

# **1 Mod\_perl 2 used to generate aggregated search results and contextually targeted advertisements.**

## 1.1 "Richard F. Rebel" <rrebel (at) whenu.com> exclaimed:

- Date: Mon, 23 Feb 2004 15:10:31 -0500
- Traffic: 12Million per day
- URL: <http://bidtxt.whenu.com/Search?templ=12%26kw=fitness%26num=3>

The two applications listed above are part of a suite of applications that provide bid text results to around 30 million client software installations around the globe. Roughly 12 million dynamically generated pages are created daily and the volume is ever increasing. We expect to increase to 50 million pages per day over the next 3 months.

Bid text results are provided by third parties via HTTP. For every impression, a connection to a remote HTTP server is initiated with query arguments derived from the initial request. Results are returned in XML format and normalized to match our internal DTD. The applications then perform selection based on various criteria, search and replacements. An XSL transform with a template is executed to create the resulting page which is then returned to the requesting client. The entire process usually executes in less than .3 seconds per request.

Due to the nature of the type of application and it's dependency on remote HTTP servers providing data, the wily nature of the internet, and the sheer volume of clients and requests, extensive use of timeouts and other protective measures is required.

With the exception of the use of XML::Sablotron, the entire application is written in perl running in mod\_perl2 which is running in Apache2 with the worker MPM.

Mod\_perl has proven to be reliable, fast, and easy to develop for, plus mod\_perl also has an active and supportive user base.

Interesting tidbits:

Hardware: 5 Penguin Computing Relion 140, 2x2.8ghz Xeon, 2GB RAM  
OS: Linux 2.4 kernel  
Software: Apache2+mod\_perl2, XML::Sablotron (libsablot)  
Capacity: 2000 dyn pages per second (peak).

## Table of Contents:

1	Mod_perl 2 used to generate aggregated search results and contextually targeted advertisements. . . . .	1
1.1	"Richard F. Rebel" <rrebel (at) whenu.com> exclaimed: . . . . .	2