WAN Optimization with Cisco's WAAS

May 2007

Introduction

Choosing the best WAN optimization solution for branch office consolidation and acceleration can be challenging due to the wide range of solutions that are currently available from more than a dozen vendors. In the course of conducting a broad market survey of enterprise requirements for WAN optimization and application acceleration[1], Ashton, Metzler & Associates (AM&A) has developed a set of evaluation criteria that can assist IT organizations in the selection of WAN optimization solutions. For example, the criteria could be used as the basis for an initial screening of vendor solutions to arrive at a short list of vendors whose solutions would be examined more carefully before a purchase decision is made. Another possible use of the criteria would be as an aid in developing a more detailed RFP to issue to vendors on the short list.

The remainder of this document is focused on a brief discussion of the selection criteria together with an evaluation of how Cisco's WAN optimization solution based on Wide Area Application Services (WAAS V4.0) meets each of these criteria. Therefore, this document provides an example of the sort of preliminary analysis that would typically be done by IT organizations in order to narrow the field of vendors to a short list of potential solutions.

The evaluation of the Cisco solution is based primarily on an examination of various documents posted on the Cisco web site, including the Cisco Wide Area Application Services Configuration Guide (Software Version 4.0.1). Throughout this document, the Cisco WAN optimization solution will be referred to as either WAAS or Cisco's WAAS.

Evaluation Criteria

The recommended criteria to evaluate WAN optimization solutions are listed in Table 1. It should be noted that this list is intended as a fairly complete compilation of all possible criteria, so a given organization might apply only a subset of these criteria for a given purchase decision. In addition, individual organizations would be expected to ascribe different weights to each of the criteria because of differences in WAN architecture, branch office network design, and application mix. As shown in the table, assigning weights to the criteria and relative scores for each solution provides a simple methodology for comparing competing solutions.

There are many techniques that IT organizations can use to complete Table 1 and then use the contents of Table 1 to compare alternative solutions. For example, the weights can range from 10 points to 50 points, with 10 points meaning not important, 30 points meaning average importance, and 50 points meaning critically important. The score for each criteria can range from 1 to 5, with a 1 meaning fails to meet minimum needs, 3 meaning acceptable, and 5 meaning significantly exceeds requirements.

For the sake of example, consider solution A. For this solution, the weighted score for each criterion (WiAi) is found by multiplying the weight (Wi) of each criteria, by the score of each criteria (Ai). The weighted score for each criterion are then summed (Σ WiAi) to get the total score for the solution. This process can then be repeated for additional solutions and the total scores of the solutions can be compared.

[1] Application Delivery Handbook, www.kubernan.com

Criterion	Weight Wi	Score for Solution "A" Ai	Score for Solution "B" Bi
Performance			
Transparency			
Solution Architecture			
OSI Layer			
Capability to Perform Application Monitoring			
Scalability			
Cost-Effectiveness			
Application Sub-classification			
Module vs. Application Optimization			
Disk vs. RAM-based Compression			
Protocol Support			
Security			
Ease of Deployment and Management			
Change Management			
Support for Meshed Traffic			
Support for Real Time Traffic			
Total Score		Σ WiAi	Σ WiBi

Table 1: Criteria for WAN Optimization Solutions

Cisco's WAAS

As shown in Figure 1, the Cisco solution is comprised of WAAS software running on Wide Area Application Engine (WAE) appliances at both branch and central sites. These appliances can either be deployed as a stand-alone device or as a module in an Integrated Services Router (ISR).

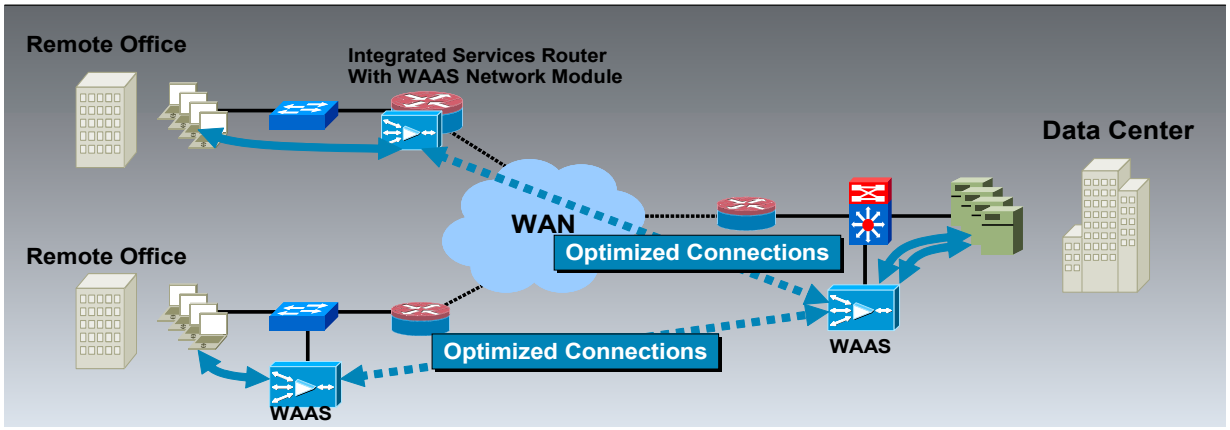

Figure 1: Cisco's WAAS

Evaluating Cisco's WAAS

This section will evaluate Cisco's WAAS based on the criteria presented in Table 1.

Performance

Description of the criteria:

Third party tests of a solution can be helpful. It is critical, however, to quantify the kind of performance gains that the solution will provide in the particular environment where it will be installed. For example, if the IT organization either already has, or is in the process of consolidating servers out of branch offices and into centralized data centers, then it needs to understand how well the WAN optimization solution supports CIFS (Common Internet File System).

As part of this quantification, it is important to identify if the performance degrades as either additional functionality within the solution is activated, or if the solution is deployed more broadly across the organization.

Evaluation of Cisco:

Cisco has sponsored a third party performance evaluation by Miercom. That evaluation included testing for the acceleration of HTTP and FTP, CIFS, Microsoft Office remote file operations, and Microsoft Exchange downloads. The test results show that under favorable conditions (large WAN latency and highly compressible content) the WAAS V4.0 solution can provide impressive degrees of acceleration. For example, the Miercom tests of file transfers using HTTP and FTP showed compression rates of 98% or higher.

When comparing Cisco's WAAS against other solutions, Miercom concluded that Cisco's WAAS software demonstrated performance parity with these other products, and that in some key metrics provided superior compression, speed and throughput. Miercom also concluded that unlike some of the other products in the market, Cisco's WAAS does not degrade under load.

Transparency

Description of the criteria:

To many IT organizations, the first rule of networking is that you should not do anything that causes the network to break. As a result, an important criterion when choosing a WAN optimization solution is that it should be possible to deploy the solution and not have anything such as routing, security or QoS break. The solution should also be transparent relative to both the existing server configurations and the existing Authentication, Authorization and Accounting (AAA) systems. In addition, the solution should not make troubleshooting any more difficult.

As part of testing products such as WAN optimization solutions, IT organizations typically test performance. AM&A recommends that IT organizations also test the transparency of these solutions.

Evaluation of Cisco:

With WAAS V4.0 all packet header information is preserved. As a result, WAAS integrates easily with the existing infrastructure including clients, servers, storage, and networking devices. WAAS is also transparent in the sense that accelerator appliances can use auto-discovery to determine whether a peer accelerator is available at the other end of the link. After auto-discovery, a pair of accelerators can auto-negotiate an acceleration policy to be applied to the application flow. If a peer accelerator is not discovered, the application flow passes-through unchanged.

Cisco is one of only a couple of vendors in this market that preserves the header information. That does not say that it is impossible for WAAS to break something. It also does not say that a solution that changes the header information will always break something. What it does say is that vendors that change header information will have to implement a workaround in order to be as transparent as a solution that does not modify header information.

Solution Architecture

Description of the criteria:

If the organization intends the solution to be able to support additional optimization functionality over time, it is important to determine if the hardware and software architecture can support new functionality without an unacceptable loss of performance.

Evaluation of Cisco:

Cisco's WAAS adheres to a layered architecture (Figure 2) that facilitates the addition of enhanced functionality and maximizes the independence between functions. Whether there is enough processing capacity to adequately service a large number of sessions with multiple simultaneous functions depends primarily on the scalability of the WAE hardware platform that is chosen. As a result of the testing that was done, Miercom concluded that Cisco's WAAS shows fairness of connection throughput and performance consistency across concurrent users and multiple workloads.

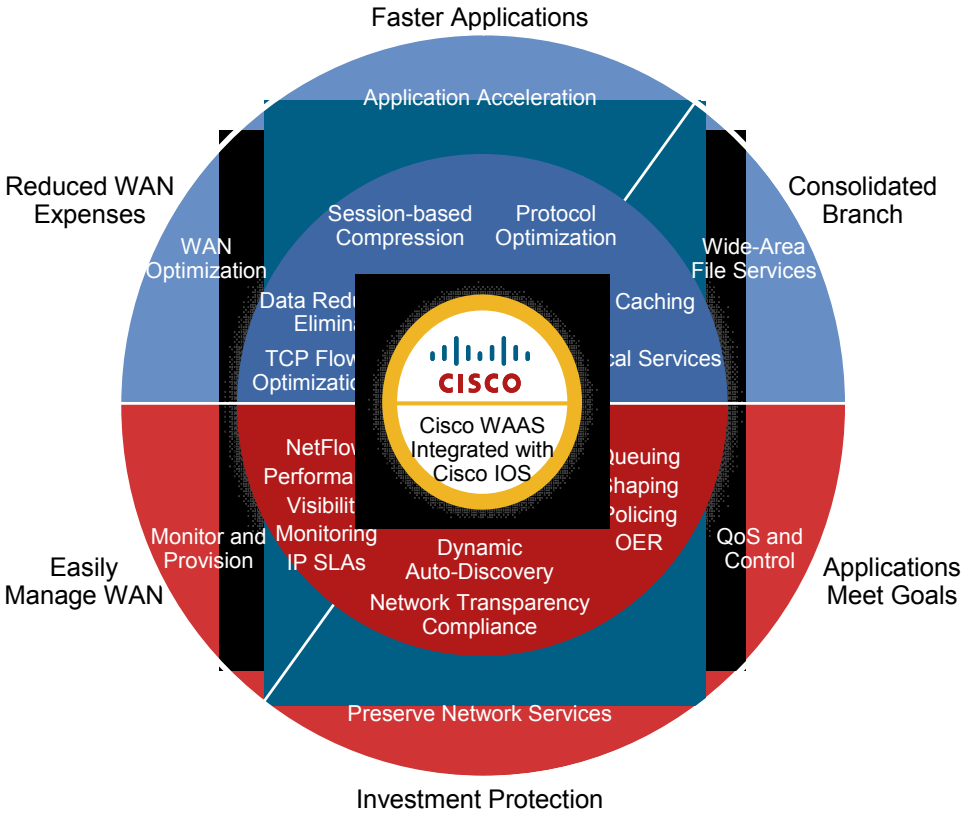

Figure 2: Cisco's WAAS Architecture

OSI Layer

Description of the criteria:

Organizations can apply many of these optimization techniques at various layers of the OSI model. They can apply compression, for example, at the packet layer. The advantage of applying compression at this layer is that it supports all transport protocols and all applications. The disadvantage is that it cannot directly address any issues that occur higher in the stack.

Alternatively, having an understanding of the semantics of the application means that compression can also be applied to the application; e.g., SAP or Oracle. Applying compression, or other techniques such as request prediction, in this manner has the potential to be more effective.

Evaluation of Cisco:

WAAS operates at both the packet and the application layers. At the packet layer, WAAS applies TCP Flow Optimization (TFO), persistent LZ (Lempel-Ziv) Compression, and Data Redundancy Elimination (DRE) to minimize the effect of latency and the amount of bandwidth consumption. At the application layer, WAAS employs application proxies and caches to overcome latency issues and minimize data flows over the WAN.

The discussion of the performance impact of applying optimization techniques at specific layers of the OSI model is somewhat academic. That follows because as mentioned, what is important about performance is how the WAN optimization solution actually performs in the environment in which it will be deployed. That being said, the fact that Cisco supports implementing techniques at multiple layers of the OSI model is positive in that it gives them greater flexibility to implement optimization techniques in whatever manner makes the most sense. This should result in enhanced performance.

Capability to Perform Application Monitoring

Description of the criteria:

Many network performance tools, such as NetQoS's ReporterAnalyzer and SuperAgent, rely on network-based traffic statistics gathered from network infrastructure elements at specific points in the network to perform their reporting. By design, all WAN optimization devices apply various optimization techniques on the application packets and hence affect these network-based traffic statistics to varying degrees. One of the important factors that determine the degree of these effects is based on the amount of the original TCP/IP header information retained in the optimized packets.

Evaluation of Cisco:

Cisco's WAAS retains the original IP header information and hence works seamlessly with all network performance monitoring tools that rely on the original IP header information between a client and server. NetQoS's ReporterAnalyzer, by virtue of its reliance on Netflow data as an input, works seamlessly when WAAS is introduced in the network. The transparent integration of WAAS with NetQoS ReporterAnalyzer allows organizations to maintain their current visibility into WAN link utilization and throughput.

The integration of Cisco's WAAS and NetQoS's ReporterAnalyzer and SuperAgent provide an accurate visibility into end-to-end response time and traffic flow data (via Cisco IOS® NetFlow) even after optimization occurs. This visibility allows enterprises to quantify the application acceleration and WAN optimization improvements offered by Cisco's WAAS by reporting pre- and post-optimized traffic accurately.

Scalability

Description of the criteria:

One aspect of scalability is the size of the WAN link that can be terminated on the appliance. More important is how much throughput the box can actually support with the relevant and desired optimization functionality turned on. Other aspects of scalability include how many simultaneous TCP connections the appliance can support as well as how many branches or users a vendor's complete solution can support.

Downward scalability is also important. Downward scalability refers to the ability of the vendor to offer cost effective products for small branches or even individual laptops.

Evaluation of Cisco:

Cisco's WAAS software runs on a family of 5 models of WAE appliance, including single processor and dual core processor models for the branch office that scale to 2X 300 GB of disk storage and 2 to 4 GB of RAM and a dual processor WAE for the central site or large branch that scales to 3 TB of storage and 24 GB of RAM. Network interfaces for all these devices are two 10/100/1000BASE-T ports.

There are two ways the WAE can be installed in the network. They can be installed off-line with WCCPv2 (Web Cache Coordination Protocol) or PBR (policy-based routing) redirect TCP traffic to the "one armed" appliance. Alternately, they can be installed in-line, with the WAE device deployed between the switch and WAN router so that all packets leaving the remote office traverse the Cisco WAE.

At the branch, the inline configuration can be scaled up by installing multiple WAE devices in series. This form of daisy-chain clustering provides active-active redundancy because the WAE has a fail-to-wire feature that bypasses the device in the event of failure. At the smaller branch, the solution scales down by using WAE modules installed in Integrated Services Routers (ISRs). Larger central sites can scale up using load-balanced arrays of WAE appliances with redirection controlled by WCCP.

As previously note, Miercom concluded that Cisco's WAAS shows fairness of connection throughput and performance consistency across concurrent users and multiple workloads and that Cisco's WAAS optimization does not degrade under load.

Cost Effectiveness

Description of the criteria:

This criterion is related to scalability. In particular, it is important to understand what the initial solution costs. It is also important to understand how the cost of the solution changes as the scope and scale of the deployment increases.

Evaluation of Cisco:

Street pricing of the Cisco solution is not publicly available information. The cost effectiveness, however, of the WAAS solution is enhanced by Cisco's position in the enterprise networking market. In particular, Cisco is in a unique position to ensure that its solutions integrate with the existing networking infrastructure. This includes routers, load balancers, and content distribution appliances. This also includes Cisco's Application Control Engine (ACE) which is used by IT organizations to maximize the performance of Web based applications as well as to maximize the availability and performance of servers.

The more seamlessly a WAN optimization solution integrates with the existing infrastructure, the more positive the impact on the overall TCO. Because of this, AM&A recommends that IT organizations test integration when they are testing other features of a WAN optimization solution such as performance and transparency.

Application Sub-classification

Description of the criteria:

An application such as Citrix or SAP is comprised of multiple modules with varying characteristics. Some Branch Office Optimization Solutions can classify at the individual module level, while others can only classify at the application level.

Evaluation of Cisco:

For QoS purposes, Network Based Application Recognition (NBAR) in Cisco routers can be used to identify and classify traffic for the various modules within enterprise applications. WAAS uses static Layer 4 port numbers, and universally unique identifiers (UUIDs) of applications that use dynamic port numbers, to classify traffic for acceleration. Therefore, any application module that can be distinguished in this way can be assigned a custom acceleration policy.

Module vs. Application Optimization

Description of the criteria:

In line with the previous criterion, some Branch Office Optimization Solutions treat each module of an application in the same fashion. Other solutions treat modules based both on the criticality and characteristics of that module. For example, some solutions apply the same optimization techniques to all of SAP, while other solutions would apply different techniques to the individual SAP modules based on factors such as their business importance and latency sensitivity.

Evaluation of Cisco:

Cisco's WAAS includes pre-defined default policies for over 150 application classifiers based on Layer 4 port numbers. The WAAS default policy is to apply the full suite of packet level optimizations (LZ, TFO, and DRE) to all enterprise application traffic, including Citrix, Oracle, and SAP. WAAS also provides the flexibility to define custom policies for specific applications. As noted above, if the application module can be classified by Layer 4 port number or UUID, it can be assigned a custom policy independent of the other modules of the same application.

Disk vs. RAM

Description of the criteria:

Advanced compression solutions can be either disk- or RAM-based. Disk-based systems typically can store as much as 1,000 times the volume of patterns in their dictionaries as compared with RAM-based systems, and those dictionaries can persist across power failures. The data, however, is slower to access than it would be with the typical RAM-based implementations, although the performance gains of a disk-based system are likely to more than compensate for this extra delay. While disks are more cost-effective than a

RAM-based solution on a per byte basis, given the size of these systems they do add to the overall cost and introduce additional points of failure to a solution. Standard techniques such as RAID can mitigate the risk associated with these points of failure.

Evaluation of Cisco:

WAAS DRE and persistent LZ compression are currently focused on TCP traffic where large disk-based data dictionaries provide the maximum compression ratios. WAAS session-based adaptive Lempel-Ziv (LZ) compression is RAM-based and minimizes the amount of bandwidth consumed per message within a session. RAM-based compression solutions are also applicable for real-time compression of UDP packet streams such as those produced by video applications and VoIP. UDP applications are not currently supported by WAAS. Cisco WAE appliances with two or more disks are configured with RAID1 mirroring pairs to protect against disk failures.

Protocol support

Description of the criteria:

Some solutions are specifically designed to support a given protocol (e.g., UDP, TCP, HTTP, Microsoft Print Services, CIFS, MAPI) while other solutions support that protocol generically. In either case, the critical issue is how much of an improvement in the performance of that protocol the solution can cause in the type of environment in which the solution will be deployed.

It is also important to understand if the solution makes any modifications to the protocol that could cause unwanted side effects.

Evaluation of Cisco:

The WAAS TCP optimization features (TFO, DRE, and LZ) provide optimizations for all applications that use TCP. In addition, Cisco's WAAS supports some application-specific acceleration features:

- Operation prediction and batching—Allows a WAAS device to transform a command sequence into a shorter sequence over the WAN to reduce roundtrips.

- Intelligent message suppression—Even though TFO optimizes traffic over a WAN, protocol messages between branch office clients and remote servers can still cause slow application response time. To resolve this issue, each WAAS device contains application proxies that can respond to messages locally so that the client does not have to wait for a response from the remote server. The application proxies use a variety of techniques including caching, command batching, prediction, and resource prefetch to increase the response time of remote applications.

- WAFS caching—Allows a WAAS device to reply to client requests using locally cached data instead of retrieving this data from remote file and application servers.

At the present time, WAAS supports this type of application-specific functionality for Microsoft CIFS and Microsoft Print Services.

As was the case with the OSI layer, the discussion of protocol support and which vendor's implementation yields the greatest performance gains is somewhat academic. The determination of which vendor's implementation yields the greatest performance gains can best be answered by testing the solutions in an environment similar to the one in which they will be deployed.

Security

Description of the criteria:

The solution must not break the current security environment, such as breaking firewall Access Control Lists (ACLs) by hiding TCP header information. In addition, the solution itself must not create any additional security vulnerabilities.

Evaluation of Cisco:

As noted earlier in connection with the discussion of transparency, WAAS preserves the native packet headers, allowing switch and router ACLs, and stateless/stateful firewalls to function normally. Because the packet payload is compressed, the functionality of any security devices (e.g., IPS) that perform deep packet inspection beyond the header in the payload will be affected. This is a limitation that applies to most WAN optimization devices that use compression.

Cisco's WAAS incorporates the following security and data protection capabilities.

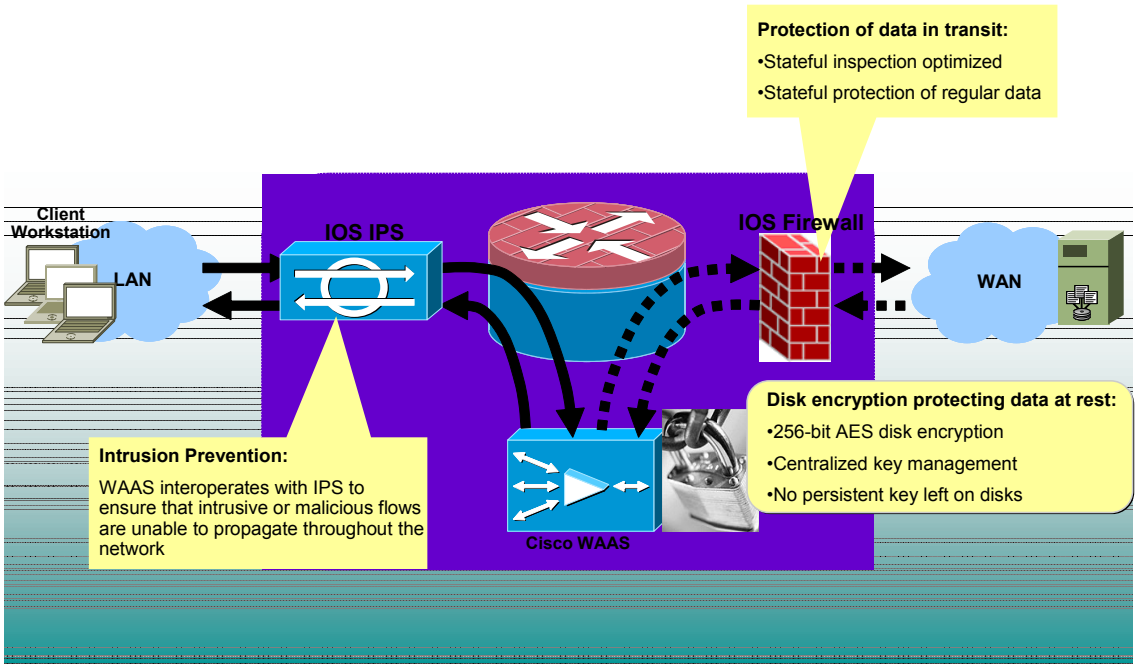

- Disk Encryption - Encryption of all cached content on the remote WAE appliance prevents unauthorized data access or theft. Federal Information Processing Standard (FIPS) 197 approved technology, the Advanced Encryption Standard (AES) 256 bit encryption – the highest commercially available encryption is used. The automated central key manager simplifies key management, provides centralized failover capability for high availability, and supports backup and restoration of keys to offline vaults for disaster recovery purposes.

- Security Compliance - Cisco's WAAS combined with Cisco IOS provides stateful firewall inspection for accelerated traffic, network virus scanning using Intrusion Prevention Systems (IPS) and encryption of data at rest, maintaining compliance with common industry-wide security standards. In addition, WAAS supports Roles-based Access Control (RBAC) to isolate users to specific capabilities and domains of management. Cisco states that WAAS is the only WAN Optimization vendor to be approved for evaluation in the Common Criteria Evaluation and Validation Scheme (CCEVS) or ISO 15408 and Payment Card Industry (PCI) 2.0 compliance.

- Stateful Firewall Protection – Because Cisco's WAAS supports stateful inspection of WAN-optimized traffic through certified interoperability between its firewalls and WAAS, IT organizations can receive equal protection for optimized and regular data through the capabilities:

  - Full stateful protection for WAAS optimized traffic

  - Full stateful functionality for non-optimized traffic

  - No static open ports needed except for management and CIFS

- Intrusion Prevention (IPS) Interoperability - Cisco's Intrusion Prevention Systems (IPS) provides virus scanning in the network while maintaining full interoperability with WAAS. This allows customers to perform WAN optimization while protecting their networks against viruses.

- Role-based Access Control (RBAC) - The Cisco WAAS Central Manager offers authentication, authorization, and accounting (AAA) integration with external authentication providers such as Microsoft Active Directory, RADIUS, and TACACS+. Customers can create profiles based on role, department, responsibility and other parameters to ensure secure access to WAAS.

Easy of Deployment and Management

Description of the criteria:

As part of deploying a WAN optimization solution, an appliance needs to be deployed in branch offices that will most likely not have any IT staff. As such, it is important that

unskilled personnel can install the solution. In addition, the greater the number of appliances deployed, the more important it is that they are easy to configure and manage.

It's also important to consider what other systems will have to be touched in order to implement the WAN optimization solution. Some solutions, especially cache-based or WAFS solutions, require that every file server be accessed during implementation.

Evaluation of Cisco:

WAAS Central Manager is a web-based central management tool that provides simplified configuration, provisioning, monitoring, fault-management, logging and reporting for up to 2,500 WAEs within a Cisco WAAS topology. Cisco's WAAS requires no modifications to applications, clients, or servers in order to provide acceleration services.

A number of factors make Cisco's WAAS relatively easy to deploy. For example, as previously mentioned, Cisco's WAAS does not rely on tunnels. For this reason and the fact that the majority of the existing network equipment is from Cisco, the Cisco WAAS solution is less likely to break the existing network than are some other solutions. Breaking the existing network is a good example of what it means to not be easy to deploy and maintain.

In addition, WAAS supports auto-discovery whereby the solution checks to see if a peer acceleration appliance exists in the packet flow between the source and the destination. If one does, an optimization policy is transparently negotiated and then applied. If not, the application flow is not changed. Auto-discovery is helpful in those situations in which a WAN optimization solution has been deployed in some sites, but not others. This functionality eliminates the need to have IT organizations implement an overlay network. Creating an overlay network is another example of what it means to not be easy to deploy and maintain.

Change Management

Description of the criteria:

Since most networks experience periodic changes such as the addition of new sites or new applications, it is important that the WAN optimization solution can adapt to these changes easily. It is preferable if the WAN optimization solution can adjust to these changes automatically.

Evaluation of Cisco:

With auto-discovery of WAE and auto-negotiation of application policies, new sites and new applications can be added to the network with a high degree of automation. In addition, the WAAS Central Manager system automates the download of new configuration files to device groups of WAEs that share a common configuration.

Support of Meshed Traffic

Description of the criteria:

A number of factors are causing a shift in the flow of WAN traffic away from a simple hub-and-spoke pattern to more of a meshed flow. If a company is making this transition, it is important that the WAN optimization solution that they deploy can support meshed traffic flows and can support a range of features such as asymmetric routing.

Evaluation of Cisco:

Due to its transparency, the operation of the WAAS solution is completely independent of WAN architecture, including partially or fully meshed architectures.

Support for Real Time Traffic

Description of the criteria:

Many companies have deployed real-time applications. For these companies it is important that the WAN optimization solution can support real time traffic.

Some real time traffic like VOIP and live video can't be accelerated because it is real time and already highly compressed. Header compression might be helpful for VoIP traffic and most real time traffic will benefit from QoS.

Evaluation of Cisco:

This Cisco approach is that header compression and QoS functions are provided by Cisco routers in the branch and central site routers. WAAS does not modify the traffic generated by real-time applications running over UDP, such as voice and video. However, for video, Cisco offers ACNS software for the WAE that provides optimization of video distribution. For VoIP, WAAS is fully compatible with the required end-to-end QoS functionality, while also maximizing the bandwidth available for voice traffic by compressing and optimizing TCP data traffic sharing the same WAN circuits.

The issue of what functionality should be done in a router and what should be done in a separate appliance is an architectural issue that is beyond the scope of this document. That being said, it is clear that whatever functionality is invoked in the router needs to be harmonized with whatever functionality is invoked in an appliance. For example, if QoS is set in the router, the appliance should not do anything to interfere with that.

Summary

The criteria listed in Table 1 were applied to the Cisco WAN optimization solution. While no attempt was made to assign a grade to the solution, that analysis did point out that Cisco's WAAS has a number of strengths. These include:

- Transparency

As mentioned, for many IT organizations the first rule of networking is that you should not do anything that causes the network to break. For example, it should be possible to implement a WAN optimization solution and not break existing functionality such as security, QoS or routing. It should also be possible to implement a WAN optimization solution and still be able to use the existing management and monitoring tools.

Given that WAAS does not change the packet headers, it provides a high degree of transparency. That transparency has an impact on TCO as well as the following factor.

- Ease of Deployment and Management

  Because it does not modify packet headers, Cisco's WAAS is less likely to break some aspect of IT than are some other solutions. As a result, WAAS is easier to deploy and manage than solutions that break something.

  In addition, WAAS supports auto-discovery whereby the solution checks to see if a peer acceleration appliance exists in the packet flow between the source and the destination. This functionality makes deployment easier by eliminating the need to have IT organizations implement an overlay network.

- Integration

  While this was not one of the criteria listed in Table 1, it impacts several of the criteria that were mentioned in that table. The point being that given Cisco's position in the enterprise networking marketplace, the chances are that the environment in which the WAN optimization solutions will be deployed will be a Cisco environment. Cisco is in the best position to ensure that its WAN optimization solution does not impact this environment.

- Performance

  Some of the comments that Miercom made about Cisco's WAAS include:

  - Under favorable conditions (large WAN latency and highly compressible content) the WAAS V4.0 solution can provide impressive degrees of acceleration.

  - The software demonstrated performance parity with these other products, and that in some key metrics provided superior compression, speed and throughput.

  - Unlike some of the other products in the market, Cisco's WAAS does not degrade under load.