F5 White Paper

Secure, Optimized Global Access to Corporate Resources

Global access to corporate applications is critical to an organization. The range of users, devices, and their locations requires stringent application access control to securely connect any user on any device from any location to wherever the application lives.

by Peter Silva
Technical Marketing Manager, Security

Contents

Introduction

The mobile workforce is expected to increase from 919 million in 2008 to 1.2 billion in 2013—a figure that will represent 34.9 percent of the worldwide workforce.[1] Workers are dispersed all over the globe, and they use a variety of trusted and un-trusted devices to request access to corporate resources from different types of networks. Users need fast, secure, and reliable access to the corporate infrastructure; but IT departments are struggling with multi-vendor access solutions and systems, policy enforcement, access regulation, security threats and vulnerabilities, and ensuring that the right user is connecting to the right applications based on context. If all that weren't enough, IT departments must also keep management and maintenance costs in line. Enterprises need unified and converged access and policy management in a globally distributed environment, so they can connect any user on any device from any location to the application, wherever it lives, securing and optimizing content delivery.

The days of expensive, single-point products that only provide an isolated solution are waning, and customers are now looking to consolidate their deployments for easier management and better return on investment (ROI) and lower total cost of ownership (TCO). Unified, optimized, and secure access is what organizations and their users and customers need—and F5 has the unified access solutions to deliver it.

Dynamic Control and Mobility

In the flow of anything, be it traffic, water, electricity, or commerce, there are specific points that can provide intelligent processing and control. Typically, these control points exist in strategic locations within a given flow; they are discrete boundaries or junctures where multiple streams of information convene for redistribution. The purpose of these points is to provide critical flow control, redirection, intelligent management, and security. The control mechanism might be simple, like a stop sign at an intersection; or it might be complex, like power distribution and control systems in modern high-rise buildings. It might be physical, virtual, or procedural. But regardless of form, without these strategic points of control, the efficient, safe, and uninterrupted flow of any system would be impossible.

At the same time, the origin and destination of traffic rarely stay static over time. As flows change, the strategic points of control often cannot adapt dynamically and thus become impediments to efficient flow. As roadways become highways, the

previously adequate control point (such as the simple stop sign) can become a dangerous bottleneck or worse. But the process of changing the type or function of a control point to meet the new demands can be long, complicated, and costly.

Today, IT departments require control points that can adapt dynamically and secure content and applications as they are delivered from a variety of locations to a multitude of users. This is especially true for global infrastructures that span the cloud and the data center. Many users are on the move, and they are accessing applications and content that may be on the move as well, from data center to cloud and back again; therefore a central policy control point is critical to managing these dynamic environments. When a control point is decentralized, IT has limited control over the flow of data. There is no contextual information with which to make intelligent decisions; and change control becomes difficult and error prone, and removes the simplicity and flexibility needed for global access deployments.

Unified access control is about securing and optimizing the delivery of applications globally to remote users by connecting them with the least amount of latency and as close to the application as possible.

Application Security at the Edge

Access can mean different things—access to an intranet web application to search for materials; to Microsoft Exchange for email; to virtualized Citrix, VMware, or remote desktop deployments or to a particular network segment for files and the full domain network. The resources themselves can be in multiple locations: corporate headquarters, the data center, a branch office, in the cloud, or a mix of them all.

When users are all over the world, globally distributed access across several data centers can help solve access and availability requirements—but both the user base and IT administrators still need a solution that is easy to use and manage. Application and access security solutions should provide centralized, strategic points of control. These Application Delivery Controllers (ADCs) are often located at the edge of the network, whether it's in-house or in the data center. The strategic points of control, positioned between users and the resources they need, can make intelligent decisions about how to handle access traffic. The F5® portfolio of application delivery and unified access products, including BIG-IP® Local Traffic Manager™ (LTM), BIG-IP® Access Policy Manager™ (APM), and BIG-IP® Edge Gateway™, provides the security, scalability, and optimization required for unified global access to corporate resources in all types of deployment environments. By

converging and consolidating remote users, LAN access, and wireless junctions within a single management interface, and providing easy-to-manage access policies, organizations can save money and free up valuable IT resources. F5's unified access solutions secure an infrastructure from within by creating a place in the network to provide security, scalability, optimization, flexibility, context, resource control, policy management, reporting, and availability for all applications.

In the traditional IT model, resources, users, and access methods were controlled by IT. Relationships among users, applications, and data were static and tightly bound; and applications were written with specific display layouts in mind. As remote and then mobile users were added, along with partners, contractors, and guests, and as IT was distributed globally, the traditional model broke down. In today's complex IT security world, it is not enough for an organization to simply know who is accessing their data and applications. Organizations must also know what device, what type of network, and which resources users are requesting. It's also critical to know whether the requested servers, networks, and applications are available and secure. With advanced dynamic control for universal global access, organizations can oversee the complexity and unpredictability of all these moving parts.

It Starts with the Endpoint

With global users requesting access to a variety of distributed applications, IT organizations must ensure they are directed to the closest or most optimal resource location that can deliver the requested content. When the user types a resource location or clicks a local program that requires a connection to a corporate server (such as email), an initial endpoint inspection can determine the location of the user, what type of network and device they are on, and other contextual metadata. With this information, the system can decide which data center location is optimal for the request. The intelligence in F5's unified access solutions can instantly direct the user to the most appropriate data center for their specific request, delivering optimal results for the user.

However, the closest location might not always be optimal for that particular user at that moment. Consider this scenario: A typical UK user's requests are usually delivered from the UK data center location, but unbeknownst to the user, the UK location is experiencing availability issues. Working in concert, the components of F5 unified access solutions are aware of the issues and will automatically reroute the user to the nearest optimal data center or cloud deployment.

Another scenario involves a data center outage that occurs while the user is downloading a file from that location. Even in this instance, F5 unified access solutions can retrieve the file from another location and fulfill the request as expected. The user is likely unaware of the reroute, and unlikely to care as long as they can get to their resources quickly and efficiently. Availability is a key metric when delivering applications to a global user base—one that F5 unified access solutions can help organizations achieve.

Another IT challenge is endpoint security. It can be difficult for IT organizations to control the endpoint security posture of the different types of machines requesting access to applications. With F5 unified access solutions, security starts as soon as the user requests a resource. Detailed endpoint host inspections enable IT to determine how much resource access to grant a specific user and device based on the corporate access policy criteria. This can range from examining the security posture of a device to inspecting and identifying the user's MAC address, CPU ID, and HDD ID, as well as determining whether the device is part of the corporate domain or a non-domain (partner, contractor, personal). Every device, whether it's an IT-issued laptop, an employee's home computer, a tablet, or a mobile smartphone, receives the same scrutiny. This is important when an organization is determining whether to grant access rights to the corporate network, and which specific resources the user can access.

Administrators can also provide remediation solutions for non-compliant devices, or place those devices in a Protected Workspace. Protected Workspace is a virtual desktop session that is completely erased when the session concludes. Administrators can even disable USB or CD-ROM read/write functionality and restrict local file access as part of a Data Loss Prevention (DLP) solution.

It Continues with Identity Management

Anonymous networks allow users to access systems via a user ID and password, but they cannot decipher exactly who the user actually is, such as an employee, guest, contractor, or partner. Anonymous networks do have visibility at the IP or MAC address level, but that information does not equate to a user's identity. Since these networks are unable to attribute IP to identity, the risk is that some information may be available to unauthorized users. These networks do not include reporting on what was accessed or where a specific user navigated to within a system.

> "Security professionals must stop trusting packets as if they were people."
>
> John Kindervag, Build Security Into Your Network's DNA: The Zero Trust Network Architecture. Forrester Research. November 2010.

Controlling and managing access to system resources must be based on identity. A user's identity, or their expressed or digitally represented identity, can include

identifiers like: what they say, what they know, where they are, what they share, who they know, and their preferences, choices, reputation, profession, or any combination of these that is unique to the user.

Unauthorized access to systems is a huge concern for companies, not only because of the potential disclosure and loss of confidential company data, but because of the regulatory compliance risks. It is critical to every business that only authenticated users gain admission to their networks, and that those users access only the resources they are authorized to see. In an enterprise infrastructure, authentication, authorization, and accounting (AAA) services is the primary method deployed to verify user identity. These systems can be complex, and managing AAA services in a web application deployment can be costly.

There are a number of ways to authenticate web users: some organizations code authentication right into the application during development; others install agents on the servers; and many have specialized access proxies. All of these methods aredifficult to manage and change, and they are not particularly interoperable. They can become costly in regard to both deployment and management since they're decentralized and every server needs individual attention. These methods can be not secure or scalable enough for a global workforce, and they may also overlook authorization and accounting, which are often required for regulatory compliance.

If an organization has a cloud deployment, adding BIG-IP APM to BIG-IP LTM or BIG-IP LTM Virtual Edition (VE) provides identity, authentication, and access control, which enables IT to consolidate the global infrastructure, reduce AAA management costs, and drive user identity into the network. BIG-IP APM integrates with an organization's AAA infrastructure, including Microsoft Active Directory, LDAP, RADIUS, Native RSA SecurID, and others, so organizations can manage access based on identity. Corporate users may authenticate against the Active Directory server, and partners and contractors may query LDAP, but all users go through the same endpoint host scrutiny prior to gaining access. BIG-IP APM centralizes web single sign-on (SSO) and access control services; offers a full proxy L4–L7 access control (at wire speeds); adds endpoint inspection to the access policy; and includes the Visual Policy Editor (VPE) feature, which provides policy-based access control along with VPE Rules, a programmatic interface for custom access policies.

With VPE, administrators can easily create and manage security policies and resources, and VPE's flowchart-style design shows exactly what types of inspections are enabled. This gives administrators complete control over which resources get delivered to which user or group of users, and enables them to control access based on device and identity. With dynamic, per-session L4 and L7 access control lists,

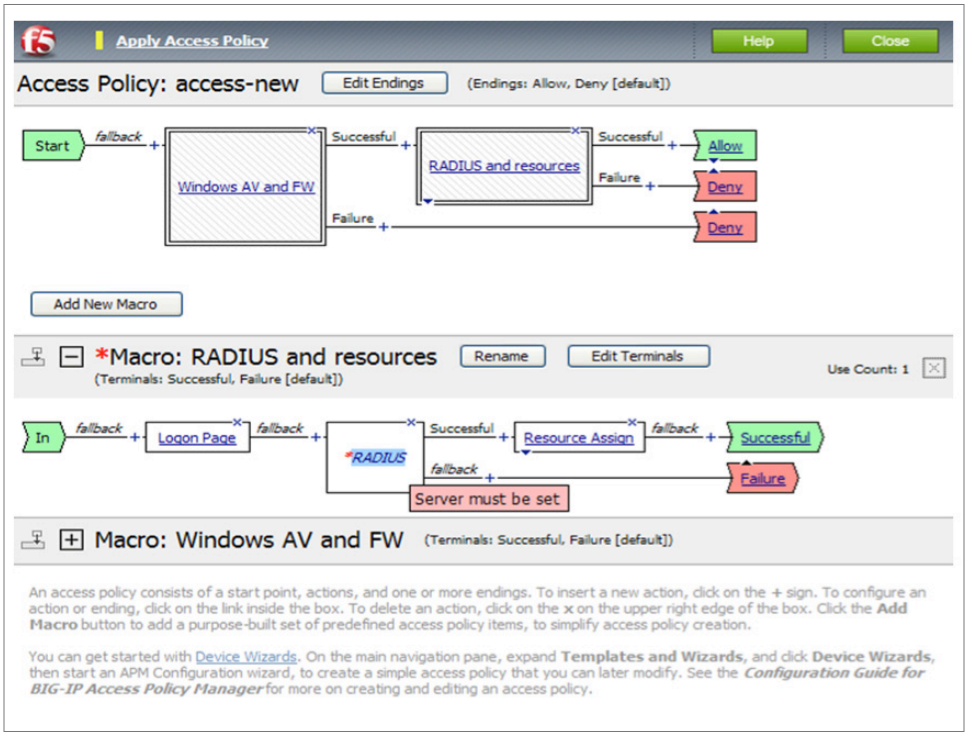

Figure 1: Access policy design using Visual Policy Editor

administrators can keep users within their functional, authorized locations at a fine-grained level—even down to a specific folder path within a web application. All this is done over SSL, so every connection is encrypted.

Administrators can enable multiple domain SSO solutions if certain users need access to multiple separate domains or multiple hosts within the same domain. VPE also supports Kerberos ticketing, which provides easy authentication design for Active Directory implementations, and it consolidates all authentication management in one Kerberos Protocol Transition (KPT) solution. Users enter their credentials only once to access resources that are spread across a distributed architecture, boosting productivity and simplifying login.

In the case of web application authentication, organizations can use BIG-IP APM to replace specialized access proxies or agents and gain superior scalability and high availability. They can also initiate an endpoint host inspection for any client requesting access to their web application, whether it's public or internal, to ensure a minimum security posture and enforce stronger authentication than is typically available for web applications. If a situation arises in which the organization needs to provide authentication and client validation prior to granting access to unsecured applications, secure authentication can be added. This authentication method is

used first to create a secure session and then to provide access to the services behind it. For instance, an organization may not want to lock down its public-facing website, but certain requests may require authentication (for example, if a user requests access to a restricted folder). With secure authentication, administrators can manage web application access efficiently; anyone can navigate the main page, but when the user clicks a "member" area, access controls raise the gate.

Global Implementation

Gaining Access to the Resources

In today's distributed architectures, applications and resources may be spread around the globe depending on an organization's needs. For example, its financial applications might be in-house, email hosted in the cloud, and CRM and ERP systems at a data center—and each installation has its own disaster recovery or backup site. Applications may move from the data center to the cloud during traffic spikes, and then back to the data center when traffic normalizes. If users are mobile and requesting applications that are also in motion, the efficient delivery of content could pose a challenge.

With dynamic access control, managing resources has never been easier. Administrators can create a single policy for all access methods or create unique policies depending on access method, device, user group, or other benchmarks. The resources themselves may be circulating throughout the infrastructure, but with specific application access criteria tied to them. Depending on host inspection results, administrators can grant full, some, or no access to the requested asset. An employee requesting access from their IT-issued laptop might get a full network access tunnel, but from their home computer, they might get restricted or limited access. Contractors or partners may have access to a certain network segment or terminal servers during business hours, but if they try to access the system after hours, they might only be able to access a dynamic webtop with only web-based applications available, for instance, ERP and CRM applications. This means that even if users and resources are in transit, the access policy will always still apply.

With users and resources globally distributed and often on the move, latency can also be a challenge. Ideally, users' requests would be addressed immediately—but this can be tricky with users and applications on opposite sides of the globe. The varying location and network of both users and applications can cause issues like latency, packet loss, and poor performance, which can have a detrimental effect

on application performance and user experience. Integrated WAN optimization and web acceleration in BIG-IP Edge Gateway minimize the effect of latency without an organization having to build a data center or co-locate equipment in a particular region. BIG-IP Edge Gateway combines remote access and optimization services on a single BIG-IP platform. Its optimization services can be used for data centers, POPs, remote sites hosting applications for mobile users, and remote branches accessing those applications.

Quality of service, particularly with VoIP, is another challenge for mobile and remote users. BIG-IP Edge Gateway offers a Datagram TLS (DTLS) mode for remote connections. TLS is the standard protocol used for securing TCP-based Internet traffic (also known as SSL); and DTLS is a protocol based on TLS that can secure the datagram transport. It is well-suited for securing and tunneling applications that are delay-sensitive. This solution reduces the required hardware in locations that may have delay-sensitive networks; provides effective application access management; and greatly improves user experience.

Secure and End to End

Most F5 unified access traffic, whether proxied or tunneled, is secured over layer 7 TLS/SSL. IPsec tunnels are also supported, which secures traffic over layer 3 using F5's iSessions technology. Two BIG-IP devices can create secure, encrypted, optimized tunnels over the WAN with iSessions. Organizations can create their own secure, optimized WAN, which adds authentication and encryption to every packet. Not only does this secure and optimize traffic between users and the applications, but organizations can secure and accelerate traffic between sites; between HQ and a branch office; between the primary data center and the cloud services facility; and between remote and local and everything in between. IPsec can secure any application traffic over an IP network. By using iSessions in the F5 unified access solution, organizations gain end-to-end security across the entire global infrastructure.

There may be instances based on user identity, security implications, or other criteria in which a full layer 3 tunnel is not appropriate or the user simply needs to access a single application, for example Exchange with their local Outlook client. In these situations, rather than granting full network access, a single F5 AppTunnel might be more suitable. AppTunnels allow organizations to create a single secure link to a specific application without having to open up full network access, so mobile users can simply click their Outlook client to get secure access to their email, no matter where they are in the world. AppTunnels are also completely WAN-optimized so

Integration with BIG-IP GTM

BIG-IP Edge Gateway is integrated with BIG-IP® Global Traffic Manager™ (GTM), so as individual BIG-IP Edge Gateway devices reach certain thresholds, BIG-IP GTM can globally load balance users to the next-best BIG-IP Edge Gateway device and provide emergency capabilities when needed. Disaster recovery, business continuity, and workforce continuity are accomplished all in one solution.

those application connections benefit from adaptive compression, acceleration, and TCP optimization to efficiently deliver content to the user.

Real-World Scenarios

Many companies are deploying Microsoft Exchange Server 2010 or migrating from Exchange 2007 to Exchange 2010. Upgrading any system is a challenge. Transitions come with a host of potential problems: some users may be on the old system and some on the new; users may be temporarily unable to access email; administrators and users might need to change their Outlook client server settings to make sure they are pointed to the proper server; replies to certain messages may bounce; the OWA URL may change; and calendars may be unavailable.

Deploying F5 unified access solutions can alleviate many of the issues associated with an Exchange migration. First, the F5 solutions allow organizations to migrate over time while BIG-IP APM authenticates users in the DMZ to ensure there are no unknown users accessing the system. Organizations can distribute a single URL, and depending on the user or group, BIG-IP APM will direct the user to the appropriate server for any Exchange iteration (OWA, ActiveSync, or Outlook Anywhere). This gives users direct access to email without requiring that they update bookmarks or other settings. Organizations can also manage email access for all devices, from all

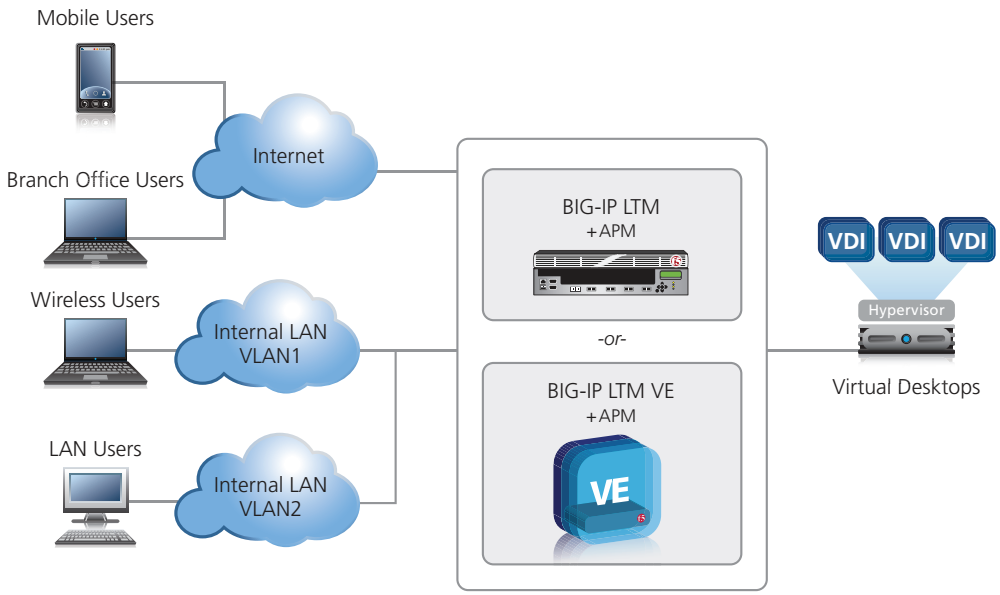

Figure 2: BIG-IP APM in a Virtual Desktop Infrastructure

locations, and on any network. After migration, F5 unified access solutions can scale to accommodate 600 logins per second, and support up to 60,000 users on a single appliance. Enabling hosted virtual desktops is also simple and secure.

The addition of the fine-grained access control in BIG-IP APM to BIG-IP LTM is a powerful enhancement to any virtual desktop deployment. BIG-IP APM optimizes, secures, and delivers a Virtual Desktop Infrastructure (VDI). The BIG-IP system improves availability and scalability by providing important load balancing, health monitoring, and SSL offloading for VDI deployments. Network and protocol optimizations help organizations manage bandwidth efficiently and in some cases, reduce the bandwidth requirements while maintaining and improving the user experience. BIG-IP APM also enables organizations to make virtual server load balancing decisions based on user identity, ensuring that users are connected to the optimal virtual instance based on their needs.

BIG-IP APM for LTM Virtual Edition (VE) can also be a 100-percent virtual remote access solution for VDI solutions, and it can be deployed as part of a hybrid cloud or disaster recovery strategy. In addition, BIG-IP APM for LTM VE will run as a virtual machine in a VMware hypervisor environment so organizations can easily add it to their existing infrastructure. As the number of users on virtual desktops grows, customers can easily transition from virtual to physical editions of BIG-IP products. BIG-IP APM for LTM VE simplifies authentication and session management for VDIs.

Simple, Detailed Management

With all this technology at IT's fingertips, managing global access should be effortless. It is also important to understand the application performance and have visibility into how remote users are accessing the applications. In BIG-IP version 10, F5 introduced Application Ready Solution Templates, which simplify the creation of virtual servers, pools, profiles, monitors, F5 iRules® scripting language, and other pertinent configuration parameters needed to optimize BIG IP deployments for specific applications like SharePoint, SAP, Oracle, and Exchange. The templates contain common default values and minimize the number of GUI clicks required to configure a BIG IP device. The common default values were determined by rigorous testing and collaboration with application vendors—they represent the optimal configuration for BIG IP devices and a given application.

New in BIG-IP version 11 is iApp, which comprises Application Services, Templates, and Analytics. iApp gives the organization a complete perspective of overall application performance. It enables an organization to control all of its ADC services

in the context of the application being deployed, and it provides infrastructure to define the components of an application service. From this definition, administrators can create the application service instance, and then for each application service, they can understand how the application is performing through application level stats and reports.

iApp's Application Services give administrators a point of control at the application level, and they provide the main point of manageability for a specific deployment. When an administrator deploys an application, iApp Application Services associate all related dependent objects, provide availability status for each object, and provide context for stats and reporting. iApp Application Services are created from and tied to a template, which enables administrators to define objects and the UI required to deploy an application. iApp Analytics provide application- and HTTP-level statistics and reporting on a per application basis. iApp Analytics collects, aggregates, and reports application-level statistics and is configurable via application templates. It presents customized views based on different objects, and statistics are polled at 5-second intervals. Administrators can set alerts through syslog, SNMP, and email, and log reports to remote reporting engines.

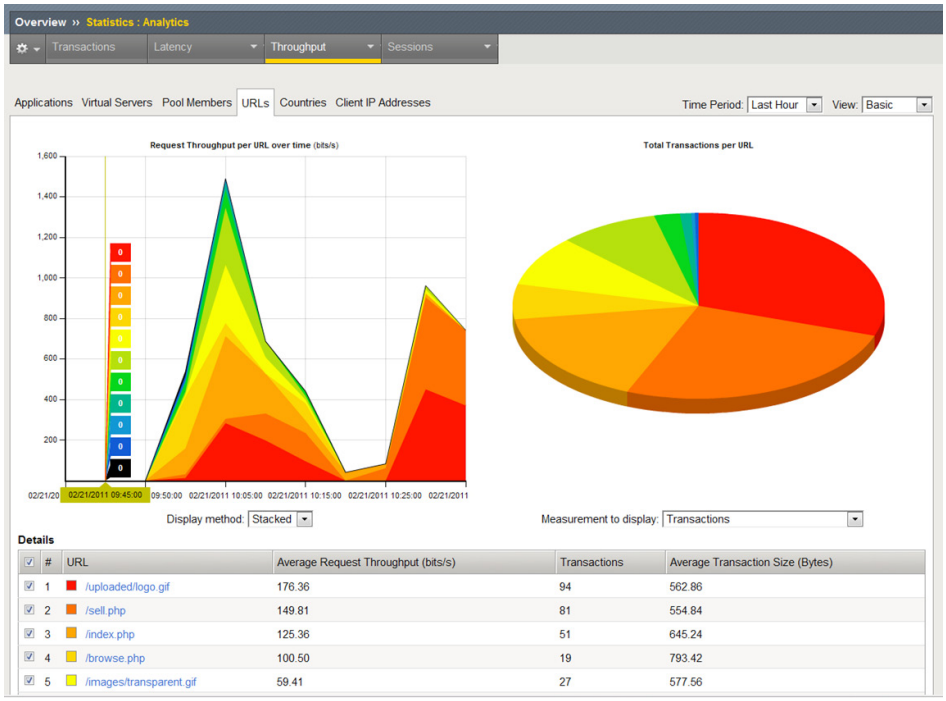

Figure 3: Per-application visibility and reporting with iApp

Performance reporting gives administrators a complete assessment of their global user base and their interactions with resources. The full interaction of the session lifecycle—user to app and back again—is available for review. The aggregation of data and on-box reporting give IT administrators and business groups alike the contextual information they need to make informed business decisions about their network, infrastructure, systems, applications, and users.

Conclusion

The mission is to connect any user from any device from any location to wherever the application lives, securing and optimizing the delivery of content. Unified access control can help an organization optimize secure application delivery to remote users around the world by connecting them as close to the application and with as little latency as possible.

F5 unified access solutions provide end-to-end availability and optimization with secure unified access. Secure mobility extends the application policy to the client, addressing the organization's challenge of identity and access management. Users get fast and secure connections to the resources they need and IT gets easy management, policy control, and detailed reporting of the user experience. Context also plays an important role to ensure users and applications come together seamlessly and securely.

If organizations have cloud deployments, BIG-IP APM for LTM VE enables them to achieve complete, optimized, secure connectivity in a virtualized environment. F5 unified access solutions secure a global infrastructure by providing security, scalability, optimization, flexibility, context, resource control, policy management, reporting, and availability for all applications from that strategic point of control within the network.

[i] IDC Report, "Market Analysis: Worldwide Mobile Worker Population 2009–2013 Forecast". December 2009, IDC #221309, Volume 1.

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

CS01-00056 0611