The Disadvantages of DSR (Direct Server Return)

Lori MacVittie, 2008-03-07

I read a very nice blog post yesterday discussing some of the traditional pros and cons of load-balancing configurations. The author comes to the conclusion that if you can use direct server return, you should.

I agree with the author's list of pros and cons; DSR is the least intrusive method of deploying a load-balancer in terms of network configuration. But there are quite a few disadvantages missing from the author's list.

Author's List of Disadvantages of DSR

The disadvantages of Direct Routing are:

- Backend server must respond to both its own IP (for health checks) and the virtual IP (for load balanced traffic)
- Port translation or cookie insertion cannot be implemented.
- The backend server must not reply to ARP requests for the VIP (otherwise it will steal all the traffic from the load balancer)
- Prior to Windows Server 2008 some odd routing behavior could occur in <2% of Windows Server installation.
- In some situations either the application or the operating system cannot be modified to utilse Direct Routing.

Some additional disadvantages:

1. **Protocol sanitization can't be performed.**
   This means vulnerabilities introduced due to manipulation of lax enforcement of RFCs and protocol specifications can't be addressed.
2. **Application acceleration can't be applied.**
   Even the simplest of acceleration techniques, e.g. compression, can't be applied because the traffic is bypassing the load-balancer (a.k.a. application delivery controller).
3. **Implementing caching solutions become more complex.**
   With a DSR configuration the routing that makes it so easy to implement requires that caching solutions be deployed elsewhere, such as via WCCP on the router. This requires additional configuration and changes to the routing infrastructure, and introduces another point of failure as well as an additional hop, increasing latency.
4. **Error/Exception/SOAP fault handling can't be implemented.**
   In order to address failures in applications such as missing files (404) and SOAP Faults (500) it is necessary for the load-balancer to inspect outbound messages. Using a DSR configuration this ability is lost, which means errors are passed directly back to the user without the ability to retry a request, write an entry in the log, or notify an administrator.
5. **Data Leak Prevention can't be accomplished.**
   Without the ability to inspect outbound messages, you can't prevent sensitive data (SSN, credit card numbers) from leaving the building.
6. **Connection Optimization functionality is lost.**
   TCP multiplexing can't be accomplished in a DSR configuration because it relies on separating client connections from server connections. This reduces the efficiency of your servers and minimizes the value added to your network by a load balancer.

There are more disadvantages than you're likely willing to read, so I'll stop there. Suffice to say that the problem with the suggestion to use DSR whenever possible is that if you're an application-aware network administrator you know that most of the time, DSR isn't the right solution because it restricts the ability of the load-balancer (application delivery controller) to perform additional functions that improve the security, performance, and availability of the applications it is delivering.

DSR is well-suited, and always has been, to UDP-based streaming applications such as audio and video delivered via RTSP. However, in the increasingly sensitive environment that is application infrastructure, it is necessary to do more than just "load balancing" to improve the performance and reliability of applications. Additional application delivery techniques are an integral component to a well-performing, efficient application infrastructure.

DSR may be easier to implement and, in some cases, may be the right solution. But in most cases, it's going to leave you simply serving applications, instead of delivering them.

Just because you can, doesn't mean you should.

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