Cloud Computing and Infrastructure 2.0

Lori MacVittie, 2008-17-10

Not every infrastructure vendor needs new capabilities to support cloud computing and infrastructure 2.0.

Greg Ness of Infoblox has an excellent article on "The Next Tech Boom: Infrastructure 2.0" that is showing up everywhere. That's because it raises some interesting questions and points out some real problems that will be need to be addressed as we move further into cloud computing and virtualized environments. What is really interesting, however, is the fact that some infrastructure vendors are already there and have been for quite some time.

One thing Greg mentions that's not quite accurate (at least in the case of F5) is regarding the ability of "appliances" to "*look inside servers (for other servers) or dynamically keep up with fluid meshes of hypervisors*".

From Greg's article:

> *The appliances that have been deployed across the last thirty years simply were not architected to look inside servers (for other servers) or dynamically keep up with fluid meshes of hypervisors powering servers on and off on demand and moving them around with mouse clicks.*
>
> *Enterprises already incurring dis-economies of scale today will face sheer terror when trying to manage and secure the dynamic environments of tomorrow. Rising management costs will further compromise the economics of static network infrastructure.*

I must disagree. Not on the sheer terror statement, that's almost certainly true, but on the capabilities of infrastructure devices to handle a virtualized environment. Some appliances and network devices have long been able to look inside servers and dynamically keep up with the rapid changes occurring in a hypervisor-driven application infrastructure. We call one of those capabilities "intelligent health monitoring", for example, and others certainly have their own special name for a similar capability.

On the dynamic front, when you combine an intelligent application delivery controller with the ability to be orchestrated from within applications or within the OS, you get the ability to dynamically modify configuration of application delivery in real-time based on current conditions within the data center. And if you're monitoring is intelligent enough, you can sense within seconds when an application - whether virtualized or not - has disappeared or conversely, when it's come back on line. F5 has been supporting this kind of dynamic, flexible application infrastructure for years. It's not really new except that its importance has suddenly skyrocketed due to exactly the scenario Greg points out using virtualization.

WHAT ABOUT THE VIRTSEC PIECE?

There has never been a better case for centralized web application security through a web application firewall and an application delivery controller. The application delivery controller - which necessarily sits between clients and those servers - provides security at layers 2 through 7. The full stack. There's nothing really that special about a virtualized environment as far as the architecture goes for delivering applications running on those virtual servers; the protocols are still the same, and the same vulnerabilities that have plagued non-virtualized applications will also plague virtualized ones. That means that existing solutions can address those vulnerabilities in either environment, or a mix.

Add in a web application firewall to centralize application security and it really doesn't matter whether applications are going up and down like the stock market over the past week. By deploying the security at the edge, rather than within each application, you can let the application delivery controller manage the availability state of the application and concentrate on cleaning up and scanning requests for malicious content.

Centralizing security for those applications - again, whether they are deployed on a "real" or "virtual" server - has a wealth of benefits including improving performance and reducing the very complexity Greg points out that makes information security folks reach for a valium.

BUT THEY'RE DYNAMIC!

Yes, yes they are. The assumption is that given the opportunity to move virtual images around that organizations will do so - and do so on a frequent basis. I think that assumption is likely a poor one for the enterprise and probably not nearly as willy nilly for cloud computing providers, either. Certainly there will some movement, some changes, but it's not likely to be every few minutes, as is often implied.

Even if it was, some infrastructure is already prepared to deal with that dynamism. Dynamism is just another term for agility and makes the case well for loose-coupling of security and delivery with the applications living in the infrastructure. If we just apply the lessons we've learned from SOA to virtualization and cloud computing and 90% of the "Big Hairy Questions" can be answered by existing technology. We just may have to change our architectures a bit to adapt to these new computing models.

Network infrastructure, specifically application delivery, has had to deal with applications coming online and going offline since their inception. It's the nature of applications to have outages, and application delivery infrastructure, at least, already deals with those situations. It's merely the frequency of those "outages" that is increasing, not the general concept.

But what if they change IP addresses? That would indeed make things more complex. This requires even more intelligence but again, we've got that covered. While the functionality necessary to handle this kind of a scenario is not "out of the box" (yet) it is certainly not that difficult to implement if the infrastructure vendor provides the right kind of integration capability. Which most do already.

Greg isn't wrong in his assertions. There ***are*** plenty of pieces of network infrastructure that need to take a look at these new environments and adjust how they deal with the dynamic nature of virtualization and cloud computing in general. But it's not all infrastructure that needs to "get up to speed". Some infrastructure has been ready for this scenario for years and it's just now that the application infrastructure and deployment models (SOA, cloud computing, virtualization) has actually caught up and made those features even more important to a successful application deployment.

Application delivery in general has stayed ahead of the curve and is already well-suited to cloud computing and virtualized environments. So I guess some devices are already "Infrastructure 2.0" ready.

I guess what we really need is a sticker to slap on the product that says so.

Related Links

Are you (and your infrastructure) ready for virtualization?

Server virtualization versus server virtualization

Automating scalability and high availability services

The Three "Itys" of Cloud Computing

4 things you need in a cloud computing infrastructure

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