

Revolutionizing The Online Media Market Through Open Source Social Publishing

examiner.com

How Examiner.com Is Being Rebuilt by Migrating to Drupal 7

Company Summary

- Organization: Examiner.com
- Industry: Media
- Type of Drupal Site: Community Driven Reference
- US Traffic Rank: 118 (ComScore, Feb. 2010)

With over 42,000 content contributors, Examiner.com has enjoyed some of the fastest year-over-year growth of any online media company. Examiner.com attributes much of its success to a model of community-driven content. Specifically, Examiner.com provides localized news coverage using a deep network of localized journalists, known as Examiners. Social media innovations, such as reader discussions, have factored into the rich give-and-take between local audiences and the local Examiners, who are knowledgeable participants in their respective area of expertise. Examiner.com's success and popularity has resulted in leading online news outlets following suit with their own community-driven content.

Examiner.com has demonstrated a long-term strategic vision to leverage social media for improving online news distribution models. This commitment has directly led to the large-scale migration of the entire Examiner.com site to Open Source Drupal.

The Struggle to Keep Pace with Online Innovation

Examiner.com has grown tremendously; both in the scope of functionality and the communities it covers, with over 240 markets throughout North America. Behind the scenes, the massive production of news content has been facilitated through custom-coded Cold Fusion pages, and SQL Server databases. As with many fast-growing online ventures, Examiner.com began to face challenges caused by its own success.

We've consistently worked towards making Examiner.com a community-driven property. For our audience, we make sure that content is timely, local, highly visible and available on demand. We've been ahead of the curve in regards to social media functionality to make content timely and relevant to localized communities.

To deliver on this model, we've pushed the limits of our Cold Fusion development. We realized a limit would arrive; though custom coding served us well in the past, it's just not a feasible plan moving forward.

Michael Meyers, CTO, Clarity Digital Group

Drupal Success Stories

The Examiner.com case study is the second in a series of papers highlighting Drupal success stories. Examiner.com is not an Acquia customer and Acquia was not involved in the project to build their website.

The purpose is of these Drupal case studies is to share details on how enterprises have built beautiful, high performance, enterprise-scale websites using Drupal.

Challenges

- Custom coding made innovation difficult
- Server licensing costs were substantial
- Search Engine Visibility was not optimized

With the writing on the wall, Michael Meyers and his team began to plan for a transition to a new platform that could meet current and future requirements.

Turning to Drupal

Michael Meyers was already familiar with Drupal, through NowPublic.com, which was acquired by Clarity Digital in the Fall of 2009. There were several factors that made Drupal suitable for this project. As a mature, community-developed content management system (CMS), Drupal would eliminate the need for Examiner.com to build and maintain a custom CMS platform. In addition to Drupal core, Drupal offers the added benefit of thousands of community developed modules, which are regularly maintained and improved through the community. The prolific nature of the Drupal.org developer community is well suited to Examiner.com's interest in offering the latest online media innovations.

While a lively community is an important factor, precedent for Drupal's reliability as a viable business solution played a key role into the decision making process. Drupal has a track record of proven success, with many highly visible name-brand companies depending on Drupal for Internet-driven business models.

A plan was established to move the SQL Server and Cold Fusion site over to Drupal, with a hybridized database solution that leverages both MySQL and MongoDB.

The Challenges and Opportunities of Migration

With Examiner.com's entire business model dependent on substantial volumes of content, the management team decided to engage in careful planning, followed by swift action. The project was initiated in mid-February of 2010, with a fully content-rich beta launch planned for the beginning of May, followed by a go-live launch date at the end of July. Examiner.com contracted Cyrve to manage one of the mission critical components of relaunching Examiner.com: data migration. Drupal data migration specialists, Cyrve has provided both expert analysis and facilitated the migration of voluminous quantities of content. Moshe Weitzman, co-founder of Cyrve, helps quantify the size of this particular migration.

"On the surface, Examiner.com boasts over 1 million discreet articles – but that only tells part of the story. These pages have related content, including over 1 million comments and hundreds of thousands of slideshows using photo images.

"Specifically, each user has an abundance of profile fields, and these fields must be brought into the new system. In Drupal, users will have a new entity called profile, using the Entity API and field API fields. This gives one some sense of not just the sheer size, but the complexity of the migration." Properly mapping this data to major Drupal objects is essential for success.

Cyrve began engaging Examiner.com by assessing and understanding the database structure, including a review of each database's schema diagram (fields and relationships). Cyrve worked with Examiner.com to map data over from SQL Server to a hybridized storage model that combines MySQL and MongoDB. Cyrve also configured the conversion of character sets to UTF-8 from their native database using PHP SQL extensions implemented in the MSSQL database driver. Mike Ryan, co-founder at Cyrve, started rebuilding the Drupal migrate module to take advantage of the most exciting aspects of the project; Examiner.com would be running on Drupal 7, currently in Alpha 2 release.

Mike explains how innovations to Drupal and the migrate module are working together to make the migration process more efficient. "Drupal 7's field API has been a godsend. Drupal 7 also features an autoloader, which takes advantage of PHP5's class autoloader. With the autoloader, the migrate module's plug-in system is clean and Object-Oriented. We've been able to refine the communication process between our module and legacy data. In the past, we would have to migrate content from SQL Server, to MySQL, and then again over to Drupal. With the updates to the module, we can now migrate directly to Drupal. This has helped address one of our biggest challenges – communicating with the client about where data goes. These updates allow us to more easily build in information about source and destination fields, and manage fields that appear on the web site."

Another contribution to the migrate module is Drush functionality. Drush adds the ability to replace complex, multi-step tasks via a single command line instruction. These module updates, submitted to Drupal.org by Cyrve, are directly benefitting the migration. Specifically Drush functionality enables processing to be more directly and swiftly initiated, monitored, stopped and rolled back – with a minimum of keystrokes and error potential.

Incremental Migration

One of the most relevant updates is support for incremental migration, including the ability to query remote databases directly. "To understand the benefit of incremental migration," explains Mike Ryan, "one should be familiar with how migrations have been handled in the past. Previously, high volume content is queued up and moved over to a new environment. Webmasters and DBAs work together, but ultimately, a switch is flipped – with a new site using a new dataset. There's an inherent risk of the site performing poorly or crashing. With Examiner.com – that's obviously not an option.

"With incremental migration, we can migrate most of the content over prior to the beta launch, including all of the millions of pages, comments and photos. Then, during the beta launch we can migrate content on a weekly and daily basis. As the go-live launch date approaches, we migrate smaller and smaller chunks of data. By the time we flip the switch on the new site, the amount of content being migrated is nominal. Most of the content has already been migrated, ensuring a much more secure and reliable process."

Benefits

- Search engine and page load performance improvements
- Ease of text and interactive media publishing for the more than 33,000 contributors
- Information is more accessible
- Increased community interaction and user engagement
- Server licensing elimination
- Business operations, such as contributor compensation, can be integrated and accelerated

Examiner.com Data Profile

- Blogs: 42,870
- Articles: 1.3 million
- BlogEntryPhotos: Over a million
- BlogComments: Over 1.5 million
- Lists: 150,000
- ListElements: 420,000
- Slideshows: Over 100,000
- SlideshowImages: Over 700,000
- Tags: 374,079

Link Preservation

Another concern being addressed is link preservation; with a new organizational architecture, Examiner.com must be able to ensure that the migration makes content more accessible, without eliminating all old paths towards navigating to popular content – potentially thwarting user interaction and search engine optimization. Jeremy Andrews, from Tag1 Consulting, is coordinating efforts between various working groups involved with the move to Drupal, and explains the process for preserving link structures. “The old URLs will be stored temporarily as fields in Drupal, until they are no longer needed. This will allow us the flexibility to move to a new model that will dramatically improve search engine optimization.”

Performance Improvements

In the process of migration, Cyrve has found opportunities to improve Drupal 7 performance, while refining the Examiner.com's data. Cyrve submitted several Drupal 7 patches that were committed. Specifically, Cyrve streamlined the saving of new nodes, users and taxonomy terms. This resulted in improved performance for all Drupal 7 sites. The migration process is helping shape Drupal 7 performance, as migration puts a unique strain on Drupal Core. Moshe Weitzman likens the process to a “canary in a coal mine,” as migration can expose performance issues much faster than incremental, day-to-day processing, and the Examiner.com project has provided valuable load testing and benchmarking.

The database architecture for Examiner.com represents a unique, emerging architecture – a hybridized relationship between MySQL and MongoDB, a NoSQL database. MongoDB will store Fields, Views, Node Status, and Watchdog tables. With auto-sharding capabilities to enable database management across multiple physical servers, MongoDB and NoSQL databases are becoming an increasingly popular architectural choice for large-scale Drupal deployments. Commercial Drupal support providers, such as Acquia, are already examining how to maximize the benefits of Mongo and NoSQL databases.

Examiner.com on Drupal

With a late July go-live date just around the corner for Examiner.com's Drupal implementation, CTO Michael Myers is already considering the benefits and long-term vision of the switch. “With the modular approach of Drupal, we'll be able to add new features and functions. With the core already reviewed, tested and implemented, reviewing modules for performance and security will be relatively straightforward.

“We aren't kidding ourselves – this is an aggressive undertaking, with an even more aggressive timetable. But in the long term, we're confident that this is the right move. Open source development and support is substantially more cost effective. We'll be saving hundreds of thousands of dollars in the first year alone, just on slashing proprietary web server licenses. Our Drupal talent will be posed to take future innovations from vision to reality much faster than we could in the past.”



Even more compelling than Examiner.com's rebirth as a Drupal powered media platform, is how Clarity Media Group intends to leverage its Drupal solution as a game-changer for the entire online news media market. The process of migration is being carefully planned and documented, and the Examiner.com Drupal sites will form the basis of a set of best practices. Once finished and refined, Examiner.com will take the solution to market, providing other online news sources with an easy-to-deploy solution that will enable better features for users, and more content sharing for distributors. Through Open Source Drupal, Examiner.com is creating a new revenue model based upon revolutionizing news distribution through social publishing.

About Examiner.com

Launched in April 2008, Examiner.com serves 240 markets across North America and is the insider source for everything local. Examiner.com feeds the passion the local community has for its favorite interests, activities, and establishments by connecting them with credible and informed contributors who write and share information with the passion and insights only a local insider can provide. Examiner.com is a division of the Clarity Digital Group, LLC, wholly owned by The Anschutz Company, a Denver-based investment company with a broad array of assets in print and digital media; live sports and entertainment; hospitality; film production and exhibition; wind energy development and transmission; as well as ranching and oil/gas exploration. For more information, visit Examiner.com.

About Tag1 Consulting

Tag1 Consulting, Inc. is a distinguished professional consulting company headquartered in sunny Florida, with an international presence providing computer consulting services worldwide. We focus on performance and scalability consulting of the entire LAMP stack, specializing on Drupal performance.

Examiner.com, on Drupal's Cost Savings

"We're confident that this is the right move. Open source development and support is substantially more cost effective. We'll be saving hundreds of thousands of dollars in the first year alone, just on slashing proprietary web server licenses."

– Michael Myers, CTO
at Clarity Digital
Group, which included
Examiner.com and
NowPublic.com



About Cyrve

Cyrve specializes in migrating your data into Drupal. We've migrated from MS SQL Server, Oracle, MySQL, Wordpress, Moveable Type, NITF, flat HTML files, and on and on. This is painstaking work, with emphasis on pain. Just hand us your source data and destination database and we'll move it for you.

About Acquia

Acquia helps organizations of all sizes build social publishing websites quickly, easily and with a lower total cost of ownership by leveraging Drupal, the Open Source social publishing platform that blends content and community. Our products, services and support enable companies to leverage the power, technical innovation and economic value of Drupal while simplifying the experience, removing the complexity and minimizing the risk. Please visit: <http://acquia.com>. Download Acquia Drupal, the completely free Drupal distribution, at <http://acquia.com/downloads>.

Contact Acquia

Acquia, Inc.

150 Presidential Way
Suite 310
Woburn, MA 01801

<http://acquia.com>

Phone:

888.9.ACQUIA
+1.978.296.5250

Email:

sales@acquia.com