

Acceleration ABCs - O is for Outages (Otherwise Known as Oops)



Dawn Parzych, 2008-13-11



Having a site unavailable is probably worse than having a slow loading page. Outages can occur for multiple reasons such as an overloaded back end server, an upgrade that didn't go as planned, or a power outage at the data centre. Whatever the reason the end result is the same dissatisfied customers (some which may never come back) and lost revenue. Nobody is immune to outages there have been some fairly high profile sites that have had outages over the last year including [Amazon](#), [Twitter](#), and [Facebook](#) - just do a search for outages and you'll find many more that have occurred over the years.

With the evolution of social media news of an outage spreads rapidly and reaches a far broader audience than previously as addressed by Stephen Pierzchala in his [Newest Industry blog](#):

"I became aware of both of these outages through a combination of FriendFeed and Twitter within minutes of them starting. This information spread quickly. And, due to the nature of these new technologies, people were able to comment on the outages, and theorize about the cause of the problems these large online firms faced."

Gone are the days when you can fix the problem and get your site back on-line before anyone realizes there was a problem.

Luckily there are solutions available to mitigate the risk of outages. A [global traffic manager](#) can direct users to a secondary data centre if a power outage occurs at the primary data centre. An [application delivery controller](#) can be used to mark a node or server as down if it fails to respond and an [F5 Networks iRule](#) can be used to redirect users to an "Under Construction" or "Site Unavailable" page as highlighted in this [DevCentral forum post](#). But is presenting a "Site Unavailable" page really the best option or what happens if you have a single application server that goes down and no secondary data centre available to fail over to? The solution: a little known feature in [WebAccelerator](#) called stand-in.

The stand-in feature allows WebAccelerator to continue serving cached content for a period of time after it has expired from cache. Typically the first request after content has expired will result in a conditional GET being issued to the server to revalidate the content if an error is received that would then be passed onto the customer. With the stand-in feature enabled if an error message is received WebAccelerator will continue serving content for a configurable amount of time. Isn't it better to serve some stale or expired content rather than no content at all?

Image courtesy of Leia.

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

©2016 F5 Networks, Inc. All rights reserved. F5, F5 Networks, and the F5 logo are trademarks of F5 Networks, Inc. in the U.S. and in certain other countries. Other F5 trademarks are identified at f5.com. Any other products, services, or company names referenced herein may be trademarks of their respective owners with no endorsement or affiliation, express or implied, claimed by F5. CS04-00015 0113