

Economic and Architectural Benefits of Hierarchical Backbone Networks

Jane M. Simmons

Corvis Corporation

7015 Albert Einstein Drive, Columbia, MD 21046-9400

jsimmons@corvis.com

Abstract

The explosive growth in data traffic will necessitate that service providers deploy networks with hierarchical logical structures. Hierarchical architectures provide significant cost savings by enabling a large degree of optical bypass and by taking advantage of cost-effective banded switch technologies.

Economic and Architectural Benefits of Hierarchical Backbone Networks

Jane M. Simmons

Corvis Corporation, 7015 Albert Einstein Drive, Columbia, MD 21046-9400

jsimmons@corvis.com

The explosive growth in data traffic will necessitate that service providers deploy networks with hierarchical logical structures. In a hierarchical network, each layer is designed to hide the details of its layer from those above, thus enabling cost effective and scalable network growth. At each layer, the appropriate traffic granularity, switching and grooming technology, and both physical and logical topologies must be selected. Significant cost benefits can be realized in both capital and operational expenditures. Multi-tiered architectures can: take advantage of cost-effective banded switching technologies as described in [1]; enable a large amount of optical bypass, which reduces the amount of required electronic termination equipment; and reduce the number of elements that need to be managed, thereby simplifying network operation. Also, with respect to restoration, a hierarchical scheme provides efficient use of restoration capacity, reduces the amount of electronic access equipment that needs to be deployed for the capacity to be effectively shared, and simplifies the restoration decision process [1].

A good example of a hierarchical architecture is that of the two-tiered scheme shown in Figure 1. In this architecture, the upper layer is comprised of localized collector rings, operating at granularities of DS-3 or OC-3, whereas the lower layer consists of a nationwide 'express mesh' operating at granularities of OC-48, OC-192, and higher. Each collector ring is typically populated with SONET Add/Drop Multiplexers (ADMs), with all nodes in the ring being logically connected (although not necessarily on all wavelengths). The functions of the collector ring are two-fold: route traffic within the ring, and collect inter-ring traffic and deliver it to the large mesh nodes located on that ring. The ring topology is suitable for this layer due to the limited geographic extent of the collector rings and the simple protection properties of rings. Furthermore, many carriers already have made huge investments in SONET-ring-based networks. The capacity of these rings is currently being exhausted by the tremendous growth in traffic, exacerbated by the inefficiencies of inter-ring routing. These legacy rings can be more effectively utilized with their role restricted to serving as the collector network – SONET ADMs are well suited for adding, dropping, and grooming traffic, and current capacity demands will be lessened by migrating the inter-ring traffic to an express mesh. Alternatively, a mesh architecture with grooming cross-connects can be used for the collector network.

As opposed to the fine granularity, limited extent nature of the collector network, the lower-layer express mesh of Figure 1 must carry large amounts of traffic for long distances. These requirements are best met with emerging ultra-long reach, high-capacity transport systems and all-optical switches. Since the collector rings perform the function of aggregation, the express mesh can carry traffic at coarser granularities – OC48 and above. Not all network nodes are included in the logical mesh; only those nodes that generate traffic at levels comparable to the granularity of this level, or those that are strategically located, should be included. Thus, depending on traffic patterns in the network, the express mesh will completely bypass a major portion of the nodes. To take full advantage of this architecture, ultra-long reach transport must be deployed such that traffic can be routed between (and through) lower-layer nodes without any intermediary regeneration. The combination of a sparsely populated express mesh (roughly 20 nodes nationwide) with ultra-long reach transport eliminates a significant portion of electronic termination in the network [2].

In addition to limiting the nodes that are included in the lower layer, it is also important to create a 'streamlined' physical topology. Thus, fiber links that do not pass by major nodes may not be included in this layer. First, not populating a fiber route eliminates the capital and operational costs of deploying amplifiers on that route. Second, streamlining the underlying topology results in larger bundles of traffic that can be routed at lower cost; i.e., restricting the diversity of routes produces more efficiently packed traffic pipes. While removing fiber routes results in greater traffic on the remaining routes, next-

generation high-capacity transport systems can accommodate such high bandwidth demands. In addition, while fewer fiber paths may result in somewhat longer end-to-end paths for some connections, this very likely does not lead to more regeneration if ultra-long reach systems are deployed. Of course, restoration options must also be considered when laying out the physical topology.

A mesh topology is most appropriate for the lower level since, through the use of switches, it creates virtual topologies that can be optimized for the routing of working traffic and the sharing of protection bandwidth. One of the most significant benefits of a hierarchical architecture, in terms of cost and scalability, is that it enables the deployment of a hierarchical switch architecture, as shown in Figure 2. The streamlined architecture of an express mesh creates large bundles of traffic that can be switched as a single unit. For example, increasing the switching granularity by a factor of B decreases the switch fabric size by a factor of somewhere between $B \log B$ and B^2 , depending on the switch technology. This directly translates to a saving in switch cost, switch size, number of physical interfaces, power, and switch states that must be managed. In Figure 2, we assume there are two levels of switches, operating at granularities of G_1 and G_2 , where G_1 is coarser than G_2 . With efficient packing algorithms, only a small portion of the traffic needs to be switched at the finer granularity; the remainder is more cost effectively switched at the coarser level. While the switches can operate in the optical or electrical domain, to best take advantage of an express mesh, optical switches should be used for the coarse granularity traffic. This enables large bundles of traffic to express through the node without any electronic termination.

Thus far, we have discussed a two-tiered approach to network design. Ultimately, what will drive the number of layers and the granularity of each layer is cost and technology. Assume that the cross-sectional traffic on any link in a particular layer of the network is T , and assume that the finest level of granularity required at that layer is G . Then the number of elements per link that are potentially switched are T/G . At a node of degree- N , this translates to a switch size of approximately $(NT/G) \times (NT/G)$. Due to limits of technology and physical space, the switch size may be limited, resulting in multiple tiers, where, for the i^{th} tier, T_i and G_i are chosen such that NT_i/G_i remains tractable. This will enable networks to continue to grow gracefully.

We have presented a qualitative argument for the economic and architectural benefits of deploying hierarchical networks, and outlined how these benefits can be best realized with new generation transport and switching systems.




Figure 1. An example of a two-level hierarchy, where a group of collector rings aggregates and grooms fine granularity traffic, and delivers traffic to a streamlined, coarse granularity express mesh.




Figure 2. A hierarchical architecture can take advantage of hierarchical switching. At the lower level, an OXC switches traffic at a coarse granularity (G_1). Only a small percentage of the traffic needs to be switched at the finer granularity (G_2).

References:

[1] Simmons, J.M., "Hierarchical Restoration in a Backbone Network," *OFC'99*, Feb. 21-26, 1999, San Diego, CA.
 [2] Saleh, A.A.M., "Transparent Optical Networking in Backbone Networks," *OFC'00*, Mar. 7-10, 2000, Baltimore, MD.