



SiteCelerate

white paper

SITECELERATE OVERVIEW

As enterprises increase their investment in Web applications, Portal and websites and as usage of these applications increase, performance degrades and bandwidth usage and costs shoot up.

To minimize the challenges inherent in delivering interactive HTTP-based applications over the web or Intranets, Arahe Solutions offers the SiteCelerate Web Acceleration System, which consolidates key application delivery functions. Deployment of the SiteCelerate System solution at your data center can simplify your IT infrastructure, improve application availability, and deliver large gains in performance for your core enterprise-class applications.

The field of content delivery has long relied on caching technology to reduce bandwidth usage and improve the end-user experience. Such caching technology has become a key element of ISP and corporate network architectures in the form of caching proxy systems. Arahe Solutions's SiteCelerate adds a new dimension to content delivery and caching — web acceleration.

Arahe SiteCelerate is an in-the-flow web acceleration system that can be deployed either as a full optimizing and caching system, or as a device sitting alongside an existing caching server. Either way, Arahe's SiteCelerate optimizes content as it is retrieved from the upstream network. Benefits of caching optimized content include increased cache capacity, reduced downstream bandwidth consumption, and reduced download times for end-users.

SiteCelerate performs four specific functions:

1. Acceleration: Deliver the best possible response times over the WAN and Internet
2. Optimization: Deliver the best response times while minimizing bandwidth use
3. Smart Dynamic Caching: Deliver content only when necessary
4. Server and SSL Offload: Deliver common functions as efficiently as possible, minimizing server and other data center resources

While a particular area may be more important in a given situation, most customers gain benefits across their entire IT infrastructure. In each of these cases, Arahe combines specific technologies with unique technologies to deliver your application at lightning speed.

THE SITECELERATE CORE

At the core of the SiteCelerate device is the application delivery engine, a real-time content processing engine that provides the basis for all higher-level performance improvement. The application delivery engine is a full Layer-7 reverse proxy that processes all content. By transforming the browser's own cache into a dynamic engine, the SiteCelerate provides capabilities to further improve performance. And unlike packet-based processing, the SiteCelerate System can optimize both within and across individual Internet sessions and dramatically improve the way that enterprise applications perform over the WAN or Internet.

The Table below provides a summary of the functions of the SiteCelerate.

Feature	Benefit
Real time optimization of all text and graphics	Dramatically reduces size and transparent compression of all content – JSP, HTML, ASP, PHP, CSS, PDF, Flash, Javascript, GIF, JPEG, BMP, PNG, Word documents and many others.
Processes all Web and Web-based applications in real time	Automatic and transparent management of HTML and XML across applications, portals, and other Internet front ends.
Manages and optimizes sessions between and across application connections	Because of its deep context sensitivity, the engine can eliminate and aggregate TCP/IP requests, overcoming basic inefficiencies of running Internet protocols over the network.
Can be adapted and configured via rules for specific applications and content	The exceptions and policy engine allow fast time-to-implementation and control at the individual page and object levels.
Supports all secure modes of operation, including Secure Sockets Layer (SSL) termination and SSL proxy modes	The engine is fully compliant and compatible with other security devices for secure operation.
Supports all enterprise architectures	Scalability and failover are critical. The engine and device support full clustering and integrate with other redundancy solutions, including WCCP routers, Layer 4 switches, load balancers, software clustering, and data center virtualization.
Dynamic cache engine	Allows more efficient caching of content at browser's end and reduce unnecessary transfers of data – including dynamic content.
SSL acceleration built-in	Allows dramatically higher throughput and capacity for SSL content
Intelligent cooperation with other SiteCelerate systems or third-party caching systems	Multiple SiteCelerate systems can work together to deliver better performance at a single datacenter or at multiple locations
Full featured Web based administration and live graphs	Allows full management of SiteCelerate and full view of live statistics at any location

The above features combine to deliver dramatic performance improvements and incredible bandwidth usage reduction across the WAN or Web.

ACCELERATE: MINIMIZE NETWORK LATENCY

IT organizations have many strategies for dealing with the performance needs of Websites; the situation is significantly more complex for interactive applications. It is not uncommon to have as many as 100 network traversals to render a single page, so even 100 milliseconds multiplies quickly. Traditional packet-based solutions can only operate within the context of a single traversal; two-box symmetric solutions can work with TCP/IP, but can do little to minimize traversals for most Web applications.

The SiteCelerate application delivery engine maintains the state of the entire application across all clients and servers. The knowledge of the context of requests allows the organization to understand and minimize network calls. For example, through its configuration, the SiteCelerate can transform data previously considered un-cacheable and eliminate the need to check with either Web or application servers. By working to aggregate Web requests and minimize unnecessary network calls, the SiteCelerate System routinely delivers between 2 and 5 times the improvement across the WAN. Moreover, these gains are accomplished for users regardless of their location, access, or client system.

OPTIMIZE: USING MINIMAL BANDWIDTH

The application delivery challenge is often not just about overcoming network latency. Invariably, organizations want to minimize their use of bandwidth either for cost or availability reasons. Many Arahe customers support application usage in areas of the world that have limited bandwidth. Some customers operate branch offices where even high-bandwidth links are shared among remote populations.

Dynamic Browser Caching

Many enterprise web applications and portal applications often mark some objects, such as images, JavaScript files, ActiveX control files, or binary files, as non-cacheable for a variety of reasons. Perhaps the object uniqueness is not determined by the object URL itself but rather by an HTTP header value or a cookie value (the object name chosen by the application is the same for all users but the object itself is different for each user). Perhaps the object itself changes frequently or unpredictably based on user actions and is considered not suitable for caching. Use of non-cacheable embedded objects causes poor application performance-it forces clients to retrieve such objects from the origin server in their entirety on each page visit. This can result in extremely slow download performance, especially for remote users with limited bandwidth.

SiteCelerate eliminates the need for users to download these objects on each request. Instead, the SiteCelerate automatically tracks the freshness of each of these objects in real time.

If a requested object has not changed, the device instructs the client to use the cached version of that particular object. If an object has changed, the SiteCelerate delivers it to the client. The SiteCelerate System delivers the object itself only when it determines that the object has changed in that specific context. As part of Arahe's suite of optimizations, this feature helps ensure optimal application response times for all users by eliminating another source of inefficiency and delay.

Image Optimization

The SiteCelerate device provides advanced image optimization capabilities that reduce image file size while preserving image quality, resulting in faster image download times, faster page renders, and more efficient bandwidth utilization.

As always, no changes to client software, server content, or server configuration are required - all configuration is accomplished using the core application delivery engine.

Compression

Using our own GZIP+ compression technology, Arahe goes beyond compression to deliver the smallest file size for all content – whether dynamic or static content.

Utilizing a unique combination of text compression, white space removal and comments stripping, file sizes are kept to a minimum.

Again, there are no changes to client software, server content, or server configuration.

OFFLOAD: MINIMIZE NECESSARY SERVER RESOURCES

It is often surprising to organizations how much server power is needed to support Web capabilities and applications. While organizations are used to sizing databases and servers for given applications, setting capacity across hundreds or thousands of users is a new process. Issues like connection management and encryption begin to impact the ability of servers to actually process core application content.

SiteCelerate can significantly reduce origin web, application and database server CPU utilization rates dramatically by off-loading connection management, SSL encryption and ensuring users connect less frequently and connect for shorter durations. Most customers see a 50% reduction in their server utilization rates.

TCP Connection Multiplexing: Offload Connection Management with Persistent Back-End Connections

With the TCP connection feature, the SiteCelerate can take responsibility and ownership for the overhead of managing network connections with browsers. The device does this by maintaining persistent TCP connections with the Web/application servers. In order to optimize overall performance as traffic levels

change, the SiteCelerate System intelligently increases and decreases the number of persistent TCP connections to the back-end servers as load conditions dictate, further increasing the overall efficiency of the application infrastructure. This frees up Web and application servers to focus solely on content generation, often resulting in a 2X or greater improvement in Web server capacity and adding headroom to allow additional server consolidation.

Caching

The SiteCelerate device uses a high-performance caching architecture to enable several of its patent-pending optimizations, including Delta Optimization and FlashForward Object Acceleration. Static caching also directly offloads server with requests for frequently requested static objects, like images and applets. This feature is fully configurable and adds to overall application performance and transaction throughput.

Adaptive and Configurable Dynamic Caching

SiteCelerate supports sophisticated content expiration policies to further accelerate enterprise application performance and improves server system scalability by enabling the SiteCelerate System itself to fulfill requests for dynamic content. Using this feature, the offload capabilities begin to offload application servers and even core databases. Today's heavily loaded application servers are forced to process millions of requests for dynamic content, even though much of the requested content is the same. These dynamic pages must be generated for each request; they typically require multiple application server queries and database queries, leading to serious bottlenecks, overloaded servers, and poor application performance.

Arahe's Configurable Intelligent Caching capability enables SiteCelerate to intelligently cache and serve dynamic page responses, eliminating the delays associated with application server and database requests and allowing more users and transactions to be served. Arahe's flexible Caching capability enables the SiteCelerate System to treat dynamic content as static and to cache multiple responses for a given URL based on specified cache parameters. These parameters can include URL query strings, HTTP headers, and cookie values. This flexibility significantly reduces application server load, resulting in more users receiving better application performance. With a simple script, even private data can be dynamically cached, leaving more resources for core transactions. Without requiring deployment of remote caches, Configurable Dynamic Caching helps achieve sub-second application response times for dynamic pages while significantly improving overall system scalability.

SSL Acceleration

The Secure Sockets Layer (SSL) protocol has become the industry standard for providing security, privacy, and confidentiality for enterprise business transactions. In fact, compared to the Internet, practically all content is transmitted using SSL. Using SSL acceleration, the SiteCelerate handles the SSL handshake with the client, decrypts Web requests from the client, transparently proxies them to the back-end Web or application servers, compresses the server response, encrypts the optimized responses, and delivers them to the client within the secure SSL connection.

One of the advantages of this approach is that back-end servers no longer need to process SSL. The SiteCelerate greatly accelerates performance under SSL because it minimizes the amount of content to encrypt- rather than entire pages, only the compressed pages need to be encrypted. Arahe's integrated SSL capabilities make the SiteCelerate System a fully transparent solution able to deliver secure content acceleration.

In addition, Arahe's URL rewriting capability enables URLs to be hidden within the HTML source by swapping them with arbitrary URL strings. This capability can be used to isolate the back-end infrastructure by preventing end users from seeing the actual URL structure used by the origin server. In addition, Arahe's FlashForward Object Acceleration eliminates the majority of object validation requests from clients on subsequent page visits, resulting in a dramatic reduction in the number of SSL-based transactions and an up to 4X increase in SSL scalability.

Single Sign-On Optimizations

As part of their security policies, many enterprises have adopted single sign-on (SSO) mechanisms such as Microsoft's NT LAN Manager (NTLM) or Active Directory to authenticate users as they log into enterprise applications. Authentication provides a mechanism for preventing user spoofing. Just as the SiteCelerate improves SSL performance by eliminating unnecessary object validation requests that require costly SSL handshakes, it significantly improves overall application performance in SSO-enabled environments by eliminating redundant authentication traffic associated with object validation requests. SSO-enabled enterprise environments that use the SiteCelerate System see dramatically improved system performance while maintaining high security.

DEPLOYMENT OPTIONS

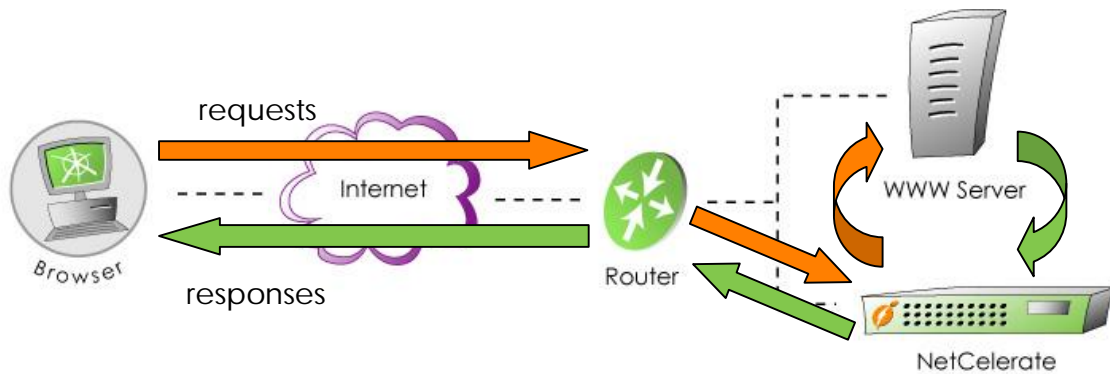
SiteCelerate fits transparently into most enterprise environments. It supports a wide variety of options for ensuring reliable, available, and scalable applications and infrastructure. The Availability Manager, a core part of the SiteCelerate System platform, includes N-way clustering, allowing organizations to choose appropriate configurations.

The SiteCelerate System supports both Active-Active and Active-Passive modes of operation. The device works in conjunction with other failover and scalability components of your IT infrastructure, such as other load balancers and content switches. Whether your organization uses software clustering, Layer 4-7 load balancing devices, or data center virtualization techniques, your teams are assured of the highest levels of performance and uptime.

SiteCelerate can be used to accelerate Web sites using several network architectures. Which architecture you choose will likely depend on your existing architecture and your requirements for availability, capacity, and security. This document presents a number of different network models in which SiteCelerate can accelerate Web sites.

Man in the Middle

A single SiteCelerate system can accelerate an individual server or cluster of servers operating under a single name or IP address. With this approach, the SiteCelerate box impersonates the Web server and handles all requests directly. Using a preconfigured private address for the Web server, it forwards any requests it can't handle on and receives and caches responses.

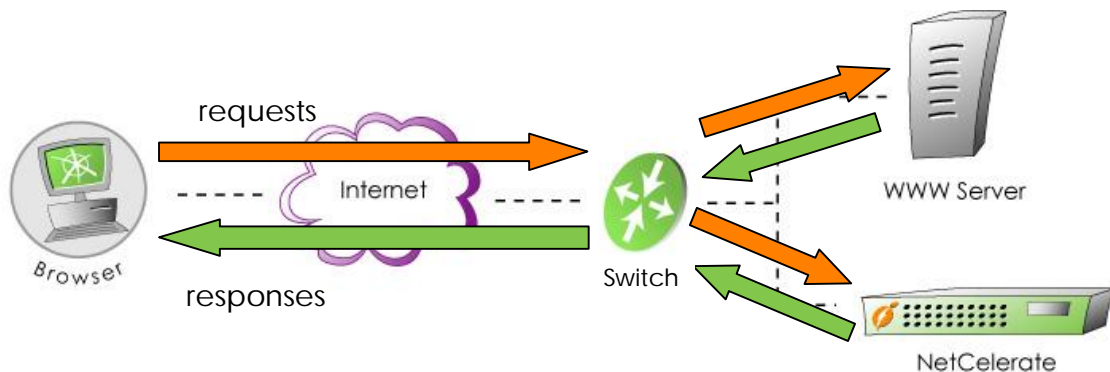


While this approach is simple, there is a downside. The SiteCelerate box must be specifically configured to forward requests on to a specific Web server or named cluster of servers. Moreover, since the SiteCelerate machine is acting as the Web server to the outside world, if it goes offline, the site will be unavailable.

The transparent configurations discussed in this document provide better fail-over.

Transparency with a Switch

A single SiteCelerate system can transparently accelerate a Web server using a switch. This method is considered transparent because the name and address of the Web server do not have to be changed and, in most cases, failure of the SiteCelerate system doesn't cause the Web site to be unavailable. If the SiteCelerate machine goes offline, the Web site will just be served un-accelerated with just a performance degradation.



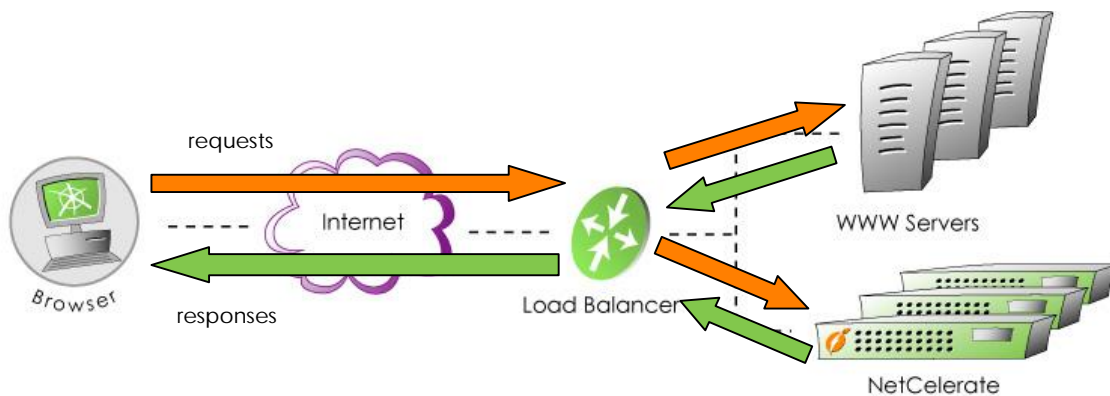
In this network configuration, a switch is provisioned to direct requests intended for the Web server to the SiteCelerate system. Exceptions to the forwarding rules allow the SiteCelerate box to make requests of the server back through the same switch and have them actually reach the server rather than be forwarded back to the accelerator. Alternatively, SiteCelerate can make requests over a private network that is not managed by the switch.

This network model is preferable to the "man in the middle" model because rules can be set to send requests directly to the Web server in the event that the SiteCelerate box is not responding.

Transparency with a Load Balancer

Clusters of Web servers are often managed with load balancing switches to ensure the request and processing load is evenly distributed across all servers in the cluster. In such a case, it's likely that just one SiteCelerate system won't have enough capacity to handle all the traffic for the cluster.

An array of SiteCelerate boxes can be configured to work from the same load balancer as the cluster of servers by either (1) establishing traffic switching rules that route requests from outside the network to the SiteCelerate array and route requests from any SiteCelerate box in the array to the cluster of Web servers, or (2) have the SiteCelerate array replace the Web server cluster as the destination address and then forward requests to the cluster of Web servers using a private address. This second option is viable because failure of any one SiteCelerate system will leave the rest of the array to service requests.

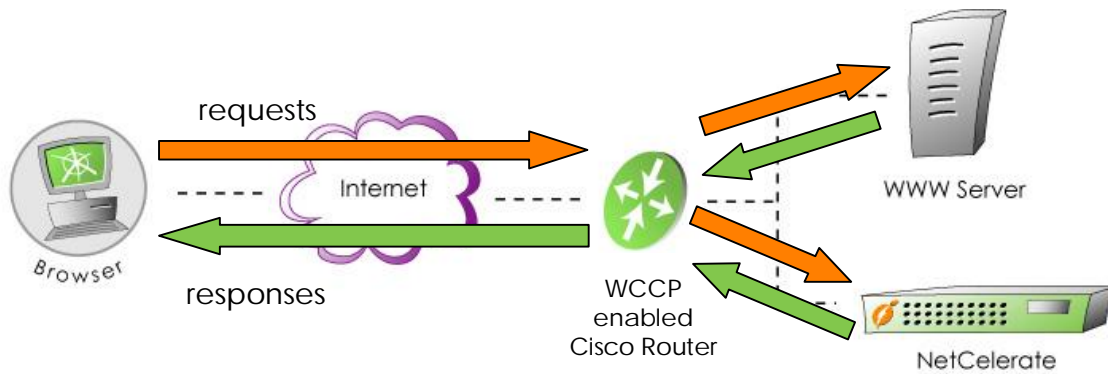


Some load balancers have cooperating daemon processes that can run on the servers being managed. These daemon processes report on specific attributes of each server to improve the load balancing capability and provide server management feedback. Without these processes, most load balancers still work by accounting for the amount of traffic they distributed. SiteCelerate does not have daemon processes for any specific load balancer and must be managed as a "black box" server.

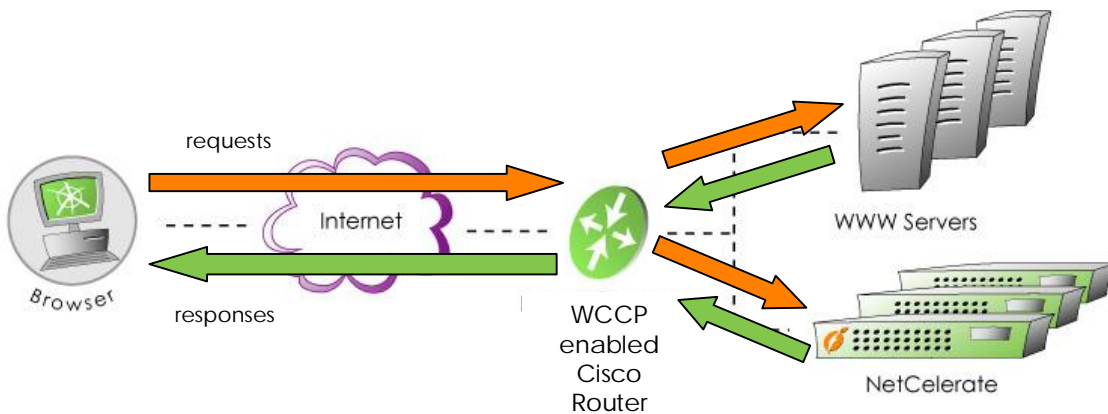
Transparency via WCCP

Many networks have Arachne routers capable of using the Web Cache Communication Protocol (WCCP). WCCP allows an Arachne router to collaborate with proxy systems to selectively distribute Web requests, manage load, and handle the event of a system going offline. Using a WCCP enabled router is often simpler than using a switch or load balancer from another vendor because the rules are already present in the router. But, the end result is the same.

For a single SiteCelerate box, the router will forward external requests to the accelerator and forward requests from the system directly to the Web server. In the very unlikely event that the SiteCelerate server goes offline, requests are routed directly to the Web server.



WCCP enabled routers can collaborate with and manage multiple SiteCelerate servers. In this case, the Arahe router interacts with the systems to discover how many are working together and to distribute the inbound requests evenly between them. In the event that one system stops responding, it redistributes the load between the remaining SiteCelerate machines in the array.




This configuration is ideal for accelerating any number of destination sites without having to explicitly configure each SiteCelerate machine or a switch about which Web addresses are to be accelerated. Every destination site managed by the router can be accelerated by the SiteCelerate array.

SSL Acceleration

Although SiteCelerate has SSL acceleration functions built-in, many enterprises may already be using SSL accelerators in their networks to offload the SSL connection management and encryption from the main Web and application servers. Because encryption randomizes the payload data in the TCP packets, Web content optimization must happen before the content is encrypted.

When dealing with an existing in-place SSL accelerator, one or more SiteCelerate systems can be positioned between the SSL accelerator and the Web server to optimize the content as it flows from the Web server and before it is encrypted by the SSL accelerator for delivery over the open Internet.



As with the non-SSL network models described previously in this document, an array of SiteCelerate systems can be used to scale beyond the capacity of one machine. Load balancers can be used to manage a cluster of SiteCelerate boxes placed between a cluster of SSL accelerators and before a cluster of Web servers.

Conclusion: SiteCelerate = Complete Application Delivery Solution

SiteCelerate is a complete enterprise application delivery solution for all Web-based (HTTP/HTTPS) HTML and XML enterprise applications, portals and websites. SiteCelerate offers the class-leading and unique technologies that minimize network latency, significantly reduce bandwidth usage, and offload web and application servers.

SiteCelerate delivers immediate impact on any application and integrate smoothly with any Web or application servers. Arahe SiteCelerate can be rapidly deployed in any environment to dramatically accelerate web application delivery to ensure split-second response time web applications and portals.

About Arahe Solutions

Arahe Technology is a privately held web technology company. Arahe has a suite of Web content optimization products and services that reduce image bulk and speed page loading—thereby lowering operating costs and improving customer satisfaction. Arahe has had numerous enterprise customers and has accelerated millions of Web pages. For more information about Arahe, please visit <http://www.arahe.com/>