Business Case for Shared Mesh Protection

Executive Summary

Current approaches to network resiliency are inadequate to meet evolving network performance and cost requirements. Existing schemes such as 1+1 protection meet the sub 50 ms performance requirement but only protect against single failures and are too costly. Best-effort approaches such as software-based GMPLS mesh restoration are cost effective, handle multiple failures but do not meet the sub 50 ms performance requirement. MPLS FRR can protect against multiple failures and achieve local 50 ms performance but requires longer time frames for end-to-end convergence and uses more costly router ports.

Several industry vendors in IETF and ITU are defining a new standards-based approach called Shared Mesh Protection (SMP) for network resiliency. Some vendors such as Infinera are implementing this solution with hardware acceleration, allowing it to deliver both the economic efficiency of shared backup paths in the more cost effective transport layer and 50 ms end-to-end recovery even in the face of multiple fiber cuts.

ACG Research conducted a total cost of ownership (TCO) comparison of SMP versus 1+1 protection. The comparison is made for the TCO of line-side 100 Gbps WDM interfaces using a national reference transport network and a five-year study. It models traffic patterns to/from data centers, cable landing sites, and metro areas. Traffic increases at 85 percent CAGR over the study period. The comparison shows that the TCO for protection resources in SMP is 27 percent less as compared to 1+1 protection. The TCO savings result from the use of shared bandwidth managed by network intelligence to protect against multiple failures versus dedicated backup resources for single failure protection used by 1+1.

KEY FINDINGS

SMP is a new standard that provides both economic efficiency and better network resiliency. When compared to 1+1 protection SMP has:

- 27% less TCO over five years
- $110 million lower total costs over five years for a meshed network of 87 nodes
- Better survivability supporting multiple failures
- Capability to enhance revenue by offering multi-tiered service levels

Introduction

Networks and the traffic they carry are evolving from ring to mesh topologies and from circuit to packet traffic loads. End-users' expectations of service availability also are increasing, which is driving operators to design their networks to meet increasingly stringent service level agreements (SLA). These SLAs reflect the growing dependence of business and society upon network availability, and even a short service outage can have severe financial impacts and impair a service provider's reputation.

Network availability, in particular, must be maintained even during natural disasters and other catastrophic events. Public safety and other disaster response services use the network to manage their operations and communicate with the public. Availability requirements are increasing from 0.9999 (4 nines) availability to 5 nines or even 6 nines availability. These requirements correspond to downtime per year of 53 minutes, 5.3 minutes, and 32 seconds, respectively. In addition, service recovery expectation is less than 50 ms, which has been the target for protected SONET/SDH systems for more than a decade.

At the same time, network costs driven by high traffic growth rates are rising rapidly, but revenues are not keeping pace. Network operators, consequently, are being pushed to reduce costs while simultaneously forced to increase network performance.

Figure 1 maps the current and proposed resiliency mechanisms by network performance (survivability, speed) and bandwidth efficiency.

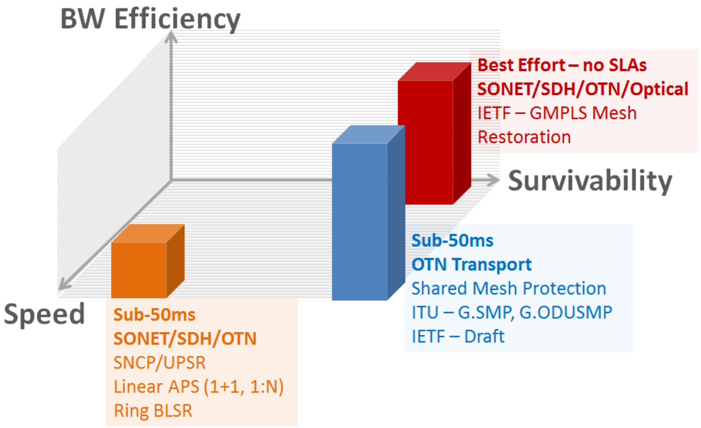

Figure 1 – Network Resiliency Techniques

Existing SONET/SDH/OTN protection schemes such as SNCP/UPSR/Linear APS/BLSR provide deterministic sub 50 ms protection but at high cost. For example, the most commonly used 1+1 protection schemes reserve one-half of all bandwidth for protection. This ensures that service to network bandwidth utilization can never be higher than 50 percent but in practice is much lower. Also, 1+1 protection supports single failures only. Resiliency schemes that use a software-driven mesh restoration approach (ASON/GMPLS) support bandwidth sharing that handles multiple failures and therefore, achieve higher bandwidth utilization, but they often provide slower and unpredictable performance within several seconds. Since guaranteed recovery times are increasingly necessary, this approach is not suitable for service providers that must deliver stringent SLAs to their customers.

Spurred by Infinera's initial contribution to IETF, several industry vendors in IETF and ITU, such as Ericsson, Huawei, Verizon, ZTE and others, are actively contributing to working committees[1] to create a new standards-based approach called Shared Mesh Protection (SMP). Deployment of mesh topologies, increases in processing speeds, and improvements in intelligent transport network control planes make SMP, a new protection protocol, feasible. Some vendors, such as Infinera, are implementing SMP with fast hardware-based lookup tables in order to handle multiple complex failure scenarios with 50 ms performance—which makes sense in the face of service providers deploying 100 Gbps meshed networks with up to 8 Tbps of fiber capacity and thousands of services. The promise of SMP is that it eliminates choosing between designing for networks that require reliable performance or for those that reduce cost. SMP, defined at Layer 1 (OTN layer), enables the more cost-effective transport layer[2] to simultaneously provide protection and multiple failure survivability.

Shared Mesh Protection Approach to Protection and Fast Recovery

A network with six nodes is used to illustrate the SMP approach to protection that maximizes bandwidth efficiency (See Figure 2).

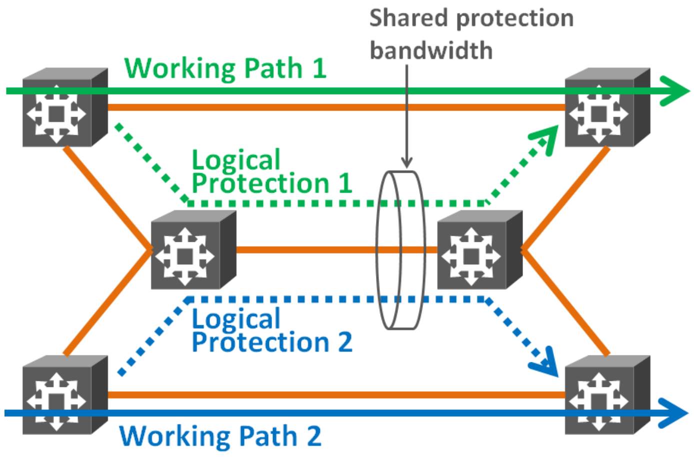

Figure 2 – Six Node Example Network with SMP

Network capacity is defined as pools of working resource and protection resource. The capacity of the protection resource pool is shared by multiple logical protection paths. Figure 2 shows the available fiber-optic connections in orange. The working paths of two Sub-Network Connections (SNC) are shown in green and blue, and two logical protection paths are shown as dashed lines. Note that the logical protection paths share not only one fiber-optic connection, but also share the same resources on that connection. This is shared protection bandwidth. This is a simple model for illustration purposes only; in large-scale mesh networks there are many alternate backup paths available.

[1] See ITU, Q9/SG15 G.smp and G.ODUSMP; IETF, MPLS working group draft-pan-shared-mesh-protection-05.txt and draft-weingarten-mpls-smp-requirements-03.txt.

[2] MPLS Fast ReRoute (MPLS FRR) also provides sub 50 ms performance and shared use of protection bandwidth. However, its local protection approach can be limiting in terms of failure coverage; SMP provides a network-wide protection mechanism. Also, MPLS FRR costs more than SMP because it employs expensive IP/MPLS router ports.

Each logical protection path is configured to register network resources when it is established, but no actual protection bandwidth is consumed until the protection path is activated in hardware. Link states (availability of protection bandwidth and paths) are maintained on the network elements. It is critical that this activation is done in hardware to achieve the sub 50 ms recovery, and SMP implemented with this technology can be deemed fast, guaranteeing a deterministic performance in large scale networks.

Figure 3 illustrates the network after logical protection Path 1 has been activated.

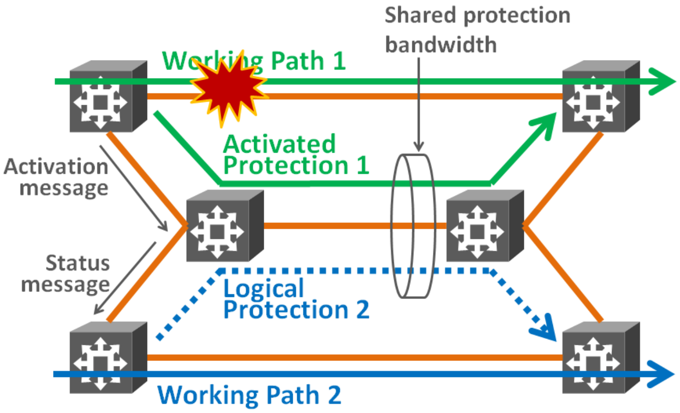

Figure 3 – Example Network with Protection Path #1 Activated

Figure 3 shows that a fault has occurred on working Path 1. This triggers an activation message that is sent to the downstream network element, which creates predefined cross-connections along the path of the activation message. Head-end cross-connection is made once the neighbor node acknowledges the activation message. A status message also is sent to the working Path 2 network elements for notification to select appropriate protect paths on subsequent failures. Once the original connection has been restored to service, traffic can revert to the original home path, and the protection path resource is released back to the protection resource pool. This dynamic aspect enables the SNC to be resilient against multiple failure scenarios not possible in the 1+1 protection scheme.

For multifailure scenarios, several levels of backup can be defined to take advantage of mesh topologies, ensuring services continue to operate on the 2nd, 3rd or Nth alternate backup. Multilevel priorities and pre-emption rules can be established to ensure guaranteed sub 50 ms protection for higher priority services while maximizing network utilization for all services.

The calculation of protection paths can be made by distributed network element-based calculations and/or by a centralized system. Additionally, planning tools can be used to simulate network failures and behavior deduced from general principles to ensure deterministic network-wide performance and to achieve global optimization of resources.

Benefits of Shared Mesh Protection

The advance calculation of protection paths combined with hardware-based protection path activation enables delivery of end-to-end sub 50 ms recovery while maximizing network utilization. This provides sub 50 ms protection similar to 1+1 protection schemes. SMP, however, provides superior survivability than 1+1 protection by leveraging the mesh topology and intelligent control plane improvements to provide multiple protection paths and support dual and arbitrary failures. SMP also lowers costs because it uses shared resources at the transport layer. For more aggressive operators, in addition to simply making their transport network more reliable, it provides an opportunity to move protection for some or all services across the network from expensive packet-based router ports using MPLS FRR to lower cost switched transport (OTN) ports.

SMP also can be used to increase revenue by offering multi-tiered service levels—a classic revenue enhancement strategy. The service levels could range from best effort service that is pre-emptible with no protection or guaranteed recovery time to a premier level service that provides protection for multiple simultaneous failures, guaranteed recovery in less than 50 ms and is not pre-emptible. Intermediate service tiers can be created by selected combinations of protection levels, recovery guarantees, and pre-emptible service designations.

TCO Comparisons for a National Reference Transport Network

A TCO comparison is made between 1+1 protection and SMP by simulating the rollout of 100 Gbps WDM interfaces over five years on a national reference transport network[3] (See Figure 4).

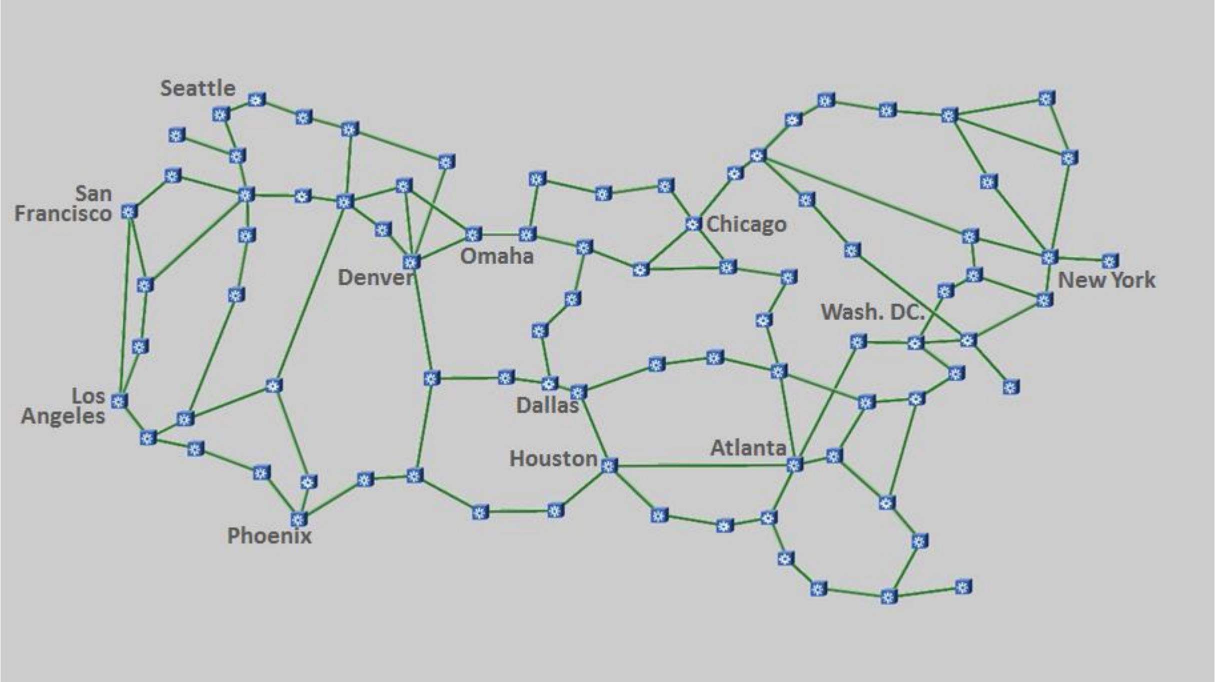

Figure 4 – National Reference Transport Network

[3] See "Total Cost of Ownership of WDM and Switching Architectures for Next-Generation 100Gb/s Networks" (http://www.infinera.com/pdfs/whitepapers/Infinera-Economics_of_Optical_Network_Design.pdf) for a full description of the national reference transport network.

The reference network depicts a U.S. long-haul network and is a composite of the major U.S. long-haul networks. Table 1 summarizes its principal features.

Feature	
Nodes	87
Links	114
Tier1 Data Centers	San Francisco, Seattle, Washington DC, Omaha, New York, Dallas
Tier 2 Data Centers	Atlanta, Los Angeles, Chicago
Cable Landing Sites	New York, Norfolk, Hillsboro, Boca Raton, San Luis Obispo

Table 1 – Principal Features of National Reference Transport Network

Each node is a traffic source and sink. The network has a highly meshed topology with an average degree of 2.62. The data centers are located and sized using public data sources and Infinera's data. The protection schemes are designed to minimize the total cycle (total of working path plus protection path).

Figure 5 provides a five-year projection of total bandwidth flow; Table 2 shows the client interface mix.

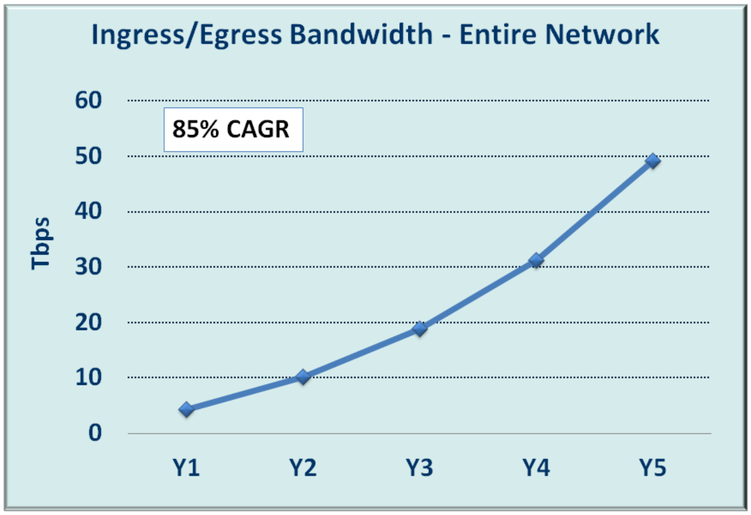

Figure 5 – Ingress/Egress Bandwidth Entire Network

Year	Sub 10G Total	10G	40G	100 Gbps
Y1	3372	1181	0	0
Y2	3868	1703	0	4
Y3	4262	2436	0	14
Y4	4624	3403	24	31
Y5	5010	4328	44	117

Table 2 – Client Interface Mix

The sources for the traffic projections include Infinera, Data Center Map, ACG Research, Infonetics, Telegeography and Ovum. Three traffic flow patterns are modeled with data centers, cable landings, and metro areas. A gravity generated traffic model is used to model flows between the nodes[4].

TCO Modeling Results

The TCO comparison between the 1+1 and SMP protection methods is made by modeling the build-out of line-side 100 Gbps WDM interfaces[5]. Identical transport network equipment is used in the design and cost computations for the two transport solutions:

1. 1+1 protected network
2. Shared Mesh Protected network

Note that fast and robust mesh-based protection is only possible using a digital OTN switch that is best implemented in an integrated OTN and WDM platform. This allows the transport layer to be more efficient and intelligent than the transponder/ROADM only architecture. To provide a meaningful comparison of the two protection techniques we consider current industry pricing for the 100 Gbps line-side WDM interfaces.

Figure 6 shows the protection-only TCO for each scheme at the end of five years. Savings of $110 million in total costs, capital expenses (CapEx) and operation expenses (OpEx) can be realized for the U.S. reference network by using SMP as compared to 1+1 protection.

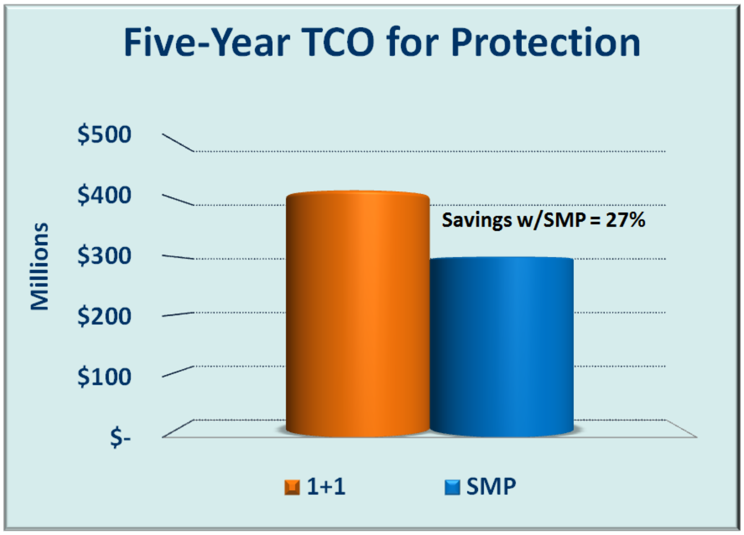

Figure 6 – Five-Year Cumulative Protection-Only TCO (CapEx + OpEx) Comparison

The shared mesh approach uses network capacity more efficiently by employing shared network connections; 1+1 protection dedicates at least half of the capacity of the network for protection. SMP uses fewer WDM interfaces and has correspondingly lower TCO.

[4] Traffic between nodes is proportional to the product of node size for each node pair and inversely proportional to the distance between each node pair.

[5] The cost of optical amplifiers that are required on longer network links is not included in the analysis because it is nearly the same for both protection schemes. Common OpEx items between the two protection schemes are excluded from the analysis. Only items that change between the schemes have been considered.

Figure 7 compares the protection-only 100 Gbps WDM interface count.

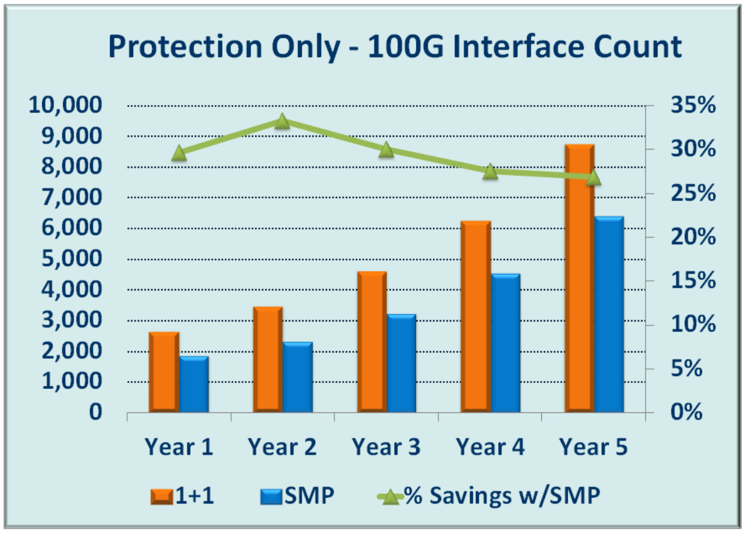

Figure 7 – Protection-Only 100 Gbps WDM Interface Count

The difference in the 100 Gbps WDM interface count between the two schemes has a direct impact on the protection CapEx (Figure 8). The model begins with savings of 30 percent in CapEx required for protection while using SMP as compared to 1+1 protection in the initial years, ending with 27 percent by Year 5. The Year 1 protection CapEx for SMP is normalized to 100. All other numbers in Figure 8 are calculated from this base to provide a gauge of the difference between 1+1 protection and to show the impact of traffic growth on protection CapEx.

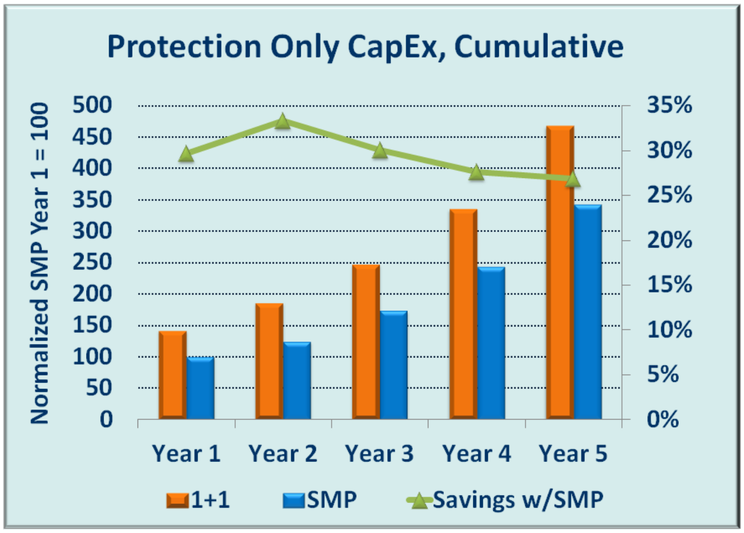

Figure 8 – Protection-Only CapEx, Cumulative

Figure 9 compares the cumulative protection OpEx required between SMP and 1+1 protection. Various elements have been considered as a part of the OpEx, including power costs, cooling costs, floor space costs, network care costs and installation costs. The Year 1 protection OpEx for SMP is normalized to 100. All other numbers in Figure 9 are calculated from this base to provide a gauge of the difference between 1+1 protection and to show the impact of traffic growth on protection OpEx. The results identify a direct correlation with the protection 100 Gbps WDM interface count as expected, more than

30 percent savings in protection OpEx for SMP as compared to 1+1 scheme in the initial years and ending with 27 percent by Year 5.

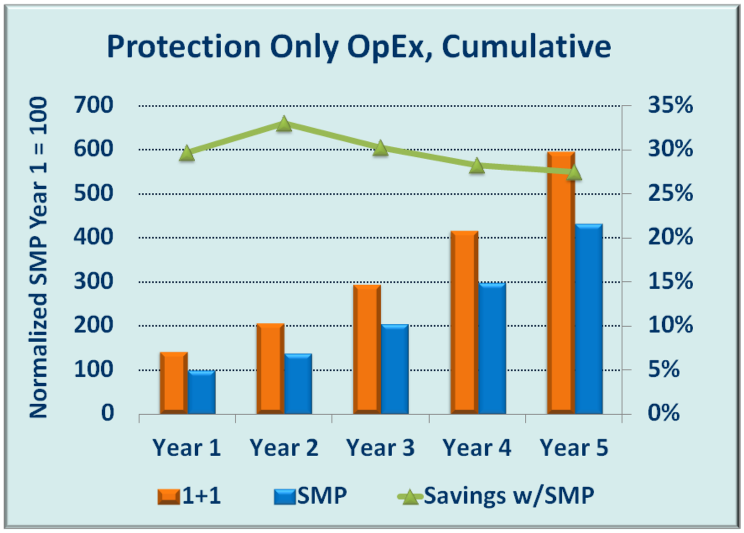

Figure 9 – Protection-Only OpEx, Cumulative

SMP saves TCO as compared to 1+1 protection in every study year. The shared mesh approach is more cost effective, maintains protection levels that are better than 1+1 protection and delivers cost performance similar to that of best-effort restoration methods.

Conclusion

Networks are evolving from ring to mesh topologies and from circuit to packet traffic loads. At the same time network designers are being challenged to meet expanding network scale and increasingly stringent SLAs. Network revenue growth, however, is not keeping pace with the cost to scale up the network and meet these SLA requirements. This dilemma is driving a new approach to network protection.

Current approaches to network resiliency are inadequate to meet evolving network performance and cost requirements. Existing sub 50 ms recovery schemes such as 1+1 protection meet the sub 50 ms performance requirements but are too costly and do not handle multiple failures. Best-effort approaches such as software-driven mesh restoration are cost effective, handle multiple failures but do not meet the critical sub 50 ms performance requirement.

SMP, a new standard being developed within the ITU and IETF, provides both better economic efficiency and performance in handling multiple failures than other network resiliency technologies. Implementing SMP in hardware ensures a sub 50 ms performance, which is necessary to meet stringent SLAs as networks scale to 100 Gbps and 8 Tbps per fiber with thousands of services.

An economic comparison made between 1+1 protection and SMP using a national reference network model shows:

- Protection TCO is 27 percent less for SMP as compared to 1+1 protection
- The cumulative TCO for SMP is lower than that of 1+1 protection for every study year

- The source of the TCO savings is the use of shared bandwidth that is 100 percent protected for single failures and priority pre-emption for multiple failures by SMP as compared to 1+1 that supports just single failures

SMP resolves the conflict between the need to minimize cost and meet SLAs[6]. The shared mesh approach is more cost effective, maintains protection levels that are better than 1+1 protection and delivers cost performance similar to that of best-effort restoration methods.

ACG Research is an analyst and consulting company that focuses in the networking and telecom space. Our best-in-class subject matter analysts have a combined 120+ years of experience and expertise in telecom segments that address both technology and business issues. We offer comprehensive, high-quality, end-to-end business consulting and syndicated research services.

[6] While not included in the calculations in this model, service providers to whom we have spoken believe that there are additional revenues based on new tiered protection schemes that can be extracted from this new capability. This area is still under investigation.