Cloud Testing: The Next Generation

Lori MacVittie, 2010-29-06

It seems only fair that as the Internet caused the problem, it should solve it.

One of the negatives of deploying an Internet-scale infrastructure and application is that until it's put to the test, you can't have 100 percent confidence that it will scale as expected. If you do, you probably shouldn't. Applications and infrastructure that perform well – and correctly – at nominal scale may begin to act wonky as load increases. Dan Bartow, VP at SOASTA, says it is still often load balancing configuration errors that crop up during testing that impedes scalability and performance under load. Choices regarding the load balancing algorithm have a direct impact on the way in which sites and applications scale – or fail to scale – and only under stress does infrastructure and applications begin to experience problems. The last time I ran a scalability and performance test on industry load balancers that's exactly what happened – what appeared to be a well-behaving Load balancer under normal load turned into a temper-tantrum throwing device under heavier load. The problem? A defect deep in the code that only appeared when the device's session table was full. Considering the capability of such devices even then, that meant millions of connections had to be seen in a single session before the problem reared its ugly head.

Today load balancers are capable of not millions, but tens of millions of connections. Scale that is difficult if not impossible for organizations to duplicate. cloud computing and virtualization bring new challenges to testing the scalability of an application deployment. Application deployed in a cloud environment may be designed to auto-scale "infinitely" which implies testing that application and its infrastructure requires the same capability in a testing solution.

That's no small trick. Traditionally organizations would leverage a load testing solution capable of generating enough clients and traffic to push an application and its infrastructure to the limits. But given increases in raw compute power and parallel improvements in capacity and performance of infrastructure solutions, the cost of a solution capable of generating the kind of Internet-scale load necessary is prohibitive.

One of our internal performance management engineers applied some math and came up with a jaw-dropping investment:

> In other words, enough hardware to test a top-of-the-line ADC [application delivery controller] would set you back **a staggering $3 million.** It should be clear that even buying equipment to test a fairly low-end ADC would be a big ticket item, likely costing quite a bit more than the device under test.

It seems fairly obvious that testing Internet-scale architectures is going to require Internet-scale load generation solutions but without the Internet-scale cost. It's only fair that if the scalability of the Internet is the cause of the problem that it should also provide the solution.

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