Who owns application delivery meta-data in the cloud?

Lori MacVittie, 2009-06-02

While the vast majority of folks are still debating what is or is not "cloud computing", there are already groups trying to get ahead of the curve by focusing on broader issues such as interoperability and portability. Indeed, by addressing the potential pitfalls associated with portability across cloud implements now rather than later, it is hoped that there won't be as many problems when it does finally become an issue.

There is a very real danger, however, that cloud interoperability and portability specifications will fail to address the very real need to include all the relevant application and network infrastructure meta-data necessary to move an application from one cloud to another. Because the network and application network infrastructure is often seen as little more than "dumb piping" it is often assumed that these vital components of a successful application delivery strategy can be simply be exchanged as easy as light bulbs. But anyone who's successful deployed a well-performing, secure application knows it takes more than just an application, and its supporting infrastructure. There's security and acceleration and optimization in the application networking infrastructure, as well, without which the application would be very much at risk for exploitation or lack of adoption for performance reasons.

The application security and acceleration policies associated with an application are often complex and are very often peculiar to the application. Those seemingly irrelevant announcements made by application delivery vendors like F5 regarding certification of solutions with specific application partners like SAP, Oracle, and Microsoft obscure the long hours spent in test environments tweaking security and acceleration policies designed to eek out the best performance with the highest levels of security possible for those applications. The seeming simple nature of the policies resulting from those efforts belies the very complex, arduous process through which those policies have been created.

Such policies cannot be created for custom applications. Base policies can address common performance and security related configuration options based on transport and application protocols, like TCP and HTTP, but they can't specifically be tuned to an application until it's actually deployed with the application.

Joyent uses F5's BIG-IP devices as the backbone to its cloud computing IaaS. The BIG-IP platform provides massive traffic handling (2-10 Gbps), while F5's powerful yet easy-to-use iRules™ scripting language provides Joyent with flexible management and deployment of its cloud computing infrastructure.

Once the application delivery network is tuned to deliver an application it essentially becomes a part of the implementation; it becomes a necessary component of the application without which security and performance can degrade. If the application is to be moved from one cloud to another, the security and delivery policies need to move ***with*** the application in order to ensure that neither security nor performance of the application is compromised.

"I've looked at the market and tried virtually everything, but there is nothing else like the F5 BIG-IP system," said Jason Hoffman, co-founder and CTO of Joyent. "BIG-IP LTM is the only application switch capable of scaling to handle the thousands of back-end systems Joyent needs to thrive. Without it, we wouldn't have a business, to be honest."

SOURCE: F5 Networks

But as Alistair Croll points out in this interview at Data Center Knowledge, the question of who owns meta-data may prevent this from becoming reality. Like the popularity of a picture on Flickr, the ownership of application network infrastructure meta-data (the security and delivery policies) is highly in question.

After all, the ability to deliver your application faster and more securely may be part of the "secret sauce" of a cloud computing provider's offering, and one of its differentiating features. If one cloud computing provider is able to accelerate the delivery of your application 20% but another can only provide 10% and performance is an essential criterion in your decision making process, then it is not advantageous for the cloud computing provider to enable the sharing of those delivery policies with other providers.

So if the application delivery network is such an integral piece of a cloud computing provider's infrastructure, it seems unlikely they'll be willing to share the relevant meta-data with other cloud computing providers, driving complete interoperability and portability efforts to concentrate simply on application infrastructure. It is unlikely that Joyent, for example, would willingly hand out the BIG-IP policies it relies on to handle billions of transactions a month to another cloud provider.

It is possible that if a specification regarding application network delivery metadata were abstracted and could be applied across application delivery network implementations, that the "secret sauce" of a cloud computing provider's offering could be maintained while still allowing portability across cloud implementations. Such a generic specification would allow the meta-data policies to be transported and applied across different cloud implementations while preserving the specific details of implementation within the cloud computing infrastructure. The choice of application delivery infrastructure would remain an integral differentiation for cloud computing providers as each implementation of the metadata would remain specific to the infrastructure provider and therefore be better or worse depending on the implementation.

But as Alistair pointed out, the real question right now is *who owns the meta-data?* If the answer is the cloud computing provider, then even attempting to formulate such an interoperability specification that bridges application delivery infrastructure implementations seems as though it would be a wasted effort.

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

F5 Networks, Inc.	F5 Networks	F5 Networks Ltd.	F5 Networks
Corporate Headquarters	Asia-Pacific	Europe/Middle-East/Africa	Japan K.K.
info@f5.com	apacinfo@f5.com	emeainfo@f5.com	f5j-info@f5.com