PERFORMANCE IMPLICATIONS OF CLOUD COMPUTING OVER SUBMARINE CABLE NETWORKS

Brian Lavallée, P.Eng., MBA

blavalle@ciena.com

Ciena, 2351 Boulevard Alfred-Nobel, Saint Laurent, Québec, H4S 2A9, Canada

Abstract: Cloud computing will have a profound effect on the global network infrastructure, including submarine cable networks that stitch continents together. The growing adoption of cloud services in healthcare, education, financial, and government sectors, served by large multi-tenant collocated data centers, continues to fuel bandwidth demand growth. As users, man and machine, are physically distanced from their storage and compute resources, the network binding them together becomes an absolutely critical component to the quality of experience. This paper reviews key design requirements related to cloud services that will be imposed upon submarine networks.

1. INTRODUCTION

Cloud computing is the latest technology trend sparking pronounced interest across a wide swath of industry sectors. While operating in hypercompetitive environments, corporations must constantly enhance the *Quality of Experience* (QoE) related to their Information Technology (IT) services, even as the number of supported users and applications steadily increases. Corporations have traditionally addressed their growing IT requirements by investing heavily in additional hardware and software, while subject to considerable internal pressures to reduce both capital and operating expenses to remain competitive and financially viable. Recent advances in storage, compute, and network virtualization are facilitating IT infrastructure outsourcing and has now become a viable option to corporations of all sizes. Consequently, *enterprise-class* infrastructure is moving further into the cloud at an accelerated pace through the commercial availability of on-demand services with robust levels of performance, security, and manageability.

Global network operators, characterized by owning and operating interconnected submarine and terrestrial networks, have been pursuing three main strategies in their quest for new revenue streams associated with cloud services.

- Networking – Providing carrier-grade network connectivity with global reach guaranteed by differentiated *Service Level Agreements* (SLA) that conceal the growing distances between users and their data centers

- Partnering – Aligning with existing cloud service providers to access their data center facilities and expertise in IT hardware/software management to offer bundled "cloud plus network connectivity" services

- Vertical Integration – Owning, operating, and growing their own geographically dispersed data centers via organic growth and/or through strategic corporate acquisitions for a complete cloud offering

Regardless of the chosen business strategy for offering differentiated cloud services, one thing remains constant; the network that connects end users to their data centers, and data centers to data centers, can be leveraged as a *key* competitive differentiator. The cloud network, connecting users to their data via submarine and terrestrial networks, must offer a similar QoE (or better) that users have become accustomed to when their data center is located in the same building or campus. This means the performance of the network that binds them together must be precisely architected and operated from end-to-end, overland and undersea, which places the global network operators contemplating the introduction of cloud services in an enviable competitive position, since they *own* the network assets and thus have far better control over its design, management, and evolution as cloud services are introduced.

According to the recent Reliability of Global Undersea Communication Cable Infrastructure (ROGUCCI) report[1], "nearly 100 percent of the world's inter-continental electronic communications traffic is carried by the undersea cable infrastructure" meaning the undersea jugular veins of intercontinental communication are a critical component to be taken into account when designing seamless end-to-end global networks that are optimized for cloud services. It also means that inevitable submarine cable faults can seriously compromise the performance and availability of cloud services, and thus must be highly resilient offering end-to-end network protection. Cloud computing places rigid performance requirements on submarine networks used to extend the reach of terrestrial networks into a global delivery vehicle for cloud services to end users. Network performance must be designed and guaranteed end-to-end with several key design requirements to be addressed, which are covered in subsequent sections of this paper.

2. CAPACITY

Cloud services are delivered to users by networks connecting them together, and must rapidly scale as the number of users adopting cloud services continues to proliferate. The increasing popularity of *Big Data*, high-speed mobile services, and video-centric content will only compound this bandwidth demand growth. To maintain pace with these bandwidth-hungry technology trends, data centers are increasingly being consolidated into monolithic physically shared locations, which are then *logically* combined into larger virtual *data centers without walls*. In both of these cases, the network between users and data centers, as well as data centers to data centers, will continue to experience significant growth in bandwidth. Coherent-based optical transmission technology addresses this growth.

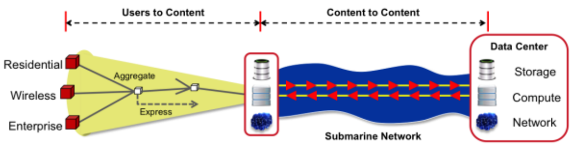

Figure 1: Global network user to data center and data center to data center traffic trends

There are two main strategies available to increase the capacity along submarine routes between continents. New submarine cables can be deployed along routes experiencing significant bandwidth demands, which is financially risky and can take many years from conception to *Ready for Service* (RFS). The faster, less expensive, and far less risky approach is to upgrade existing submarine cables from legacy 10 Gb/s channel rates to 40 Gb/s and 100 Gb/s coherent-based transmission technology rates, with the chosen rate determined by economics, target applications, and route distances. As coherent-based optical transmission technology continues to advance at a healthy pace, the de facto channel rate for submarine cable channel upgrades in 2013 will shift to 100 Gb/s away from 40 Gb/s, as a result of better economics, improved network performance, and matched line rate compatibility with 100 Gb/s Ethernet router/switch ports. It should be pointed out that upgrading existing submarine cables, although the more feasible option today, will not always be the case as submarine cables continue to age and will eventually approach their *End of Life* (EOL) status, which means new submarine cables will indeed have to be deployed in coming years.

Coherent-based optical transmission performance is superior over *uncompensated* submarine and terrestrial line systems. This means even better performance is possible when compared to existing *compensated* submarine cables that were designed and optimized many years ago for legacy 10 Gb/s transmission technologies. New submarine cables will benefit greatly from coherent-based optical transmission technology applied over newer D+ (positive dispersion) submarine-grade fiber with larger cores and without inline compensation. Compensation for *chromatic dispersion* and *Polarization Mode Dispersion* (PMD) is corrected in the more cost-effective electrical domain using mathematical techniques and *Digital Signaling Processing* (DSP) technology. Using a single optical fiber type over the entire submarine network, rather than the traditional multi-fiber architecture used to compensate for chromatic dispersion, lends itself to significant simplifications in designing, repairing via splicing, and inventory planning of cable maintenance vessels and shore-based submarine cable repair depots. The same coherent-based optical transmission technology can be used over terrestrial-based networks as well leading to further *economies of scale*. These tangible business and technology benefits are why savvy submarine cable planners will surely leverage this transmission technology for future submarine wet plants, since its value propositions are simply to hard to ignore.

Figure 2: Programmable modulation capabilities of agile coherent-based transponders

3. AVAILABILTY

Enterprise-class IT services were traditionally delivered by in-house data centers over *Local Area Networks* (LANs) that connected users (employees) to their data and applications. LANs, given the *local* nature of the network, are relatively easy to design, upgrade, troubleshoot, and repair. Physical access to LAN equipment and wiring is relatively quick and easy, especially when compared to locating, troubleshooting, and repairing submarine cable faults located hundreds of kilometers away across rough seas at depths of several kilometers. Multiple submarine earthquakes in the Asia-Pacific region over the past few years have emphasized the rather hostile environments of submarine cables, which can be covered with tons of sediment and debris after associated subsea landslides. What is quite clear is that if *enterprise-class* cloud services are to continue to proliferate over increasingly longer distances, submarine cable networks operating in inherently hostile subsea environments must incorporate robust resilience to maintain the same or better network availability that cloud-based users will expect. It is simply unfathomable for a submarine network that connects several large data centers overseas to go down for hours, regardless of the extent of the damage incurred by manmade causes and/or natural disasters. The best technology for ensuring the very highest levels of network availability in the most hostile of operating environments is an *intelligent mesh network*.

Submarine networks have traditionally been protected using *ring-based* topologies using SDH-based MS-SPRing (Multiplex Section-Shared Protection Ring) technology. Although this network protection technology will maintain availability during a single submarine cable fault, severe degradations or outages result when multiple faults occur, which is often the case when submarine earthquakes strike leading to several submarine cables damaged in multiple locations over a relatively short timeframe. Achieving a network availability of 99.999% (five 9s) is possible with a ring-based topology; but there is a far better alternative that offers improved network availability. Properly designed intelligent mesh networks with multiple available protection paths can achieve a network availability as high as 99.9999% (six 9s) resulting in only ~ 31 seconds of total annual network downtime and represents an order of magnitude improvement when compared to traditional ring-based network protection. The number of diverse paths chosen is determined by a variety of factors such as the risk level of the region (ex. amount of shipping traffic and/or the frequency and severity of earthquakes experienced annually), business economics, and targeted customers. Savvy global network operators have recognized the resilience of intelligent mesh networks, which is a key reason why they have been deployed over the past decade in submarine and terrestrial networks. It is also why enterprises themselves have purchased unprotected services over multiple submarine cables and then implemented their own mesh network to ensure the availability of their data services, even while operating in high-risk regions of the world.

Beyond increased resilience and improved network availability, mesh protection offers numerous other advantages to network operators, which further explains its growing popularity and acceptance. Some of these enviable values enabled by intelligent mesh networks include faster turn-up of highly differentiated services, latency-based routing, improved overall network bandwidth utilization that reduces operating costs through shared protection bandwidth, and a network-wide connection database that is updated in near real-time. Mesh network *control plane* intelligence also lays the foundation for *Software-Defined Networks* (SDN) and *Performance-on-Demand* (PoD) capabilities where applications autonomously and dynamically manage underlying network connections. All of these benefits afforded by implementing control plane intelligence will help proliferate cloud-based services on a global scale. The alternative is to continue building out cloud networks based on peak demands that over time can lead to a substantial amount of wasted network resources and higher operating costs that reduce competitive positioning in the long term.

4. SECURITY

Security is a primary concern that must be addressed and allayed if the adoption of enterprise-class cloud services is to achieve its maximum potential. Different stakeholders can have very different perspectives on security and where in the cloud it should be applied. Data at *rest* and *inflight* data can be addressed using various types of encryption with the actual encryption scheme decided by end user requirements. Inflight data can be encrypted from end-to-end using AES-256, for example. Regardless of the inflight encryption scheme chosen, the associated latency incurred must be kept to a strict minimum, which is achieved by implementing the encryption engine in hardware, preferably within the transponder cards. Since the level of security required is highly subjective, it is beyond the scope of this paper to provide the various advantages and disadvantages of the various encryption schemes available today.

Coherent-based optical transmission technology has side benefits, as it pertains to increasing the *obscurity* of inflight data given the nature in which the technology is commercially implemented. When compared to legacy *Intensity Modulated Direct Detection* (IMDD) transponders, coherent-based transponders provide additional data obscurity benefits due to unique frame structures, interleaving patterns, soft-decision *Forward Error Correction* (FEC), and encoding techniques. The technology decisions chosen to ensure coherent-based optical transmission technology provides network operators with maximum performance, reliability, and agility at the lowest cost, also ensure that wavelengths transported over submarine cable networks are more *secure* than legacy implementations. One of the most vulnerable points in the network from a security perspective are nodes where propagating signals are converted back into the electrical domain and is why strict physical access guidelines are implemented in both *Cable Landing Stations* (CLS) and inland data centers, for example. Although securing physical access to these locations will continue forward, it is easier to do so at a manned inland data center compared to a remote, unmanned CLS. As such, maintaining signals in the optical domain as it passes through the CLS is another benefit enabled by coherent-based optical transmission technology, which is one of the many benefits of an *Optical Bypass* (OB) configuration.

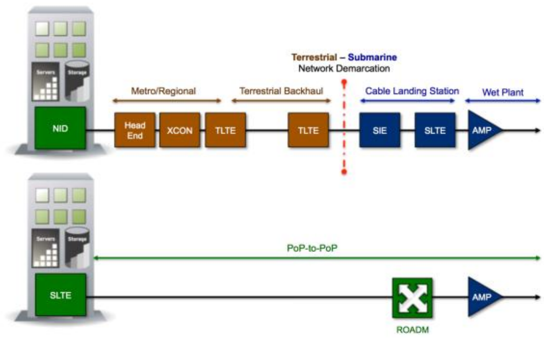

Figure 3: *Cable Landing Station* simplification leveraging *Optical Bypass* architectures

5. LATENCY

There is far more to coherent-based optical transmission technology than the line rate itself. The full value of this innovative transmission technology is realized when the entire ecosystem of related technologies are taken into account. The coherent technology ecosystem includes soft-decision FEC, enhanced OEO subsystems, high performance DSP, high performance analog control, and advanced modulation schemes. Although this ecosystem was initially developed for long haul terrestrial networks, due to a wealth of ecosystem-enabled benefits, it quickly found its way into regional, metro, and submarine networks, all using a common technology base. Of particular interest is the greatly extended reach of this ecosystem of technologies that allows global network operators to think beyond the *beach* and design their networks from inland Point-of-Presence (PoP) to inland PoP. Gone are the days of being forced to design three networks; namely, the *beach-to-beach* submarine network using one set of transmission technologies, and the two terrestrial backhaul networks from the beach to the inland PoP using another set of transmission technologies. Global network operators can now leverage the coherent ecosystem of technologies and design networks from *PoP-to-PoP* (ex. data center to data center) while keeping the data within the inherently more secure optical domain as it passes through the *beach* by leveraging another technology initially developed for terrestrial networks, the *Reconfigurable Optical Add/Drop Multiplexer* (ROADM) configured in an OB architecture.

As submarine and terrestrial networks were traditionally designed and deployed using quite different technologies, primarily due to the transoceanic reaches and 25-year life spans related to submarine networks, a clear submarine to terrestrial demarcation point was created within the CLS. Although a proven reliable architecture, it is now seen as a highly inefficient and often unnecessary architecture given the advent of coherent-based transmission technology, which can be deployed over both submarine and terrestrial networks using a common technology base. The increased reach often yields sufficient link budget to extend reach well beyond traditional beach-to-beach network designs of the past meaning the transported data can be kept in the optical domain across all of the three network segments by using a ROADM in the CLS. The ROADM eliminates the traditional submarine-terrestrial demarcation point, along with a substantial amount of equipment leading to significant savings in power, space, complexity, cost, and latency. The latency savings are achieved by eliminating multiple *Optical-Electrical-Optical* (OEO) stages over the PoP-to-PoP route, primarily through the elimination of equipment in both the CLS and the terrestrial backhaul network. Although the main contributor of latency in a submarine network is the actual length of the submarine cable itself and the propagation of light, minimizing the latency within the CLS and all the way back to the inland PoP can make the difference in meeting the stringent latency requirements of cloud-based services. Intelligent mesh networks, governed by a control plane, support latency-based routing to ensure that newly created connections will meet specific application latency requirements, as well as the protection path during inevitable network faults.

6. ORCHESTRATION

Network orchestration, which refers to the dynamic creation and deletion of on-demand services, requires a *network hypervisor* to *virtualize* underlying submarine (and terrestrial) network assets, similar to compute and storage virtualization available today. The network hypervisor allows cloud operating systems to access network assets from PoP-to-PoP, while completely concealing the complexity of the underlying network architecture. Virtualizing global network assets interconnecting data centers not only maximizes bandwidth utilization over time by way of statistical multiplexing, it also allows for new and innovative applications to be developed and offered to end users leading to highly coveted, differentiated services, which is the goal network operators strive towards to grow revenues streams.

The network hypervisor, acting as the *gateway* into mesh network *control plane* intelligence, executes commands on behalf of the cloud operating system (OS) to dynamically create/delete connections based on various PoD SLA criteria such as Bit-Error-Rate (BER), latency, packet loss/jitter, and availability. Coherent-based transponders used with OB are key orchestration building blocks operating at the optical layer. Current coherent-based transponders offer programmable modulation schemes for different reach vs. capacity applications, as well as programmable soft-decision FEC where the provisioned error-correcting performance is *right-sized* to the required reach, which minimizes the amount of incurred latency. Agile ROADMs configured in an OB configuration allow incoming wavelengths to be switched/routed to either the inland PoP or onwards along the submarine cable network to the next CLS served, all under software control. ROADMs and OB can also provide optical layer protection as well leveraging geographically diverse terrestrial backhaul routes. Although these optical layer building blocks are directly programmable, they will likely be accessed through a network hypervisor with the underlying control plane managing network assets at layer 0, 1, and/or 2, which will facilitate SDN development and 3rd party control.

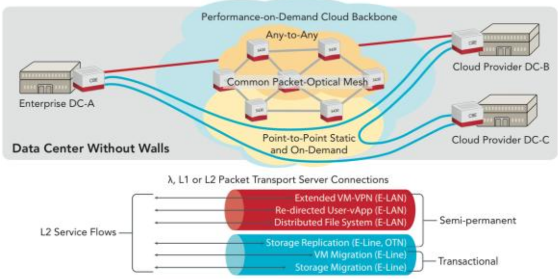

Figure 4: Varied, dynamic cloud-based traffic patterns and characteristics

7. GLOBAL CLOUD NETWORKS

The evolution of the cloud toward expanded enterprise IT utility is driving the creation of a true *data center without walls* concept that logically integrates multiple network operator and enterprise data centers while supporting very different types of traffic flows along with their associated performance requirements. It is precisely these different types of dynamic traffic flows, associated performance requirements, and evolving business dynamics that will impact submarine networks much as it is already impacting terrestrial networks, especially as data centers are moved further away from end users, whether they be man and/or machine. Eliminating the legacy demarcation point between submarine and terrestrial networks is a fundamental move towards facilitating cloud services being offered on a global scale. Other strategies, such as *TCP application acceleration* and *Content Delivery Networks*, will also be leveraged as distances between users and their data continue to increase as cloud service providers consolidate data centers to fewer, more manageable locations while expanding coverage to increase revenues streams.

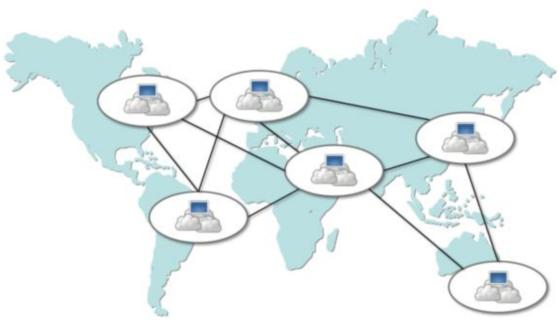

Figure 5: Global cloud network interconnected by intelligent optical mesh networks, overland and undersea

Coherent-based optical transmission technology quickly and reliably increases network capacity and reach, while laying the foundation for elastic bandwidth via programmable modulation and soft-decision FEC. OB simplifies network designs from CLS to inland PoP to reduce latency while offering dynamic ROADM-based wavelength switching. Intelligent mesh networks offer improved network resilience and associated availability. Control plane *intelligence* facilitates elastic bandwidth allocation by acting as the connection management engine. It abstracts the complexity of the underlying network assets by executing connection management requests from applications, such as cloud OS, via the network hypervisor gateway. By combining coherent-based transmission technology, optical bypass, intelligence mesh networks, and network hypervisor functionality, cloud computing demands on a larger more global scale can be realized – anywhere and anytime, thus breathing new life into global networks, overland and undersea. As legacy demarcation points that once clearly separated distinct submarine and terrestrial networks are eliminated, managing global end-to-end network connections is thus facilitated allowing global network service providers to offer and guarantee cloud-centric SLAs, which fully leverages ownership of their global network assets.

8. REFERENCES

[1] Proceeding of the ROGUCCI Study and Global Summit, Issue 1, Revision 1, 2010, IEEE Communications Society