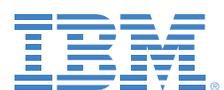# Network Services and the Cloud

To fully maximize the potential of virtualization and the cloud, organizations must take a more holistic approach to their Application Delivery Network. The F5 dynamic services model provides a foundation of reusable services that can understand the context of each transaction regardless of user, device, or location. The IBM Services portfolio leverages F5 products to align business requirements with technology.

by Craig Scarborough
Principal Solutions Engineer

Contents

Introduction

The network plays a preeminent role in cloud infrastructures. Cloud computing uses the network to gain on-demand access to computing resources, and the network becomes the conduit for enormous computing capability. This critical role of the network in cloud computing demands that you "get the network right"—that you configure it to achieve the right levels of performance, security, availability, responsiveness, and manageability.

IBM has a proven history of designing, building, and supporting complex infrastructures as well as emerging technologies like today's cloud configurations. IBM can provide the insights and capabilities you need to assess, plan, design, implement, and manage the network infrastructure to support your client's chosen cloud computing environment with on-demand IT.

F5 Networks®, a key IBM partner in this application delivery space, can provide the tools an organization needs to ensure the successful rollout of public, private, or hybrid clouds. F5's product portfolio enables the delivery of applications efficiently, securely, and on-demand, so the enterprise can reduce costs and gain the efficiency and agility of a cloud-based infrastructure.

Application Delivery in the Cloud

The primary focus of (and litmus test for) any IT infrastructure is application delivery and performance. Application delivery, whether coming from the cloud or a traditional data center, requires the network to make quick, intelligent decisions to ensure consistent availability, secure access, and optimized application response time.

The variable nature of a cloud computing environment requires that its supporting infrastructure be dynamic, programmable, and automated. Unlike a traditional data center, cloud computing environments come with application delivery challenges that can vary depending on the implementation. Cloud computing requires a dynamic control plane that enables you to integrate the various network and application delivery network components—switches, IP address management, load balancing, application security, acceleration solutions, and more—both inside the data center and in the cloud. This dynamic control plane must be able to:

- Intercept application and data traffic.
- Interpret the context of the traffic.

- Instruct the cloud on how to handle traffic appropriately while still maintaining its availability, security, and performance.

The Dynamic Services Model

As organizations progress from enterprise virtualization to cloud maturity, one thing remains constant: the challenges of a traditional data center are amplified. The increased level of complexity, coupled with reduced control over applications and data that leave the corporate data center, brings many organizations to a point of diminished returns long before they achieve the goal of on-demand IT.

Businesses are looking for a way to leverage the flexibility of cloud infrastructure while containing its inherent complexity. Organizations want freedom from the bounds of their own data centers, but need to maintain the same levels of control over security, optimization, and management in the cloud that they can with static, dedicated servers. To accomplish this, businesses need new types of control points in the data center that can dynamically and intelligently adapt to change.

A dynamic services model includes reusable services that understand context and can provide control of information flows regardless of application, virtualization, user, device, platform, or location. It creates a holistic ecosystem that elicits more value from each existing system component (application delivery, security, optimization, virtualization, and management) as it integrates with that ecosystem.

One key characteristic of the dynamic services model is that you can deploy strategic points of control to create a coordinated mediation layer between users, applications, and the data they access. The dynamic services model is the final abstraction between the users and the applications. This layer presents a consistent interface and set of services that can be used to access all applications and data regardless of their current location. It has an intelligent understanding of context—who is accessing what from where and why—that it uses to determine the optimal connection between users, applications, and data. Finally, it uses a deep understanding of context to inform the underlying application and data elements about the current resource requirements and to instruct the infrastructure to adapt in real time to ensure optimized application and data access.

The dynamic services model describes the ultimate strategic point of control—the combination of a stop sign, metered stoplight, express lane, and high-speed bypass. This strategic point of control can dynamically change its nature and the nature of

the surrounding infrastructure based on the unique needs of each individual:
it can instantly turn a two-lane road into a superhighway.

The F5 Dynamic Control Plane

F5® refers to its implementation of the dynamic services model as the dynamic
control plane architecture.

The dynamic control plane architecture is the extension of F5's entire suite of
application and data delivery solutions; it integrates and coordinates with the
virtualization, management, security, and application components often already
present in an organization's infrastructure.

Flexibility in the Cloud with F5 Application Delivery Controllers

Many organizations have deployed application delivery solutions for their business-
critical applications. These deployments have enabled high availability and offloaded
security and optimization, and have promoted simplified application management
and control. However, when an organization wants to move those applications—
whether temporarily to handle outages or demand spikes or permanently—to
external hosting providers and/or cloud providers, how does it proceed if the service
provider does not offer the same application delivery solution? Organizations must
either find a way to translate the existing application delivery configurations to
the provider's solutions—and hope that they provide the same functionality and
performance—or they must purchase and pre-position their own Application
Delivery Controllers in the provider network. This significantly reduces the flexibility
of the solution and increases the cost of moving from one provider to another.
It is easier to re-route traffic if you do not have to build the detour from scratch
every time.

F5 customers can deploy application delivery solutions on a choice of platforms,
including a virtual platform that has the same functionality and configuration as
its hardware-based counterpart. This gives F5 customers the flexibility and ease of
control that is required for a dynamic services model. Organizations can deploy F5
Application Delivery Controllers in dedicated data centers and, at a moment's notice,
deploy a virtual version with the exact same functionality for security, optimization,
and availability in any service provider's network. The organization's unique policies

and configurations can be migrated with the application as it moves to the new service provider. With parity of functionality, organizations can guarantee an identical user experience while significantly reducing costs associated with changing cloud providers. And because all of the application delivery functions are integrated, it's easy to automatically provision new virtual devices.

Long-Distance Migration of Live Applications

Moving the applications themselves is also a challenge. While hypervisor vendors like VMware provide tools that enable organizations to move applications from one piece of physical hardware to another without affecting users, these tools have limited capability to move applications between data centers—or to external service providers—even with dedicated connections. Ideally, organizations want the flexibility to move applications at any time, to any location without affecting users. However, WAN or Internet connections are often too slow to accomplish this feat. Even if the tools themselves do not fail because of excessive delay, the time it takes to move the application makes it nearly impossible to continue servicing user requests. This is particularly true if the application has a back-end disk image that must also be migrated: while the application might move to the new location, the delay caused by having to access its data store in a remote data center would bring the user experience to a stand-still. Lastly, moving the application does not mean that users' requests automatically transfer to the new location. Traffic redirection is required to avoid lost connections and poor performance.

F5's dynamic control plane integrates multiple F5 technologies, as well as other vendors', to enable real-time dynamic provisioning of services. If the orchestration engine recognizes that compute resources are below set thresholds or using excessive resources, it can dynamically move the application to hardware with additional resources and notify the F5 solution of the new or added instances. An organization may decide to do a live application migration between data centers or to a cloud service provider because an application instance lacks sufficient resources to run locally.

In long-distance live migration, the dynamic control plane applies the appropriate services to ensure consistent delivery. F5 devices dynamically create secure, optimized tunnels between devices. These tunnels serve multiple functions. They can significantly reduce the time required to move a machine and its disk image between data centers, so you can use the same migration tools whether the target is in the next rack or around the world. The tunnels can also reroute

existing user connections to the new application location with little to no effect on user performance.

This process is possible only because the F5 dynamic control plane applies services intelligently, and is inherently open to collaboration and integration with the greater ecosystem.

Accurate, Real-Time User Performance Measurement

Users' increased mobility and the numerous connections they use to access corporate applications and data make it difficult to accurately monitor the true user experience. While organizations like Gomez provide sophisticated solutions to manipulate the data stream for real-time performance measurement, not every application is customizable enough to allow the addition of these measurement components, even when an organization has full control of its applications and data. When applications are outside of the corporate data center, this inability to track performance is especially problematic. If you cannot measure the real user experience, it is difficult if not impossible to ensure that co location and cloud peering-point facilities are maintaining the SLAs they promise. Ideally, all applications, regardless of type or location, should be monitored in real time so organizations can use the resulting knowledge to optimize their delivery dynamically.

As a strategic point of control point between applications and users, an F5 Application Delivery Controller offers a universal tagging point to automatically inject tracking components into all traffic to all applications—not just those than can be customized. And, since this component is available in physical and virtual instances, it can provide this function beyond the confines of the corporate data center.

IBM Integrated Communications Services

IBM networking services professionals can help organizations realize the full benefits of a cloud computing environment with:

- IBM solutions that integrate with F5 technologies in the dynamic control plane to create a holistic cloud deployment model.

- Cloud computing insights based on decades of experience with virtualization, extensive research and development activities, and active involvement with industry standards organizations.

- Experience gained from a range of IT environments and by operating more than a dozen IBM cloud computing centers.

- A global pool of skilled networking professionals with in-depth technical expertise and a deep understanding of the evolving demands of network, server, storage, and desktop virtualization.

- A vendor-independent approach using best-of-breed technologies like F5 Application Delivery Controllers.

- Proven methods and approaches to the optimization of networking infrastructures.

Networking Strategy and Optimization Services

Network Application Optimization

Network Application Optimization is a set of services that helps organizations understand how their business-critical applications are affected by the network infrastructure. It can also help them predict the performance of a new or changed application; optimize end-to-end application performance over their network; determine alternative solutions using WAN and/or application acceleration technologies; and assist with selection of applications for migration to cloud computing.

For cloud computing, the Network Application Optimization services leverage an issue-based consulting methodology and empirical analysis to help clients make informed business decisions. By prioritizing applications for migration to cloud computing based on IT infrastructure complexity and risk, these services help clients determine if the application workload characteristics are suitable for cloud computing environments. They result in the delivery of a scorecard, recommendations on optimization opportunities, and a roadmap for migration.

Network Infrastructure Optimization

Network Infrastructure Optimization is a set of services that focuses on helping clients who have selected a cloud computing delivery model and identified the applications that will be delivered using this model, to determine how ready their networking infrastructure is to support private or hybrid cloud computing.

This network readiness assessment will help establish what updates are required for the design, infrastructure, and operation of the organization's network as it relates to the applications to be delivered, focusing on:

- Financial aspects
- Performance and capacity
- Security
- Flexibility
- Availability
- Management
- Organization

These services result in a report on the readiness of the client's networking environment to support the anticipated delivery of its applications via cloud computing, based on its business objectives and cloud computing readiness criteria.

Network Integration Services

Application and Network Performance Optimization

Application and Network Performance Optimization (ANPO) is a set of design and implementation services that focuses on specific applications to provide industry-leading solutions that can lower your infrastructure bandwidth and server costs while significantly improving application performance across even large geographical areas.

Web-based instances are often limited in scale, performance, and reliability by the IT infrastructure they run on. Improved reliability and performance translates directly to productivity gains and cost reductions.

IBM addresses these challenges at the infrastructure's core—where your network, storage, servers, and security come together—to create a tightly linked, application-

fluent infrastructure that's specifically optimized to run your applications. IBM network consultants address real-time business demands and help optimize infrastructure with key technology partners like F5.

Conclusion

As an enterprise progresses from virtualization as an emerging technology to cloud maturity, IT and stakeholders must respond to an endless variety of demands. Rapid response to change requires a holistic solution that enables flexibility and simplicity, but without forfeiting context and control.

The dynamic services model presents a new perspective on IT architecture by providing the foundation for reusable services that understand context and can provide control regardless of application, virtualization, user, device, platform, or location. It adds value to traditional, static, dedicated-server application infrastructure, enables more modern virtual implementations, and becomes a critical requirement of cloud services deployment.