

Innovation to Transform Market for Enterprise Networks

Conflicting Performance and Cost Requirements for Enterprise WANs Resolvable with Architectural Innovation

yankee
group

www.yankeegroup.com

by Sandra Palumbo | August 2008

Executive Summary

Current and emerging trends for business communications and IT applications place enterprise interoffice networks, or wide-area networks (WANs), under a stress of having to simultaneously support a set of seemingly conflicting requirements. This situation calls for profound modernization of architectural models for enterprise network services to enable a new standard for WAN performance and cost-efficiency that both meets the customers' business needs and delivers them direct bottom line benefits at once.

Regarding the conflicting requirements for WAN services, on the one hand, at the time of pressure for more efficient IT spending, there is a need to rapidly scale up the enterprise WAN capacities to support new and emerging high-bandwidth enterprise applications such as video conferencing and other real-time dynamic interaction among multiple locations. Moreover, several ERP applications and real-time data replication or rapid data backup requirements further increase the WAN data throughput demands. Consequently, there is a need for enterprises to get multiple times more WAN capacity per unit cost. To reduce the cost for WAN services per unit of access capacity, service providers are offering packet-layer-shared WANs (i.e., virtual private networks or virtual private LAN services), which connect the sites of multiple customers over shared packet-switched network infrastructure. While such packet-layer-shared WAN services can improve the service provider network utilization, and avoid the need for per-customer-dedicated network capacity, they suffer from non-deterministic performance (e.g., unpredictable throughputs and latencies), ongoing security risks and complicated management, and therefore are not able to satisfy all the current—and much less the emerging—high-performance communications and IT applications requirements.

On the other hand, while there is need to rapidly multiply the WAN capacities per unit cost, the emerging mix of applications will demand high quality of service (QoS) and security guarantees that traditionally have been most effectively met with premium-priced dedicated WAN services. Unfortunately, the traditional dedicated-capacity-based WAN services do not provide the level of cost-efficiency that would make them an economically viable alternative, given the exploding volumes of traffic that need to be supported.

As none of the prevailing WAN technologies are able to effectively address the full set of requirements for enterprise networking of the near future, the industry is poised to undergo a profound transformation both in terms of technological standards and business models, with the winners being defined by their willingness to acknowledge the architectural limitations of the current WAN technologies and by proactively pursuing the market-transforming innovations that will emerge given the unsustainability of the economics with the conventional WAN products and services.

Accordingly, key findings and recommendations of this Yankee Group Report include:

- The majority of enterprise IT staff time is still being spent on non-value-adding routines such as maintaining, optimizing and upgrading IT/network infrastructure. This calls for innovations that enable more intelligent and highly automated, adaptive technology infrastructure for enterprises, provided ideally as managed services, to enable enterprises gain a better return on their IT spending by allowing them to allocate more of the IT staff time for more value-adding, customer-facing and differentiating functions.

Continued on next page

This custom publication has been sponsored by Optimum Communications Services, Inc.

© Copyright 2008. Yankee Group Research, Inc. All rights reserved.

This Yankee Group Report is published for the sole use of Yankee Group clients. It may not be duplicated, reproduced or transmitted in whole or in part without the express permission of Yankee Group, Prudential Tower, 800 Boylston Street, 27th Floor, Boston, MA 02199. For more information, contact Yankee Group: info@yankeegroup.com; phone: 617-598-7200. All rights reserved. All opinions and estimates herein constitute our judgment as of this date and are subject to change without notice.

Executive Summary (continued)

- None of the prevailing WAN architectures—dedicated and packet-layer shared—is able to provide all of the identified key requirements for an effective WAN solution: reliability, security, performance, manageability, scalability, flexibility and cost-efficiency. The traditional WAN technologies leave a trade-off, forcing the enterprise customers in any situation to forfeit some of these key requirements to achieve others. However, with the convergence of communications and requirements for real-time multi-location interaction, all the above criteria are quickly becoming absolute necessities for the modern enterprise. This creates an urgent need for architectural innovation allowing enterprises to break free from the affordability-versus-performance trade-offs with the traditional WANs.
- Innovation addressing the above pressing needs among the large enterprises is in deed happening. However, the much-needed innovation tends to be led by intellectual-property-driven startup companies, such as Optimum Communications Services, Inc. (OCS) and Matisse

Networks, rather than the incumbent vendors. To bring these innovations to their customers, the service providers and systems integrators must be willing to create technology and channel partnerships with the innovative startups to incorporate the cutting-edge innovations in their enterprise telecom/IT solution offerings. Ignoring the unmet customer requirements, or innovations allowing to address them, causes a provider to run the risk of being locked to obsolete technologies and business models amid changing enterprise network markets. Likewise, the enterprises should not accept the status quo with the trade-offs of the WAN solutions, but instead the enterprises' IT leadership need to press their network providers to bring them network solutions that utilize the latest innovations to deliver to the enterprise the type of WAN solution that meets their actual business needs at affordable cost.

Table of Contents

I. Enterprise WAN Trends and Drivers	3
Methodology	4
II. What Enterprises Want from WAN Providers	4
III. Choosing the Right WAN Architecture	5
Enterprise WAN Architecture Options	5
Traditional Dedicated WAN Solutions	5
Packet-Layer-Shared WANs	5
The Need for Innovation	7
Enterprise Grids	8
IV. Conclusions and Recommendations	11
Summary Conclusions	11
Recommendation for Service Providers and Systems Integrators	11
Recommendations for Enterprises	11

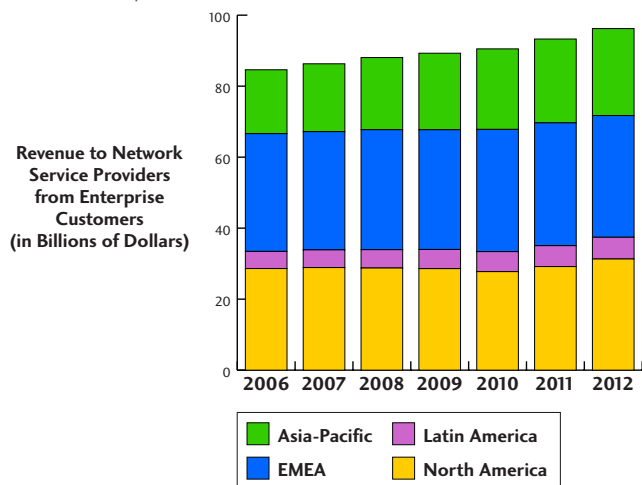
I. Enterprise WAN Trends and Drivers

As enterprises become increasingly dependent on high-performance network connectivity among their operational facilities and global workforce, customers and suppliers, Yankee Group continues to see stable growth in WAN services market globally, from \$85 billion in 2007 to \$95 billion in 2012 as shown in Exhibit 1.

Thirty-nine percent of large businesses currently outsource network operations and maintenance to third parties, with another 29% currently considering the outsourcing of these functions, according to the Yankee Group *Anywhere Enterprise—Large: 2007 US Managed IT and Communications Services Survey* (see Exhibit 2). Expansion of operations, corporate security, business continuity/disaster recovery and implementation of business applications are some of the top business issues influencing communications decisions for 2008. Internal management of the WAN and communications environment is not a top issue and is often viewed as a necessary evil rather than a business enabler. This fact explains the strong desire among large enterprises for assistance in managing and maintaining their network infrastructure. Many businesses that chose outsourcing or managed services did so because they believe it adds value to the business by enabling their IT teams to better focus on the core business and support the more specific end-user needs, rather than addressing ongoing network operations issues.

Exhibit 1
Global WAN Services Market Forecast

Source: Yankee Group, 2008



Note: The geography refers to the location of the enterprises' operational headquarters. Dollar figures projected assume zero inflation.

Among enterprises that use service provider managed WAN solutions, the top reasons for their decision to outsource their WAN management were the pursuit of improved cost-efficiency and reliability of enterprise networking (see Exhibit 3). Managed services inherently allow the customer to reduce their network capital expenditures while facilitating more strategic allocation of enterprise IT resources.

Exhibit 2
Enterprise Interest in Outsourcing Network Operations and Management

Source: Yankee Group *Anywhere Enterprise—Large: 2007 US Managed and IT Communications Services Survey*

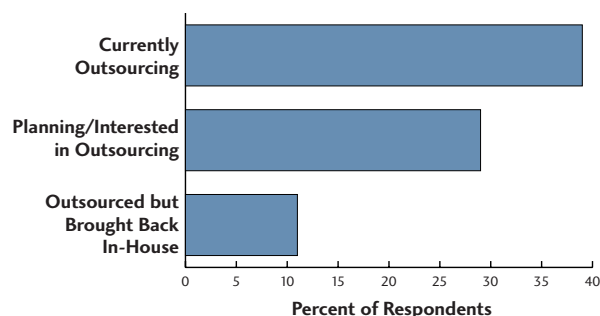
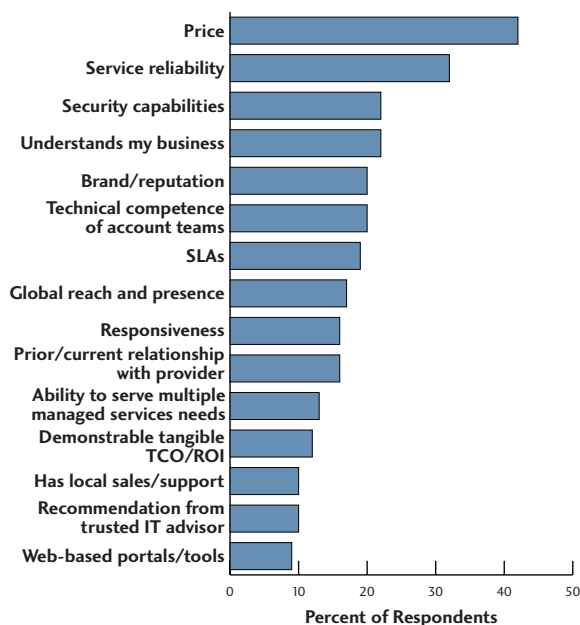


Exhibit 3
Price, Reliability and Security Top Reasons for Choosing Managed Services Providers

Source: Yankee Group *Anywhere Enterprise—Large: 2007 US Managed IT and Communications Services Survey*



Methodology

This Report leverages Yankee Group's vast existing research and analysis of enterprise WAN requirements and solutions, large business survey data on the use and adoption of communications services, and five target multinational corporate (MNC) interviews specifically covering traditional and innovative alternative WAN architectures.

II. What Enterprises Want from WAN Providers

Enterprises are looking to simplify the management of their technological operating environments. According to Yankee Group data, enterprise IT organizations spend 54% of their time maintaining and optimizing/upgrading their IT and networking infrastructure (see Exhibit 4). That the majority of enterprise IT personnel time currently has to be spent on the such non-value-adding routine functions means that there is an urgent need for "self-operating" enterprise technology infrastructure or managed services that would allow companies to allocate their IT resources in a more productive manner (e.g., to develop applications that improve the competitiveness of the enterprise).

Aside from the overall desire to move to more intelligent yet simpler managed services, enterprises also seek the following key requirements in their WAN solution:

1. **Reliability:** Deterministic, pre-definable network operation with enforceable service availability (e.g., 99.9% or 99.999% of contract duration).
2. **Security:** Certainty that an organization's WAN capacity and communications over the WAN stay protected from unauthorized uses and users.
3. **Network performance:** The capability to support real-time services such as high-quality voice or video, as well as high-bandwidth services such as data backups.
4. **Manageability:** The ability to ensure that the network operates in the intended way.
5. **Flexibility and scalability:** The ability to continue using a chosen architecture despite a changing or growing business.
6. **Cost-efficiency:** The amount of guaranteed inter-site network capacity per a given cost budget.

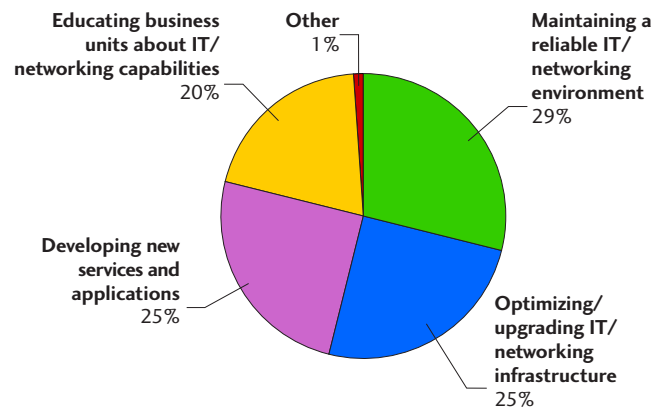
The above key WAN solution selection criteria are reflected in current RFPs for WAN operations and management contracts being issued by MNCs. The primary research for this Report indicated that reliability and security were the most important WAN architecture selection criteria, while cost-efficiency was given the lowest importance among the respondents. However, as shown in Exhibit 3, price is still the single most important criterion when it comes to actually purchasing WAN services. Moreover, according to our primary research, the enterprises and service providers implicitly assume that dedicated-capacity-based WANs—while known to provide better reliability, security and performance—are in many case disfavored simply because they are considered prohibitively expensive.

These findings highlight the fundamental desirable characteristics versus affordability trade-off that exists between the conventional WAN architectures. In particular, while many of the traditional WAN technologies can achieve a subset of the above criteria, none of them excel in each of them—as would be required to meet the actual needs of the enterprises. Consequently, the enterprise customers currently need to prioritize some of the above essential criteria at the expense of others, leaving them with a WAN that is suboptimal in one or more ways. Typical trade-offs with prevailing WAN solutions include, for example, that if the enterprise prioritized security and reliability over any other criteria, they would go for a dedicated-capacity-based WAN, and ended up with a solution that was inferior to internet-based alternatives in terms of cost-efficiency and flexibility of network access. Alternatively, if another enterprise had to prioritize cost-efficiency and flexibility over other criteria, it would choose an internet-based or other packet-layer-shared WAN service, which generally is inferior to dedicated WAN in terms of performance, security, reliability and manageability.

Exhibit 4

More than Half of Corporate IT Staff Time Is Consumed by Non-Differentiating Functions of IT/Network Infrastructure Maintenance, Optimization and Upgrading

Source: Yankee Group Anywhere Enterprise: Large—2007 US Managed and IT Communications Services Survey



Though not supported by the prevailing WAN technologies, enterprises actually want a secure, simplified WAN architecture that does not carry extraordinary costs to sustain or scale. As one enterprise we spoke with stated, “We want affordable, sustainable capacity...and we want to ensure that applications get their share of bandwidth; we don’t over provision. Thus, the WAN architecture has to support the business and not become a millstone around its neck.”

III. Choosing the Right WAN Architecture

Enterprise WAN Architecture Options

To comprehensively address the network services needs of enterprises, network technology developers and service providers and integrators need to come up with truly innovative solutions that do not just focus on certain subset of the essential requirements for WAN solutions, but instead allow them to break free from the trade-offs with the traditional WAN architectures. In the next sections, we analyze traditional as well as emerging innovative WAN solutions from the perspective of how they perform against all the essential WAN solution requirements by enterprises.

Traditional Dedicated WAN Solutions

Conventional per-customer-organization dedicated-capacity-based WANs provide a straightforward, albeit not cost-optimized, way to deliver deterministic-QoS network connectivity among the customers sites connected to the WAN. With traditional network technologies, at least one dedicated, constant bandwidth Layer 1 (L1) network connection is needed in the WAN per each one of the connected sites.

One common, dedicated WAN implementation architecture is so-called hub-and-spoke, where all the WAN access sites are connected to a central hub for packet-layer switching. However, to eliminate single points of failure in the network, which is often necessary for the WAN to reach its required level of availability, at least two mutually protecting instances of the hub switching facility need to be implemented. Moreover, because in many WANs there are subgroups of mutual interest (e.g., national networks within a global corporate WAN), many dedicated WANs implement a hierarchy of various stages of regional, national and international hubs. Such hierarchical hub-and-spoke versus a simple global hub-and-spoke architecture is often implemented to limit the transmission network capacity requirements and associated costs at the expense of increased complexity of the network.

In any case, because the dedicated WAN capacity is based on per-customer-site dedicated, fixed bandwidth L1 connections and per-customer dedicated packet layer switches/routers, the network utilization is not optimized. Effectively, dedicated WANs provide static network capacity allocation even though the inter-site capacity demand variations can be highly dynamic. Such static capacity allocation requires the network to be overengineered; essentially such dedicated, fixed bandwidth L1 point-to-point connection-based WAN needs to be overprovisioned ahead of time for all possible traffic distribution scenarios, though only one of those traffic load scenarios will actually materialize at any given time.

Advantages

- Dedicated architectures—hub-and-spoke, for example—provide a straightforward way to achieve high reliability, security and performance.

Disadvantages

- Non-optimized, statically allocated usage of the network capacity leads to high network capacity requirements and therefore high cost (in terms of the cost per Mbps of network access capacity).

Characteristics of Appropriate Enterprise Adopters

- Requirements for deterministic network performance (i.e., for strict network security, reliability and QoS guarantees, which the dedicated WANs can achieve via isolation of different customers’ traffic in their separate L1 subnetworks).
- Customer can afford the premium price for the dedicated, non-optimized capacity-based WAN.

Packet-Layer-Shared WANs

The rapidly increasing cost of traditional dedicated WANs with the increase of the number of enterprises’ network access sites, international reach and bandwidth requirements have created a market opportunity for more flexible and lower cost (in terms of the cost per Mbps of access capacity) packet-switched WAN services. The prevailing packet-layer WANs, or Layer 2/3 virtual private networks (L2/3 VPNs), seek to achieve improved network utilization, and therefore lower costs for the service provider by collapsing traffic from multiple (and up to all) customers of the provider to a single packet-switched pool of network capacity (see Exhibit 5 on next page).

With such packet-layer-shared WANs, traffic belonging to multiple different enterprise WANs share—and compete for—capacity in the same packet-layer switches/routers and for bandwidth in the L1 connections between the switches/routers. The service providers have

often preferred to share network resources among multiple customers, unless the customer expressly or via SLA terms required dedicated capacity, in order to reduce their costs per a given WAN service contract. Sharing of the network infrastructure at the packet level naturally can improve the utilization level of the network capacity, and accordingly the service cost per unit of network access capacity to enterprise normally is lower with the packet-layer-shared WANs than with traditional dedicated WANs.

However, in packet-layer-shared networks, traffic patterns of any one enterprise WAN will affect the performance of all other WANs implemented over the same pool of shared network capacity. The fact that with conventional packet-layer WANs several independent WAN applications from different customers use the same shared network capacity causes that none of the individual enterprise WANs gets deterministic QoS.

Advantages

- Lower cost per unit of network access capacity compared with traditional dedicated WANs.

Disadvantages

- Non-deterministic QoS; unpredictable performance.
- Exposes an enterprise’s WAN traffic at packet level to multi-user shared networks, making the packet-layer-shared WANs vulnerable to ongoing network security threats.
- Complicated packet layer routing/switching network administration and route control because of the need for the WAN switches/routers to perform multi-domain packet switching/routing between multiple customer and service provider domains.

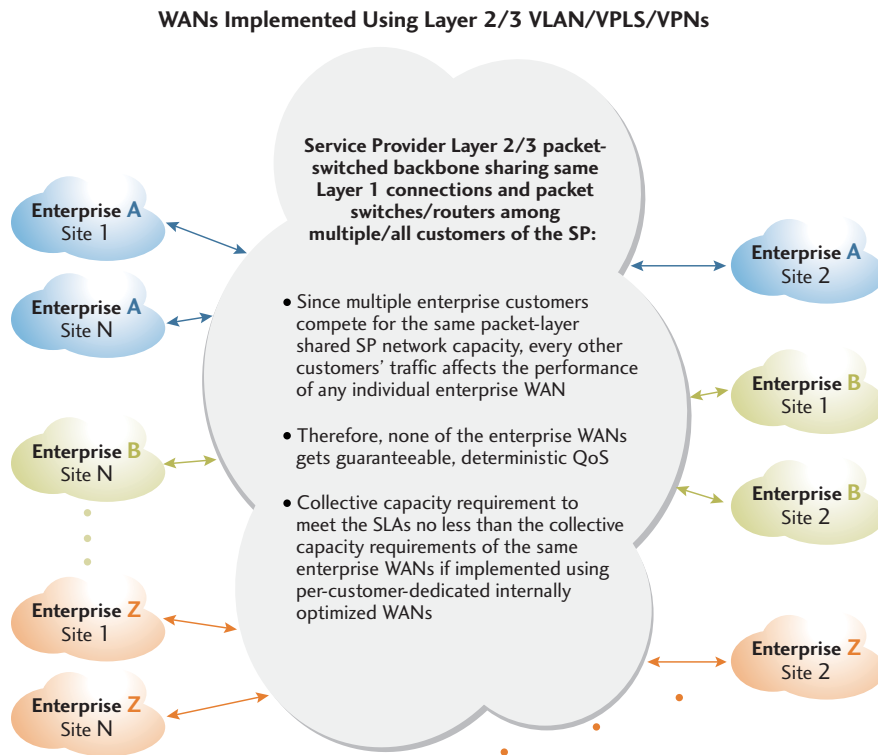
Characteristics of Appropriate Enterprise Adopters

- The customer can tolerate unpredictable inter-site network throughputs and latencies.
- Can afford the cost and complexity of providing communications security at application level, as the packet-layer shared WAN are inherently non-secure.

Exhibit 5

With Conventional Packet-Switching WAN Services, Multiple Customers of the Service Provider Share and Compete for the Same Pool of Network Resources, Causing Non-Deterministic Performance for any Individual Enterprise WAN

Source: Optimum Communications Services and Yankee Group, 2008



The Need for Innovation

Exhibit 6 summarizes how the traditional WAN architectures compare against each other, as well as against an assumed ideal WAN, with respect to key WAN solution selection criteria. It is shown that the capabilities of the traditional WAN architectures to meet these WAN selection criteria are fully non-overlapping. That is, it is not possible with any of the traditional WANs to achieve each of the desired attributes of an enterprise WAN. Therefore, architectural innovation is needed to overcome the architectural deficiencies of the traditional WANs to deliver an enterprise network solution that allows an enterprise to simultaneously achieve all the essential desired attributes of its WAN in one package.

As shown in Exhibit 6, while traditional dedicated WAN solutions can achieve good reliability, security, performance and manageability, they are not sufficiently flexible, scalable or cost-efficient for most large enterprise customer needs. On the other hand, while packet-layer shared WANs can achieve better flexibility, scalability and cost-efficiency than traditional dedicated WANs, they lack the essential reliability, security and performance qualities.

It is also seen that with the conventional WAN alternatives, there are certain interdependencies between their functional attributes. For instance, the use of dedicated L1 capacity (called leased lines) has

traditionally meant that there will not be any dynamic capacity allocation optimization, confining the traditional dedicated WANs to suboptimal capacity utilization levels, thus making the conventional dedicated networks not sufficiently economical for delivering high-bandwidth, dynamically burstable packet traffic. Also, a service provider-managed packet-switched WAN service has traditionally meant that the WAN is implemented over a multicustomer shared packet-switched network, causing complex L2/3 (and even higher level) protocol network administration and difficult QoS enforcement across the multiple operator domains involved.

With such undesirable interdependencies inherent to the prevailing WAN technologies it is not possible to implement a WAN solution that would deliver all the essential WAN selection criteria. Consequently, enterprises have been left no choice but to accept WAN solutions that are suboptimal with respect to at least some of the essential characteristics of enterprise networks. However, trying to support the emerging mix of business communications and IT applications, requiring simultaneously hard real-time QoS guarantees and radically improved cost-efficiency, is not sustainable. As a result, innovation needs to be brought to the enterprise WAN market so all of the traditionally irresolvable, conflicting requirements can be effectively addressed.

Exhibit 6

None of the Traditional WAN Alternatives Can Achieve All of the Desired WAN Attributes

Source: Optimum Communications Services and Yankee Group, 2008

		Per-customer-dedicated L1 network connections	Per-customer-dedicated packet switches/routers	Direct customer controllability of the packet switches/routers	Service provider deployed and operated packet switches/routers	Bandwidth efficiency via burstable, demand-driven capacity allocation
For which WAN selection criteria is this factor essential?	Reliability	✓	✓	✓		
	Security	✓	✓	✓		
	Performance	✓	✓			
	Manageability			✓		
	Flexibility, Scalability				✓	✓
	Cost-Efficiency				✓	✓
Traditional Dedicated WAN Customer-dedicated L1 leased lines connecting customer operated routers/switches		Yes	Yes	Yes	No	No
Packet-Layer-Shared VPN Customer network configured at packet layer over publicly shared service provider network		No	No	No	Yes	Yes (locally per a given shared link)
"Ideal" Customer-dedicated, self-optimized pool of L1 capacity with customer-controllable dedicated switches		Yes	Yes	Yes	Yes	Yes (globally for any customer network)

Enterprise Grids

A few innovative startups—including Optimum Communications Services, Inc. (OCS) and Matisse Networks—are developing dynamic Layer 1/0 based solutions that appear to overcome some major limitations seen in more mature WAN technologies. Because of the wider application of the technology, this part of the report focuses on OCS' Enterprise Grid™ WAN solution, based on OCS' Intelligent Transport Network™ (ITN) technology platform. Matisse is targeting only the metro networks.

Enterprise Grid offers transparent packet-switched network connectivity over per-customer-dedicated pools of internally optimized L1 capacity, which provides the bandwidth, security, performance and scalability of a wavelength with the bandwidth efficiency of packet switching. Based on patented Adaptive-Mesh network architecture, Enterprise Grid provides a mesh of dynamic bandwidth L1 circuits for carrying variable bandwidth packet traffic flows between the grid access points. The Adaptive-Mesh automatically adjusts to the real-time inter-site traffic load variations by dynamically optimizing the capacities of the mesh of L1 circuits through the customer's dedicated network bandwidth pool, thereby continuously optimizing the WAN traffic throughput. Despite the novel adaptive mesh of L1 circuits within the Enterprise Grid, the customer access interfaces to Enterprise Grid are standard point-to-point links (e.g., MPLS over SDH/SONET or Ethernet links commonly supported by the customer WAN access equipment). Enterprise Grid thus is designed for direct interoperability with the existing enterprise (or provider edge) L2/3 switches and routers.

The real-time traffic load adaptive L1 capacity allocation-based network throughput maximization of Enterprise Grid significantly reduces the amount of physical network capacity required to support given enterprise WAN contract SLA. Consequently, the real-time self-optimizing Adaptive-Mesh architecture of Enterprise Grid makes the use of per-customer-dedicated pools of L1 capacity more economical than the use of multicustomer shared capacity with conventional packet layer WANs—which at physical layer still rely on the inefficient, non-adaptive capacity L1 connections.

The use of dedicated L1 subnetworks for each individual enterprise customer WAN with the Enterprise Grid model significantly simplifies the packet forwarding across Enterprise Grid WANs, compared to conventional packet layer shared WANs, such as regular L2/3 VPN or VPLS networks. In particular, the packet forwarding engines of any given Enterprise Grid network only need to do simple single-domain packet forwarding instead of complex multi-domain packet switching or routing that the conventional packet switching WANs have to support. This unique feature of Enterprise Grid brings

about significant network operations simplification and security benefits by eliminating the need to exchange and sync the packet-layer switching and routing information between the enterprise and the service provider network domains.

Moreover, OCS' Enterprise Grid implementation allows fully packet-layer transparent packet forwarding across the service provider network based on predefined rules for the packet-forwarding engines in Enterprise Grid to interpret the packet forwarding instructions (e.g., MPLS labels, Ethernet VLAN tags, IPv6 flow IDs) placed in the packet headers by the enterprise's routers/switches connected to the Enterprise Grid. Thereby, Enterprise Grid WANs do not need any packet-layer routing, switching or forwarding tables, which further simplifies the WAN implementation and management, while eliminating any chance of the service provider routing tables being misconfigured or out of sync (which problem would exist in conventional packet-switching WANs). Furthermore, the packet-layer transparent, any-to-any connectivity offered by Enterprise Grid among the interconnected customer network access points allows the customer switches/routers to handle all their packet layer (L2 and above) protocol interactions directly with each others through the transparent Enterprise Grid WAN, further simplifying the network operations.

Accordingly, the Enterprise Grid model provides each enterprise customer with its own dedicated capacity based packet-switched network service with deterministic QoS, built-in security and real-time optimized data throughput, while not requiring collectively any more capacity for a group of enterprise customers than what would be necessary with conventional packet-layer-shared WANs to meet the same customer SLAs (see Exhibit 7 on next page).

According to network case studies by OCS, the Adaptive-Mesh architecture of Enterprise Grid can reduce fiber-optic network wavelength capacity requirements by a factor of up to 20-to-1 compared with alternative, equal-quality networks relying on conventional, non-adaptive-bandwidth physical layer connections. For instance, the Adaptive-Mesh architecture can provide non-oversubscribed packet-switched connectivity among up to 20 10-Gbps network access points over a single 40-Gbps optical wavelength ring, whereas when using alternative, comparable-quality technologies, a pair of mutually protecting core switches/routers and a dedicated optical wavelength loop (whether 10-Gbps or 40-Gbps capable) for each of the 20 access points (i.e., a total 20 optical wavelengths) would be needed. When thus designed using Adaptive-Mesh, an up to 20-times improved network cost-efficiency can be achieved compared with traditional, similar quality network alternatives.

In addition, the Adaptive-Mesh architecture enables in the order of tenfold reduction in packet processing and switching equipment capacity requirements compared with conventional hop-by-hop packet-switched shared network solutions. To achieve optimized network bandwidth utilization, Adaptive-Mesh nodes (unlike conventional packet-switching nodes) do not need any packet-layer functionality for their transport (ring) network interfaces, but only for their access interfaces. Therefore, in case of, e.g., doubled 1-Gbps packet-switched access to 10-Gbps transport ring, whereas a conventional packet-switching node needs to provide packet-level processing and switching for both of its 10-Gbps ring interfaces (10 Gbps + 10 Gbps) as well as its access interfaces (1 Gbps + 1 Gbps) for a total of 22 Gbps of packet processing and switching capacity, an Adaptive-Mesh node in the same case needs to provide packet-forwarding for only for its access interfaces for a total of 2 Gbps of packet processing (and no packet-switch) capacity required. Thus, the Adaptive-Mesh architecture in such cases reduces packet level processing capacity requirements per a node by a factor of ten or higher.

Importantly, the network cost-efficiency gains achieved with the Adaptive-Mesh architecture of OCS do not come with any performance drawbacks. On the contrary, the use of adaptive-bandwidth L1 circuits for packet transport across multiple nodes in Adaptive-Mesh networks causes a significant reduction in packet-hop counts compared with conventional packet-switching network technologies that seek to improve network utilization via performing packet-level switching at each of the network packet-switched network access points for all the traffic.

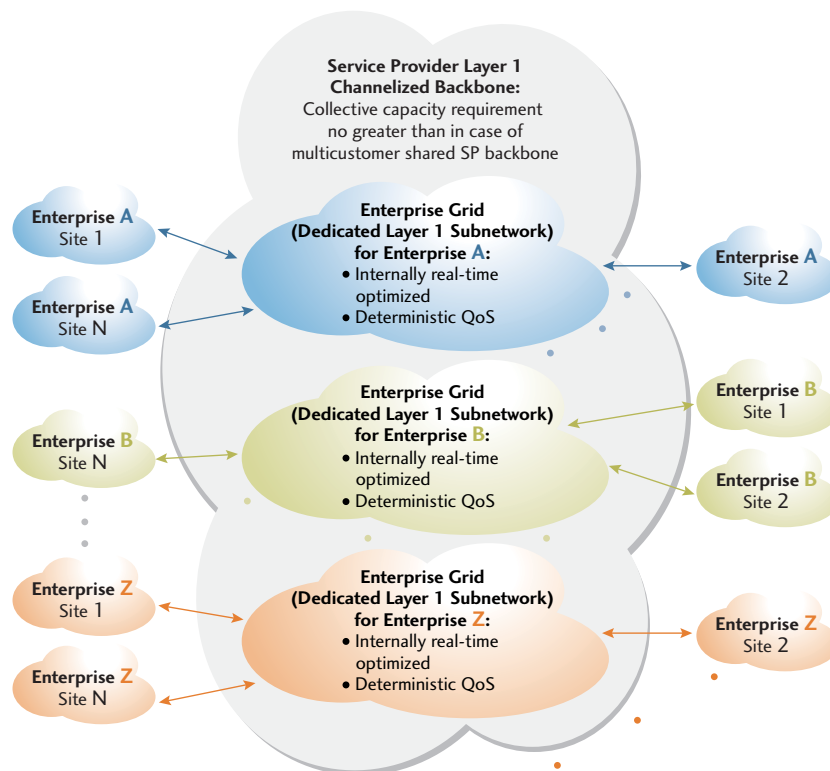
Because every instance of packet switching (i.e., packet hop) along a route of a packet across the network is point of potential congestion, each packet hop increases delay, jitter and packet discard probability. Therefore, an increased number of packet hops will exponentially increase the probability that a given packet will not reach its destination, either at all or not with in its allowed time budget. Because such packets that did not reach their destination (on time) nevertheless do consume network bandwidth, the potentially high level of network utilization with conventional packet-switched networks does not translate directly into high level of on-time data throughput; in fact, the opposite may be the case. Therefore, the conventional packet-switching networks can only support relatively low levels of sustained traffic

Exhibit 7

Enterprise Grids Achieve Optimized Network Cost-Efficiency without Compromising QoS

Source: Optimum Communications Services and Yankee Group, 2008

WANs Implemented Using Enterprise Grids



loads, especially when they are used to support services that do not tolerate high delays, jitter or packet loss rates. In any case, the linear increase in the number of packet hops with conventional packet-switching WAN technologies with the increase in the number of network nodes will exponentially limit the effective data throughputs of such networks. Accordingly, the minimization in packet hops achieved by Adaptive-Mesh architecture of Enterprise Grid causes exponentially increasing QoS and on-time traffic throughput differences between Adaptive-Mesh and conventional packet-switching networks, to the benefit of Adaptive-Mesh.

Ultimately, the Enterprise Grid solution can act as a much-needed game-changer in the enterprise networking arena: Its real-time self-optimizing Adaptive-Mesh architecture makes the dedicated L1 capacity based, highest performance network connectivity at least as economical as any other WAN alternative. The new architecture will enable the industry to break free from the trade-off between desirable qualities and affordability of WAN services that exists between the conventional WAN architectures.

Advantages

- Enterprise Grid eliminates the requirement for monitoring and predicting traffic loads as the self-optimizing Adaptive-Mesh architecture of Enterprise Grid adapts automatically to traffic load variations. Moreover, because with an Enterprise Grid WAN the customer's traffic does not go through any packet switches or routers shared with any other customers, the L2/3 network administration is also greatly simplified compared with traditional packet-layer WANs/VPNs. These aspects of Enterprise Grids allow releasing IT staff from network operations routines, enabling the enterprise IT departments to focus on more value-added activities.
- There is no sharing of packet switches or L1 connections among different customers within an Enterprise Grid WAN. This provides SLAs supported by dedicated L1 capacity based QoS guarantees (such as for inter-site minimum/available throughput, maximum latency and jitter guarantees).
- Enterprise Grid offers high levels of security because of traffic separation at L1, meeting the needs of customer organizations with strong data security requirements.
- As discussed previously, the Adaptive-Mesh network architecture of Enterprise Grids enables:
 - Up to 20-fold optical wavelength capacity utilization efficiency gain versus alternative, equal-performance network technologies
 - In the order of 10-to-1 reduction packet processing capacity requirements per node compared with conventional packet-switching network technologies

- Architectural minimization of packet-hop counts, increasing QoS and on-time data throughput differences between the Enterprise Grid and conventional packet-switching WANs (in favor of Enterprise Grid)

Disadvantages

- OCS' Enterprise Grid is offered only at access rates of OC-12c/STM-4 and higher (including GbE, OC-48c/STM-16 and 10GbE/OC-192c/STM-64).
- Requires availability of fiber-optic network infrastructure.
- Some of the advantages of Enterprise Grid solutions rely on patented or patent-pending intellectual property (IP), allowing the owner of IP (presently OCS) to arbitrate a greater profit margin on the Enterprise Grid services compared to the profit margins taken by providers for traditional WAN services. In other words, the patent-protected nature of the technologies enabling the cost-efficiency gain of Enterprise Grid allows the owner of the enabling IP to retain a potentially significant portion of the savings created by Enterprise Grid versus the least expensive (equal quality) alternative WAN implementation.

Characteristics of Appropriate Enterprise Adopters

- A minimum of five corporate sites or network access points with higher than 100 Mbps of network services capacity requirements.
- Requiring the performance, security and reliability of dedicated L1 capacity based solution, but not willing to invest in and to deploy and operate its own physical wide-area network.
- Looking for a service provider managed WAN.
- Need to improve cost-efficiency of IT spending, especially in cases requiring the enterprise to lower capital expenditure with WANs, while increasing the WAN throughput and performance.
- Requirement for strict data security and off-site storage, e.g., for business continuity, competitive or regulatory reasons such as enterprises in the high-tech, healthcare and financial verticals.

IV. Conclusions and Recommendations

Traditional enterprise WAN architectures have fallen into one of two mature technologies: either dedicated L1 capacity-based WANs (leased lines and dedicated packet switches/routers), or packet-layer-shared WANs L2/3 VPLS/VPNs), each providing their own, and non-overlapping, advantages and disadvantages. But since the emergence of these architectures, an increasing number of enterprises would need from their WAN solution the desirable aspects of each of the traditional

WAN architecture, without their associated drawbacks. As a result, the time is ripe for architectural innovation that breaks free from the trade-offs built in to the traditional WANs, enabling enterprises to get affordable, flexible network connectivity with high performance, reliability and security suitable for the full mix of prevailing and future enterprise communications and IT applications.

Summary Conclusions

- **Game-changing innovations void traditional assumptions regarding enterprise networking.** Several traditionally held assumptions, such as that dedicated-capacity-based high-quality WAN solutions would unavoidably be more expensive than shared capacity based WANs, are no longer valid because of the emerging innovations regarding the underlying WAN technologies. For instance, adaptive bandwidth L1 techniques are making the use of high-performance, per-customer-dedicated L1 capacity more economical than the conventional sharing of network capacity at packet level among multiple customers.
- **Relying on old assumptions for network decision-making in the face of market-transforming innovations carries risks that can be greater than that of adopting innovations.** Ignoring the opportunities enabled by innovation exposes service providers and enterprises to the risk of being stuck with technological standards and supplier relationships that are substandard both in terms of performance and cost-efficiency compared to the level of efficiency achieved by the competing enterprises that embrace new architectural and business models enabled by the game-changing innovations.
- **Seek architectural innovations that enable “smart simplification” of enterprise networking.** Successful innovation does not need to involve any additional complexity for the user. Instead, the best innovations in network technologies allow at the same time to simplify and streamline network implementation and operation while providing multiple times more bandwidth per dollar spent.
- **Architectural innovations open a way to reach a new standard of cost-efficiency and performance of enterprise networking.** Innovations with the underlying technologies of WAN solutions are enabling the industry to break free from the quality versus affordability trade-offs inherent to the traditional WAN technologies. For instance, the emerging adaptive bandwidth physical layer network connectivity-based WANs can at the same time deliver a performance and cost-efficiency that is superior both to the traditional dedicated as well as shared capacity based WANs.

Recommendations for Service Providers and Systems Integrators

- **Look at more innovative ways to meet the technology and communications needs of enterprise customers.** Actively pursue partnerships with upstream technology suppliers or specialized applications service providers with innovative technology-based solutions, to be integrated and resold as complete business IT and communications solutions package for enterprises.
- **Adopting innovations that enable rapidly scaling up revenue-generating service delivery capacity while simplifying the delivery infrastructure is key to improving competitiveness and profitability.** Current business communications market economics are hardly sustainable. Enterprises need multiple times more capacity per dollar, as well as strict QoS guarantees, to support their emerging real-time dynamic, multi-location collaboration, etc., high-bandwidth, high-QoS communications and IT applications. The provider that is fastest to deliver the required next-generation enterprise network solution at new standard of performance and efficiency at affordable price stands to win plenty of market share.

Recommendations for Enterprises

- **Determine the best management option for your WAN environment.** Service provider managed WAN solutions may provide valuable time and productivity benefits, and enable you to take advantage of more advanced architectures without having to worry about having the right in-house expertise.
- **Negotiate with your service providers for contracts that allow for easy migration and expansion.** The needs of your business will continue to shift and evolve, requiring an adaptable technology platform. Having flexible architecture and service provider relationships will enable faster and easier adaptation to changing business needs.
- **Press your service provider or integrator to deliver you a network solution that addresses each of your actual business needs, without cost-efficiency and flexibility being sacrificed for security and performance, or vice versa.** There are innovations being developed to address your actual network service requirements, but the service providers and integrators need active and clear customer input for them to justify the business case of breaking the status quo and delivering the innovations to the enterprise customers.

Yankee Group

Yankee Group has research and sales staff located in North America, Europe, the Middle East, Africa, Latin America and Asia-Pacific. For more information, please contact one of the sales offices listed below.

Corporate Headquarters

Prudential Tower
800 Boylston Street
27th Floor
BOSTON, MASSACHUSETTS 02199
617-598-7200 phone
617-598-7400 fax
info@yankeegroup.com

Europe

56 Russell Square
LONDON WC1B 4HP
UNITED KINGDOM
44-20-7307-1050 phone
44-20-7323-3747 fax
euroinfo@yankeegroup.com

Yankee Group | the global connectivity experts™

A global connectivity revolution is under way, transforming the way that businesses and consumers interact beyond anything we have experienced to date. The stakes are high, and there are new needs to be met while power shifts among traditional and new market entrants. Advice about technology change is everywhere—in the clamor of the media, the boardroom approaches of management consultants and the technology research community. Among these sources, Yankee Group stands out as the original and most respected source of deep insight and counsel for the builders, operators and users of connectivity solutions.

For 37 years, we have conducted primary research on the fundamental questions that chart the pace and nature of technology changes on networks, consumers and enterprises. Coupling professional expertise in communications development and deployment with hundreds of interviews and tens of thousands of data points each year, we provide qualitative and quantitative information to our clients in an insightful, timely, flexible and economic offering.

Yankee Group Link

As technology connects more people, places and things, players must confront challenging questions to benefit from the changes: which technologies, what economic models, which partners and what offerings? Yankee Group Link™ is the research membership uniquely positioned to bring you the focus, the depth, the history and the flexibility you need to answer these questions.

Yankee Group Link membership connects you to our qualitative analysis of the technologies, services and industries we assess in our research agenda charting global connectivity change. It also connects you to unique quantitative data from the dozens of annual surveys we conduct with thousands of enterprises and consumers, along with market adoption data, comprehensive forecasts and global regulatory dashboards.

Yankee Group Link Research

As a Link member, you have access to more than 500 research reports and notes that Yankee Group publishes each year. Link Research examines current business issues with a unique combination of knowledge and services. We explore topics in an easy-to-read, solutions-oriented format. With the combination of market-driven research and built-in direct access to Yankee Group analysts, you benefit from the interpretation and application of our research to your individual business requirements.

Yankee Group Link Interaction

Our analysts are at your further disposal with data, information or advice on a particular topic at the core of a Link membership. We encourage you to have direct interaction with analysts through ongoing conversations, conference calls and briefings.

Yankee Group Link Data

Yankee Group Link Data modules provide a comprehensive, quantitative perspective of global connectivity markets, technologies and the competitive landscape. Together with Link Research, data modules connect you to the information you need to make the most informed strategic and tactical business decisions.

Yankee Group Consulting

Who better than Yankee Group to help you define key global connectivity strategies, scope major technology initiatives and determine your organization's readiness to undertake them, differentiate yourself competitively or guide initiatives around connectivity change? Our analysts apply Yankee Group research, methodologies, critical thinking and survey results to your specific needs to produce expert, timely, custom results.

Yankee Group Signature Events

Yankee Group conferences, webinars and speaking engagements offer our clients new insight, knowledge and expertise to better understand and overcome the obstacles to succeed in this connectivity revolution.

www.yankeegroup.com

The people of Yankee Group are the global connectivity experts™—the leading source of insight and counsel for builders, operators and users of connectivity solutions. For more than 35 years, Yankee Group has conducted primary research that charts the pace of technology change and its effect on networks, consumers and enterprises. Headquartered in Boston, Yankee Group has a global presence including operations in North America, Europe, the Middle East, Africa, Latin America and Asia-Pacific.