DevCentral Top 5: May 27, 2014

John Wagnon, 2014-27-05

And here we are again...that fateful time when we admire and, yes, celebrate the amazing contributions of our DevCentral authors. The easy part is writing about all the great content; the hard part is picking just 5 articles to highlight. Nonetheless, here they are in no particular order...the DevCentral Top 5:

Mitigating sslsqueeze and other no-crypto, brute force SSL handshake attacks

If you don't know David Holmes, you pretty much need to stop whatever you are doing and start reading his stuff. In this article, he wows us again with his security expertise and shows us why we all love the flexibility (and programmability) of F5 technology. David wrote an article back in 2011 on an SSL Renegotiation DOS Attack and showed how a client can take advantage of a server in the SSL handshake because an SSL handshake requires at least 10 times more processing power on the server than on the client. Well, David recently found a new class of SSL attacks where the client doesn't do any cryptography at all...the client just sends a bunch of pre-canned packets that look like an SSL handshake. In this attack, the server uses ***100 times*** more processing power than the client! David was able to get a copy of one of the tools (called **sslsqueeze**) that runs this attack and use it against a real, physical BIG-IP with cryptographic offload accelerators. He was able to take a $200 used computer and overload the BIG-IP with fake SSL handshakes. The good news is that he wrote an iRule (actually 2 iRules for different scenarios) to mitigate this attack. You gotta love the flexibility that iRules provide! The security world is an interesting one...attackers will always find a new vulnerability to exploit, and good guys will find a way to stop them. Thank goodness David Holmes is a good guy.

Hybrid DDoS Needs Hybrid Defense

It's not if you get DDoS attacked, but when. Lori MacVittie publishes another masterful write-up where she outlines a DDoS approach that many top analysts recommend. In the world of DDoS protection, it's best to implement off-premise detection and mitigation with on-premise protection. A hybrid solution provides the resiliency and scale of cloud based solutions with the granularity and always-on capabilities of on-premise solutions. Many times, when an organization is under attack, the answer is to shut down computationally expensive services to prevent overall service outages. But, this means IPS, firewalls, anti-fraud detection, etc are sometimes eliminated for the sake of keeping the network running. With a hybrid approach, an organization can take advantage of additional capacity in the cloud but still maintain the flexibility of protecting against more frequent and easily managed attacks. Where do you turn for the technology to make all this happen? The good news for everyone is that Defense.Net just joined forces with the F5 family. By combining the cloud-based services of Defense.Net and the on-premise protection of the F5 firewall, organizations will be better armed to detect and mitigate DDoS attacks at the network and application layers...simultaneously!

The BIG-IP GTM: Configuring DNSSEC

Is it awkward that I'm including one of my own articles in this edition of the Top 5? It's only weird if we let it be weird. So, I'm cool with it if you are. OK, but for real, F5 does some seriously awesome work when it comes to DNS services. We all know that DNS was originally built back in the 1980s, and it was designed with some inherent trust features that bad guys could exploit. When a user types a web address in his browser, he expects to be reliably directed to the correct website. If an attacker is able to manipulate the response from a DNS server, he could send the unsuspecting user to a malicious site that is full of malware. This vulnerability (among others) created the need for a more secure DNS experience. DNSSEC addresses the security problem by validating the response of DNS servers. This is done through a trust relationship that is built with a series of security keys. As you can imagine, validating each DNS response can be computationally expensive, so it's nice if you have custom-built, high powered hardware and software to do this job for you. Well, the BIG-IP GTM does just that. It will authoritatively answer your DNS requests, but it will also sign DNSSEC validated responses. You'll need to configure a few things on the BIG-IP to make this happen, but this article shows you all the steps needed to take care of that. So get out there and configure your BIG-IP GTM and let it handle all your DNS needs.

Is OpenStack Ready For Production?

Ranjeet Sonone answers THE question that runs through the mind of anyone diving deep on OpenStack. Way back in 2012, OpenStack lacked solid networking design and required significant resource allocation to be implemented in a true production environment. By 2013, OpenStack was gaining momentum and user success stories were helping it grow in popularity. But even then, you needed lots of system knowledge and scripting skills to assemble all the moving parts. So, many people were watching as 2014 approached to see if OpenStack was truly ready for prime time. The answer was given at the recent OpenStack summit in Atlanta, GA where 4,500 networking professionals heard about real world OpenStack success stories. OpenStack is now sufficiently rich in service offerings, and the core components are now stable for production environments...so, yes, it's ready for production! F5 has been investigating OpenStack for several years now, and we joined the OpenStack Foundation last year. Customers are now using F5 plug-ins in their OpenStack labs and sharing workloads between the public cloud and their private cloud. In the end, Ranjeet shows how F5 ensures that our customers have access to our custom built solutions when evaluating cloud platform integrations. Great job Ranjeet!

Heartbroken and then Redeemed

The Heartbleed vulnerability blew up a little over a month ago, and F5 was right there to help mitigate this significant problem. After the initial press from the Hearbleed bug had calmed down a little, a security researcher named Yngve Pettersen discovered hundreds of new Internet hosts that appeared to be vulnerable to Heartbleed. He called these newly vulnerable hosts "heartbroken" servers. He also noted that a specific characteristic on these servers suggested that they might be F5 devices. Not good, right? Well, our worldwide security evangelist David Holmes contacted Yngve and asked about his research. Yngve was kind enough to share his data with us. After all, he wasn't trying to give F5 a black eye, he was simply stating the results of his research. David (and a host of other F5 security experts) analyzed the data and found that actually ***none*** of the devices were F5 equipment. That's good news! Unfortunately, Yngve had already posted a technology blog about his initial findings, but he was gracious to remove the F5 references from his blog and offer an apology to F5 and F5's customers. I think that's really awesome stuff. Yngve is a really smart guy who discovered some very interesting data, but he is also a true professional in that he is willing to admit when he makes a mistake (as we all do). So, kudos to David for asking the clarifying questions, and kudos to Yngve for being a true professional.

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