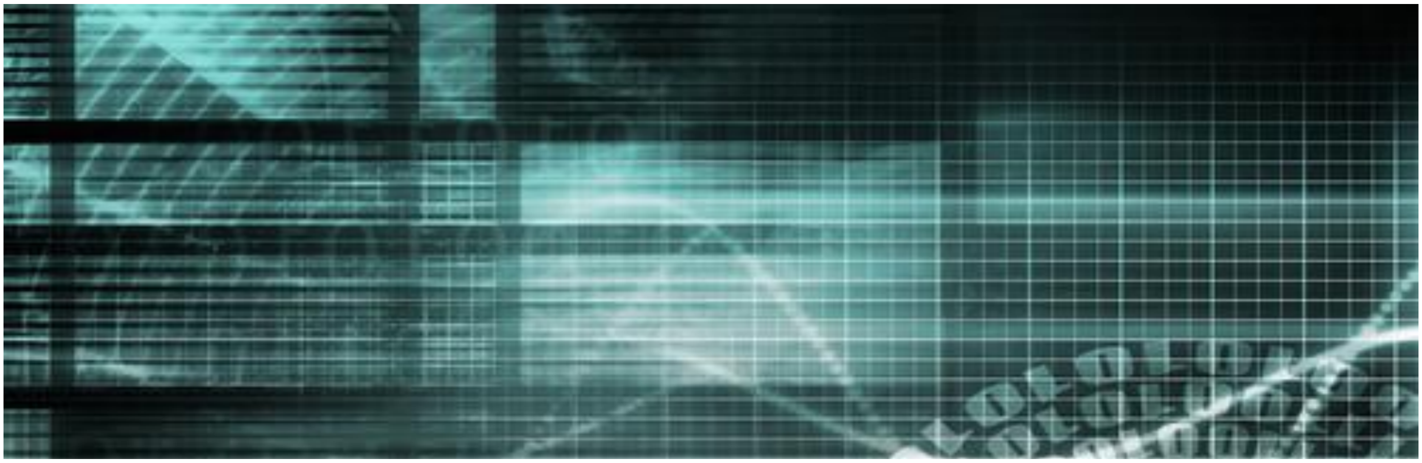

Executive Summary

Traffic requirements are growing rapidly because of the widespread acceptance of online video services, cloud computing, and mobile broadband. Wide Area Networking (WAN) costs also are rising with traffic growth in part because of suboptimal network utilization efficiency. At the same time service creation processes are lengthy and service providers' responses to competitive threats such as over-the-top (OTT) video and cloud-based services have been sluggish, resulting in slow revenue growth. A root cause of rising WAN costs and slow revenue growth is poor WAN management information flows and many manual work steps.

Software defined networking (SDN) in the WAN offers the opportunity to drive down costs through increased operational efficiency, increased service creation velocity, and differentiated and personalized network services. Cisco WAN Automation Engine (WAE) implements SDN in the WAN for service providers' networks. It is a platform that provides real-time visibility, analysis and control across multivendor network infrastructure and services. Real-time software abstraction of the network allows applications to gain visibility and control of the network through web programming techniques rather than through device-specific embedded programming techniques. This strengthens service providers' network control and lessens their dependence on systems vendors.

ACG Research analyzes four use case examples to demonstrate the financial benefits of SDN in the WAN as implemented by WAE to service providers. Two dynamic bandwidth management examples identify opportunities for service providers to sell occasional, high-bandwidth services to enterprises and quantify financial metrics, including net present value (NPV), gross profits, and return on investment (ROI). Two traffic engineering examples describe opportunities to apply automated traffic load management to increase service velocity, reduce operating expenses, and improve network utilization.

KEY FINDINGS

Cisco WAE implements SDN in the WAN for service providers' networks. Four use case examples found:

- 94% ROI in Year 1 for a bandwidth calendaring service offer
- 54% five-year revenue CAGR for a bandwidth on-demand service offer
- 45% TCO savings for a TE tunnel balancing solution for transoceanic transport
- 48% TCO savings for a TE multilayer traffic rearranging solution for a Tier 1 national network

Introduction

Traffic requirements are growing rapidly because of the enthusiastic acceptance of online video services, cloud computing, and mobile broadband. Wide Area Network (WAN) costs, in turn, are rising with traffic growth because of suboptimal network utilization efficiency. Capacity planning and network deployment processes are lengthy and disjointed—a major contributor to low efficiency. The root cause is poor information flows among processes and many manual work steps. The lack of a global view of traffic, topology, costs, and policy continues to create complexity as interaction with the network requires detailed knowledge of networking protocols and an awareness of each vendor's implementation.

Profitability and business viability are further threatened by slow revenue growth. As is the case with WAN costs poor information flows and manual steps are impairing service providers' abilities to respond to rapidly changing competitive threats (for example, over-the-top [OTT] video, cloud-based services) and to accelerate service innovation and deployment.

Software defined networking (SDN) in the WAN offers the opportunity to drive down costs through increased operational efficiency—network utilization in particular—increased service creation velocity, and differentiated and personalized network services. Specifically, SDN and automation are used in a hybrid centralized/decentralized control and automation solution to improve WAN utilization and create and deliver profitable new services.

Cisco WAN Automation Engine (WAE) implements SDN in the WAN for service providers' networks. Four use case examples are described and their benefits are calculated:

- Dynamic Bandwidth Management Use Cases
  - Bandwidth Calendaring
  - Bandwidth on Demand

- Traffic Engineering (TE) Use Cases
  - TE Tunnel Balancer
  - TE Multilayer Traffic Rearranger

WAE in Cisco Evolved Services Platform

The Cisco Evolved Service Platform is a comprehensive software engine that makes network services and applications easy to consume. It is composed of three functions: 1) Orchestration Engine, 2) Catalog of Virtual Functions, and 3) Service Broker. WAE provides the orchestration engine function for the WAN. It is a multivendor platform that hides implementation complexity and provides visibility, analysis, and control across the network infrastructure and services. WAE extends the Path Computation Element (PCE) architecture (See IETF RFC 5440) by adding comprehensive historical, real-time, and predictive models to the traffic engineering decision making process, providing insight into the historical and present network, as well as the forecasted future network after new demands or applications are added and failure scenarios considered.

WAE constructs a near real-time model of the network and exposes the network as a set of abstractions accessible via a set of RESTful APIs. This allows programmers to interact with the network by simply considering services, locations, and demands. There is no requirement to understand devices, operating systems (OS), and protocols. The complexity of the underlying network is hidden while simple abstractions are exposed and leveraged using web programming techniques.

Figure 1 shows the WAE architecture and workflow:

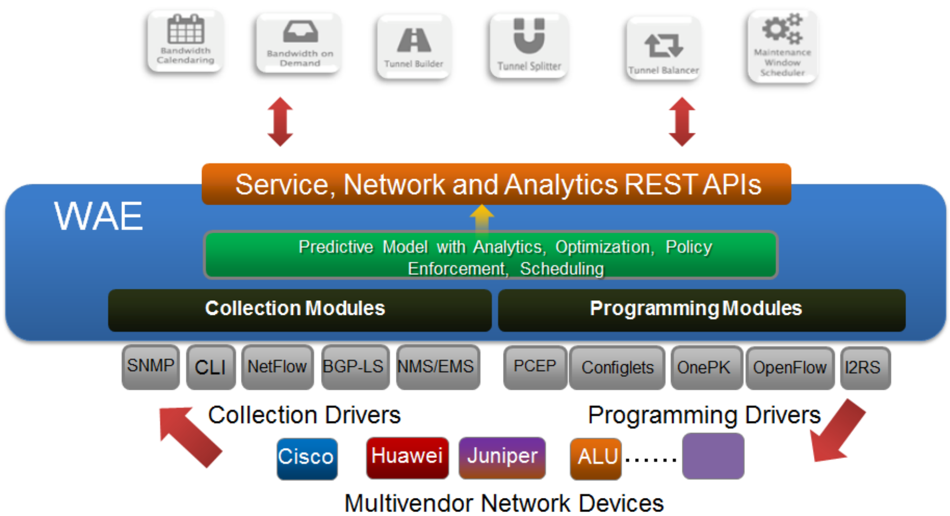

Figure 1 – WAE Architecture and Workflow

WAE has a modular, open design and implements integrated functions, including collection, analytic processing, optimization, and prediction, scheduling and deployment. Collection mechanisms provide visibility into the network infrastructure. They use standards-based approaches that support SNMP, PCEP, NetFlow, BGP-LS, and CLI. The analytic processing capability analyzes in near real-time thousands of network topology models and scenarios to consider new demands and applications or to generate customized reports for traffic trends. Additionally, WAE performs global or tactical IGP or LSP optimization to run networks at higher levels of sustained utilization and resiliency. A built-in scheduler reserves bandwidth and schedules configuration for future consumption. The deployer configures and modifies Label Switched Paths (LSP) using a set of flexible programming options, including PCEP, Configlets, NETCONF/YANG, OpenFlow, and I2RS, thus offering complete device-agnostic network control. For more than 10 years, these algorithms have been field tested and validated in MATE design (See page 9 side bar).

Dynamic Bandwidth Management Use Cases

There is an unfulfilled demand to provide enterprises with occasional high-bandwidth services such as Hybrid Cloud backup and storage, data center to data center replication, disaster recovery or workload migration. The demand is unfulfilled because service providers' networks are often provisioned to peak traffic loads because of an inability to time shift or space shift traffic. For most service providers, the average network utilization rate leaves approximately 50 percent of capacity unused. This unused capacity is money left on the table because of the inability to offer unused capacity dynamically to currently unreached market segments as well as the existing customer base. Unused capacity can be offered using scavenger or reserve classes of service. These classes of service provide options such as where to present demand: load placement; when to present the demand: calendaring; and how to route the demand: policy. In addition, service providers' business models require that bandwidth is dedicated to the enterprise for a lengthy period (typically three years) to cover the cost of deploying the bandwidth. It is difficult, therefore, for enterprises to make a long-term expense commitment for occasional high- bandwidth usage, resulting in enterprises suffering network service degradation during these occasional peak service usage periods.

WAE makes it operationally and economically feasible for the service provider to offer an attractive price for short-term bandwidth additions. WAE is able to find and deliver unused network capacity quickly by applying its extended PCE solution through the northbound/southbound APIs that link business logic to WAE and through the abstraction, virtualization, and automation capabilities of WAE.

Figure 1 shows how WAE is used to find and fulfill a request for occasional bandwidth.

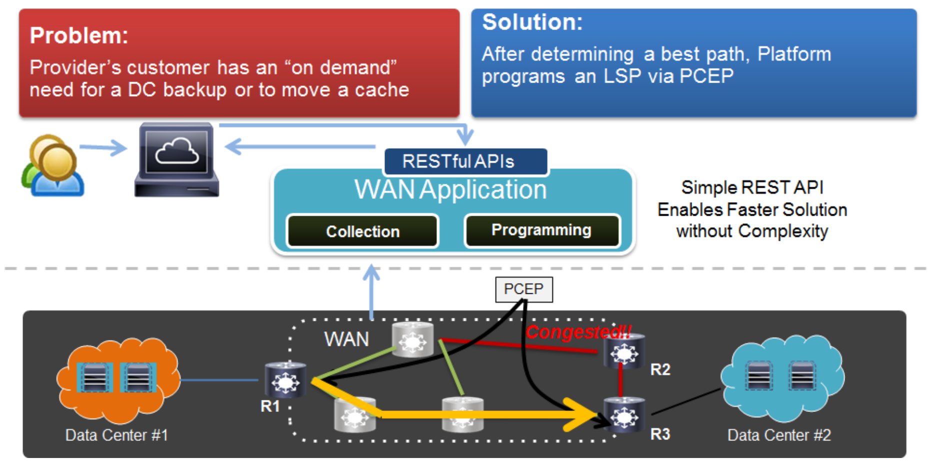

Figure 2 – WAE Dynamic Bandwidth Management

The service provider issues a request for occasional bandwidth that is passed to the WAE platform via a REST API. The WAE Collector gathers information on the current network state that is then passed to the WAE analytics engine. Traffic engineering is performed via PCE to determine how to optimally fulfill the bandwidth request. The deployer then defines an LSP, selects a path subject to policy, takes into account failures and other reserved bandwidth, and communicates to the affected network elements.

Bandwidth Calendaring Use Case

A North America service provider has identified an opportunity to sell periodic dedicated bandwidth for data center to data center data replication to a user base of 15,000 small to medium businesses and 9,000 enterprises. The service is a bulk data transfer with guaranteed bandwidth and quality of experience (QoE) and uses an "as a service" delivery model. The customer is billed on a usage basis, which is much more affordable than the price of permanently nailed-up bandwidth (61 percent savings in this use case). This provides an affordable solution that eliminates occasional network performance degradations. Examples of such applications are cloud backup and disaster recovery.

This capability to put up and tear down bandwidth rapidly gives the service provider a new revenue opportunity to attract small and medium businesses with limited MPLS VPN budgets, improves the service provider's network utilization, and provides an offering with an incremental cost that is low compared to the incremental revenue.

Inputs and assumptions include:

- 3% service adoption in Year 1
- 7% service adoption in Year 3
- 50 hours of data replication per enterprise per month
- 6 Terabytes of data transferred per month in Year 1
- 250 Mbps data transfer rate
- 25% data center traffic five-year CAGR

Figure 3 shows the revenue projection.

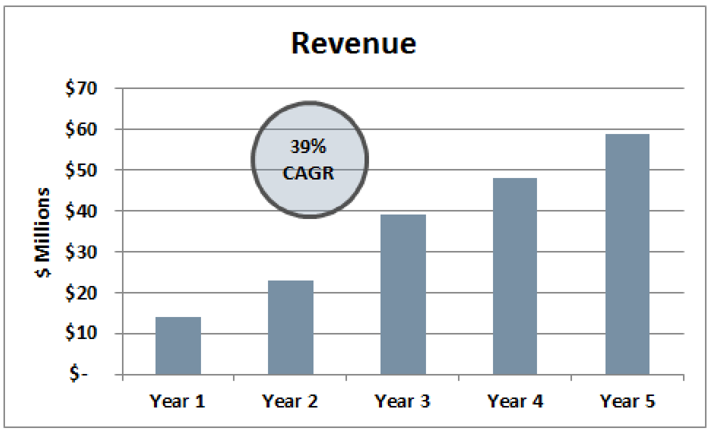

Figure 3 – Revenue Projection for Bandwidth Calendaring

Bandwidth calendaring revenue has a 39 percent CAGR, net present value of $38 million for five years and a return of investment (ROI) of 94 percent in the first year.

Bandwidth on Demand Use Case

A North America service provider has identified the opportunity to sell automated (through a web portal) bandwidth provisioning for on-demand over Committed Information Rate bandwidth capacity to its 9,000 enterprise customers. The bandwidth could be used, for example, for a disaster recovery or workload migration use case with VMware vMotion bandwidth requirements. The service is delivered with guaranteed bandwidth and QoE, using a "as a service" delivery model. It provides the enterprise an affordable alternative to incurring network performance degradation during peak usage periods. The service provider adds new revenue, improves its network utilization, and penetrates additional small and medium business accounts.

Inputs and assumptions include:

- Enterprise saves 87% compared to the cost of buying additional Carrier Ethernet VPN service with a three-year contract
- 2% service adoption in Year 1
- 7% service adoption in Year 5

- 20 hours of disaster recovery per enterprise per month
- 1 Gbps data transfer rate
- 25% data center traffic five-year CAGR

Figure 4 shows the revenue projection.

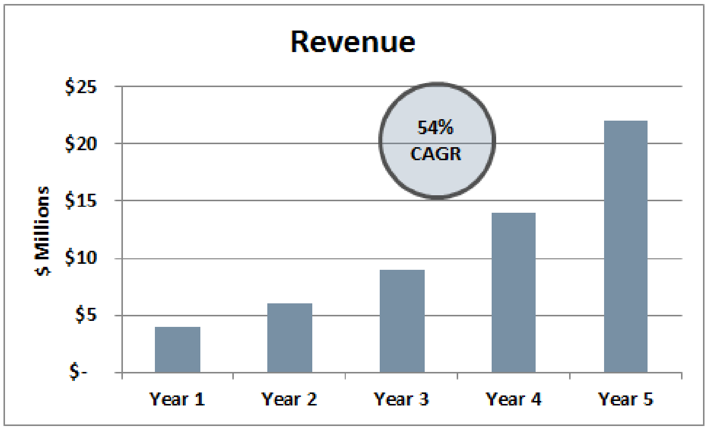

Figure 4 – Gross Profit Projection for Bandwidth on Demand

Bandwidth on demand has revenue of 54 percent CAGR for five years with a net present value of $39 million for five years and a ROI of 55 percent in the first year.

Traffic Engineering Use Cases

Automated traffic load management increases service velocity, reduces operating expenses, and improves network utilization by accelerating planning and traffic engineering processes using WAE and by using predictive modeling techniques incorporated into the Cisco MATE software portfolio.

Figure 5 illustrates the use of WAE to automatically do policy-based path planning; the sidebar describes the predictive modeling capabilities of MATE.

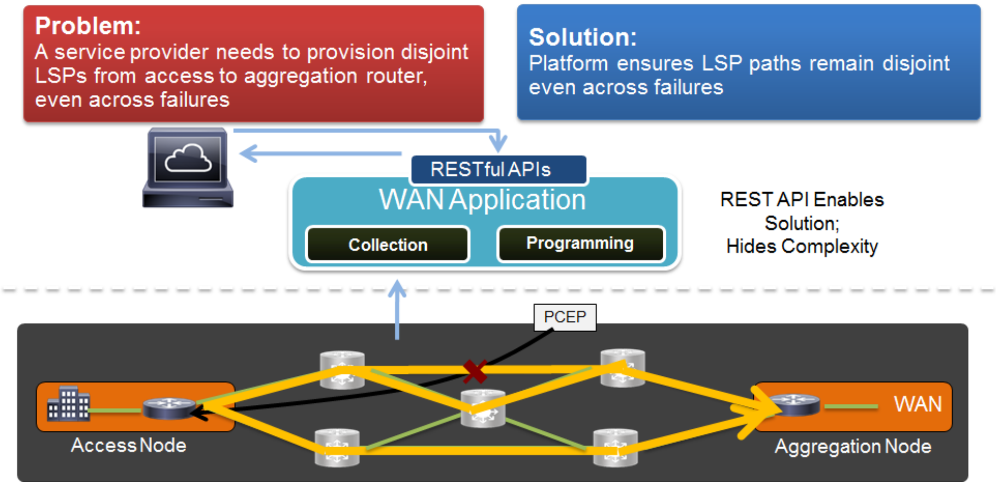

Figure 5 – WAE Policy-Based Path Planning

The service provider defines the requirements to provision disjoint LSPs between nodes, for example, between access to aggregation, edge to edge, and access to access by using a high-level TE Manager application. The requirement is communicated to the WAE platform via the REST APIs. The WAE Collector gathers information on the current network state that is then used by the analytics engine and the MATE predictive modeling software to find the most efficient disjoint LSP paths through the network. The deployer then defines an LSP, which it communicates to the affected network elements.

TE Tunnel Balancer Use Case

A South America ISP uses transoceanic network links between Miami and Rio in Brazil. It seeks to optimize its network utilization on these costly transoceanic links by automating traffic load management so that that traffic capacity allocations respond rapidly to changing traffic patterns and fiber path failures.

Comparison inputs and assumptions include:

- 220 Gbps aggregate connectivity pool
- 22 10 G links
- 4 Fiber paths
- 28% traffic CAGR
- 8 routers deployed in Year 1
- 3% CAGR PMO router deployments
- 28% CAGR PMO router 10GE port deployments
- PMO link utilization 40%
- $65,000 monthly recurring charge for 10GE transoceanic link

Figure 6 shows the total cost of ownership (TCO) savings.

MATE Portfolio

The Cisco MATE Portfolio provides visualization, analysis and optimization tools to deliver the critical information necessary to design, plan, and operate IP/MPLS networks. Central to the MATE portfolio is a predictive model that simulates and analyzes how traffic traverses the network and the impact that traffic has on the infrastructure. For more than 10 years, the MATE portfolio has been validated in large-scale deployments.

The MATE portfolio consists of three software products:

- MATE Collector software gathers and maintains information on infrastructure elements, topology, operational states, and traffic statistics for network planning and analytics.
- MATE Design software provides an integrated system for designing, engineering, and planning IP/MPLS networks.
- MATE Live software delivers (in near real-time) in-depth network analytics for both current (live) and historical data that is used to make business and technical decisions.

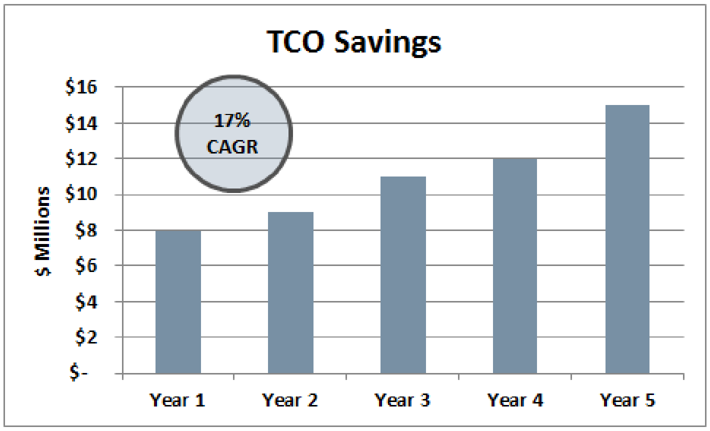

Figure 6 – TCO Savings for TE Tunnel Balancing

The one-year ROI is 258 percent. This shows very rapid recovery of the costs of the WAE and tunnel balancer application software licenses. TCO savings has a five-year CAGR of 17 percent; five-year net present value is $41 million. The WAE solution increases link utilization from 40 percent to 71 percent in a protect for failure mode where a fiber path failure has no negative impact on links as the links are automatically redistributed across other fiber paths with no change in service quality or risk. Also, Mean Time to Repair is reduced from days to minutes. Five-year TCO is reduced by 45 percent as compared to the PMO. Ninety-seven percent of this savings is attributable to the reduced cost of the transoceanic network links.

<u>Impact of Evolving From MATE to WAE</u>

The MATE software portfolio produces 32 percent of the TCO savings; the automation capabilities of WAE contribute 13 percent for a total TCO savings of 45 percent as compared to the PMO. The evolution to WAE enables the service provider to capture an additional 41 percent savings uplift in the first year because of automation. MATE permits increased utilization of the transoceanic links by providing predictive models that foresee the impact of various traffic and outage scenarios on the network infrastructure and producing optimal traffic engineering solutions that allow link utilization to be increased without compromising network performance or assuming additional risks of failure. However, MATE does not automate the capacity deployment process.

In this use case WAE incorporates the MATE predictive models and design processes and couples them to the WAE automated deployment process. The automation capabilities of WAE, therefore, shorten the time to deploy capacity as well as Mean Time to Repair. This increased service velocity permits further increases in the link utilization levels with no change in the risk of failure. This is the source of the additional 12 percent contribution of WAE to the overall TCO savings.

Multilayer Traffic Rearranger Use Case

A Tier 1 service provider seeks to optimize capacity utilization, improve service velocity, and reduce network operations costs by automating capacity planning for the packet and transport layers on its national network. The service provider's existing protection scheme is compared to a WAE protection scheme. The existing protection scheme has no visibility across the optical and IP layers. This results in inefficient provisioning of optical wavelengths, transponders and ports to support all failure types. The existing protection scheme uses pure IP restoration and leads to high TCO compared to the WAE solution. WAE provides multilayer traffic rearranging. WAE eliminates overprovisioning by providing visibility across network layers and, thus, is able to exploit available resources after a port or optical failure occurs. With the ability to rearrange traffic across functional resources in an agile manner and optimize IP topology, WAE enables significant cost savings as less spare capacity needs to be provisioned in the network.

Inputs and assumptions used in the comparison include:

- 10,715 Gbps aggregate connectivity pool in Year 1
- 75 routers deployed at study initiation
- 3% CAGR PMO router deployments
- 28% CAGR PMO router 10GE port deployments
- 50% PMO link utilization
- Blended TCO per 10G Port, $4,283
- Lease price per 10G of transport/mile/month, $5.85

Figure 7 shows the TCO savings.

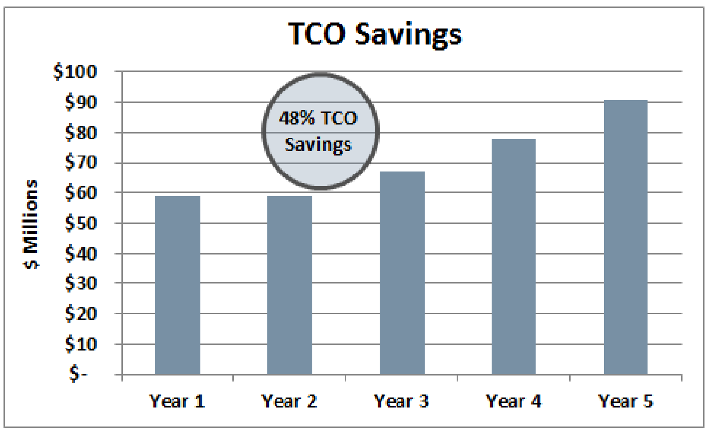

Figure 7 – TCO Savings for TE Multilayer Traffic Rearranger Use Case

The investment in the WAE solution is rapidly recovered with one-year ROI of 902 percent with a five-year net present value of $280 million in TCO savings. Transport cost reductions account for 62 percent of the savings; network operation expense reductions contribute 33 percent of the savings; the remainder is due to network capital expense savings. The automation and near real-time optimization capabilities of WAE enable the increased utilization by:

- Reducing the capacity planning cycle from months to days, which reduces traffic capacity forecasting errors
- Automated and integrated planning of the packet and transport layers, which simplifies and accelerates the planning process and eliminates a bias toward over forecasting capacity requirements
- More accurate identification of failure impacts that limit the need to provision extra capacity to protect against planning uncertainty

Conclusion

WAN costs are tracking traffic growth in part because of suboptimal network efficiency utilization. One inefficiency source is service providers' overdependence on systems vendors for controlling the behavior and operating expenses of their networks. At the same time revenue growth is slow as are service providers' responses to competitive threats. This is because service creation, capacity planning, and network deployment processes are lengthy and disjointed. The root cause is poor information flows among processes and many manual work steps because of a lack of simple network to application interfaces.

SDN in the WAN offers the opportunity to drive down costs through increased operational efficiency—network utilization in particular—service creation velocity, and differentiated and personalized network services. Cisco WAE implements SDN in the WAN for service providers' networks.

Two dynamic bandwidth management and two traffic engineering use cases are described and analyzed to quantify the financial contributions of WAE to service providers.

- Bandwidth calendaring: Selling periodic dedicated bandwidth for data center to data center data replication
- Bandwidth on demand: Selling automated bandwidth provisioning for on-demand over Committed Information Rate bandwidth to enterprises
- TE tunnel balancer: Automating traffic load management on transoceanic transmission links
- TE multilayer traffic rearranger: Automating capacity planning for multilayer optimization of the packet and transport layers on a national network

Table 1 summarizes the financial benefits of the dynamic bandwidth management use cases as compared to the PMO.

Bandwidth Management Use Case	NPV ($ Millions)	ROI 1 Year	Five-Year Revenue CAGR
Bandwidth Calendaring	$38	94%	39%
Bandwidth on Demand	$39	55%	54%

Table 1 – Bandwidth Management Use Case Results

Table 2 summarizes the financial benefits of the traffic engineering use cases as compared to the PMO.

Traffic Engineering Use Case	NPV ($ Millions)	ROI 1 Year	Five-Year TCO Savings
TE Tunnel Balancing	$41	258%	45%
TE Multilayer Traffic Rearranger	$280	902%	48%

Table 2 – Traffic Engineering Use Case Results