F5 Synthesis: What Makes a Services Fabric?

Lori MacVittie, 2013-17-12

#SDN #SDAS What makes a services fabric, well, a fabric?

Earlier in 2013 when discussing the evolution of application delivery and F5 I posed a question to which, at the time, didn't have an answer. It was simply "What's Next?"

As the time has gone by, the answer to that has become clear. What's next is an evolutionary step in the history of F5 and, more broadly, application delivery. What's next is the evolution from application delivery networks to services fabric.

The ballooning volume of connections, data, users, devices and applications has a very real impact on the data center. It forces changes in both the application and the network architecture to enable businesses to act and react rapidly, despite the need to coordinate across an increasingly complex web of devices, networks and applications. To enable businesses today means enabling IT, and to do that requires an evolution of the entire data center toward a more flexible, dynamic and ultimately service-focused approach to provisioning and management.

That means application delivery architectures must also evolve beyond traditional models. Application delivery networks must evolve into services fabric that support a more integrated, collaborative approach to managing what is increasingly a software-defined and application-driven world.

So what makes a services fabric a "fabric"?

The industry accepted definition of "fabric" generally refers to a topology in which a set of hardware and/or software network elements are connected with each other and offer a consistent set of services. Fabric implies a means of distributing network functions across the elements, often in real-time.

The definition of a services fabric is not all that different.

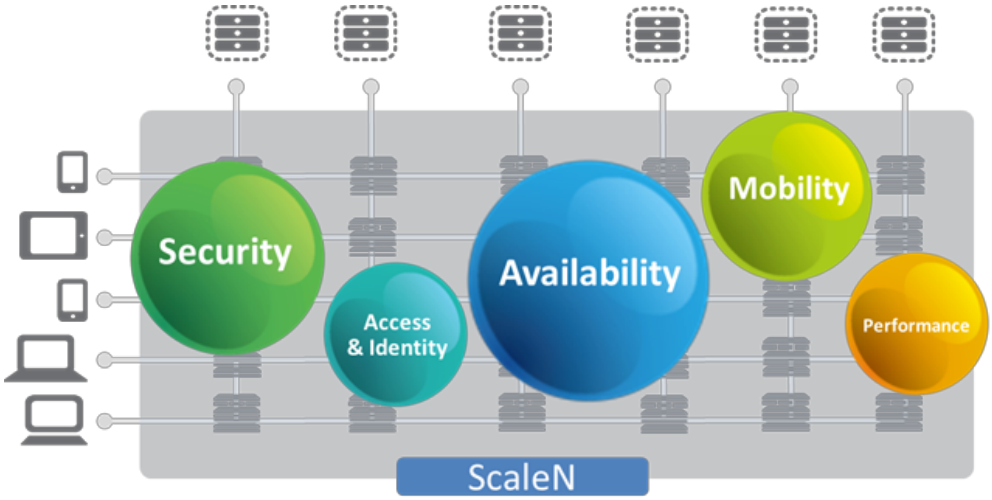

At the heart of F5 Synthesis is a High Performance **Services Fabric:** a topology in which a set of application service elements (F5 platforms deployed on any combination of hardware and software) are connected with each other (Device Service Clustering) and offer a consistent set of services (mobility, cloud, platform, DNS, access and security). ScaleN enables the distribution of service functions across the F5 services fabric to achieve the economy of scale and performance expected of fabric-based technology.

This is not just redundant pairs (or multiple sets) of devices clustered together so that traditional aggregation protocols can be used to distribute traffic across them. While such technologies can certainly be used to emulate a fabric-like model, the result is neither as robust nor reliable as a true **services fabric**.

What's key to being considered a **services fabric** is twofold: it must be all active and service-focused.

1. **All-active.** All elements must be actively participating in the fabric.This is necessary to realize the real-time reliability of services expected. If a problem arises with a service on one element, the fabric must be able to move that service - in real time - to another element. To achieve that, all elements must be active and able to take on new services. Similarly, if a service must be scaled to meet surging demand, it must be possible to provision another service instance elsewhere in the fabric. Thus, an all-active model is a requirement.
2. **Service-focused.** The purpose of a network fabric is to deliver network services. The purpose of a service fabric is to deliver services, specifically for application delivery, **application** services. Applications experiencing a surge in demand require not only that application resources be scaled to meet that demand, but its associated **application** services as well. A services fabric must be able to scale an application's services in tandem with demand. Conversely, problems affecting one service should not affect another. When issues arise in a network fabric, a new path is found to mitigate the problem. When issues arise in a services fabric, a new chain (path) needs to be found to avoid disruption. Thus, a services fabric must be focused on service-level availability, not just device-level availability.

There's nothing all that special about clustering ADCs together into a mesh of devices. Plenty of folks can do that. What's special about the F5 High Performance Services Fabric is that it moves beyond a focus on managing ADCs and marrying pairs of ADCs to applications and focuses on application services. It creates a fabric, a pool of services resources, atop which application services can be provisioned, managed, migrated, and extended.

Service mobility is imperative to enabling a variety of emerging architectures and deployment models. When applications are bound to a specific ADC instance rooted to a topology, it inhibits the ability to move that application or even extend it into the cloud. F5 High Performance Service Fabric, in conjunction with F5 Synthesis' Intelligent Services Orchestration model eliminate such constraints, and ensure that the fabric is not a constraining factor. Rather, it is an enabler of more fluid, mobile architectures and models.

The future of application delivery is a services fabric based architecture.

Additional Resources:

1. F5 Synthesis: The Time is Right
2. F5 and Cisco: Application-Centric from Top to Bottom and End to End
3. F5 Synthesis: Software Defined Application Services
4. F5 Synthesis: Integration and Interoperability
5. F5 Synthesis: High-Performance Services Fabric
6. F5 Synthesis: Leave no application behind
7. F5 Synthesis: The Real Value of Consolidation Revealed

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

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

F5 Networks
Asia-Pacific
apacinfo@f5.com

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

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