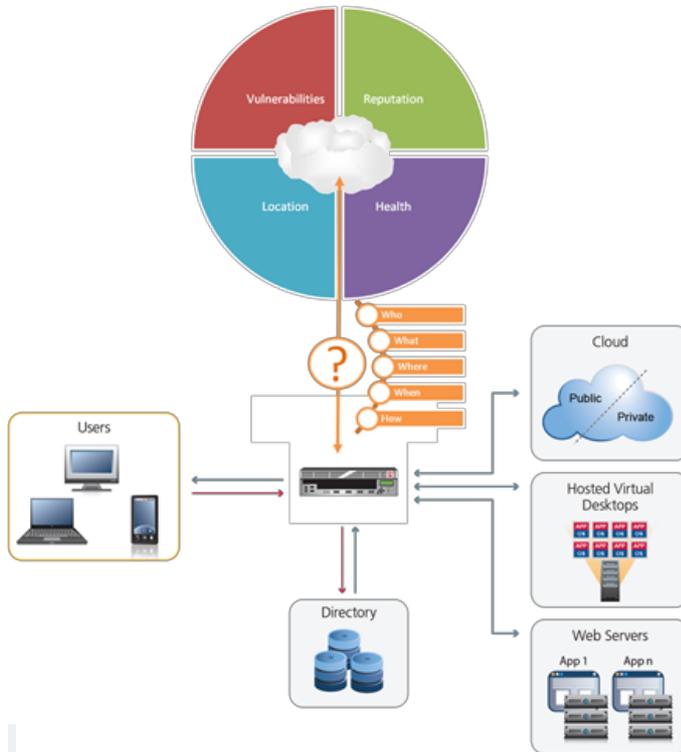


When Big Data Meets Cloud Meets Infrastructure



Lori MacVittie, 2012-02-05

#stirling #interop #infosec #bigdata Bridging the Gap between Big Data and Business Agility



I'm a huge fan of context-aware networking. You know, the ability to interpret requests in the context they were made – examining user identity, location, client device along with network condition and server/application status. It's what imbues the application delivery tier with the agility necessary to make decisions that mitigate operational risk (security, availability, performance) in real-time.

In the past, almost all context was able to be deduced from the transport (connection) and application layer. The application delivery tier couldn't necessarily "reach out" and take advantage of the vast amount of data "out there" that provides more insight into the conversation being initiated by a user. Much of this data falls into the realm of "big data" – untold amounts of information collected by this site and that site that offer valuable nuggets of information about any given interaction.



Because of its expanded computing power and capacity, cloud can store information about user preferences,

which can enable product or service customization. The **context-driven variability provided via cloud** allows businesses to offer users personal experiences that adapt to subtle changes in user-defined context, allowing for a more user-centric experience.

-- "The power of cloud", IBM Global Business Services

All this big data is a gold mine – but only if you can take advantage of it. For infrastructure and specifically application delivery systems that means somehow being able to access data relevant to an individual user from a variety of sources and applying some operational logic to determine, say, level of access or permission to interact with a service.

It's collaboration. It's integration. It's an ecosystem.

It's enabling context-aware networking in a new way. It's really about being able to consume big data via an API that's relevant to the task at hand. If you're trying to determine if a request is coming from a legitimate user or a node in a known botnet, you can do that. If you want to understand what the current security posture of your public-facing web applications might be, you can do that. If you want to verify that your **application delivery controller** is configured optimally and is up to date with the latest software, you can do that.

What's more important, however, is perhaps that such a system is a foundation for integrating services that reside in the cloud where petabytes of pertinent data has already been collected, analyzed, and categorized for consumption. Reputation, health, location. These are characteristics that barely scratch the surface of the kind of information that is available through services today that can dramatically improve the operational posture of the entire data center.

Imagine, too, if you could centralize the acquisition of that data and feed it to every application without substantially modifying the application? What if you could build an architecture that enables collaboration between the application delivery tier and application infrastructure in a service-focused way? One that enables every application to enquire as to the location or reputation or personal preferences of a user – stored “out there, in the cloud” – and use that information to make decisions about what components or data the application includes? Knowing a user prefers Apple or Microsoft products, for example, would allow an application to tailor data or integrate ads or other functionality specifically targeted for that user, that fits the user’s preferences. This user-centric data is out there, waiting to be used to enable a more personal experience. An application delivery tier-based architecture in which such data is aggregated and shared to all applications shortens the development life-cycle for such personally-tailored application features and ensures consistency across the entire application portfolio.

It is these kinds of capabilities that drive the integration of big data with infrastructure. First as a means to provide better control and flexibility in real-time over access to corporate resources by employees and consumers alike, and with an eye toward future capabilities that focus on collaboration inside the data center better enabling a more personal, tailored experience for all users.

It’s a common refrain across the industry that network infrastructure needs to be smarter, make more intelligent decisions, and leverage available information to do it. But actually integrating that data in a way that makes it possible for organizations to actually codify operational logic is something that’s rarely seen.

Until now.

-
- [Mobile versus Mobile: 867-5309](#)
 - [Why Layer 7 Load Balancing Doesn’t Suck](#)
 - [BYOD–The Hottest Trend or Just the Hottest Term](#)
 - [What Does Mobile Mean, Anyway?](#)
 - [Mobile versus Mobile: An Identity Crisis](#)
 - [The Three Axioms of Application Delivery](#)
 - [Don’t Let Automation Water Down Your Data Center](#)
 - [The Four V’s of Big Data](#)
 - [The Challenges of SQL Load Balancing](#)
-

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