

The Fastest Hyperconverged Infrastructure Platform

Cut your infrastructure costs by 70%

Run bare-metal performance MySQL and PostgreSQL in AWS without the costs



High Performance MySQL and PostgreSQL on AWS

MySQL and PostgreSQL are two of the most popular open-source relational databases, and are utilised heavily for mission-critical, highly transactional workloads, such as CRMs and eCommerce. Database optimisation at the infrastructure level can yield huge speed improvements, leading to better application performance and lower infrastructure costs.



More resources = faster performance?

Due to the volume of transactions and IO, databases are usually bound by storage performance. Because of this, throwing as much memory at the database as possible is highly desirable to maximise caching and minimise disk reads/writes. However, this approach hits a plateau as databases get larger, and is mostly beneficial to read performance.



Storage IOPS in AWS

When deploying a database in the AWS cloud, storage is typically provided by EBS (Elastic Block Store). EBS GP2 storage delivers a limited 300 to 16,000 IOPS per volume. Provisioned IOPS (PIOPS) are also available, giving up to 64,000 IOPS per volume (32,000 IOPS per volume on non-Nitro instances), but come at a prohibitive cost – see Total Cost Analysis below. These volumes can be striped together using RAID0 to maximise performance, however there is then a hard limit of 80,000 IOPS aggregate per EC2 compute instance. In addition, only certain very large instance types can achieve that maximum IOPS figure – see callout – requiring you to pay for large amounts of unused compute.

AWS offers high performance local NVMe storage on certain instance types, which can theoretically achieve millions of IOPS, however the native AWS hypervisor can typically only achieve x% of the bare-metal performance, and the storage is not persistent or resilient – not suitable for a database workload.

Implementing master-slave configurations will also improve read performance, but requires significant administrative housekeeping, and does not improve database write times, which constitute roughly 30% of database queries in an eCommerce use case – far higher during the critical checkout process.

Maximising storage IOPS is therefore critical to good database performance.

AWS instances that support 80k IOPS

Nitro (64k IOPS per vol)

c5d.18xlarge, c5n.18xlarge, i3en.24xlarge, m5.24xlarge, m5d.24xlarge, p3dn.24xlarge, z1d.12xlarge, r5.24xlarge, r5d.24xlarge,

Non-Nitro (32k IOPS per vol)

g3.16xlarge, h1.16xlarge, p3.16xlarge, x1.32xlarge, x1e.32xlarge



As of 28/6/19 from <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSOptimized.html>



Sunlight – when performance matters

Sunlight’s Hyperconverged Computing Infrastructure stack is a new way to deploy IOPS intensive workloads in the AWS cloud – offering all the performance of bare metal with all the manageability of virtualisation. Sunlight servers are deployed on-demand in AWS, singly or in clusters, on bare metal servers. Once live, you can deploy Sunlight VMs, storage and networking through the Sunlight Dashboard. Native AWS AMIs can be deployed without changes on the Sunlight stack.

Sunlight VMs can individually access more than 1M IOPS from local NVMe storage, managed as a Virtual SAN by Sunlight SDS, which also provides resilience and persistency (when deployed in a 2-node+ cluster), and this can be backed off to EBS storage for backup.

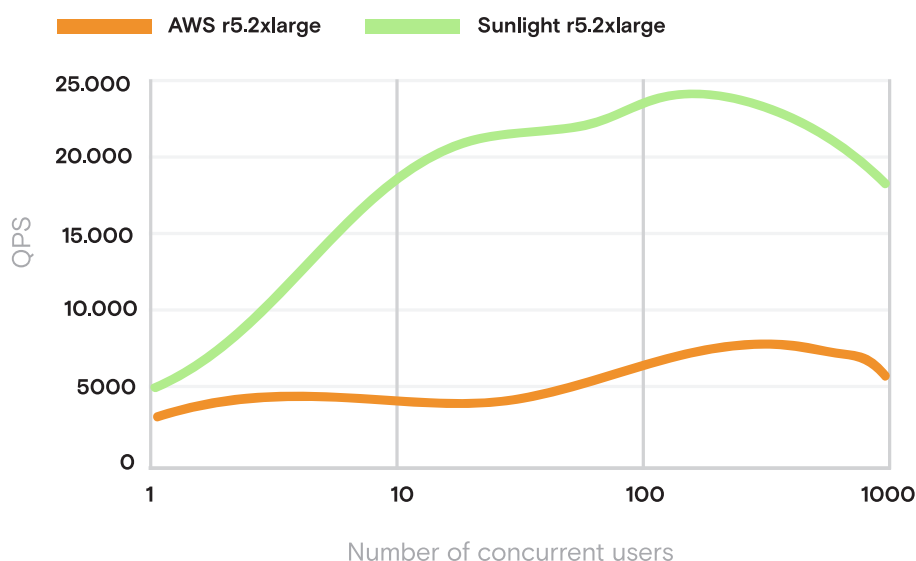
This gives orders of magnitude better IOPS performance than is possible with native AWS instances – saving between 80 and 95% of the total instance plus storage costs when using PIOPS.

Benchmark testing of MySQL on native AWS instances vs Sunlight-in-AWS instances shows a consistent 400% increase in query-per-second performance.

Because Sunlight is so resource efficient, the hardware footprint underlying a database implementation can be drastically reduced. Rather than the database needing to run in a large native AWS VM to maximise IOPS (e.g. r5.24xlarge), it can run on a much smaller Sunlight VM – e.g. a r5.2xlarge equivalent. This leads to far more efficient resource utilisation.

QPS of MySQL database with different number of concurrent users

Database buffer size: 51.2GB (80% of 64GB RAM); Database size: 230GB.



Total Cost of Ownership Analysis for Database Tier

The following table shows the equivalent monthly costs for a 1 year reserved instance paid upfront.

The AWS native case is for a R5.24xlarge with 2x1800GB EBS volumes with 40,000 PIOPS each – for 80,000 IOPS total. The Sunlight case is for a R5d.metal with 4xNVMe drives locally attached – delivering at least 1M IOPS total.

	EC2	EBS	PIOPS	Sunlight	AWS Support	Total
AWS native	\$2596	\$450	\$5200	0	\$757	\$9003
Sunlight	\$2972	0	0	\$601	\$225	\$3797

In this case, Sunlight is 42% of the cost of the native AWS implementation for over 10x the IOPS performance. This does not take into account the fact that Sunlight is only using 1/12 of the capacity of the bare metal server, capacity which can be used to provision additional Sunlight VMs for the rest of the application components. Taking this scaling factor into account, the total cost for the Sunlight implementation is far lower.

Next Steps

Contact Sunlight to arrange a trial of Sunlight on AWS and see how your database implementation can be accelerated by removing IOPS roadblocks to database performance.

Take the Sunlight Challenge



If you'd like to see how **Sunlight** can solve your performance problems with a better ROI than any other hyperconverged infrastructure platform, get in touch for a free trial. www.sunlight.io/freetrial

