



Get a Grip on Hosting Costs for Your High Volume Website

Improve performance, scalability and availability, while reducing business risk and costs with the Drupal Open Source social publishing platform and Acquia Hosting services

Executive Summary

You've built a fantastic Drupal website. Traffic growth charts are encouraging. Conversion rates are above industry averages. Lead numbers and revenue from the web site are growing faster than forecast.

And yet the pressure remains high to reduce costs and improve the profitability of operations wherever possible in your company. Though you have met or exceeded your commitments to the business, the CEO and the Board of Directors still want more.

Perhaps your budget cycle is about to start up, or your hosting contract is up for renewal. You need to ensure you have a grip on the total cost of hosting your website, and to recommend alternative approaches which will cut those costs while improving service levels (performance, scalability and availability).

You know that success on the web doesn't come without cost. But there are significant opportunities to dramatically reduce those costs. You can deliver dynamic, highly interactive "social" websites, and handle volumes of millions of page views per month and up. And you can do so with high performance, 100% availability, at a fraction of the cost most companies are paying today.

There are significant infrastructure costs associated with hosting a high volume website. Whether those costs are carried internally through hardware and staff in your datacenter, or through outsourced managed hosting solutions – the numbers add up quickly. And these options present two fundamental questions: 1) is it truly core to your strategy to be in the website hosting business? And 2) do third-party hosting services have the expertise necessary to deploy, manage and tune your specific site applications?

In this paper, we will present a plan to achieve significant reduction in the total cost of your website operation, without sacrificing quality, availability, or agility. By partnering with Acquia for enterprise-grade managed hosting services, companies can reduce the web operations costs by over 90%.

A Million Page Views a Month: Architecture and Cost Drivers of a High-volume Website

A web infrastructure designed to manage very high traffic volume, whether managed on-site, co-located with an ISP, managed off-site by an ISP, or cloud-based, has a generic architecture which we can use to identify where the key cost drivers are as we prepare to discuss alternative approaches in greater depth.

Let's dispense with content management and social publishing software license and maintenance fees right away: this is the first place to look for significant reduction in your total web infrastructure costs. Rather than repeat the analysis here, please refer to "TCO for Open Source Social Publishing,"¹ Acquia's report which illustrates how you can move to the Drupal social publishing platform, obtain enterprise-grade technical support from the world's best-known Drupal experts, and reduce costs by 89% to 98% compared to many proprietary web publishing solutions – without sacrificing functionality or performance.

Next, let's look at the "moving parts" of hardware. The workhorses of the website are the web servers, and database servers which handle the execution of application logic, retrieval of information, and presentation to users in the browser. As site traffic grows, and web transactions and interactions become more critical to the business, the infrastructure needs to expand in order to:

- Meet user demand in the event of a spike in traffic ("the Digg Effect"²) or continuous organic growth of site activity. Web servers and database servers may need to be "clustered" so that these sudden or sustained loads can be handled without losing valuable site traffic.

1 <http://acquia.com/community/resources/library/tco-open-source-social-publishing>

2 http://en.wikipedia.org/wiki/Digg_effect

- Ensure that page views are evenly distributed across available computing power in the clusters (using load balancers);
- Ensure that query and transaction load can be fulfilled (via high-performance, redundant high-availability database clusters);
- Enable recovery of critical site data in the event of an unforeseen disaster anywhere in this now complex and costly website infrastructure (entire system backup services and storage).

Servers suitable for web, application, and database server use in this infrastructure, range in price from \$1000 for a low-end unit to \$6000 and up for a high-performance blade server³. Storage array pricing varies similarly, with few enterprise-quality systems available below \$8000 per 12-bay unit⁴.

For each of these technologies, there are specific skills required in order to install, operate, and maintain this hardware and software. As examples, according to salary.com, a web security administrator earns an average salary of \$93,000⁵; system administrators with four to five years experience earn on average \$114,000⁶.

However, if you genuinely require 99.95% availability for your site: you will need resources to support load testing of your website in a staging environment; you will require 24x7 monitoring and emergency response; you will need performance tuning and expertise to handle security patches and software updates. As you can see, the costs mount up rapidly.

The requisite site infrastructure (including staff) to implement a robust solution could easily exceed \$250,000 in the first year. In this abstract example, however, the hardware and staff are likely to be considerably underutilized – as web activity and traffic tends to have peaks and valleys, and both hardware and humans can be “on the bench.”

In response to this dubious set of costs, a variety of third-party hosting solutions have been made available from thousands of service providers.

3 <http://www.dell.com/poweredge>

4 <http://www.cdw.com>

5 <http://www.salary.com>, web security admin

6 Ibid.

Hosting Options: An Overview of Managed and Cloud-based Hosting Services

The first step toward optimizing the cost and control balance in web hosting is acceptance of the fact that expertise in this area is not core to most business strategies. Granted, there are organizations of such size and resource wealth that they manage their own web infrastructure – in fact some large companies generate their own electrical power and manage their own telephone systems. But for 99% of businesses, web hosting should be considered a utility service, available from many providers, who have expertise and capital infrastructure focused specifically on this problem area.

Managed Hosting

With managed hosting offerings, the customer relinquishes some control compared to in-house, or Do-It-Yourself (DIY) hosting, but in exchange also has the opportunity to reduce certain costs and improve service level. With managed hosting, the customer does not have complete control and access to the server hardware, as they would with self-hosted, dedicated, or collocation approaches. In tech-speak, the customer gives up “root” access (Linux) or “administrator” access (Windows). They retain control over what applications execute in that environment (their content management system and social publishing software, for example).

The customer gives up flexibility to some extent – they cannot make configuration tweaks to the lowest level of the server setup – but they gain in other areas.

With managed hosting, the service provider can focus on monitoring performance, ensuring underlying software (operating system, web server, database) is up to date with patches and fixes, and attend to security threats. They can offer upgrades to their customers, should growth in site volume trigger the need for more web server or database server capacity (through clustering, for example), or load balancing. Customers reduce costs associated with system administration (the hosting provider now takes care of these headaches), and gain a website infrastructure which is secure, “always on,” and up to date with the latest operating system and middleware technology.

Managed hosting draws a hard line at the web application, however. That is, while the hosting provider can make sure that Linux, the Apache servers, the MySQL databases and PHP are all running – they provide no support for the application that is running atop this infrastructure. Whether that is an Open Source social publishing platform like Drupal, or other software for web content management, CRM, e-commerce – responsibility for these applications all rests with the user. In-house staff, or consultants with the appropriate software expertise, must be available 24x7 should there be application problems

Cloud Hosting

Cloud computing has created a major breakthrough in web hosting. During the 1990s, as internet access and usage expanded exponentially, companies sprinted to get more and more content on the web and to build out their e-commerce capabilities. In the process, an astonishing inefficiency emerged. Companies and their hosting partners designed web infrastructure for the occasional (one might even say rare) spike in activity that might occur in a given business cycle. Preparing for the few days or weeks associated with “the big holiday rush,” compute resources operated at as low as 10% of capacity at all other times.⁷

Cloud computing and cloud hosting services remove this market inefficiency by providing dynamic provisioning of compute resources. Users can configure their servers for average traffic volumes, and then dynamically expand the resources supporting their web site in the cloud as and when needed. “The Cloud” can deliver this granularity of service by aggregating the computing needs or large pools of users in secure, redundant datacenters offering business continuity and disaster recovery service quality. Using massive computing horsepower and virtualization technology, cloud providers can easily improve the performance and availability over the entire pool of their customers.

However, like managed hosting options, standard cloud hosting products draw the same hard line between the hardware and basic web infrastructure (the LAMP stack), and the applications that are driving business on the web. The user remains solely responsible for 24x7 service coverage of the web applications themselves.

Sunshine State News and the Perils of a Successful Scoop

Sunshine State News is an online media start-up blending rich media and local relevancy to cover business and political activity in Florida. In May, 2010, Sunshine State News reporter Lane Wright broke a now-infamous story about a Florida State Senator who was caught on video looking at pornographic content on the Senate floor using his State-issued laptop computer.

The story spread like wildfire – placing unanticipated and enormous load on the burgeoning startup’s Drupal infrastructure. Working with Acquia for both support and hosting, Sunshine State news was able to flexibly expand their cloud-hosted environment, and meet the demand within 30 minutes.

“In our excitement, we hadn’t realized that we were exposing our site to an exponential increase in traffic. Acquia took it all in stride, and within a half-hour, we were operating at peak performance.”

- Lane Wright, Sunshine State News

⁷ “ Like most computer networks, Amazon’s uses as little as 10% of its capacity at any one time just to leave room for occasional spikes.” Jeff Bezos’ Risky Bet, Business Week, November 13, 2006

Mind the Gap: Tradeoffs with Generic Hosting Services

Though hosting offerings are evolving in the right direction for high-volume website owners, there remains an “expertise gap” between the customer and the hosting provider. Here’s why.

The Prevalent Web “Stacks”

The “common denominator” of hosting platform stacks, by and large, extends from the network and system hardware “up the stack” to the operating system and middleware. The Open Source stack is known as “LAMP:” it includes Linux (OS), Apache (web server), MySQL (database) and PHP (development environment). The alternative stack from Microsoft includes: Windows (OS), IIS (web server), SQL Server (database) and .net (development environment).

Standardization on these two platforms has enabled great efficiency (via competition) to emerge in the web hosting market. Millions of applications and sites have been developed on these stacks, and can be freely moved from one host to another with great portability.

Still No Panacea

However, even cloud hosting, by itself, is no panacea. Individual organizations are free to build as simple or complex an application environment on these standard web applications stacks as their imaginations and budgets permit.

Each application – whether deployed on LAMP or Windows stacks – has its own design center and architecture. Each application (licensed, Open Source or custom-built in house) has been tweaked and tuned to use CPU, memory, and the database according to its specific needs.

Furthermore, it is not uncommon for a website to be the front end for many applications, such as a content management system, customer support forum, e-commerce platform, blogging platform, or customer relationship management package. What happens when problems occur?

Dealing with Performance Problems, Downtime, Service Outages

Support and service from the hosting provider is available to ensure proper configuration, operation, performance and security of the generic stack components (OS, web server, database) – but they are often flat-out refuse to address the needs of a specific application or component running on that generic stack in the event of an outage or performance problem.

As a result, users are left to provide that expertise on their own staff, or to leverage the support available from their vendors. This often leads to finger pointing and standoffs as users try to pinpoint accountability for problems among the host, their own IT staff, and vendor technical support teams – and always leads to high costs.

Currently, Drupal experts with three to five years' experience can command salaries between \$75k and \$100k, with senior and principal roles earning over \$120k. Conversely, daily rates for Drupal experts are in the \$1500 - \$2000 per day range. Either way, a plan that offers 24x7 coverage is going to require a \$200k+ annual investment.

A Match Made in the Cloud: Acquia Hosting – Optimized, Managed Hosting Service for Dynamic Drupal Websites

The adoption of Drupal as a strategic platform for social websites has been astonishing. Drupal's online community, Drupal.org, has over 600,000 members; downloads regularly exceed 250,000 per month; there are at least 400,000 active, live web sites built on Drupal;⁸ and there are over 4500 extensions to Drupal core functionality, known as Drupal modules.

More and more organizations are seeing the productivity benefits driven by Drupal's breadth of functionality.⁹ And now they can realize compelling total cost of ownership benefits thanks to new offerings from Acquia.

Chapter Three: An Interim Step for the DIY or Managed Hosting Drupal Site

By nature, the Drupal community is large and entrepreneurial. Acquia Enterprise Select partner Chapter Three, based in San Francisco, CA, provides professional Drupal consulting, design, development, online marketing, and training services.²² They have leveraged that service experience to create a pre-built Drupal hosting solution called Pantheon. With Pantheon, Drupal administrators can download and deploy a pre-configured single-server Drupal hosting stack that will boost performance to as much as 2,000+ pages per second.²³

Pantheon includes Open Source technologies such as Varnish, Memcache, and the high-performance Drupal distribution, Pressflow.²⁴ Because it is an Open Source offering, users of Pantheon must provide the management and administration skills necessary to manage and scale out this solution as their site grows over time. Also, while Pantheon addresses performance, your IT department of website host will retain responsibility for high availability and disaster recovery strategies.

8 This is a conservative number; since Drupal is open source, not all sites are 'registered,' and go undetected.

9 "The Surest Path to Social Publishing," Acquia, 2009, <http://acquia.com/resources/whitepapers>

22 <http://www.chapterthree.com/>

23 <http://getpantheon.com/>

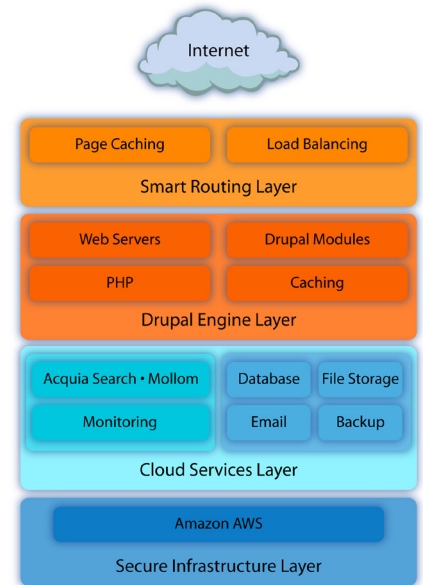
24 <http://getpantheon.com/mercury/technical-info>

Acquia recently launched Acquia Hosting¹⁰, a highly available cloud-based hosting platform tuned for Drupal performance and scalability.¹¹ The combination is a cost and productivity tour de force for Drupal users.

Acquia Hosting, the Drupal-optimized Managed Hosting Service, offers high-performance and fully managed hosting, optimized for Drupal and simplified for users.

Acquia Hosting provides:

- high-performance routing and caching to meet high volume throughput needs
- highly-tuned Drupal-specific services ensuring optimal performance of LAMP resources, backed by a capacity reserve for dynamic deployment as needed
- automated control of cloud services for high availability including application monitoring and infrastructure
- a services layer architected for high performance and high availability of database, email, and Acquia Search¹² services
- and the highest levels of security, fault-tolerance and operational controls possible in the cloud¹³



Acquia's Hosting Features

Optimized for Drupal Scalability and Performance

Acquia Hosting delivers state-of-the-art performance and scalability for dynamic Drupal-based website by offering both technology and expertise and each of four layers of service.

¹⁰ <http://acquia.com/about-us/newsroom/press-releases/acquia-provide-high-availability-managed-hosting-enterprise-drupal->

¹¹ <http://acquia.com/products-services/acquia-hosting>

¹² <http://acquia.com/products-services/acquia-search>

¹³ For more details on Acquia Hosting, see [our web site](#).

Smart Routing Layer

Beginning at the front end, where HTTP request first reach your website infrastructure, Acquia Hosting offers page caching and load balancing services. At this time¹⁴, Acquia Hosting is using Varnish¹⁵ caching technology to deliver this service.

For dynamic load balancing across the Acquia hosting grid of servers, Acquia has chosen nginx¹⁶ to distribute HTTP request load across web servers.

Together, these technologies ensure that Acquia Hosting sites can serve tens of thousands of web pages per second, thus ensuring high performance at extremely high spikes in site activity.

Drupal Engine Layer

The Drupal Engine Layer delivers the bulk of Drupal's dynamic, social publishing functionality, and is optimized aggressively using several technologies so that users can add and reduce their hardware resources to handle changing site traffic conditions and benefit from elasticity on the Acquia Hosting infrastructure.

In this Engine Layer, the Apache web servers, PHP & Drupal modules all live. Any combination of Drupal core, custom, or contributed modules is supported. Another form of caching – OpCode caching – is implemented at this level as well.

Acquia's implementation is laser-focused on Drupal site performance. This layer is configured with a maximum of memory allocated for PHP execution. OpCode caching, intelligently applied, accelerates PHP performance by delivering precompiled scripts, rather than requiring parsing and compiling of PHP with each request.

Acquia Hosting supports Subversion (SVN) repositories, the Apache revision control system, for developer workflow¹⁷. A staging environment is also provided so that final testing can be executed in parallel with production, and when it's time to "go live," SVN ensures that all of your production servers are properly updated with the latest revisions.

14 Acquia is always evaluating alternative technologies in search of the highest performance, most reliable solution, and may in the future offer alternatives to the technology components offered today as part of Acquia Hosting.

15 <http://www.varnish-cache.org>

16 <http://wiki.nginx.org/Main>

17 Subversion became a top-level Apache project in February, 2010

Cloud Services Layer

In the Cloud Services Layer, Acquia has assembled a set of back-end resources to further boost your site performance and reliability.

Acquia's MySQL database implementation leverages Drupal-specific expertise to identify and tune Drupal-specific database interactions. In addition, the MySQL infrastructure is configured for master-master replication, ensuring redundancy in case of any database system or hardware failure.

Furthermore, Acquia Hosting provides hot backup: in addition to nightly backups (stored on geographically distinct data centers if required for disaster recovery) of both code and databases, Acquia captures four-hour snapshots from production servers to cover any gaps in partially-complete backup jobs.

Acquia Hosting builds on Acquia Network support services – such as Mollom for spam blocking and content protection. Acquia Search services are also hosted on Acquia Hosting infrastructure, offering Acquia Hosting customers higher performance access to these robust and high performance search capabilities.

Finally, Acquia hosting provides Exim¹⁸ email services to round out the Cloud Services layer¹⁹.

To Acquia Network's standard monitoring services, Acquia Hosting adds proactive monitoring – site monitoring services from Acquia's 24x7 team of Drupal and hosting experts. Whether you require the deployment of additional hardware resources to handle a sudden spike, or performance tuning advice to make adjustments prior to expanding your server footprint, Acquia is there with the manpower and the experience, around the clock. In the end, the user gains extremely fine control over capacity planning – a critical ability in today's budget-conscious environment.

Secure Infrastructure Layer

Acquia Hosting is currently available on the Amazon Web Services (AWS) secure and fault-tolerant infrastructures.

¹⁸ <http://www.exim.org>

¹⁹ Exim provides adequate performance for moderate volume email delivery loads; users requiring extremely high volume email delivery should consult with Acquia on the choice of a volume email service provider.

Acquia Hosting

Acquia experts manage the operating system (OS) and OS updates, Amazon provides the balance of the underlying infrastructure. AWS currently boasts SAS70 type II certification – ensuring the highest standard of security and operational controls are in place.²⁰ AWS also offers 99.95% uptime: among the highest available on the internet today.

Acquia Hosting also offers SSL certificate services, so that data can be encrypted if necessary between your site's servers and your users.

Why choose Acquia Hosting?

No More Finger Pointing

First, organizations running Drupal-based sites with Acquia Hosting achieve instant simplification of their technical support process. They have one vendor to deal with when they need support for Drupal platform functionality from Acquia Network, or hit speed-bumps at lower levels of the website hosting infrastructure.

Tuned Specifically for Drupal

Acquia Hosting only supports Drupal sites. Acquia hosting professionals are constantly advancing their expertise in how Drupal deployments can best benefit from tuning to the LAMP stack, building out high-availability server and database configurations, and more. There is no greater concentration of Drupal administration expertise in the world.

Happy Site Visitors

Site visitors will experience high performance, high availability – and the breadth of social publishing functionality that only Drupal sites can provide. With ever-increasing business value being tied to customer satisfaction and “stickiness” to an organization's website, Drupal social publishing plus Acquia Hosting equals breakthrough business success on the web.

Understanding the Value of Acquia Hosting and Support

“Acquia is staffed with very knowledgeable Drupal experts. With previous Drupal support providers, troubleshooting was limited. It was based more on trial and error guesswork, and not a true understanding of the underlying systems. Identifying the source of problems took a long time; meanwhile we were suffering increasingly frequent episodes of website instability.

“Fortunately for us, Acquia has demonstrated superior expertise in a very simple way - they excel at discovering and solving problems almost immediately... It's as if Acquia support engineers have an instinctual understanding of what is causing problems; they would return with effective solutions in record time, sometimes even before we had identified the problem.”

- - Robert Wise, System Administrator, Mother Jones

²⁰ Acquia is always evaluating other cloud technologies for Acquia Hosting, and may support additional providers (beyond Amazon) in the future.

Dramatically Lower Drupal Operations Costs

Acquia Hosting is available with an Acquia Network Enterprise subscription, with pricing starting at \$1,300 per month. Thus, 24x7 support, Drupal best practices guidance, proactive performance monitoring, 100% availability through failover, disaster recovery – can be had from Acquia – the Drupal experts – starting at \$16,000 per year. This represents a 92% cost savings just on staffing the necessary system administration and application skills in-house.²¹

Conclusion

The never-ending drive for greater productivity in business is inevitable. The push for productivity creates innovation in technology and fuels the demand for ever-lower business operating costs.

The web and new web solutions like the Drupal social publishing platform continue to grow and bear more of the marketing and business development load for companies. As a result, your website and web hosting environment will continue to “attract attention” from executives. You will need to demonstrate your web properties’ impact on the business, and to justify the costs associated with it.

The variety and sophistication of hosting offerings has evolved rapidly, and continue to advance. And while they offer companies a variety of options for their web sites, there remains a critical and hard “line in the sand” between lower level website infrastructure (operating system, database, hardware, routers, firewalls), and the web applications (like Drupal) which actually engage users with your organizations.

With Acquia Hosting, that line in the sand is wiped away, and companies that have chosen Drupal for their social publishing needs now have an extremely cost-effective hosting solution which brings Drupal expertise together with enterprise-grade web infrastructure to provide a truly unified offering for high volume, websites – and a truly dynamic, engaging experience for site visitors. The combination: competitive advantage.

About Acquia

Acquia empowers enterprises with the open-source content-management system Drupal. Co-founded by Drupal’s creator in 2007, Acquia helps customers manage their growth and scale their online properties with confidence. Acquia’s software, consultation, cloud infrastructure, and services enable companies to realize the full power of Drupal while minimizing risk, as it’s done for Examiner.com, Al Jazeera, and over 700 others. See who’s using Drupal at <http://showcase.acquia.com> and learn more at <http://acquia.com>.

© Copyright 2011, Acquia, Inc.

Acquia, Inc.
25 Corporate Drive, 4th Floor
Burlington, MA 01803
USA

www.acquia.com
sales@acquia.com

+1.781.238.8600

²¹ Based on conservative cost estimate of \$180k, or 1.5 full time equivalent staff at \$120k annually.