

# Achieving Zero Downtime and Accelerating Performance for Drupal

## Executive Summary

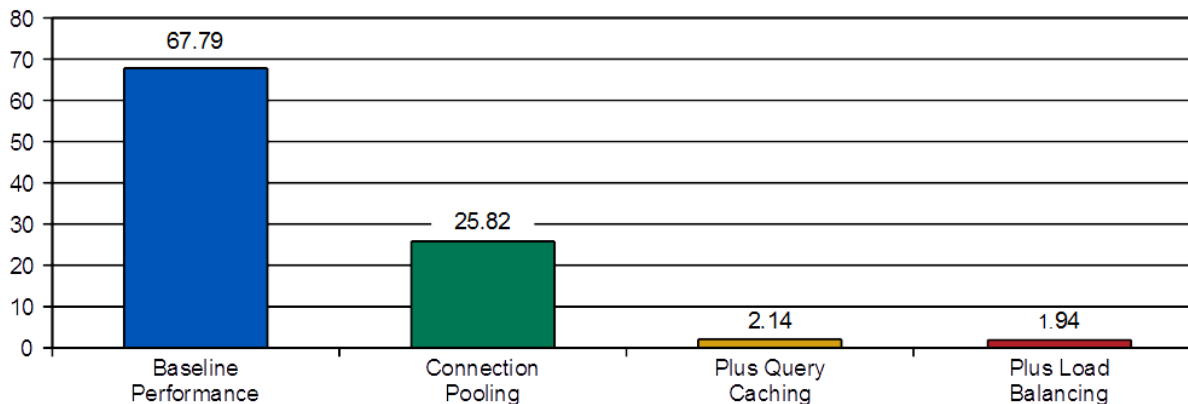
Drupal serves more web pages per capita than any other content management system (CMS). Powering high-traffic sites such as NASA, NBC, and Twitter, Drupal's highly detailed administrative and management controls are popular with those running websites for government and security-conscious organizations.

Traditionally, organizations have directly modified applications to improve database performance and uptime. For Drupal, that scenario is challenging for many reasons, including keeping up with frequent security-critical platform updates. Code-level application tuning consumes considerable resources: for every new security update, one needs to re-certify these custom tunings.

Database load balancing management software by ScaleArc provides a compelling alternative. ScaleArc deploys transparently between Drupal and MySQL to improve performance and availability. This architecture abstracts the application from the database, making it possible for system administrators to take servers offline, patch them, and "rotate" them back into service – all while the site is still operating.

ScaleArc has developed database load balancing software for MySQL, SQL Server, and Oracle database servers. Deployed in conjunction with MySQL, the ScaleArc software instantly improves Drupal performance, allowing scaling to support millions of customers without any changes to the platform.

Testing shows a reduction of average query response time from nearly 68 milliseconds (ms) to under 2 ms as successively more ScaleArc features are enabled:



Providing such functionality and response time improvements would require modifying Drupal to add connection pooling, load balancing, and query caching code. But ScaleArc software enables these features in-line, allowing organizations to use Drupal “out of the box” and gain over an order of magnitude improvement in application database performance – in a matter of a few hours.

This application note details how organizations can leverage the ScaleArc load balancing software to quickly and easily improve Drupal performance.

## Drupal Overview

Drupal is a popular website content management system (CMS), especially for high-traffic web sites. Serving 75% more page views per capita than WordPress,<sup>1</sup> individual sites running Drupal serve more page views than individual sites running any other CMS.<sup>2</sup> Thus Drupal is the de-facto standard for enterprise-grade web content.

Drupal is written in PHP to run on either MySQL or Postgres database servers. This combination provides a capable, expandable, and feature-rich platform for hosting complex, dynamic sites responsive to the needs of end customers and site administrators.

## Drupal Challenges

Although Drupal presents a simple interface that shields users from its details, optimizing it for improved application performance is difficult. Among the challenges:

- **Architectural complexity:** The web server, MySQL database, and Drupal engine all work together to deliver a content-rich experience to the end customer. Each server has its own nuances and presents its own optimization opportunity.
- **Lack of performance insights:** With about 75 tables in the database, architects and developers must be aware of dozens of factors that affect performance and often lack the tools needed to determine the specific bottlenecks for a given Drupal site.
- **Resources needed for database tuning:** At the database layer, optimizing MySQL requires specialized development. Improving performance via database replication requires patching Drupal to selectively issue read and write requests to different servers depending on the replication architecture,<sup>3</sup> diverting valuable engineering resources.

## Database Load Balancing Software – The Simple Alternative

Just as web load balancing software abstracts applications from web servers and optimizes HTTP performance, database load balancing software improves SQL application availability and performance without changing the application. This software inserts transparently between applications and SQL databases – enabling high-availability (HA) and auto-failover

---

<sup>1</sup> Based on the [OpenSourceCMS](#) statistical sample.

<sup>2</sup> Drupal Market Position, March 2014, [W3Techs](#).

<sup>3</sup> Using Replication for Scale-Out, Section 16.3.3, MySQL 5.6 Reference Manual.

capabilities, increasing application response time, and providing real-time SQL traffic analytics.

ScaleArc offers database load balancing software for MySQL, SQL Server, and Oracle databases. Leveraging this software in MySQL environments offers an ideal solution for improving Drupal performance. Among its capabilities, the ScaleArc software provides functionality such as:

- **Read/write split:** ScaleArc automatically directs read requests to replication slaves and write requests to replication masters. No change within Drupal itself is required. This approach allows you to use your hot spare slave to serve read traffic, using otherwise idle capacity and improving performance.
- **Load balancing:** ScaleArc automatically balances SQL traffic among a pool of database servers using a dynamic load balancing algorithm that sends more requests to the server that is most likely to respond the fastest. Drupal does not need to track how busy (or free) various servers are within the pool, and geo-aware load balancing occurs automatically.
- **Automatic failover:** ScaleArc can automatically detect server failures and either redirect load to a secondary master or promote a standby slave to the master role. It does so in a transactionally consistent manner, keeping your data safe.
- **Connection pooling:** ScaleArc automatically pools requests from multiple Drupal sites such that database connections are used in the most efficient manner possible. As a result, you can add many more web servers without having to worry about running out of database connections.
- **Query result caching:** ScaleArc automatically caches SQL queries on a per-user, per-query basis. Repeated requests for the same page, such as a static “about us” page, may be subsequently retrieved from memory cache rather than disk. Internal testing shows that caching about a dozen common Drupal queries for even one minute provides a dozen-fold performance boost.
- **Full SQL logs:** ScaleArc automatically logs all successful and exceptional activity between Drupal and a collection of MySQL servers to offer a single point for data collection.
- **Real-time analytics:** The ScaleArc software includes visual graphing tools to help administrators and developers analyze site operations dynamically in real time.

## Benefits of ScaleArc Software

The ScaleArc software offers site administrators and developers the following benefits:

- **Immediate performance gains.** The combination of read/write split, load balancing, connection pooling, and caching traffic between Drupal and MySQL offers an immediate performance improvement, well beyond that available with a single-point optimization approach.
- **Application transparency.** Implementing the ScaleArc software requires only administrative changes. No code change is required. Once MySQL server replication is configured, optimizing another Drupal site requires only a connection reference change.
- **Development and operations de-coupling.** Administrators may determine which database servers are masters and allocate servers based on seasonal performance load without involving Drupal. Developers can use the simpler single-threaded design pattern even if administrators find a more complex architecture is needed.
- **High availability.** ScaleArc offers a fault-tolerant, high-availability mode that can provide 100% uptime. This capability is crucial for 24x7 website operation.
- **Reduced cost.** Since ScaleArc optimizes performance at the SQL query level, it is effective across the entire database infrastructure with all storage engines simultaneously. Site administrators also save money because they don't need to invest in versions of MySQL that perform thread pooling.
- **Analytics and logging.** Detailed reporting provides up-to-the-minute problem and status metrics, going well beyond simple "go/no go" testing. This real-time data provides proactive insight for avoiding congestion problems before they happen.

## See the Gains in your Environment

To get a better feel for how the ScaleArc software works, you can watch a [brief animation](#) on our website. The best way to see what ScaleArc can do for your Drupal environment is to try it out. You can request a [free trial](#) on our website – deployment is simple, and you'll be amazed by the results.

Also check out our detailed [Drupal Deployment Guide](#). It provides step-by-step instructions for deploying ScaleArc to improve availability and performance on Drupal.



2901 Tasman Drive, Suite 205  
Santa Clara, CA 95054  
Phone: 1-408-780-2040  
Fax: 1-408-427-3748  
[www.scalearc.com](http://www.scalearc.com)



ScaleArc is the leading provider of database load balancing software. The ScaleArc software inserts transparently between applications and databases, creating an agile data tier that provides continuous availability and increased performance for all apps. With ScaleArc, enterprises also gain instant database scalability and a new level of real-time visibility for their application environments, both on prem and in the cloud. Learn more about ScaleArc, our customers, and our partners at [www.ScaleArc.com](http://www.ScaleArc.com).

© 2015 ScaleArc. All Rights Reserved. ScaleArc and the ScaleArc logo are trademarks or registered trademarks of ScaleArc in the United States and other countries. All brand names, product names, or trademarks belong to their respective holders.