

## A Performance Solution for Document Distribution: Enabling Real-Time Document Distribution for Global Enterprises

**Overview** Web enabling applications has been a tremendous success for today's enterprises, offering enhanced user productivity, lower costs, and greater overall efficiencies. But the rush to the web has brought to light all new performance issues. At first, fixing these problems was a relatively simple affair—tune the application, add hardware, or purchase additional bandwidth for the corporate backbone network. But as user populations grew, especially via new remote users outside of the corporate LAN, additional problems emerged that could not be solved by the old methods.

**Challenge** Traditionally, project collaboration within the enterprise was a simple affair. Documents were shared over a LAN locally, with remote partners included via e-mail and/or fax. During the past decade however, the adoption of more sophisticated document versioning, change control, and workflow systems radically changed how project collaboration was accomplished. Although these advancements improved overall enterprise productivity, they created an insatiable need for global access from both web and client/server environments.

Many different solutions from a variety of vendors soon emerged to meet the demand, and the “best” solution really depended on the specific needs of the enterprise itself. For enterprises with large collaboration projects involving remote participants, suppliers and customers, web-enabled applications were the answer. All members of a given project could access documents in real time and collaborate as if they were in the same location. The trouble is, this put huge demands on the enterprise infrastructure, and many users (especially those on a WAN) found the promise of the Internet to be unfulfilled—both application and network performance were unacceptably slow due to a variety of bottlenecks within the system.

These bottlenecks are associated with the volume of document delivery on the network, and are exacerbated in collaborative environments where fast response time is a requirement. Documents take considerable time to deliver because applications either do not, or cannot, deliver documents to browsers optimally. Even though browsers can accept compressed documents in zip form, web based applications typically do not compress documents before delivery due to unpredictable incompatibilities on the user's client machine. Furthermore, even if a client application will accept a compressed document, the standard compression accepted by browsers is not effective on documents with significant embedded image content, as images are already compressed.

In addition, network resources are wasted on documents that are repeatedly accessed. Organizations routinely solved the “repeat access” problem by deploying simple edge caches in remote offices that, in turn, delivered documents directly from the cache in the remote office and eliminated retransmission over the WAN. However, in collaborative environments, edge caches don't offer significant benefits for variety of reasons:

- (1) Many documents require authentication before delivery, so they can't be served directly by the edge cache.
- (2) Edge caches lack knowledge of when documents in the cache are out of date, and could potentially deliver stale content to users.
- (3) In many web applications that serve documents, URLs for document access are not consistent—many applications generate unique URLs for the same document and the same user—so traditional caching by a document URL gives no benefit.
- (4) Support for document compression in the edge cache is limited—edge caches lack the ability to compress or decompress documents depending on browser and document type—therefore requiring more time to deliver the documents to the user the first time they are pulled into the cache.

**Document Collaboration Performance Challenges**

- Document authentication and validation requirements
- Non-predictable updates
- Inconsistent URLs
- Limited compression

As a result of these limitations, enterprises are discovering that edge caching “solutions” are ineffective for collaborative environments.

**Two Approaches to Document Repositories: Central and Distributed**


Enterprise document management solutions (known as repositories) generally involve two different approaches—central and distributed. Both topologies have distinct advantages, and many corporate environments feature a mix of both types. Central systems offer the benefit of ease of control and management, while distributed systems offer superior fault tolerance and (oftentimes) improved performance through local access.

An enterprise employs a central repository in order to eliminate the hassles involved with maintaining multiple sites—a single central site is simpler and significantly less expensive to deploy and manage/maintain. Users local to the central repository realize the inherent benefits while being able to quickly access and update the documents they need. However, such an implementation can often impede remote users when it comes to those same benefits. The system ends up being slow and inefficient due to the limitations of web-based document access via a WAN.

A distributed repository infrastructure employs multiple locations to share the burden of repository duties, and there are many good reasons to distribute repositories across the enterprise. Corporate business practices or local regulations may require that certain documents be maintained in remote offices. In addition, in order to minimize WAN traffic and optimize performance, documents that are most frequently accessed or updated by an office will be replicated in a repository local to that location. However, replicated repositories are inherently complex, costly to deploy, and typically require significant ongoing IT support.

Regardless of the approach, enterprises usually deploy a document repository application such as a web-enabled Document Management System (DMS) or a robust enterprise portal to manage workflows, signoffs, version control, notification, common document storage and similar tasks. An enterprise portal enables a company to store and retrieve data for such activities as problem solving, dynamic learning and strategic planning. Portals have become an important component of the enterprise infrastructure as they provide a rich collaborative environment as well as an advanced means to organize corporate knowledge.

Yet as advanced as a DMS and portal systems are, they generally treat network factors as external



and therefore do not try to resolve problems linked with slow access performance.


**Solution F5 WebAccelerator Technology: LAN Performance That Leverages the Existing Enterprise Network Infrastructure**

The F5 WebAccelerator and F5 WebAccelerator Remote offer LAN-like performance to remote users of web-based document repositories, whether the enterprise employs a topology consisting of a central repository, distributed repository, or a combination of both. Unlike other solutions that improve project collaboration performance, F5 WebAccelerator does not require fundamental alterations to the network topology. The technology offers a highly flexible framework that enables distributed users to efficiently access web-based document repositories. Two prime deployment modalities are supported: asymmetric and symmetric.

- With asymmetric deployments, F5 WebAccelerators are installed in the data center in front of the DMS or portal application serving documents from the repositories.
- With symmetric deployments, F5 WebAccelerators are installed at both ends of remote segments, in front of document repositories and in remote offices.

Both deployments combine F5 WebAccelerator’s patented and patent-pending Application Smart Caching server-side technology and Express network-side technologies to overcome traditional distributed document access bottlenecks, offering tremendous advantages to the enterprise:

- 1) Optimized document delivery—The servers’ Express technologies speed delivery of documents by reducing the effects of network latency and limited bandwidth, thereby optimizing documents for delivery over the network as well as leveraging browser and edge caches to accelerate repeat views of the same documents.
- 2) Document authentication and authorization support—the F5 WebAccelerator relies on the document repository application as the sole arbiter of authentication and authorization policies, and no changes are necessary to existing repository applications. All document requests are forwarded to the application. If the request is valid and the content has not changed, the document will be served from the F5 WebAccelerator or browser cache.
- 3) Document URL aliasing—When the same document is accessed by multiple URLs, the F5 WebAccelerator will recognize that fact and deliver the document from its cache if possible, eliminating redundant transmission of the document over the network.
- 4) Improved application scalability and performance -- When application scalability and server latency are of real concern, F5 WebAccelerator Application Smart Caching can stage some or all documents in its cache, eliminating the need to load the application down with repeated requests.



F5 WebAccelerator asymmetric deployment

An authentication plug-in architecture exists to validate document requests before delivering them directly from the F5 WebAccelerator.


**Symmetric F5 WebAccelerator deployments offer the following additional benefits:**

- (1) Remote office document staging—Once initially accessed, documents are automatically staged in remote offices and delivered with LAN performance. Documents are only transmitted over the WAN when they are new or changed. All document requests are forwarded to the application for access checks, and if the request is valid it will be served from the WebAccelerator cache in the remote office.
- (2) Optimized document compression—During initial document views, servers transmit documents in highly compressed forms between each other: different compression techniques can be applied to documents based on document type. This is particularly beneficial for WANs with low to medium bandwidth.
- (3) Optimized WAN network performance—Central and remote F5 WebAccelerator’s maintain persistent connections between them, thereby eliminating the effect of TCP “slow start”, reducing connection churn, and optimizing overall WAN throughput.
- (4) Static edge caching—Static objects associated with the DMS or portal are staged at the edge server. This eliminates WAN requests for these frequently used objects.

**F5 WebAccelerator Technology Benefits Summary**

Asymmetric Deployment Benefits	Symmetric Deployment (Additional Benefits)
Optimized document delivery	Remote office document staging
Document authentication and authorization	Optimized document compression
Express document URL Aliasing	Optimized WAN network performance
Improved application scalability and performance	Static edge caching

This flexible deployment framework enables enterprises to keep their existing mix of central and/ or distributed document repositories as the F5 WebAccelerator will match the existing business processes—there is no need to retrofit the network in order to accommodate the WebAccelerator system.



As an example, consider a large pharmaceutical corporation with main headquarters in San Diego and two remote R&D centers in London and Singapore. The main headquarters hosts a central collaborative portal that is accessed globally, especially by the Singapore R&D center. The Singapore location, in turn, hosts a web-based document repository of its own for business and regulatory reasons. Documents in Singapore are also accessed globally. And there is a large base of users in other global remote offices that access both repositories.

This diagram shows a suggested deployment of F5 WebAccelerators at all three locations of the company. At headquarters, the F5 WebAccelerator sits in front of the portal and works in symmetrical fashion with the servers in Singapore and London in order to accelerate document



**Conclusion** When it comes to project collaboration, modern enterprises are caught off guard by the increasing demands placed on their infrastructures created by the proliferation of the Internet and advanced intranets/extranets. And in their rush to meet these demands they inadvertently create even more bottlenecks that threaten their ability to successfully compete. Many solutions that resolve one problem ultimately create a greater problem elsewhere.

F5 WebAccelerator is the only acceleration solution that resolves all document repository network access problems, whether local or remote, irrespective of the enterprise environment. The result? All participants in enterprise collaborative projects realize unparalleled performance and access. Bottlenecks are permanently removed and companies are finally able to achieve the performance and productivity that the enterprise systems were supposed to deliver in the first place.

For more information on F5, document repository acceleration technologies, and the F5 WebAccelerator, go to [www.f5.com](http://www.f5.com).

**About F5** F5 Networks is the global leader in Application Delivery Networking. F5 provides solutions that make applications secure, fast and available for everyone, helping organizations get the most out of their investment. By adding intelligence and manageability into the network to offload applications, F5 optimizes applications and allows them to work faster and consume fewer resources. F5's extensible architecture intelligently integrates application optimization, protection for the application and the network, and delivers application reliability – all on one universal platform. Over 9,000 organizations and service providers worldwide trust F5 to keep their applications running. The company is headquartered in Seattle, Washington with offices worldwide. For more information go to [www.f5.com](http://www.f5.com).