



Optimum Communications Services: finally a way out of the zero-sum game?



In a previous post, [The new network dogma: Has the wheel turned full circle?](#), I wrote about the now common perception that the mix of technologies that forms the basis of a service provider's Next Generation Network (NGN) is considered to be done and dusted. This is because network technologies are now wholly dominated by IP-based protocols. In itself the wide-scale adoption of IP has been transformational for the carrier industry, but at the heart of this domination lays a major 'zero-sum game' dichotomy where the gains are offset by an equal, if not greater, loss.

What has caused this dichotomy? Today's IP and MPLS networks have to manage bursty and unpredictable IP traffic over a layer-1 network that, regardless of the protocol, is composed of fixed bandwidth connections. This has led to the creation of a family of ever-more complex protocols that aim to 'micro-manage' IP traffic at layers 2 / 2.5 and 3. In the above posting, it was suggested that today's IP networks could be considered to be "over engineered" and often seems to act like children with tantrums.

In practice, many network service providers have balked at deploying these complex protocols and have turned to the 'simple' concept of over-provisioning where bandwidth is thrown at the network to ensure adequate Quality of Service (QoS) for their business customers. However, over-provisioning maybe considered to be simple to deploy and manage, but it comes together with innately high deployment costs and low efficiency. This is because the expensive network equipment port and transport bandwidth headroom that is required to handle bursts is not used for much of the time. It seems that a network architect or planning engineer is in a no-win situation and have to choose between the costly simplicity of over-provisioning or the complexity of a micro managing the network.

This no-win situation is further compounded by other factors. First, there is the shift in the mix of applications carried by IP from non real-time file transfer type traffic such as emails and web pages, for which IP was originally intended for, to real-time services such as VoIP or streamed video services. This shift demands sophisticated application QoS control capabilities, increasing the cost and complexity of providing IP services.

On top of this change, the overall volume of service traffic continues to increase apace, at much greater rate than the customers' spending can increase. These factors, on their part, would require massively scalable IP service delivery infrastructure with radically reduced cost per Gbps of service capacity. This demand for simplicity and cost-efficiency seems to be in direct contrast to the need to guarantee deterministic QoS for the emerging realtime and interactive IP applications.

Taken together, it appears that these pressing requirements cannot be solved by either of the traditional means of over-engineering (leading to unmanageable complexity) or by over-provisioning (leading to decreasing utilization and increasing capacity costs).

Thus, even if the NGN technology plans were once considered done deal, the end-user application demands and the harsh economics cause there to be, as much as ever, need for innovations that concentrate on delivering both network cost-efficiency and performance improvements. I would suggest that if a more open approach was taken on network architecture and business model options then competitive advantage *can be* gained for those service providers who are willing to consider new approaches. Perhaps through innovation, the concept of 'brute-force' over-provisioning can be made to be seen as one of yesterday's architectural strategies?

In this vein, I'd like to take the opportunity to write about one company that has taken a highly innovative approach to address improved efficiency in next generation networking. Moreover, their business model is that of a service provider rather than a technology vendor. You don't often come across a service provider that bases its business on internally developed and patented innovative technology, [Optimum Communications Services](#) (OCS) ⁽¹⁾ is one that has done just that. To quote the opening paragraph on their home page:



The real-time self-optimizing, self-configuring broadband VPN and multi-service backbone network solutions of Optimum Communications Services (OCS) combine the performance and security of a private network with the economies of shared Internet infrastructure. With OCS' Intelligent Transport Network™ (ITN) solutions, performance and cost are not a trade-off; the efficiency of ITN results in both maximized performance and minimized cost.

The last sentence should be enough to intrigue any service provider CTO, COO or CFO and want to find out more.

OCS' business model

It's always difficult to know whether to talk about technology or business models first as it depends on whether the reader is a technologist or a commercial manager. In this case however, I feel that it is quite important to explain the business model first as OCS have chosen to be a **service provider** rather than be a **technology provider** (these two approaches are not necessarily mutually exclusive).

Although OCS have developed innovative and patented technology ⁽²⁾, they plan to bring this to market by offering a wholesale network connectivity and/or optimisation services that benefit from their technology developments. They offer these in two principal flavours:

Wholesale service: With customers who do not desire to invest in, own or operate their own fibre-optic networks and customarily buy their required network connectivity services from network service providers, OCS will operate as a traditional **service provider**. They will bring a set of services to the market that will benefit their customers from both cost-efficiency and service improvement perspectives. OCS will manage the customer's network services enhanced through the use of their own technology innovations applied to off-the-shelf leased optical wavelengths. Examples of organisations could be systems integrators, who could use OCS' service as a more cost-efficient alternative to traditional leased lines or wavelengths in an enterprise end-customer's Wide Area Network (WAN) application, or they could be value-added service providers such as CDNs, ASPs, wireless service providers using OCS' service to interconnect their POPs, data-centres, switching centres and network access points. In this case, OCS is acting as a wholesale service provider.

Optimisation as service: For organizations with a significant investment in their own fibre optic network capacity, OCS can provide a network throughput maximization service as a **service provider** and will operate and manage the its optimisation-layer technology on behalf of the customer though, in this case, over fibre-optic wavelength capacity provided by the customer.

Background

Even after more than a decade of packet traffic driven network technology and service offering development, and while most of the network traffic already is packet based i.e. inherently of variable bandwidth nature, it is strange but true that the variable bandwidth packet traffic is still being delivered over fixed-bandwidth physical layer network connections, regardless of the protocols used. These planning and technology deficiencies unavoidable lead to a high level of structural inefficiency because a high percentage of costly bandwidth remains unused for most of the time. This wasted bandwidth is called **Stranded Bandwidth** and can be found in any network even if is an all IP network or the strategically 'ideal' Next Generation Network (NGN) and represents money down the drain for most operators.

So what are the technology innovations that lie behind these service offerings?

OCS' technology

OCS uses its patented technology as a principal embedded component in its wholesale and optimisation services so OCS does not sell technology as such. However, it is worthwhile looking at how OCS achieves such significant network efficiency improvements.

OCS' approach was to go back to basics and realise that the root cause for the ever-increasing complexity of traffic management at packet layers is the need to try to mitigate the impacts of structural inefficiency originating from the use of non-variable bit-rate layer-1 connections for delivering variable bandwidth packet-switched services.

Instead of continuing down the path of increasing complexity via various packet-layer (i.e. layer 2 and above) network micro-management techniques, a more elegant and more sustainable, architectural solution could be achieved by delivering a self-optimising layer-1, to achieve all the efficiencies there are to gain with packet-switching, while at the same time providing per-customer-dedicated connectivity that is deterministic with high QoS and security. Specifically, if the most modern existing layer-1 standard, SDH, could be enhanced through the addition of real-time dynamic bandwidth allocation driven by the real-time traffic requirements between each access port of a network, then it would be possible to achieve a number of benefits, one of which is to eradicate stranded bandwidth. OCS calls this concept Adaptive-Mesh networking, implemented as OCS' Intelligent Transport Network™ (ITN) technology.

An Adaptive-Mesh based network is fully real-time self optimising so there is no 'underutilised' or stranded capacity, as the capacity allocation across the entire network is continuously optimised.

The purchase of OCS' commercial services enables globally maximized network revenue-generating service delivery capacity and provides an improved network performance using several mechanisms:

Dynamic bandwidth allocation at network physical layer (layer-1): An Adaptive-Mesh network provides a full mesh of dynamic bandwidth layer-1 connections among a group of network ingress and egress points. An Adaptive-Mesh network provides non-oversubscribed any-to-any connectivity among the customer sites that it interconnects, via providing a sub-pool of network capacity for transport of data from the network ingress points per each egress point of the network. Each such sub-pool within an Adaptive-Mesh has the same total bit rate as its corresponding network egress point. These sub-pools work as wide-area multiplexers in that they provide a dynamic-bandwidth layer-1 connection from each of their source nodes to their destination node, and are called **Adaptive-Concatenation Multiplexer Buses (AMBs)**. OCS has implemented AMBs to support such dynamic bandwidth source node specific circuits with STS-1 capacity increments that are re-optimised for every new STS row period, i.e., at rate of 9 (rows per frame) times 8000 (frames/second) = 72000 optimization cycles per second, providing per second bandwidth granularity of 50Mbps/72000 i.e. finer than 1kb/s. Thus, for a 10Gbps destination access point of Adaptive-Mesh connecting to a 10Gbps port connected to a customer's switch/router, there will be an STS-192 AMB, which provides a dynamic bandwidth, Adaptive-Concatenated STS-X channel for each of the source access points of the Adaptive-Mesh that needs direct layer-1 connectivity to the given destination.


In addition to the dynamically multiplexed layer-1 data plane, the AMB provides a real-time control plane to optimise the allocation of the STS-1 timeslots of the AMB among its source node specific STS-X circuits. To accomplish that, the destination node of AMB receives, for every new STS row period, capacity demand info from the source points of the AMB, and in response the destination node re-optimises the allocation of the

STS-1 timeslots on the AMB among the source points, ensuring that all source points get their requested portion of the AMB capacity at least up to their fair share, and as much of the total capacity as possible while fairly meeting also other requests.

The AMB control plane signalling is carried in the otherwise unused STS-1 overhead fields, and the same optimisation process is repeated for every STS row cycle on the STS-N bus, causing the capacity allocation of each AMB, and thus the entire Adaptive-Mesh network, to be always optimised and the network traffic throughput thus continuously maximised.

Thus the Adaptive-Concatenation of OCS' Adaptive-Mesh, **apparently for the first time ever**, has enabled supporting burstable traffic (the desired aspect of Variable Bit Rate requirements) with minimum throughput guarantees (the desired aspect of Constant Bit Rate requirements).

The entire optimisation process of Adaptive-Mesh is furthermore fully hardware automated and transparent for the customer nodes which are in full control of the Adaptive-Mesh capacity allocation, as it is the traffic loads between the customer nodes that drive the capacity allocation optimisation of Adaptive-Mesh, and the customer has the ability to control the traffic loads among its nodes connected to an Adaptive-Mesh network.




Importantly, this method of layer-1 dynamic bandwidth allocation *does not* require packet processing / switching at the backbone network interfaces but only at the access port's bandwidth (e.g., avoids need for packet layer logic at 10Gbps ring interfaces in case of 1Gbps access, as well as avoids need for packet level logic for 100/160Gbps transport interfaces in case of 10/40Gbps access). This dynamic bandwidth layer-1 circuit based packet transport of ITN significantly, in the order of 10:1, reduces the cost of


packet layer processing compared to the use of a traditional packet-switching networks, which need to process packets at the backbone line speed as shown in the diagram on the left.

The diagramme also shows why a typical layer-3/2 router/switch connected to 10Gbps backbone and 1Gbps access interfaces requires a total packet processing and switching matrix capacity that is multiple times greater than the sum of its access port capacities and that most of the packet layer logic capacity is consumed by the backbone network interfaces. As backbone capacities rise to 100/160Gbps and beyond, this requirement can make conventional packet-switching equipment very, very expensive to develop and manufacture and that would unavoidably be reflected in the end-user price.

Innovative topology: A second opportunity for optimisation lies at a topology level. For example, take a simple closed 5-node network that is functioning as a service provider's optical network. The requirement is that each node has a total traffic exchange capacity with other nodes of doubled 10Gbps, any breakdown of which aggregate between the four source/destination node specific connections needs to be dynamically supported (based on real-time inter-area traffic load variations) in order to not block revenue-generating traffic. As shown in the diagram, implementing this connectivity requirement with conventional, non-adaptive network physical layer, would require in total ten such bi-directional mesh connection, each of which would need to support 10Gbps in order to not limit the network throughput. The doubling of the mesh for protection purposes would require a total 20 such layer-1/0 connections, which conventionally would be mapped to 20 optical wavelength loops on a fibre ring.



The efficiency gain of OCS' ITN solution for this example lies in the fact that it can provide an equivalent fully meshed dual-provisioned non-blocking connectivity using only a **single** wavelength resource, where the conventional architecture needed **twenty** of them on the same circular fibre route. Simple mathematics says that the minimum aggregate ring bandwidth required to support 5 sites with the required 10Gbps access with non-blocking connectivity is $M(N-1)/2$ where M = port bandwidth and N = node count. Therefore $10(5-1)/2 = 20\text{Gbps}$ is, in theory, sufficient, provided that the network capacity is continuously kept optimally allocated among the source-destination node connections, ITN, as so far unique and patented technology, is able to meet this condition. With dual connectivity, as is the case with this example, 40Gbps of ring bandwidth is adequate and such a signal carried over a single wavelength. Thus OCS' ITN realises a 20:1 reduction in physical network capacity costs. The financial consequence could mean savings of several million Pounds or Dollars per year, if not more, for even modest size network contracts.



Why should this be so? In this five site example, on any section of the 40Gbps OC-768 ring, both 20Gbps halves of the ring can support **two individual 10Gbps STS-192 Adaptive-concatenation Multiplexer Buses (AMBs)**, each of which provides a dynamically optimised bandwidth STS-X circuit per each source node along the AMB to its destination node (ITN Interface Module, IM) and there are two destination nodes to be reached on any given ring section along either ring direction, as per the diagramme. Thus, with AMBs, non-oversubscribed any-to-any connectivity among five sites with 40Gbps of traffic exchange capacity at each can be provided over 20Gbps of ring capacity, and 40Gbps ring can support two such (protecting) AMB mesh networks.

The key result of this is that ITN is able to provide a full $2 \times 2 \times 10\text{Gbps} = 40\text{Gbps}$ of traffic throughput between any of the five network sites — both without the packet processing of conventional hop-by-hop packet-switching network, and without the multiple times higher number of wavelengths a traditional solution would require. **This means that a fully layer-1 meshed non-blocking network has been implemented at a fraction of the price that would be needed without using OCS' ITN.**

Cut-through routing: The advent of **Multi-Protocol Label Switching (MPLS)** standards defined a new mechanism for packet forwarding in routers and enabled the rollout of **connection-oriented** networks. Connection-oriented data networks emulated the way that paths were set up in the **Public Switched Telephone Network (PSTN)** networks. This involves setting up a path *prior* to sending packets across the network ([The rise and maturity of MPLS](#)) and enabled the removal of the ATM layer from live networks that originally provided this capability.

Without MPLS, Carrier Ethernet or ATM, both the link layer and routing headers of each packet needed to be analysed at each node it passed through to determine where next to send it. Packets certainly arrived at their destination reliably (dependent on re-transmission by layer-4 protocols) but with unpredictable delays. We are all familiar with this from the Internet where packets sometimes need to transit 20 nodes or 'hops' before arriving at your computer. This may be acceptable for non-real-time services such as email but is a killer for real-time services such as VoIP or video conferencing.

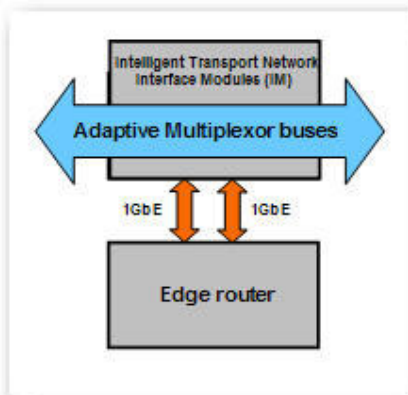
With layer-2 or 2.5 switching or forwarding techniques such as FR, ATM, Ethernet or MPLS, the layer-3 routing and processing can be avoided at intermediate network nodes, allowing more deterministic packet delivery with normally fixed routes with reduced delay variation. However, layer-2 or 2.5 packet forwarding techniques cannot avoid packet (or cell) layer processing, queuing and switching at the intermediate nodes and for all interfaces, without losing the bandwidth efficiency and flexibility advantages (vs. traditional layer-1 multiplexing) that those techniques were developed for. Such hop-by-hop packet-switching, besides wasting most of the packet processing logic for transit traffic at each node, moreover unavoidably exposes the packet traffic to consequences of variable levels of congestion (i.e. increasing delay, jitter and packet-loss probability with each packet hop), as well as for ongoing security concerns in such packet-layer shared

backbone environments. Thus, the conventional layer-2 or 2.5 forwarding overlay network solutions could be considered as half-solutions, with their set of trade-offs compared to e.g. L3 VPNs.

The Adaptive-Concatenation based adaptive-bandwidth layer-1 of ITN however allows avoiding layer-2+ packet layer functionality altogether at the intermediate nodes along a path of a packet without losing any of the maximum achievable bandwidth efficiency or flexibility of packet-switching, and without having to mix traffic from unrelated customer or application networks in the same physical layer connections or packet processors. Moreover, since OCS' innovation is based on the physical layer of the networks, its benefits apply just the same for all services delivered over any layer-2+ protocol stacks.

This means that packets can pass through any number of intermediate nodes at full line speed without being slowed down or the service quality in any way degraded. The number of hops could be as low as two or three whereas without this capability it might be twenty or thirty with L2/3 techniques targeting the same bandwidth efficiency. What is also interesting is that the QoS experienced by these packets delivered over Adaptive-Mesh networks is of the very highest achievable in any telecommunications network as it is derived from the direct L1 circuit based transport across multiple intermediate nodes, which otherwise would need to perform delay, jitter and packet drop probability increasing L2/3 processing, in order to target similar bandwidth efficiency as is achieved with the Adaptive Concatenation.

Compatible access: OCS' **Intelligent Transport Network Interface Modules (IM)** units can forward packets at any packet layer protocol forwarding identifiers (e.g. MPLS Labels), and support common standard router/switch interfaces such as Packet-over-SONET Ethernet. OCS' services are straightforward and can be deployed in multiple scenarios such as the two described below.



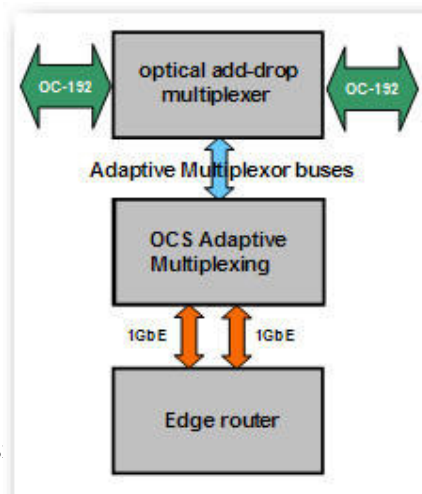
Corporate, ASP, CDN or wireless service provider backbone networks: An OCS Adaptive-Mesh can be deployed for a application, information or communications service provider backbone or a corporate Wide Area Network (WAN) application over any fibre optic network and interconnect any layer-3 or layer-2 routers / switches of the customer.

As an OCS' WAN solution requires only the *minimum theoretical network capacity necessary* to meet specified inter-site throughput etc. QoS requirements for a customer network, it is able to provide a highly cost-optimised solution.

Another benefit is the cost reduction accruing from not needing backbone bandwidth ports on a WAN's edge routers in case of ring topology, or from elimination of need for core routers or switches in case of hub-and-spoke topology. In the example in the diagramme above, only two 1Gbps ports are required connected to OCS' IM units over the WAN interfaces. As previously presented, conventional packet-switching network architecture would require in the order of ten times more packet processing and switching capacity per node, due to the need to handle traffic at packet-level also on the (e.g.10Gbps) backbone network interfaces.

Telecom service provider backbone network upgrades: With a service provider who currently owns at least some optical network capacity, the OCS' Adaptive-Mesh will work as a self-operating optimisation layer, between standard WDM Optical Add-Drop Multiplexors (OADMs) and the edge routers. (Note: A service provider does not need to add any new equipment to their networks due to OCS' Adaptive-Mesh based network upgrade, as OCS provides the Adaptive-Mesh connectivity as a managed service, not as an equipment vendor)

An OCS solution often reduces fibre-optic wavelength capacity requirements for a given inter-router throughput requirements to



a small fraction, such as to 1/10 or less, multiplying the time in years before the service provider would need to invest in additional WDM/fibre capacity, or alternatively, allowing economical use of per-customer-contract leased fibre/wavelength capacity via the Adaptive-Mesh based optimisation.

In case of upgrading packet ring topology for Adaptive-Mesh, the edge routers do not need to be upgraded for the new backbone network transport rate blades, while the network traffic throughputs get multiplied, so a significant reduction in router cost could be experienced in addition to the other architecture benefits discussed above. With upgrade to Adaptive-Mesh architecture from packet ring, each edge router is relieved from having to process the pass-through ring traffic, with the entire capacity of each edge router getting more effectively used for its local access traffic.

If the service provider buys in additional wavelengths to extend the reach of their network backbone or take over an enterprise's network through an outsourcing contract, the use of OCS' Adaptive-Mesh as an up-front-cost-free, transparent optimisation layer could result in considerable savings over the use of the traditional architectures.



Benefit inventory

Before all carriers jumped on the IP bandwagon in the late 90s, those that were prepared to 'think the unthinkable' and look towards trying a new approach were the first to see the benefits. This not only applied to network technologies but to applications as well. This is well demonstrated by the transition from frame relay services to IP-VPN services ([Where now frame relay?](#)). It's an easy decision to follow the crowd and it's often not a bad thing to do. It's safe and it's difficult to get things completely wrong. But when everyone has followed the same path there is little differentiation between competitive players. Everyone has the same services and makes the same promises and prices fall due to natural competition in an over-crowded market. It's now very difficult for any carrier to say "Our network is the best!" when that is also the strap-line for all their competitors as well.

There comes a time when forward thinking individuals and companies need to start to think the unthinkable again as was the case in the late 90s and to look again how their company could be made to stand out

from the crowd. Maybe some of the benefits listed below for the use of OCS' ITN could deliver some of that now hard to achieve customer-visible differentiation?

Flexibility: An Adaptive-Mesh service can be adopted by any organisation having a large network including enterprises or service providers whether they own a fibre-optic network or not, while producing similar, in the order of 10-fold cost efficiency gain.

Serious complexity reduction: An alternative approach to over-provisioning is the deployment of fine-grain traffic management techniques such as Deep Packet Inspection and micro flow policing. It could be considered that this really adds to the complexity of an IP / MPLS network which could make it less transparent and ultimately more unreliable. Because of this concern, a high proportion of carriers err on the side of simplicity and stick with over-provisioning. **It seems today that a network architect or planning engineer is in a no-win situation and have to choose between the costly simplicity of over-provisioning or the complexity of a micro managing the network.** OCS could provide an alternative way forward without the weaknesses of either approach, *without having to replace existing infrastructures, but actually making optimised use of it*

Potential customers of OCS

Examples of potential customer types below shows that there is no shortage of potential beneficiaries.

Large enterprises: Large Multinational Companies (MNCs) who operate their own private inter-site networks and buy in wide area network optical connectivity services in the form of SDH / SONET or wavelength services — in particular, media, entertainment, financial and high-technology sector organisations that require the highest levels of performance, QoS and security. Enterprises using OCS services could benefit from a significant reduction in bought-cost per Gbps and improved performance.

Business parks: Newer business parks often have a optical fibre networks centrally managed by the landlord, and/or plans for such. The use of Adaptive-Mesh architecture in such campus networks for communities of interest could significantly reduce implementation and management costs that could be passed on to their tenants.

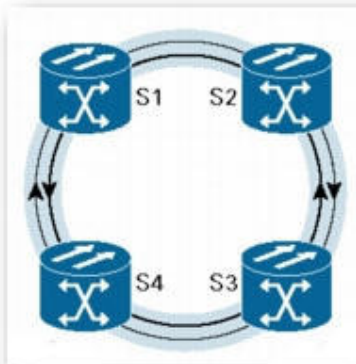
Software as a Service (SaaS), Content Distribution Services and Search: All business-class applications service providers operating at a global level currently have to operate and manage their own global networks. The reason for this is that if they are to provide the required performance and QoS for their paying corporate customers, they need to bypass the Internet and deliver services using their own private networks where they are able to control and manage the network. For these types of applications, Adaptive-Mesh allows the ASPs to get, as a managed service, the performance and controllability of private network with the Gbps/\$ cost-efficiency of Internet transit service.

System integrators: System integrators provide network outsource services to enterprises and must demonstrate that they can manage customers' networks at a lower costs while improving performance and quality. In practice, this is a tough act to deliver on and will always present a major challenge to the successful outsourcer. The adoption of an OCS' service, as a more cost-efficient, but equal, high-QoS alternative to traditional, leased line based services, could provide an additional weapon to help win new contract bids and really reduce the operating costs of already won out-sourced networks.

Virtual Network Operators (VNOs): VNOs provide connectivity services such as IP-VPNs to enterprise customers in competition to traditional carriers and system integrators. The difference is that they do not own their own networks and buy in connectivity services from global, regional or incumbent carriers. For VNOs, OCS Adaptive-Mesh architecture could provide a dedicated, private network like high quality backbone solution integrated over leased wavelengths, with multiple times higher cost-efficiency achieved due to that Adaptive-Mesh only requires a small fraction of the wavelength capacity, compared to the alternative architectures, to meet the given SLA terms for the wholesale customer contract. In many cases, the use of OCS' Adaptive-Mesh services can allow the customer such as VNO to get dedicated capacity based, premium QoS wholesale network services at cost per Gbps lower than even the bulk Internet transit services.

Mobile operators: Most mobile operator networks now manage their own private MPLS data networks. This is not only because of the rising use of mobile Internet data services but also due to their business plan needs to offer triple or quadruple play services that includes consumer Internet accesses. Moreover, traditional SDH connectivity services are used to connect cell sites back to the central Mobile Switching Centre ([MSC](#)) and represents a significant component of their costs. OCS' Adaptive-Mesh services allow multiplying the capacities of mobile core networks without associated capital costs, while reducing the level of ongoing operating costs, and while allowing efficient multi-service integration over high-performance, realtime optimised, dedicated layer-1 capacity based backhaul and backbone networks.

Fixed network operators: Clearly any traditional carrier who operates a traditional fibre optic network could utilise OCS services to help reduce the level of CAPEX and OPEX - particularly when looking at extending their network.



One alternative approach

Resilient Packet Ring (RPR), originally derived from Cisco's Dynamic Packet Transport (DPT), RPR is now subject to industry standardisation called IEEE [802.17](#). Based on a pair of 'counter direction' fibre optic rings, RPR offers hop-by-hop packet-switched connectivity on fibre or wavelength ring. It is based on a type of packet-layer (i.e. L2+) bandwidth

allocation called **Spatial Reuse Protocol** and appropriate sub 50ms restoration objectives.

RPR is particularly aimed at metro networks, however one of the key issues is whether its deployment is cost competitive to traditional hub-and-spoke approaches, which consume more layer-1/0 bandwidth, but save on router/switch hardware costs and provide better QoS due to fewer packet-hops. In relation to OCS' approach, which also provides minimised hop count based premium QoS, RPR does not achieve any higher bandwidth efficiency, so its architectural difference with Adaptive-Mesh is reduced to that RPR requires multiple times more packet processing and switching capacity at its nodes, causing higher cost and inferior scalability. Also RPR networks have multiple times higher packet-hop counts, possibly hurting QoS control. Unlike RPR, or any other purely packet-switched techniques, OCS does not need to process packets at backbone rates and uses an inherently higher QoS, more scalable and lower cost method of providing dynamic bandwidth allocation optimisation at layer-1.

Being in effect a packet based solution, RPR provides a packet buffering **Class of Service (CoS)** capability for all traffic on the rings to manage traffic in case of congestion. OCS take a much more fundamental way of segregating traffic by placing traffic with different destinations into *completely* different sets of aggregated NxSTS-1 'lanes' in the network called **Adaptive-concatenation Multiplexer Buses (AMBs)**. Using this approach, it is quite easy for potential user organisations to understand how it works and have more faith that a deterministic QoS can be achieved.

Naturally, the fundamentals of this comparison of Adaptive-Mesh to RPR applies to any other packet-ring type techniques, or generally any node-by-node packet-switching networks such as Carrier Ethernet, just the same.

Roundup

In [The new network dogma: Has the wheel turned full circle?](#) I wrote about possible opportunities to simplify enterprise and service provider networks through looking at new innovative approaches. The world has pretty much converged on a single vision and protocol set for Next Generation Networks (NGN) centred on network intelligence being based in layer-2.5 and above. Although this has much merit, I detect that many carriers want to keep their networks as simple as possible and have actively resisted deploying some of the newer protocols aimed at micro-managing individual packet flows. Maybe, layer 2, 2.5, 3 (and higher level) protocol approaches *are* actually getting too complicated?

OCS have introduced a concept that has much merit which could lead to considerable network simplification and all the benefits that could bring. This is in no way a battle based on religion. An OCS' solution, based on real-time, network byte-timeslot granular data throughput maximisation, is highly complementary to macro-level, and non/near realtime layer-2.5+ traffic engineering techniques, and since OCS' innovation is based on the fundamental physical layer (layer-1) of networks, its use benefits all services running over any layer-2+ protocol stacks. Importantly, non-carrier networks, such as enterprise WANs, could also benefit from many of the features inherent in an Adaptive-mesh network that is not available to them at present.



If you would like to learn more about OCS, visit their [News page](#) where you can download a number of papers that go into more detail than is possible here. If you are an enterprise networker, the Yankee Group report below is highly recommended.

[Innovation to Transform Market for Enterprise Networks](#)

⁽¹⁾ I should take the opportunity to declare a personal interest in OCS in that I am an advisor to the company. However, OCS has not in any manner sponsored, validated or endorsed this column and all opinions expressed or implied (other than direct quotes) are independent opinions of the author, and not of OCS.

⁽²⁾ OCS' technology patents can be downloaded from here: <http://www.pat2pdf.org> once there, search for US patent numbers 7333511, 7254138 and 7349414.

* * *