



Overcome All Application Performance Bottlenecks

Organizations depend on applications to support business operations and drive revenue. At the same time, users are demanding more from those applications, including faster load times and access across a wide variety of devices. Poor performing applications can result in reduced employee productivity, abandoned shopping carts, and missed recovery point objective and recovery time objective (RPO/RTO) targets.

F5® BIG-IP® Application Acceleration Manager™ (AAM) combines the application delivery features previously available in BIG-IP® WAN Optimization Manager™ (WOM) and BIG-IP® WebAccelerator™. BIG-IP AAM overcomes network, protocol, and application issues to help you meet application performance, data replication, and disaster recovery requirements presented by cloud, mobile applications, and video distribution. By offloading your network and servers, BIG-IP AAM decreases the need for additional bandwidth and hardware. Users get fast access to applications, and you gain greater revenue and free up IT resources for other strategic projects.

Key benefits

Improve the user experience

Improve end user experience through multiple optimization technologies applied at all layers of the application delivery chain.

Optimize data center efficiency

Consolidate devices and services to deliver optimized apps. Reduce the application load from servers and the network by offloading CPU-intensive processing tasks.

Streamline Application Delivery Optimization

Quickly support and optimize legacy and emerging protocols/standards (SPDY, FTP, UDP, HLS). Optimize delivery of any application content to any device without recoding apps.

What's Inside

- 2 Integrated Application Delivery Optimization
- 2 Data Center Optimizations
- 3 Transport Optimizations
- 4 Application Delivery Optimizations
- 7 F5 Application Ready Solutions
- 8 Flexible Deployment Options
- 10 The Power of the BIG-IP System
- 12 The BIG-IP Application Acceleration Manager Architecture
- 12 BIG-IP Application Acceleration Manager Platforms
- 12 Virtual Platform
- 13 F5 Services
- 13 More Information

Integrated Application Delivery Optimization

BIG-IP Application Acceleration Manager is built natively on the F5 TMOS® unified architecture, enabling the integration of application delivery with web performance and WAN optimization technologies. This enables traditional acceleration technologies like SSL offloading, compression, caching, and traffic prioritizing to combine with technologies like image optimization, video delivery optimization, and byte-level data deduplication, thereby reducing complexity in your data center.

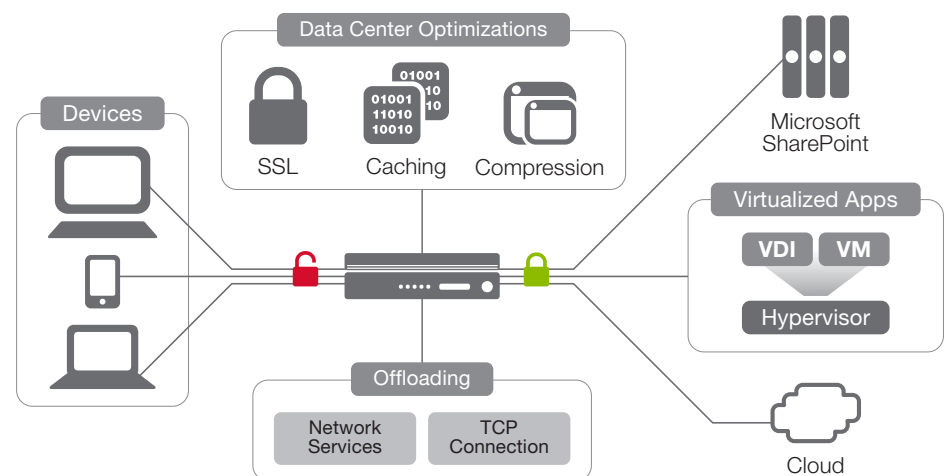
BIG-IP AAM makes use of the F5 iControl® application programming interface (API) and F5 iRules® scripting language capabilities, giving you unprecedented flexibility and control in scaling, managing, and optimizing your BIG-IP system.

BIG-IP AAM can optimize a wide variety of protocols delivered to either a client browser, desktop application, or another BIG-IP device, depending on the deployment. Optimizations are divided into data center optimizations, including server and network optimizations, transport optimizations, and application delivery optimizations, including application protocol and web performance optimizations.

Data Center Optimizations

BIG-IP Application Acceleration Manager optimizes the data center to help with the ever-changing demands on IT infrastructure, such as large amounts of data, including videos, and the use of mobile devices. Data center optimization can help to reduce application load from servers by offloading CPU-intensive tasks like encryption, caching, and compression and reduce bandwidth by sending less data over the network. The end result is a more efficient infrastructure.

BIG-IP AAM can improve the performance of WAN application traffic by optimizing application protocols, prioritizing traffic, optimizing TCP from clients to servers, and reducing the amount of data sent over the WAN, helping to prevent costly bandwidth upgrades. Quality of service (QoS) technologies ensure that critical or time-sensitive applications receive priority over others to maximize performance over the WAN. They provide granular control of traffic based on enterprise needs, enabling you to manage and prioritize bandwidth per application and improve QoS for critical applications over the WAN.



BIG-IP Application Acceleration Manager optimizes the data center.

Symmetric data deduplication

With symmetric data deduplication, BIG-IP Application Acceleration Manager delivers a highly advanced level of WAN optimization. This provides significantly more bandwidth for applications and effectively expands WAN capacity to improve response times and increase throughput. Redundant data is no longer transferred across the network through the use of pattern matching and byte caching technologies. Symmetric data deduplication ensures high-speed application performance and reduces the amount of data transferred over the WAN by up to 99 percent.

Solid state drives deduplication

Data duplication can be done in memory or hard drive disks. Typically, memory-based deduplication is recommended due to the slow I/O performance of standard hard drives. However, for large volumes of data, deduplication using solid state drives (SSD) can have up to a three time improvement in replication time over memory-based deduplication.

BIG-IP Application Acceleration Manager running on the BIG-IP 11000 platform, with support for four 600 GB SSDs, is the ideal choice for the high-volume requirements of data center to data center replication or as the head end of a hub and spoke deployment.

SSL acceleration

BIG-IP Application Acceleration Manager offloads computationally intensive SSL encryption and decryption, reducing server processor utilization by up to 50 percent. It consolidates private key creation and storage, SSL certificate management, and FIPS SSL support. BIG-IP AAM standalone devices run on the F5 TMOS operating system and include the maximum available TPS for that specific hardware platform.

Parking Lot

The Parking Lot feature in BIG-IP Application Acceleration Manager queues multiple requests for the same new or expired cached object, and then sends only one request to origin web server. When the object is retrieved, BIG-IP AAM responds to all the requests. This reduces the load on the servers when a flood of requests come in at once.

Transport Optimizations

BIG-IP Application Acceleration Manager improves the capacity of application servers and the efficiency of network protocols by offloading intensive processing tasks such as SSL encryption, optimizing application, and network protocols. Optimization features include the following.

Symmetric adaptive compression

Symmetric adaptive compression ensures the fastest data reduction for any traffic between BIG-IP systems. Symmetric adaptive compression automatically selects and uses the appropriate deflate, bzip2, or LZO compression algorithms (or no compression if the data cannot be compressed) to maximize bandwidth usage and throughput. In addition, symmetric adaptive compression can use BIG-IP hardware compression where available to provide unprecedented scalability.

Forward error correction (FEC)

Forward error correction (FEC) is a method for controlling errors in transmitted data over high packet loss communication channels. Data is sent in a redundant manner, enabling the receiving end to correct any potential errors or corrupted data without requiring a retransmission. FEC can be enabled between two BIG-IP devices or from a BIG-IP device to an edge client, significantly improving application performance on high packet loss networks.

HTTP protocol optimizations

BIG-IP Application Acceleration Manager maintains high user performance levels by optimally tuning each HTTP and TCP session for each user's connection conditions. Optimizations for Microsoft NTLM authentication protocol enhance access to protected resources.

Bandwidth Controller

Bandwidth Controller provides the ability to manage the amount of bandwidth a device, subscriber, or application receives. Traffic can either be enforced or marked, identifying and flagging packets that are exceeding bandwidth.

TCP optimization

When application performance suffers, IT managers often assume that adding bandwidth will solve the problem. But TCP throughput degrades significantly on the WAN, particularly on high-latency, intercontinental links, so adding bandwidth is often ineffective.

To overcome inherent protocol limitations, BIG-IP Application Acceleration Manager uses adaptive TCP optimization, which combines session-level application awareness, persistent sessions, selective acknowledgements, error correction, and optimized TCP windows. This enables BIG-IP AAM to adapt, in real time, to the latency, packet loss, and congestion characteristics of WAN links, to fully utilize available bandwidth and accelerate application traffic (for up to 20 Gbps LAN-side, TCP optimized throughput).

Application Delivery Optimizations

Application Delivery Optimization is a holistic way of looking at all the pieces in the delivery chain that need to be optimized from the transport mechanism to the application protocol. BIG-IP Application Acceleration Manager solves application delivery issues by optimizing the TCP stack and the application protocol and ensuring the best use of bandwidth.

Application performance on the WAN is affected by a large number of factors that can't be solved by adding bandwidth alone. Performance is limited by factors such as the natural behavior of application protocols that were not designed for WAN conditions, application protocols that engage in excessive handshaking, and the serialization of the applications themselves.

CIFS acceleration

Microsoft's remote file access protocol, common Internet file system (CIFS), is standard on Windows clients and servers and is commonly used to provide complete read/write access to files across data centers and branch offices. CIFS is a "chatty" protocol and not designed for high latency WAN environments. F5's CIFS acceleration provides intelligent read-ahead and write-behind plus other techniques to help mitigate the effect of WAN latency. This provides

significant reduction in transfer times and bandwidth usage, improving performance of enterprise information transferred over the WAN.

MAPI acceleration

Message Application Programming Interface (MAPI) is the email protocol used by Microsoft Exchange Server and Outlook clients. Use of symmetric adaptive compression and symmetric data deduplication dramatically improves performance and reduces bandwidth usage for customers using Microsoft Exchange, especially when sending email attachments.

HLS delivery optimization

HTTP Live Streaming (HLS) is the protocol used by a number of devices to view both live and on-demand video. HLS breaks the video down into segments that can be cached for multiple users. HLS can be optimized by caching the individual segments or by controlling the bitrate that is made available to end users.

SPDY gateway

SPDY is an emerging new application-layer protocol developed by Google that augments HTTP by improving the inefficiencies related to connection management and data transfer, with the goal of improved performance. It supports multiple streams within a single TCP connection, compresses the HTTP headers, and allows for prioritization of requests.

Because requests are interleaved on a single channel, the efficiency of TCP is much higher: fewer network connections need to be made, and fewer, but more densely packed, packets are issued. These benefits would specifically help in the mobile use case, given the typical slower mobile connection.

F5 provides a SPDY gateway in TMOS to convert SPDY requests to HTTP to backend web servers. This takes advantage of the optimizations without requiring disruptive and potentially costly upgrades to application infrastructure.


Web performance optimization

BIG-IP Application Acceleration Manager solves web content delivery issues by modifying the data and reducing the number of round trips required to fully display a web page. The result is significantly decreased download times, reduced bandwidth usage, and lower costs for using enterprise web applications in remote office and mobile deployments.


Mobile users face additional challenges due to the proliferation of different types of mobile devices, from smartphones to tablets, which have different operating systems and browsers. The additional latencies due to the extra hop from cell towers and WiFi hotspots make matters worse. Users end up with a range of page download times, all which are typically worse than what users get at the office or home.

To resolve these performance issues, BIG-IP Application Acceleration Manager uses a number of techniques to improve the end user experience. These optimizations do not require any server side installations, client side software, or changes to users' browsers.

Before F5



With F5



Application Delivery Optimization reduces the number of round trips required to deliver a web application.

Dynamic Content Control (DCC) is a group of capabilities in BIG-IP Application Acceleration Manager that control users' browser behavior to improve end user experience, ensure the best use of bandwidth, and prevent repetitive or duplicate data from being downloaded. By reducing the amount of conditional requests and data transmitted between the browser and the web application, DCC reduces the effects of WAN latency and errors.

DCC includes these main features:

- **Intelligent Browser Referencing™**—Reduces the number of requests and speeds page rendering times by managing object expiration dates and storing frequently requested objects in the browser cache. Ensures that the browser only downloads truly dynamic and unique content by eliminating the download of repetitive data and browser conditional requests for static data that is incorrectly considered dynamic.
- **Content reordering**—Optimizes the order of when JavaScripts and Cascading Style Sheets (CSS) are loaded to speed up the appearance of page rendering.
- **Content inlining**—Reduces the number of requests by inlining JavaScripts, CSS, and images directly into HTML, eliminating the need to perform additional GET requests. This optimization is beneficial for content that will be viewed only once or for mobile devices that have limited cache sizes.
- **MultiConnect**—A form of domain sharing that enables browsers to open more simultaneous connections between the browser and web application for increased parallel data transfers. MultiConnect is extremely effective on high latency/high bandwidth networks such as satellite and mobile networks.
- **Dynamic linearization**—Enables users to display PDF pages or jump to specific pages and view them without having to wait for the entire document to download first.

Dynamic Data Reduction (DDR) reduces bandwidth utilization and improves page load times by reducing the amount of data that needs to traverse the WAN or Internet. F5 BIG-IP Application Acceleration Manager offers the following DDR functions:

- **Image optimization**—Reduces size of images by lowering the quality, stripping out unnecessary metadata, and converting the image format. For mobile devices, this optimization can be more beneficial given the smaller screen sizes and slower mobile connections.

- **Minification**—Removes white space and comments from JavaScripts and Cascading Style Sheets, reducing the size of the files. Useful for situations where compression cannot be performed.
- **Dynamic caching**—Caches data that may seem dynamic (contains query parameters, cookies, or session IDs) but is actually static data or changes in an identifiable pattern. By fully inspecting every aspect of HTTP requests, controlling caching behavior, and invalidating cached data, BIG-IP Application Acceleration Manager caches a high percentage of data from dynamic web applications while maintaining proper application behavior. BIG-IP AAM cache can scale up to 1 TB, depending on the hardware platform.
- **Dynamic compression**—Compresses dynamic data from web applications and ensures that compression is used only when it will improve performance. Dynamic compression is different from standard compression implementations because of its high efficiency and its ability to avoid widespread browser compression bugs. Even dynamic content requiring unique session IDs within every link on the page can be delivered and compressed, often with zero compression overhead.

F5 Application Ready Solutions

F5 works with some of the world's largest software vendors to bring you F5 Application Ready Solutions, a complete set of resources that simplifies the design, deployment, and management of your applications across the network. F5 Application Ready Solutions are designed, engineered, tested, and documented with BIG-IP Application Acceleration Manager—along with F5's integrated product line—in a variety of real-world environments.

F5 Application Ready Solutions reduce the time, money, and errors associated with deploying and managing mission-critical, enterprise applications. Only F5 offers this comprehensive set of essential, application-specific tools.

Application acceleration policies

Pre-defined, validated web acceleration policies enable you to quickly configure and deploy BIG-IP Application Acceleration Manager to optimize your application acceleration right from the start. These policies can be used as built-in templates to enable you to customize BIG-IP AAM for your specific web applications.

Validated web application acceleration policies that ship with BIG-IP AAM include Microsoft SharePoint, Oracle Portal, SAP Portal, Microsoft Office Outlook Web Access, Oracle E-business Suite 11 and 12, Oracle Siebel CRM, and many more. Generic policies are also available for custom and less common applications that do not have a pre-defined policy. BIG-IP AAM configurations and policies can also be managed and updated using F5 iApps® templates.

Application Ready Solution guides

Each specific Application Ready Solution guide provides a comprehensive overview, details how to ease your application deployment, and shows you the specific results you can achieve with your BIG-IP Application Acceleration Manager implementation.

Deployment guides

Detailed, step-by-step procedures walk you through deployment from day one. Every procedure has been thoroughly tested and optimized in real-world environments to achieve top performance. Each deployment guide contains a comprehensive set of configuration scenarios to cover your specific needs.

Active user community

An active, collaborative community on F5 DevCentral™ offers feedback, documents, and tips for a successful deployment. Dedicated Application Ready Solution pages provide application-specific content, including downloads, help and forum discussions, links to related podcasts, and more.

Flexible Deployment Options

BIG-IP Application Acceleration Manager can be deployed in multiple modes to suit your existing infrastructure and network topology, and to simplify deployment.

Core and advanced acceleration options

BIG-IP Application Acceleration Manager Core offers acceleration as a core component of BIG-IP® Local Traffic Manager™ (LTM). Compression, Bandwidth Controller, F5 iSession® network tunneling, and SPDY gateway capabilities are available as part of every BIG-IP LTM platform. The full BIG-IP Application Acceleration Manager product provides advanced application protocol optimizations.

Cost-effective asymmetric deployment

BIG-IP Application Acceleration Manager can be placed in the data center in an asymmetric deployment to achieve performance improvements of two to five times. In addition, deploying in a remote site for caching offload can speed up local requests for specific recurring high volume data and applications. Unique to BIG-IP AAM, asymmetric web acceleration offers immediate, significant return on investment (ROI) for a moderate investment.

Asymmetric topologies can be either inline or one-armed. When deploying in an inline topology, BIG-IP Application Acceleration Manager is installed in the data path behind the WAN router, in either a routed or bridged configuration.

With one-arm mode using policy-based routing (PBR), BIG-IP Application Acceleration Manager can be deployed to optimize traffic based on specific policies on the router, making this deployment method extremely flexible for application needs. One-arm mode using the Cisco-developed Web Cache Control Protocol v2 (WCCP) and other methods can be used to deploy BIG-IP devices with a single connection to a switch or router. With WCCP support there is no need to change network topology.

Symmetric deployment for maximum acceleration

Symmetric deployments can provide acceleration of up to 10 times over unaccelerated applications. In a symmetric implementation, BIG-IP Application Acceleration Manager is deployed at the data center and at one or more key remote locations or data centers. By serving unchanged content directly from the remote device, symmetric acceleration further eliminates the effects of high latency connections. The result is maximum performance

acceleration and additional decreases in bandwidth usage. The role of the device whether it is fronting an application or remote from an application can be configured on a per application basis not on a per device basis.

As the foundation for site-to-site communication, the F5 iSession network tunneling feature secures and accelerates data traveling over the WAN. Any two BIG-IP devices can be deployed symmetrically to create a site-to-site secure connection to improve transfer rates, reduce bandwidth, and offload encryption for more efficient WAN communication. Through iSession, all data can be symmetrically encrypted between two BIG-IP devices using either SSL or IPsec, providing site-to-site data security. SSL throughput is based on the level of your BIG-IP hardware platform.

Clustering to scale

BIG-IP Application Acceleration Manager devices can be clustered to create very large arrays to scale capacity as your web application acceleration needs grow.

Creating a private content delivery network (CDN)

Many organizations may choose not to use commercial content delivery network (CDN) providers because their content is internal, dynamic, and confidential or they do not want to pay the recurring costs. Deployed symmetrically in conjunction with other F5 solutions, BIG-IP Application Acceleration Manager enables your organization to create its own private enterprise CDN. This provides your enterprise websites with high availability and performance, content control, and denial-of-service (DoS) attack protection. It can also help you reduce OpEx costs and meet regulatory compliance.

E-commerce stand-in capability

When e-commerce web servers go down, BIG-IP Application Acceleration Manager can ensure high availability by “standing in” and continuing to serve static content that is already cached. BIG-IP AAM can prevent lost or abandoned shopping carts and hand off to financial transaction servers for processing.

Product module or standalone solution

BIG-IP Application Acceleration Manager is available as a product module on BIG-IP Local Traffic Manager or as a standalone solution on any of the hardware appliance platforms.

Acceleration and security in one

You can accelerate and secure web applications by running BIG-IP Application Acceleration Manager, BIG-IP® Access Policy Manager®, and BIG-IP® Application Security Manager™ concurrently on the same BIG-IP device. This saves the cost of extra hardware, rack space, and energy consumption, while simplifying deployment through consolidated and centralized access to the management interface.

Dynamic discovery

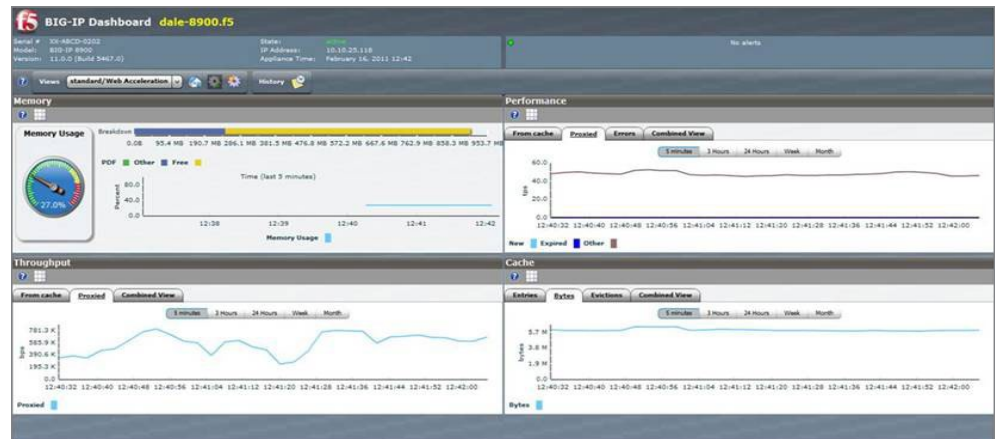
BIG-IP Application Acceleration Manager drastically reduces configuration time by discovering remote BIG-IP device peers and the networks that they serve. Once a remote BIG-IP device has been discovered and a secure connection is established, the BIG-IP device then updates available networks for WAN optimization. Servers and clients that communicate across the WAN can be added or removed, without having to reconfigure the BIG-IP devices.

The Power of the BIG-IP System

BIG-IP Application Acceleration Manager, as part of the BIG-IP system family, includes the following features.

Performance dashboard

The performance dashboard offers a detailed “on-box” monitoring and reporting tool, giving administrators a quick look at real-time data, performance, and bandwidth gains for traffic optimized with BIG-IP Application Acceleration Manager. The easy-to-use GUI provides a faster, intuitive way to find the information you need: historical statistics, log based alerts, remote peer status, health statistics, and more.



The performance dashboard provides real-time data for traffic optimized with BIG-IP Application Acceleration Manager.

F5 TMOS plug-ins

Native integration with TMOS plug-ins gives BIG-IP Application Acceleration Manager faster performance and better stability under high load. This full compatibility with BIG-IP Clustered Multiprocessing (CMP®) enables it to run on multi-core systems.

iRules flexibility

F5 iRules, a TCL-based scripting language to control the behavior of BIG-IP devices, can be used with BIG-IP Application Acceleration Manager. An example is using an iRule to eliminate round trips due to URL redirection. The iRule would detect URL redirects and serve the “final” URL content, reducing the additional round trips from browser to web server.

NTLM authentication support

The NTLM authentication protocol requires frequent re-authentication with the application server and can significantly affect web application performance. Native NTLM authentication optimization is now part of the TMOS OneConnect™ feature, which enables greater performance scalability when accelerating NTLM-enabled web applications.

Resource provisioning

BIG-IP Application Acceleration Manager resource provisioning automatically allocates CPU, memory, and disk space for the modules licensed on the BIG-IP system, based on the provisioning options chosen. This makes optimal system resource allocation easier, and an enhanced UI provides graphical representation of the allocations. Often BIG-IP modules can be enabled without requiring a system reboot.

Evaluation licensing

For existing BIG-IP customers, this feature enables customers to evaluate BIG-IP Application Acceleration Manager and other BIG-IP product modules without needing to re-license the BIG-IP device.

Logical Volume Manager (LVM)

Unlike normal disk storage, Logical Volume Manager (LVM) virtualizes physical disks into logical volumes that allow disk partitions to be resized as needed without having to reinstall TMOS or requiring system downtime in order to migrate data to a larger disk partition. The result is increased flexibility and improved performance for BIG-IP Application Acceleration Manager disk-based caching.

The BIG-IP Application Acceleration Manager Architecture

Running as a module on BIG-IP Local Traffic Manager or as a standalone appliance, BIG-IP Application Acceleration Manager uses F5's unique, purpose-built TMOS operating system. TMOS is an intelligent, modular, and high-performing full proxy operating system that optimizes, secures, and accelerates your web applications.

BIG-IP Application Acceleration Manager Core features include:

- Symmetric adaptive compression
- SPDY gateway
- Bandwidth Controller
- Dynamic compression
- Caching
- Compression
- TCP Express
- OneConnect
- Dynamic caching/deduplication
- Multi-protocol optimizations (HTTP, FTP, MAPI, UDP)
- Forward error correction
- Parking Lot (GET request queuing)
- MultiConnect
- PDF Dynamic Linearization
- Pre-defined and generic acceleration policies for ease of configuration
- Performance dashboard
- Flexible deployment (symmetric and asymmetric)

BIG-IP Application Acceleration Manager features include:

- Intelligent Browser Referencing (IBR)
- Image optimization
- Content reordering
- Scalable clustering
- E-commerce stand-in capability
- BIG-IP APM, ASM, and AAM layering
- iApps support

BIG-IP Application Acceleration Manager Platforms

BIG-IP Application Acceleration Manager is available on hardware appliances or VIPRION® modular chassis and blade systems designed specifically for application delivery. F5 systems enables simple on-demand scalability as your Application Delivery Network grows. See the BIG-IP System Hardware and VIPRION Datasheets for specifications and details.

Virtual Platform

BIG-IP Application Acceleration Manager Virtual Edition (VE) offers the flexibility of a virtual software solution for web performance optimization. Running on your choice of hypervisor and hardware, BIG-IP AAM VE can help you meet the needs of your virtualized environment in the data center or at remote sites.

F5 Services

F5 Services offers world-class support, training, and consulting to help you get the most from your F5 investment. Whether it's providing fast answers to questions, training internal teams, or handling entire implementations from design to deployment, F5 Services can help you achieve IT agility. For more information about F5 Services, contact consulting@f5.com or visit f5.com/services.

More Information

To learn more about BIG-IP Application Acceleration Manager, use the search function on f5.com to find these and other resources.

White paper

[Application Delivery Optimization](#)

Blogs

[Programmable Cache-Control: One Size Does Not Fit All](#)

[Random Acts of Optimization](#)

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 www.f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com

