
Open Source Microstock Agency

How Stock Photo Agency YayMicro.com was created using only open source technology

by Bjorn Sjogren, CTO at YayMicro. Oslo September 2010



BEGINNING

Two years ago, three entrepreneurs set out with a common mission; creating the best microstock agency in the world. With little resources, but lots of faith, energy and a common belief in how a microstock agency should be made, they all quit their job and started creating [yaymicro.com](#) from scratch. They literally had nothing: No founding, no code, no stock photos. Today YayMicro.com features more than one million stock photos and is commonly regarded as the easiest to use microstock agency. This is a fairy tale story showing how far the open source community has come, and how you can create something out of nothing if you just believe.

MICROSTOCK

Microstock photography is an offshoot of traditional stock photography. Microstock companies source their images via the Internet from a wide range of photographers and sell their images at a very low price compared to traditional stock photography agencies.

At YayMicro.com anyone can register as a photographer, upload their pictures and, if the photos are approved by YayMicro's reviewers, sell them to customers around the world. When an image is sold through the website, Yay receives 50% of the sales price and the photographer receives 50% in commission. In contrast to traditional stock agencies, licensing rights at Yay is super simple. Yay offers a Royalty Free license that gives the customer permission to use an image on anything the customer want, as many times as he wants, for as long as he wants. The only limitation is on reselling rights, which can also be bought, but at a higher price.

TECHNOLOGIES



Since resources were scarce, and software is expensive, the people at Yay were pretty much forced to look at free alternatives for anything they did. Microsoft, Oracle, IBM and Apple were out the window. Instead, every piece of technology was chosen from the open source community. Here is the technology YayMicro.com was built on:

Operating System - CentOS

The enterprise-class linux distribution was chosen because it has proven to be one of the most reliable and stable distributions. It's one of the most popular server distributions, meaning it is easy to find solutions and support online. Yay is currently running CentOS release 5.5 (Final) that was released in May this year.



CentOS was an easy choice for Yay, since the developers had experience from RHEL (Red Hat Enterprise Linux). The community around CentOS provides the support needed, and the functionality is just as good as RHEL. In fact, CentOS can be regarded as the free edition of the same software that makes Red Hat an enterprise solution. It is incredibly stable, efficient, and secure. It provides the same level of security as other enterprise Linux versions. All updates also coincides with a release of RH, ensuring consistent compatibility.

Database Management System - PostgreSQL



Obviously, for a microstock agency with millions of images, selecting the right DBMS is crucial. There were two clear alternatives Yay had to choose between; mySQL and PostgreSQL. Traditionally mySQL has been assumed to be faster, but less full-featured than PostgreSQL. However, this assumption is to a high extent based on history and old stereotypes, and for the most part outdated. PostgreSQL has become much faster recent years, while mySQL has implemented a number of functionalities in its last releases.

Both PostgreSQL and mySQL have their own loyal supporters, focusing on advantages in their system while talking mostly about faults and problems with the other. In the end both systems are excellent systems, and both would probably have been good choices for Yay. The choice might have come down to two factors that might or might not have been correct.

First, PostgreSQL have a rumor for scaling better with large amounts of data. This was hard to test initially, but with millions of millions of records, the database is still super responsive, even at high concurrency levels.



Secondly it might have factored in that Yay felt PostgreSQL was the more “open” of the two open source systems, a belief that was somewhat “confirmed” in 2009 when Oracle bought mySQL.

Search Engine - Apache Solr

One of the most important aspects of a successful stock photo agency is a good search engine. If users can’t find their picture in seconds, they’ll go somewhere else. Technically, searches could be done straight to the PostgreSQL database, but even though it’s super fast, it’s nowhere near what Yay needed. Consider this; each image is tagged with 50 keywords, the database holds more than a million images. Each search should take into account the relevancy of search phrase, orientation of images, how many people is in the picture, country of image, explicit images and categories. It should return a ranked list of a search, giving you exactly result from 350 to 400 if that is what you ask for. Consider thousands of people searching the website for images at the same time. Apache Solr is up for it.



Solr was created by CNET Networks in 2004, and donated to Apache in 2006, thereby going open source. Although Solr is still a young project, it has become extremely popular and proved to be a really strong search engine. YayMicro is the only microstock agency that uses this modern search engine, and as a result features the fastest search in the industry. Other users of Solr include NASA, CNET, Yellow Pages, Digg.com and SourceForge.

JavaScript Library - jQuery

JavaScript is what made web pages dynamic, and jQuery is what made JavaScript accessible. That’s a bold statement, and might be an exaggeration, but it is meant to show how important jQuery has been to the development of YayMicro.com. “jQuery: The Write Less, Do More, JavaScript Library” is their slogan, and it’s very fitting. Javascript has never been this easy.

Adding a simple line of code (importing the jQuery Library) does two things. It makes it easy to write javascript code, and it makes it possible to include tons of jQuery compatible plugins. Whenever new functionality like image previews, login boxes, button effects or image galleries have to be implemented, jQuery is there to help. It’s not an exaggeration that implementing the preview function on the page took less than three hours, and most of that time was used browsing for the right plugin. Released just four years ago in 2006, it’s already used by more than 31% of the 10.000 most visited websites. That’s no coincident.

Yay uses the latest release, 1.4.2 of jQuery. Updates are done as soon as new versions are proven to be compatible with all used plugins.

Web Application Development Framework - Stripes

Web Application Development with Java has historically been responsible for many headaches for developers who had to spend 80% of their time configuring, and only 20% of the time were spent on actual code writing. When Struts went open source in 2000, it got a little better for developers, but there was still a lot of time spent on configuration. WebWork 2 and Spring-MVC are better, but nothing really compares to how easy it is to develop Java EE applications with Stripes. Auto-discovered ActionBeans, transparent file upload capability and a great Localization system are just some of the things that make Stripes so easy to use.

On the Stripes website, they claim that: “Stripes aims to provide an experience similar to owning Apple hardware, Sony TVs and luxury German cars (without the price premium); things just work, they feel right and every now and again you get that “Oh, it does that too? Awesome!” feeling.” Stripes really does this, and for the diehard Apple fans in Yay, the Apple reference alone is reason enough to use it.

Image software - ImageMagick

An image agency needs a stable and extensive image converter. ImageMagick provides all the needs, and does so with a BSD license. All image conversions and vector graphics are done with ImageMagick and GhostScript. ImageMagick is also easily adjustable to Yay’s server environment.

ImageMagick has support for over 100 formats, tons of functionality and extensive security features, including memory and thread error detection. All the stock photos, more than one million, have in some way been handled with ImageMagick





Other technologies

To make YayMicro.com work, there are too many technologies involved to write in-depth about all of them. Here are the nuts and bolts of the open source microstock agency that is YayMicro.com.

- Eclipse

The open source toolset Eclipse is used for developing the Java EE applications that are used.

- Apache ActiveMQ

Apache ActiveMQ is a message broker that enables applications to talk to each other.

- Bucardo

Bucardo is the open source Database Synchronization system built for PostgreSQL

- Bacula

Bacula is the backup application that makes everybody at Yay sleep well.

- Apache Subversion

Subversion is the revision control system used. A more modern version of the all popular CVS.

- Wordpress

blog.yaymicro.com is running on Wordpress, (which means Yay is also using mySQL)