The Potential Ramifications of Platform-Based Vulnerabilities on Cloud Computing

Lori MacVittie, 2012-08-02

#infosec #adcfw #cloud *Alternate title: How to take out an entire PaaS cloud with one vulnerability*

Apache Killer.

Post of Doom.

What do these two vulnerabilities have in common? Right, they're ***platform***-based vulnerabilities. Meaning they are vulnerabilities peculiar to the web or application server platform upon which applications are deployed. Mitigations for such vulnerabilities generally point to changes in configuration of the platform – limit post size, header value sizes, turn off some value in the associated configuration.

But they also have something else in common – risk. And not just risk in general, but risk to cloud providers whose primary value is in offering not just a virtual server but an entire, pre-integrated and pre-configured application deployment stack. Think LAMP, as an example, and providers like Microsoft (Azure) and VMware (CloudFoundry), more commonly adopting the moniker of PaaS. It's an operational dream to have a virtual server pre-configured and ready to go with the exact application deployment stack needed and offers a great deal of value in terms of efficiency and overall operational investment, but it is – or should be – a security professional's nightmare. It's not unlike the recent recall of Chevy Volts – a defect in the platform needs to be mitigated. The only way to do it, for car owners, is to effectively shut down their ability to drive while a patch is applied. It's disruptive, it's expensive (you still have to get to work, after all), and it's frustrating for the consumer. For the provider, it's bad PR and negatively impacts the brand. Neither of which is appealing.

A vulnerability in the application stack, in the web or application server, can be operationally devastating to the provider – and potentially disruptive to the consumer whether the vulnerability is exploited or not.

STANDARDIZATION is a DOUBLE-EDGED SWORD

Assume a homogeneous cloud environment offering an application stack based on Microsoft ASP. Assume now an exploit, oh say like Post of Doom, is discovered whose primary mitigation lies in modifying the configuration of each and every instance. Virtualization of any kind provides a solution, of course, but introduces the possibility of disruption in the impact to consumer applications from the configuration change. A primary mitigation for the Post of Doom is to limit the size of data in a POST to under 8MB. Depending on the application, this has to potential to "break" application functionality, particularly those for which uploading big data is a focus. Images, video, documents, etc… These all may be impacted negatively, disrupting applications and angering consumers.

Patching, of course, is preferred, as it eliminates the underlying vulnerability without potentially breaking applications. But patching takes time – time to develop, time to test, time to deploy. The actual delivery of such patches in a PaaS environment is a delicate operation. You can't just shut the whole cloud down and restart it after the patches are applied to the base images, can you? Do you wait, quiesce the vulnerable images and only force the patched ones when new instances are provisioned? A configuration-based mitigation, too, has these same issues. You can't just shut down the whole cloud, apply the change, and reboot.

It's a delicate balance of security versus availability that must struck for the provider, and certainly their position in such cases is one not to be envied. Damned if they do, damned if they don't.

Then there is the risk of exploitation ***before*** any mitigation is applied. If I want to wreak havoc on a PaaS, I may be able to accomplish simply by finding one with the appropriate platform vulnerable to a given exploit, and attack. Cycling through applications deployed in that environment (easily identified at the network layer by the IP ranges assigned to the provider) should result in a wealth of chaos being wrought. The right vulnerability could take out a significant enough portion of the environment to garner attention from the outages caused.

Enterprise organizations that think they are immune from such issues should think again, as even a cloud provider is often not as standardized on a single application platform as an enterprise is, and it is that standardization that is at the root of the potential risk from platform-based vulnerabilities. Standardization, commoditization, these are good things in terms of many financial and operational benefits, but they can also cause operational risk to increase.

MITIGATE in the MIDDLE

There is a better solution, a better strategy, a better operational means of mitigating platform-based risks.

This is where the role of a flexible, broad-spectrum layer of security applies. One that enables security professionals to broadly apply security policies to quickly mitigate potentially disastrous vulnerabilities. Without disrupting a single running instance, an organization can deploy a mitigating solution that detects and prevents the effects of such vulnerabilities. Applying security policies that mitigate such vulnerabilities ***before*** they reach the platform is critical to preventing a disaster of epic (and newsworthy) proportions.

Whether stop gap or a permanent solution, by leveraging the application delivery tier of any data center – enterprise or cloud provider – such vulnerabilities can be addressed without imposing harsh penalties on applications and application owners, such as requiring complete shutdown and reboots.

Leveraging such a flexible data center tier insulates the platform from exploitation while insulating customers from the disruption required to mitigate immediately on the platform layer, allowing time to redress through patches or, at least, understand the potential implication to the application from the platform configuration changes required to mitigate the vulnerability.

In today's data center, time is perhaps the biggest benefit afforded to IT by any solution, and yet the one least likely to be provided. A flexible application delivery tier capable of mitigating threats across the network and application stack without disruption is one of the few solutions available that offers the elusive and very valuable benefit of time. Providers and enterprises alike need to consider their current data center architecture and whether it supports the notion of such a dynamic tier. If not, it's time to re-evaluate and determine whether a strategic change of direction is necessary to ensure the ability of operations and security teams to address operational risk as quickly and efficiently as possible.

- At the Intersection of Cloud and Control…
- The Full-Proxy Data Center Architecture
- The Pythagorean Theorem of Operational Risk
- The Future of Cloud: Infrastructure as a Platform
- Infrastructure Architecture: Whitelisting with JSON and API Keys
- If Security in the Cloud Were Handled Like Car Accidents
- VU#903934 – Post of Doom
- F5 Friday: Zero-Day Apache Exploit? Zero-Problem

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