

Load Balancing: Welcome to the 21st Century



Lori MacVittie, 2008-11-06

Kevin Saitta, a solution consultant, has a [nice blog post on architecting a Microsoft BizTalk 2006 R2 solution](#).

Unfortunately, amidst the goodness, is a statement regarding load balancers that needs to be corrected. Kevin is not alone in his beliefs regarding load balancers, unfortunately, I've seen a lot of posts lately that seem to indicate that folks out there still have a [circa 1999 knowledge set regarding the capabilities of load balancers](#).

Kevin writes

Load Balancer

A load balancer balances the load between servers, but more importantly, if one server goes down, the load balancer sends all requests to the server(s) left running. An example of hardware load balancing is [BIG-IP from F5 Networks](#). An example of software load balancing is the [Network Load Balancing Service](#) included in Windows.

NOTE: Load balancers only understand server level availability and have no understanding of the application that run on the servers. This means that if the application is unavailable the load balancer will still deliver request to the server, hence the need for a monitoring solution as SCOM.

It is the "NOTE" that needs to be addressed, especially given the reference to [BIG-IP](#), which is absolutely capable of understanding applications that run on the servers and, in fact, does understand the applications and their protocols and is capable of optimizing, accelerating, and securing those applications.



[BIG-IP](#) is capable of recognizing the difference between the server being available, and the application being available *and* has the ability to understand when the application is serving the right content or not. [BIG-IP](#) can even dig into the payload, the application messages, and apply policies based on the application-specific nuggets of information, rather than just HTTP headers or even application(HTTP)/transport layer (TCP) protocols. Add in [iRules](#) and [BIG-IP](#) becomes an *application delivery platform*: extensible, dynamic, and extremely

intelligent.

The problem is that load balancers *aren't just load balancers* any more. They have [evolved into application delivery controllers](#), and are capable of much more intelligence and integration with the applications they are tasked to deliver and are just as capable of intelligent monitoring of applications to ensure they are actually available.

Kevin's assertion that a load balancer will deliver requests to the server even if it's down is wrong, at least in the case of [BIG-IP](#). While Windows NLB is certainly incapable of this particular functionality (it's called *network* load balancing for a reason), [application delivery controllers](#) like [BIG-IP](#) and its competitors *are* capable of intelligent monitoring, dynamic routing, and recognizing when an application is no longer available or responding correctly.

As [Deb](#) pointed out in [response to a previous post](#), there are plenty of examples and samples for [BIG-IP](#) which illustrate the powerful, flexible, application aware monitoring capabilities of [BIG-IP](#).

- [Sample Monitors in the Adv Design & Config codeshare](#)
- [Tech tip: Passive Application Monitoring with LTM](#)
- [Health Monitor page in the Adv Design & Config wiki](#)
- [HTTP Monitoring with POST](#)
- [Monitoring & Management Forum on DevCentral](#)
- [Passive Monitoring - Maintaining Performance and Health \(white paper\)](#)

To be fair to Kevin, who wrote a really informative and interesting post, this is not the first blog or forum post I've seen in the last two weeks that improperly tags modern load balancers as "dumb" network devices incapable of understanding application layer protocols and messages. There's a perception - and a wrong one at that- that load balancers are still just simple layer 4 devices that add very little value to the network or application architecture.

As it's part of my job to educate folks on topics like this, I guess I'm not doing very well, am I? I'll get to work on that, pronto, after my [upcoming mini-vacation](#).

For starters I'd like to say to everyone out there interested in load balancing: welcome to the 21st century, where application delivery controllers have replaced load balancers, and where intelligent, application aware devices are capable of much, much more than their predecessors.

Stay tuned, there's much, much more to come.

Imbibing: Coffee

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