Blue Coat SG Client Technical Overview

White Paper

INTRODUCTION

Two major business trends have been driving IT organizations to rethink the technical solutions they have in place for delivering applications to end users:

- Consolidation of applications and data resources
- Push towards 'anytime anywhere' access to information

First, for several reasons – such as cost savings and regulatory compliance – organizations are consolidating their applications and data, such as centralizing all resources in a data center. With the increased geographic distance – e.g., mail server is now three thousand miles away versus thirty feet – and the well-known latency issues of many of today's business applications and their underlying protocols (e.g., CIFS), end users – both those in remote offices as well as mobile employees – are experiencing poor and often unacceptable application performance.

Second, organizations are increasingly leveraging the web to provide 'anytime anywhere' access to information and applications through increased use of internally-hosted web applications as well as ASPs. Additionally, more employees than ever are mobile and untethered from the corporate LAN, using the web for remote access to enterprise resources. The public networks, however, present obvious security risks for end users and enterprise data as well as being unpredictable with respect to performance.

As a result, IT organizations must rethink how they deliver applications to end users that are beyond the reach of corporate LAN to provide the security and application performance that they require. That is, organizations must deploy solutions that can deliver the necessary security and acceleration instrumentation to the endpoints of their users.

This white paper provides a technical overview of the SG Client architecture along with background information on the market drivers and the current technologies available for delivering security and acceleration services.

COMPUTING PARADIGM SHIFT & IMPACT ON APPLICATION DELIVERY

The use of the web to provide 'anytime anywhere' access – while seemingly making the computing environment simpler – presents several technical challenges to IT departments tasked with supporting web users:

- Web applications are available to virtually any computing endpoint, not just corporate laptops.
- Web applications separate end users from devices so users and their policies cannot be tied to a specific endpoint.
- Web applications separate applications from devices. Web applications are delivered on demand and are not preinstalled.
- Web applications have evolved from static web HTML to new technologies, such as Web 2.0 technologies (e.g. web services, service-oriented architecture (SOA) and AJAX (Asynchronous JavaScript, XML, etc.)) and deliver all types of content – dynamically-generated content, images, pictures, streaming video, etc.

The movement to the web combined with the consolidation of applications and data resources presents a computing paradigm shift that presents new challenges for IT departments and their ability to deliver applications:

- **Existing Applications.** Legacy – both client-server and web – applications that were designed for reliable LAN environments must be changed or artificially "optimized" to adapt the new distributed environment. For example, applications that require a high number of connections, have inefficient protocol negotiations or are highly sensitive to data transfer rates must be improved – either by re-architecting or through third party technologies – to provide consistent performance and an acceptable user experience.
- **New Applications.** The shear number of possible connecting endpoints, the on demand delivery of web applications and the consolidation of enterprise data require new applications to be developed without making any assumptions about underlying network environment. Applications must work for all network environments – wired and wireless, LAN and WAN, low capacity and high latency, high bandwidth and high packet loss, etc.
- **Data Security.** Data this processed that assumes that the computing environment – end users, devices, applications and the network – is secure must be rethought. Authentication, authorization, data privacy and data integrity must be added to the data independent of endpoint device and application.
- **Accessibility.** Availability and reliability of data and applications to every authorized user and all available endpoint devices must be reconsidered for vast and heterogeneous network environment where devices are unknown and networks unpredictable.

This new computing paradigm requires new solutions for delivering applications with the security the enterprise demands and the end user experience the business requires.

TECHNOLOGIES FOR ADDRESSING SECURITY & ACCELERATION ISSUES

First, an overview of the various technologies used today to address application performance and content security is necessary before introducing the specific challenges of delivering acceleration and security to the endpoints.

Technologies implemented to address application performance and data security are similar to solutions designed to solve other networking or security issues in their implementation at different layers within the network-based computing environment, often addressing specific layers of the OSI model.

Not surprisingly, each of the layer-specific approaches has its merits as well as its drawbacks. Generally speaking, technologies implemented at the lower layers have less "intelligence" than solutions at higher levels (for example, the difference between knowing a solution is TCP versus HTTP). And, generally speaking, higher layer solutions can discern more context (i.e. see the overall picture) than those implemented at lower layers.

For the acceleration of applications and for the security of application content, there are four layers at which solutions are traditionally developed:

- Network layer
- Transport layer
- Application layer
- Content layer

Network layer solutions address acceleration and security problems for IP packet delivery, such as packet QoS and different packet priority queuing techniques. Also, packet-type-based bandwidth allocation techniques are used and new techniques, such as packet payload reduction and caching, are being. For security, IPSec is used to secure packets authenticity, privacy and integrity.

Transport layer solutions have been developed to improve the performance and security of transports, mainly the TCP/IP transport. Most performance technologies focus on flow control, such as improved algorithms for TCP window scaling; congestion detection and control; latency detection and control; and packet acknowledgment and retransmission control. Compression and caching are also used at the transport layer. For security, SSL is the most popular for protecting content at the transport layer.

Application layer (or application protocol layer) solutions have emerged recently for WAN optimization / application acceleration. Techniques such as protocol optimization, data pre-fetching and caching are now common. Protocol optimization is devised to remove application-specific protocol deficiencies, such as the chattiness of the protocol (e.g., MAPI and CIFS), sequencing of messages (send message, await confirmation), the frequency of short messages, etc. Dynamic data pre-fetching and caching are also used. For security, application filtering is done to control URL access and block malicious content.

One noticeable difference between the above layers and those of the OSI model is the inclusion of the "content layer," a layer which has recently become popularized by industry analysts and technology vendors. The content layer is a layer above the application layer of the OSI model (layer 7).

The main reason for adding the content layer is the increasing popularity of web-based applications. From the OSI perspective, the web communication protocols, namely HTTP and HTTPS, belong to the application layer. For web applications, HTTP and HTTPS are being used as a virtual transport to deliver higher layer protocols and applications. Understanding and processing the content delivered via HTTP and HTTPS requires an abstraction layer above the OSI application layer, which is now represented by the content layer. Furthermore, the concept of the content layer does not only apply to HTTP and HTTPS, but it also to other application protocols such as CIFS, MAPI, SIP and SOAP. *(NB: the content layer is not an official layer of the OSI model, but it basically subdivides the application layer for web applications so the protocol and application content can be discussed separately.)*

The content layer provides a layer of abstraction to separate web communication protocols – such as HTTP – from the applications content. Most acceleration techniques at the content layer leverage the knowledge obtained by analyzing the application content and apply data reduction and data security techniques specific to the various content types.

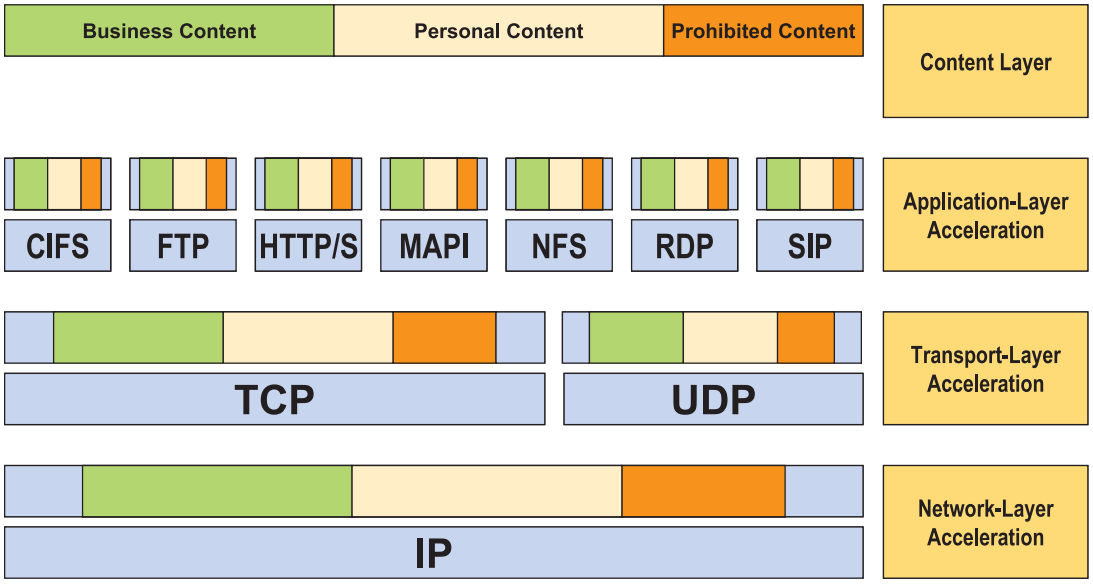

Figure 1: Endpoint solution without content layer intelligence. Because lower layer solutions don't have any intelligence about the content the lower layer acceleration solutions indiscriminately accelerate all content.

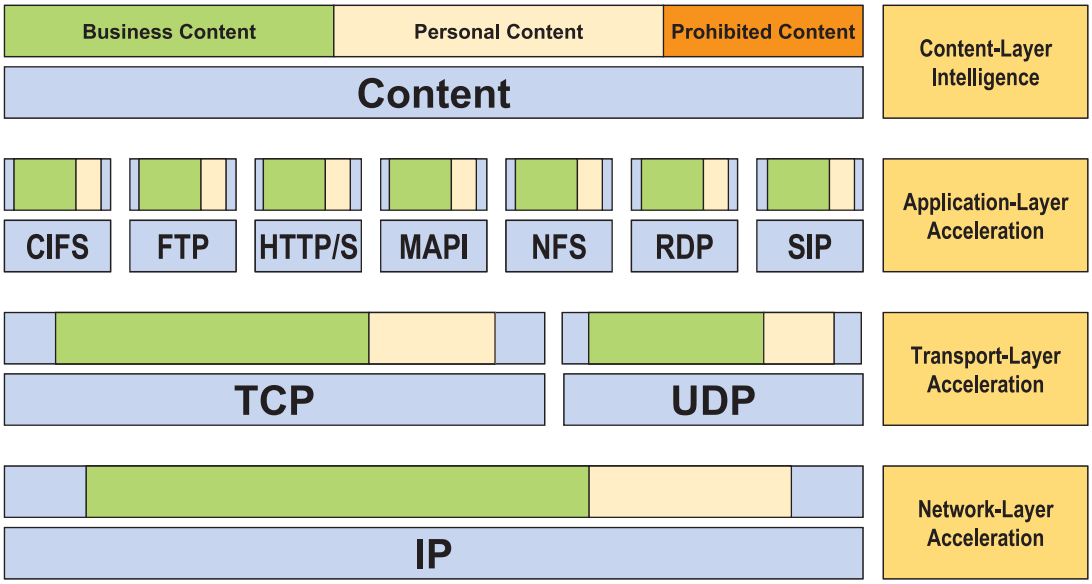

Figure 2: Endpoint solutions with content layer intelligence. With content intelligence, only the appropriate content is accelerated by the lower level technologies.

CHALLENGES & OPPORTUNITIES FOR ENDPOINT SERVICES

The acceleration and security technologies described above have been implemented by various vendors across a broad spectrum of products.

Most of these products are implemented as intermediation gateways – which are deployed as standalone gateways and / or communicate with one or more other gateways to provide point-to-point acceleration and security services. These gateways often do not require any changes to the applications or to the endpoint devices and can, if necessary, co-exist transparently with other gateway solutions. They are primarily implemented within WAN environments to provide application acceleration and security services to users – such as branch employees – connected directly to the LAN.

As discussed earlier, many users are mobile and using the public networks to access internal and externally-hosted applications. For these users, gateway solutions, for example, provide minimal value. To be sure, there are several acceleration techniques, such as pipelining and object caching for web applications, that accelerate applications without touching the user's machine, but most of the technologies require endpoint instrumentation.

When compared with gateway solutions, endpoint solutions present different challenges as well as presenting new opportunities for acceleration and control that are not available in gateway-only solutions. There are the obvious client software challenges of endpoint software provisioning, installation and ongoing management. There are also the issues with supporting a diverse set of endpoint devices – such as laptops, PDAs, and smart phones – that use various operating systems. Additionally, when virtually any device is a potential enterprise computing device, organizations – and consequently technologies – must anticipate encountering various endpoint configuration and network environments. Finally, different types of users will have different types of security policies, which are determined by the user's device and network.

As described earlier, the use of the public infrastructure and the business requirement for anytime anywhere access to applications and data dramatically increases the total number of endpoints and introduces new security risks and acceleration requirements (such as accelerating over wireless links). In short, when delivering applications into unknown environments, there are new requirements that ultimately impact the architecture and features for an acceleration and security client.

While there are challenges (i.e. it's significantly more difficult than simply putting acceleration technologies into a client format), endpoint software presents new opportunities to add acceleration and technology features that extend beyond the solutions implemented at the various OSI layers described above. That is, endpoint instrumentation provides the ability to add security and acceleration features at the content layer.

Most technologies implemented at the network, transport and application layers are all applicable at the endpoint. By layering network drivers, endpoints can have network layer acceleration and security functions such as packet QoS and IPSec. By layering transport providers, endpoints can have transport layer acceleration and security functions such as optimized flow control, data reductions and SSL. By adding application specific proxies, endpoints can have application layer acceleration and security services such as object caching, protocol optimization and content filtering.

Beyond the traditional gateway technologies, there are new areas where endpoint solutions can be extended to provide additional functionality. First, the endpoint is the only place where network traffic can be classified before entering the network to be delivered to the application. Second, only at the endpoint can network traffic be accurately associated with a specific application and user. This level of visibility and content awareness allow endpoint solutions to perform content-oriented application acceleration and security functions. For example, acceleration and security functions such as application bandwidth management and QoS, content reduction, content filtering, content protection and content right management can all be implemented at the content layer – all of these functions are beyond what is technically possible at the gateway.

BLUE COAT SG CLIENT ARCHITECTURE

Blue Coat SG Client is designed to address the acceleration and security challenges of the endpoint.

Blue Coat SG Client uses an architecture similar to a service-oriented architecture (SOA). Blue Coat SG Client uses a policy-oriented architecture (abbreviated in this paper as POA). Unlike SOA where services are delivered from application servers upon the service requests from endpoints, POA delivers acceleration and security services to endpoints based on endpoints operating policy. Instead of sending service requests to application servers in SOA-based web computing, SG Client sends endpoint operating specifications (terms, provisions) to the Blue Coat SG Client Manager (CM).

One of the key elements in SG Client POA is its ability to supporting multiple types of endpoint devices. Where gateway-based solutions make assumptions about the connecting device and primarily support corporate-managed devices, SG Client POA is designed to support both managed devices as well as unmanaged endpoints – those beyond the control of corporate IT.

The foundation of SG Client is the client service framework (abbreviated in this paper as CSF). CSF consists of two components.

- On-demand service agent
- Persistent service agent

The on-demand service agent is a web service to deliver security and acceleration functions to the end user when endpoint device is connected to the enterprise network.

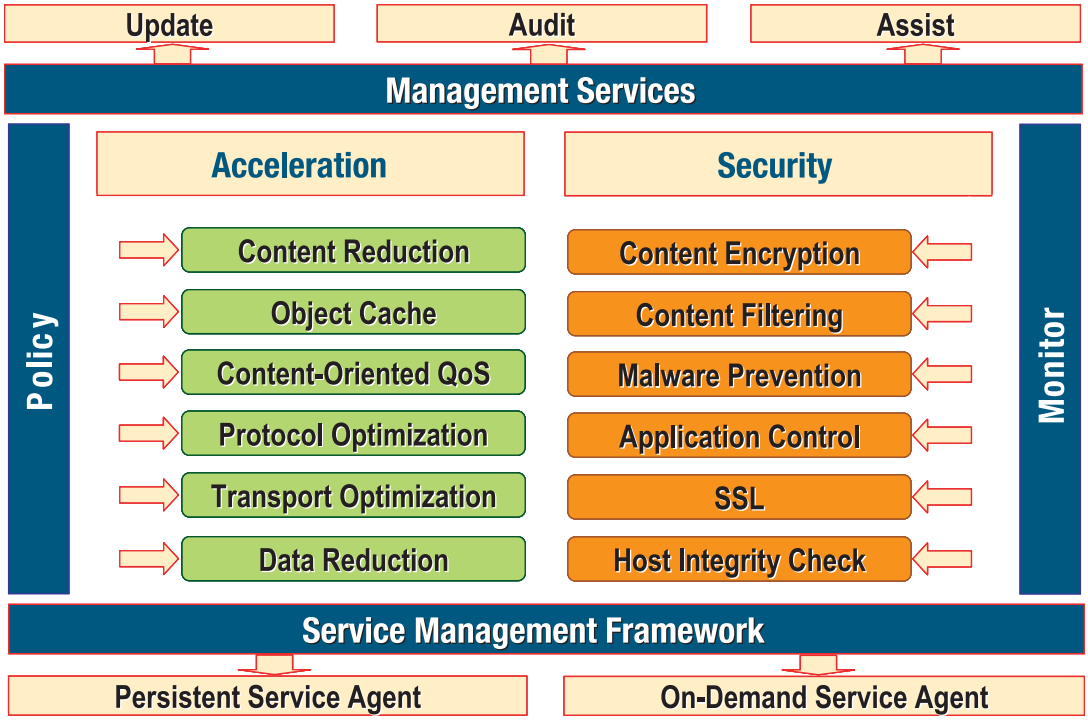

Figure 3: Blue Coat SG Client policy-oriented architecture.

Unlike the on-demand service agent which is completely transient and does not require client software installation, the persistent service agent is a piece of web installed client software specifically targeted at endpoints that require additional services that are not available to the on-demand service agent. The persistent service agent adds two special services.

- System resident (embedded) services
- User offline services

Storage and patch services are examples of system resident services. The user offline services are application services that are performed when endpoint devices are not connected to the enterprise networks. Offline file services, application content filtering and application data protection are examples of user offline services.

Before delivering and enabling endpoint services, SG Client gathers the following endpoint operating specifications:

- Hardware and software configuration
- Network location and networking environment
- End user identity and security configuration

Based on the information received from the endpoint and the policies defined by IT, the SG Client Manager delivers a set of security and acceleration function to the endpoint. The set of acceleration and security functions available to connecting endpoints include:

1. Application protocol optimization
2. Application data compression and caching
3. Transport layer optimization and SSL data encryption
4. Endpoint integrity monitoring
5. Location-aware load balancing and failover
6. Content-layer application filtering, object encryption and information protection
7. Content-layer bandwidth management and data delivery priority
8. Content-layer business process performance monitoring

Numbers 1 through 4 are commonly understood technologies and are not further explained in this paper. Number 5 through 8, however, represent security and acceleration functions that are only possible at the endpoint.

Location-aware load balancing and failover allows endpoints to connect to (or failover to) the closest geographic gateway to minimize the actual distance the application data must travel.

Content-aware services are only possible at the endpoint where network traffic can be accurately segmented by content type, application, business process and user. Content aware operations listed below can be performed by SG Client.

- Bandwidth management
- Data delivery priority
- Application filtering
- Object encryption
- Information protection
- Process performance monitoring

Such content-aware operations are especially important for endpoints where the operating environment is unpredictable.

BLUE COAT TECHNOLOGY ADVANTAGE

Delivering a flexible solution that can deliver a variety of security and acceleration functions to the endpoint requires unique technologies to address:

- Diversity of devices – hardware, operating system and application combinations
- Diversity of operating environment – managed devices, unmanaged devices, predictable and unpredictable networking environments
- Diversity of users – knowledge workers, task-oriented workers, customers and partners

The architectural challenge is that it is not feasible to predict all the possible operating environments and deliver every possible control. It is also not feasible to constantly update and roll out new client versions to adapt in evolving computing environments.

To address this, Blue Coat has developed patent-pending Connector technology. Connector technology was developed as client middleware that allows new acceleration and security functions to be added to the SG Client without modifying the underlying client, i.e. without releasing a new version of the client.

Connector enables SG Client to address two critical requirements for the endpoint: the ability to support virtually any endpoint and the ability to deliver a wide and extensible suite of acceleration and security functions.

In a very high level, Connector is client software that performs the following two functions:

- Interception of application operations within application operating space
- Proxy application function requests through Connector's function providers

Applications rely on a computing device's operating system (OS) to perform its functions. A common technique to add new controls to applications is called OS layered shimming. All operating systems use a layered software module approach to deliver their services to the applications. Different layers are responsible for processing different types of application data and functions. By inserting software modules into different layers, the shimmed software can exert different controls on the application data and functions operating at shimmed layers. For example, anti-virus software inserts a software module into an operating system's file system layer. The shimmed software would then scan all the files on the system to find and remove virus files. As another example, IPSec software inserts a module into OS' IP packet layer. The shimmed software would encapsulate original IP packets into encrypted packets. Connector takes a different approach. It injects software module directly into application process space when the application has or is being loaded into operating system memory.

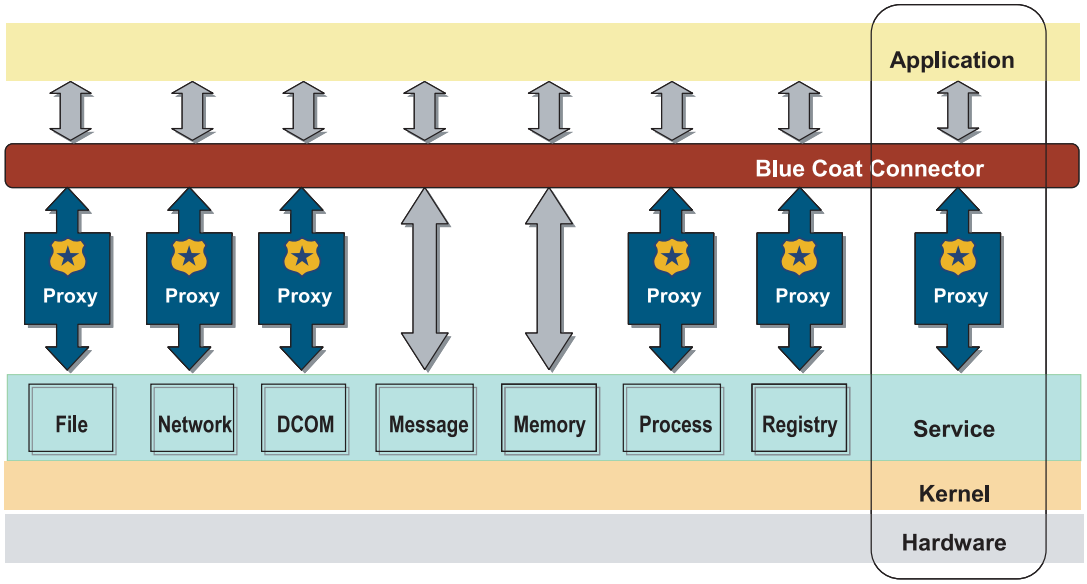

Figure 4: Blue Coat Connector in the context of an endpoint computing platform. The Blue Coat logos represent instances where the Connector is proxying services and operations used by the application.

Software modules are inter-connected through well-defined interfaces. Applications that need services or functions provided by other software modules such as third party tools and OS services must use these interfaces to get the services. To control certain application behaviors or usages, Connector identifies the interfaces that the application uses for its operations and the proxies those interfaces to Connector interfaces.

CONCLUSION

With trends such as data consolidation dragging down application performance and the drive towards anytime anywhere access to information and applications, organizations must rethink the technologies used to deliver applications to end users. Gateway-based solutions serve a subset of users, but many users – such as mobile users, partners, and customers – get left behind, suffering from poor application performance and exposing sensitive personal and corporate data. Enterprises needs an endpoint solution that can deliver the variety of security and acceleration services to a growing number of devices and users.

With its unique policy-oriented architecture and connector technology, Blue Coat SG Client is capable of meeting application acceleration and security challenges in the new era of distributed computing, web computing and data consolidation.

420 North Mary Ave.
Sunnyvale, CA 94085
www.bluecoat.com

1.866.30.BCOAT
408.220.2200 Direct
408.220.2250 Fax

Blue Coat secures Web communications and accelerates business applications across the distributed enterprise. Blue Coat's family of appliances and client-based solutions – deployed in branch offices, Internet gateways, end points, and data centers – provide intelligent points of policy-based control enabling IT organizations to optimize security and accelerate performance for all users and applications.Blue Coat is headquartered in Sunnyvale, California, and can be reached at 408.220.2200 or www.bluecoat.com.