

Percona Server with XtraDB Puts Flash in ideeli's Flash Sales

A Percona Case Study




Abstract

[ideeli](#) is a members-only online shopping club that offers time-limited “flash sales” on designer brands. As traffic increased, ideeli’s MySQL database experienced serious performance problems at peak load. Simple queries ran for orders of magnitude longer than they should have. The solution was to upgrade to [Percona Server with XtraDB](#). After upgrading, ideeli has enjoyed higher throughput and stable performance, even though the application has become busier and the traffic to the database has nearly doubled.

The ideeli Application and Architecture

ideeli is one of the first companies to build a very large-scale Web application on the Ruby on Rails platform, and pioneered many scaling techniques that are now common practice for Ruby on Rails applications. Founded in 2006 and based in New York, ideeli has grown rapidly, even during times of economic downturn. The ideeli application handles an unusual and demanding load pattern: it is all but idle much of the time, but when a time-limited event begins, traffic increases to many times the off-peak load. “We experience a self-created denial of service attack every day,” jokes Mark Uhrmacher, ideeli’s founder and CTO. The following graph is a typical day of ideeli’s traffic, with the peak traffic reaching as much as 40 times the off-peak.



“Since our first shopping event in June 2007, we have grown at an extraordinary rate, and today we’re a 9-figure business,” Uhrmacher continues. That’s a lot of growth in three and a half years!” Due to this highly dynamic usage and continued rapid scaling needs, ideeli hosts its application in the Amazon AWS cloud, and spins up EC2 servers as needed to handle the load. The ideeli application is approaching 100 EC2 instances at the time of writing. Traffic is load-balanced by HAProxy.

The application uses memcached, which is such a core piece of the infrastructure that ideeli maintains special code modifications to the Ruby On Rails framework to handle cache expirations smoothly. In a high-traffic environment, a cache expiration can

potentially cause a spike of requests to the back-end servers, especially the database server. The ideeli team’s solution to this is a combination of out-of-band cache warming processes refreshing the cache at four times the actual TTL (time-to-live) rate, and a parallel set of cache keys to coordinate processing and prevent a “dog-pile effect” if a cache key expires. The lack of persistence in memcached is an issue, and ideeli is considering alternatives.

At the database tier, ideeli uses a master-master architecture, with only one node actively accepting writes at any time. The other node is on standby in the event of a failure. Prior to the engagement with Percona, the database server version was MySQL 5.1.37-1ubuntu5. This was upgraded to Percona Server with XtraDB version 5.1.43. The database servers are running on EC2 high-memory quadruple extra large instances with 68.4GB of memory and 8 virtual processors. At the time of writing, this is the largest EC2 server offered.


The database traffic is mostly simple queries. Most queries read or write a single table, and the most frequent and expensive queries access only a single row in one table. JOIN queries are rare and do not access many tables. Although the application overall is extremely read-heavy, the read workload is largely absorbed by the caching layer, so the database sees a more balanced ratio of reads to writes, and the peaks are not as dramatic. In contrast to the 40x variations in the application’s overall traffic, peak throughput at the database is only about three times as high as off-peak.

The ideeli application is complex, yet ideeli’s team has built it from simple components. The team is distributed across the United States and even worldwide—wherever ideeli can find talented people. The ideeli software development workflow includes practices such as continuous integration and

unit testing to keep the code clean and elegant. The fact that the application shields the database from some of the complexity is particularly important for performance, because the database performs better with a simpler workload. This enables the application to serve very high traffic volumes while providing an outstanding experience to users.

Database Performance Problems

The performance problems on ideeli's MySQL database servers were characterized as "spikes" of slowness. This was easily visible in the HAProxy logs, where requests became queued for intervals of 5 to 15 seconds several times a day at peak load. These events were also visible in [New Relic](#), the tool that ideeli uses for application performance monitoring. New Relic graphs revealed that the query response time during the queuing was very slow. The following image illustrates a typical "spike" in the New Relic graphs.



During these periodic performance problems, the database seemed to simply stop responding. Trivial queries that should have been extremely fast were responding as slowly as eight seconds or more. System-level statistics on resources such as CPU and disks showed that virtually nothing was happening during these times, eliminating the operating environment as the source of the problems.

The Percona consultant assigned to help ideeli improve database performance reviewed the system and found that during normal operation, nothing was amiss. The consultant reviewed the server configuration and suggested some changes, as well as providing a top-20 list of queries that were consuming resources. Some of them were candidates for limited amounts of optimization, but overall it was clear that the server's workload was well-optimized and the queries could not be the cause of the performance problems.

However, it is not possible to draw conclusions about what happens during an unusual event by

looking at average-case performance. The consultant used special diagnostic tools to gather carefully scoped performance and operational data during the "spikes" of slowness. The trigger condition was the number of actively running queries. During normal operation, the database server never had more than 100 queries trying to run simultaneously, but during the momentary outages, the output of `SHOW GLOBAL STATUS` showed the `Threads_running` counter to be higher than 100.

After capturing and analyzing several sets of diagnostic data, the consultant determined that the database was suffering from two problems:

- **Query cache scalability limitations.** The query cache performs poorly under very high concurrency. Percona captured the problem in action and recommended reconfiguring the query cache. However, the severe performance spikes at peak load remained.
- **Internal contention inside InnoDB.** InnoDB is the most widely used transactional storage engine for MySQL. Although it is fast and reliable, it is prone to contention under high-concurrency workloads. Percona identified the exact "signature" of the problem, which is solved in XtraDB, Percona's own version of the InnoDB storage engine.


After careful diagnosis, the ultimate fix for the bottleneck was as simple as upgrading to the Percona Server with XtraDB. ideeli evaluated this change over a period of time, and then upgraded.

Performance After the Upgrade


When asked how Percona Server with XtraDB is working, ideeli's CTO Mark Uhrmacher said simply, "we're impressed with it." He noted that queueing at the load balancer is dramatically reduced, as seen by analysis of HAProxy's log files.

The Percona consultant measured performance for before-and-after comparisons. Percona was initially engaged in June 2010, and a comparison of traffic in early June with traffic in late August showed that not only is the queueing problem at the load balancer solved, but the database server is also delivering much higher throughput than before. In June, the

database was faltering under its peak traffic, but after the switch to Percona Server with XtraDB in August, it was serving peaks that reached nearly twice as many queries per second with no quality of service issues. The following chart shows the throughput for the database server, in queries per second, for a typical day in June versus a typical day in August on Percona Server with XtraDB:



The next graph shows the average response time in microseconds for the same time ranges:



It is clear that Percona Server with XtraDB is achieving much higher throughput, while simultaneously keeping response time smooth and predictable, typically about half a millisecond. The average response time over the course of the day was 721 microseconds per query before the upgrade. This decreased to 538 microseconds per query after the upgrade. This fast and consistent response time is essential for producing a good user experience for ideeli's online shoppers.

About Percona Server with XtraDB

Percona Server with XtraDB is a 100% free and open-source drop-in replacement for the MySQL database server. It is backwards compatible with the MySQL

server, and an upgrade is as simple as installing Percona Server and restarting the database. Percona Server with XtraDB delivers immediate performance and scalability improvements, enabling it to take advantage of today's multi-core, large-memory servers with fast disks and solid-state storage. It also includes XtraBackup, a companion tool for creating online, non-blocking backups of XtraDB and InnoDB data files.

Percona Server with XtraDB is designed for users who need their database to provide increased performance and scalability, more configurability, and better instrumentation. Percona chooses high-value, safe enhancements that solve important needs for conservative users by deviating from the stock MySQL database server as little as possible. The result is a database server that:

- Performs better
- Scales to larger hardware
- Eases operational and administrative tasks
- Enables inspection of server internals
- Supports scientific performance optimization
- Reduces downtime and server restarts

About Percona, Inc.

Percona's international team of MySQL experts provides 24x7 support, consulting, training, and engineering services for MySQL, Percona Server, and MariaDB. Percona's founders are world-famous for their expertise in MySQL performance and scaling. They are the authors of *High Performance MySQL, 2nd Edition*.

If you are interested in Percona's products or services, we invite you to contact us through our website at <http://www.percona.com/>, or to call us. In the USA, you can reach us during business hours in Pacific (California) Time, toll-free at 1-888-316-9775. Outside the USA, please dial +1-208-473-2904. You can reach us during business hours in the UK at +44-208-133-0309.

Percona, XtraDB, and XtraBackup are trademarks of Percona Inc. InnoDB and MySQL are trademarks of Oracle Corp.