

Understanding LAMP

A hitchhikers tutorial on the Open Source community



LAMP is an acronym coined in recent years to describe a common web architecture in the open source world. At its most basic, it describes a web application built on Linux, Apache, MySQL and PHP, Perl, or Python, in reality it describes any architecture that relies on an open source operating system, open source web server, open source database, and open source programming language for its implementation.

Advantages

Cost

The most obvious advantage is cost. Since each layer in the stack is based on an open source solution, the entire stack can be implemented for the cost of development plus the cost of hardware. No piece of the stack requires software licenses.

Agility

The LAMP stack is arguably easier to maintain and expand, and quicker to adjust to business requirements.

Scalability

The LAMP architecture is horizontally scalable, meaning it grows as you add hardware to it. No single piece of the architecture is a bottleneck because each piece of the stack grows on its own, and is loosely coupled to the other pieces in the stack.

Uses

Production Customer Facing Websites

The vast majority of so-called Web 2.0 sites are built in a LAMP style. In addition, large established sites like Amazon use a LAMP like architecture built on top of open source tools.

Internal Applications

The inexpensive nature of the software, along with the relatively low cost of hardware makes LAMP an ideal architecture for internal applications.

Web Services

As more companies rely on SOAP, XML, and REST web services for communication, LAMP provides an easy way to expose

The LAMP Architecture

Linux

The base operating system for a LAMP system is typically the Linux operating system. It's also common to see FreeBSD, and increasingly Sun Solaris, especially OpenSolaris. Sun Microsystems saw the advantage of the LAMP architecture and open sourced Solaris in an effort to increase their footprint in the web services market.

Apache

Apache is the traditional web server in the LAMP stack. Increasingly, many sites are opting for Lighttpd with FastCGI. Apache, however, has a longer history, better support, and more systems administrators comfortable with its configuration and maintenance. In addition, Apache works as a application agnostic front-end, meaning it can support J2EE applications as easily as a PHP or Python application.

MySQL

MySQL is the most popular of the open source relational databases. It's fast, easy to maintain, and has a large user base. It supports clustering, replication, and has a variety of backup solutions. PostgreSQL is another open source database that's gaining traction in the market. In addition, Oracle, SQLServer, and DB2 can be accessed by a typical LAMP stack, although it's not as common.

Application Layer (Perl / PHP / Python / Ruby)

When the acronym LAMP was first created, Perl and PHP were the 'P'. Python was added later, and increasingly LAMP architecture is applied to Ruby applications. Regardless, a LAMP application assumes the use of an open source programming language, typically in the form of a framework.

Frameworks

The following frameworks have gained a significant amount of traction in the open source market. They typically only affect the Application layer (the P in LAMP), but will often dictate decisions made in the other layers of the stack.

Django (Python)

Django is a Python based framework that runs significant portions of the Washington Post Online, <http://chicagocrime.org> as well as recent work by PSC at Leapfrog Online. Django's strengths are in quickly developing custom content heavy sites. It includes an ORM (object relational mapper, a way of abstracting a relational database that makes development easier and more manageable), a templating language, database caching system, and full support for integration into the Apache web server.

Ruby on Rails (Ruby)

Ruby on Rails is a Ruby framework that runs basecamp.com, twitter.com, and is widely considered to be the current favorite "Web 2.0" framework. That said, deployment is an issue that the Rails community is just now addressing. Currently, deployment requires a significant investment in sysadmin time, and deployments are typically non-traditional. However, Rails is an excellent framework that has significantly changed the web landscape.

PHP

PHP is an embedded templating / programming language similar to ASP. Many developers consider the language not as robust as Ruby or Python, but PHP is a very capable web application language. It integrates very well with Apache and is well understood by the vast majority of Systems Administrators.

Hardware Sizing

For sizing purposes, it's reasonable to assume that a Linux server with two CPU's, four gigabytes of memory and average disk capacity can support between 100 and 150 concurrent users. This average scales proportionally, meaning the more hardware you add to the mix, the more concurrent users you can support. It's reasonable to say that a LAMP solution can handle 100 - 150 concurrent users per average size server, that two such servers can support around 300 concurrent users, etc. This report does a good job of explaining the benchmarks that support these numbers:
<http://wiki.rubyonrails.com/rails/pages/Framework+Performance>

Scaling

Scaling of LAMP applications has recently become an issue for many growing websites. Currently, there aren't many resources online for estimating sizing needs for Python and Ruby based websites.

Note: For scaling and other assistance in learning more about Open Source, contact PSC at 1-800-592-8003 or send an e-mail to info@psclistsens.com