White Paper

Top Considerations When Choosing an ADC

The Application Delivery Controller (ADC) chosen as the foundation for a flexible, efficient application delivery strategy can significantly affect network performance, availability, and security while providing opportunities to add business value.

by Lori MacVittie
Senior Technical Marketing Manager

Contents

Introduction

Driven by financial and operational pressures to add value, take advantage of cloud computing, improve security, and address concerns such as performance and availability, IT organizations frequently must evaluate the data center infrastructure for ways to meet goals that may seem mutually exclusive. In addition, as cloud computing and consumerization transform the traditional server (application) tiers into mobile, virtualized containers, applications and servers are abstracted from the infrastructure, severing their easy integration with the systems typically used to provide for availability, performance, and security. As a result of these trends, more organizations are recognizing a critical fourth tier within the data center architecture—a flexible and highly scalable tier in which application delivery concerns such as security, performance, and availability can be readily addressed.

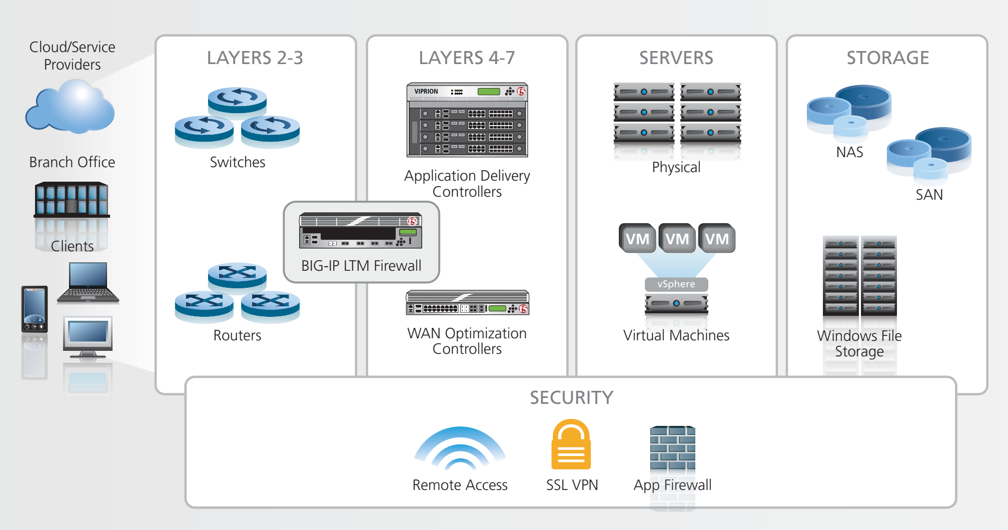

Figure 1: The application delivery tier delivers a flexible, dynamic operational platform upon which security, performance, and availability can be efficiently managed.

The application delivery tier is based on an Application Delivery Network, a set of services that address and mitigate the operational risks imperiling the successful deployment and delivery of applications. At the heart of the Application Delivery Network is the Application Delivery Controller (ADC).

While most often associated with load balancing—a core technology used to address availability and performance issues—the modern ADC has evolved to include features and functionality that span the three primary operational risks IT departments must address daily. No longer merely load balancers, ADCs today provide services to mitigate security threats, ensure availability, and improve performance within the data center and into the cloud.

Because of its strategic location in the data center network, the selection of an ADC should involve careful consideration of both functional and financial factors. While financial considerations are relatively obvious and intuitive, the core functional considerations may ultimately have longer-lasting effects on the IT department's ability to design and deliver the infrastructure services required of applications today—and tomorrow.

ADC Functional Selection Criteria

Potential ADCs should be evaluated against a variety of functional criteria when being considered as the core of an infrastructure's application delivery tier. The most important criteria may be grouped in categories with overlapping implications for performance, scalability, and security:

- Performance
- Scalability
- Visibility
- Security
- Manageability
- Flexibility

Each entails metrics and perspectives that go beyond traditional definitions of hardware or network infrastructure performance measures.

Performance

Performance has long been a primary concern of both IT practitioners and executives. This concern generally begins with applications that support and drive the business and ultimately extends into the supporting infrastructure. Because an ADC is logically deployed between end-users and applications to provide delivery services, its performance is a critical consideration.

Performance traditionally has been measured in terms of speeds and feeds, in line rates and packets per second. This type of measurement is appropriate for packet-processing devices such as routers and switches, but it fails to accurately represent the performance of infrastructure components that are primarily concerned with the delivery of applications. Connection capacity and decisions per second are paramount to understanding the ability of an ADC to support the performance of modern applications and infrastructure, as it is often one of these two factors that becomes a bottleneck, ultimately impairing the performance of applications.

Connection capacity

A variety of factors influence the connection needs of today's applications, and all are driving up the number of connections necessary to meet demand while maintaining performance. Connection management is one of the most common causes of application performance problems, as managing connections requires not only consumption of resources that then cannot be used to process requests but also consumes additional time per request as the application searches longer and longer lists of connections to identify the most efficient one.

Performance is therefore directly related to connection management on web and application servers. An ADC can mitigate this issue by mediating for the servers and limiting the number of connections that must be opened on the server without negatively affecting the number of concurrent users that can be served. This means the ADC must be able to sustain voluminous numbers of connections without negatively affecting performance. Connection capacity becomes even more critical as application-layer attacks increasingly bypass traditional security measures and threaten the infrastructure with an overwhelming number of connections.

F5® ADCs provide outstanding connection capacities. F5 BIG-IP® Local Traffic Manager™ (LTM) can handle up to 192 million concurrent connections and 320 Gbps of throughput, managing them with various timeout behaviors, buffer sizes, and other security and application options. This capacity enables BIG-IP LTM to manage the volume of a traffic onslaught—whether incurred by an attack or a flash-crowd of seasonal visitors. The ability of BIG-IP LTM to handle such a vast number of connections while maintaining superior performance is due to its unique internal architecture, which was designed and developed to ensure optimal performance and capacity.

Transactions per second

Given behavioral changes in applications—such as identifying mobile devices and the increasing use of Asynchronous JavaScript and XML (AJAX) to deploy real-time updates—as well as the increasing reliance on APIs for mobile applications, it is important to evaluate an ADC on its ability to make decisions at an acceptable rate. Even simple uses of an ADC, such as application switching or layer 7 routing to simultaneously manage mobile and traditional versions of web applications, can have a profound effect on overall performance.

Consider the difference between a simple HTTP request and response in which the request is nothing more than a GET request paired with a zero-byte payload response, a POST request filled with data that requires processing not only on the application server but on the database, and the serialization of the response. The metrics that describe the performance of these two requests will almost certainly show that the former has a higher capacity and faster response time than the latter. The same is true of an ADC that must perform page routing, determine access permission, or scrub data for compliance purposes.

Performance tests that measure only surface abilities to pass packets or open and close connections are simply not enough to understand the performance of an ADC. It is important to compare ADC options from the perspective of decisions-per-second rather than surface-layer protocol-per-second measures.

With no standard definition of "decision," however, or a test methodology defining which decisions should be tested, it is up to the IT organization to define criteria and evaluate the performance of ADCs configured to perform application-layer decisions. Vendor-provided test data will not suffice when vendor definitions of "transaction" vary, sometimes wildly.

SSL

The use of SSL to secure an entire web application (as opposed to only login or order pages) is a best practice increasingly adopted across the web, as demonstrated by sites such as Facebook and Twitter. While the traditional SSL metric of RSA operations per second is still valid, it is no longer the only measurement necessary to understand SSL performance. When an entire site is secured with SSL, both transactions per second (TPS) and bulk encryption rates become important performance figures to consider, particularly as the industry moves toward standardizing on 2048-bit key lengths. Measuring TPS will assist in gauging the number of users that can be supported concurrently, but bulk encryption rates will determine more accurately how much traffic can be supported without degrading performance, as bulk encryption metrics determine how much secure data can be exchanged and at what rate. The longer sessions are open, the more significance bulk encryption metrics gain compared to TPS, because the transaction overhead only happens during session setup and teardown.

This is why F5 integrates cryptographic acceleration hardware into its BIG-IP hardware platforms. While solutions without such hardware-assisted functionality can process a relatively high volume of secured connections, they do not succeed nearly as well when performing bulk encryption during the ensuing session. Because the handshaking process measured by traditional SSL TPS metrics occurs only at the beginning of a session, and a session may last for quite some time, the bulk encryption rate becomes a much bigger factor in the responsiveness—and capacity—of the application during actual use.

Failure to evaluate the connection capacity, decision speed, and encryption performance of an ADC with respect to its intended use in the data center will have financial ramifications, since performance issues may require scaling of the ADC or the application. This means additional cost regardless of the ADC form factor, as expansion of either applications or ADCs always requires some sort of hardware and thus incurs both hard operational costs and soft managerial costs.

This is particularly true of increasingly popular "pay-as-you-grow" scalability strategies that employ licensing limitations on hardware platforms. While initially appearing more cost effective, these strategies do not always provide for the scalability of all performance criteria. Licensing restrictions do not affect the underlying hardware capacity, and it is the hardware capacity and performance that are always the most constraining factors for overall performance. As hardware utilization increases, capacity degrades, albeit more slowly in some cases than in others. Consequently, scale-by-license approaches incur increasing costs per transaction.

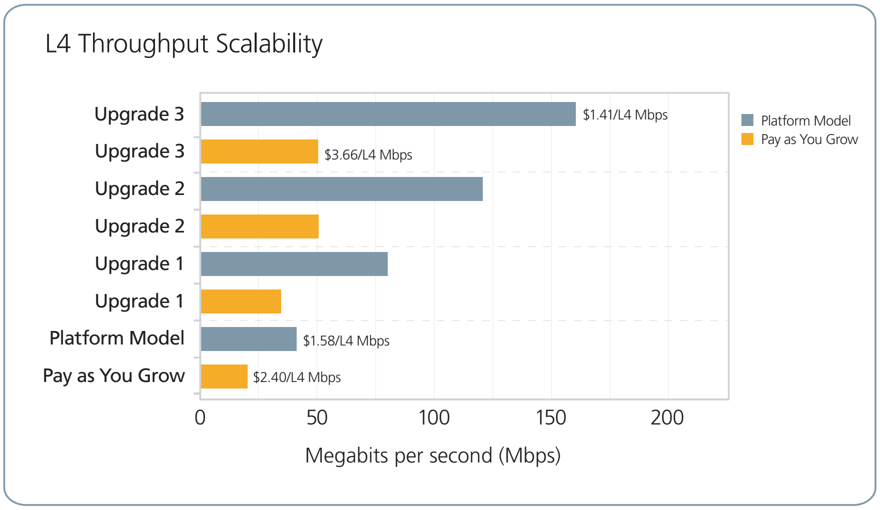

Figure 2: The layer 4 throughput of a pay-as-you-grow (licensing-based) scalability model increases at a slower rate and higher per-transaction cost than in a platform-based (hardware) scalability model.

For example, consider the case of a typical mid-sized organization that anticipates a future need to scale and compare the cost per transaction as licensing is increased in a scale-by-licensing (pay-as-you-grow) strategy with a straight platform-based scalability model in which additional hardware is added at each growth step. The cost of the pay-as-you-grow, mid-sized model is lower, starting at $48,000 and topping out at $183,000 at the third upgrade. A comparable mid-sized platform-based model begins at $63,000 and reaches $225,000 with the third upgrade. Unfortunately for pay-as-you-grow customers, performance does not increase in the same relative proportions, so the pay-as-you-grow model provides just over twice the performance after the third upgrade for more than four times the cost. By comparison, the platform-based model increases in performance faster than it increases in cost, delivering a four-factor improvement in performance at less than four times the cost.

Similarly, metrics for layer 7 requests-per-second (RPS) also exhibit uneven scalability with respect to costs for the pay-as-you-grow model, providing only 1.5 times the performance (from 1,000 to 1,467 RPS) for four times the price. The platform-based model, however, shows approximately linear gains of four times the performance (from 1,000 to 4,000 RPS) at less than four times the total cost.

This non-proportional performance scaling directly affects the cost per transaction, which is a common financial metric used to evaluate infrastructure because it directly translates into business costs and can be used to adjust pricing and facilitate expense estimation. Using the same costs and performance metrics as in the example above, the pay-as-you-grow model begins at a cost of $2.40 per L4 megabits per second of throughput, but this cost increases to $3.66 by the third upgrade, a sharply increasing trend. Conversely, the platform-based model begins at $1.58 per L4 megabits per second, which decreases to $1.41 by the third upgrade. The rising cost per transaction for a pay-as-you-grow strategy is also seen in layer 7 RPS metrics, which rise from $0.05 to $0.12 per RPS, while costs in the platform-based model remain constant at $0.06 per RPS.

This disparity is not one that is often considered up front, as it is usually the initial capital investment that is most important in the purchase decision. The oversight, however, almost always proves to be problematic, since most organizations soon need additional capacity and performance, and thus the long-term costs of pay-as-you-grow strategies result in a much poorer performance return on investment than with a platform-based model.

Scalability

Scalability is an important facet of both availability and performance. Scalability, which includes a lot of seemingly unrelated technologies, focuses primarily on increasing available resources to meet demand. Just as important, however, is the ability to failover from one application instance to another or one data center to another—or to the cloud. Failover capacity is important, as it's often critical to business continuity.

These two capabilities—failover and scaling—are more interrelated than they first might appear. They also rely on a third capability, visibility, to provide accurate, actionable data upon which an ADC can base its routing decisions and which it can share with other infrastructure components responsible for failure and application scaling tasks.

Scalability models and failover strategies

For some time now, an N+1 model has been the most prevalent choice for high availability (HA) architectures. The N+1 model assumes any number (N) of "active" components, each with an independent, dedicated secondary (standby). This is obviously inefficient, as there are always components sitting idle—unused resources.

BIG-IP LTM avoids this inefficiency, breaking out of the N+1 model by eliminating the tight coupling between the primary and standby components and allowing the primary to fail over to any available and appropriately configured component. The result is the achievement of active-active-active^{N} configurations in which all applications and services have failover and scalability support with the least amount of idle resources. BIG-IP LTM does this via the F5 Scale$^{N™}$ architecture, which is composed of two core F5 technologies:

- **F5 virtual Clustered Multiprocessing (vCMP™).** The fundamental technology making it possible to deploy individual BIG-IP LTM "guest" instances that enable fault-isolation, version independency, and on-demand hardware-layer scalability
- **Device Service Clusters.** A Device Service Cluster (DSC) is a group of two or more BIG-IP LTM devices in a trust relationship that can share resources and ensure high availability for application delivery. DSCs allow the targeted failover of application instances by defining clusters of BIG-IP LTM devices—in any form factor and across locations—to enable on-demand scalability and ensure availability. A DSC can contain devices with different application

delivery modules, allowing for flexible provision and scaling of application delivery services across the data center.

With Scale^{N}, any available component configured to be a part of the DSC can serve as a secondary for a failing component. What's more important today, however, is the ability to eliminate failover requirements at the device or component level and a move upward to failover at the application layer. Application-layer failover provides a level of fault isolation not previously offered by traditional HA architectures, which assume an "all or nothing" approach to failover; that is, if one application triggers a failover event, all applications will be affected. That's not so with Scale^{N}, which allows individual applications to failover or purposefully move between components in a configured DSC. For even more flexibility, components can be physical or virtual and need not be identical hardware or have identical configurations.

The flexibility of this new scalability model means organizations can use cloud computing in a variety of ways, including for architectures like virtualized and bridged models, to ensure resiliency and scale across and within environments. An organization with a hardware-only model in the data center can take advantage of Scale^{N} to failover or scale out using software or virtualized components, in the cloud or remote data centers, as easily as it can employ virtualized instances using vCMP technology in the primary data center. vCMP makes it possible to deploy individual BIG-IP LTM instances that enable fault-isolation, version independency, and on-demand, hardware-layer scalability.

Additionally, F5 VIPRION® hardware technology enables this ADC to scale, on-demand, without disruption when additional blades are added to its chassis. The addition of a modular performance blade causes the ADC to automatically and non-disruptively add the new resources, enabling growth without imposing additional financial and operational costs caused by the need to buy, configure, and connect additional hardware.

Automation and integration

Cloud computing and virtualization have brought to the fore important questions regarding the way in which we scale not only applications but infrastructure. In addition to the clear benefits that cloud computing and virtualization bring to companies that use them, they have driven other advances that benefit even those who don't. Among the most important advancements are those in automation and integration. Both directly or indirectly provide big efficiency gains and associated reductions in capital and operating expenditures by enabling the codification of operational processes used to deploy and scale applications. As a result, IT gains reliability and assurance of successful deployments, whether in the data center or in a cloud computing environment.

Increased automation requires integration of components responsible for scalability with the infrastructure, which are generally associated with leading virtualization and management vendors such as HP, IBM, VMware, and Microsoft. The ADC is most often responsible for scaling of applications—whether those applications are deployed in virtual containers or on physical machines. When choosing an ADC, then, it is important to consider the level of integration and support for various automation, orchestration, and virtualization solutions, particularly those in the realm of provisioning. The BIG-IP platform integrates with and supports all major virtualization platforms, and F5 has established partnerships with most leading application and management providers. These intimate strategic and technical partnerships provide a firm foundation for innovative and timely solutions with smooth integration that enables powerful automation and orchestration across environments. F5 ADCs easily support heterogeneous environments, bringing organizations the benefits of operational consistency: lower management costs, repeatable deployments across environments, and consistent enforcement of security policies. By contrast, an ADC that is specifically designed to deliver Citrix XenDesktop, for example, but not VMware View, is likely to be relegated to a niche role in the data center, with the return on that investment greatly reduced over time as the organization diversifies its application and virtualization portfolios to meet specific business and operational needs.

F5 integration partners and BIG-IP platform integration cases include:

- **IBM SmartCloud**
- **Hobsons**
- **Microsoft Virtualization**
- **VMware vMotion**
- **HP Cloud Maps**

Similarly, it's also important to consider future integration with increasingly popular methods of automation like PuppetLabs' Puppet and Opscode's Chef, which rely on scripting-based technologies. These solutions take advantage of open, standards-based APIs like F5 iControl®, as well as platform-specific scripting options such as the F5 TMSH (TMOS® Shell) command-line interface. The F5 DevCentral™ community actively participates in providing solutions specifically for these emerging DevOps toolsets[1].

Visibility

Visibility has always been important to application delivery, but its importance grows as applications become more fluid, moving not only from machine to machine but perhaps location to location. Given the increased complexity of multi-tier, multi-server, and geographically dispersed applications, the ability of an ADC to combine

[1] Automating Web App Deployments with Opscode Chef and iControl

these multiple sources of health information into single application views—and intelligently use the individual status points to control the flow of traffic—has become paramount to ensuring efficient scalability while simultaneously addressing potential performance and availability issues.

Visibility is also integral to automation and integration, enabling the automated scaling out, and back down, of applications across environments. Visibility is the key to detecting attacks and preventing their negative effects, such as the unnecessary scaling of applications and infrastructure (and the associated costs). In short, visibility is a crucial capability of an ADC that should not be treated as a checkbox item, but rather investigated fully to ensure comprehensive views and functionality.

Version 11 of BIG-IP LTM introduced F5 iApps,™ which help organizations provide and manage application delivery services from an application-centric perspective. Along with iApp Templates comes iApp Analytics, an application-centric view of performance and capacity-related data. This data provides a holistic view of performance across network, client, and server components and facilitates drilling down into a specific application and digging through data to determine where performance problems may be originating. iApp Analytics includes per URI reporting, which is critical for understanding API usage and impact on overall capacity. Being able to tie metrics back to a specific business application enables IT staff to provide business stakeholders with a more accurate operational cost, which permits more accurate ROI analysis and proactive consideration of potential growth issues before they negatively affect performance or availability.

The importance of visibility to scalability

Scalability is more than simply increasing capacity; it should be viewed as the ability to meet demand and maintain performance for an application. This includes scaling down as well as out, particularly in cloud computing environments where elasticity is a means to contain costs and enable more efficient use of computing resources. Depending on an organization's current and future cloud initiatives, it may be advantageous to include the ADC in emerging cloud computing frameworks.

F5 is the only ADC supported by CloudStack, one of the more popular open source cloud frameworks today.

Visibility of both demand (users and requests) and supply (computing resources) is necessary to ensure that the availability and performance goals specified by business and operational requirements are met and kept in proper balance. Achieving such visibility means taking advantage of technologies beyond simple network and application protocol health monitoring to establish current capacity based on application-specific parameters.

When capacity limits are approached, the ADC should not only log that event appropriately but also communicate it to provisioning management systems to ensure that capacity is increased or decreased in a timely manner to ensure availability. This means broadly supporting external systems and providing the means by which a variety of monitoring and analytical systems can access data collected by the ADC. BIG-IP LTM is the only ADC that can use information from both the server and the client to make provisioning decisions.

Organizations also should evaluate the ability of an ADC to validate application behavior; monitor connections with respect to known limits and performance characteristics based on real-time conditions; and share data with appropriate management and provisioning systems.

BIG-IP LTM monitors and evaluates the health of applications via a wide variety of mechanisms, including content-level verification to ensure correct execution. BIG-IP LTM provides even greater value when deployed in global application delivery architectures with BIG-IP® Global Traffic Manager™ (GTM), incorporating both global and local load balancing as real-time conditions for capacity, performance, and availability can be communicated across environments and application instances.

The importance of visibility to security

Visibility is a key capability not only for detecting attacks but for subsequently ensuring they do not affect application capacity. As an integral data center partner for many of the largest organizations across many vertical industries, F5 Networks has experienced firsthand the effect of modern, multi-layer attacks on its customers. An ADC is by definition fluent in application protocols, but modern attacks have expanded beyond simple exploitation of protocols and standards. Therefore, today's ADC should also be able to monitor and analyze behavior indicative of an application layer attack.

BIG-IP LTM monitors protocol, behavior, and data to detect attacks and prevent attackers from causing resource consumption that can result in outages or increased costs due to unnecessary scaling. BIG-IP LTM decodes IPv4, IPv6, TCP, HTTP, SIP, DNS, SMTP, FTP, Diameter, and RADIUS communications, enabling sophisticated analysis based on protocol as well as payload. This allows BIG-IP LTM to detect anomalies indicating an attack in progress and to take appropriate action.

Additionally, BIG-IP LTM can intercept and inspect application server responses. During an attack, servers may disclose information via a stack trace or server error conditions. By recognizing these potential sources of valuable information, BIG-IP LTM affords IT staff an opportunity to redress the possible leak through sanitization of the response, redirection to another server, or presentation of a customized response. Visibility is the key to enabling this functionality, which is why it should be considered a core capability of any ADC.

Security

Because of its location in the data center architecture, an ADC is uniquely positioned to provide security in a variety of ways to protect not only the application but the computing and network resources upon which applications rely.

Increasingly, this strategic location requires advanced security such as data center, network, and web application firewall services. As an intermediary between clients and services, an ADC offers a cost-effective and processing-efficient solution for deploying security services. From the client perspective, the ADC is the endpoint and thus a more appropriate point for network and data center firewall services than an upstream or downstream device. Similarly, the ADC must necessarily intercept and examine requests and responses to perform advanced load balancing and application routing functions, which provides it the opportunity to examine the content in depth and to ensure it is free of infection or malicious code.

Some regulations and standards, including the U.S. FIPS 140-2 Level 2 standard, require security that can only be provided by hardware. BIG-IP devices are FIPS 140-2 Level 2 Certified.

In addition, some government and industry regulations, certifications, and standards require the additional security that can only be provided by hardware. For example, the U.S. government's Federal Information Processing Standards (FIPS) 140-2 Level 2 standard and above require security mechanisms such as tamper-evident hardware. Compliance is not possible using a software or virtual form-factor ADC without additional hardware and costs. Specific security requirements such as this are therefore important considerations in ADC purchase decisions.

If an ADC will be used as a firewall, certifications such as the International Computer Security Association (ICSA) certification can help assure companies that the product they choose is secure in ways beyond the mere addition of access control lists to a load balancer. An ADC capable of providing certified services also benefits the organization by presenting a common management platform for all application delivery services—including security—that reduces the hard and soft costs associated with more disjointed architectural options requiring multiple solutions.

For example, BIG-IP LTM can detect the number of layer 7 connections per second per client and impose various rate-limiting schemes that have proven effective in mitigating layer 7-based attacks. BIG-IP LTM further includes native firewall services capable of providing core network-layer protection with a much higher connection capacity than traditional firewalls.

ADCs with programmatic interfaces that allow fast, flexible responses to emerging threats—such as zero-day vulnerabilities or new application-layer attack techniques—enable IT teams to respond immediately to mitigate risks without lengthy waits for patches or hot fixes. The F5 iRules® scripting language is such a programmatic interface. It is supported by F5 engineers and professional services as well as the 80,000-strong DevCentral community, which consistently develops, shares, and refines iRules that then may be signed by F5 to validate their integrity.

The BIG-IP product family was certified by ICSA as a network firewall in December 2011. The F5 web application firewall, BIG-IP® Application Security Manager™ (ASM), is also ICSA certified.

The programmatic ability of iRules provides a flexible means of enforcing protocol functions on both standard and emerging or custom protocols. Via iRules, BIG-IP LTM can be directed to enforce protocol compliance and perform traffic steering and related actions, rate limiting, and response injection. iRules has been designed specifically for firewall-focused services, enabling organizations to react to zero-day or emerging vulnerabilities for which a patch has not yet been released, such as in the case of the 2011 publication of an Apache Killer protection iRule.[2] When evaluating ADCs, organizations should consider this type of programmatic capability in addition to ADC support for pre-packaged policies that encapsulate security best practices and standards.

Manageability

Manageability used to mean providing a command-line interface, GUI, and sometimes SNMP management information bases (MIBs). Today, thanks to the increasingly distributed nature of data center models and the demand to automate and orchestrate IT processes to improve efficiency and scale operations, manageability means much more. Modern manageability requires integration with automation and orchestration systems as well as the ability to codify configuration tasks on every data center component to maximize efficiency and reduce application deployment time.

To be manageable, an ADC must not only provide the means to integrate with orchestration and automation platforms, it must also provide better packaging of tasks to promote agile operations and consistency across environments. Configuring an ADC for even the simplest of tasks, such as load balancing, requires creation and management of many configuration objects and steps. Reducing the time and effort required to perform these tasks is paramount to realizing efficiency gains.

[2] F5 Friday: Zero-Day Apache Exploit? Zero Problem

BIG-IP LTM delivers a dynamic control plane is designed to achieve these goals as a core component of its platform. Use of iControl and TMSH enables both custom and pre-packaged integration with automation and orchestration systems, while iApps offers an elegant, deployment-focused packaging system that enables the rapid deployment of applications with far less risk of human error.

The use of iApps to define and manage applications and their ADC configurations on BIG-IP LTM further enables on-demand replication—across environments and into cloud-based implementations—of the security, performance, and availability policies associated with those applications. This aspect of manageability ensures operational consistency across deployments, regardless of location or environment.

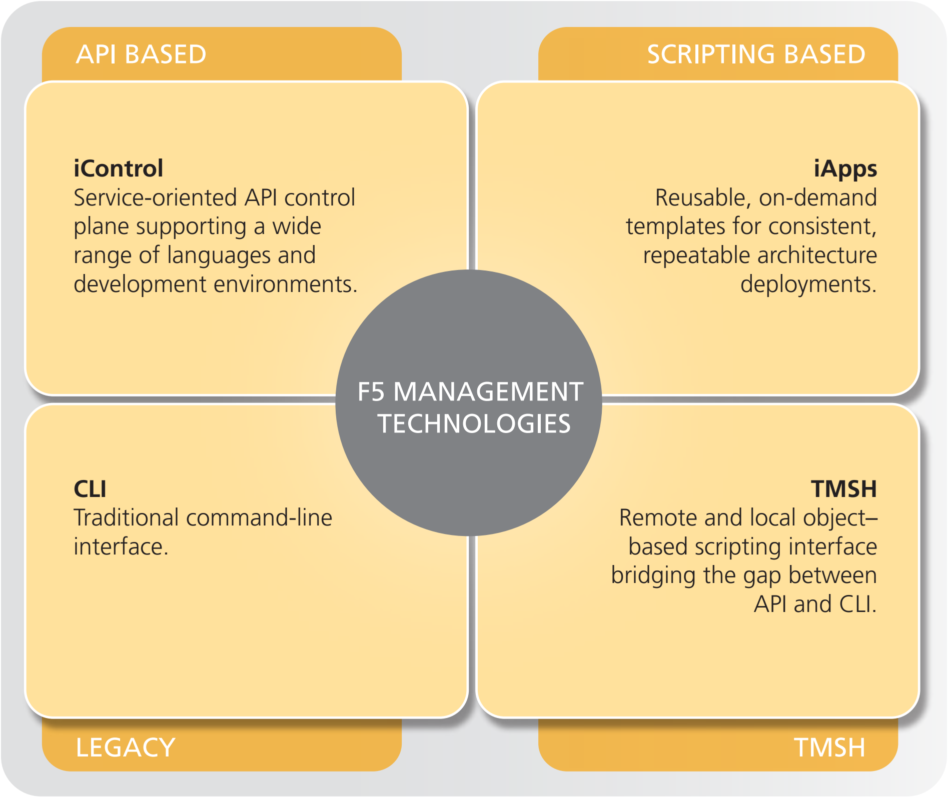
API BASED
SCRIPTING BASED
iControl
Service-oriented API control plane supporting a wide range of languages and development environments.
iApps
Reusable, on-demand templates for consistent, repeatable architecture deployments.
F5 MANAGEMENT TECHNOLOGIES
CLI
Traditional command-line interface.
TMSH
Remote and local object–based scripting interface bridging the gap between API and CLI.
LEGACY
TMSH

Figure 3: The F5 management plane provides a variety of options to ensure that both pre-packaged integration and customization are available to meet specific organizational needs.

The concept of shared policy management for efficiency extends to consolidation. While many define consolidation as simply aggregating many like devices into a faster, bigger hardware platform, F5 views consolidation as also including the use of shared infrastructure to deploy and manage like application delivery services. An

ADC should be able to consolidate web application security, access management, load balancing, and acceleration services onto a single, shared, and consistently managed platform, not only to reduce performance degradations caused by an architecture composed of multiple solutions, but to reduce the time and costs associated with managing multiple solutions.

All application delivery services should be available on a unified, consistent platform through which IT staff can integrate, automate, and replicate policies in an on-demand and highly agile manner to achieve the greatest possible efficiency. An ADC should provide a consistent operational experience across all application delivery services.

When issues do arise, as they are wont to do, the ADC provider should offer an array of services designed to help troubleshoot and resolve the problem as quickly as possible. In addition to guided support options, F5 maintains a variety of supportive services designed to enable self-service troubleshooting and resolution. F5 iHealth® is a self-service, focused portal enabling customers to troubleshoot, optimize configurations, and obtain valuable information if issues are escalated to a guided support option. DevCentral maintains a number of forums and groups in which experienced F5 engineers and customers are able to advise, assist, and provide support in an open community.

Flexibility

Commoditization is one of the primary drivers of cost reduction. By addressing 80 percent of data center needs, a solution can dramatically decrease costs. Unfortunately, the strategic nature of the ADC and the rapid rate at which new needs arise (thanks to consumerization and cloud computing) mean that commoditization can become an inhibitor rather than an enabler of solutions.

An ADC must provide both a cost effective means of addressing common problems and the ability to react to threats and issues that arise from emerging technologies and evolving threat spectrums.

Policy flexibility

Operational efficiencies are primarily gained in modern data center models through the use of standardized or template-based policies. The question with an ADC becomes "Whose standards are used to codify these policies?" This is an important question to ask with respect to ADC policy creation and management because the diversity of applications delivered by an ADC and the number of variables involved cross vertical industry lines as well as organizational peculiarities. A standardized template encoded by the vendor may well address most organizations' needs, but for those organizations not covered, such templates can be frustrating or altogether useless.

Similarly, security policy codification will fall behind quickly after implementation as attackers are evolving faster than the market's natural product revision cycle. While being able to check a box to protect against exploits is certainly desirable, the ability to protect applications and infrastructure against zero-day exploits is even better.

A combination of standardization and flexible customization is the best approach to meeting the need for both customizable and efficient, standardized templates. F5 iApps provides this combination, offering pre-packaged templates as well as a framework that enables organizations to codify their own policies in a standardized way. Backed by a vibrant and active community of IT practitioners on DevCentral, iApps offerings combine the best of standardized support with flexible, customized templates.

The iRules scripting language further enables deeper customization and application management by facilitating deep content inspection and manipulation. IT staff can use iRules to respond immediately to emerging threats, address unique application issues as a stop-gap measure between patches or application updates, and craft innovative solutions and architectures that provide significant business and operational value.

An ADC without the ability to support a variety of policy management approaches is ultimately relying on infrastructure upgrade cycles to improve support and security. But an ADC with the flexibility to adapt independently of upgrades provides a much better return on investment in the long term.

Architectural flexibility

A flexible ADC should be able to support new technologies and deployment models without requiring new solutions. Cloud integration models, for example, are variations on existing themes that rarely require a brand new, and ultimately costly, product. For organizations taking advantage of an ADC in cloud computing scenarios, what changes is the architecture. Thus, when evaluating an ADC, it is more important to examine its ability to support a wide variety of architectures than it is to look for a companion "cloud" product to provide functionality that almost certainly already exists within the core product. Such companion products are often the result of vendor cloud-washing and are counterproductive, more often than not increasing complexity and thus negatively affect performance, return on investment, and failure rates.

Some technology will always be required to support new models, and in the case of cloud computing, these technologies are networking standards such as EtherIP (RFC 3378) and IPsec, neither of which are new technologies but are becoming requisite components of new architectures required to support the integration of cloud services with the data center. An appropriate ADC solution will be compatible with these technologies as well as existing capabilities such as deep content inspection, global server load balancing, and deep visibility to enable cloud and data center integration.

F5 has long been able to integrate remote infrastructure and resources into an overarching data center architecture, and the BIG-IP platform makes use of both EtherIP and IPsec. F5 also supports cloud computing and virtualization, both highly flexible frameworks, by offering its solutions as virtual editions. Virtual editions are available for all BIG-IP products and can be integrated into architectures to better provide the scalability and portability of application delivery services without sacrificing operational consistency across environments.

Application flexibility

The way in which an ADC manages delivery policies is important, but so are the applications it supports. An ADC should be able to support a broad set of applications over a variety of protocols. The ADC evolved from the need for large-scale web applications, and its focus has predominantly remained on web-based applications and protocols. The need for scalability on today’s Internet, however, has moved beyond these simple protocols and now encompasses a broad spectrum of transport and application layer protocols such as SIP, SMTP, MMS and SMS, UDP, virtual desktop infrastructure (VDI)-related protocols, DTLS, and more.

The ability to control and deliver any application is critical. It is not enough to focus on one specific application or workload; a flexible ADC must be able to simultaneously deliver a variety of applications, each with its own unique policies and workload types.

Conclusion

An extensible, integrated application delivery platform is the foundation of future data center architectures. Whether integrating cloud computing resources or providing a flexible infrastructure tier through which emerging mobile and VDI applications can be delivered, an ADC provides the critical application delivery services required to support the security, availability and performance requirements of today's demanding and highly dynamic data centers.

The best long-term returns on an ADC investment will be achieved by organizations that carefully consider not only financial criteria but also a variety of operational criteria in the following categories:

- Performance
- Scalability
- Visibility
- Security
- Manageability
- Flexibility

The ADC an organization ultimately selects will have a significant effect on the efficiency, agility, and responsiveness of its IT infrastructure.

Offering outstanding value against each of the key selection criteria, the F5 BIG-IP product family has been designed for performance, cross-environment visibility, agility, flexible management, on-demand scaling via hardware or software, and easy extension into cloud computing environments to enable dynamic data center models and add value today and in the future.

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