Application Security in the Cloud is still Cloudy

Lori MacVittie, 2013-12-09

#infosec #cloud An interesting statistic raises an even more interesting question.

IBM shared a security related infographic via Twitter recently and in looking through the statistics (most of which are attributed to 2011 research, by the way) I happened to catch a statement claiming "The average company is attacked 60,000 times a day."

IBM notes that "average" is average for the study, which consisted of mostly large enterprises, and while I'm certain there are still experts who would dispute this claim (it's higher! it's lower! That's only an average of a subset of a selection of a ...) for me it raised an interesting question with respect to attacks and cloud-based applications: If your application is deployed in the cloud, how do you if/when it's being attacked? Perhaps more importantly, though, is whether or not you ***should*** know. After all, "the cloud" is taking care of all that infrastructure and networky stuff under the covers for you, right? And part of that "stuff" is addressing attacks.

NETWORK versus APPLICATION

Certainly there are many instances of organizations pointing out that their cloud provider is, in fact, dealing with ***network*** layer attacks.

We could spend an entire post detailing attacks (there've been many) but that's not the point. In general, the attacks that cloud providers are able to detect and thus address (even reactively, if not proactively) are ***network*** layer attacks, not application layer attacks.

You know, the kind of attacks that target vulnerabilities in either your application or its platform (web / application server) and sometimes in its network stack (think SYN floods and such). Cloud providers would (logically and correctly) point out that they have no control over the application or its platform (unless it's a PaaS or SaaS provider) and thus it is not their responsibility to monitor for attacks against those components.

And they'd be (mostly) right. Even in the cloud, the security of your application is still your responsibility because it's your application. Potential vulnerabilities in the application layer - whether data, logic, behavior or code based - are yours to address, not the provider. Where things get murky is at the protocol (HTTP & TCP) layers [pdf], where exploitation can be used with relatively few resources to successfully execute a DDoS attack against your application and yet there is virtually no way for the application instance to recognize such an attack.

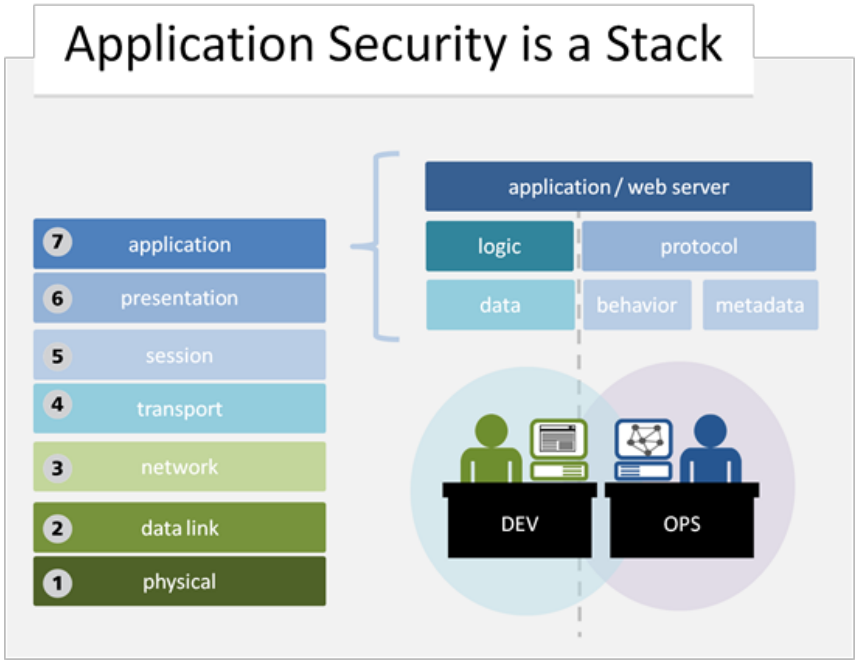

That's because detecting some types of application layer attacks requires visibility into both client characteristics as well as its behavior. A client connecting from a 100Mbps capable LAN that purposefully tries to use a tiny TCP window size or seems to be receiving data at a rate far below what it is capable of just might be part of an attempt to consume and keep busy application resources and in doing so with enough clients, successfully deny service to legitimate clients.

This kind of attack (which is often part of a holistic and much larger (not to mention coordinated) attack designed to distract organizations from true intent) is not detectable by an application instance unless it has end-to-end visibility. In the cloud, that's just not possible. Visibility into client characteristics is typically no deeper than a client IP address stuffed into a custom HTTP header. You can't infer an attack from an IP address, and transfer speed means little without other variables against which to measure the normality (or abnormality) of the behavior.

Even deploying traditional services to address (web application firewall, application delivery firewall) may not provide the visibility required because those services are being deployed atop the abstracted, service-based infrastructure and are treated more like application instances than infrastructure services in need of visibility into the network. Certainly these services will assist in mitigating many application layer attacks that focus on logic, platform, or data exploitation, but they may not be able to fully analyze the protocol layer because of the many layers of abstraction between them and the variables they need.

Thus, while aspects of application security in the cloud are certainly the responsibility of the developer (or organization), there are other aspects that require the assistance of the provider. Either in the form of services that have access to the information necessary or a means of sharing that same information with services deployed atop the infrastructure.

Which leaves us where we are today: application security in the cloud is still cloudy. It's definitely an area in which there is room for new services and solutions.

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