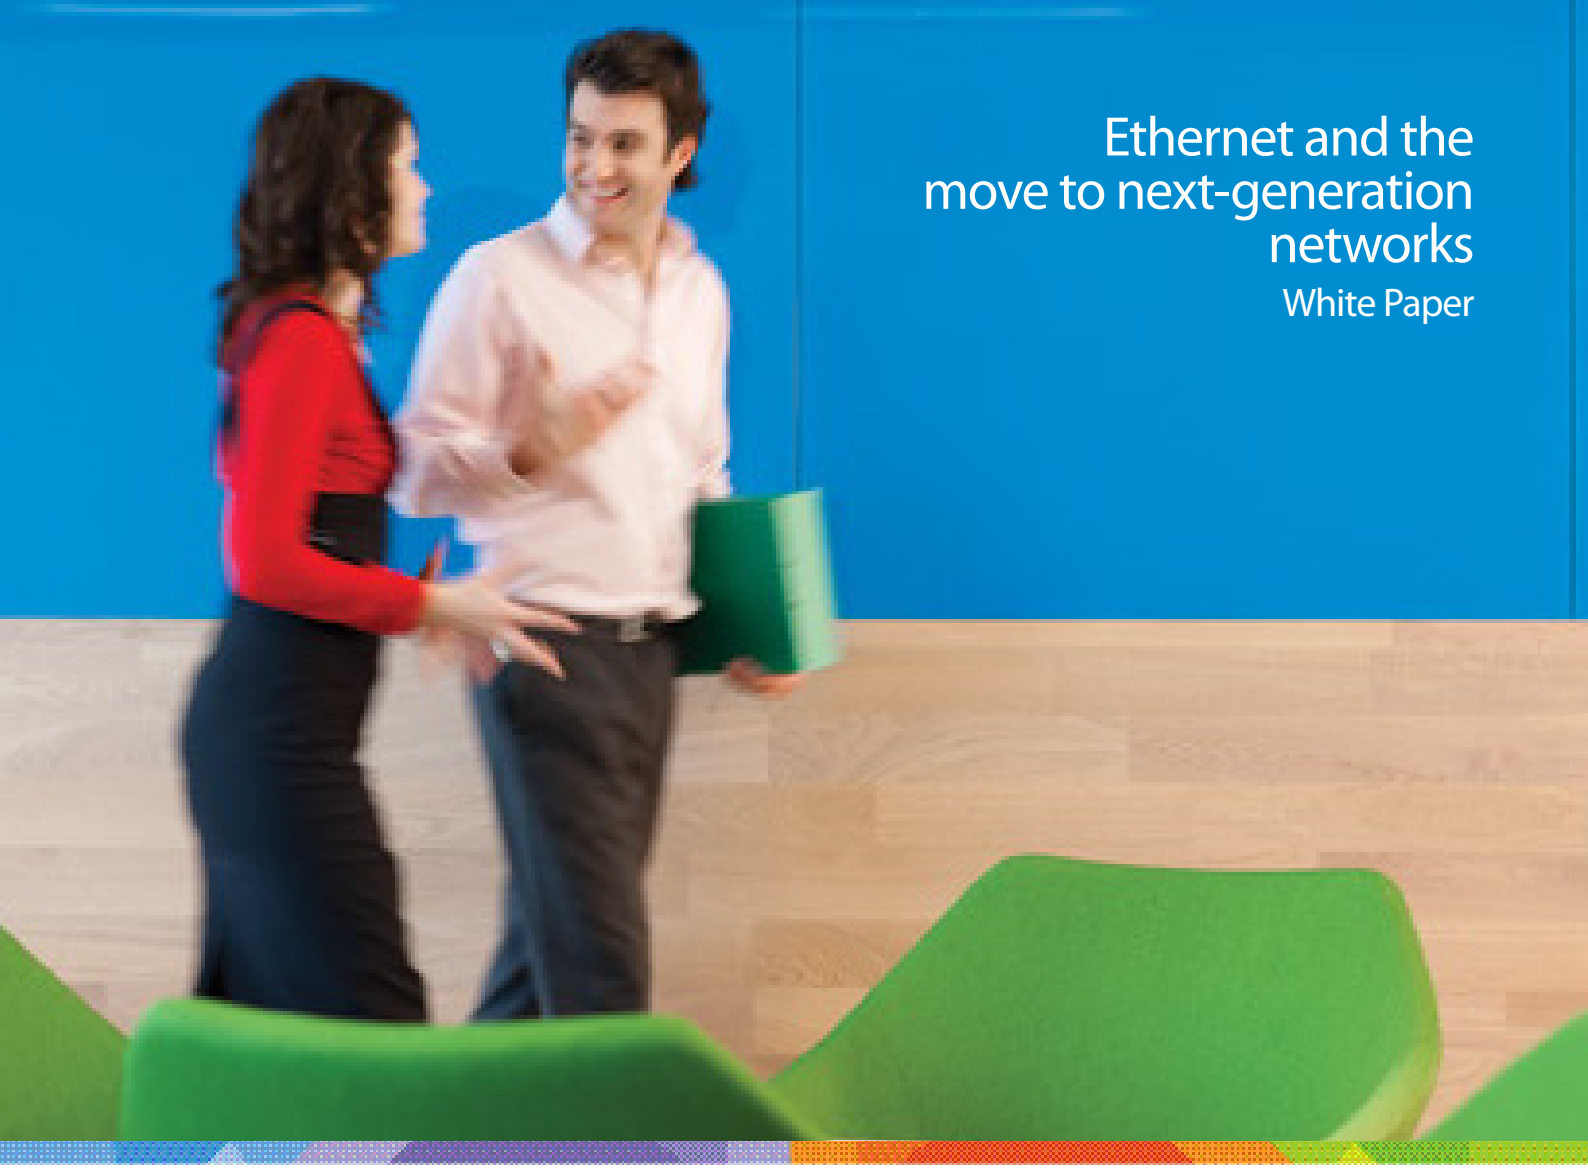

Ethernet and the move to next-generation networks

White Paper

Introduction

Just over five years ago Synchronous Digital Hierarchy (SDH) was the undisputed technology of choice for high speed data transmission. Today this is no longer the case, as new technologies that enable next-generation networking are beginning to make their presence felt.

SDH won its crown because it was a universal standard in digital transmission. With both optical and electrical interfaces, easy and cost-efficient traffic add / drop and cross connect capability it enabled the deployment of powerful management capabilities within the network.

The key to SDH's manageability was its highly structured nature, subdividing transmission paths into various traffic bandwidth payloads and handling traffic according to these. However, when SDH is measured against the management flexibility provided by Ethernet, it is clearly limited.

KCOM

Whilst SDH management traditionally takes place at the network, Ethernet can be managed at the network and its edge where customers are connected. This offers compelling management advantages for service providers: for example complex solutions (such as "office-in-a-box" services for instance, with internet connectivity, voice PBX and data connections) can be pre-configured off-site and delivered to the customer for almost immediate deployment when a single Ethernet link is provisioned.

This has several obvious benefits. Carriers can easily deploy services without needing to pre-configure the network and service 'tweaks' and even bandwidth upgrades can be simply handled by service providers. Essentially all the carrier has to do is provide the pipe between the services. Furthermore, Ethernet is today increasingly gaining trust as a secure transmission medium.

There is comfort in familiarity, however, and the self-healing and low latency properties of SDH are often perceived as important for customer peace of mind and security. However, long distance Ethernet provision has similar self-healing properties to SDH in its underlying core MPLS (Multiprotocol Label Switching) network. In most SDH deployments, the last mile (core edge to the customer) is generally a single spur, a "flattened ring" or a non-protected tail circuit. In these instances SDH does not provide significant advantages over Ethernet in terms of resilience.

Most carrier NGN Ethernet services are now delivered over an MPLS core network. It is the native medium for IP and provides the foundation for the move to integrated next-generation networks, without the need to code or decode transmissions.

It's a fact that Ethernet adoption is accelerating at exponential rates as incumbent carriers look to close down legacy TDM networks. This move is most evident among mobile operators who are looking for hybrid network solutions and who are taking advantage of the much higher bandwidth to cost ratio offered by Ethernet.

Green light for MPLS Ethernet

The rapid adoption of Ethernet has, however, brought with it a few problems of its own, not least in terms of product definition and confusion about what the various methods of Ethernet provision can deliver.

Back in the early days of Ethernet networks in campus environments they were designed as a dedicated fibre solution. This persisted for many years as Ethernet continued to be viewed as a Local Area Network or short haul data service, limited by attenuation to (at best) 25km.

Some Ethernet tails are still provided as Layer 1 services, although this is a rather expensive approach to using finite fibre resources. Most network operators now provide both core bandwidth and local access as Layer 2 services, using MPLS transmission. Once out of the core network, local access tails are still limited by range, but national and international Ethernet services are now common whereas previously they were seen as the exception.

Historically, some customers have been resistant to MPLS delivered Ethernet for a number of, originally, valid reasons. However, developments in MPLS technology and the way in which carriers deploy their core MPLS networks have overcome many of those objections.

Often, an MPLS delivery is perceived as being contended in the core network, because of the way the core network allocates bandwidth as a virtual path rather than a fully "nailed up" circuit. Although some carriers operate a contended core delivery service, assuming that the pool of core bandwidth will be enough to handle requirements at any given time, some carriers will guarantee an uncontended circuit path through the network, offering a contended option as a lower cost solution.

Customers have also sometimes misunderstood the nature of throughput bandwidth over Ethernet and have been puzzled to find that they are not getting what they believe they have paid for. The truth is that no transmission technology can deliver the full nominal data rate across a circuit path – the laws of

physics would make this impossible – and there is always a certain level of "overheads to deliver the end-end communication".

These overheads are proportionally more with Ethernet than, say SDH. However the much higher "bang per buck" in terms of the bandwidth available at a given price, makes Ethernet a very economic way to deliver data services. Long distance 100Mb services, for example, cost a fraction of the price of a 34Mb SDH circuit on the same route.

A more pressing concern for customers may be that of frame size, in other words the size of the data package that customers want to transmit. Historically, MPLS Ethernet struggled to handle larger frame sizes but these days 2000MTU can easily be delivered over MPLS Ethernet, far in excess of the commonly required 1560MTU. So-called "jumbo frames" of up to 9000MTU would, however usually still require a non-MPLS transmission protocol.

It is also not uncommon for end-users, particularly in the financial sector, to specify delivery via Ethernet over SDH (EoSDH). The truth is, however, that very few carriers actually provide this end-to-end and incumbent carriers often do not specify whether the last mile is delivered via a Layer 1 or Layer 2 service.

In reality, other than providing a committed path through the network (something that MPLS carriers can also do), there are few advantages to providing an EoSDH solution. This is because the service will still conform to the characteristics of an Ethernet delivery, regardless of the bearer, and non-protected SDH access provides little more in terms of resilience than an Ethernet tail.

In many cases, EoSDH is perceived as a "security blanket" rather than an improvement on MPLS. For some customers it will still make sense to have services provided over a non-MPLS transmissions medium, possibly incorporating the use of DWDM in the network core. For the majority of Ethernet services, this will be a premium product as at the sub 500mb level it has higher cost implications than MPLS.

Interestingly, however, non-MPLS delivery can provide a lower cost option for higher order services that are increasingly in demand. No doubt, however, this will be subject to further change as technologies improve.

A broad reach for new services

So, high bandwidth services are available nationwide at much more economic prices than before. The question for service providers, then, is how do they gain access to these services? Some typical options they might explore are building a backbone network of their own to backhaul local access tails to Points of Presence (PoP) sites, establishing supply agreements with multiple carriers depending on who has a network and where, and economically managing interfaces with network providers and supplying these circuits onwards to their customers.

There are, of course many options and some will make sense to certain service providers but not others, but perhaps the most cost effective solution is to work with a provider who can supply end to end services at a competitive price across the widest possible geographic spread.

Taking delivery of services across a Network to Network Interface (NNI) means further flexibility and cost savings in terms of service integration and onward delivery to customers in, say, a data centre. Effectively a large bearer Ethernet circuit, usually 1 Gigabit, the NNI carries individual circuits as VLANs, much in the same way as an SDH bearer can be subdivided into multiple Virtual Containers (VC's).

The advantage lies in not having to take delivery of each circuit as 10 or 100 baseT and having to provide the relevant hardware and cabling. Instead, a 2Mb service only takes up 2Mb of the VLAN capacity on the NNI. Furthermore, service providers can effectively create a new PoP for themselves without having to have network connectivity via a NNI hub delivered into a new location. This

significantly reduces the cost of additional circuits into that location and effectively brings what was off-net, on-net.

If a service provider is operating a neutral facility, rather than providing service to its customers through its own network, the addition of VLAN switch equipment to the NNI hub enables cabling directly out to individual customer racks.

Multi-site strategy

A further area of opportunity is enterprises that need networks between different sites. The obvious option may be to deliver connectivity using a VPLS (Virtual Private LAN Service) meshed structure, which provides a pipe into each location supplied from a central core of switched bandwidth. This network structure can be an efficient use of bandwidth, keeping the number of dedicated connections to a minimum, however the switched core element may be proportionately more expensive than a fixed network solution.

In contrast, a meshed hub infrastructure, switched at the edge, with a hub in each location, can provide a much more compelling option. Once the hubs are established, incremental bandwidth between those hubs is much more economic as there are no access tail costs and the available bandwidth between sites is uncontended. This is essential for enterprises that have high levels of critical data travelling between sites simultaneously.

Copper last mile reaches a growing market

Cost will always be a factor for some customers and where the cost element is more important than consistency of throughput, companies with multiple geographically distributed sites may opt instead, to rely on a DSL based VPN.

Using ADSL services, however, has some major downsides. By its very nature, ADSL is a contended service and at peak hours, even with "business class" services, experienced bandwidth speeds can suffer. Those expecting Fibre to the Cabinet (FTTC) rollout to improve speed may be in for some surprises as this only improves transfer speeds from the DSLAM in the exchange to the end-user. With the growing demand for bandwidth-hungry services including MMP gaming and the advent of domestic cloud computing, the concentration of investment in the access infrastructure rather than the core, means more pressure on the available bandwidth. Business consumers may start to see a sort of inverse Moore's law coming into effect for a period with prices going up and experienced throughput coming down.

An alternative is to look at copper delivered Ethernet, also known as Ethernet First Mile (EFM). This is an entirely viable alternative to fibre delivery for some types of customers such as home workers who have definable bandwidth requirements and need permanent connections. In fact, this market segment is increasingly growing as companies realise the savings that can be gained with employees working from home.

There are some limitations to EFM. Maximum circuit bandwidth is 20Mb and the availability of speeds will deteriorate the further the end user is from an EFM equipped exchange. Repair SLAs are also lower and it does not have the flexibility in terms of upgrade path that fibre does. Clearly, it's not a premier product, but as a delivery mechanism for uncontended full duplex bandwidth, where the costs stack up, there will be a growing market for it, particularly in metropolitan areas.

Looking ahead

The options outlined so far illustrate how carriers can take advantage of current technologies to deliver new services and improve the delivery of existing services. In a sense, they also reflect new emerging dynamics and provide a roadmap for the future.

Broadly speaking, as we move towards a true next generation environment, we are moving to a position where the core network will increasingly become a more generic cloud and it will be less

easy to define, and less important which network is carrying the service at which point.

Kcom believes we are moving towards a network that is entirely resilient and one which also provides mirrored services. Data services will become more standardised (or, perhaps to be more accurate, increasingly generic and transparent) but those carriers who are highly integrated with the network cloud will offer greater service reliability, benefit from economies of scale and also deliver value-for-money services. In fact, carrier diversity will be characterised by premium product diversity.

Importantly, by pushing the services out to the router, service providers, data centres and even end users will become virtual network operators. In short, carriers will no longer sell individual services; rather they will offer core generic capacities to customers.

Customer self-service, carrier capacity

One consequence of these developments will be increased online service provisioning enabling end users and channel partner customers to price, specify and configure the services they want. Self-provisioning will become an increasingly attractive option for these customers. The carriers who are quick to recognise this and allow their customers to do more with their own services will ultimately be the more successful carriers.

Partnering with Kcom

Kcom is in a unique position to enable the transition to Next Generation Networks. We can offer true national coverage throughout the UK via a strategic partnership with BT. In addition our own national network and regional metro networks, provide full in-house carrier diversity, enabling our channel partners to deliver a fully bespoke network to their customers.

Kcom offers the widest range of Ethernet services designed for channel partners. And, with other network options including dark fibre, managed wavelength services, SDH leased line and 21CN DSL, we can offer the fastest, most up-to-date services to our customers.

Within the context of a changing network landscape, Kcom is well-placed to put its partners at the forefront of the transition towards the truly converged network future.

CONTACT US:
please visit kcom.com or call 0800 915 5226
or email us channelpartners@kcom.com

July 2011
Part of the KCOM Group