SOLUTIONS PLAYBOOK

September 2016 - March 2017

SECTION CONTENT

APPLICATION DELIVERY

CLOUD (PRIVATE AND PUBLIC)

SECURITY

SERVICE PROVIDER

MANAGED SERVICES

F5 NETWORKS

SOLUTIONS PLAYBOOK

APPLICATION DELIVERY

CONTENT

01 Global traffic balancing
02 Local traffic balancing
03 Balancing environments through multi-homing
04 Business continuity with DPC minimum back-up Services
05 Business continuity for systemic institutions
06 Business continuity and Access congestion control
07 Server resources cost control
08 Firewalls scalability
09 DNS scalability solutions
10 Federation and simplification of VDI environments
11 Integration of hyperconvergent platforms – Nutanix
12 Migrating Exchange environments
13 Optimization of DC access based on geolocation
14 Optimized cache balancing (CARP)
15 TCP optimization
16 Rewriting gateway domains
17 HTTP/2 gateway
18 IPv6 gateway
19 Geo-location solution with EDNS
20 HA and DC solution in HUB & spoke topologies
21 Interconnect solution for overlapping networks
22 Data centre persistence solution
23 Customized portals solution
24 DNS64 solution
25 Multi-Language solution in web environments
26 NAT64 solution
27 Intelligent packet brokering solution
28 Cisco ACE replacement

GLOBAL TRAFFIC BALANCING

PROBLEM

The globalisation of companies plus the increasing trend to distribute applications across multiple data centres and cloud services providers has come at a time when working practices have also changed to encourage a more flexible distributed and mobile work force.

The result is the need to enable customers, employees, and automated services to connect to this newly distributed applications based on a number of business parameters such as service availability, geo-location, response times, costs, etc. Simply providing a DNS record is no longer enough.

ALTERNATIVES

- Lack of business continuity in cases where a complete data centre goes down.
- Only being able to implement active-passive strategies with the consequent underutilisation of infrastructure and high failover costs (involving time and processes).
- Excessive latencies and degradation of service for roaming users.

F5 | DNS SOLUTION

The DNS module from F5 (formerly known as GTM), allows you to distribute client traffic across multiple locations (either DC or cloud services), based on multiple business metrics. Thus users located across the globe requiring access to an application can use a single service name (FQDN). Depending on their location, the applications availability, the DC's availability, link speeds, number of hops, etc the user is directed to the most suitable location to service their needs. This allows:

- The guarantee of the continuous service through the constant monitoring the status of the applications in each location.
- Reduction in latency issues by directing the user to the closest or fastest responding instance thereby improving the user experience, response times and productivity.
- Deployments of applications in active-active data centres to deliver highly available soultions.
- The optimal use of data centres (or the Cloud) depending on cost, response time, etc.

REFERENCE ARCHITECTURE | GLOBAL TRAFFIC BALANCING

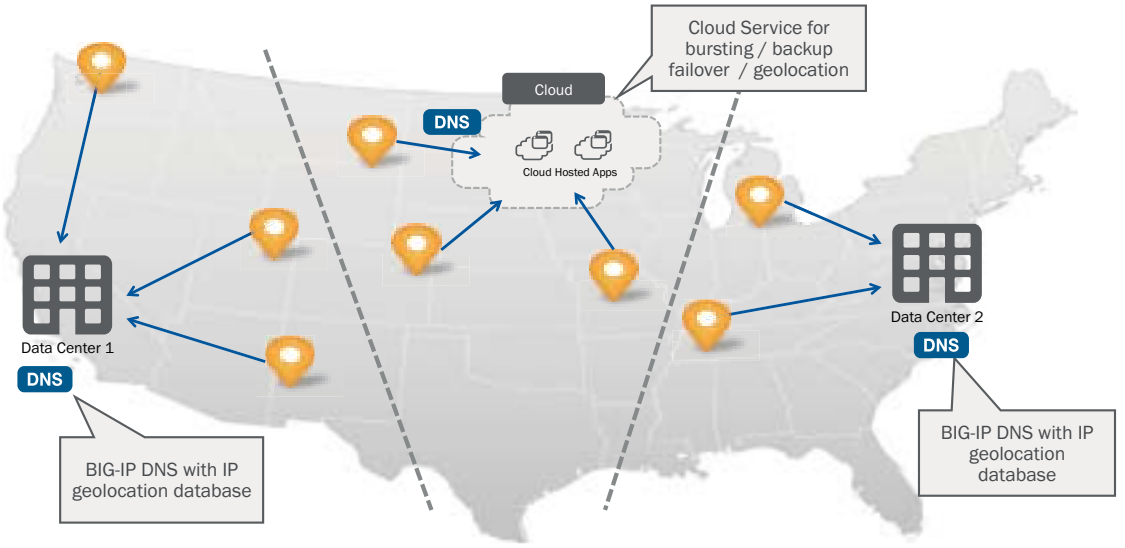

BALANCING LOCAL TRAFFIC

PROBLEM

One of the main problems with both internal and customer facing application is how to deliver service in an efficient manner as the number of users grows at a rapid rate. A rapid growth in customer base or spikes in traffic can put stress on servers and increase response times or adversely affect the availability the service. The result is a poor or unstable service to the user.

ALTERNATIVES

- Manage multiple input lines and divide the requests manually in different services.
- Use bigger servers which is neither cost effective or flexible.
- Using more traditional and less intelligent balancing mechanisms, which impacts the efficiency of the service and restricts how scalable the solution can be.

F5 | LTM SOLUTION

With advanced monitoring of data centre resources and a choice of load balancing algorithms you can make intelligent decisions on how to distribute traffic to provide a highly available solution with unprecedented scalability.

With the BIG-IP LTM (Local Traffic Manager) solution, while continuing to present a single service (FQDN) to users, it will distribute the traffic evenly across the resources that are available to run the application.

Even in environments where specific bespoke or dated applications are used, F5 enable custom monitors to be used to deliver the same level of scalability and performance without the need to re-write or re-engineer the current solution.

REFERENCE ARCHITECTURE | BALANCING LOCAL TRAFFIC

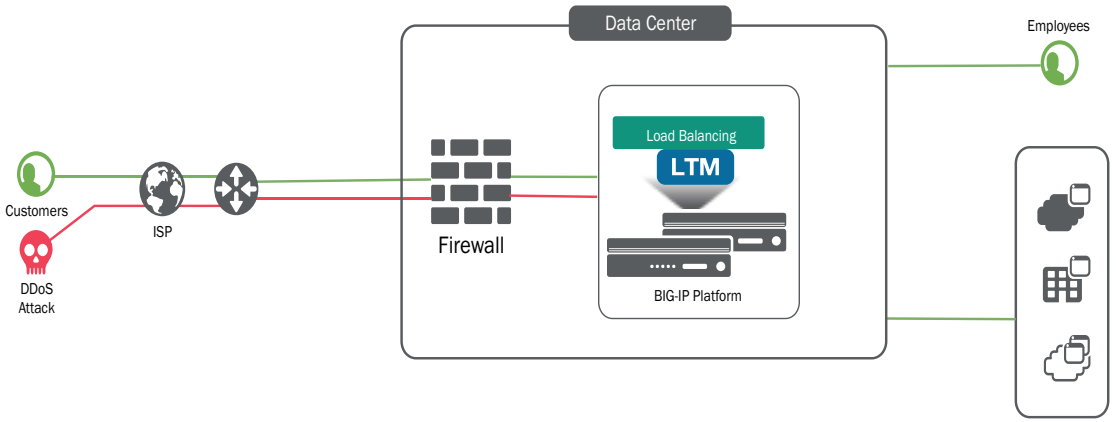

BALANCING ENVIRONMENTS THROUGH MULTI-HOMING

PROBLEM

Mechanisms for deploying applications have changed due to the use of cloud services which offer new models for IT consumption, allowing companies more flexibility in their service models, time to market, cost management, etc.

Other similar alternatives can include the development of private clouds, generating new models of deployment in hybrid architectures making it possible to make use of the best advantages of each type of deployment.

For this reason, it is increasingly common for applications to be deployed using “multi homing”, which requires there to be monitoring of how services are consumed based on new business parameters such as costs, proximity, availability, etc. These are parameter that have not been widely considered to date.

ALTERNATIVES

- Manual configuration on a user-to-user basis or mandatory connection to a DPC.
- Lack of business continuity if there is a complete failure of a DPC.
- Active-passive strategies with consequent under-utilization of infrastructure and fail-over costs (time and processes implications).

F5 | LTM + DNS SOLUTION

F5’s DNS module (formerly known as GTM) makes it possible to distribute traffic from users or clients between different data centres or cloud services based on a wide variety of business metrics.

These capabilities are complemented by the use of the LTM (Local Traffic Manager) module, which allows you to redirect traffic between different locations through the visibility of the traffic at layer 7. The combination of both modules gives you fine grain control over how a user can connect to an application regardless of its location or the architecture being used.

REFERENCE ARCHITECTURE | BALANCING ENVIRONMENTS THROUGH MULTI-HOMING

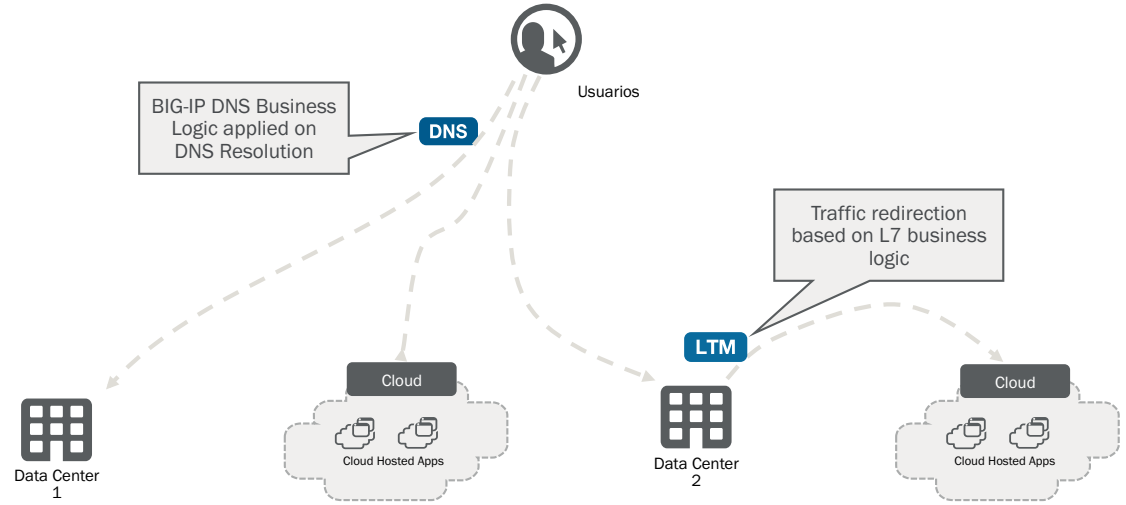

BUSINESS CONTINUITY WITH DPC MINIMUM BACK-UP SERVICES

PROBLEM

Business continuity cannot be impacted in a contingency situation. However, the deployment of back-up centres with full capacity can be a very expensive investment, not forgetting the difficulty of complying with the regulatory requirements.

Ensuring access to the services of the back-up centre should be able to be undertaken on a selective basis for those users who have SLA's that demand it.

ALTERNATIVES

- Manual procedures do not ensure business continuity, and require «human» intervention.
- Building back-up «mirroring» centres impacts on CAPEX/OPEX.

F5 | DNS SOLUTION

The DNS solution enables intelligent traffic-balancing between different data centres in both Active-Active and Active-Passive deployments.

Once critical services requiring business continuity have been determined and users who must have guaranteed access to these services have been defined, F5 provides the automated mechanisms to direct all users to the primary DC then in the case of a failure automatically redirect selected users to the back-up centre. This user-level granularity is critical for compliance with SLAs.

By using data structures in LTM (Local Traffic Manager) (where the specific information for «VIP» users is stored), it is further possible to determine the different services to which these users have access in the back-up centre. If a user does not have access to a service, it will display a message along the lines of «service temporarily unavailable».

By using iControl, it is possible to integrate the management of these data structures with third party tools, which makes it possible, for example, to import data files for «VIP» users into the solution.

REFERENCE ARCHITECTURE | BUSINESS CONTINUITY WITH DPC MINIMUM BACK-UP SERVICES

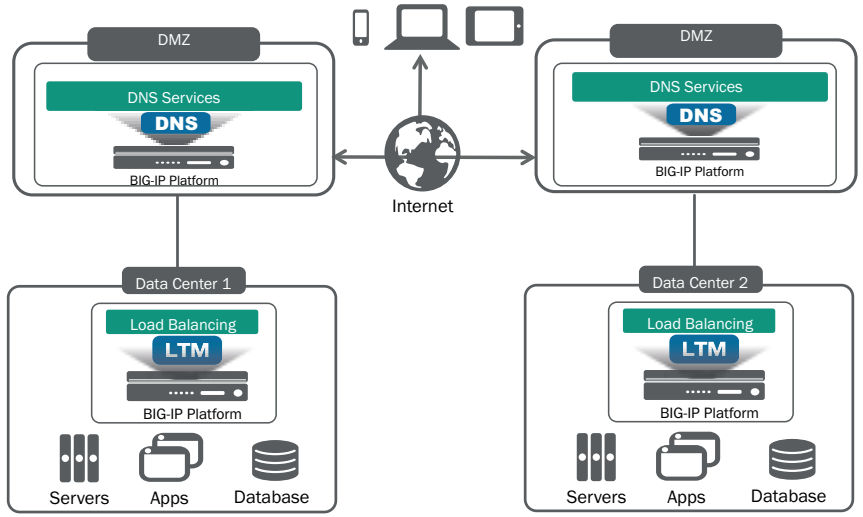

BUSINESS CONTINUITY FOR SYSTEMICALLY IMPORTANT FINANCIAL INSTITUTIONS

05

PROBLEM

Companies required to comply with the Basel III and Sarbanes-Oxley regulations face the risk of incurring penalties when operating distributed data centres dispersed between different continents without the capability to meet failover times.

The requirement is to have an RTO (RECOVER TIME OBJECTIVE) of less than an hour to ensure the business cannot stop for more than an hour. Additionally the business must have an RPO (RECOVER POINT OBJECTIVE) of less than 24 hours to ensure data consistency must be able to be recovered in less than 24 hours.

ALTERNATIVES

- Manual redirecting of IPs to cover the RTO. This is time consuming and are prone to delays due to network convergence.
- Increasing the number of lines or expanding the bandwidth of existing lines to cover RPO. This is a very expensive solution.

F5 | DNS + AAM SOLUTION

Using the DNS product, F5 go beyond RTO standards by enabling the recovery or forwarding of a service to another DC both automatically and in real time.

Using F5 AAM (Application Acceleration Manager) enables companies to achieve RPO standards while reducing costs (less bandwidth) by applying compression and deduplication techniques to remote copies and back-ups. We are certified in the most important replication solutions: TrueCopy (HDS), SRDF (EMC) and IBM. These techniques reduce both the bandwidth and the time required to back-up and to replicate an asynchronous copy affected by latency.

REFERENCE ARCHITECTURE | BUSINESS CONTINUITY FOR SYSTEMICALLY IMPORTANT FINANCIAL INSTITUTIONS

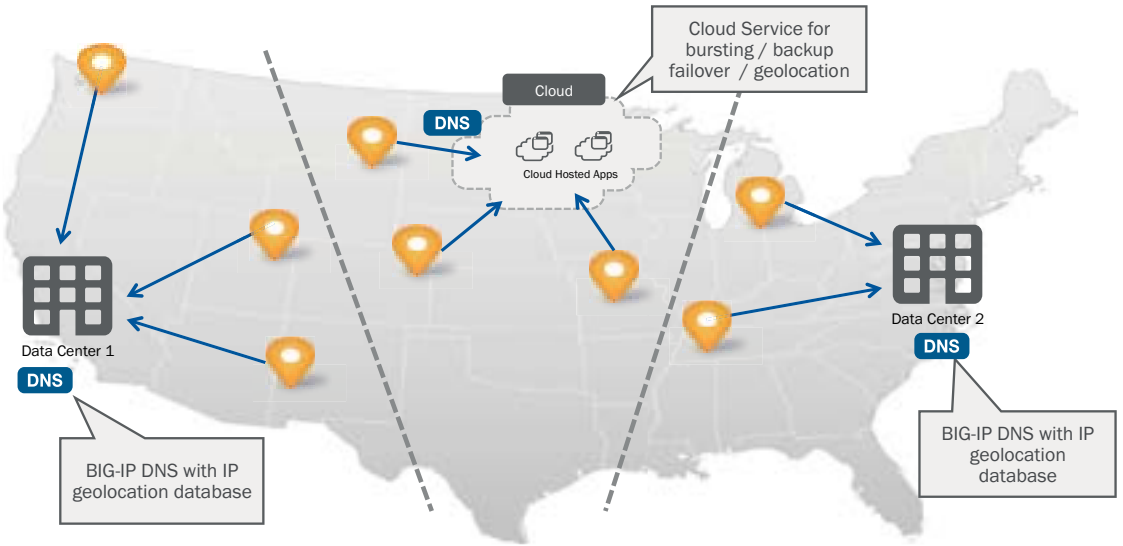

BUSINESS CONTINUITY AND ACCESS CONGESTION CONTROL

PROBLEM

Sometimes companies are in a position where the use of an online service goes beyond the capabilities of their server infrastructure. This can happen for a number of reasons such as marketing campaigns, online ticketing releases, product promotions, online flight-search services, etc. During these events, there will be a significant spike in the number of requests directed towards a particular web service.

On all of these occasions, the availability of the service can be compromised (similar to suffering a DDoS attack but with legal traffic) and users suffer a bad experience affecting their impression of the company. Users rarely return to sites after being affected by poor quality of service and poor quality of the user experience.

ALTERNATIVES

- Outage during peak traffic.
- Investment in additional hardware to process peak traffic correctly.
- Users making repeated attempts to access the service, which can extend the denial of service period.

F5 | LTM SOLUTION

Thanks to LTM (Local Traffic Manager) standard functionality plus the programmability of F5 (iRules), it is possible to establish mechanisms for the management of traffic spikes, which are well above the capabilities of the servers.

Users who want to access a service at a time of excess demand can be handled directly by BIG-IP, which queues the user (first come, first served) in the form of a «waiting room». This absorbs the flood of requests, ensuring service availability and preserving the quality of the user experience.

REFERENCE ARCHITECTURE | BUSINESS CONTINUITY AND ACCESS CONGESTION CONTROL

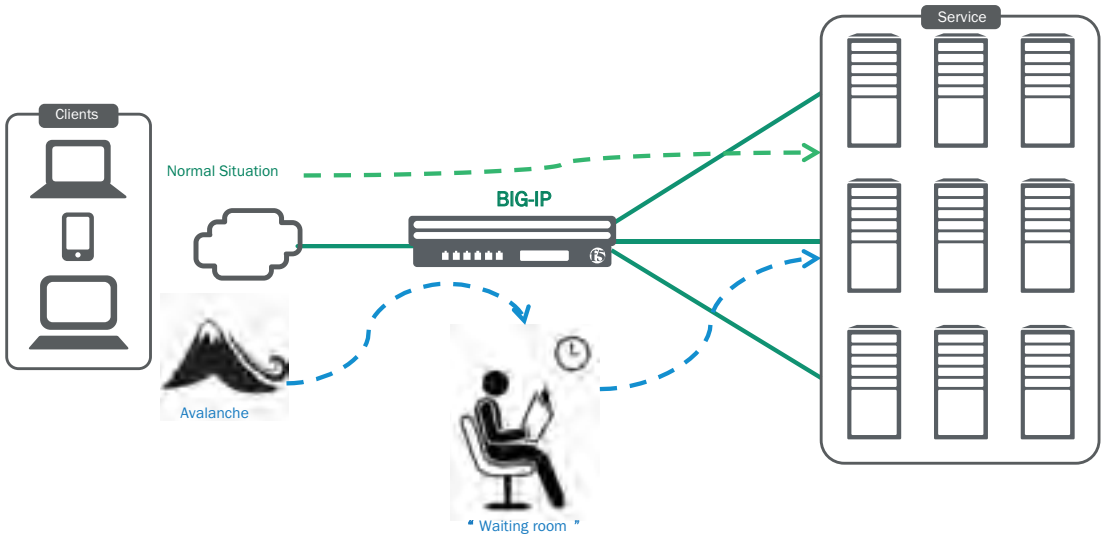

COST CONTROL OF SERVER RESOURCES

PROBLEM

Each year, companies commit to providing more services via the internet with the consequence if significantly increasing spend on new and powerful servers. This means that they are not able to address projects in other areas, such as security, for lack of funds.
This expenditure is increased by the fact that increasingly-complex applications are being developed, which make intensive use of CPU resources and the RAM on servers. In addition, operations such as TLS/SSL encryption or compression, consume many more resources when running on generic operating systems and generic hardware.
Moreover, the current standard of using 2048-bit encryption keys in internet communications has increased CPU consumption on servers by almost 80%, compared with the 1K-bit keys which were used previously.
If we then factor in the energy and hosting costs of these new servers, the costs skyrocket and require specific monitoring.

ALTERNATIVES

- More powerful servers which significantly increase overall cost of a solution.
- Use of specific devices to cache, compress and SSL terminate. This results in more devices to manage and build.

F5 | LTM + AAM SOLUTION

This F5 solution incorporates hardware and software component dedicated reducing the impact of intensive tasks. Techniques used include:

- Balancing/intelligent distribution of traffic. F5 can distribute traffic based on advanced rules, for example to permit one group of servers, (which are less powerful but have more storage) to be used to serve static images and another group to manage dynamic requests, as this imply greater cpu processing.
- Offloading compression and SSL from the servers. This frees up numerous CPU cycles on the servers. Moreover, Crypto Off-loader technology, allows us to scale the performance of SSL traffic
- Caching. F5 caches many objects, which removes the need to make repeat requests for them to the servers. The AAM (Application Acceleration Manager) module makes it easier for users to cache objects in their browsers, removing the need for users to make repeat unnecessary requests.
- Multiplexing Connections. OneConnect is able to open a number of connections to the servers, and multiplex the user connections, thus relieving servers of the management of TCP connections (pooling).

In this scenario, the ROI on the procurement of F5 equipment is very high; it reduces the cost of new servers by up to 60%.

REFERENCE ARCHITECTURE | COST CONTROL OF SERVER RESOURCES

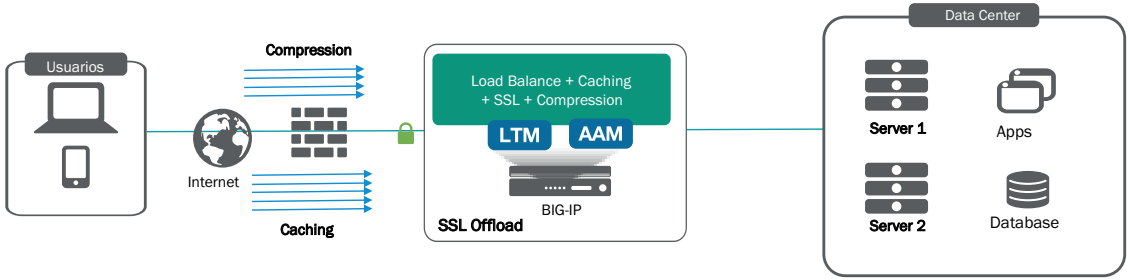

FIREWALLS SCALABILITY

PROBLEM

The traditional approach of deploying firewalls/ NGFWs in clusters does not scale in proportion to the needs of the users and services they protect. In most cases, the scalability of these security solutions is based solely on the over-sizing of the platform and not on a real capacity for increased performance on demand. This has a strong financial and operational impact.

This lack of flexibility means that it is often necessary to replace a security platform completely in order to increase performance, with consequent economic and operational impact. Moreover, the number of connections per second and simultaneous connections is the «Achilles heel» of these security platforms. Both factors are critical in perimeter security environments and often it is the same firewalls/NGFW, which suffer from compromised performance (for example under DDoS).

ALTERNATIVES

- Over-sizing of the firewall / NGFW security platform, with consequent impact on CAPEX/ OPEX.
- Reduction of these platforms' ROI due to the impossibility of real growth on demand.

F5 | LTM SOLUTION

The LTM (Local Traffic Manager) module by F5 makes it possible to balance any element of security, maintaining session persistence and avoiding flows of asymmetric traffic between the elements which make up the group of firewalls/NGFW, using a «sandwich»-type architecture. This architecture allows the deployment of additional elements based on service needs. These new elements can have a different level of performance from the existing ones (and even be from a different manufacturer!). They can also be deployed without loss of service.

All the intelligence required to control the distribution of traffic and to manage the availability of elements is delegated to F5. This intelligence allows you to monitor and protect the resources of the firewalls/ NGFW in order not to compromise performance.

REFERENCE ARCHITECTURE | FIREWALLS SCALABILITY

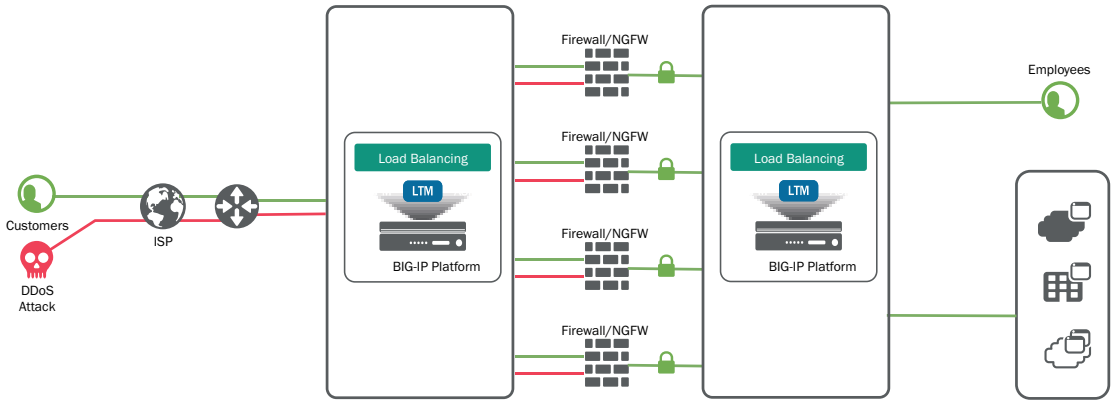

DNS SCALABILTY SOLUTIONS

PROBLEM

The DNS protocol has become a critical point in the operation of the internet, and its importance is set to grow with the progressive adoption of IPv6, in which IP addresses have a length of 128 bits (instead of 32-bit IPv4). The large increase in the number of mobile devices and the imminent arrival of the IoT, have also contributed to the great increase in the number of DNS requests which these solutions must resolve.

Some 41% of the time, loss of web service infrastructure is due to DNS-related problems, so it is essential to maintain the availability of the DNS service. The loss (or degradation) of the service adversely affects service users, leading to loss of revenue and loss of productivity for users attempting to access corporate resources such as e-mail.

ALTERNATIVES

- Risking losing the DNS service, thus endangering business continuity and the productivity of the company's employees, is difficult to justify.
- Adding more and more DNS servers, without intelligent balancing which can monitor server performance, is also not particularly advisable.

F5 | DNS SOLUTION

The F5 DNS solution makes it possible to scale existing DNS solutions in an effective and safe way, making use of diverse capacities, such as balancing and monitoring traditional DNS systems. DNS Express is a technology which is proprietary to F5 and which makes it possible to transfer zones from the traditional DNS infrastructure to a BIG-IP device, where it is served from RAM and with hardware acceleration. Another great advantage of DNS Express from F5 is that it is a proprietary implementation, not based on BIND.

In addition, IP multicast can be used by F5 DNS to distribute the DNS services between multiple data centres. The nearest or fastest connection to the originating query responds to the request.

REFERENCE ARCHITECTURE | DNS SCALABILITY SOLUTIONS

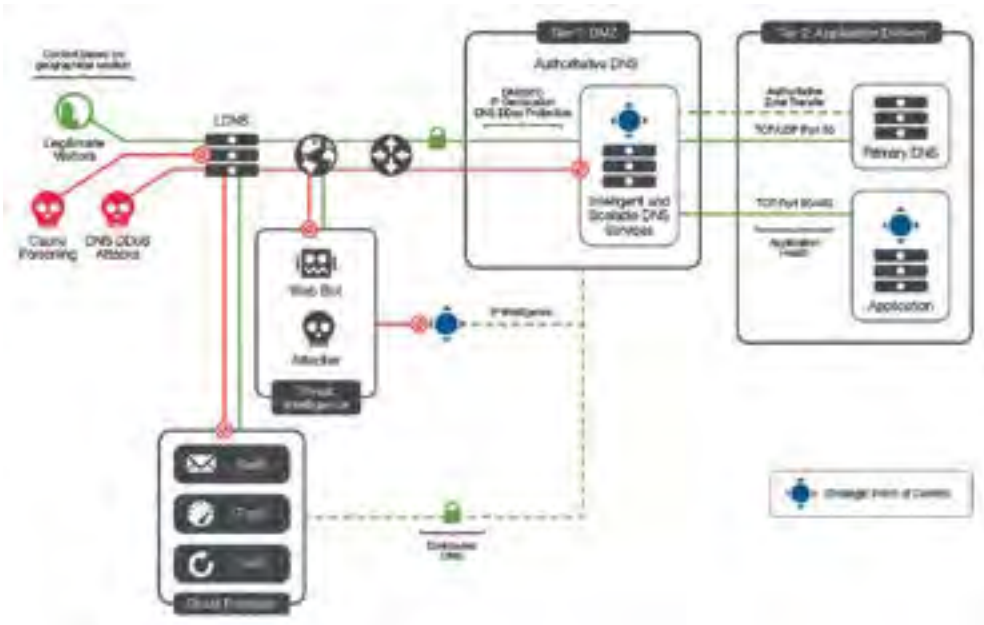

FEDERATION AND SIMPLIFICATION OF VDI ENVIRONMENTS

PROBLEM

VDI environments are often used to make it possible for users to connect to enterprise environments without the need for heavy, expensive end user devices. However their use entails other problems implicit in the architectures of this type of technology.

Each VDI environment requires a user authentication element, another element to present the options available for each user and finally an element, which acts as a «broker» and which manages access to the VDI.

In an environment with multiple VDI solutions there is the need to replicate this architecture for each one of them adding expense for the initial implementation and increasing cost of ownership through the need to maintain multiple environments.

ALTERNATIVES

- Having VDI solutions from multiple vendors, which represents an increase of CAPEX and OPEX, which is difficult to justify.
- The use of proprietary elements in VDI solutions reduces the flexibility of the solution from the users' perspective and can increase costs.

F5 | APM SOLUTION

The F5 APM (Access Policy Manager) solution authenticates users and establishes SSL-secured VPN tunnels to the network, replacing the security-gateway solution in VDI environments.

In addition, APM presents the user with different options on the screen, consolidating functionality onto a single platform therefore saving the need for this element to be incorporated in the VDI solution.

APM is independent of the VDI solution implemented, making it possible to standardize the «look and feel» of the solution to the client, presenting the same format and options on the screen regardless of the VDI solution deployed.

REFERENCE ARCHITECTURE | FEDERATION AND SIMPLIFICATION OF VDI ENVIRONMENTS

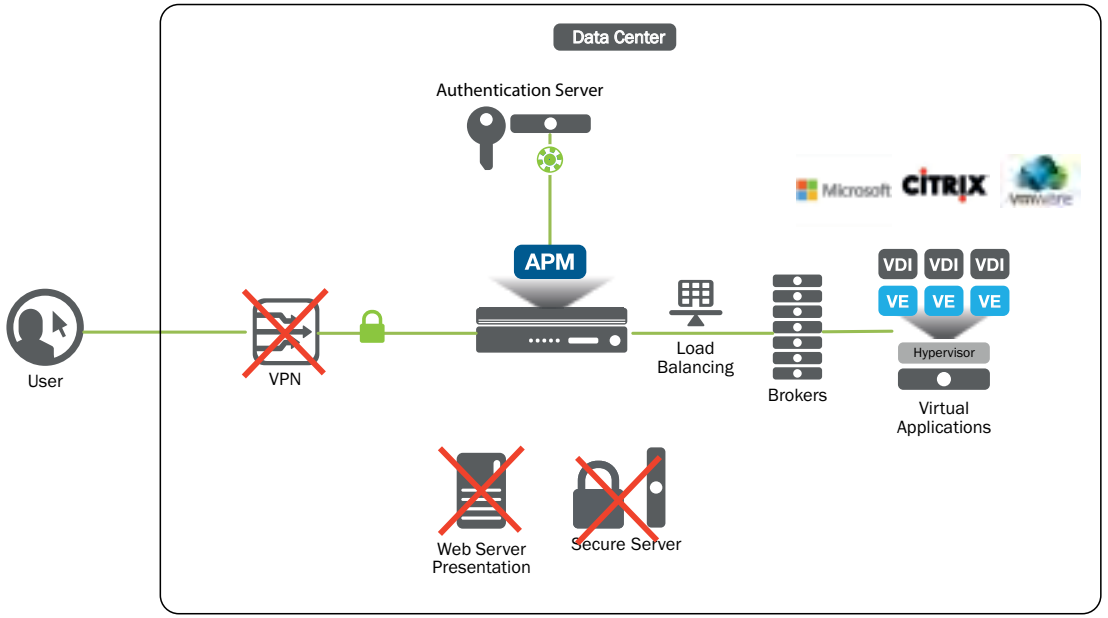

INTEGRATION OF HYPERCONVERGENT PLATFORMS – NUTANIX

PROBLEM

Integrated systems have become a popular choice of architecture in companies, thanks to their promise of rapid deployment, simplified management and efficient operation. The first wave of «converged infrastructure» technology provided a combination of computing and preconfigured storage. This provides organizations with CAPEX and OPEX savings when compared with the traditional infrastructure approach in which the organizations obtained their hardware infrastructure and their software independently of each other. Converged infrastructures also offer a simplified model for expansion based on 'infrastructure blocks', which provide linear scalability.

Recently, hyper-convergent platforms have abstracted hardware settings by increasing the (software-defined) virtualization of the infrastructure, such as storage and networks. These solutions have the potential of offer even greater cost savings and greater flexibility.

The evolution of these architectures, however, has changed the fundamental nature of the applications which run on them. Applications still require application services on layers 4-7 (e.g. application delivery, the security of the network and application layers or acceleration services). You will still need to supply these services on top of your converged or integrated systems. The challenge is to design the best way to insert these services into these new architectures.

ALTERNATIVES

- The trend is that 54% of companies currently without hyper-convergence are planning a deployment in the next 24 to 36 months.
- Adding the additional L4-7 services provided by F5 ensures Fast, Available and Secure application can be delivered on this architecture.

F5 | LTM + AAM + APM + ASM SOLUTION

By implementing solutions by F5 and Nutanix, organizations can gain the benefits of web-scale IT architecture without sacrificing application availability, performance and security.

The combined solution improved data centre security by:

- Harnessing the agility and scalability of the cloud and combining it with the security of on-site solutions.
- Protecting critical applications, no matter where they reside, with a comprehensive data centre security solution.
- Preventing unauthorized access to the network.

REFERENCE ARCHITECTURE | INTEGRATION OF HYPERCONVERGENT PLATFORMS – NUTANIX

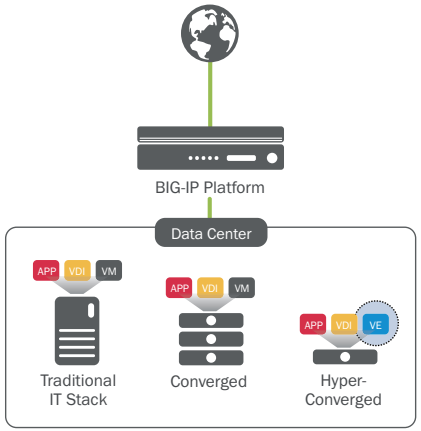

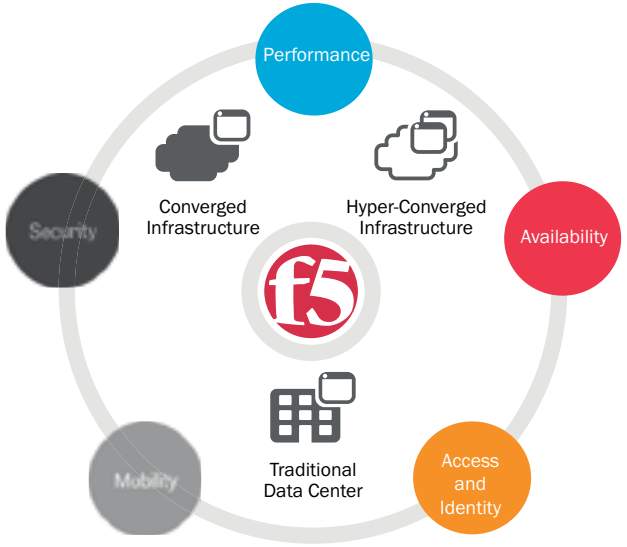

MIGRATING EXCHANGE ENVIRONMENTS

PROBLEM

The migration of Exchange versions and environments can be very complex. Customers have to make these migrations gradually and need to synchronize changes made on servers with changes in client computers in order to accomplish a successful migration.

If the Exchange service has no default configuration, this migration can be even more complex, for example if two-factor authentication has been added and therefore HTTP access cannot be migrated in a way which is transparent to the customers.

The case of migration to Exchange 2013 presumes a change in the customer's infrastructure. We must take into account the protection of remote access, the preparation of the infrastructure in order to consolidate server functions, virtualization and the configuration of network. All this represents a major challenge for companies, which are also facing a loss of service until the migration has been completed.

ALTERNATIVES

- Manually migrating users from one environment to another, plus synchronizing with changes in the back-end servers.

F5 | APM SOLUTION

The F5 solution performs the pre-authentication of users using an APM (Access Policy Manager) solution. This pre-authentication may incorporate other factors, such as additional authentication certificates, one time passwords (OTPs), or security checks on the client computer.

Once the user has been authenticated, the APM checks the directory for the configuration of the user's mailbox and can use this to redirect access to the relevant group of servers. This can be used, for example to migrate from Exchange 2007 to 2013 without changing anything on the client computer and without the need to stop the service.

And thanks to our partnership with Microsoft, the release of Exchange can be done easily by using a template on BIG-IP by F5, which is specific for Exchange and which will allow us to make the configuration in a much more agile way.

REFERENCE ARCHITECTURE | MIGRATING EXCHANGE ENVIRONMENTS

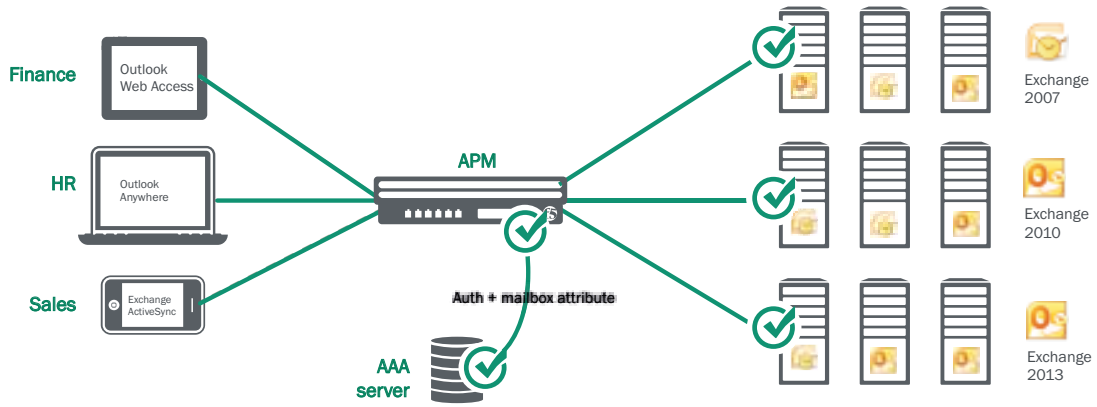

OPTIMIZATION OF DC ACCESS BASED ON GEOLOCATION

PROBLEM

Global companies often have different data centres distributed geographically. It is widely assumed that the most optimal access point for a user is always the closest data centre, since such communication involves the fewest jumps.
However the closest data centre is not always the most responsive one. For example, if we know that the launch of a new marketing campaign will flood the data centre located in Europe in the morning, but we have a second data centre in America, which is underused during that time-slot, then it makes sense to direct users of that campaign to there and to make use of the existing investment without incurring additional costs.
Some organizations are now beginning to consider cost optimization when considering how to maintain their processing capacity. Some companies are starting to make the decision to turn off or turn on services in different locations depending on the weather (50% of the energy consumption of a data centre is for cooling), or based on the use of electricity tariffs, as an example, consumption is low during the night hours, resulting in a constant shift of computing resources between different time zones using a strategy known as «follow the moon».

ALTERNATIVES

- Failure to make use of a global rolling method such as geolocation presupposes inefficiency in business costs, investing where capacity is needed and ignoring unused resources elsewhere.

F5 | DNS SOLUTION

The DNS solution enables intelligent traffic-balancing between different data centres in both Active-Active and Active-Passive deployments.

Once certain criteria (business, technical or cost as described above) have been determined and the users and applications involved have been defined, F5 provides automated mechanisms to enable this access, redirecting requests from these users to the optimal DPC at all times.

The solution also has a geolocation database, which includes specific methods for global balancing (e.g. Round Robin, topology, Global Availability, CPU, server capacity, Packet Rate, Completion Rate, Least Connections, Persistence Round Trip Time, VS score, QoS).

REFERENCE ARCHITECTURE | OPTIMIZATION OF DC ACCESS BASED ON GEOLOCATION

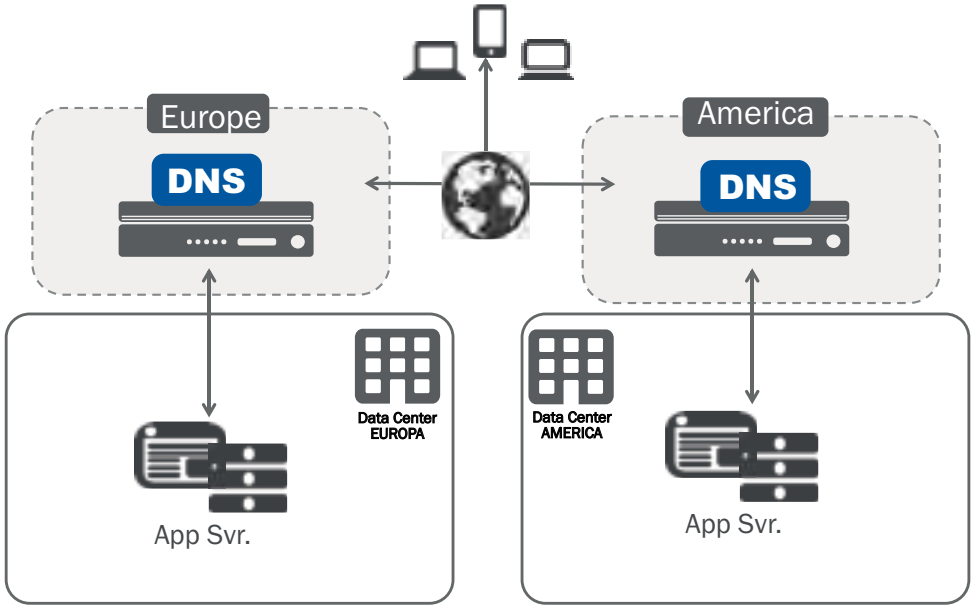

OPTIMIZED CACHE BALANCING (CARP)

PROBLEM

In services dedicated to serving content, such as balancing proxy caches, it is important to optimize the ratio of cached content. Traditional algorithms stay at the origin and are not effective for these scenarios.

Take, for example, a video-on-demand service; if the content of video A is in cache 1 and this has a service outage, then the traditional algorithm, which is based on the hash of origin, is recalculated, losing the information regarding the cached content and causing all, or part of the content to be re-ordered from the original content repository, causing unnecessary network usage/bandwidth and delay in the content-serving service.

What's more, if cache 1 rejoins the service, the balancer does not remember that video A was accessible through that cache. Therefore, a more intelligent algorithm is required for this type of service, one which stays in the requested video, to maximize the ratio of caching and optimize response times and the success rate in caching content, this also creates a more scalable service

ALTERNATIVES

- No other manufacturer of ADC is capable of implementing the CARP algorithm. The alternative is to have a content-delivery system, which is less efficient in terms of scalability, network consumption, cache efficiency and response time.

F5 | LTM SOLUTION

F5 implement the CARP (Cache Array Routing Protocol) algorithm, which is highly recommended for services which balance proxy caches. This algorithm stays at the destination and is not recalculated when a cache node joins or leaves the pool of caches. Thanks to CARP, the balancer «remembers» in which cache the requested content resides, maximizing the percentage of cached content and minimizing network usage and response time.

Furthermore, this algorithm does not need to save anything in memory (stateless selection), which means that it consumes very little of the balancer's resources and does not need to be replicated in a high-availability system.

REFERENCE ARCHITECTURE | OPTIMIZED CACHE BALANCING (CARP)

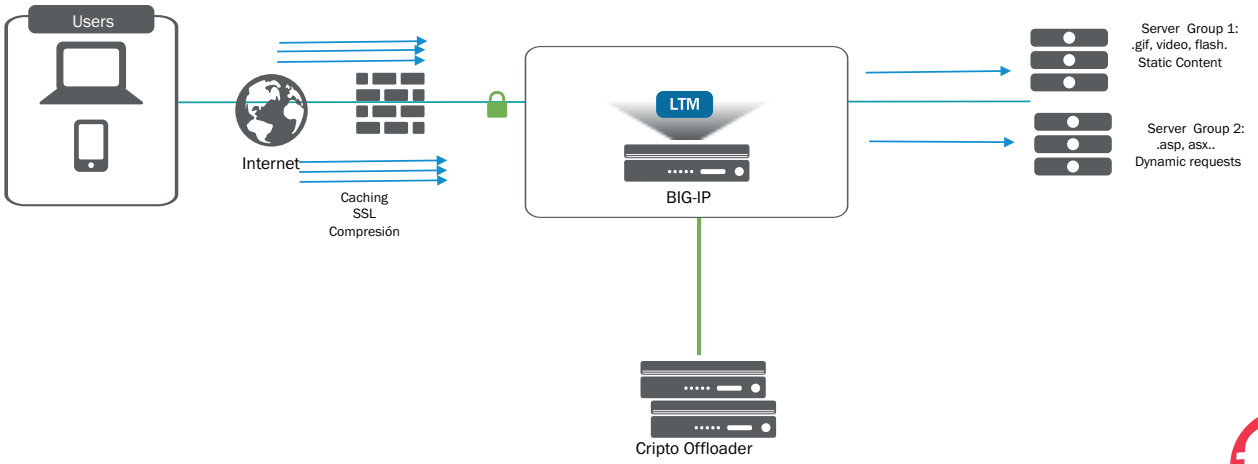

TCP OPTIMIZATION

PROBLEM

The TCP protocol was created in the mid-1970s with the aim of providing connection-oriented communications and thus to guarantee that data packages would be delivered without errors and in the same order they were transmitted. Mechanisms of congestion control and flow control are an essential feature of TCP and their goal is to maximize the rate of data transfer and to avoid network congestion. In wired networks, TCP identifies packet loss through network congestion.

The wireless data networks (2G, GPRS, 3G, 4G, etc.) have a very high rate of package loss (compared to the wired networks), which has nothing to do with network congestion, and also have highly-variable latency (300ms on 3G and 50ms on 4G). These two features mean that navigating the web from mobile devices can be unsatisfactory for users.

ALTERNATIVES

- The increase in mobile devices, and the intensive use made of them, makes it necessary to improve the users' web-browsing experience (HTTP / HTTPS) through these devices.
- Applying generic optimization profiles to all users/subscribers and to all types of data networks (fibre, cable, 3G, 4G, etc.) is not a real option, due to the very different characteristics of these networks. Sometimes the cure is worse than the disease.

F5 | LTM + PEM SOLUTION

Thanks to its Full-Proxy technology, BIG-IP LTM (Local Traffic Manager) is the ideal platform for implementing TCP optimization. LTM can apply specific optimization to profiles depending on the type of wireless data network (2G, 3G, 4G, etc.) to which users connect, in order to maximize the experience of those using mobile devices. These profiles can be applied at the subscriber level and change dynamically (for example, when subscribers roam between 3G and 4G networks), thanks to the PEM (Policy Enforcement Manager) module.

In addition to supporting the most common TCP-optimization algorithms (Vegas, Westood, Illinois, H-tcp), F5 has developed its own algorithm called Woodside.

Optimizing TCP (L4) also optimizes the protocols used in the superior layers (for example HTTP, HTTPS and SPDY) without having to make modifications at the application layer.

REFERENCE ARCHITECTURE | TOP OPTIMIZATION

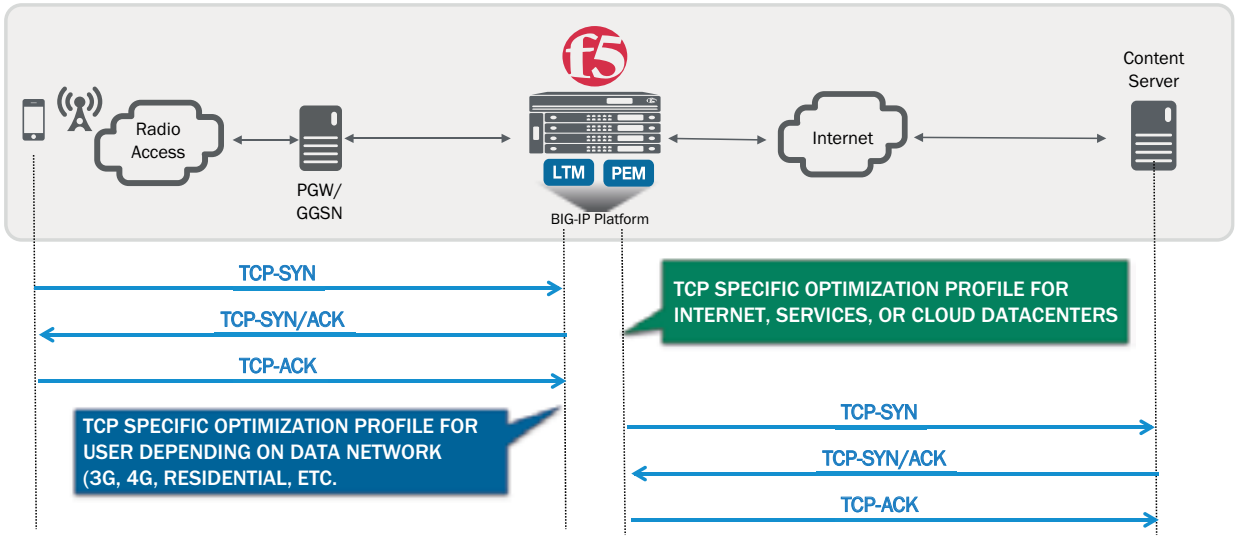

REWRITING GATEWAY DOMAINS

PROBLEM

When companies and organizations publish services on the internet/intranet, it can give rise to the following problems relating to domain usage :

- Proprietary applications: in most cases, programmers do not contemplate the need to differentiate between internal and external addresses and introduce direct references to internal URLs, which are not reachable from the outside.
- Commercial applications: often it is necessary to use a pair of reverse proxies for each application, these perform the rewrites and sometimes undertake some pre-authentication work. Each application typically requires the deployment of at least two proxies, which are also not compatible between different versions of different commercial solutions.

ALTERNATIVES

- Deployment of multiple systems of reverse proxies on a per-application basis, such as Microsoft TMG, ISA servers, BlueCoat, or other business solutions such as WebMarshal, WebSEAL, etc.
- The problem with these solutions is that they do not allow for the consolidation of functionality, there is no guarantee of the level of performance, they require general-purpose operating systems and do not support recent encryption requirements, etc.

F5 | LTM SOLUTION

The F5 LTM (Local Traffic Manager) product enables real-time rewriting of domains and applications without the need to deploy new elements or reverse proxies.

F5 also allows for efficient management, which is optimized at a single point on the encrypted traffic, undertaking the downloading and compression of traffic through hardware which is specifically designed for this purpose.

REFERENCE ARCHITECTURE | REWRITING GATEWAY DOMAINS

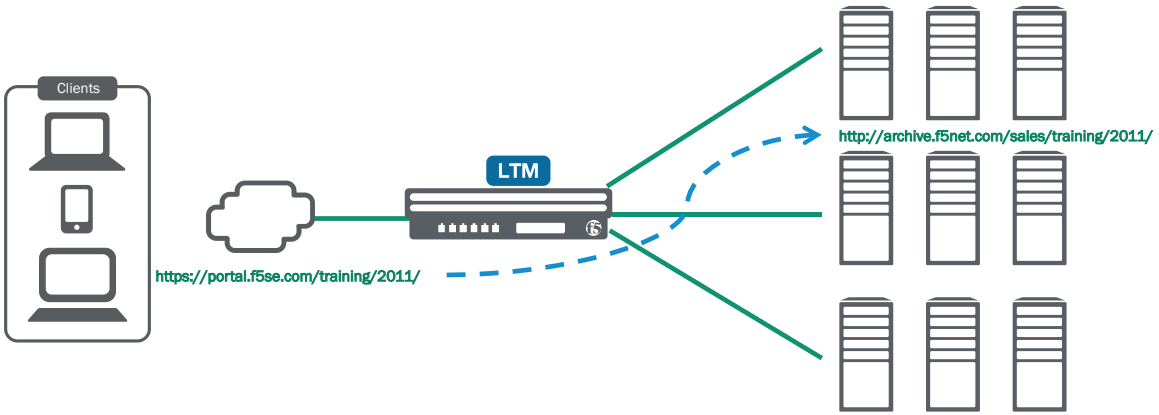

HTTP/2 GATEWAY

PROBLEM

HTTP/1.1 has been in existence since 1999. Since then, the internet has changed radically. There are now more than one million websites, the average item of content is 45 times heavier than in 1999 and users are increasingly demanding services such as video-streaming (HTML5, etc.). HTTP/1.1 was not intended to support this evolution.

Content providers are forced to invest continuously in their platforms without substantially improving the user experience. Up until now, it has not been possible to address the real source of the problem - the inefficiency of HTTP/1.1. A change is needed.

ALTERNATIVES

Not adopting efficient protocols such as HTTP/2 means that web services cannot benefit from the best user experience. This means that the performance of these web services is at a disadvantage compared to the services offered by those who have already adopted these improved protocols. This impacts negatively on the business.

F5 | LTM SOLUTION

HTTP/2 solves many of the problems which have been present since 1999. The need to optimize online web traffic has led to the creation of new protocols such as SPDY and HTTP/2.

F5 offers the native capability of translating the HTTP/2 information originating from the client web browser into HTTP/1.1, which currently runs on web servers (Apache, IIS, etc.).

Although HTTP/2 solves the problems of efficiency seen in HTTP/1.1, it is a binary protocol, which renders obsolete the current reporting platforms and web analytics. The translation undertaken by F5 in a way which is transparent to the HTTP/1.1 infrastructure, means that the service life of these platforms can be extended, thus protecting the ROI.

REFERENCE ARCHITECTURE | HTTP/2 GATEWAY

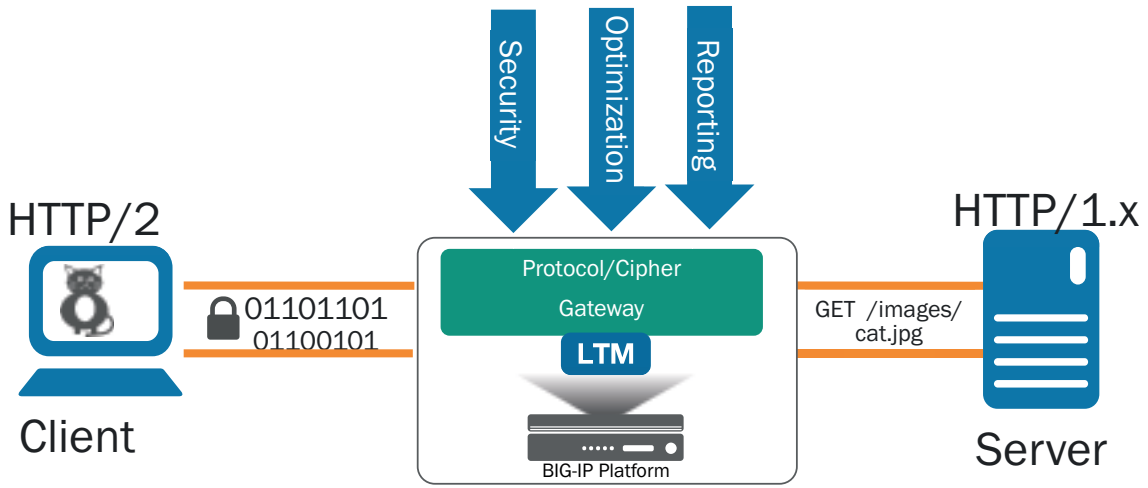

IPV6 GATEWAY

PROBLEM

In January 2011, IPv4 addresses began to run out, which gave rise to a need to undertake a process of transition to the new IPv6 protocol. Today many «legacy» systems and devices still do not support IPv6.

Conversely, many operators are already providing their subscribers with IPv6 addressing; publishing content in IPv6 has become a reality. It is not easy to implement an infrastructure in which both protocols co-exist until the full migration to IPv6 is complete, and this has high implementation and re-engineering costs.

ALTERNATIVES

- Undertake a migration of the entire infrastructure to IPv6, which has an impact on CAPEX/OPEX, as well as on the operation and security of the infrastructure itself.
- Big companies like Google, Facebook, Microsoft and Apple are driving users to migrate to IPv6, so ignoring this change will end up reducing visibility and business value.

F5 | LTM SOLUTION

The F5 LTM (Local Traffic Manager) solution supports IPv6 natively, and acts as a gateway between the two protocols, allowing a dual-stack architecture which translates the traffic in a bi-directional manner (IPv4 - IPv6).

LTM makes it possible to publish services in IPv6 while the client infrastructure is still on IPv4, allowing a gradual migration, without any impact to users and ensuring the ROI of existing infrastructure.

REFERENCE ARCHITECTURE | IPV6 GATEWAY

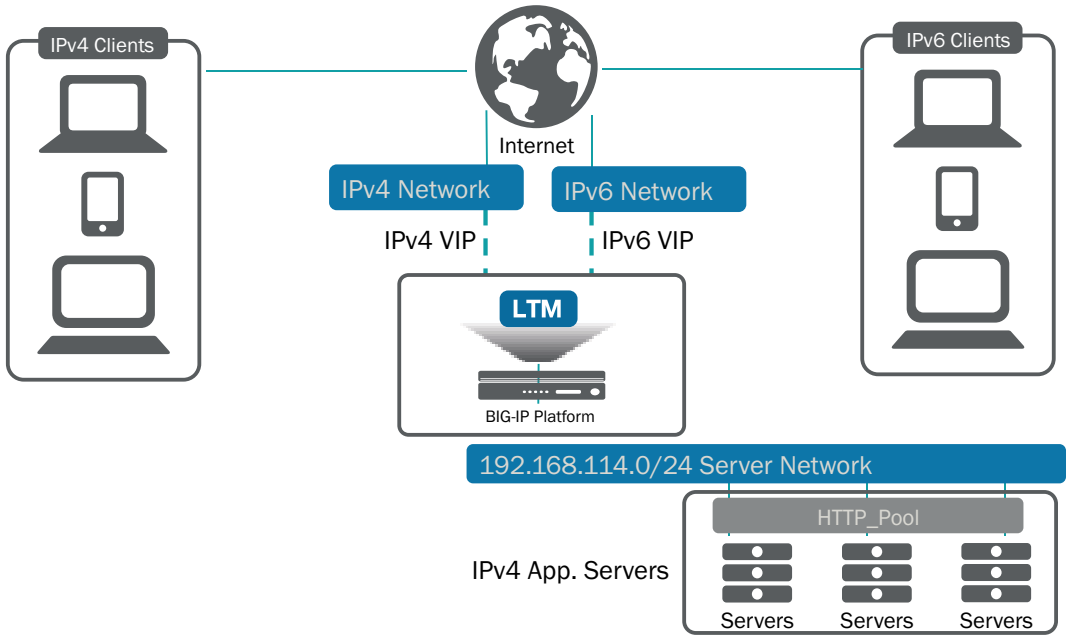

GEO-LOCATION SOLUTION WITH EDNS

PROBLEM

Since its development (in the early 1980s), DNS has been improved with new features while maintaining compatibility with earlier versions of the protocol. In 1999 there was a proposal to expand the DNS to support new parameters and response codes, allowing longer answers, while maintaining a framework compatible with previous implementations. EDNS (RFC6891) adds information to the DNS messages in the form of pseudo-Resource Records (pseudo-RR) of OPT type, which are included in the additional data section of the DNS message.

When users use global DNS servers (such as Google), it can impact on the functionality of services based on geo-location, because requests to the authorization servers are made from these global DNS servers. In the absence of the actual location of the original client, DNS responses may not be optimal, and may even invalidate these services.

ALTERNATIVES

- The use of global DNS servers is a common practice among internet users, so not implementing these mechanisms of extending DNS (EDNS) jeopardizes the accuracy and performance of services based on geo-location.
- Using CDN, which can lead to increased operating costs for the service.

F5 | DNS SOLUTION

F5's DNS module can insert, read and operate the EDNS' OPT records.

Thanks to F5's DNS module, it is possible to insert an OPT record in the client's original DNS request, which contains the public IP address of the computer. The decision will be based on the IP address contained in the OPT record (instead of using the IP address in the original request) and the correct resolution will be returned to another OPT record.

REFERENCE ARCHITECTURE | GEO-LOCATION SOLUTION WITH EDNS

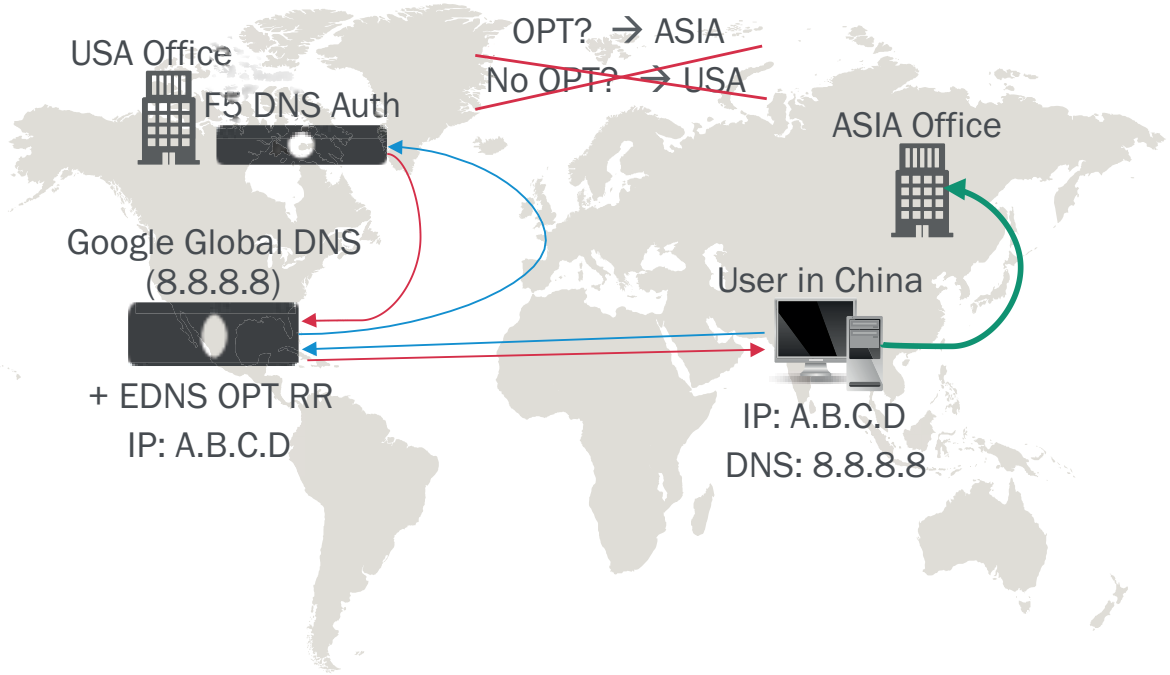

HA AND DC SOLUTION IN HUB & SPOKE TOPOLOGIES

PROBLEM

Implementing high-availability mechanisms in Hub & Spoke topologies can be technically complex and can also involve significant investment in equipment and high operating costs.

ALTERNATIVES

- Providing redundancy for all the equipment at each site data centre or remote site (spoke) implies a high potential cost of purchasing equipment.

F5 | LTM + vCMP SOLUTION

F5's vCMP (Virtual Clustered Multiprocessing) technology makes it possible to segment a device, assigning physical resources (CPU and RAM) to various guests ensuring their performance and isolation.

The solution is based on the deployment of a single F5 BIG-IP devices in remote locations (spokes), and provides redundancy by using a single F5 BIG-IP device in the central data centre (hub), which uses vCMP to segment the unit into various «guests», one for each remote site.

In case of the failure of one of the devices in the remote data centres, the corresponding vCMP «guest» will take up an active role, thus ensuring continuity of service.

REFERENCE ARCHITECTURE | HA AND DC SOLUTION IN HUB & SPOKE TOPOLOGIES

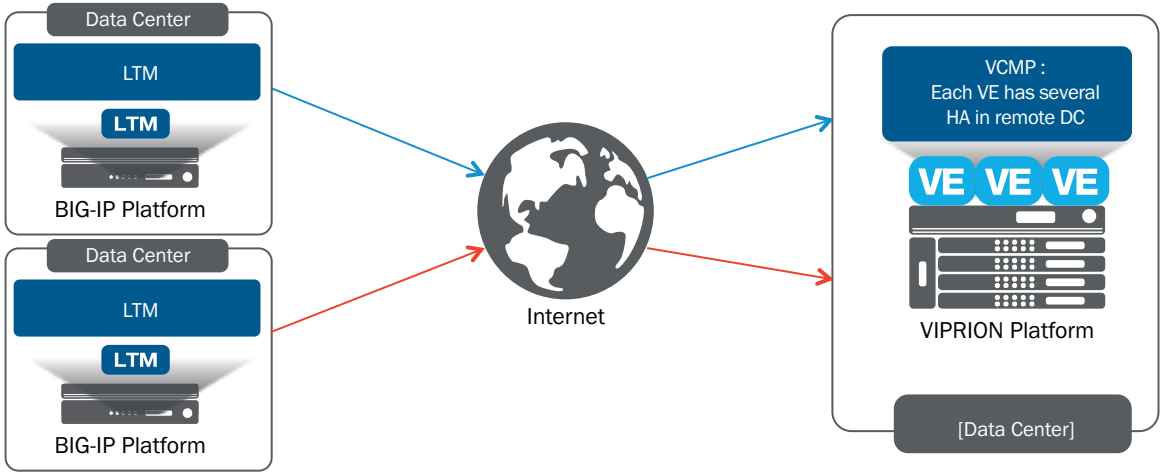

INTERCONNECT SOLUTION FOR OVERLAPPING NETWORKS

PROBLEM

The interconnection of networks with overlapping IP addressing is very common in mergers and/or acquisitions. This interconnection creates several problems for IT departments:

- Interconnections should be operational within the time limits, with no impact to the business/ service, and controlled costs.
- The existence of «legacy» services in which the IP addressing cannot be changed easily or which use an IP address which is embedded in applications (instead of using DNS resolution).
- Need to control the flows of traffic between the overlapping networks of both environments.

ALTERNATIVES

- Re-engineering of IP addressing. Places serious demands on time, money and resources, and does not guarantee a solution for all scenarios.
- NAT solution: are expensive, have little flexibility (traceability, logging, etc) and have scalability issues in terms of pools of IP addresses and ports. Standard NAT solutions have a high impact on CAPEX/OPEX.

F5 | LTM + AFM SOLUTION

LTM makes it possible to create simple rules to translate between the overlapping networks, keeping the host addresses and ports of origin, and controlling the traffic between the two networks with minimal intrusion and in a way which is transparent to applications and users.

This solution has unlimited scalability (there are no pools of IP addresses or ports) and near-zero impact on network performance.

It allows the configuration and export of logs to be completely defined by administrators, to ensure the traceability of connections.

Dynamic routing protocols (OSPF, BGP, etc.) are supported, allowing deployments to be integrated with the existing infrastructure in both environments.

Finally its role in the network means it is the optimal place to enforce security policies (AFM - Advanced Firewall Manager) to control the flow of traffic between the networks in both environments.

REFERENCE ARCHITECTURE I INTERCONNECT SOLUTION FOR OVERLAPPING NETWORKS

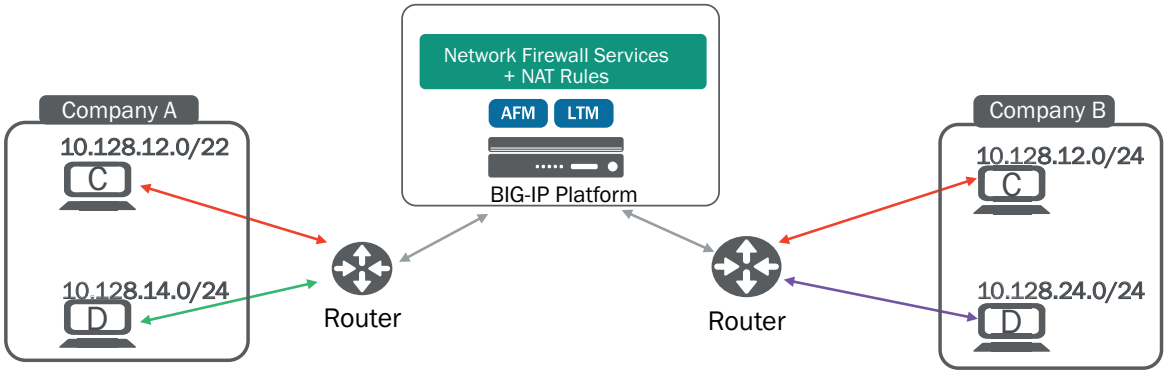

DATA CENTRE PERSISTENCE SOLUTION

PROBLEM

There are applications which require session persistence, even when the service is offered from multiple data centres. When users are roaming and may be using separate public LDNS services, it is necessary to establish mechanisms that ensures the user remains in the original data centre during their session.

ALTERNATIVES

Currently there is no alternative to this solution. At this time, companies are not able to maintain the user session so the session is reset.

F5 | LTM + DNS SOLUTION

Thanks to the persistence functionality offered by BIG-IP DNS, we are able to set limits on IP ranges, permitting the next DNS operator to maintain the same persistence.

LTM also contributes to the persistence of the data centre by inserting a cookie in the browser to keep the information regarding the data centre and session.

REFERENCE ARCHITECTURE | DATA CENTRE PERSISTENCE SOLUTION

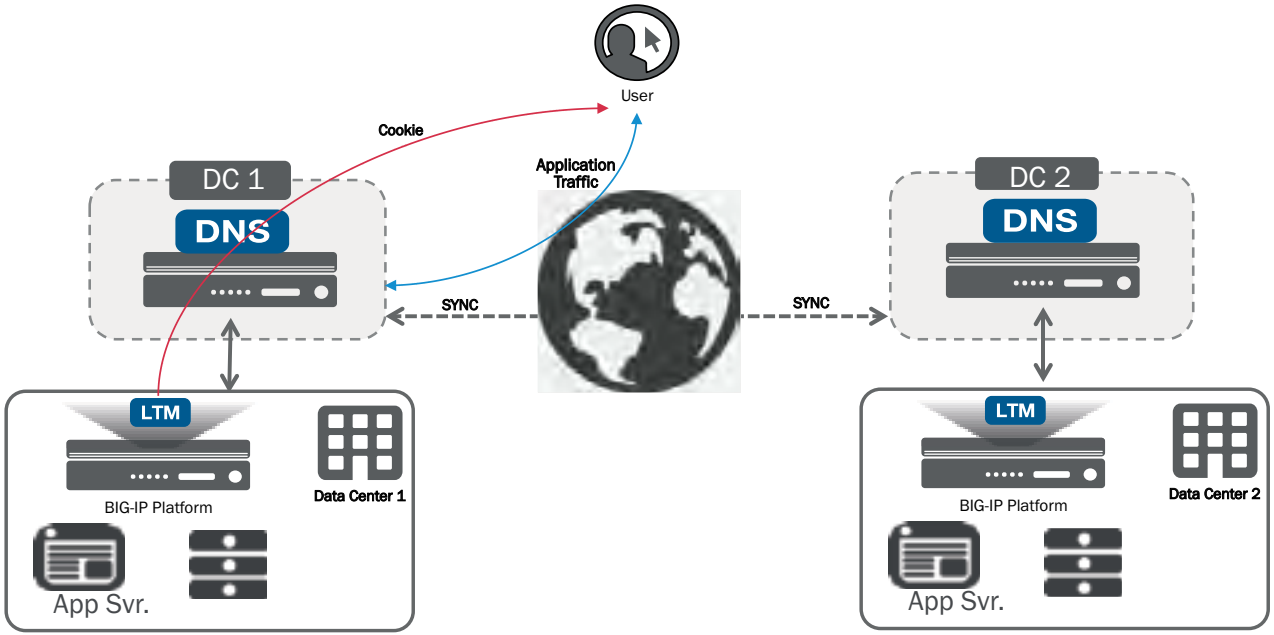

CUSTOMIZED PORTALS SOLUTION

PROBLEM

In many cases, it is necessary to treat different customers in different ways according to their profile, whether it be the type of access to the available applications, or to different advertising blocks, according to the products or services which the client has contracted.

The management of this logic is very complex and has an impact on the programming of applications, thus requiring development time, testing, associated costs, etc.

An example of this customization might be that if a user has a «gold» loyalty card, the banners and advertising offers which are shown to them may be personalized and the content shown will be different to that seen by users who have a «silver» loyalty card. Implementing this logic requires dynamically adapting content in the application and updating these adaptations can be extremely complex and/or costly.

ALTERNATIVES

- Ad-hoc developments, which are non-reusable and which require applications to be modified.
- Development costs and time required to perform these customizations.

F5 I LTM SOLUTION

F5 (LTM) makes it possible to redirect the user traffic to individual services according to the client/user profile. This feature does not impact the development elements of the application and the fact that it is performed ahead of the applications means that it is totally transparent to them.

Moreover, the fact of not requiring a schedule, allows a huge reduction of «Time to Market» and reduces the costs of implementation and deployment.

REFERENCE ARCHITECTURE I CUSTOMIZED PORTALS SOLUTION

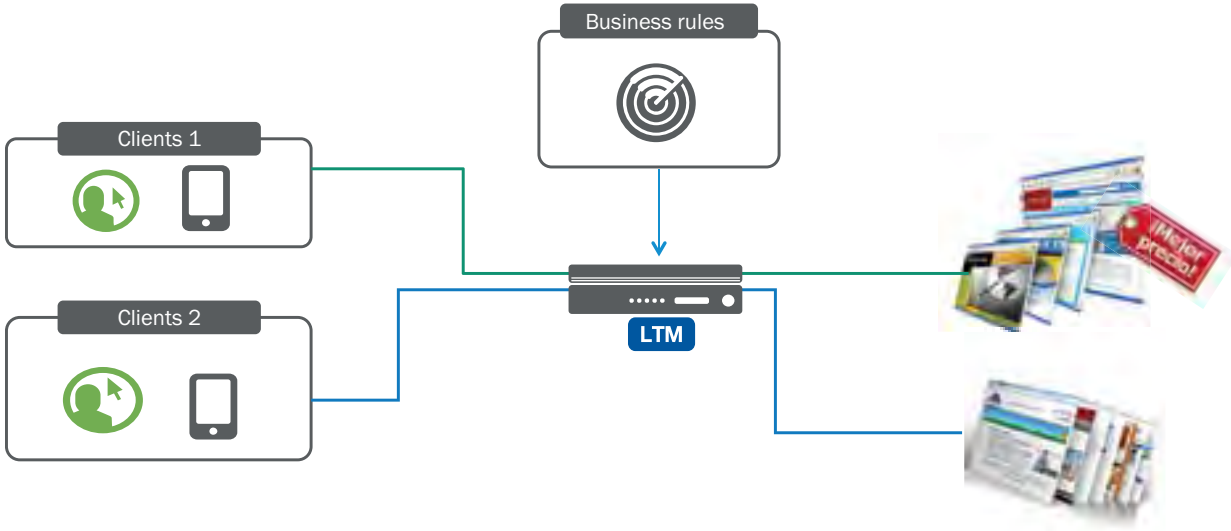

DNS64 SOLUTION

PROBLEM

The progressive introduction of IPv6 means that the coexistence of IPv4 and IPv6 networks is a reality which will continue to exist for quite some time.

Migration to an IPv6 world should be a transparent process for users and this transition should definitely not lead to the loss of access to IPv4 services for users who already have IPv6 addressing.

While the configuration of IPv6 among users is gradually becoming a common practice, not all services are available in IPv6 (and many never will be). Communication between IPv6 and IPv4 worlds has become a problem.

ALTERNATIVES

- Blocking IPv6 users from «legacy» IPv4 services is not a real alternative.
- Updating and configuring traditional DNS platforms to support DNS64 can be a complicated and expensive task.

F5 I DNS SOLUTION

DNS64 is a mechanism for creating DNS records in the form of AAAA (IPv6) from records of type A (IPv4). This domain name resolution between IPv6 and IPv4 is a key element in the communication between clients which only use IPv6 and servers which only use IPv4, without requiring changes to any of them.

The DNS module can perform the DNS64 functions in a manner which is both autonomous and transparent to the standard DNS solutions, which have already been deployed; alternatively it can also be integrated and used to complement another third party DNS64 solution. In the latter case, the BIG-IP intercepts, manages and optimizes DNS client requests in IPv6 before sending them to external DNS64 servers.

REFERENCE ARCHITECTURE I DNS64 SOLUTION

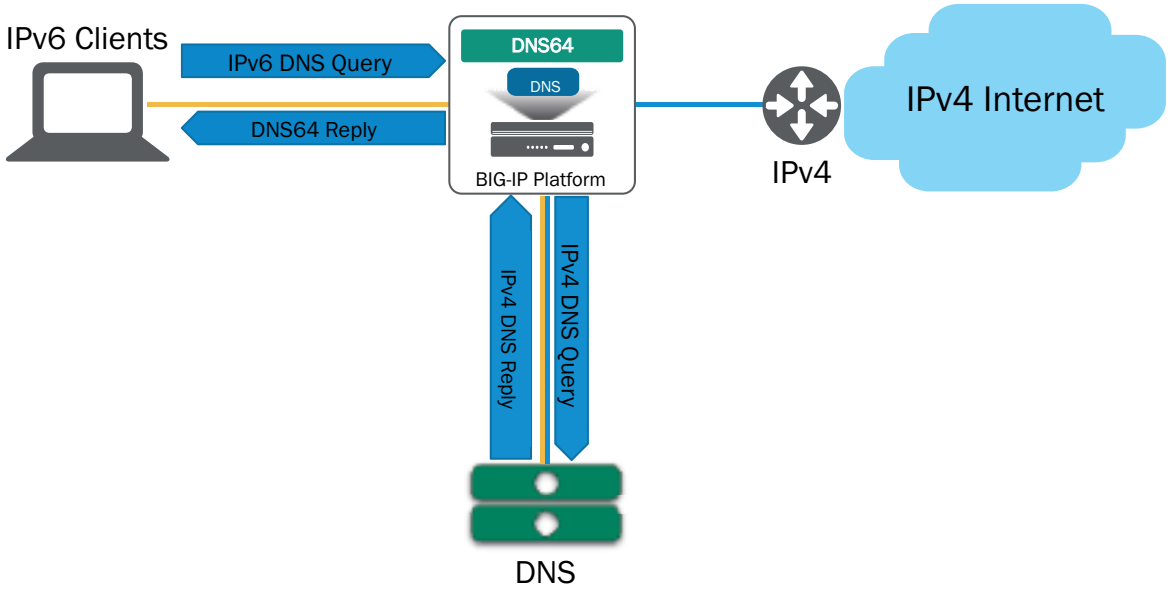

MULTI-LANGUAGE SOLUTION IN WEB ENVIRONMENTS

PROBLEM

The end goal of technology in a company is to be aligned with its business. Therefore this technology needs to be programmable, adaptable and powerful, so that it can move into alignment in a short time. Two examples to illustrate this requirement would be the creation of a multilingual website or the insertion of personalized advertising according to the customer profile in a marketing campaign. In the first example, what is intended is that the application changes the displayed language according to the different criteria of application users, such as geolocation, the browser's header user-agent, or any field in the user repository (LDAP, Active Directory...).

In the second example, we want to publicize one of the company's products (e.g. for banks, car, life or home insurance, mortgages, loans...) by inserting an advertising banner in our application. This banner needs to have a certain degree of intelligence, so that it can use the profile of the client in question to determine which products they have already purchased and show ones which they have not yet bought and/or those which are compatible with what they already have.

This type of intelligence in technology is vital to achieving these clear business objectives within the company.

ALTERNATIVES

- While it is true that certain manufacturers have tried to plagiarize iRules technology, they are far from what has actually been achieved by F5, as the scripts are much more powerful and operate at many layers of traffic. Moreover, none of them has a community which can compare to DevCentral.
- The alternative to iRules technology is applying this intelligence in the same application, which requires an extended time commitment for the development and maintenance of new code.

F5 | BIG-IP SOLUTION WITH iRULES

iRules technology in BIG-IP, is a very powerful scripting language based on TCL. It makes it possible to implement the intelligence required in the examples described above in a short time and little programming knowledge.

iRules technology works in all BIG-IP modules (LTM, APM, ASM, AAM, AFM...) and of course it is one of the main reasons why F5 is leading the ADC market.

It is also important to have a strong and active community to support this technology, which in this case is DevCentral. This community has more than 100,000 registered users, who exchange experiences, ideas, documentation and code for F5's most programmable technologies: iRules, iControl, IAAP, iRules LX, etc, etc.

REFERENCE ARCHITECTURE | MULTI-LANGUAGE SOLUTION IN WEB ENVIRONMENTS

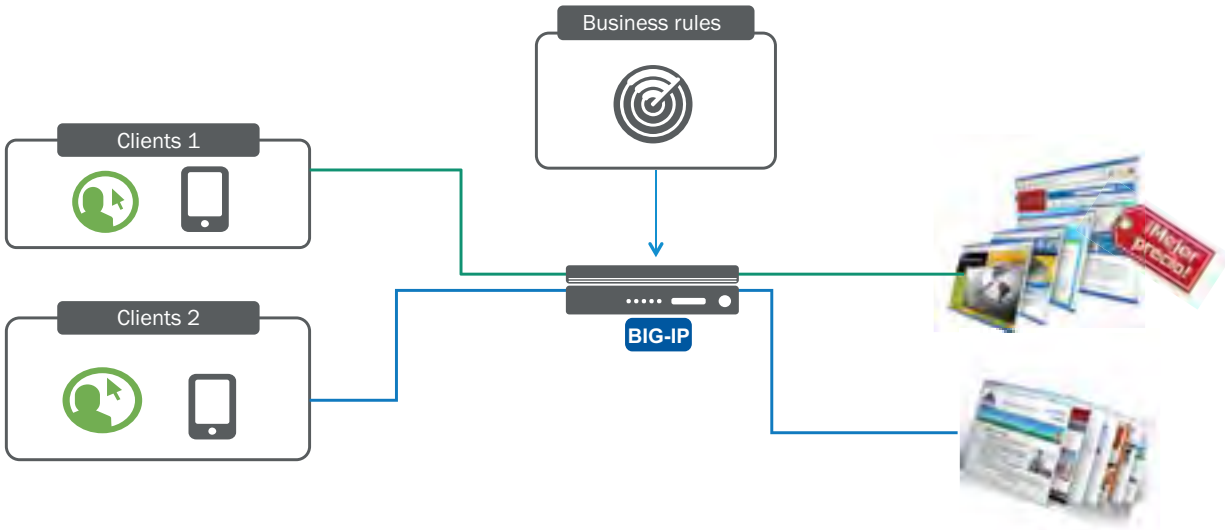

Adapted content

NAT64 SOLUTION

PROBLEM

The progressive introduction of IPv6 means that the coexistence of IPv4 and IPv6 networks is a reality which will continue to exist for quite some time.

Migration to an IPv6 world should be a transparent process for users and this transition should definitely not lead to the loss of access to IPv4 services for users who already have IPv6 addressing.

While the configuration of IPv6 among users is gradually becoming a common practice, not all services are available in IPv6 (and many never will be). Communication between IPv6 and IPv4 worlds has become a problem.

ALTERNATIVES

- Blocking IPv6 users from «legacy» IPv4 services is not a real alternative.
- The alternative is to implement NAT64 in traditional NAT/CGNAT platforms, losing the consolidation capabilities offered by F5's BIG-IP.

F5 | LTM/CGNAT SOLUTION

IPv6 clients who want access to «legacy» IPv4 services need a solution which translates the IPv6 address into an IPv4 address. This solution is NAT64.

The NAT64 solution only applies to traffic from IPv6 clients accessing services in IPv4. Both the LTM module and the CGNAT module support NAT64 functionality.

REFERENCE ARCHITECTURE | NAT64 SOLUTION

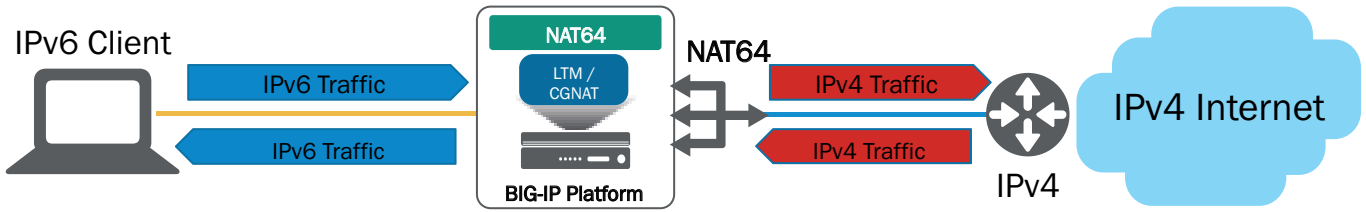

INTELLIGENT PACKET BROKERING SOLUTION

PROBLEM

Different flows of traffic in the data centre need to be sent and/or copied to platform-specific solutions (IPS/IDS, antivirus, content/URL filtering, parental control, etc.), depending on the needs of users.

Specific technical solutions do not scale in an efficient manner and costs skyrocket.

There is little flexibility when deploying and integrating new solutions within the data centre.

ALTERNATIVES

- In-line configuration of all elements of the data centre, which adds unnecessary latencies, reduces technical scalability and significantly raises costs. Also, being in-line, all elements become critical points of failure.
- Static configuration of the services (with various mechanisms) which provides neither the flexibility nor the intelligence necessary for efficient deployment.

F5 | PEM SOLUTION

PEM (Policy Enforcement Manager) is a tool developed by F5, which makes it possible to configure the functionality of «intelligent packet brokering» in an easy and scalable way, according to the origin or type of traffic, routing and/or to copy separate flows of traffic to different platforms (monitoring, antivirus, IPS, filter URLs, parental control, etc.).

It is possible to configure flows of online traffic (IPS, etc.) and/or offline traffic (IDS etc.), which means that each specific solution can be sized and scaled independently and in a way which is both technically and economically effective.

PEM provides the necessary monitoring to check the availability and performance of all of the solutions integrated in the various flows, allowing these to be dynamically changed.

REFERENCE ARCHITECTURE | INTELLIGENT PACKET BROKERING SOLUTION

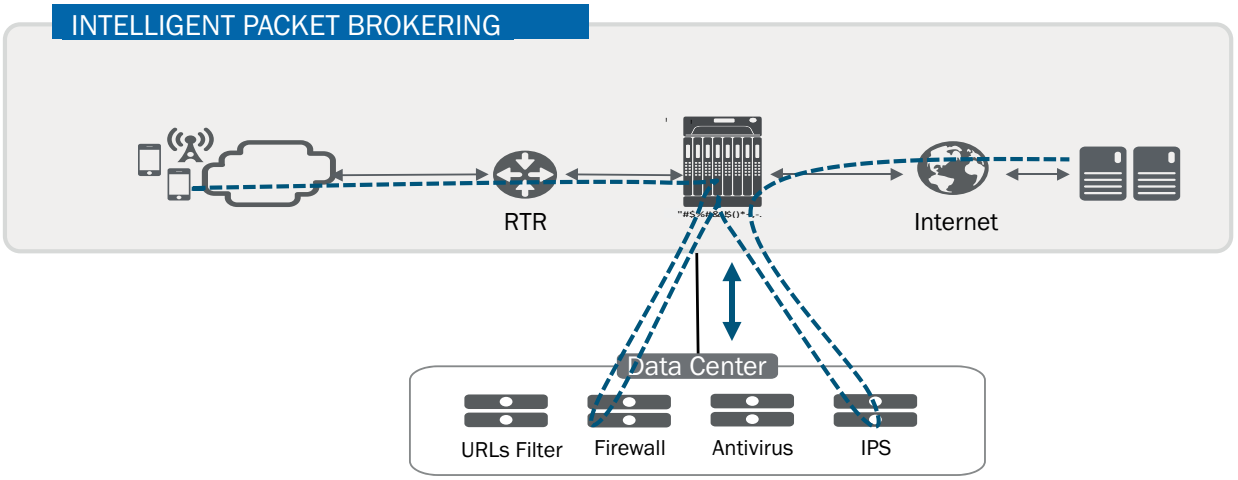

CISCO ACE REPLACEMENT

PROBLEM

In September 2010, Cisco decided not to continue developing the next generation of their ACE load-balancing products, with the result that companies were faced with difficulties with regards to the ongoing support and maintenance of the platforms.

It should also be taken into account that the Cisco ACE solution was focused on L2-L4 load balancing, is not compatible with IPv6, has limited support for functions like caching and compression and does not have templates for application deployment.

In a nutshell, customers today have no guarantee that Cisco ACE equipment can provide the functionality of a load balancer or support for platforms.

ALTERNATIVES

- Continue using the Cisco ACE solution. This creates the problem of having a team without direct support from the manufacturer, causing a serious security risk to the customer's infrastructure.

F5 | LTM SOLUTION

F5 BIG-IP Local Traffic Manager™ (LTM) is the leading solution in the ADC (Application Delivery Controller) market.

The LTM converts the network into an agile infrastructure for application delivery. It acts as a full proxy between the users and the application servers, and creates an abstraction layer to protect, optimize and load-balance application traffic. This functionality provides the flexibility and control needed to add applications and servers, eliminate downtime, improve application performance and meet your security requirements.

The advantage of acquiring the F5 solution is the ability to consolidate all the functionality onto one computer (either virtual or physical).

Thanks to the strong alliance between Cisco and F5, migrating the Cisco ACE to any F5 platform is a simple and low-risk operation.

REFERENCE ARCHITECTURE | CISCO ACE REPLACEMENT

VIPRION 4480 w B4300 Blades
10,000,000 Layer 7 Requests/Sec

Cisco ACE30 16G
160K L7 RPS x **62 Blades** = 10M L7 RPS

CLOUD
(PRIVATE AND PUBLIC)

CONTENT

29 Cost control ‘in The Cloud’ bursting
30 Federation of identities in cloud environments
31 Integration architectures SDN - Cisco ACI
32 Integration with SDN - VMware architectures
33 Integration with Cloud Access Brokers (CASBs)
34 Migration to cloud environments
35 Protection of services in cloud environments

COST CONTROL ‘IN THE CLOUD’ BURSTING

PROBLEM

Due to their business needs, some companies plan for surges or peaks in demand, which can be problematic. It is possible that the DPC will not be able to absorb these peaks in traffic caused by the increase in users attempting to access the service at the same time.

This is difficult to justify if these business needs or intermittent peaks affect the company’s well-being.

ALTERNATIVES

- Manually direct traffic into the cloud, this is slow and costly requiring human intervention and monitoring.

F5 | LTM SOLUTION

Some companies specify a certain amount of redundant infrastructure resources to meet these peaks in traffic. This has a negative impact on the return on investment (ROI). The solution for this is called Cloud Bursting. Cloud Bursting means that an application which is hosted on-site or in a private cloud can automatically move to a public cloud in critical situations where there is peak traffic, this means that the outsourcer can provide the extra capacity needed by the customer, who only pays for what they need.

The Cloud Bursting solution by F5 automates and coordinates the implementation of application services in the on-premise and cloud infrastructure, always managing the dynamic redirection of workloads to the most appropriate location. Thanks to this automation, it offers up to 80% time savings, 62% cost savings and 54% SLA compliance.

REFERENCE ARCHITECTURE | COST CONTROL ‘IN THE CLOUD’ BURSTING

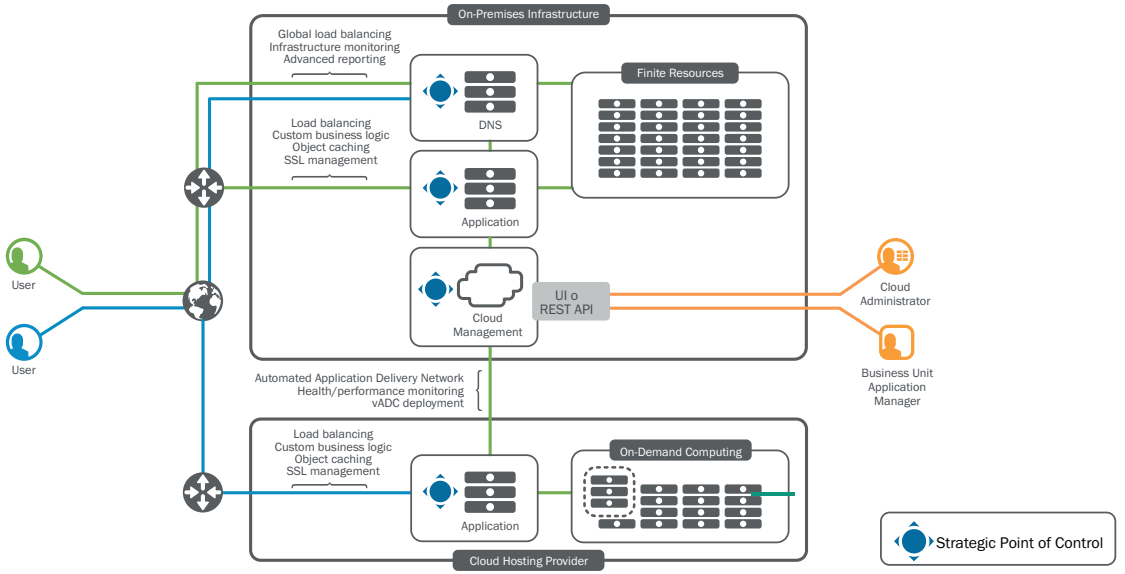

FEDERATION OF IDENTITIES IN CLOUD ENVIRONMENTS

PROBLEM

Hybrid organizations implement distributed architectures which can span multiple security domains. At any given time, a user could access the corporate data centre, the organization's cloud infrastructure and even a third party SaaS web application.

This has led to a lack of access control at the profile level as well as at the management level. This lack of control is due to the differentiation of access based on user profiles and permissions. The complexity of managing user access is due to the fact that there are multiple applications, passwords and forms of authentication.

ALTERNATIVES

- Different proprietary development solutions, which trigger higher maintenance costs and reduce the solution's flexibility.
- Specific solutions for application, which do not offer scalability and do not solve the problem of complexity of access for users.

F5 | APM SOLUTION

F5's APM (Access Policy Manager) solution provides a single point of access to applications wherever they are, simplifying, consolidating and managing secured access.

APM can offer SSO services and/or identity federation and generate reports on the type of traffic and user access.

In addition we support granularity at the level of profiles and user permissions, centralizing and applying access control policies in a single point of network infrastructure.

We support all types of devices and operating systems.

REFERENCE ARCHITECTURE | FEDERATION OF IDENTITIES IN CLOUD ENVIRONMENTS

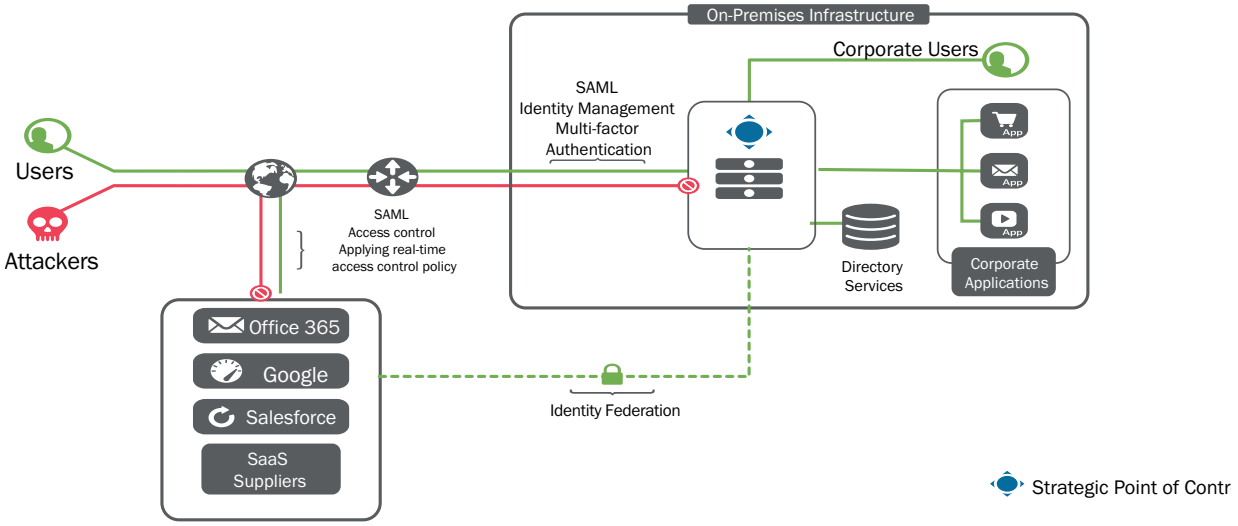

INTEGRATION WITH SDN CISCO ACI

PROBLEM

It's critical for clients that any network infrastructure on which business applications are hosted, is sufficiently flexible and sufficiently automated to minimize the time-to-market of any service. It is necessary to minimize the maximum time required to deploy applications or services.

Today, provisioning a new virtual machine can take a matter of minutes, but configuring the network associated with that virtual machine (setting up VLANs, routing, QoS policy, firewall rules, configuring load balancers) can take hours, or even days.

ALTERNATIVES

- Alternatives to VMware NSX include Cisco ACI, Nuage, and OpenStack.
- Alternatives to F5 are only able to provide very basic load balancing as a service (LBaaS) functionality.

F5 | LTM SOLUTION

The solution is to automate the provision of network services through Software-Defined Networking (SDN) technology. In SDN architectures, you can program and automate the creation of a network, orchestrate this creation and integrate it with external systems.

The integration between F5 and CISCO ACI makes it possible not only to provision the network (connectivity), but also to provision complete services and applications. This is thanks to F5's ability to provide services in an automated way using iApp technology, BIG-IQ devices and custom packages.

This is to say that F5 can not only provision a load balancer as a service (LBaaS) but also multiple capabilities, such as WAF, authentication, DNS, etc.

REFERENCE ARCHITECTURE | INTEGRATION WITH SDN CISCO ACI

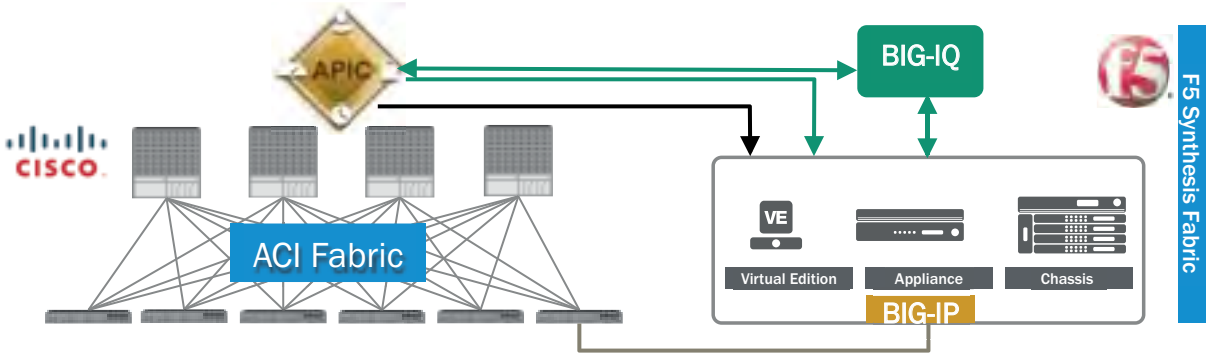

APIC to BIG-IP Integration Model

APIC to BIG-IQ Integration Model

INTEGRATION WITH SDN VMWARE ARCHITECTURES

PROBLEM

It's critical for clients that any network infrastructure on which business applications are hosted, is sufficiently flexible and sufficiently automated to minimize the time-to-market of any service. It is necessary to minimize the maximum time required to deploy applications or services.

Today, provisioning a new virtual machine can take a matter of minutes, but configuring the network associated with that virtual machine (setting up VLANs, routing, QoS policy, firewall rules, configuring load balancers) can take hours, or even days.

ALTERNATIVES

- Alternatives to VMware NSX include Cisco ACI, Nuage, and OpenStack.
- Alternatives to F5 are only able to provide very basic load balancing as a service (LBaaS) solutions. F5 can also automate the creation of many advanced services in line with the requirements of applications.

F5 | LTM SOLUTION

The solution is to automate the provision of network services through Software Defined Networks (SDN) technology. In SDN networks, you can program and automate the creation of a network, orchestrate this creation and integrate it with external systems.

By integrating F5 with VMware NSX, it is possible not only to provision the network (connectivity), but also to provision complete services and applications, thanks to F5's ability to provide services in an automated way using iApp technology, BIG-IQ and the insertion of VMware NSX services.

This is to say that F5 can not only provision a load balancer as a service but also multiple capabilities, such as WAF, authentication, DNS, etc.

In addition, NSX can be used to create instances of F5 VE virtual machines and these will be automatically configured and licensed.

REFERENCE ARCHITECTURE | INTEGRATION WITH SDN VMWARE ARCHITECTURES

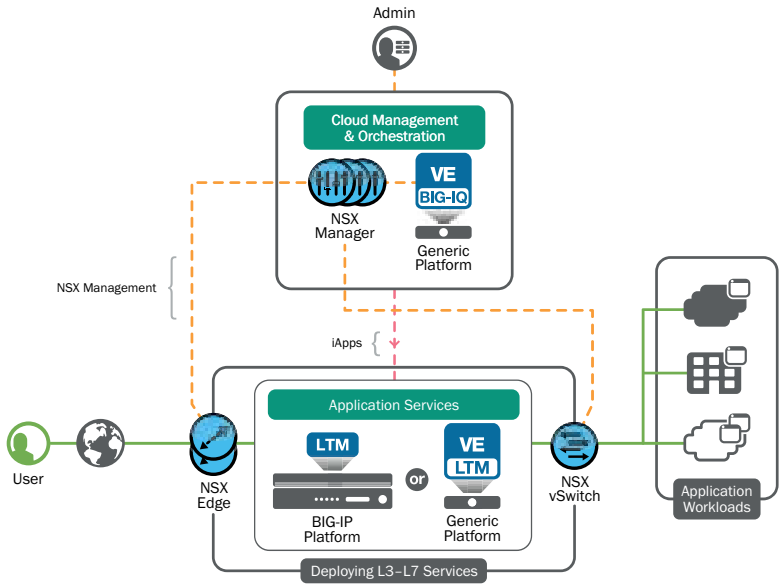

INTEGRATION WITH CLOUD ACCESS BROKERS (CASBs)

PROBLEM

There are many companies that decide to move corporate applications «into the cloud», and this generates new security problems. The concepts of security and perimeters change, applications and data are no longer protected in our data centre and, most importantly, we lose visibility into what users do. Therefore the likelihood of establishing a common usage policy is lower and we need to find new solutions

If any user can access data anytime, anywhere and through any device, there needs to be a new understanding of security which allows us to protect data and to ensure the lawful behaviour of users.

ALTERNATIVES

- The alternative is to use several solutions and integrate them, for example proxy chaining and NAC, this increases complexity and costs

F5 | APM + SWG SOLUTION

With the F5 solution we can solve many of the problems we face in migrating applications to the cloud.

Application visibility in SSL traffic: We can analyse and control traffic by integrating as a proxy and encryping/ decrypting SSL/TLS traffic as required, making it possible to establish efficient policies and to prioritize the use of business applications.

Authentication, Authorization and Auditing of users: Native support for popular methods to authentication and control of access to applications according to user, group, location, device, etc.

Data security through robust encryption designed to protect data and transactions in cloud applications.

REFERENCE ARCHITECTURE | INTEGRATION WITH CLOUD ACCESS BROKERS (CASBs)

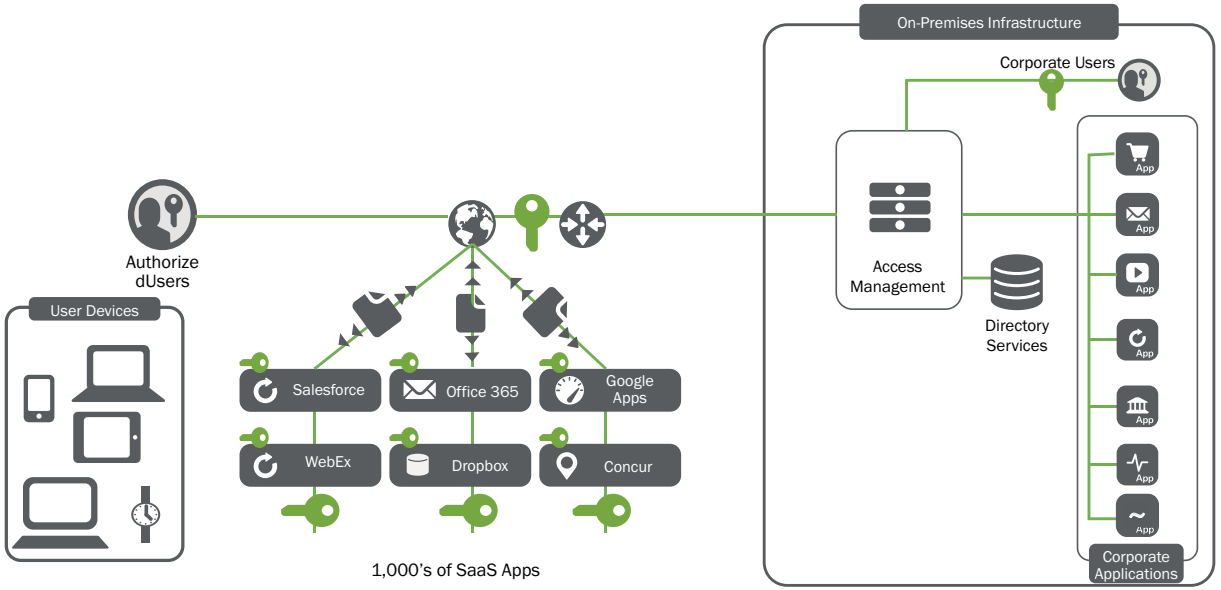

MIGRATION TO CLOUD ENVIRONMENTS

PROBLEM

Companies contract cloud services for reasons of excess traffic or to reduce operating costs. The problem arises when an application is transitioned from a local data centre to a cloud environment or when we want on-premises and cloud instances of an application to co-exist, we call this a hybrid cloud.

Companies face migration challenges and user experience may be impaired.

ALTERNATIVES

The alternative is to manually handle the migration, stopping the local data centre services and starting cloud instance at appropriate times, while in parallel managing static DNS entries. This practice is not only cumbersome and complicated, but may also incur service outages. When services are critical to the company, this practice is not advisable.

F5 | LTM + DNS SOLUTION

Thanks to the GSLB functionality which is included in the BIG-IP DNS module by F5, it is possible to define both data centres and the services in each one of them.

BIG-IP DNS is able to advertise services in the preferred data centre for each client and can prioritize data centres depending on whether services may be migrating to a different location (possibly to the cloud).

Rules are used to analyse requests from users and to steer sessions to the optimal data centre or to one which contains an available instance of the required service.

By integrating F5 in SDN environments and through the use of APIs such as iControl REST, we can automatically provision services on demand in cloud environments, undertaking a seamless migration without stopping services for users.

REFERENCE ARCHITECTURE | MIGRATION TO CLOUD ENVIRONMENTS

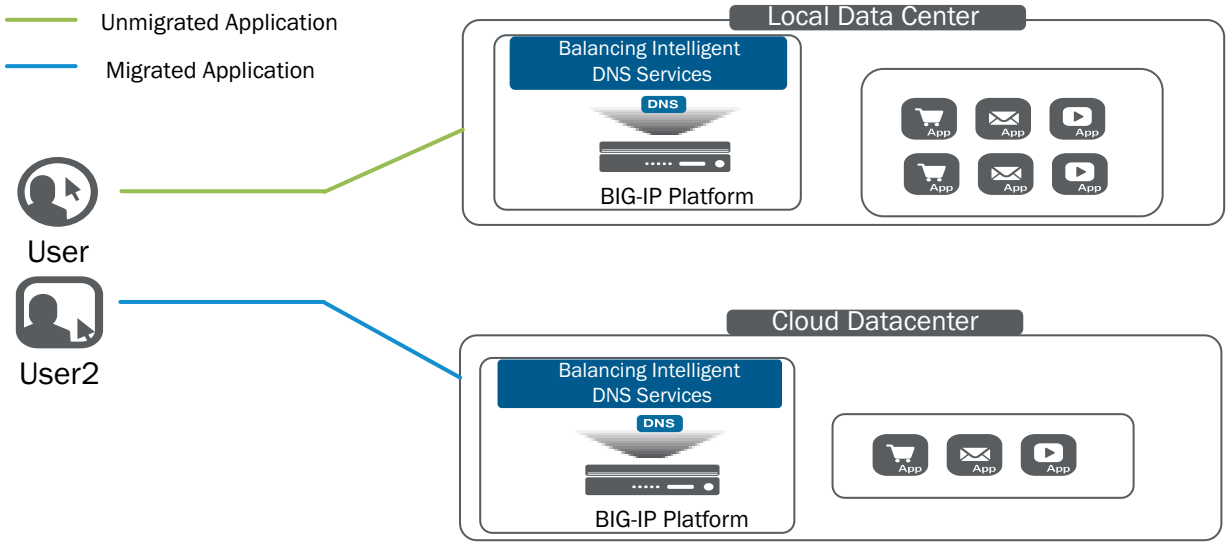

PROTECTION OF SERVICES IN CLOUD ENVIRONMENTS

PROBLEM

Many companies in the media environment have decided to use cloud solutions for publishing content. This requires a pay-per-use solution that ensures 24/7 service availability, flexibility of storage, computing capacity and link size, according to business needs. It is also quite common to use more than one cloud-service provider so that there is more flexibility if circumstances change. While this scenario greatly facilitates the business environment and improves simplicity, it generates new requirements which bring about a need for advanced solutions. One example of this is application security, how can we secure application data or prevent DDoS attacks?

How can we apply common security policies to all public content and applications?

ALTERNATIVES

- Contracting security solutions and services or cloud service providers.
- Make use of CDN-type solutions which allow us to secure web traffic at an earlier point.
- Both solutions are much more expensive and more complex to manage.
- The solutions offer little flexibility or versatility.

F5 | SILVERLINE SOLUTION

With F5's cloud security solutions (Silverline), we can apply policies to web services from a common point and through a fully-managed and flexible service, for which you pay only for what you use.

We eliminate the problems associated with the use of more than one cloud service provider, integrating our solution to give proxy access to any published service.

We eliminate the management complexity involved in any type of WAF solution by delegating management to expert staff at F5, without losing visibility and information on the usage of applications.

REFERENCE ARCHITECTURE | PROTECTION OF SERVICES IN CLOUD ENVIRONMENTS

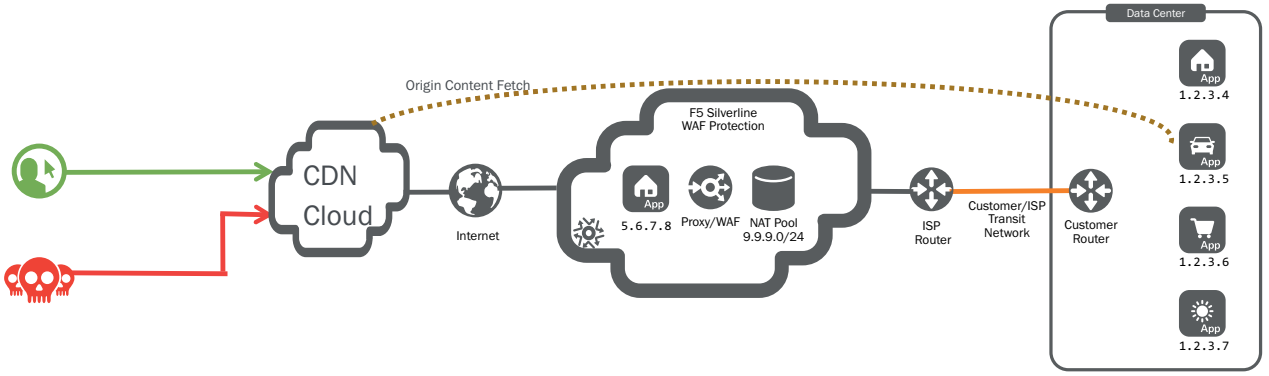

SECURITY

CONTENT

36 Advanced user access control
37 Regulatory compliance PCI-DSS
38 DDoS attack detection multi-homing scenarios
39 L7 DDoS attack detection for activating cloud protection
40 Phishing attack detection
41 Detection and mitigation of L7 DDoS attacks for encrypted traffic
42 Integration with anti APTs – FireEye
43 Integration with MdM – Airwatch
44 Integration with MdM – MobileIron
45 Integration with DAST and WAF solutions
46 HSM integration solutions
47 Optimization of Google ranking using TLS/SSL
48 Reverse proxy gateway
49 Anti-Bot protection in OWA environments
50 Protection with smart reputation lists
51 DNS infrastructure protection
52 Protection of Public-Service portals
53 DNS protection through DNSSEC
54 Transparent fraud protection for e-commerce transactions
55 Juniper SSL VPN replacement
56 Microsoft Forefront TMG replacement
57 Replacement of DNS platforms based on bind
58 Always-on solution for mobile environments
59 Kerberos authentication solution for mobile devices
60 Remote VPN SSL access solution
61 SSL/TLS offloading solution
62 SSL/TLS visibility solution
63 Single Sign-On solution
64 Web Application Firewall (WAF) solution
65 Proxy solution for safe navigation for employees
66 Enhanced Web Application Firewall (e-WAF) solution

ADVANCED USER ACCESS CONTROL

PROBLEM

Services require user identification to prevent unauthorized access.

The increasing frequency of attacks means that more intelligence is required to allow user access to a particular resource, a password is not enough, the user's context also needs to be considered, where, how, when, etc.

ALTERNATIVES

- One alternative to this technology is to implement this logic in the application itself, which requires personalizing it, reprogramming, code, testing, etc. This alternative, besides being cumbersome, is not replicable in other applications.
- Another alternative may be the use of proxies from specific manufacturers but these lack scalability (and HA), they are not compatible between different software manufacturers and also lack the same capabilities as BIG-IP APM. (This requires one pair of proxies per manufacturer with different policies etc., (IBM, SAP, Oracle, MS, etc.) and intermediate elements with little ability to manage encryption traffic, points of failure, OPEX++, etc.

F5 | APM SOLUTION

F5's Access Policy Manager makes it possible to establish granular access policies for different services, which leads to personalized access control for each application and for each group of individuals.

Access to the application takes into account the context of the user, who they are, from where they are requesting access and to which group they belong.

This solution validates the user against multiple types of repositories (LDAP, RADIUS, AD, etc.), checks the security status of the device, undertakes authentication using advanced mechanisms such as SAML, Kerberos, NTLM, etc. and supports multi-factor authentication to applications according to the user's context.

REFERENCE ARCHITECTURE | ADVANCED USER ACCESS CONTROL

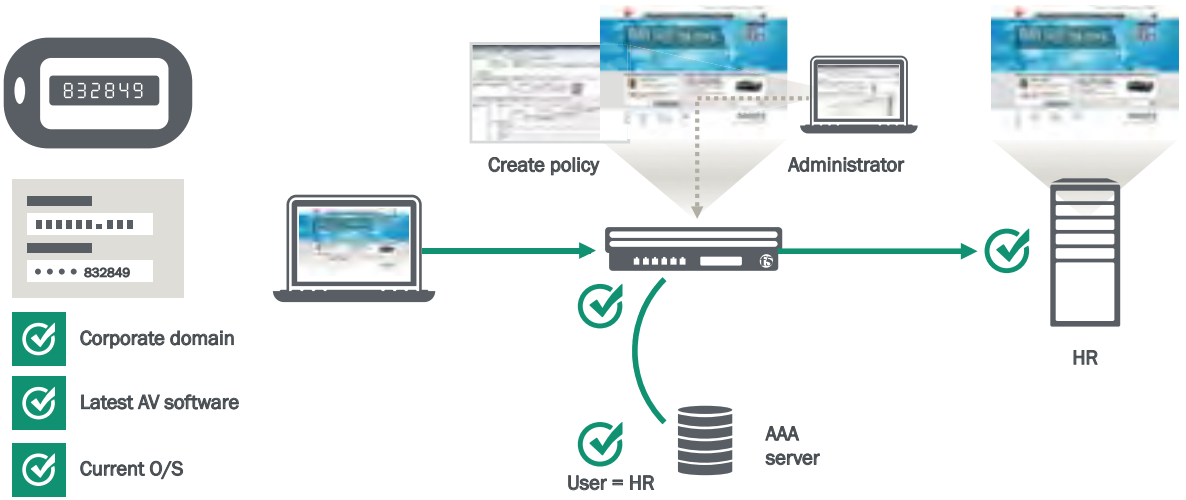

Perform check of user access based on multiple parameters: compliance, AD group, device, geolocation, etc.
Add Double factor authentication based on context
Integration with MDM, etc.

REGULATORY COMPLIANCE PCI-DSS

PROBLEM

A consortium of the leading companies in the Payment Card Industry (PCI) created a set of Data Security Standards (DSS) for organizations that use payment gateways on their websites and process or store such data. These companies must meet a minimum level of security to help prevent fraud. Companies which do not comply with this rule may be subject to fines and even lose the use of the payment gateway.

Of the twelve requirements imposed, one of the most critical points is in paragraph 6, which deals with the security of the application. Initially, the standard required the development of secure code and that this be audited periodically, etc. Given the difficulty of complying with this point, when the standard was first reviewed it was amended to incorporate the option of using a dedicated security solution (Web Application Firewall). A WAF, in addition to being optimal, offers more security.

ALTERNATIVES

- Applying all the points of the PCI-DSS standards during the process of developing the application is much more expensive and insecure. A dedicated device which is kept regularly updated is unlikely to contain exploitable security vulnerabilities, unlike code which is audited every six months.
- Competing solutions help «enforce» compliance with the regulations, but keep in mind that F5's solution is the most widespread in the market and, due to its strategic location in the client architecture (within the ADC), it is technically the most appropriate.

F5 | ASM SOLUTION

F5 ASM (Application Security Manager) makes it possible to deploy a WAF (Web Application Firewall) module in a fast and simple way and can thus comply with PCI-DSS standards. ASM is deployed in a way which is transparent to applications, thus protecting against web application attacks such as OWASP and DDoS.

ASM also incorporates a detailed report on which aspects of PCI-DSS the application passes/fails.

REFERENCE ARCHITECTURE | REGULATORY COMPLIANCE PCI-DSS

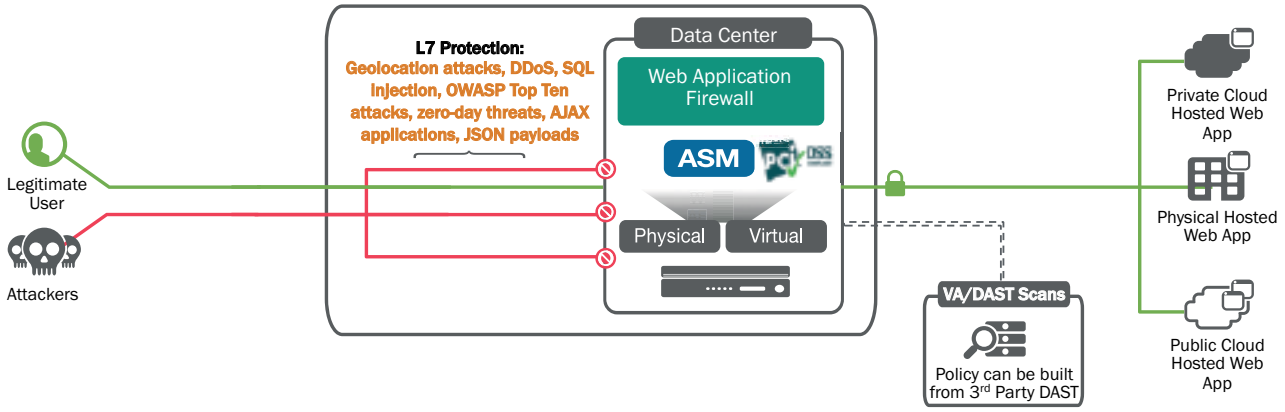

DDoS ATTACK DETECTION MULTI-HOMING SCENARIOS

PROBLEM

Customers with more than one outbound internet service provider (ISP) face the dilemma of deciding whom they should contract for protection against DDoS attacks.

Contracting a protection service from all providers is unlikely to be an economically-viable solution. On the other hand, contracting the service from a single supplier requires inspection on-premises of the aggregate traffic from all suppliers.

ALTERNATIVES

- Contracting a DDoS protection service with all suppliers is not economically profitable, nor does it improve the quality of service.
- Placing the CPE extension for the detection of attacks at each ISP greatly increases the costs and complexity of the solution.

F5 | AFM + ASM SOLUTION

F5's AFM (Advanced Firewall Manager) and ASM (Application Security Manager) modules make it possible to see DDoS attacks from L3 to L7 across the aggregate of client internet lines, so we can deploy a strategy of contracting services for protection against volumetric DDoS attacks with a single ISP.

As soon as it is identified that an attack is taking place through any of the ISPs, the F5 platform mitigates the attack on layers L3 to L7, restricting it to the line capacity of each of the suppliers, minimizing the use of defensive restraints.

In the case of an actual attack, the F5 platform may use «Cloud Signalling» to notify the ISP with which it has contracted for DDoS protection, leaving it to announce the prefixes for the services to other ISPs, ensuring service availability.

This solution is also relevant to the scenario of attacks on encrypted traffic.

REFERENCE ARCHITECTURE | DDoS ATTACK DETECTION MULTI-HOMING SCENARIOS

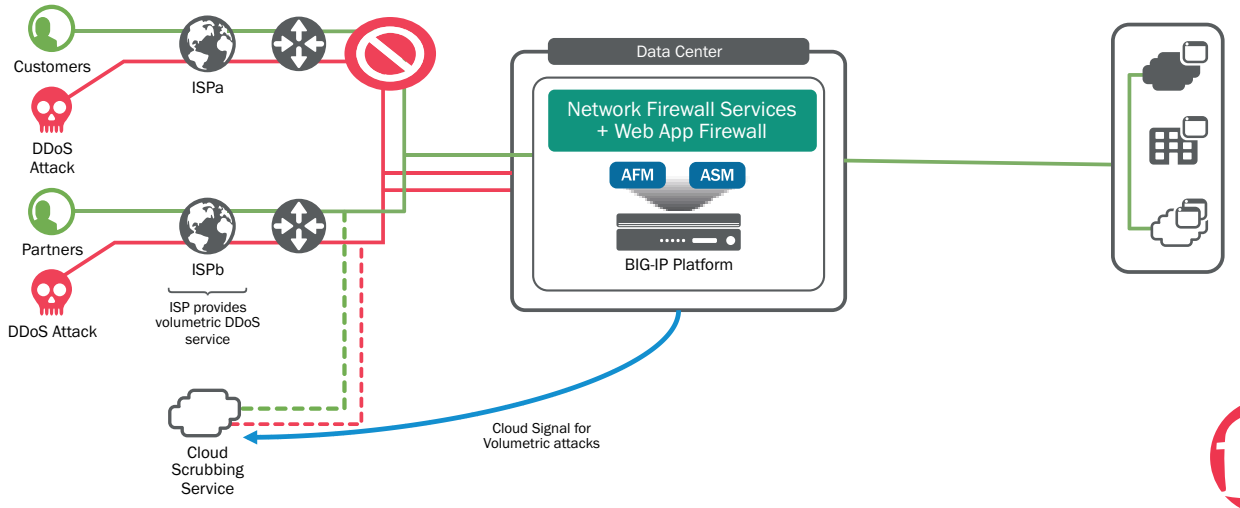

L7 DDoS ATTACK DETECTION FOR ACTIVATING CLOUD PROTECTION

PROBLEM

- Detection of DDoS attacks in the Service Provider's network is performed using the NetFlow protocol, L3 (Network) and L4 (Transport).
- Detection by NetFlow is ineffective against «Low & Slow» attacks in L7 (Application), although it is an effective way to mitigate them once the traffic has gone online and passed through the mitigation centres (scrubbing-centres) following the BGP/DNS announcement to redirect traffic.

ALTERNATIVES

- Network operators can provide equipment at the customer's site to detect L7 DDoS attacks, but this CPE is dedicated to this purpose and greatly increases the cost of the service.
- The alternative is always to pass the traffic online via the service operator; this is much more expensive and greatly increases the latency of all traffic, even when there is no attack of any kind.

F5 | ASM SOLUTION

F5 ASM (Application Security Manager) solution is deployed at the entrance to the data centre; it identifies all types of DDoS attacks which are specific to L7 and protects applications. Through its «Cloud Signalling» functionality, the ASM can signal the existence of an attack on the SP service, thereby activating in-cloud protection and mitigating the attack before it reaches the data centre.

ASM is able to mitigate these DDoS attacks specific to L7, restricting them to the available bandwidth in the data centre, saving the customer the cost of activating the cloud-protection service during those attacks which do not exceed the available bandwidth.

REFERENCE ARCHITECTURE | L7 DDoS ATTACK DETECTION FOR ACTIVATING CLOUD PROTECTION

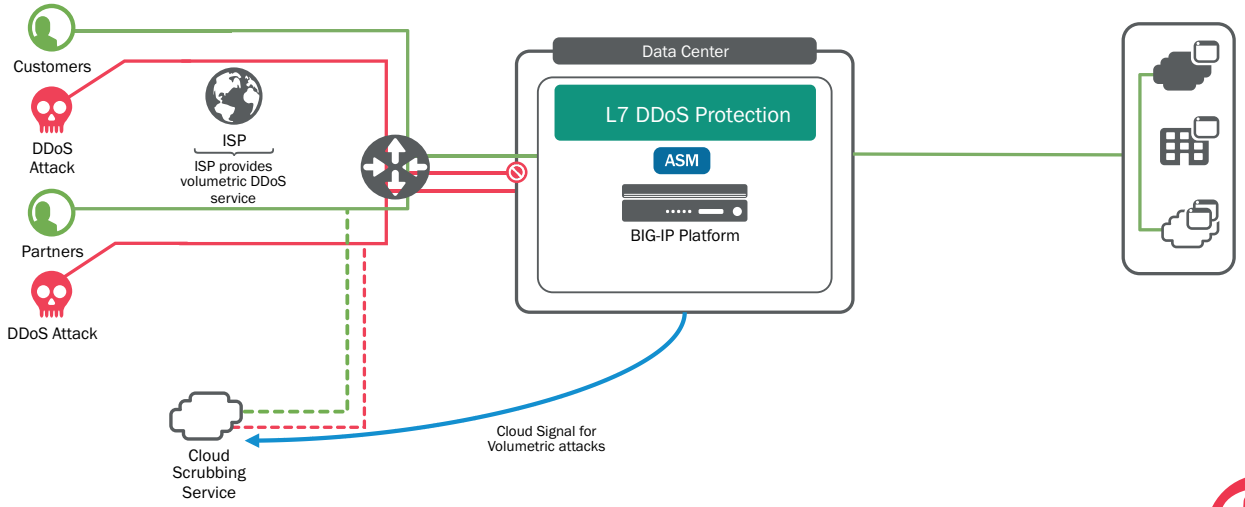

PHISHING ATTACK DETECTION

PROBLEM

Every day there is an increase in the level of economic fraud in e-commerce transactions. One of the most common methods of attack in this area is the phishing attack. There are many variations of this technique, but basically it consists of copying a company's website and redirecting potential victims to this copy in order to steal their credentials and thereby be able to spoof them. Broadly speaking, phishing has 3 phases: copying the original website; publishing the fraudulent copy and stealing the victim's credentials. It is also important to be able to close the fraudulent website as soon as it has been detected to avoid there being new victims.

ALTERNATIVES

- Although there are other products and services, which are able to detect phishing, F5's technology is the only one to be completely transparent in its implementation (there is no need to modify the application, or install anything on the client) and also the only one which works on all 3 layers as described above.

F5 | WEBSAFE SOLUTION

F5's anti-phishing technology is unique in the way it works and in how it is implemented. It transparently injects an obfuscated JavaScript code which protects itself against possible removal. This code is capable of verifying whether the application is the original or has been issued fraudulently, and if it detects phishing, it sends an alert to the company which owns the application or to F5. This technology is the only one capable of undertaking detection at all 3 layers described above: if the application is copied; if it has been published on a fraudulent site and, most interestingly, if a victim has entered this "phished" site.

In addition, the cyber-security specialists at F5 are able to shut down these fraudulent websites within a few minutes to avoid online fraud and to protect the image of the corporation.

REFERENCE ARCHITECTURE | PHISHING ATTACK DETECTION

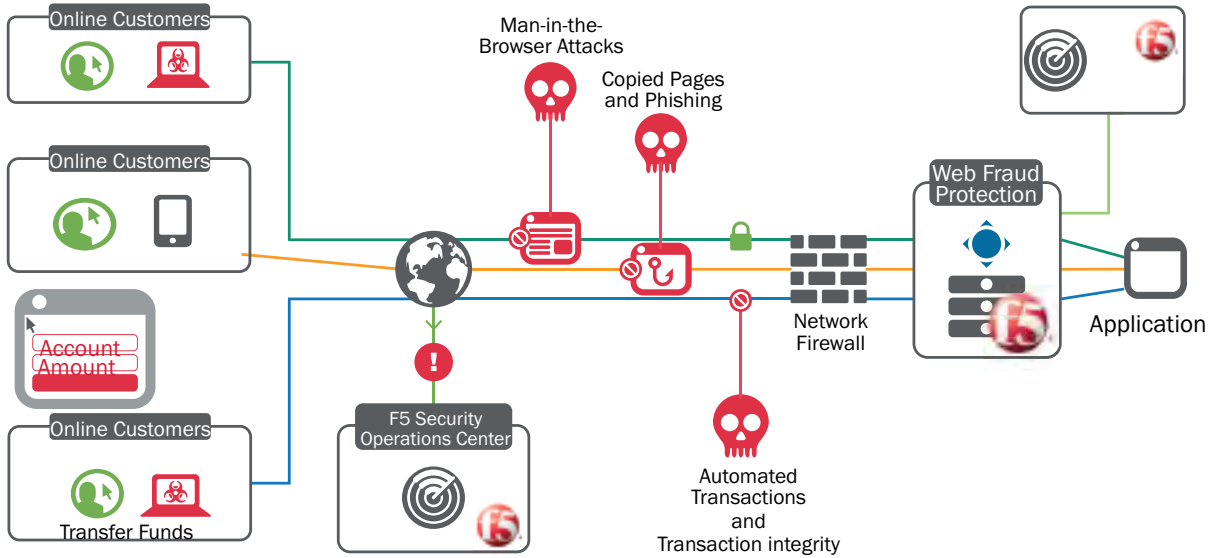

DETECTION AND MITIGATION OF L7 DDoS ATTACKS FOR ENCRYPTED TRAFFIC

PROBLEM

The «Low & Slow» attacks in L7 on encrypted traffic (HTTP over SSL) are not detectable by the protective services offered by the SP or in the cloud, unless we provide private encryption keys to the service provider.

ALTERNATIVES

L7 detection and protection services on encrypted traffic offered by SPs or cloud-based solutions require private encryption keys. Providing these keys to a third party may violate security policies or the privacy of the service.

F5 | ASM SOLUTION

The ASM (Application Security Manager) solution makes it possible to decrypt HTTPS traffic locally (importing private encryption keys) and analyse the resulting HTTP traffic, thereby identifying all types of DDoS attacks specific to L7 and protecting applications.

Once the traffic has been cleaned, it is possible to re-encrypt traffic to application servers, even with a different key length to the original, complying with the service's security standards.

REFERENCE ARCHITECTURE | DETECTION AND MITIGATION OF L7 DDoS ATTACKS FOR ENCRYPTED TRAFFIC

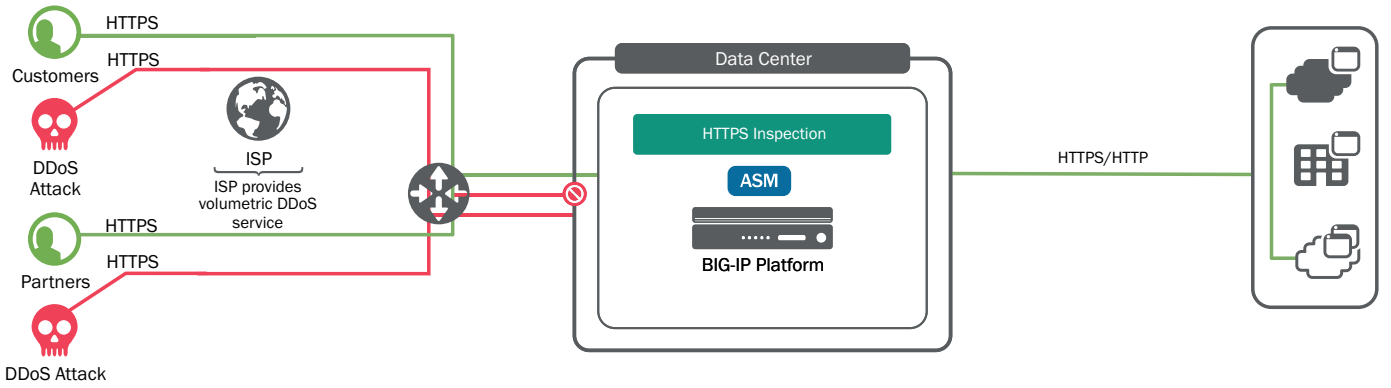

INTEGRATION WITH ANTI APT SOLUTIONS – FIREEYE

PROBLEM

Some classes of attack have become so sophisticated that they are no longer known simply as attacks but as APTs (Advanced Persistent Threats), or targeted and persistent attacks, with an intelligence able to outwit firewalls and traditional devices. Countering these requires specific solutions which take responsibility for investigating these new attacks, simulating the entire flow of the attack and seeing the damage they can cause, this is the role played by FireEye.

SSL/TLS-encrypted traffic is a problem for these solutions, since they are not able to decipher it to examine the content, which means that they cannot raise an alert if there is an infection or attack.

Moreover, FireEye devices need to be scaled and load balanced by a device that is capable of interpreting the traffic and sending it to a device at layer 2.

ALTERNATIVES

- Leave the FireEye equipment without SSL visibility
- Scale the FireEye by putting one in each of the VLANS and using physical interfaces for each one.

F5 | LTM + APM + SWG SOLUTION

With the advanced SSL/TLS functionality of BIG-IP LTM, we can load balance FireEye devices and bypass their limitations around encrypted traffic. This makes it possible to install a pool of FireEye devices to serve an entire enterprise.

With the SSL functionality provided by BIG-IP, we can decrypt traffic and send it to FireEye for analysis, and then encrypt it once again before it exits to the internet.

If we were to pass all traffic through FireEye this could either cause it to be overloaded, or to think that it was being flooded in an «illegal» process on some domains. We therefore use SWG (Secure Web Gateway) to classify the traffic and to distinguish between different categories, bypassing domains we do not want to investigate.

APM (Access Policy Manager) functionality is used to authenticate users and to determine whether or not it is appropriate to send them to be investigated by FireEye.

REFERENCE ARCHITECTURE | INTEGRATION WITH ANTI APT SOLUTIONS – FireEye

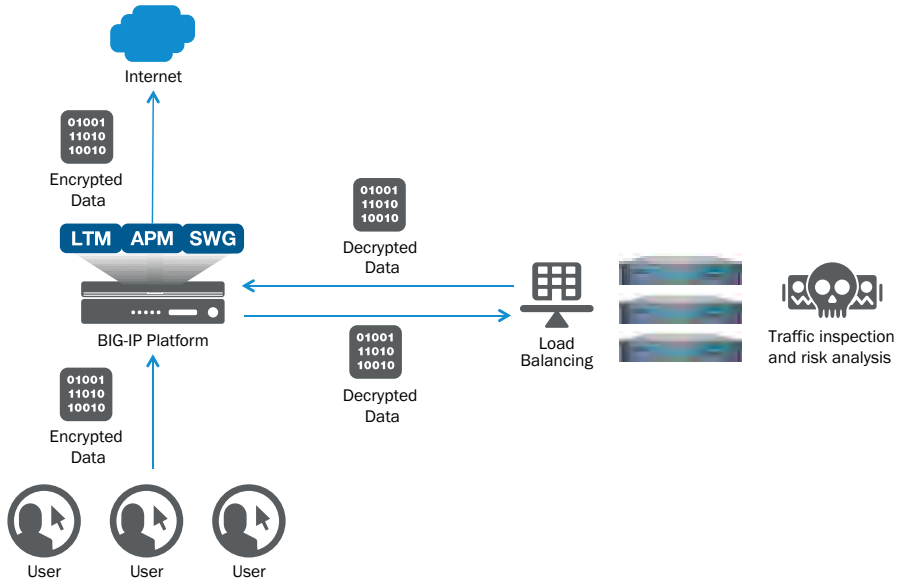

INTEGRATION WITH MDM AIRWATCH

PROBLEM

In recent years, access to corporate resources from mobile devices has become an increasingly common practice. Some companies even allow employees to use their own devices (BYOD), but this involves a security problem both in terms of the confidentiality of the data and in terms of the protection of users' mobile devices.

MDM solutions provide assurance that mobile devices comply with a set of security policies when connected to corporate resources and at the same time they provide mechanisms for the company to manage these devices to ensure the confidentiality of data and to erase confidential data in case of loss or theft. MDM solutions provide security for mobile devices, but they do not address the needs of network access from such devices.

ALTERNATIVES

- Let MDM solutions perform their function independently, doubling the security management of access from mobile devices.
- Double authentication in two different solutions, with consequent «inconvenience» to the user.

F5 | APM SOLUTION

The F5 APM (Access Policy Manager) module makes it possible to converge all authentication mechanisms onto a single point (the BIG-IP) and this device then ensures compliance with security policies across the network.

APM allows for integration with MDM solutions (such as AirWatch) so that once the device has been identified, APM interrogates the MDM to see if the device complies with the security policies required for access.

The solution makes it possible to use SSO for enterprise applications.

If the device does not meet the established safety features, we can deny access or mandate "re-enrolment" with the MDM.

REFERENCE ARCHITECTURE | INTEGRATION WITH MDM - AIRWATCH

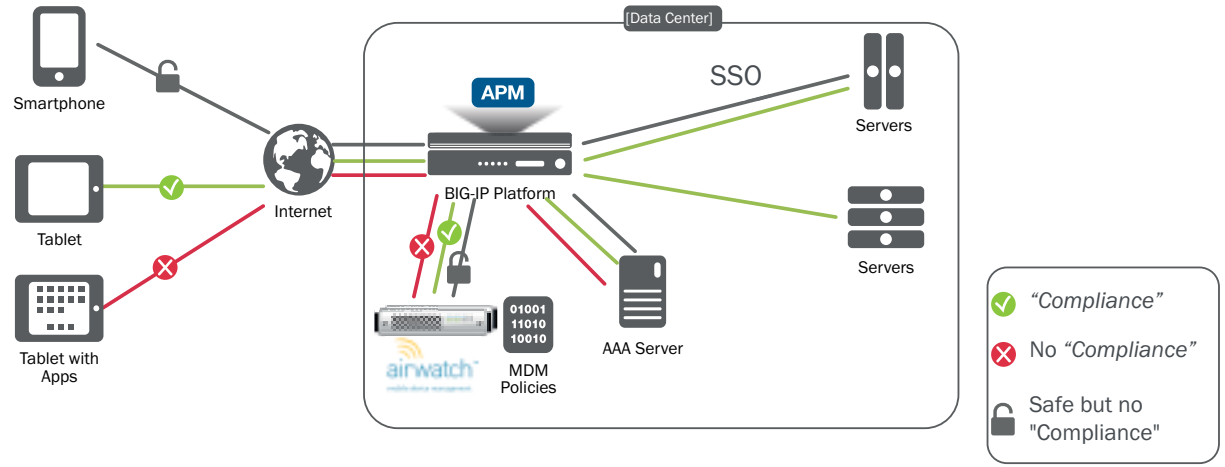

INTEGRATION WITH MDM MOBILEIRON

PROBLEM

In recent years, access to corporate resources from mobile devices has become an increasingly common practice. Some companies even allow employees to use their own devices (BYOD), but this involves a security problem both in terms of the confidentiality of the data and in terms of the protection of users' mobile devices.

MDM solutions provide assurance that mobile devices comply with a set of security policies when connected to corporate resources and at the same time they provide mechanisms for the company to manage these devices to ensure the confidentiality of data and to erase confidential data in case of loss or theft. MDM solutions provide security for mobile devices, but they do not address the needs of network access from such devices.

ALTERNATIVES

- Let MDM solutions perform their function independently, doubling the security management of access from mobile devices.
- Double authentication in two different solutions, with consequent «inconvenience» to the user.

F5 | APM SOLUTION

The F5 APM (Access Policy Manager) module makes it possible to converge all authentication mechanisms onto a single point (the BIG-IP) and this device then ensures compliance with security policies access the network.

APM allows for integration with MDM solutions (such as MobileIron) so that once the device has been identified, APM interrogates the MDM to see if the device complies with the security policies required for access.

The solution makes it possible to use SSO for enterprise applications.

If the device does not meet the established safety features, we can deny access or mandate "re-enrolment" with the MDM.

REFERENCE ARCHITECTURE | INTEGRATION WITH MDM - MOBILEIRON

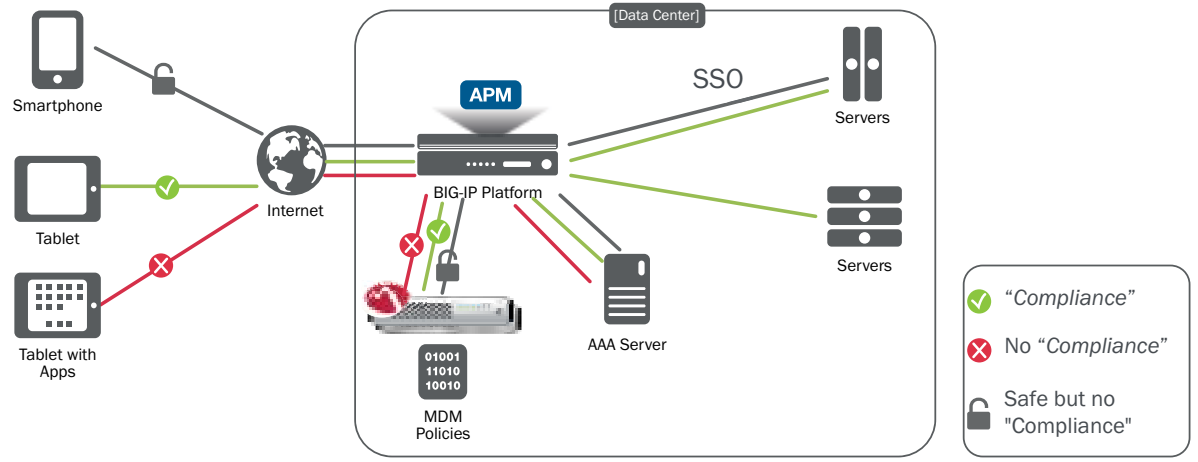

INTEGRATION WITH DAST AND WAF SOLUTIONS

PROBLEM

WAF technology can be deployed in several ways in a data centre, it can range from a fully-automated service based on blacklists, which are themselves based on signatures, or to the most sophisticated option of granting access using whitelists based on rules which are highly customized for each application. The first option does not provide a sufficient level of security for many customers seeking protection for more than OWASP attacks and the second option requires a higher operating cost, which is more than many customers can afford.

Similarly, when vulnerabilities are discovered, the time-frames for implementing new rules in an organization's security equipment are greatly extended by the organization's internal processes, which often greatly hinders the reaction time when faced with such events.

There is an intermediate possibility for a low cost solution, which improves the degree of security and provides signature-based protection through the automated integration of Dynamic Application Security Test (DAST) scanning tools (such as Rapid7, Qualys, Faast, Tenable, WhiteHat, etc.) with automatic and dynamic import of new rules in the WAF.

ALTERNATIVES

- Trusting only in databases of attack signatures leaves us more vulnerable to sophisticated attacks which require more dedicated supervision.
- Using a human-supervision model requires dedicated resources which can significantly raise operating costs.
- Not automating the process of generating new signatures will massively extend the time it takes to remedy incidents.

F5 | ASM SOLUTION

The F5 ASM (Application Security Manager) module makes it possible to deploy a WAF quickly and easily, in a way which is transparent to applications, and thereby to protect them against web application attacks on L7, such as OWASP and DDoS. It also includes a feature for monitoring compliance with PCI-DSS regulations.

ASM supports importing vulnerability reports from third-party DAST-type tools to supplement protection policies.

F5's scheduling capabilities make it possible to automate the process of parsing the report from the DAST tool, generating new recommended signatures and implementing any new security controls in production.

REFERENCE ARCHITECTURE | INTEGRATION WITH DAST AND WAF SOLUTIONS

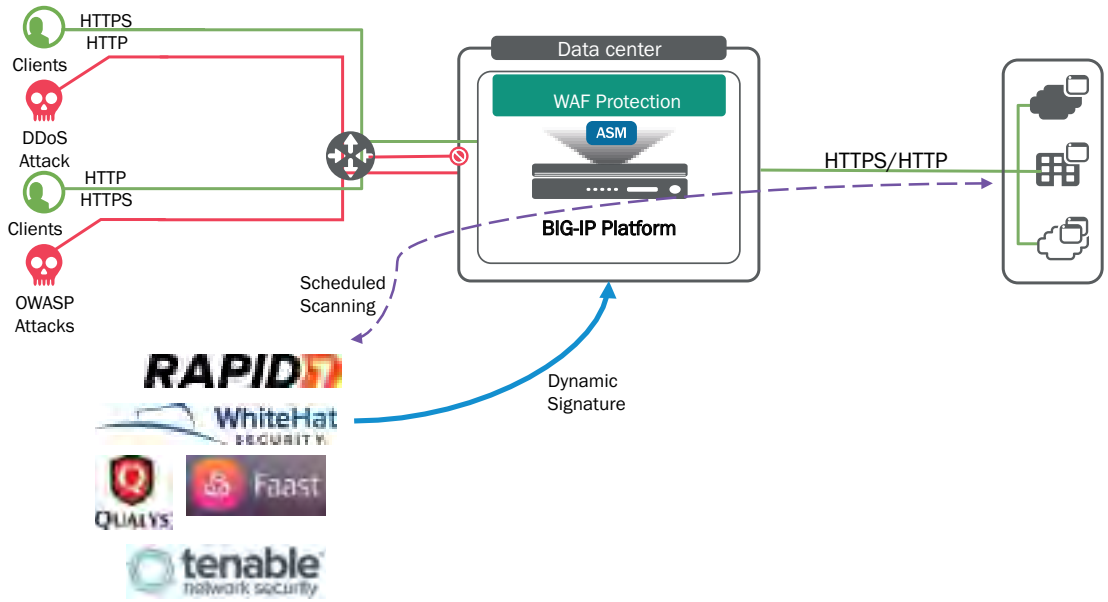

HSM INTEGRATION SOLUTIONS

PROBLEM

In scenarios where security is critical, many companies use external certificate repositories (HSM) from manufacturers such as Thales or SafeNet. These solutions guard the private key of the certificates, preventing them leaving their custody. This helps to avoid these private keys being compromised. In these scenarios, it is necessary for the teams involved in the service(s) using these certificates to be able to integrate with these repositories, to the extent that they can encrypt/decrypt the traffic, albeit without the private key.

In addition to enhancing security, the HSMs simplify the management of certificates and are a way to achieve FIPS 140-2 security certification, which is required for certain critical environments.

ALTERNATIVES

- F5 is the only manufacturer of an ADC, which has the functionality to integrate with external HSMs and to coexist with them in high-security environments where the certificates are in a single external repository, and thus fulfil FIPS standards. Even a VE could work in a FIPS environment.
- The alternative is to use an ADC with internal HSM equipment (FIPS), apart from the expensive price, this solution no longer follows the standard practices of companies which use HSMs and can therefore no longer be the sole repository of corporate certificates.

F5 | LTM SOLUTION

All BIG-IP equipment, including VEs, are able to integrate with HSM equipment, achieving FIPS 140-2.

There are certain appliances, such as the BIG-IP 7000 and 10000 series, that already have FIPS certification because they include an internal HSM for the safe custody of the keys.

REFERENCE ARCHITECTURE | HSM INTEGRATION SOLUTIONS

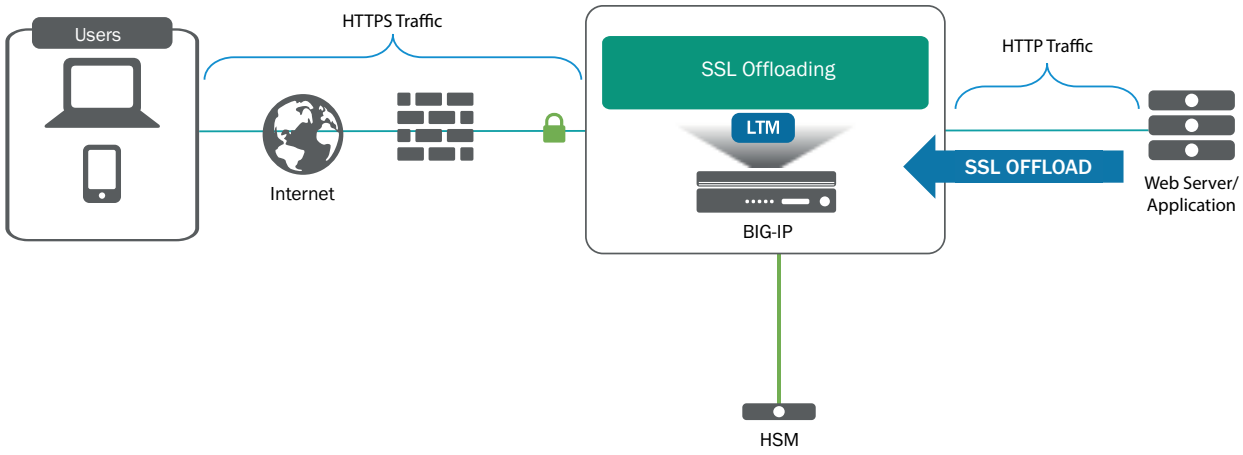

OPTIMIZATION OF GOOGLE RANKING USING TLS/SSL

PROBLEM

SSL is an encryption protocol which secures internet connections using public key cryptography to encrypt the data transfer between a sender and a recipient and thus ensure confidentiality.

One of its most common uses is to encrypt HTTP to form HTTPS; this is used to encrypt more than 70% of web transactions over the internet.

Older versions of SSL are highly vulnerable. The subsequent evolution of SSL (called TLS) is significantly safer, particularly in its latest versions and one of the criteria used by Google for ranking web pages is the result of their «SSL assessment *». Fifty per cent of companies get a grade of «C», which is very poor. It is far from easy to implement TLS in old web infrastructure. * Https://www.ssllabs.com/ssltest/index.html.

ALTERNATIVES

- Renew the entire infrastructure of the web service and dimension it to support the latest encryption standards; this entails considerable investment in terms of cost and implementation time.

F5 | LTM SOLUTION

F5's BIG-IP LTM solution is a full-proxy architecture that incorporates hardware acceleration and supports more encryption protocols than any other platform on the market.

It is the perfect solution for positioning between users and applications, communicating with the user via the latest standards of encryption. To ensure the security of connections and to maintain the server-facing encryption without modification, F5 LTM becomes a SSL/TLS gateway.

BIG-IP LTM considerably improves the security of the service and as a beneficial side effect it also improves internet positioning (SEO) without incurring large investments in a way which is both almost immediate and transparent to today's applications.

REFERENCE ARCHITECTURE | OPTIMIZATION OF GOOGLE RANKING USING TLS/SSL

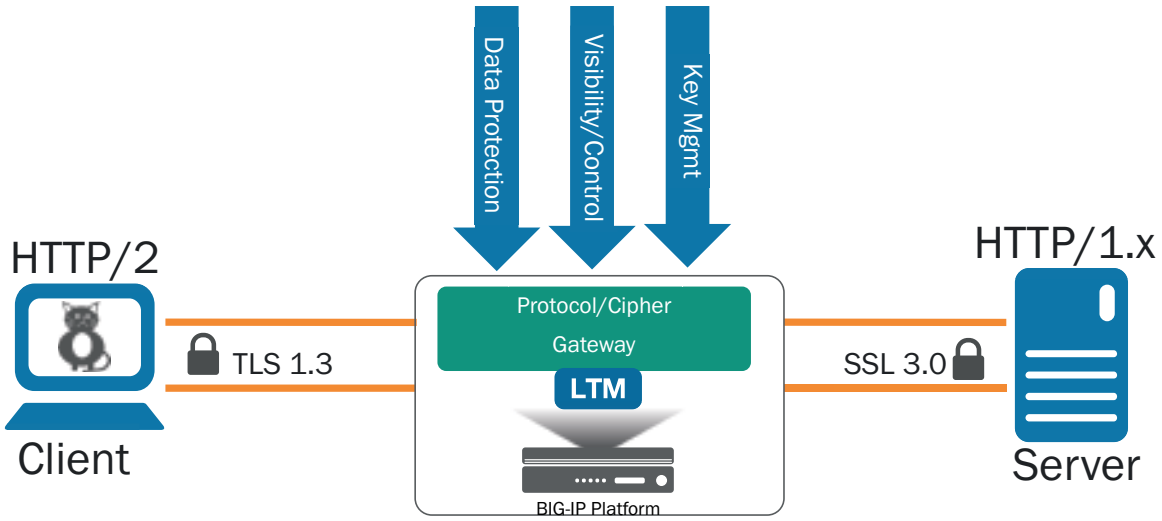

REVERSE PROXY GATEWAY

PROBLEM

It is quite common to use gateways or reverse proxies to provide external users with access to certain types of applications for internal use. Common examples of this include the use of MSFT TMG or ISA Server to provide access to Exchange/SharePoint or Apache servers to rewrite URLs, but also many companies use such solutions to improve the security of the access; HTTPS traffic is decrypted in these devices and forwarded «in the clear» to servers or authentication services, especially with the use of digital certificates.

Based on these previous examples, it is easy to conclude that we are talking about specific solutions according to application type, leading to environments which are very heterogeneous and components which nobody takes forward, develops or manages regularly. More points of failure and more complexity.

ALTERNATIVES

- Partially handled by a reverse-proxy.

F5 | APM SOLUTION

F5's APM (Access Policy Manager) module makes it possible to resolve all configuration scenarios where the use of solutions such an access gateway or reverse proxy is required.

User Authentication. Native support for use of both the main and the most varied authentication methods.

Publishing services and rewriting URLs. Publishing services externally and forwarding to internal services and servers, hiding the real internal addresses and URLs.

Encryption/Decryption of SSL traffic. F5 incorporate SSL accelerator cards in our hardware devices to perform these functions optimally.

Publishing of MSFT services. Optimizes, accelerates and secures the release of Microsoft services, replacing and improving the architecture and TMG or ISA Server solution.

Deployment of SSL-VPN solutions to cover user access to applications and corporate resources from any device and from any location.

REFERENCE ARCHITECTURE | REVERSE PROXY GATEWAY

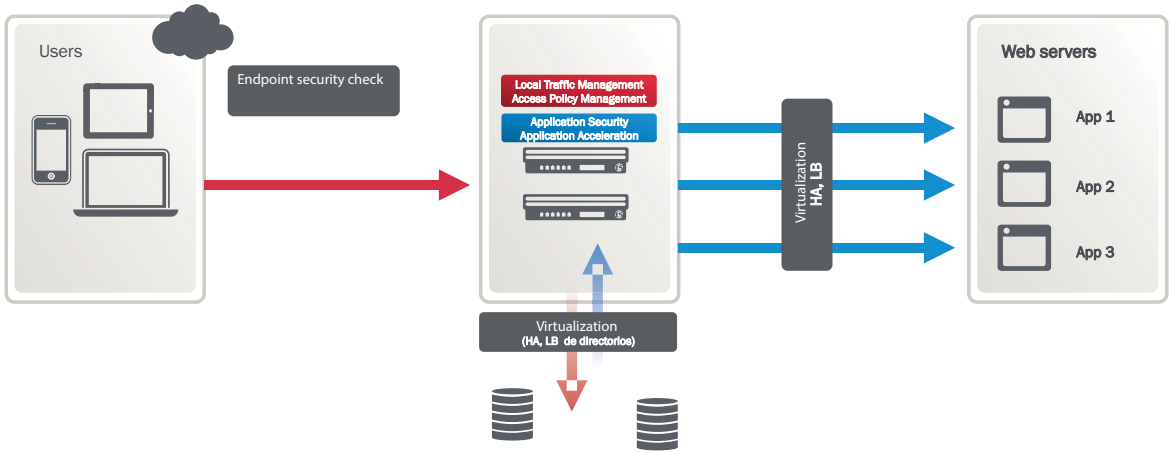

ANTI-BOT PROTECTION IN OWA ENVIRONMENTS

PROBLEM

E-mail is one of the most used (and most vulnerable) applications in enterprise environments. DoS/DDoS attacks on Exchange web portals (OWA) are based on the idea of preventing access to any user authenticated by the Active Directory (AD). This attack works by trying to authenticate a user who enters an incorrect password more than 3 times, which results in the AD automatically activating a method of protection to block access to that user. Of course, when talking about a DDoS attack, this attack is played out across a lot of users (sometimes the attack involves all users registered in the company AD).

The big problem is that such attacks presuppose that no one company can review, manage and analyse all the information in Exchange, plus there are dangers inherent in exposing the information in the mailboxes.

ALTERNATIVES

- You can configure the Active Directory to increase the number of attempts permitted before a user is blocked or leave it unprotected, which is not a solution which can protect users from these DDoS attacks

F5 | APM SOLUTION

F5 proposes adding a level of management with advanced user authentication. This solution (APM - Access Policy Manager) will look to consolidate all users of the applications (both internal and external).

APM can provide a user-access screen with the appearance you need for the application (for example, Exchange) and can manage all user authentication for the application. In this way we can control the number of authentication attempts permitted to a user (less than that permitted to the AD) in order to avoid blocking users in the AD. In the event that this number of attempts is exceeded, it is possible to present other security mechanisms such as two-factor authentication, OTPs, Captcha, etc.

REFERENCE ARCHITECTURE | ANTI-BOT PROTECTION IN OWA ENVIRONMENTS

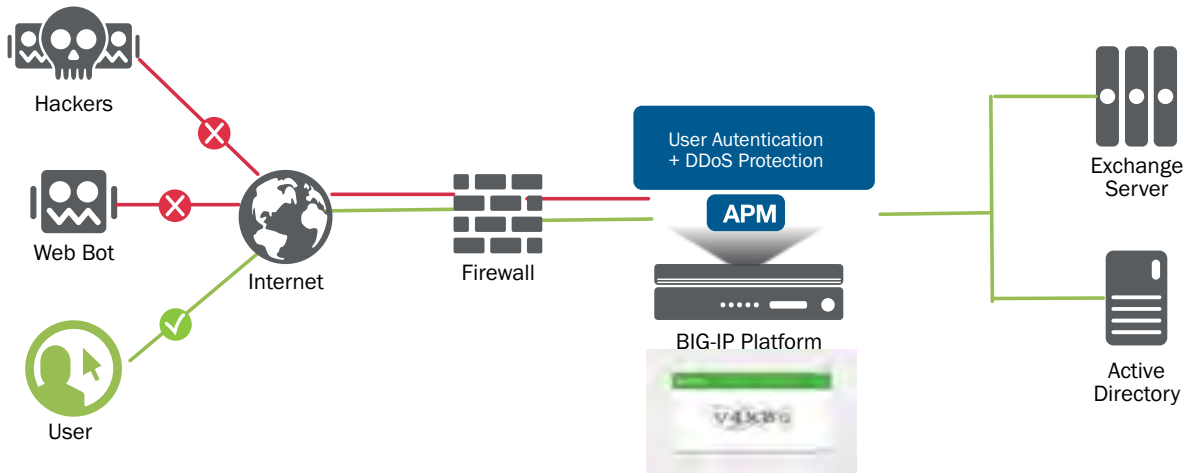

PROTECTION WITH SMART REPUTATION LISTS

PROBLEM

Companies' application servers face daily attacks both from botnets, such as scanners, and from phishing. At the same time, networks must prevent outbound connections to C&C (Command & Control) servers or sites with malware. Detecting and blocking these threats is becoming increasingly difficult, given the need to maintain the high availability of applications.

Applications and websites should be available for legitimate users, but at the same time, companies must be able to prevent illegitimate and malicious access. Infected endpoints attempt to connect to company networks and this causes a loss of network resources. Users can therefore connect to sites with malware. Companies need to prevent the infection of their endpoints, bearing in mind that the threats are persistent and increasingly advanced.

ALTERNATIVES

- NGFW are not scalable in the event of a DDoS attack, nor are they able to filter DDoS IP vectors.

F5 | IP INTELLIGENCE SOLUTION

One of the most effective and economical mechanisms for reducing spam is the real-time blocking of certain mail traffic on the basis of reputation lists.

The F5 IP Intelligence solution adds a layer of security with an updated overview of both incoming and outgoing IP addresses. Companies can protect their data centre assets through the real-time identification of malicious or objectionable IP addresses and configure BIG-IP policies to block their access. The result is a reduced load on the back-end, including firewalls and servers.

IP Intelligence can be implemented in any BIG-IP device or integrated with the ASM (Application Security Manager) (F5's WAF) to have a detailed view of the activity of suspicious IP addresses.

REFERENCE ARCHITECTURE | PROTECTION WITH SMART REPUTATION LISTS

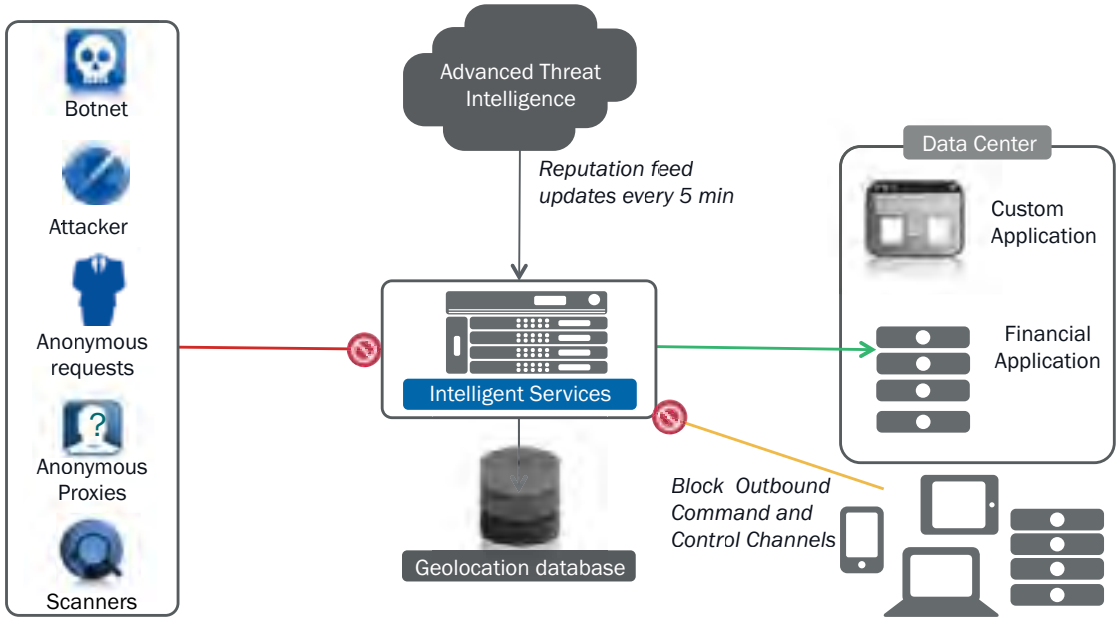

DNS INFRASTRUCTURE PROTECTION

PROBLEM

Companies use the DNS service to provide users with access to web applications. If the DNS service is not available (or performance is compromised), access to these applications cannot be guaranteed.

It is critical to optimize and secure DNS infrastructure to ensure that users can be provided with a service. The operation of this DNS infrastructure requires the ability to respond to a large number of requests per second and the ability to upscale quickly becomes critical when you have to handle thousands of domain names.

It is also necessary to ensure user protection and the integrity of the service against DDoS attacks, DNS cache poisoning and tunnelling.

ALTERNATIVES

- BIND-based solutions are expensive to operate due to the frequent updates required to counter the constant threats.
- Traditional DNS solutions are difficult and costly to scale, there are no flexible and/or advanced integrated security solutions specific to DNS.

F5 | DNS + AFM SOLUTION

BIG-IP DNS + AFM (Advanced Firewall Manager) enables organizations to optimize, secure and monetize their DNS infrastructure. This solution provides LDNS caching services, resolving at carrier-class level and with high performance, and is a hyper-scalable Authoritative DNS solution which includes DNS firewall services, using hardware to mitigate DDoS attacks on the DNS. BIG-IP DNS + AFM provides intelligent and scalable DNS infrastructure, which enables a rapid response to mobile users. By managing customizable monitors and GSLB services, organizations can allocate appropriate resources to answer DNS requests and provide the best user experience. BIG-IP DNS also enables a DNS64 environment for IPv6 environments, with a fault-tolerant infrastructure, optimizing traffic and increasing the quality of service to users, thus protecting the brand and the reputation of the business.

In addition, BIG-IP DNS + AFM protects the DNS infrastructure against malicious attacks generated by infected users and unwanted DNS requests and responses. Intelligent DNS by F5 has a firewall which inspects and validates the protocols and discards or refuses to accept unsolicited responses. It also mitigates attacks by blocking access to malicious domains.

Finally, BIG-IP DNS provides statistics and reports, and includes functionality for high speed logging for DNS to facilitate capacity planning, service optimization and the monetization of services.

REFERENCE ARCHITECTURE | DNS INFRASTRUCTURE PROTECTION

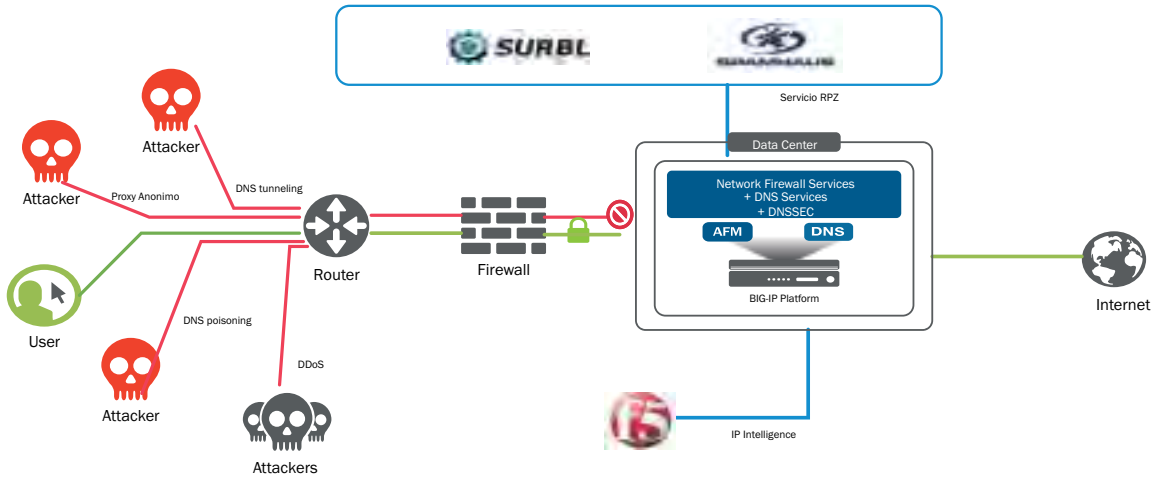

PROTECTION OF PUBLIC-SERVICE PORTALS

PROBLEM

Cities with more than 20,000 inhabitants often provide public services in the form of a «single public-service portal». Online transactions (such as paying property tax, etc.) require a transactional system. Online payments and personal information about the payee are critical data, which means that the city administration has to take care to ensure the security and privacy of electronic transactions.
In addition, the municipalities have a system of access for internal and external employees and in some cases also for citizens. This service requires user identification to prevent unauthorized access, and to control the access to different services and applications based on user profiles. The ever-increasing frequency of attacks necessitates the application of higher-levels of intelligence to enable a user to access a particular resource – a password is no longer enough. It is crucial to consider the user's context, where, how, when, etc.
This same problem is faced by councils which offer these services to smaller municipalities.

ALTERNATIVES

- For the protection of the web application, the alternative is to implement safer application development policies, which may result in organizational conflicts, major delays in the development of applications and a substantial increase in development costs, while taking no action poses a huge reputational risk.
- One of the alternatives for user control is to implement this logic in the application itself, which requires personalizing it, reprogramming, code, tests, etc. Besides being cumbersome, this alternative is not replicable in other applications.

F5 | ASM + APM SOLUTION

The twin issues of securing the web application and securing user access are solved by supplementing the infrastructure with two functionalities by F5.

The F5 ASM (Application Security Manager) module makes it possible to deploy a WAF in a way which is transparent to applications, protecting against web application attacks such as OWASP and DDoS. With this module, municipalities and provincial councils can ensure their citizens that all the information they need to fill or complete in their web applications will be secured and not be in danger of being detected by attackers, they protect the web from L7 DDoS attacks, ensuring that the administering body's image will not be damaged. The ASM offers reporting which contains all the detailed information on the application's level of security. In addition, F5 supports integration with the leading developers of DAST tools, which means that your report can be included quickly, making it possible to avoid most OWASP attacks.

APM (Access Policy Manager) module allows you to set granular access policies for different services, so that the access control can be customized to each application and to each group of individuals. Access to the application takes into account the context of the user, who they are, from where they are accessing the application, to which group they belong. This solution makes it possible to validate the user against multiple types of repositories (LDAP, RADIUS, AD, etc.) to check the security status of the device, to authenticate using advanced mechanisms such as SAML, Kerberos, NTLM, etc. and to introduce two-factor authentication to applications according to the user's context.

REFERENCE ARCHITECTURE | PROTECTION OF PUBLIC-SERVICE PORTALS

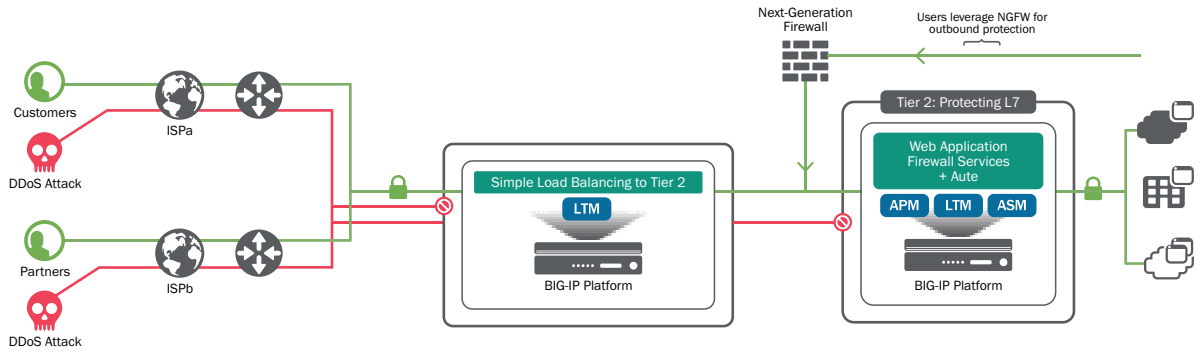

DNS PROTECTION THROUGH DNSSEC

PROBLEM

Originally, the DNS (Domain Name Service) protocol lacked any security protocol.

This has made this critical service vulnerable to attacks such as cache poisoning, which basically consists of spoofing responses to DNS queries to maliciously redirect users. There are several forms of attack which exploit the vulnerability of cache poisoning, such as the «Birthday Attack». The attacker's objective is to modify the contents of the DNS cache such that it does not return the actual IP address of a FQDN, but another IP address, usually that of a server which the attacker controls, and where they can host malicious content.

ALTERNATIVES

- Not using DNSSEC exposes our DNS servers to the threat of cache poisoning.
- Working with restricted access lists greatly complicates the management of the solution, especially with regards to updating whitelists and blacklists.

F5 | DNS SOLUTION

The BIG-IP DNS module natively incorporates the DNSSEC protocol extension, which allows us to use digital signatures to authenticate responses sent from the DNS, as well as signing the answers from our own DNS.

While DNSSEC does not encrypt data, cryptographic authentication prevents users/customers from suffering cache poisoning attacks relating to the protected DNS zone.

BIG-IP uses hardware to perform DNSSEC signing, this results in a solution that is highly scalable.

Additionally, it supports offloading certificates using our own cryptographic cards or we can integrate with external HSMs so that they can confirm the certificates.

REFERENCE ARCHITECTURE | DNS PROTECTION THROUGH DNSSEC

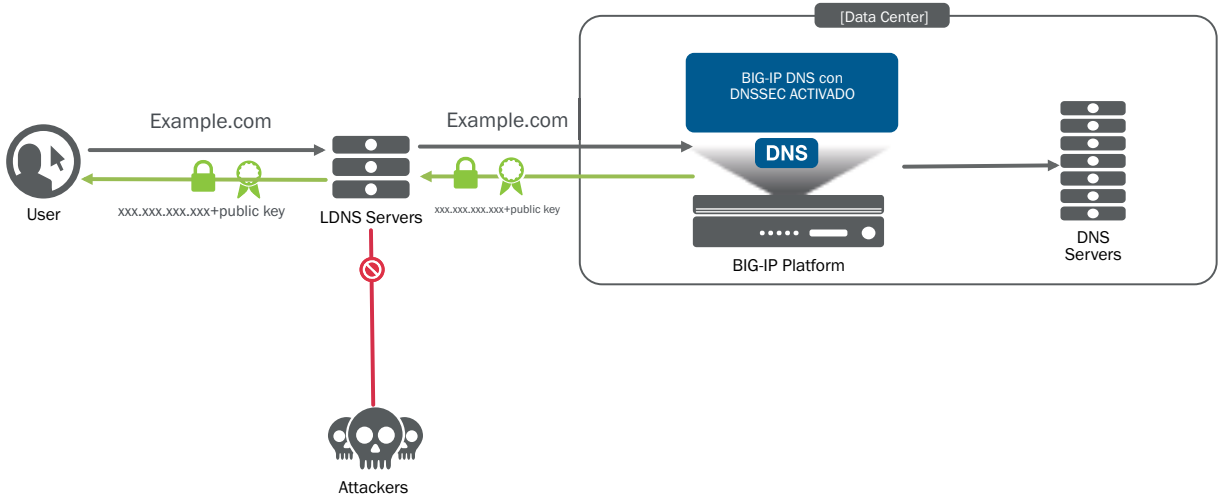

TRANSPARENT FRAUD PROTECTION FOR E-COMMERCE TRANSACTIONS

PROBLEM

Every day there is an increase in the level of economic fraud in e-commerce transactions. This fraud is facilitated by two factors: the vulnerabilities of web applications (code level) and web malware installed on users' browsers.

Faced with such attacks, users and companies face two problems: The economic impact of the attack (at both user and enterprise level) and the reputational impact of the attack (at enterprise level).

ALTERNATIVES

- Other solutions require the installation of clients, so integration is intrusive and may affect the performance of web applications.

F5 | WEBSAFE + ASM SOLUTION

F5 is the only manufacturer to offer a complete and transparent solution to protect against economic fraud in the web environment and for the protection of users and web applications:

ASM (Application Security Manager) is the WAF solution which protects against all vulnerabilities and DDoS attacks at the web-application level. ASM protects the web infrastructure from the vulnerabilities inherent to software development and/or in the platforms (Apache, IIS, Tomcat, MySQL etc.).

WebSafe is a clientless anti-fraud solution, enabling seamless integration with existing web applications and all existing web browsers.

The solution safeguards the ROI, avoiding economic losses and the loss of information relating to users, applications and web platforms, thereby safeguarding the company's reputation.

REFERENCE ARCHITECTURE | TRANSPARENT FRAUD PROTECTION FOR E-COMMERCE TRANSACTIONS

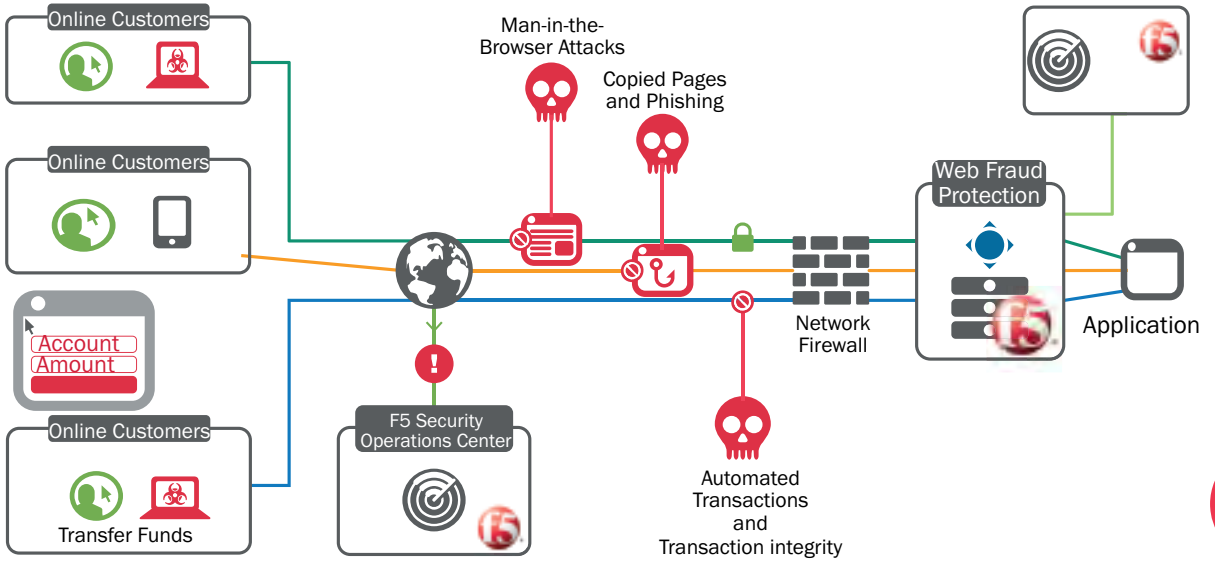

JUNIPER SSL VPN REPLACEMENT

PROBLEM

Juniper has sold its business line to Pulse Secure, which has led to two problems:

- Firstly, there is great uncertainty in the evolution of their SSL VPN access solution.
- Secondly, there are difficulties with regards to the ongoing support and maintenance of the platforms.

The Pulse SSL VPN solution has limited support for integration with third-party tools, for example authentication systems such Latch, MDMs, etc.

ALTERNATIVES

- UTM/Next-generation firewalls
- These do not allow granular control, only VPN-SSL-VPN. They do not consolidate either authentication or access.
- Limited support types of authentication and strong authentication.

F5 | LTM + APM SOLUTION

BIG-IP APM (Access Policy Manager) is the market-leading solution, which unifies SSL VPN services and the management of authentication and user accesses, integrating SSO services and Federation of identities services into the same solution.

It provides visibility and control of user and application traffic and supports all types of devices and operating systems.

The solution is much more scalable than the competition (up to 200,000 simultaneous SSL VPN clients per chassis). There is native support for VDI (Microsoft, VMWare, Citrix) and support for most authentication mechanisms and strong authentication (NTLM, Kerberos, SAML, digital certificates, tokens, OTPs, etc.).

It supports full integration with MDM solutions such as AirWatch and MobileIron.

REFERENCE ARCHITECTURE | JUNIPER SSL VPN REPLACEMENT

- Secure VDI
- Advanced AAA & Endpoint security
- Traffic management
- Scalable –200,000 concurrent users
- Offload/replace Security Server

(5) MAG6611 with 4 SM360 blades = 40,000 CCUs

(1) VIPRION C4800 = 200,000 APM CCUs

MICROSOFT FOREFRONT TMG REPLACEMENT

PROBLEM

In September 2012 Microsoft announced its decision to discontinue the product «Forefront Threat Management Gateway» (TMG), formerly known as Microsoft Proxy Server or Microsoft ISA Server.

The TMG solution acts as a reverse proxy for publishing applications (mainly Exchange, Lync, SharePoint, etc.) to remote users. TMG has two components: One, a very basic traffic manager to balance HTTP and HTTPS and two, an authentication client to consult with different user directories (LDAP, Radius, Kerberos, etc.)

For businesses, having a product in a production environment without a guarantee of continuity and support is a critical problem, so it is necessary to consider its replacement.

ALTERNATIVES

- Continuing to use the Microsoft TMG solution is not a viable alternative. Using computers without direct support from the manufacturer poses a serious security risk to applications published through the Forefront TMG platform and ultimately to the client's network infrastructure.

F5 | LTM + APM SOLUTION

Microsoft recommends F5 to its customers, as the leading solution in the market.

F5's BIG-IP Local Traffic Manager (LTM) is the leading solution in the ADC (Application Delivery Controller) market. BIG-IP LTM converts the network into an agile infrastructure for application delivery. It offers scalability, high availability and high performance with the BIG-IP platform. APM (Access Policy Manager) allows us to configure the security policies necessary to provide granular control of user access and remote devices.

One advantage of acquiring the F5 solution is the ability to consolidate all functionalities into one device, thus simplifying the infrastructure and providing the functionality of load-balancing and user-access authentication.

REFERENCE ARCHITECTURE | MICROSOFT FOREFRONT TMG REPLACEMENT

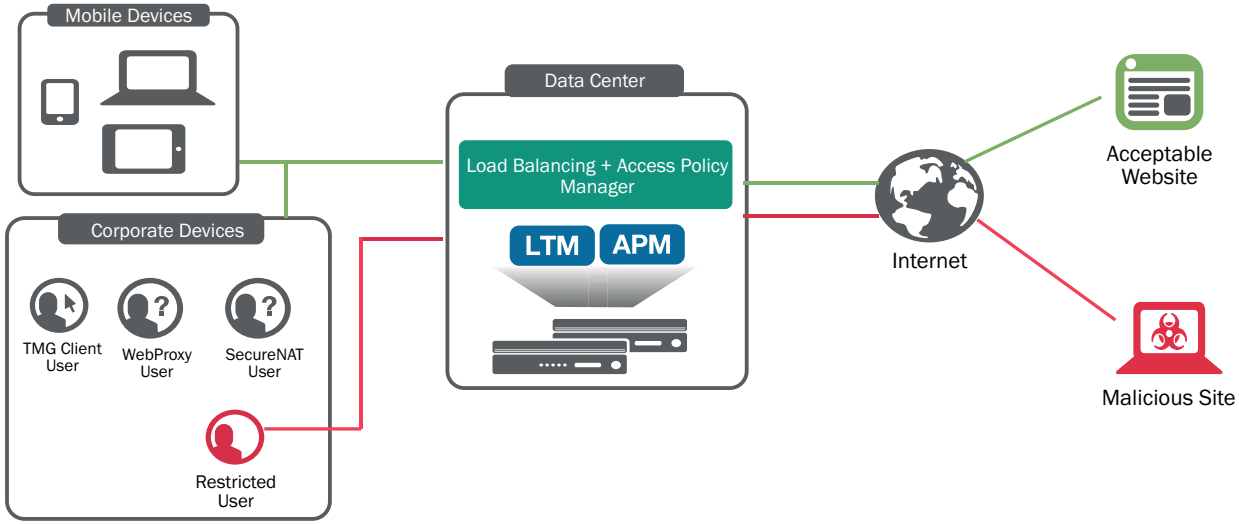

REPLACEMENT OF DNS PLATFORMS BASED ON BIND

PROBLEM

BIND is the most-deployed DNS software on the internet and is standard on many Unix systems. Designed in 1980, it has now evolved to version 9. BIND is based on open-source code and the list of security vulnerabilities is public, allowing any attacker to exploit them and attack the DNS service. This forces BIND administrators to patch the BIND platform on a continuous basis, resulting in interruptions in service and, ultimately, it makes BIND less reliable as a DNS solution.

In 2013 and 2014, vulnerabilities such as DNS Amplification affected DNS services based on BIND in a particularly virulent way, as many of them are deployed in production with default parameters. BIND is based on a general-purpose operating system, making it difficult to scale the solution and limits its performance, which makes it particularly sensitive to DDoS attacks.

ALTERNATIVES

- Continuing to use DNS systems based on BIND is an inexhaustible source of problems and security vulnerabilities, which can affect not only the company or organization, which uses them, but can be the source of DDoS attacks (such as DNS Amplification) on third-parties. This triggers the operating costs of the DNS platform and impacts your company's reputation.
- The performance and scalability of solutions based on BIND is much lower than that of solutions based on proprietary developments with dedicated hardware acceleration. This makes them vulnerable when faced with large numbers of requests for DNS resolution, even when they are legitimate.

F5 | DNS SOLUTION

F5's DNS solution is based on proprietary development and is not vulnerable to the security flaws which are continually discovered in the BIND code.

Being hardware accelerated, F5 DNS is a hyper-scalable solution, reaching rates of 40 million RPS (responses per second) and protecting businesses and organizations against DDoS attacks.

F5's DNS supports and accelerates DNSSEC, unloading the demanding validations of this protocol onto hardware acceleration to ensure fast responses.

REFERENCE ARCHITECTURE | REPLACEMENT OF DNS PLATFORMS BASED ON BIND

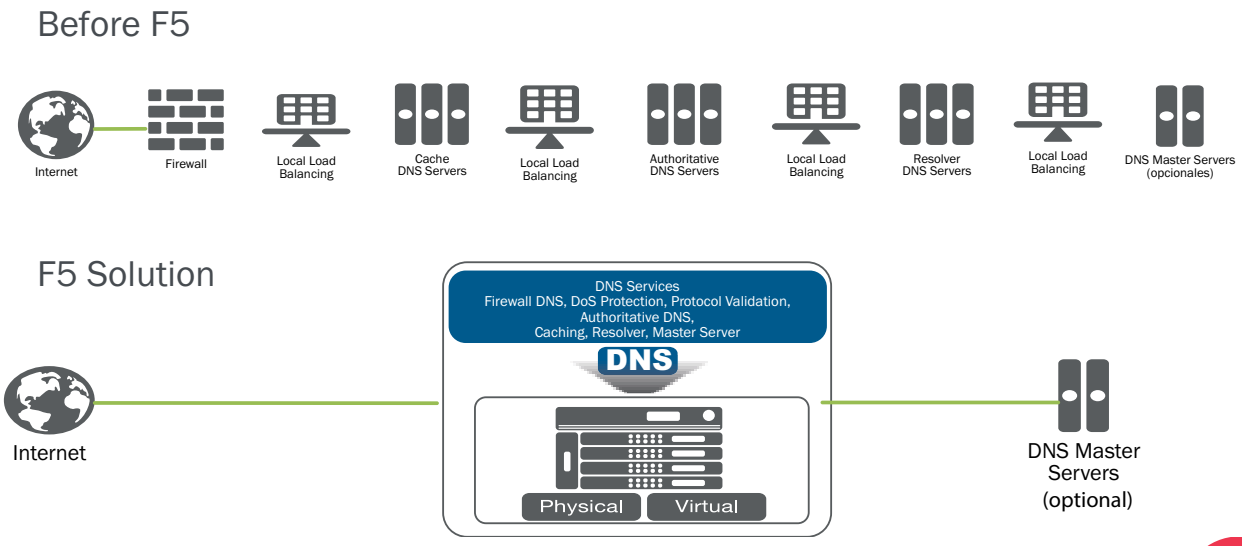

ALWAYS-ON SOLUTION FOR MOBILE ENVIRONMENTS

PROBLEM

Roaming mobile users require the use of stable and secure VPN connections. One of the biggest problems users encounter is loss of connection, particularly when roaming. A user seeking access from a mobile device can be connected to a Wi-Fi network at one moment and a 3G/4G connection at the next. Each time the connection breaks, the VPN client need to reconnect.

ALTERNATIVES

- In-house development.

F5 | APM SOLUTION

APM (Access Policy Manager) module by F5 makes it possible to deploy VPN-SSL access solutions to cover user access to applications and corporate resources from any device and from any location.

With the deployment of the F5 Edge VPN client for mobile and desktopAPMusers can configure always-on connectivity. Such functionality enables automatic connection from the client device to the corporate network without requiring the client to take any action. As soon as there is an active data connection, the Edge Gateway client automatically establishes a tunnel to the corporate network.

REFERENCE ARCHITECTURE | ALWAYS-ON SOLUTION FOR MOBILE ENVIRONMENTS

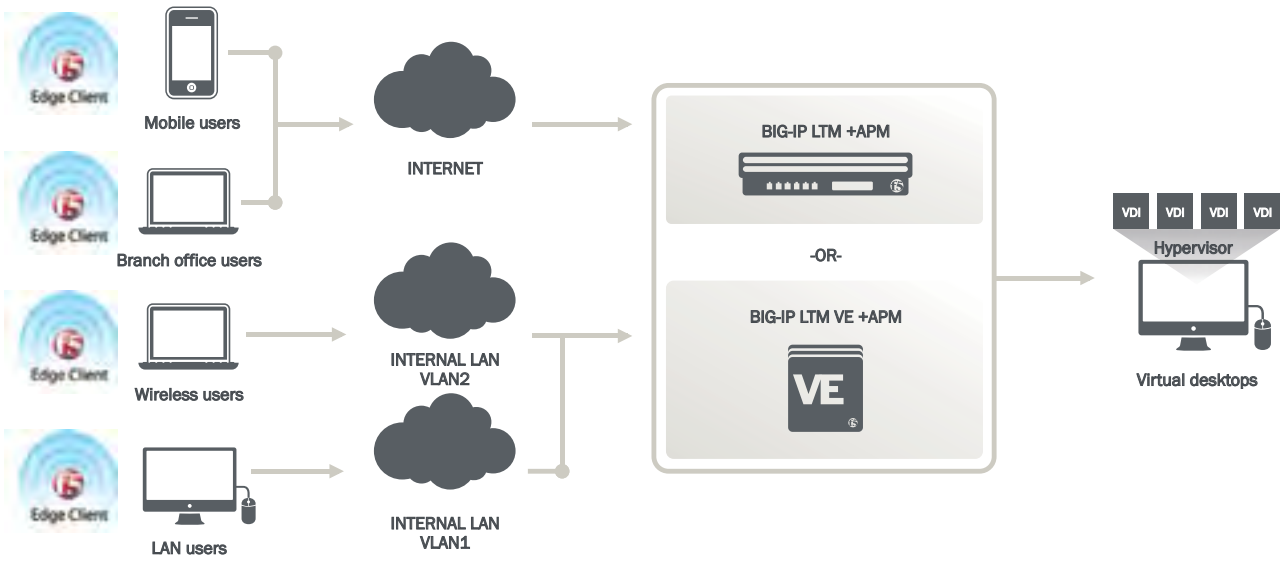

KERBEROS AUTHENTICATION SOLUTION FOR MOBILE DEVICES

PROBLEM

The Kerberos authentication protocol is recognised as one of the most reliable and security authentication protocols available today as it enables access to specific services (SPN) for a finite time (the duration of a Kerberos ticket). In addition, it is a way to implement SSO in the company, as it requires no user interaction. Therefore, Kerberos authentication is the most-highly recommended by security consultants.

The problem is that Kerberos is limited to Microsoft client devices in which the user logs in directly to the domain. The high proliferation of mobile devices requires a similar authentication solution for these users to authenticate themselves securely in external services deployed by the corporation.

ALTERNATIVES

- There are only two vendors who can support the implementation described above (with Kerberos authentication certificates) for mobile devices: F5 with APM module and Microsoft with the TMG proxy. TMG has now be discontinued by Microsoft..

F5 | APM SOLUTION

APM (Access Policy Manager) module implements both standard Kerberos authentication and Kerberos Constrain Delegation (KCD). KCD facilitates Kerberos authentication for users who are not logged on to the internal Windows domain, i.e., external users who want to access internal services.

For this, APM needs to first authenticate the user and, optionally, the device of the external user. APM then makes an internal request to the KDC (Kerberos Domain Controller) for a Kerberos ticket in the name of that user. A widespread and reliable solution is to authenticate external users with mobile devices on the basis of a certificate installed on these mobile devices. APM is able to validate the certificate, extract the necessary user information and request a Kerberos ticket (Kerberos Delegation) on their behalf to provide these external users with secure access. This implementation not only provides a strong Kerberos authentication solution for mobile devices, but also transparent user authentication (SSO).

REFERENCE ARCHITECTURE | KERBEROS AUTHENTICATION SOLUTION FOR MOBILE DEVICES

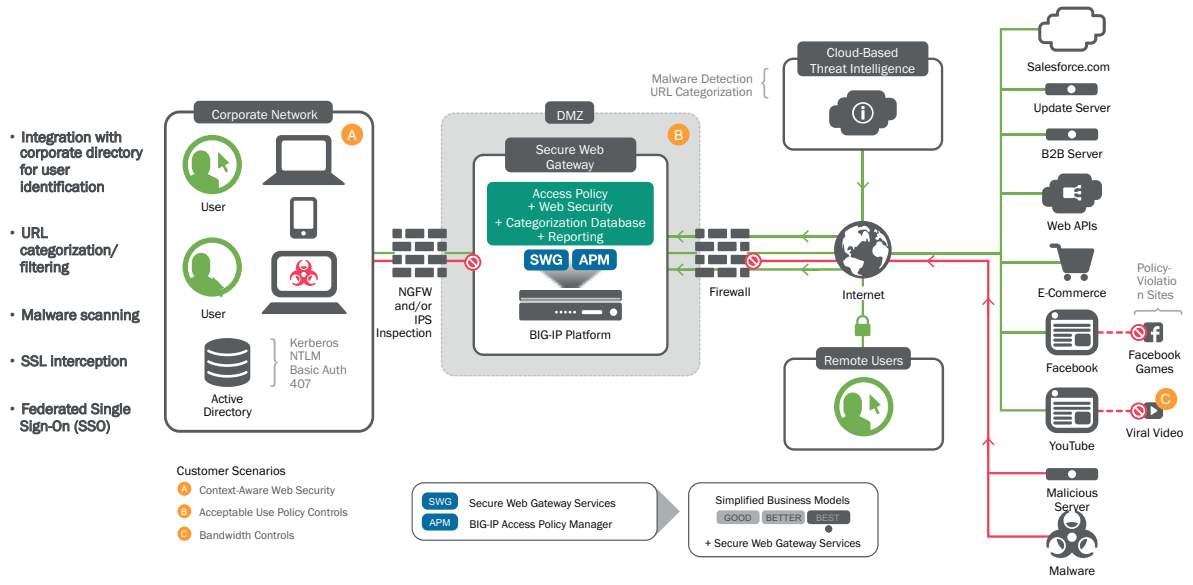

REMOTE VPN SSL ACCESS SOLUTION

PROBLEM

Organizations need to provide remote user access to corporate resources (as well as those which are less corporate, such as an extranet) to facilitate teleworking, access to internal applications, applications which require access to multiple ports, local IPs, or other variations.

The traditional approach of installing local agents on devices involves additional management overheads, versioning, and multiple constraints due to the traditional IPSEC encapsulation (which is often blocked when accessing the internet via a proxy or when within 3rd party environment).

ALTERNATIVES

- Using the client's IPSec tunnels does not work in certain environments (behind a web proxy, firewalls, etc.), which reduces the versatility and value of the solution.
- Solutions which require a software client to be installed on the access devices do not always cover the needs of the environment (e.g. BYOD), and also require the clients to be managed, which reduces the scalability and agility of the solution.

F5 | APM SOLUTION

F5's Access Policy Manager facilitates flexible models for remote access to corporations so that each user group is presented with only the resources for which the user has permission.

These connections may be further customised according to the user's context so that the user is presented with a portal of applications or remote access depending on their context. For example, if the user does not have updated antivirus software they can be given access to a remediation portal where they can update it, but not allowed access through an IP tunnel.

When providing access via an SSL-encrypted VPN tunnel remote access can be supported even from behind a web proxy or firewall.

APM offers the option of using a software VPN client which facilitates the configuration of automatic session establishment and supports most operating systems, both desktop (MS Windows, Mac OS X, Linux) and mobile (Android, IOS, Windows Phone...).

REFERENCE ARCHITECTURE | REMOTE VPN SSL ACCESS SOLUTION

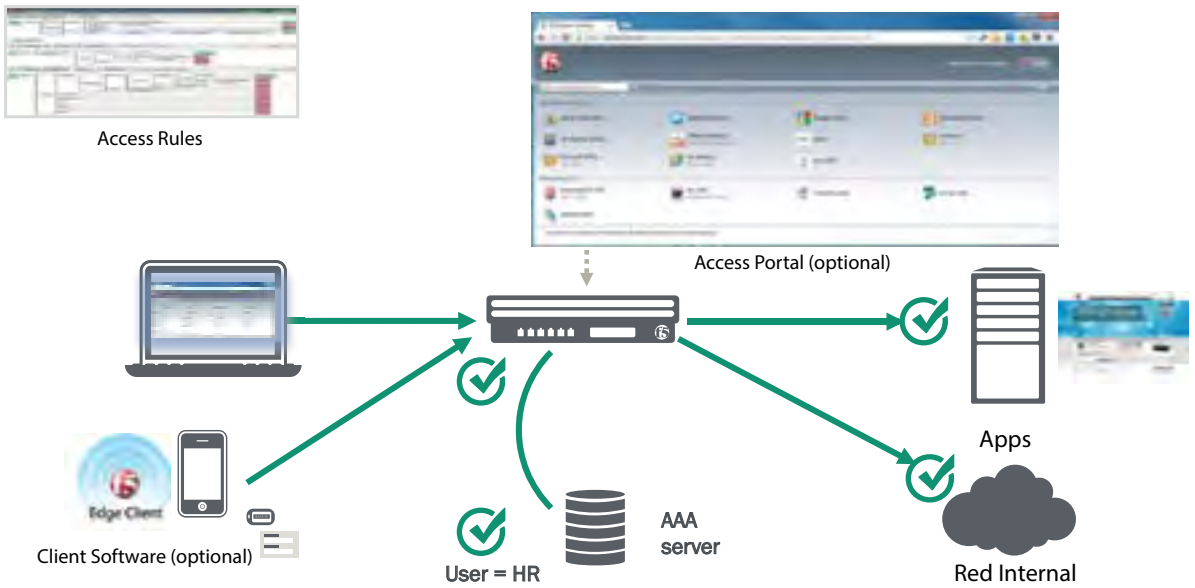

SSL/TLS OFFLOADING SOLUTION

PROBLEM

The use of HTTPS (HTTP over SSL/TLS) communications over the internet grows daily. OTT (Over-The-Top) content services, social networks, web browsers, peer-to-peer services, etc. - the list is almost endless. Publishing unencrypted web services on the internet is already an unusual practice - and an unsafe one. The arrival of HTTP/2 (where encryption is «optional») has once again brought the issue back into debate.There are currently no implementations of browsers which support HTTP/2 without encryption. Using SSL/TLS to encrypt and decrypt communications, although secure, is very costly in terms of CPU usage and jeopardizes the performance of the web servers themselves. The migration from 1024-bit keys to the new of 2048-bit keys (as recommended by NIST in 2011) has resulted in fivefold (x5) increase of the computing resources web servers require to handle the same number of SSL/TLS sessions, reducing the number of connections per second that those same servers can manage by 80%. This model does not scale.

ALTERNATIVES

Increasing the number of front-end or application servers to address the increase in resources required to process the SSL/TLS traffic is an alternative which is expensive both economically and operationally.
In addition, the independent management of the certificates in each of the servers is not efficient, and increases the risk of errors in cases where it is necessary to manage potentially dozens of certificates per server.
Acceleration solutions based general-purpose hardware liberate the servers from the computational load but are limited by the same factors and do not scale properly.

F5 | LTM SOLUTION

The F5 platform has hardware dedicated to the acceleration of encryption and decryption of SSL/TLS freeing the server from the processing load generated by performing cryptographic key exchanges.

SSL/TLS offloading enables organizations to migrate 100% of their communications to SSL/TLS to increase security with no impact to the front-end or application servers and provides a common repository into which to consolidate the certificates, thereby centralizing their management.

In environments where end-to-end encryption is required, it is possible to use smaller keys (for example 1024-bits) for connections from BIG-IP to the internal application servers and use larger keys (2048-bits or higher) when connecting to external clients.

F5 offers models which meet the FIPS standard, and can be integrated with third-party HSM solutions.

REFERENCE ARCHITECTURE | SSL/TLS OFFLOADING SOLUTION

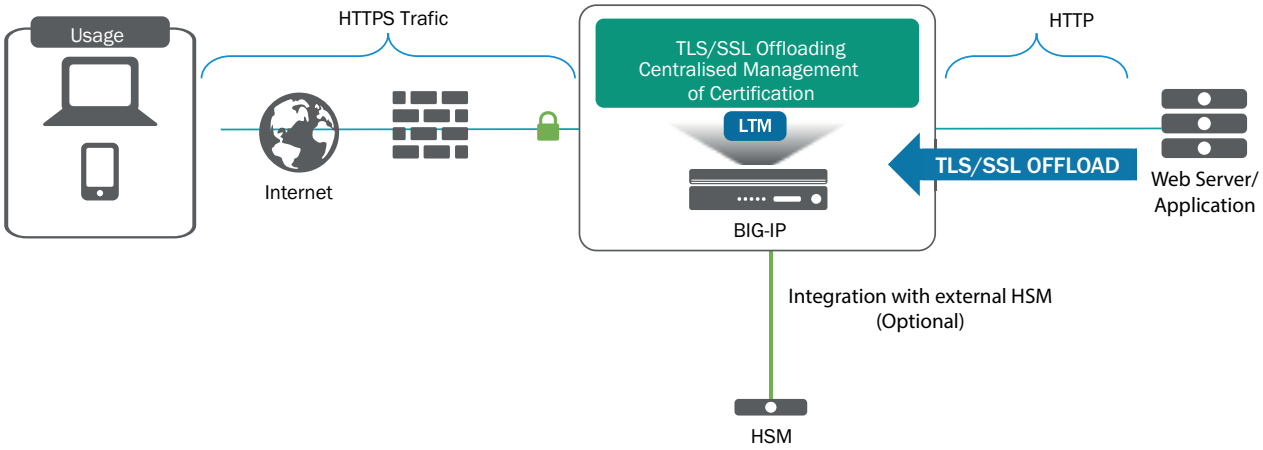

SSL/TLS VISIBILITY SOLUTION

PROBLEM

SSL/TLS is an encryption protocol which secures internet connections using public key encryption and certificates to encrypt communications and verify the identity of the endpoints, thereby ensure confidentiality and integrity. One of its most important uses is to secure otherwise unencrypted HTTP. This is used, for example, to secure web pages in e-commerce applications but is also used by attackers to circumvent network security devices.

The task of encrypting/decrypting HTTPS traffic impacts significantly on the performance of the network elements which are used to inspect the traffic (firewalls, IDS/IPS, antivirus, URL Filtering, DPIs...). On average, the performance of a NGFW falls by 81% when it has to inspect encrypted traffic. So we have two problems: one of visibility and another of resource efficiency.

ALTERNATIVES

Traditional security service architecture is usually sub-optimal and is generally based on solutions with a high CPU consumption (and are therefore costly) and cannot inspect more secure key exchanges such
as Elliptic Curve Diffie-Hellman (ECHDE).

F5 | LTM SOLUTION

F5's full proxy architecture means that separate SSL/TLS connections can be established between the client and LTM and between LTM and the application server. This enables LTM to see in to previously encrypted communications and deliver the traffic in the clear to the web servers or network elements, which are used to analyse them (firewalls, IDS/IPS, antivirus, URL Filtering, DPIs...).

In this way, we can provide the necessary visibility and consequently reduce the size and cost of the equipment which receives this traffic (FireEye, PAN, CheckPoint, SourceFire, etc), thereby also facilitating horizontal scalability.

Unlike other solutions on the market, which are based on software SSL man-in-the-middle techniques, F5 uses hardware acceleration to process SSL traffic (SSL offloading), this delivers great performance at an affordable cost.

REFERENCE ARCHITECTURE | SSL/TLS VISIBILITY SOLUTION

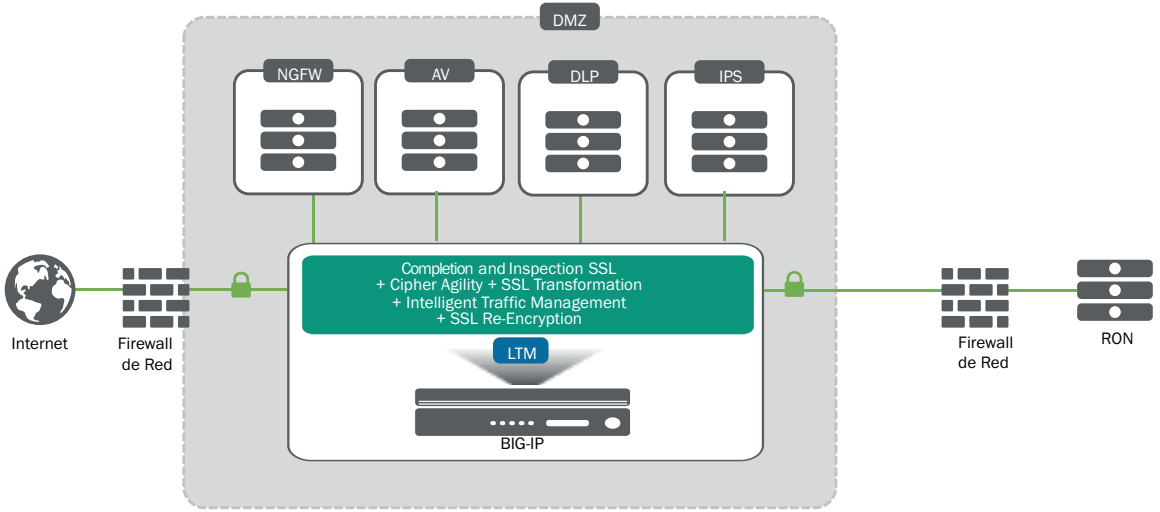

SINGLE SIGN-ON SOLUTION

PROBLEM

«Single Sign-On» is a conceptual authentication mechanism which allows the user to access multiple applications with one set of unique credentials and so is convenient, practical and safe.
Today, the number of applications we use (including corporate applications) has increased significantly. Accessing these applications requires users to enter their credentials again and again representing a decrease in productivity and in the quality of the user experience.
In addition, Software as a Service (SaaS) solutions, like Salesforce, Office 365, SharePoint Online, etc., create a security problem because users end up using the same passwords for all of them and may even write them down so as not to forget them.

ALTERNATIVES

- Failure to implement SSO mechanisms in our applications or SaaS solutions leads to a decrease in productivity and security for users and the company.
- Using SSO solutions based on clients (applications or devices) requires these clients to be managed and introduces potential compatibility problems (between applications, browsers, etc.).

F5 | APM SOLUTION

F5's Access Policy Manager (APM) makes it possible to authenticate the user using multiple parameters, depending on the context, to allow or deny access to a particular application.

Once the user has been authenticated by APM, their credentials are automatically sent to applications as the user accesses them, eliminating the need to re-enter them. APM sends these credentials in the format that each application requires, as if the user was doing so him-/herself.

APM implements SSO mechanisms in a centralized manner which is transparent to applications.

For example, it would be possible for APM to authenticate a user by inspecting a digital certificate and for APM then to authenticate that user to one application using a web form and to another application using Kerberos authentication.

REFERENCE ARCHITECTURE | SINGLE SIGN-ON SOLUTION

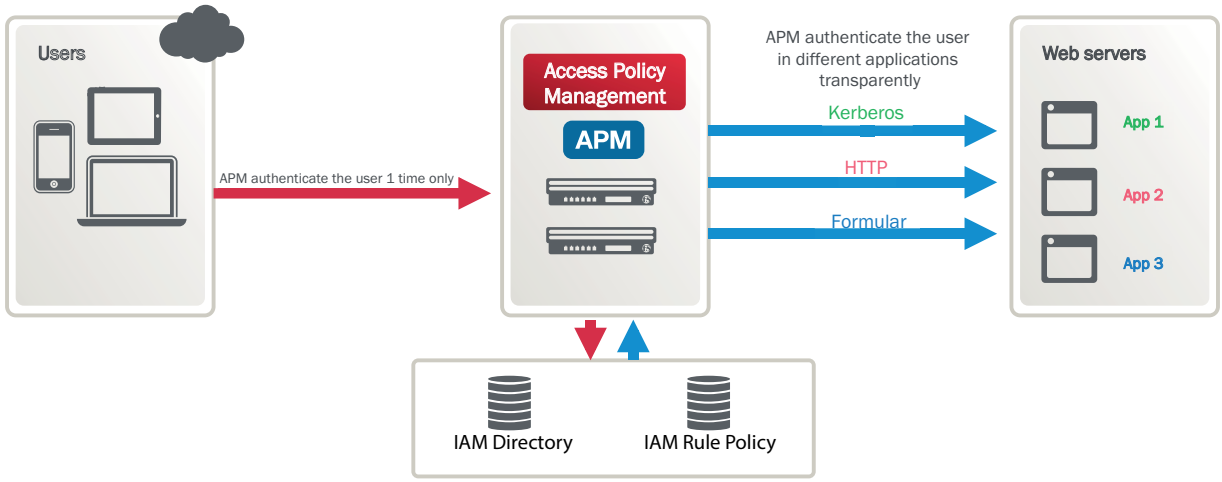

WEB APPLICATION FIREWALL (WAF) SOLUTION

PROBLEM

In many cases, the security of a company depends greatly on the security of web applications which are developed by other departments and which you cannot control or, worse, are developed by third parties outside your organization which place more emphasis on functionality and speed of development than on security. Typical attacks such as those which are known as the OWASP Top Ten (SQL-injections, cross-site scripting...) may result in the leakage of valuable information from the company with major economic and even legal consequences for the organization. Similarly, web (L7)-specific DDoS attacks may result in loss of service for our customers. In addition, the PCI-DSS regulations require the use of web-protection devices for the obfuscation of sensitive content.

ALTERNATIVES

- Implementation of more secure policies for application development, which may result in organizational conflict, long delays in the development of applications and a substantial increase in development costs.
- Taking no action poses a huge risk and means non-compliance with regulations.

F5 | ASM SOLUTION

F5's ASM (Application Security Manager) module makes it possible to deploy a WAF quickly and easily, in a way which is transparent to applications and thereby to protect against web-application attacks such as Cross-Site-Scripting and SQL-injection. It also includes a module for monitoring compliance with PCI-DSS regulations.

After an initial learning period in which the ASM comes to recognize the web application's normal patterns of operation, the WAF protection is activated, this is specific to each URL and the user can specify the threshold at which the automatic protection is triggered.

The equipment can be configured to work with whitelists or with blacklists.

REFERENCE ARCHITECTURE | WEB APPLICATION FIREWALL (WAF) SOLUTION

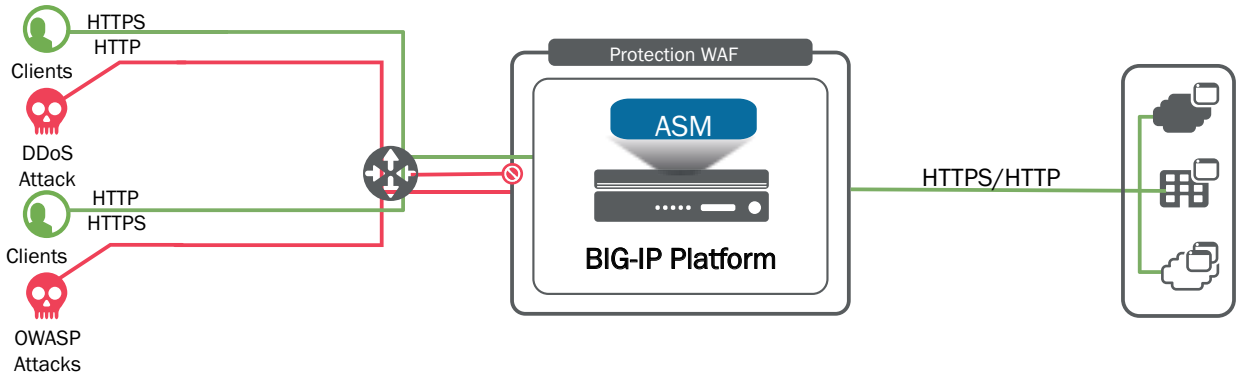

PROXY SOLUTION FOR SAFE NAVIGATION FOR EMPLOYEE

PROBLEM

The pattern of internet usage has changed in recent times. The use of social networks, IM messaging and collaborative tools has mushroomed.

Internet use by employees can be productive and to the benefit of the business or for leisure or personal purposes, which requires granular control of these accesses.

Another risk is the appearance of malware from web pages or from documents accessed via browsers.

Finally, it is also important to control the use of webmail (Hotmail, Gmail, etc.) which can be used to send (leak) confidential corporate information.

In a nutshell, today, companies need to monitor the internet usage of all their users in order to secure their assets.

ALTERNATIVES

- F5 is the only solution on the market which provides access control of both «inbound» traffic (APM: VPN, SSL, VDI, Mobility) and «outbound» traffic (safe browsing).
- The combination of F5 & Websense, outperforms other proxy-caching technologies, both in terms of performance and in the level of their intelligence with regard to detecting malware.

F5 | APM + SWG SOLUTION

F5's safe web proxy solution (SWG: Secure Web Gateway), combines high performance and scalability with the most advanced techniques for detecting malware in real time and for filtering content. For the latter, F5 chose technology by Websense, a leading vendor in the field of content inspection and incorporated their databases of categories, as well as their advanced inspection engine (ACE Technology by Websense) within F5's devices. This allows you to apply a secure and granular policy to users accessing social networks, IM, webmail, etc.

APM (Access Policy Manager) is able to authenticate users in multiple ways and can also combine authentication techniques: AD, NTLM, Kerberos, Radius, tokens, certificates, SAML, etc. With F5's SSL Visibility technology, you can even inspect encrypted traffic transparently and select the SSL traffic to inspect in a granular manner, with respect for regulations such as the Data Protection Act (e.g. banking or health services).

APM and SWG also offer a powerful solution for logging and reporting, this makes it possible to create advanced reports on user navigation, making it easier to monitor.

REFERENCE ARCHITECTURE | PROXY SOLUTION FOR SAFE NAVIGATION FOR EMPLOYEE

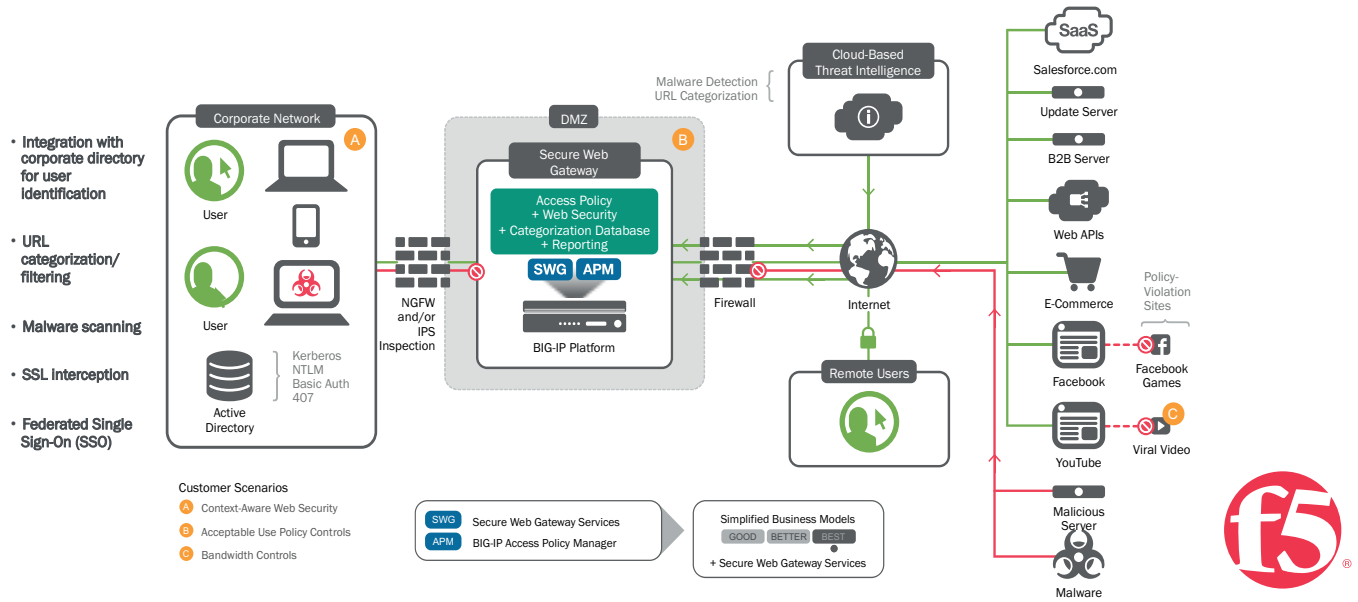

ENHANCED WEB APPLICATION FIREWALL (E-WAF) SOLUTION

PROBLEM

While a WAF solution protects our applications from malicious use through potential DDoS or vulnerable web application attacks, it cannot protect users of our applications against phishing attacks, RATs, man-in-the-Browsers, keyloggers...

This protection is particularly complex, since we have no access to our customers' devices (fixed and/or mobile) and therefore cannot install any clients on these devices, which also use a multitude of different operating systems and browsers.

Theft of user credentials, phishing attacks and automated malicious transactions greatly harms the reputation of organizations and it can lead to significant economic losses with legal implications.

ALTERNATIVES

- Solutions from other vendors require the installation of clients on the server and/or on the client machine which is not feasible in most cases due to the interaction between the client and the application software itself (in the case of servers) or due to the use of devices, which are outside our organizational remit (for customers).
- Trusting only in the WAF module does not protect the users of our applications.

F5 | ASM + WEBSAFE SOLUTION

The combination of F5's ASM (Application Security Manager) and WebSafe modules protects both enterprise applications and those who use them without the need to install any clients, modify applications, or manage end-points either fixed or mobile.

REFERENCE ARCHITECTURE | ENHANCED WEB APPLICATION FIREWALL (E-WAF) SOLUTION

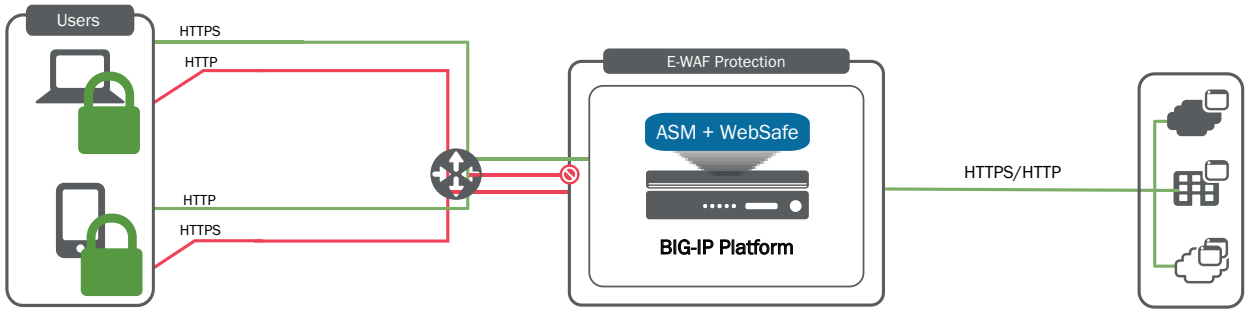

SERVICE PROVIDER

CONTENT

67 Optimization of VoIP solutions
68 VoLTE and IMS Protection
69 Replacement of Citrix ByteMobile
70 PEM analytics solution
71 CGNAT solution
72 Interconnected signalling solution
73 Solution for monetizing OTT services
74 Intelligent DNS firewall solution
75 Gi firewall solution
76 LTE-Roaming solution
77 SIP/IMS signalling solutions
78 On demand bandwidth availability solution

OPTIMIZATION OF VoIP SOLUTIONS

PROBLEM

Numerous organizations use VoIP solutions in which different devices are manually registered in a data centre, which was designated at the point where the VoIP solution was designed, without taking into account principles such as geolocation, the server-load associated with the service, etc...

Signalling and call control rely mainly on SIP (Session Initiation Protocol) in conjunction with protocols such as SDP, which describes the multimedia information in the session, such as the IPs, ports and codecs to be used in the communication.

In general, these solutions were not intended to take into account the location of the customer's phone nor the condition of the organization's global distributed VoIP system.

ALTERNATIVES

- In business environments as globalized as the ones which exist today, there is a particular interest in the optimal use of the distributed VoIP systems, which already exist and in improving their use from the point of view of redundancy and reduced management costs.
- Alternatives based on manual reconfigurations and the use of dynamic routing protocols are not valid solutions in the context of the current requirements for these solutions.

F5 | LTM + DNS + AFM + APM SOLUTION

F5 allows VoIP devices (phones, smartphones, PCs, etc.) to be registered in an intelligent manner, always redirecting the registered device to the most optimal data centre.

F5 has the native capability to translate information originating from the client SIP element and to make the modifications necessary for the proper functioning of these elements, undertaking the functions of a SIP proxy and adding security features and optimization.

Using AFM (Advanced Firewall Manager) and APM (Access Policy Manager) technologies F5 can improve security and performance in the context of access to distributed contact-centre solutions.

REFERENCE ARCHITECTURE | OPTIMIZATION OF VoIP SOLUTIONS

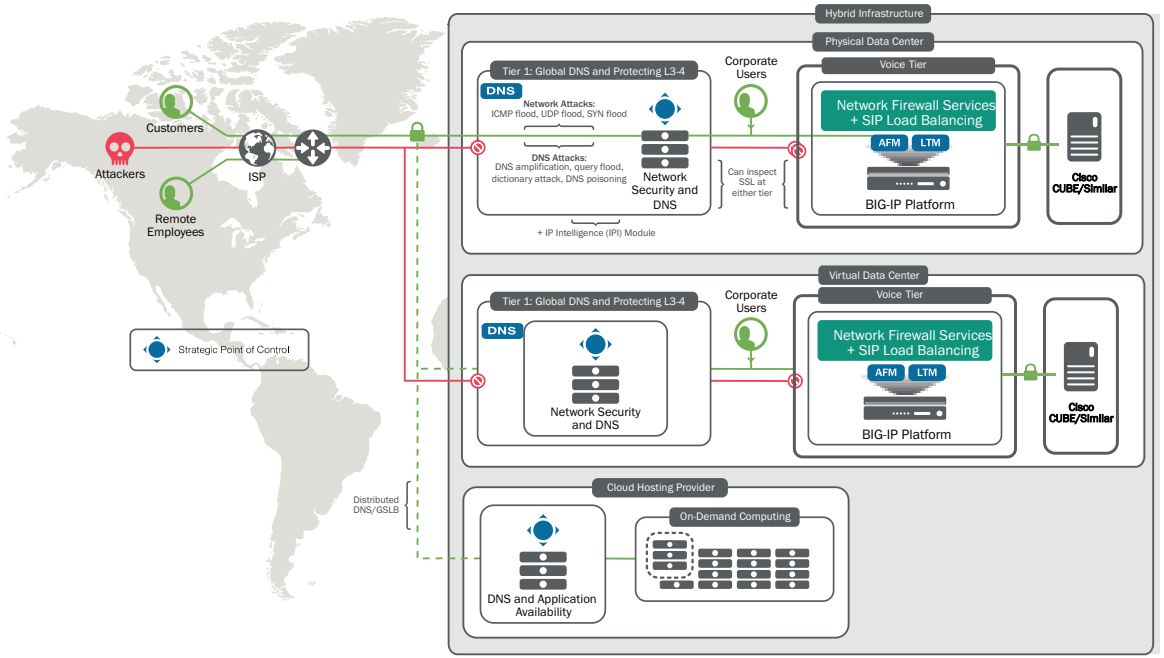

VoLTE AND IMS PROTECTION

PROBLEM

VoLTE (Voice over LTE) is a strategic service for operators who are trying to consolidate all of their services into an all-IP network in order to reduce operating costs and increase flexibility. As this technology becomes more popular (in the USA about 40% of calls are made on it), attacks are increasing and they are focussing on the associated signalling. Therefore, securing VoLTE means focussing on protecting and controlling signalling protocols, including Diameter and SIP and this becomes a critical task.

ALTERNATIVES

- Less scalable and more dispersed solutions which are difficult and expensive to manage, while increasing energy consumption.

F5 | SDC SOLUTION

F5's SDC (Signalling Delivery Controller) helps operators to develop VoLTE services in a way which is both agile and secure. The SDC and the BIG-IP platform which supports SIP ALG (Application Layer Gateway) provide continuity for the VoLTE service, protecting against unauthorized access, unexpected peaks in traffic, signalling storms and spoofing sessions. The F5 FW with SIP ALG monitors SIP messages and only allows RTP flows when it can validate the SIP channel control, thus protecting user traffic. Combining a firewall, Application Delivery capabilities, DDoS protection and the control of signalling across the P-CSCF, this solution is capable of securing and properly distributing traffic, even in periods of unexpected stress.
F5 provides the highest level of protection in the industry, together with the highest connection rates, offering more than a terabit of bandwidth and 1.2 billion concurrent connections. At the same time, the F5 solution is energy efficient and consumes approximately 50% less energy than similar competing solutions.

F5 enables a fast and secure VoLTE service. It provides the highest possible security.

It supports high numbers of connections (per second and concurrent).

REFERENCE ARCHITECTURE | VoLTE AND IMS PROTECTION

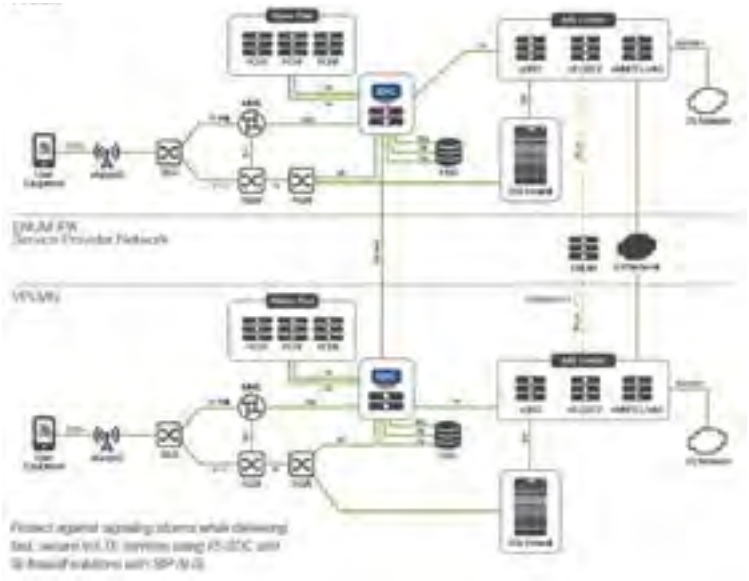

REPLACEMENT OF CITRIX BYTEMOBILE

PROBLEM

Traditionally, operators were able to familiarize themselves with the type of traffic generated by their users and could make decisions based on this knowledge. A typical example was the optimization of video traffic, one of the most bandwidth-demanding types of traffic and one which had the greatest impact on the user experience and the consumption of network resources. This was the main purpose of platforms such as Citrix ByteMobile. However, increased traffic encryption, led by the 'internet giants' (Facebook, Google, YouTube...) has made it impossible for operators to maintain this level of visibility, and therefore video-optimization platforms can no longer be justified due to their high investment and operating costs. At the same time, other optimization mechanisms, such as TCP-centred ones, have become more relevant, these have evolved substantially and are unrelated to encryption traffic.

ALTERNATIVES

- Continuing with ByteMobile which has high OPEX and low effectiveness.
- Giving up on optimization and therefore renouncing the improvement to the user experience and the efficiency of network resources, which, in simple terms, means accepting a loss of competitiveness when compared with other operators.

F5 | LTM - PEM SOLUTION

F5's BIG-IP LTM/TCP optimization offers the most advanced TCP optimization algorithms, making it possible to improve the user experience and to reduce the consumption of network resources. In addition, when combined with the 'subscriber awareness' module (PEM - Policy Enforcement Manager), it enables optimisation according to the status of each network user (for example, applying a different optimisation mechanism to a user on 3G versus one on 4G, adapted to the particular characteristics of each access method).

F5 TCP OPTIMIZATION:

- Improves the user experience by adequately managing the TCP connection to the user and to applications, optimizing each one independently.

F5 TCP PEM (POLICY ENFORCEMENT MANAGER):

- Determines the traffic corresponding to each user (for example, via RADIUS accounting or DHCP) and acts on it in a convenient way.
- Makes it possible to generate value-added services.

REFERENCE ARCHITECTURE | REPLACEMENT OF CITRIX BYTEMOBILE

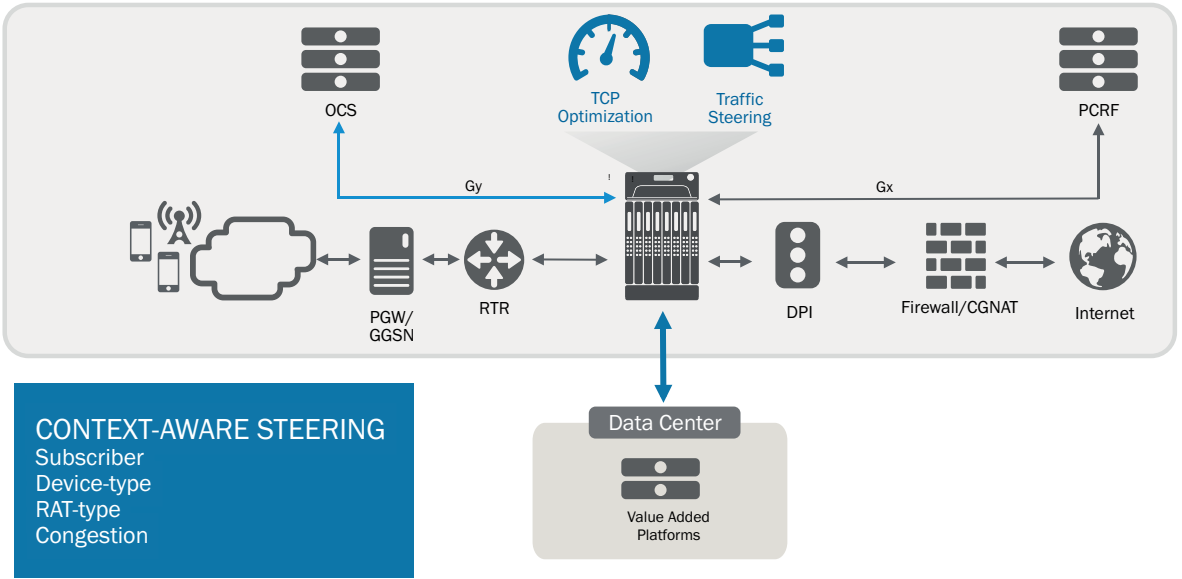

PEM ANALYTICS SOLUTION

PROBLEM

New services allow operators to form a closer relationship with their users and to gain a better understanding of their needs and preferences. Until recently, operators were not able to obtain detailed information on how their customers used the network or what their favourite applications were. This lack of data left them unable to offer customized services and obliged them to offer general and simpler services instead. Increased competition has forced operators to be more imaginative in their approach to their customers. Solutions which provide visibility into the applications used by users provide an opportunity to create innovative services which will satisfy most customers.

Solutions which provide visibility into the applications used by users provide an opportunity to create innovative services which will satisfy most customers.

ALTERNATIVES

- Competing with simple plans, differentiating on price alone.
- Tools which are complex and difficult to integrate

F5 | PEM SOLUTION

BIG-IP PEM (Policy Enforcement Manager) classifies traffic according to the type of application, making it easier to offer personalized services to customers and, consequently to generate new revenue and increase user satisfaction. It has the ability to support pricing on a per-application basis and can manage quotas, which makes it possible to offer plans, which are based on the real needs of customers.

For example, there will be users interested in VoIP packages. You can offer these users a plan, which allows unlimited VoIP for an additional fee. If users want a business package, you can offer a service which enables business applications without impacting data costs.

Analytics makes it possible to provide multiple types of services based on the specific characteristics of users, facilitating the generation of additional income, improving the user experience and encouraging customer loyalty.

REFERENCE ARCHITECTURE | PEM ANALYTICS PROTECTION

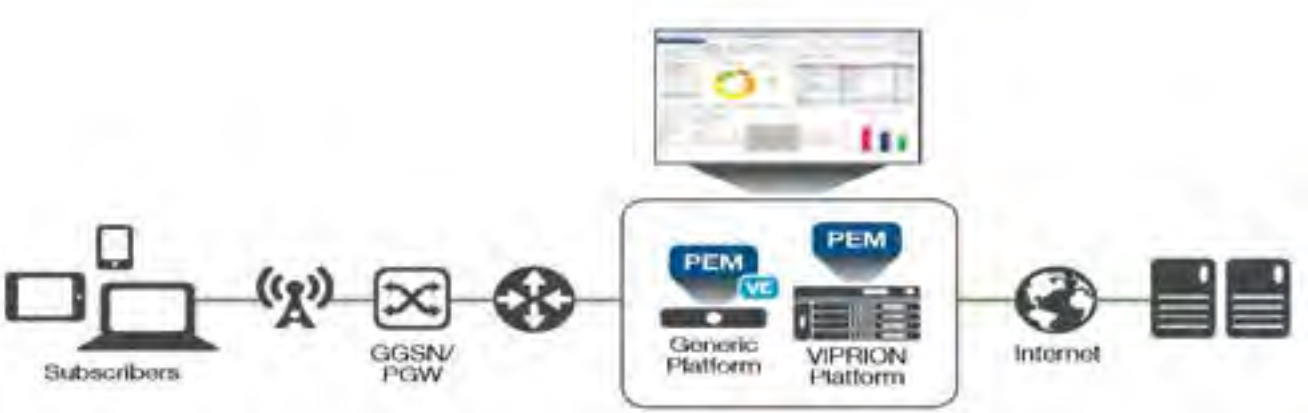

CGNAT SOLUTION

PROBLEM

The global proliferation of mobile devices and the IoT (Internet of Things) has led to the depletion of IPv4 addresses. Although the growth rate of IPv6 traffic is increasing (it is estimated that in 2018 40% of mobile traffic will be on IPv6), it is estimated that the use of IPv4 will continue for quite some time (due to «legacy» services, applications and devices, which do not support IPv6).

The challenge for operators is to support and manage content and devices using IPv4, while simultaneously facilitating the transition to the new devices, services and applications in IPv6.

ALTERNATIVES

- Other NAT solutions, which are less scalable and less flexible, with high costs.
- Giving up IPv6 services, isolating a growing source of traffic and moving away from the needs of users.

F5 | CARRIER-GRADE NAT SOLUTION

F5's BIG-IP Carrier Grade NAT (CGNAT) offers a powerful and scalable solution which enables operators to migrate successfully to IPv6 while continuing to support IPv4, ensuring the interoperability of both worlds. F5 CGNAT provides high-speed- logging capabilities as well as support for IPFIX. This compresses NAT-logging, which significantly reduces the amount of data generated and therefore reduces the total cost of storage and processing. In addition, it supports multiple deployment models and functionalities (such as DS Lite, 6RD, MAP; NAT44, NAT64, 464XLAT, PBA, Deterministic NAT, Hair-pinning, etc.) as well offering extensive support for ALGs (Application Layer Gws), which is essential for applications such as VoIP, SIP services, etc.

REFERENCE ARCHITECTURE | CGNAT SOLUTION

It allows full connectivity and a smooth migration to IPv6

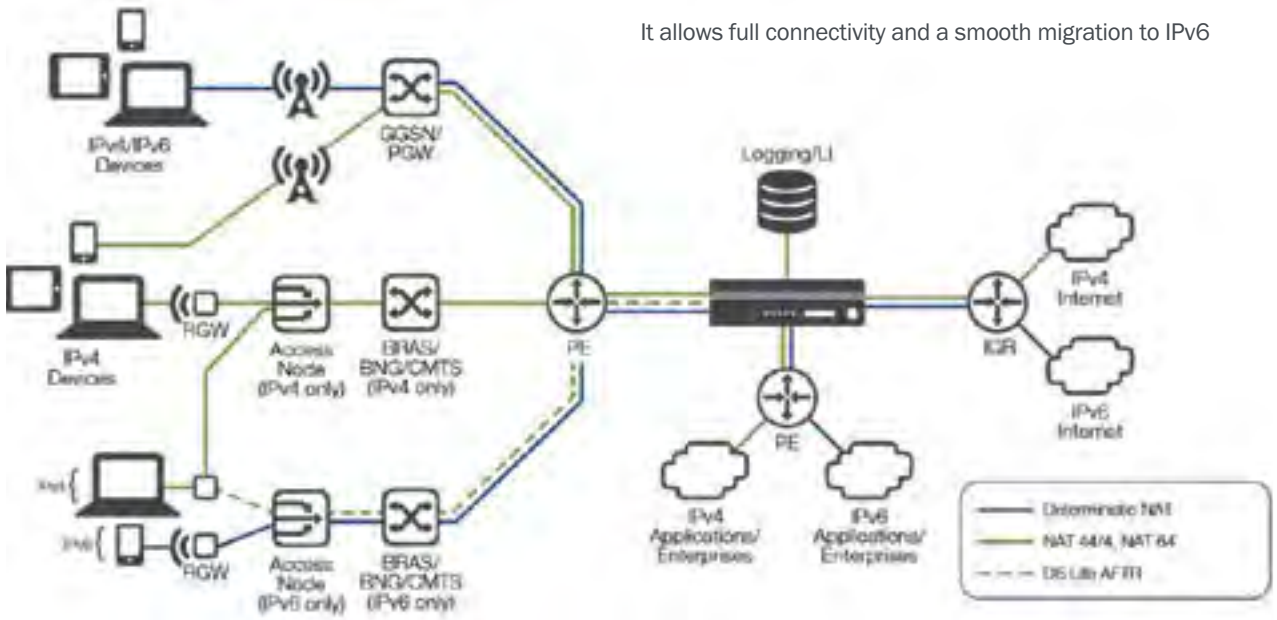

INTERCONNECTED SIGNALLING SOLUTION

PROBLEM

It is essential for operators to maintain the interconnection between 2G/3G networks and 4G LTE networks in general and the SS7 MAP and the Diameter worlds in particular.

Without this communication, operators would need to replace the entirety of their ageing infrastructure with LTE networks.

When new LTE or IMS elements equipped with the Diameter protocol are introduced, we must ensure that this is undertaken in a way which works efficiently with the other elements in the network, no matter how old they may be or from what manufacturer they have come.

ALTERNATIVES

- Replacing the old infrastructure. Expensive, inefficient and impractical in the short term.
- Creating parallel infrastructures. Little flexibility, difficult to operate.

F5 | SDC SOLUTION

F5's SDC (Signalling Delivery Controller) provides interworking functionality through its Diameter gateway to ensure interoperability in multivendor- and multitechnology-environments. The SDC provides connectivity between RADIUS and Diameter, allowing AAA services, based on RADIUS, to communicate in the Diameter world as well as supporting the HLR authentication of Wi-Fi traffic. It also supports connectivity between the SS7 protocol and Diameter. Thus there is an immediate solution to issues caused by the coexistence of legacy infrastructure with the new LTE infrastructure.

REFERENCE ARCHITECTURE | INTERCONNECTED SIGNALLING SOLUTION

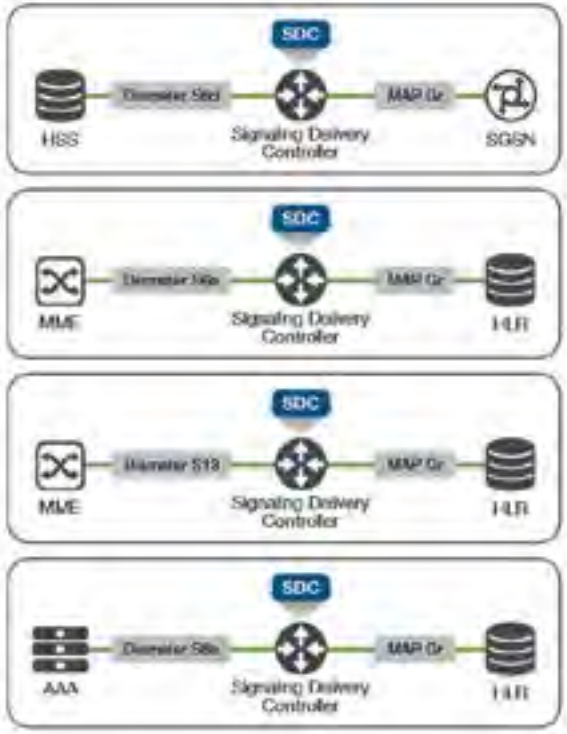

SOLUTION FOR MONETIZING OTT SERVICES

PROBLEM

Companies with OTT (Over-The-Top) services put pressure on telecommunications operators to deliver applications with high bandwidth consumption, which stress the capacity of these networks without generating additional income. Telecom operators can take a more proactive role and develop partnerships with OTTs which allow them to generate new revenue.

ALTERNATIVES

- Generic solutions on common network equipment, which degrade network performance and are not easily integrated.

F5 | LTM + PEM SOLUTION

With BIG-IP PEM (Policy Enforcement Manager) it is possible to detect and classify specific applications and implement appropriate policies on them, such as a higher QoS (Quality of Service) or excluding application-derived traffic, which is above the user's contracted limit. For example, it can detect a video-streaming application and determine whether the user has paid for the premium package which guarantees a certain QoS for that application, while other users have content delivered on a best-effort basis. Another example is that the traffic generated by a user on Facebook, or any other application, is excluded from the contracted level of data consumption. In these scenarios, you can form business alliances with relevant OTTs and earn extra income from them.

REFERENCE ARCHITECTURE | SOLUTION FOR MONETIZING OTT SERVICES

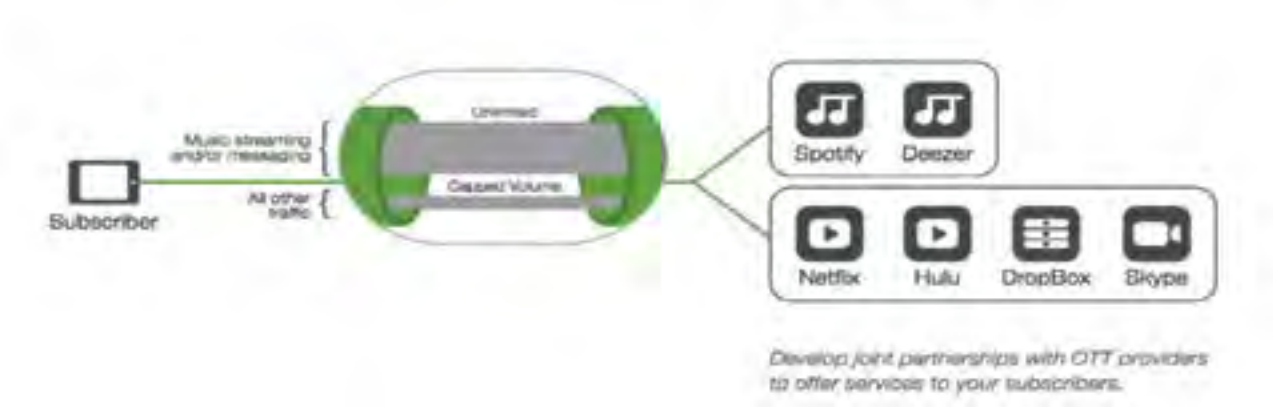

INTELLIGENT DNS FIREWALL SOLUTION

PROBLEM

Companies use the DNS service to provide users with access to web applications. If the DNS service is not available (or performance is compromised), access to these applications cannot be guaranteed.

It is critical to optimize and secure DNS infrastructure to ensure that users can be provided with a service. The operation of this DNS infrastructure requires the ability to respond to a large number of requests per second and the ability to upscale quickly becomes critical when you have to handle thousands of domain names.

It is also necessary to ensure user protection and the integrity of the service against DDoS attacks, DNS cache poisoning and DNS tunnelling.

ALTERNATIVES

- BIND-based solutions are expensive to operate due to the frequent updates required to counter the constant threats.
- Traditional DNS solutions are difficult and costly to scale, there are no flexible and/or advanced integrated security solutions specific to DNS.

F5 | DNS + AFM SOLUTION

BIG-IP DNS + AFM (Advanced Firewall Manager) allows service operators to optimize, secure and monetize their DNS infrastructure. This solution provides LDNS caching services, with carrier-class-level resolution and high performance and is a hyper-scalable Authoritative DNS solution, which includes a DNS firewall services with hardware to mitigate DDoS attacks on the DNS. BIP-IP DNS + AFM provides an intelligent and scalable DNS infrastructure, which makes it possible to provide a rapid response to mobile users. The management of customizable monitors and GSLB services, means that service operators can allocate adequate resources to DNS requests and respond in a way which provides the best user experience. BIG-IP DNS also enables a DNS64 environment for IPv6 environments, with a fault-tolerant infrastructure, optimizing traffic and increasing the quality of service to users, thus protecting the brand and the operator's reputation.

In addition, BIG-IP DNS + AFM protects the DNS infrastructure against malicious attacks generated by infected users and unwanted DNS requests and responses. Intelligent DNS by F5 has a FW which inspects and validates the protocols and discards or refuses to accept unsolicited responses. It also mitigates attacks by blocking access to malicious domains.

Finally, BIG-IP DNS provides statistics and reports, and includes functionality for high speed logging for DNS to facilitate capacity planning, service optimization and the monetization of services.

REFERENCE ARCHITECTURE | INTELLIGENT DNS FIREWALL SOLUTION

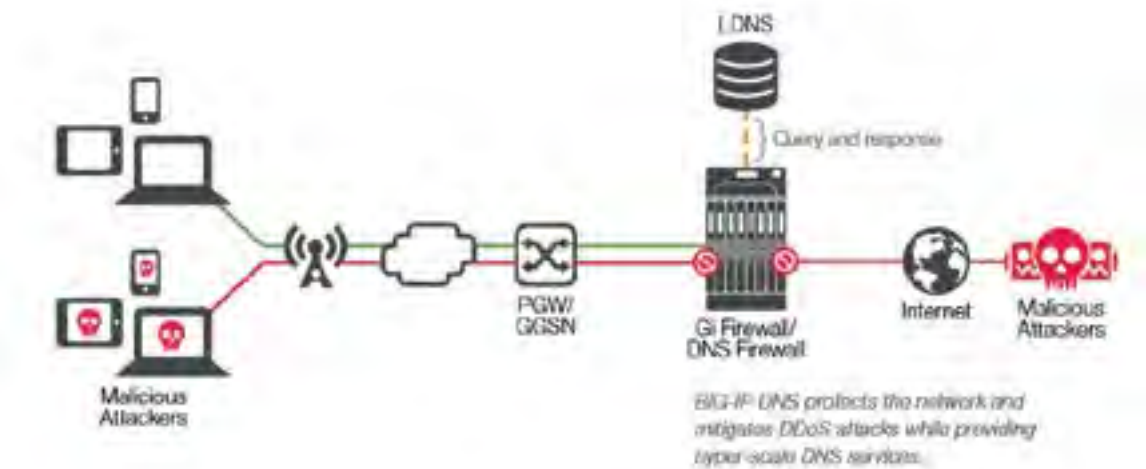

GI FIREWALL SOLUTION

PROBLEM

As mobile operators and other service providers start to base their networks entirely on IP, such as, for example, the 4G LTE networks, intrusions and attacks become more common. Operators have to defend themselves constantly against security threats to ensure the availability of their most valuable resource - the network. This increases costs and operational complexity and has a negative impact on network performance and the user experience.
The progressive introduction of IPv6 addressing in subscriber networks presents the need to protect these users, whose traffic can be routed directly to and from the internet without the necessity of elements such as CGNAT

ALTERNATIVES

The use of traditional firewalls has proved itself to be a poor alternative because of its low level of performance and high operating costs. Its drawbacks include:

- poor performance in mobile environments or where there is a large number of connections.
- poor energy efficiency.
- poor space efficiency.
- high cost per user.

The use of CGNAT platforms to protect IPv4 environments puts the CGNAT platform itself in danger of flooding-type attacks. Furthermore, in IPv6 environments, user traffic does not require NAT, which makes it necessary to have an element specifically for protection – the Gi Firewall.

F5 | SOLUTION WITH AFM

BIG-IP Advanced Firewall Manager (AFM) defends the mobile infrastructure and mobile users from attacks, regardless of their origin. Its capabilities include the mitigation of large-scale DDoS attacks, such as network floods, port scans and sweeps, or connection floods. By detecting and stopping these attacks, BIG-IP AFM can prevent the congestion and overloading of the network-control panel and of the wireless access. The level of protection against DDoS attack vectors increases with each new version of the F5 operating system (TMOS). In many BIG-IP platforms, the functions related to protection against DDoS attacks are accelerated by the use of specialised hardware.

BIG-IP AFM, a firewall which is certified by ICSA Labs, offers the protection of a full-proxy firewall, terminating and completely inspecting client connections in the face of threats. This ensures the availability of the network and a better user experience.

REFERENCE ARCHITECTURE | GI FIREWALL SOLUTION

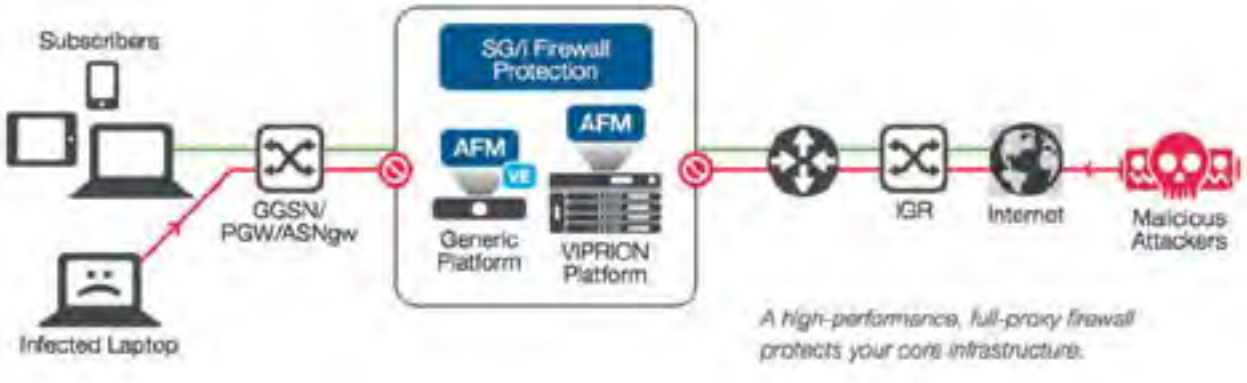

LTE-ROAMING SOLUTION

PROBLEM

Communications service providers (CSPs) face a major change in the way voice services and traditional data services are delivered to users. The globalization of communications requires that services are available regardless of where in the world the users happen to be.

The deployment of LTE networks means that operators have to provide roaming services to their users, including providing connectivity to mobile users moving between LTE and 2G/3G networks. This complex process entails managing the specific requirements of routing, scalability and security, while ensuring quality of service (QoS).

ALTERNATIVES

- Multiple solutions from different manufacturers with high OPEX.
- Solutions with less flexibility, which are difficult to adapt to changing needs.

F5 | SDC SOLUTION

The SDC (Signalling Delivery Controller) solution ensures that the technical needs of roaming are met, while facilitating additional income. To solve the complexity of maintaining the connection between users moving between different networks (2G/3G and 4G), the SDC also functions simultaneously as a IWF (thus facilitating interworking), Diameter Gateway and DEA (Diameter Edge Agent), providing connectivity, routing, translation capability and security. The SDC makes it possible to connect different mobile networks, wholesale roaming providers and IPX (IP eXchange) operators.

F5 SDC facilitates:

- A reduction in the time required to activate the connection with a roaming partner.
- The implementation of new services.
- Network growth and increased revenues

REFERENCE ARCHITECTURE | LTE-ROAMING SOLUTION

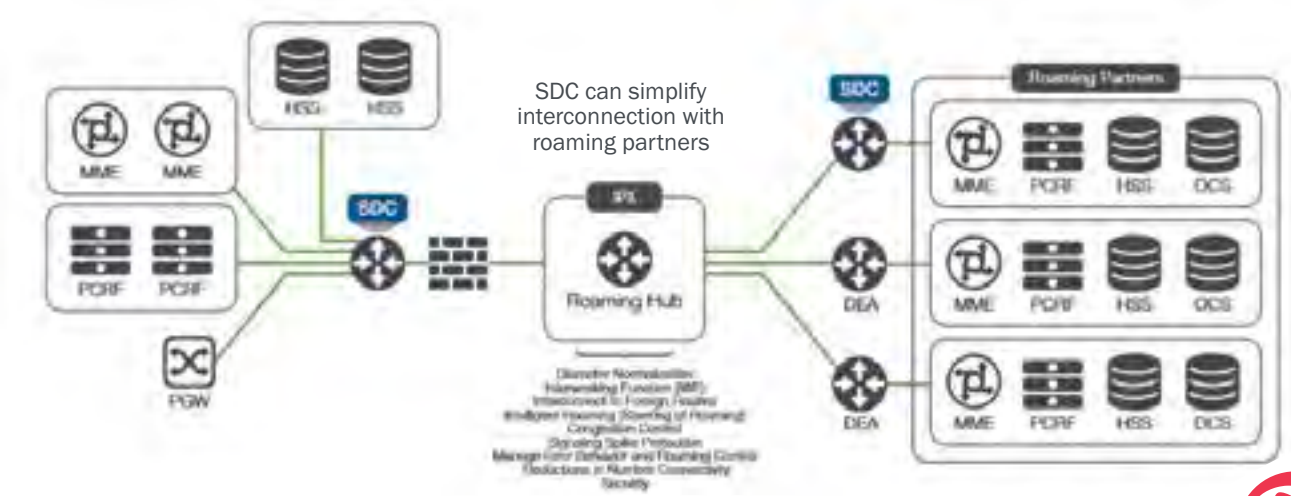

SIP/IMS SIGNALLING SOLUTIONS

PROBLEM

The adoption of LTE is making it possible for operators to provide high-speed services, multimedia communications (RCS) and VoLTE with the aim of improving their portfolio of solutions and services and increasing their ARPU. To achieve this, they have implemented IMS architectures with SIP as the main signalling protocol, but migrating completely to IP-based networks has its challenges, including security. The nature of the protocols is based on open standards, making them vulnerable to attacks on the networks and services, including DDoS-type attacks, stealth floods, etc. and malformations of the SIP protocol. Furthermore, this migration to full-IP networks also represents a potential problem with regard to managing capacity and performance, since customer usage is continually growing and we need to ensure that all services are continually available.

ALTERNATIVES

- Multiple solutions from different manufacturers with high OPEX.
- Solutions with less flexibility and scalability, which are difficult to adapt to changing needs, making it difficult to plan ahead.

F5 | SDC SOLUTION

The SDC (Signalling Delivery Controller) solution provides SIP solutions within a high-availability and scalable system for IMS network infrastructure, it includes devices such as the X-CSCF and SBCs (Session Border Controllers). The SDC manages SIP traffic and ensures service availability through the continuous monitoring of SIP servers and applications. It also handles the SIP standards, ensuring that there are no interoperability problems between different IMS services. It also increases security by detecting and discarding malformed or failed SIP communications.

The F5 SDC solution:

- Ensures the interoperability of SIP requests and responses. Scales up to millions of simultaneous calls.
- Improves the robustness of the solution to carrier-class level, including the synchronization of sessions and total capacities against failure, withour any loss of connections.

REFERENCE ARCHITECTURE | SIP/IMS SIGNALLING SOLUTIONS

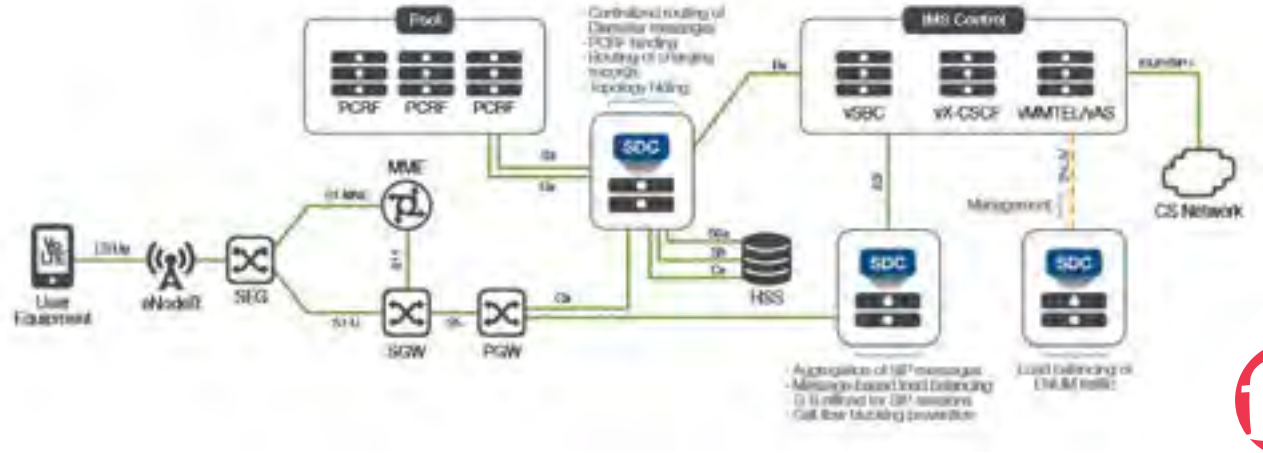

ON DEMAND BANDWIDTH AVAILABILITY SOLUTION

PROBLEM

The increasing demand for bandwidth forces operators to increase the capacity of their networks to accommodate this demand. However, the peaks in traffic which stress the capacity of the transmission network only happen at certain times of the day, while at other times there is ample availability of bandwidth which means that resources lie idle and therefore there is a loss of efficiency and there is the definite potential for operators to suffer economic losses.

ALTERNATIVES

- Oversize networks.
- More restricted service portfolios, i.e. less competition, lower incomes.

F5 | SOLUTION PEM

Thanks to F5 PEM (Policy Enforcement Manager), operators can differentiate their offerings, adapt them in real time to the bandwidth requirements of their customers and manage the associated fees. A user may, for example, only need an increase in bandwidth for a specified period of the day. This can be used by the operator to charge the user an additional fee for extra bandwidth during that period, which means additional income. It is also possible to monitor the performance of the network and offer attractive rates to encourage higher consumption during times of low utilisation.

F5 PEM provides:

- The opportunity to optimize operator resources. Better user experience.
- Additional income for operators.

REFERENCE ARCHITECTURE | ON DEMAND BANDWIDTH AVAILABILITY SOLUTION

Bandwidth increase based on the time of day

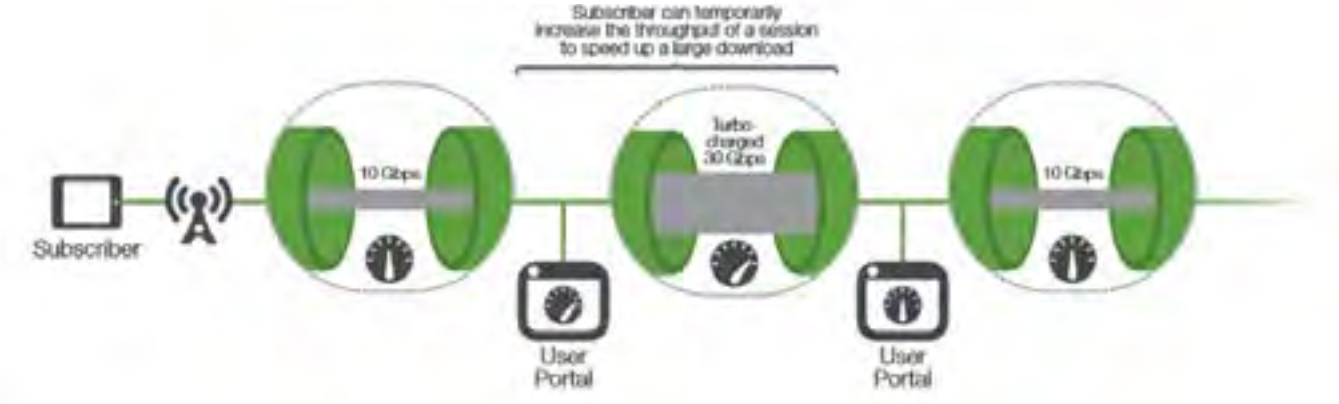

MANAGED SERVICES

CONTENT

79 MSP offering a DDos as a service solution
80 MSP offering a DNS as a service solution
81 MSP offering an enhanced WAF as a service solution
82 MSP offering a WAF as a service solution
83 MSP offering a Federation as a service solution
84 MSP offering a LB as a service solution

MSP OFFERING A DDOS AS A SERVICE SOLUTION

PROBLEM

DDoS attacks are a growing problem for organizations. The cost of generating a DDoS attack is minimal in comparison with the potential cost of loss of service and possible methods of remediation and recovery of service. This is the reason for the emergence of new «hacktivist» groups, which carry out attacks on organizations under any pretext they consider to be legitimate. Furthermore, the types of attacks are becoming increasingly complex, from attacks on the network layer to attacks on the application layer, the DNS tier and even exploiting vulnerability in the business logic of the application.
In MSP environments this problem is even greater because an attack on a single client can compromise access for many customers.

ALTERNATIVES

- Deployment of on-premises DDoS-prevention solutions which offer no other added value.
- Inability to manage volumetric attacks, delegating the service to the line operator, with consequent cost overruns and tie-in to the operator.
- Inability to prevent attacks on the application layer or the encryption layer, with consequent loss of service.

F5 | AFM + ASM + SILVERLINE SOLUTION

The BIG-IP platform can prevent DDoS attacks at multiple levels. Attacks on the network layer (ICMP Floods, SYN Floods, UDP Floods, DNS Floods, etc.) are directly managed by BIG-IP's hardware accelerated DDoS offload technology (AFM (Advanced Firewall Manager) and DNS/GTM licenses required). Attacks on the application layer can be managed by ASM (Application Security Manager), which detects malicious activity and performs an access-control check on the service to verify that the request is being made by a human, (by performing browser finger printing or injecting a CAPTCHA or invisible JavaScript challenge, etc.).

Volumetric attacks are mitigated by the hosted SilverLine platform which receives malicious traffic and undertakes fully-automatic filtering in the cloud via F5's expert service which is available 24/7. All of this is managed through a user portal which provides a real-time source of information on the countermeasures and actions taken.

REFERENCE ARCHITECTURE | MSP OFFERING A DDOS AS A SERVICE SOLUTION

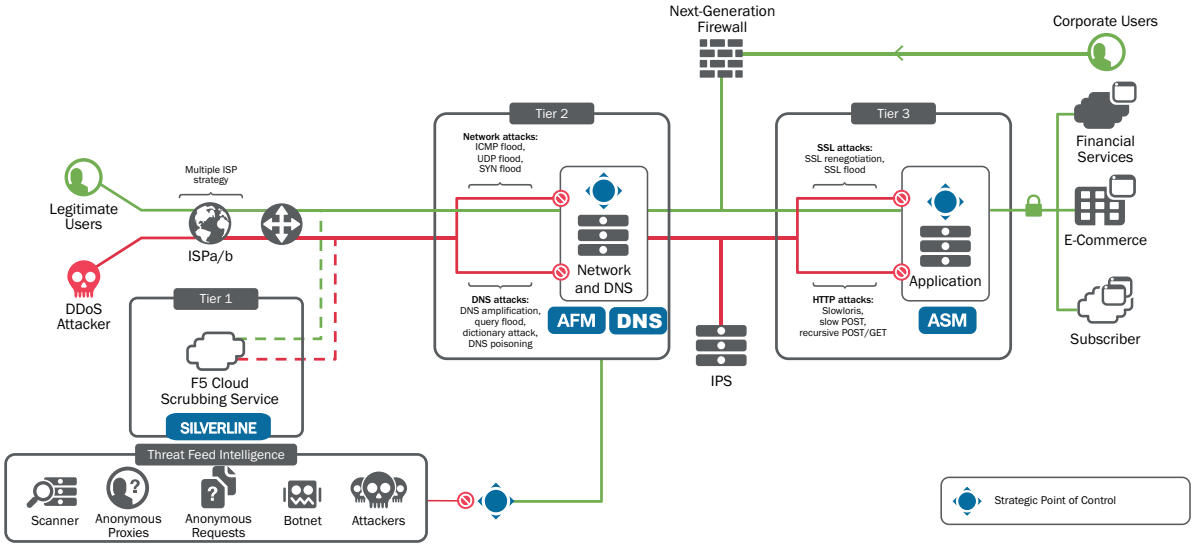

MSP OFFERING A DNS AS A SERVICE SOLUTION

PROBLEM

There is an ever-increasing number of attacks on DNS infrastructure - both specific attacks on the DNS protocol, such as DDoS attacks, and attacks on the services themselves - which have a global impact on applications. MSPs need to manage DNS services for their customers, to ensure their security and availability. This can mean balancing the DNS service and/or providing a fully-managed DNS service and/or provide intelligent balancing services according to availability and/or the migration and management of traffic-bursting for services.

This should be governed by mechanisms which make it possible to reduce costs, implement and monitor SLAs, undertake customer-level customization and use the capabilities of orchestration to fit this service into the provision and operation of services.

ALTERNATIVES

- Multiple layers for the service delivery.
- GNU/BIND solutions which need continuous patching and have exposed vulnerabilities.
- Very limited scalability based on balancing.
- Inability to generate value-added services or correlate service-availability information on the DNS service.

F5 | DNS SOLUTION

F5's BIG-IP platform makes it possible to deploy multi-tenant DNS services on demand, thus offering unique capabilities in security and providing intelligent name resolution based on multiple parameters such as availability, location, continuity of service or migration of services to third parties or the cloud.

Open REST APIs permit service integration with other manufacturers enabling a natural integration into SDN/SDDC environments such as Cisco APIC, VMWare NSX or OpenStack.

The consolidation of functions enabled by F5 also reduces operating costs by consolidating the functions of DNS, Firewall, VPNaaS, and/or LBaaS, with a drastic improvement in performance over any other competitor.

REFERENCE ARCHITECTURE | MSP OFFERING A DNS AS A SERVICE SOLUTION

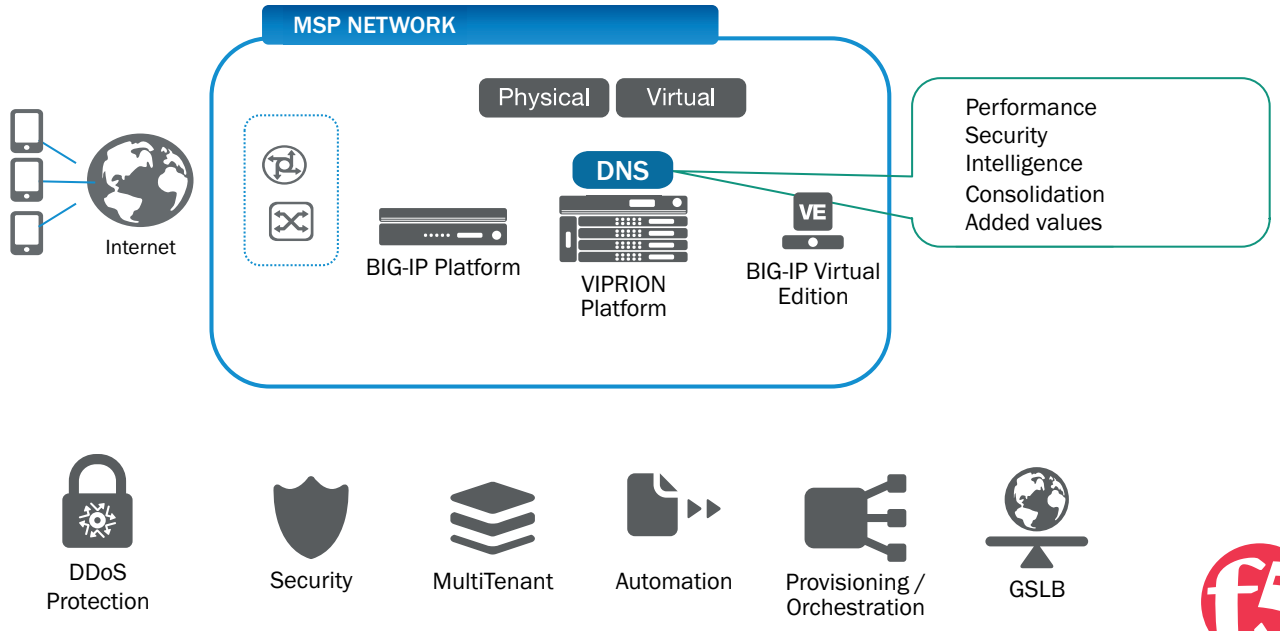

MSP OFFERING AN ENHANCED WAF AS A SERVICE SOLUTION

PROBLEM

While a WAF solution protects our applications from malicious use through potential DDoS or application layer attacks, it cannot protect users of our applications against phishing attacks, RATs, man-in-the-Browsers, keyloggers...

This protection is particularly complex, since we have no access to our customers' computers (fixed and/or mobile) and therefore cannot install any clients on these devices, which also use a multitude of different operating systems and browsers.

Credential theft, phishing attacks and financial fraud greatly harms the reputation of organizations and it can lead to significant economic losses with legal implications.

The deployment of E-WAFaaS services is a very interesting business opportunity for MSPs, offering more value to the customer and generating additional revenue and customer loyalty to the MSP's services.

ALTERNATIVES

- Solutions from other vendors require the installation of clients on the server and/or on the client machine, which is not feasible in most cases due to the interaction between the client and the application software itself (in the case of servers) or due to the use of devices, which are outside our organizational remit (for customers).
- Trusting only in the WAF module does not protect the users of our applications.

F5 | ASM + WEBSAFE SOLUTION

The F5 ASM (Application Security Manager) module makes it possible to deploy a WAF quickly and easily in a way which is transparent to applications and thereby to protect them against web application attacks such as Cross-Site-Scripting (XSS) and SQL-injection (two of the most common methods of data breaches). It also includes the ability to produce reports for compliance with PCI-DSS regulations.

The combination of F5's ASM and WebSafe modules protects both enterprise applications and those who use them without the need to install any clients, modify applications, or access devices, either fixed or mobile.

REFERENCE ARCHITECTURE | MSP OFFERING AN ENHANCED WAF AS A SERVICE SOLUTION

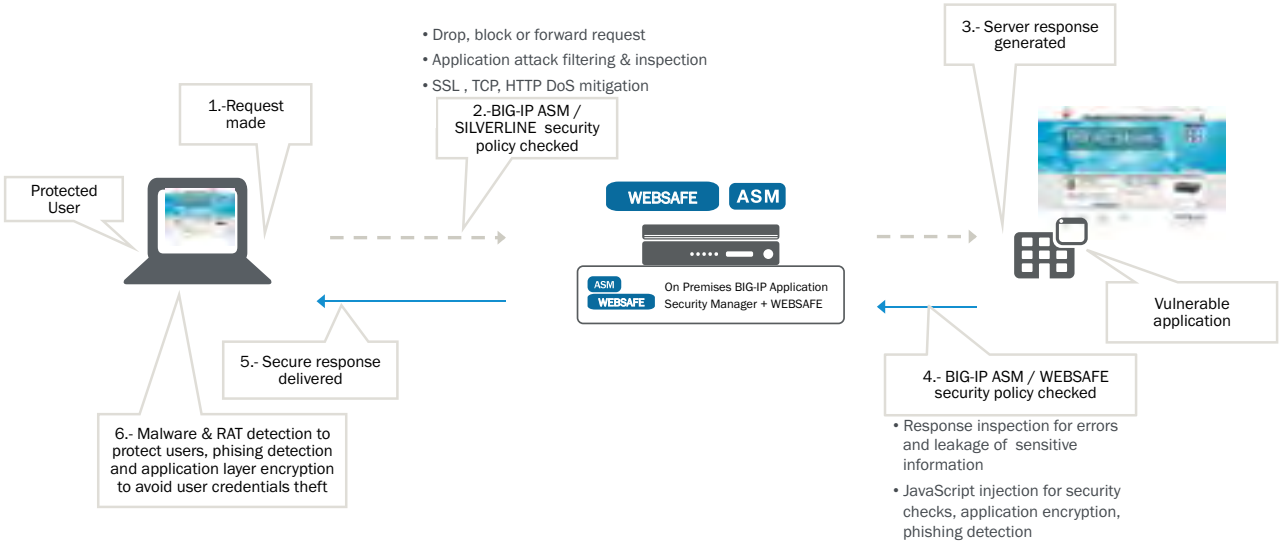

MSP OFFERING A WAF AS A SERVICE SOLUTION

PROBLEM

The security of a company often depends on the security of web applications, which have been developed by third parties and which are vulnerable to known forms of attack (SQL-injections, cross-site scripting…). These attacks may result in valuable information being leaked from the company, with major economic and even legal consequences for the organization. Similarly, web (L7)-specific DDoS attacks result in a loss of service for our customers. In addition, PCI-DSS regulations mandate the use of web-protection devices for the obfuscation of sensitive content.

The deployment of E-WAFaaS services is a very interesting business opportunity for MSPs, offering more value to the customer and generating additional revenue and customer loyalty to the MSP's services.

ALTERNATIVES

- Implementing more secure application-development policies, which may result in organizational conflict, long delays in the development of applications and a substantial increase in development costs.
- Taking no action poses a huge risk and means non-compliance with regulations.

F5 | ASM/SILVERLINE SOLUTION

F5's ASM (Application Security Manager) module makes it possible to deploy a WAF quickly and easily, in a way which is transparent to applications and thereby to protect against web-application attacks such as OWASP and DDoS. It also includes a module for monitoring compliance with PCI-DSS regulations.

After an initial learning period in which the ASM comes to recognize the web application's normal patterns of operation, the WAF protection is activated, this is specific to each URL (FQDN) and the user can specify the threshold at which the automatic protection is triggered.

Similarly, the equipment can be configured to work with whitelists or with blacklists.

Another option for the protection of applications, is the use of the SILVERLINE service, this is a managed service provided remotely on a 24/7 basis and supported by experts from F5.

REFERENCE ARCHITECTURE | MSP OFFERING A WAF AS A SERVICE SOLUTION

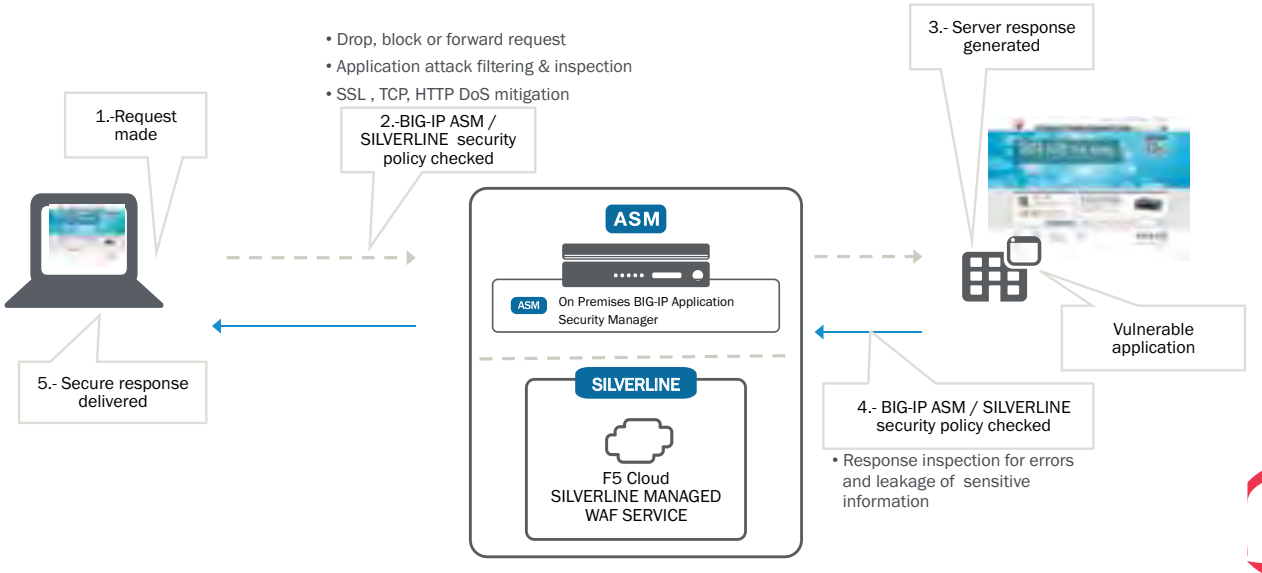

MSP OFFERING A FEDERATION AS A SERVICE SOLUTION

PROBLEM

Outsourcing services to SaaS (Software as a Service) platforms, such as Office 365, SalesForce, Concur, Workday, Google Apps, etc. is becoming more common by the day.

Integrating all these services requires access control and control of the identity of the users who consume those services. Allowing the user to manage their access and their password to the service, will generate multiple problems from the point of view of security (repeat passwords, change management, different policies, etc.). Synchronizing corporate directory services means giving all service providers visibility of internal user credentials, which is neither secure nor appropriately scalable and therefore is not an option.

ALTERNATIVES

- Exposing the internal directory to external service providers.
- Business solutions which are not compatible between different vendors (e.g. ADFS) and which lack advanced authentication capabilities (such as two-factor authentication).
- Allowing users to configure and access external services via specific passwords for each service, each of which has its own policies on password strength and renewal.

F5 | APM SOLUTION

The BIG-IP Access Policy Manager (APM) platform makes it possible to create a federation of identities between different sites. Thus, there are two roles in the provision of a specific service, the role of the «Service Provider» (SP) and the role of the «Identity Provider» (IdP). The federation between the two permits the identification of users without the need to know their credentials.

BIG-IP APM can perform both roles: it can control access to an external service (in the case of an MSP serving their customers) and/or it can act as a local IdP, creating an advanced system of authentication for SaaS services, independent of the chosen SaaS platform. In this way it is possible to avoid exposing the local directory to the MSP / SaaS provider, while also offering SSO functionality.

REFERENCE ARCHITECTURE | MSP OFFERING A FEDERATION AS A SERVICE SOLUTION

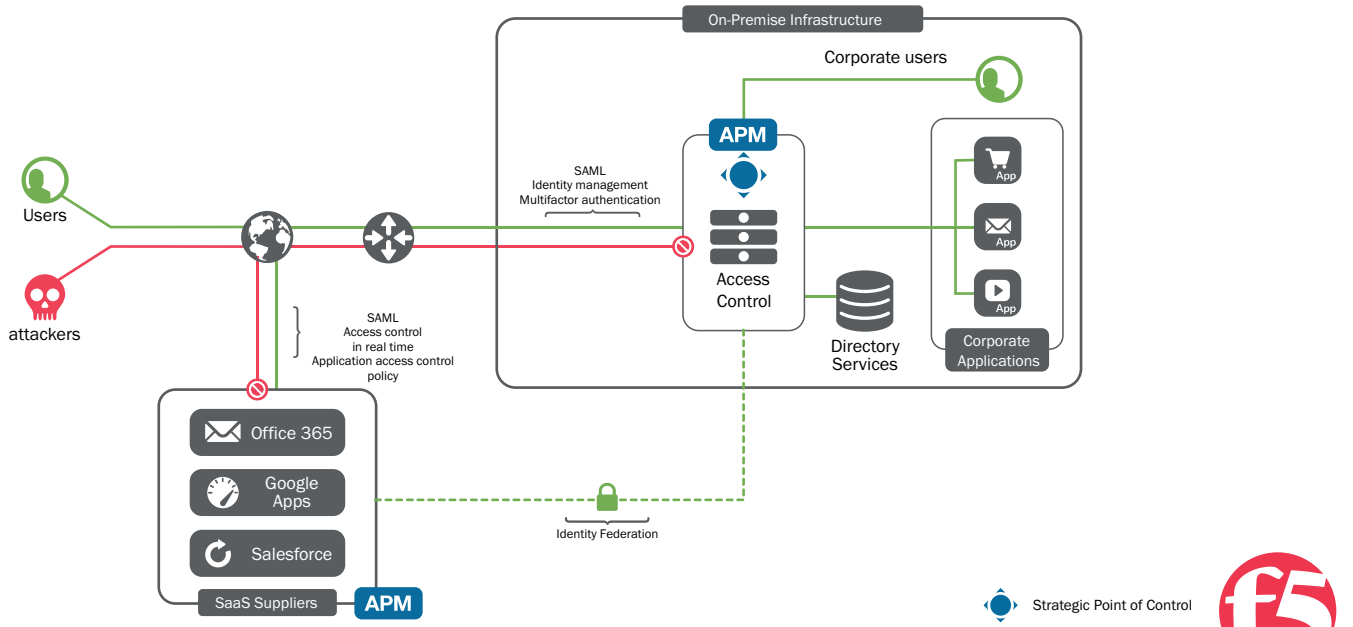

MSP OFFERING A LB AS A SERVICE SOLUTION

PROBLEM

MSPs need to provide value-added services to their customers and to provide availability and scalability to solutions and to the deployed infrastructure.
This should be governed by mechanisms which make it possible to reduce costs, implement and monitor SLAs, undertake customer-level customization and use the capabilities of orchestration to fit this service into the provision and operation of services.

ALTERNATIVES

- Ad-hoc solutions, created through in-house development or based on GNU, lack the capacity and power of the BIG-IP platform both in terms of functionality and in terms of performance and have multiple hidden costs with regard to supporting and upgrading these platforms.
- Other commercial solutions lack F5's flexibility in terms of integration, licensing models and the multi-tenancy capabilities which are native to the BIG-IP platform

F5 | | LTM + BIG-IQ SOLUTION

F5's BIG-IP platform makes it possible to deploy load balancing services on demand, with intelligence at L7, using a single platform, with multi-tenancy capability.

Open REST APIs permit service integration with other manufacturers, enabling a natural integration into SDN/ SDDC environments such as Cisco APIC, VMWare NSX or OpenStack. The deployment of services through iApp templates make it possible to create an application catalogue to facilitate the provision of services to multiple different customers thereby reducing both operating costs and the possibility of human error.

The licensing service is versatile, in addition to traditional models, there are alternatives such as license pooling or subscription licences.

BIG-IQ further enhances the centralised management of multiple F5 platforms.

REFERENCE ARCHITECTURE | MSP OFFERING A LB AS A SERVICE SOLUTION

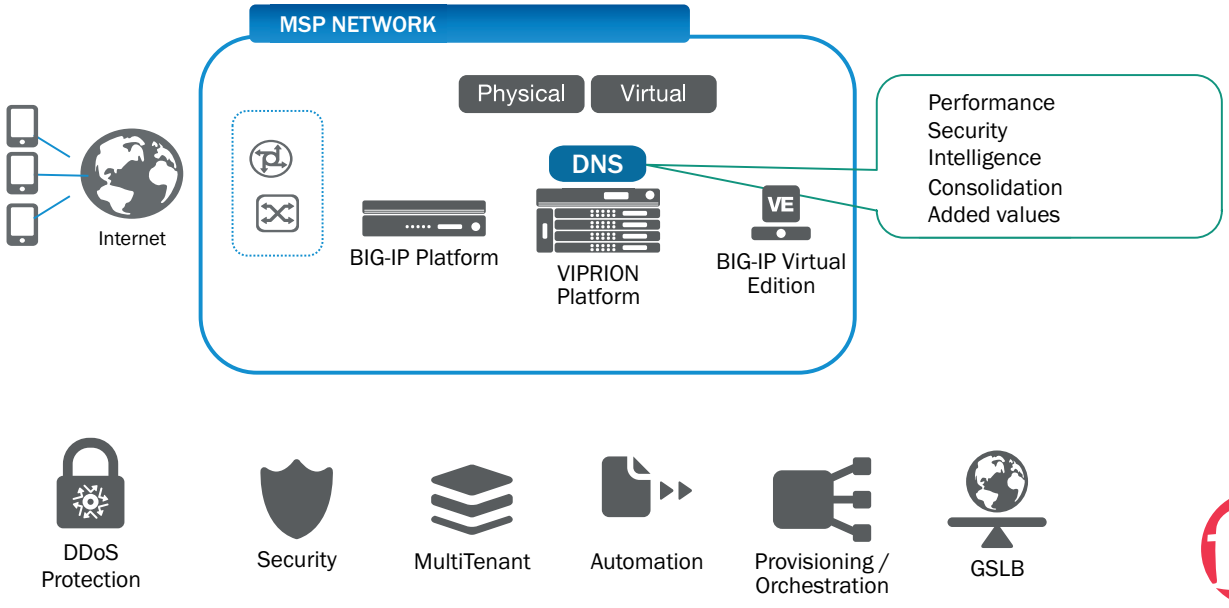

SOLUTIONS PLAYBOOK

September 2016 - March 2017