

Project Acceleration: Conclusions



Colin Walker, 2013-25-04

Note: As of 11.4, WebAccelerator is now a part of BIG-IP Application Acceleration Manager.

This is tenth and final article in a series on DevCentral's implementation of WebAccelerator. Join Colin Walker and product manager Dawn Parzych as they discuss the ins and outs of WebAccelerator. Colin discusses his take on implementing the technology first hand (with an appearance each from Jason Rahm and Joe Pruitt) while Dawn provides industry insight and commentary on the need for various optimization features.

- [Part 1 - The Basics](#)
- [Part 2 - TCP Optimization & Compression](#)
- [Part 3 - WebAccelerator Building Blocks](#)
- [Part 4 - Intelligent Browser Referencing](#)
- [Part 5 - Image Optimization](#)
- [Part 6 - Content Reordering](#)
- [Part 7 - Tips, Tricks, & Gotchas](#)
- [Part 8 - Test Infrastructure](#)
- [Part 9 - Programmatic Performance Testing with HTTPWatch](#)
- [Part 10 - Conclusions](#)

Well there you have it, you've now got an application fully up and running behind WebAccelerator. You've read through the series, you've tweaked, tuned, toggled, fiddled and corrected to your heart's content. At least, in my little theoretical world where you've followed along through this series and used it to help leverage new and ever increasingly awesome features for your application. That's what I have, anyway. I have an app running behind WA, with all the bells and whistles turned on, and I've learned a fair bit along the way. So at the end of the day, what are my thoughts and feelings on WebAccelerator?

Tl;dr version – I like it, and it works!

Dawn Says...



I am biased, I admit it. I think that WebAccelerator is one of the greatest products out there, but then again I've been working with it for almost 8 years. After all these years I still get excited when I talk to customers and hear about the benefits they saw as a result of implementing WebAccelerator. We covered a lot of the basics and some of the advanced features of the solution but this series is just scratching the surface of what can be done and what is to come. Some of the capabilities available in WebAccelerator today that we didn't cover in the series:

- Multi-Connect or Domain Sharding – splitting the resources on an HTML across multiple domains to enable a browser to download more content simultaneously.
- Invalidations – sending triggers to invalidate content when an action is performed on an application such as commenting on a blog post.
- SPDY – How to quickly enable SPDY support for browsers that support it without changing your back end infrastructure.
- Symmetric deployments – How to create your own enterprise CDN with geographically dispersed devices.

I can't talk about the future here but stay tuned for some exciting new features in the application delivery optimization space.

Less brief version – WebAccelerator is an amazingly powerful tool that can do some pretty fantastic things to improve the responsiveness, consistency, availability and functionality of your HTTP(S) based application. It is far more powerful than I expected when I first dug into it, honestly. With that power comes some complexity, and it took a bit of time to become fully comfortable with the different features, options and configuration settings necessary to really get the most out of the product. A little time spent with an expert, some reading, a couple of videos on the web later, and I felt confident enough to wade into the fray and tackle things myself. After some time doing so, I can say that the learning curve feels steep at first, but plateaus nicely. This is a very nicely manageable product once you get the hang of it. Even while you're learning, the one-click and done features alone can be massive savings and improvements, depending on what you already have in place.

Jason Says...



This journey was collaborative from the get-go and I am grateful to Colin and Dawn for taking this effort from testbed to production and even further, for getting it all in print so we could share the work with the community. Everything short of SPDY went through the testbed, though not everything ended up in our production policy. Recall Dawn's twofold advice "art, not science | test, test, and retest" from earlier in this series? There are features that enhance user performance, features that save bandwidth, and features that save server resources. our priorities were user

performance first, bandwidth savings second, and server resources last, so our final policy didn't enable everything under the sun, just the right blend to maximize user performance without spending far too many cycles fine tuning for those last few milliseconds. Below, you'll find a few resources to help discover the right policy for your applications in your testbed.

If you want my take on it, I give WebAccelerator the nod if you're looking for a comprehensive HTTP acceleration solution. When combined with LTM (and perhaps even GTM) it becomes a near endlessly powerful tool to control and improve user experience.

Resources

- [WebAccelerator Test Policies](#)
- [iRules for Policy Switching](#)
- [HTTPWatch Test Scripts](#)
- [Results Slide Deck](#)
- [Project Acceleration Wrap-up Video](#)

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