

# Technologie Extraordinaire

mod\_perl is the power behind many of the Internet's busiest and most advanced web sites. Listed here are success stories from people using mod\_perl; also, world-wide statistics of mod\_perl usage

Last modified Thu Jan 29 08:16:35 2004 GMT

Table of Contents:

- ▶ 1. Sites Running mod\_perl  
A collection of big and small sites enjoying mod\_perl
- ▶ 2. Success Stories  
Success reports from people using mod\_perl
- ▶ 3. mod\_perl World Wide Deployment Statistics  
Discover how mod\_perl is used worldwide, through graphs and numbers.

# 1 Sites Running mod\_perl

This document contains information to give you an idea of where and more importantly, **how** mod\_perl is being used. If you have an interesting mod\_perl application, let us know about it!

Of course, the sites described on this page are not the only sites running mod\_perl. Since mod\_perl inserts its name and version into Apache's `Server` header, we are able to see others based on the Netcraft survey results.

Also remember that many services are using mod\_perl as their back-end; therefore these aren't reported by Netcraft, since the back-ends usually cannot be accessed by the Netcraft scanner. For the same reason, if you try to verify whether some of the services listed here are running mod\_perl, you may not be able to, because the back-end is not accessible. The listing is made of reports from people who develop and maintain these services, so they know.

In addition to the sites listed in this document, you might also be interested in the sites running these application toolkits/servers using mod\_perl:

- Embperl
- Apache::ASP
- AxKit
- Mason
- Slashcode

---

**Helping out Perl developers:** The first public site to run mod\_perl was PAUSE, the Perl Authors Upload Server. This was a birthday present given to PAUSE by Andreas König on August 20, 1996.

Many other sites helpful for Perl developers are using mod\_perl. Search engines like Search.cpan.org and Randy Kobes' CPAN search use mod\_perl. This is also the case for the popular Use Perl and Perl Monastery web sites. The Perl.org news server's web interface also uses mod\_perl, as well as the Perl jobs site (which is also where we get our mod\_perl jobs listing from).

**Movies/music:** Take a look at how Rob Hartill uses mod\_perl to help you find out anything and everything there is to know about movies at the Internet Movie Database. Besides using mod\_perl to speed up the interactive database queries, mod\_perl steps in during the first stage of a request, mapping URIs to cached query results if present and mapping based on language preference if the user has presented one. See the success story.

Mark A. Downing tells us: I have been running an Amateur Radio callsign database (with 800k records) on my web page for nearly two years, originally with WebSQL. But due to the lack of portability, I rewrote my scripts using sybperl. Now with mod\_perl, I have successfully cut the time to complete a lookup from nearly 5 seconds to under 1 second (It takes longer to display the data than to do the query). This was accomplished by creating persistent database connections (to sybase) using mod\_perl, and Apache is now able to establish those connections upon startup. No longer do I have to wait for my original scripts to connect and gather data.

Mojam is a new Internet music media company with the goal of attracting the largest audience of music listeners and players anywhere. Mojam is different than RollingStone or MTV because it focuses on helping new bands get their music out to the listeners by posting their show dates, music clips, and news

releases. mojam.com is 100% Apache mod\_perl running Mason to dynamically deliver all of it's pages.

**Banking:** Inteligo, the Polish internet bank, recently migrated its transactional web service (the application used by the bank clients to make different kinds of payment orders, check account balances etc) from complicated Java-based solution to the mod\_perl application. The application implements web frontend to the business services implemented by the main bank system and accessed via the bank middleware. It is worth mentioning that the application constitute the main access channel for the bank clients. Read the whole story.

**Technology news:** WIRED Digital uses mod\_perl (on linux and solaris) for several applications. It is used widely throughout HotWired, Wired News, Webmonkey and Suck.com as a replacement for CGI scripts, and to control the HotWired member pages. Mod\_perl also runs two servers that redirect requests for external pages from within WIRED sites. WIRED Digital regards mod\_perl as an important and highly valuable addition to the server development toolset, and will continue to consider mod\_perl as a strong candidate for solutions.

CMPNet, a technology information network, uses mod\_perl to generate 70% of its pages - over half a million hits a day. Our network includes TechWeb, a technology news site. We switched to mod\_perl because we couldn't stand writing and debugging NSAPI code for Netscape servers anymore. Needless to say the productivity improvement was immeasurable. Recently our company evaluated several top commercial web publishing platforms (Vignette's StoryServer, INSO's DynaBase) for a possible switchover. But in the end we stayed with our mod\_perl architecture and agreed to standardize our company's internet operations on it!

Rob Malda tells us that Slashdot - "News for Nerds", uses a combination of Perl and MySQL. Slashdot runs under mod\_perl (using SlashCode) which keeps things nice and speedy.

**Macromedia:** The creators of Flash and other interesting web technologies and programs use mod\_perl for their website.

**Adobe:** The creators of Photoshop, Illustrator, the SVG Viewer, and a lot of high-quality software packages use mod\_perl to run their high-traffic website.

**Search engine:** Scandinavia Online uses mod\_perl for the Kvasir search engine. Kvasir is Norway's most popular Internet directory.

mod\_perl scripts which search the archives of CPAN, CTAN, Linux, teTeX (a Unix TeX system), and MuPAD (a symbolic math program) are available at theory.uwinnipeg.ca. These scripts query an mSQL database via various criteria, and employ the CPAN multiplexer code to choose a nearby mirror of the archive, if available and desired.

The AtHand Network Yellow pages serve their content thanks to mod\_perl.

**Picture database:** Gerald Richter and ECOS are using mod\_perl (with Embperl) for a picture database. This contains pictures from touristic information Rheinland-Pfalz. It's intended as press information, to reduce the need of sending photographs around. You can view and search the picture via the internet or via a direct dial-in. The download is only available for valid users and via direct dial-in. Also it's possible to maintain the database via the web and insert new pictures and descriptions, change or delete them. It is

running Apache 1.3, mod\_perl, Embperl, DBI, DBD-Pg, PostgreSQL 6.21 on Linux 2.0.34.

www.stonehenge.com provides an online easy-to-navigate archive of the 150+ magazine articles written by Perl Hacker Randal Schwartz, along with an organized collection of the 20,000 pictures he's taken with his digital camera all over the world, using mod\_perl and Template Toolkit's Splash widget library for rapid development and access.

ArtToday has a collection of over 600,000 keyworded images of all types. Customers find images using keyword and category searches. They serve about 250,000 raw hits daily.

Randy Ray uses Apache+mod\_perl for his Software Configuration Management team's site within US WEST IT. About 1/3 of the data the server sends out is CGI-generated. After the conversion to mod\_perl, some existing CGI scripts running unchanged via `Apache::Registry` showed measurable speed increases of as much as 723%. All of the SCM CGI scripts use the `Image::Size` library to add `HEIGHT` and `WIDTH` attributes to `<IMG>` tags. As `Image::Size` caches the dimensions of each files as it is first read, the persistent dataspace will virtually eliminate the step of computing image sizes.

**City guide:** CitySearch.com provides online city guides for more than 100 cities worldwide, city-search.com helps people find and plan what they want to do and then lets them take action, offering local transactions such as buying event tickets and making hotel and restaurant reservations online. Its traffic exceeds 100,000,000 page views a month. Of course it's running under mod\_perl.

**Online communities:** Webpersonals uses `HTML::Embperl` running under mod\_perl to drive this nice match making site.

Alvar Freude uses mod\_perl on Assoziations-Blaster. It is a "truly interactive text network", written completely in mod\_perl. For a quick, non-technical overview take a look here.

About one year ago, it was running on Microsoft IIS with ActivePerl and some PHP, in the meantime it has been completely rewritten as an Apache module, using MySQL as a database. With this, I speeded up the execution time from ~3 Seconds to ~10 milliseconds for each Blast-Page (OK, OK, the old machine had very bad hardware, now we use only semi-bad hardware: Pentium II 350, 320 MB RAM with Soft-RAID 0 under Linux).

The blast\_engine includes links to the texts in realtime, and the statistics are created in real time: english version, (german version, with much more traffic).

The blaster uses the speed benefit of keeping the complete keyword list in memory (more then 5 MB for the german version), for the non-linear real-time linker I use a ~50 line regexp .-) The HTML files are compressed on-the-fly with `Compress::Zlib`, so we keep bandwidth (and transmission time to the users) small.

Jayme Cox explains that at Broderbund Software, we have a site running mod\_perl to keep persistent database connections open between our Apache web farm and our Oracle database server. We have a game site that checks our Oracle DB for a customers email address and lets them download additional game maps if they have registered the software. Using persistent DB connections increased the perceived response time by over 200%.

The Fantasy Sports site, which puts you in the position of manager of a group of athletes, runs mod\_perl to server its dynamic content.

**Ovulation calendar:** Conceptionstore.com's ovulation calendar runs on HTML::Embperl under mod\_perl.

**DNS provider:** Hammernode Internet, a no-cost DNS provider serving thousands of zones, uses mod\_perl to power both its dynamic web site, as well as the server end of its published API interface for client programs. "We're very satisfied with mod\_perl, and very thankful for it."

**Statistics:** Rick Mangi and Thaumaturgix, Inc. use mod\_perl to provide a method for gathering stats on web usage including a logging proxy server module for their customers.

ValueClick, the well-known advertising provider, uses mod\_perl to serve its advertisements at high speed while still maintaining good statistics of what is served.

**Authentication:** Jason Bodnar at Cox Interactive Media, explains: Right now we're using mod\_perl for authentication on two sites (soon to be a third): Austin360 and GoBig12. We started out using Netscape servers and dbms for authentication. We were writing all our demographic info to flat files but that got out of hand so we moved the demo info to Informix. This was okay but it meant keeping two separate databases (dbms or Berkely DB with username/password and Informix with demographics). Not fun. So when we switched to Apache (for performance reasons) I was able to consolidate all the info and do authentication out of Informix thanks to mod\_perl, Apache::DBI, Apache::AuthenDBI and DBD::Informix. It makes life much simpler! We're also eventually going to be running our Eats Database (list over 1400 restaurants in Austin) and our movie database with mod\_perl. I'm sure we'll find alot more uses for it in the future.

**Link between 5 different database systems:** Magirus Datentechnik GmbH is a German company employing around 200 people. Magirus has offices located in Germany, Switzerland, Italy, Austria and needs a powerful Intranet System for it's internal information flows. We're using mod\_perl to do it. Our mod\_perl applications make the link between 5 different database systems (Perl power and mod\_perl permanent database connections) to allow users to get price-lists, make offers and orders, get the status of a client, etc. A principle advantage of this kind of configuration is that we only need to install a Web Browser on the user's side. These tools are available from both Intranet or secured Internet connection. Without mod\_perl, the average response time for the top 5 applications is between 3 and 9 seconds. This is unacceptable to our end-users. With mod\_perl, the response time seems (for the end-user) to be immediate (depending of the client's speed). This is only working on a private-network and so we can't give you an address to try it. For more information, contact Philippe Froidevaux.

Information about the collection persists in an Oracle 7 database, and keyword searches happen via a custom application written using a Verity search engine. All of this is glued together using Perl. Our hardware consists of a single Sun Ultra with lots of storage (about 150GB) and an unnecessarily large monitor.

We switched to Apache/mod\_perl after becoming frustrated with the Netscape Commerce Server performance. Although I don't have hard numbers, I would estimate a factor of 2-5 times CGI performance. Using mod\_perl and Apache we've turned our "extremely loaded" server into a "comfortably loaded" server, even allowing us room for some software development. Mod\_perl saved us from having to buy another Ultra!

1 Sites Running mod\_perl

**Replacing CGI:** PBS Online is using mod\_perl to improve the speed of its heavily loaded servers, having replaced CGI for games, navigation control, and commerce.

---

Is **your** site running mod\_perl? Let us know!

# Table of Contents:

<b>Technologie Extraordinaire</b>	. . . . .	1
<b>Sites Running mod_perl</b>	. . . . .	3
1 Sites Running mod_perl	. . . . .	3