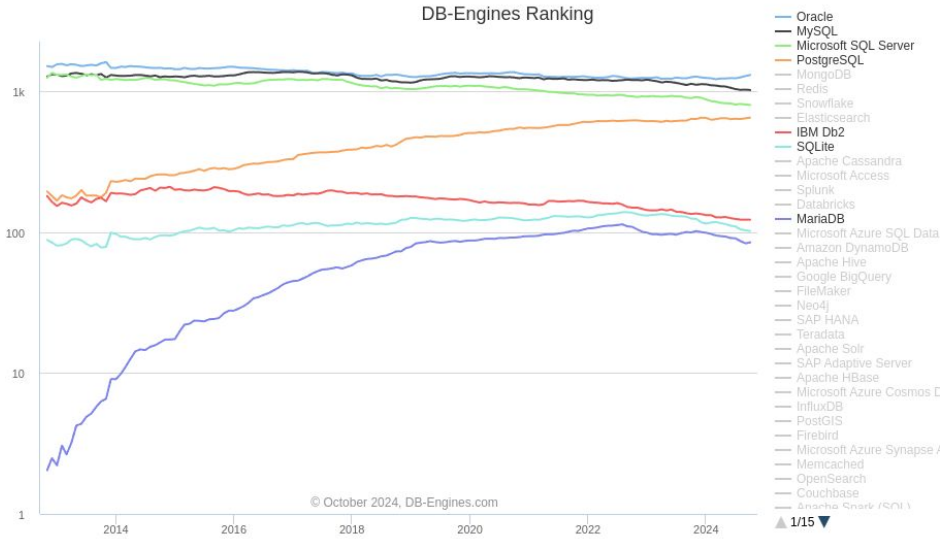

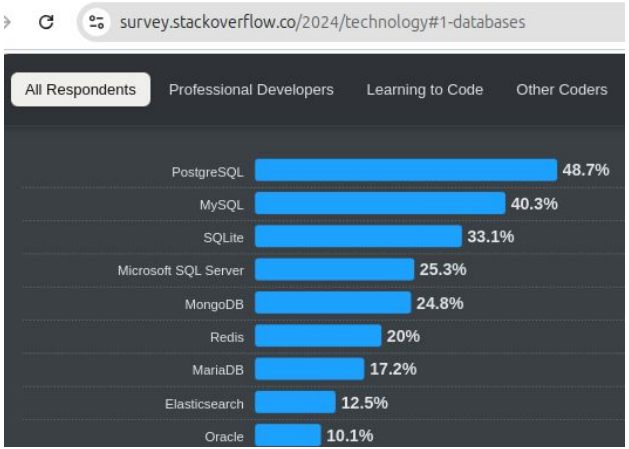

The Register

Open source PostgreSQL named DBMS of the year by DB-Engines

Already more than 37 years old, the relational system continues to gain popularity

Lindsay Clark

Wed 3 Jan 2024 // 17:00 UTC

Peter Zaitsev • 1st
Entrepreneur | Driving Success with MySQL, MariaDB, MongoDB & Postgr…
3d •

"Postgres is eating the database world" - Do you agree ?
https://lnkd.in/ecUyy5Rj #postgres

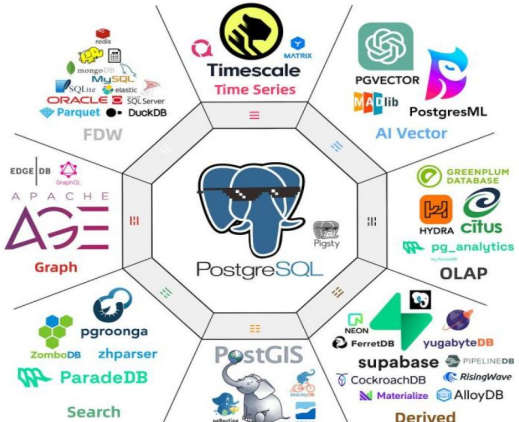

Postgres, postgres, postgres

It’s already awesome!

But wait... what if one could make it even better? Like 5x better - in about 3 minutes?

Talking about $$$ here :)

The Elephant in the room - managed database costs...

Spot instances!

Amazon EC2 Spot Instances

Run fault-tolerant workloads for up to 90% off

Get started with Amazon EC2 Spot Instances | Contact Sales

Save up-to 90% on On-Demand prices.

Take advantage of the massive operating scale of AWS and run hyperscale workloads at a significant cost savings.

Launch, scale, and manage Spot Instances with ease through AWS services or integrated third parties.

Build sustainable solutions by leveraging unused EC2 capacity.

ChatGPT

What is the estimate percentage of really business critical databases?

The percentage of **business-critical databases** can vary significantly depending on the size, industry, and operational needs of a company. However, based on industry surveys and studies, it is estimated that **15% to 25%** of an organization's databases are considered **business-critical**.

Amazon EC2 Spot Instances let you take advantage of unused EC2 capacity in the AWS cloud and are available at up to a 90% discount compared to On-Demand prices. You can use Spot Instances for various stateless, fault-tolerant, or flexible applications such as big data, containerized workloads, CI/CD, web servers, high-performance computing (HPC), and test & development workloads.

A selection of non-critical workloads

- Dev / staging environments
- Feature testing “forks”
- Backup / DR restore testing
- Production snapshots for ad-hoc reporting and exploration with mutations
- Testing effects / feasibility of HW upgrades
- Testing critical maintenance operations like VACUUM FULL-s
- ...

Spot is actually not that scary

If to use the Spot Instance advisor tool by AWS

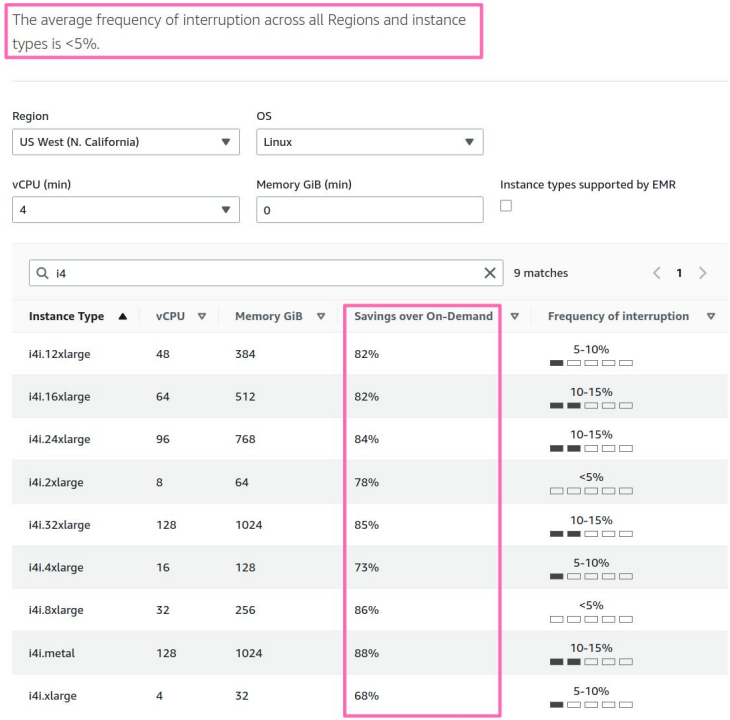

The average frequency of interruption across all Regions and instance types is <5%.

Region: US West (N. California)

OS: Linux

vCPU (min): 4

Memory GiB (min): 0

Instance types supported by EMR

i4

9 matches

Instance Type	vCPU	Memory GiB	Savings over On-Demand	Frequency of interruption
i4i.12xlarge	48	384	82%	5-10%
i4i.16xlarge	64	512	82%	10-15%
i4i.24xlarge	96	768	84%	10-15%
i4i.2xlarge	8	64	78%	<5%
i4i.32xlarge	128	1024	85%	10-15%
i4i.4xlarge	16	128	73%	5-10%
i4i.8xlarge	32	256	86%	<5%
i4i.metal	128	1024	88%	10-15%
i4i.xlarge	4	32	68%	5-10%

Meaning - **on average, one can expect to run a few months uninterrupted!**

Stateful databases + stateless Spot ?

Gets messy of course…those pesky DBs namely have “state”…

Wouldn’t it be nice if someone else deals with the annoying details ?

What about something like:

PG Spot Operator - in “one-liner” mode

psql "$(**pg_spot_operator** --region=eu-north-1 --ram-min=64 --storage-min=500 \

--storage-type=local --tuning-profile=analytics --instance-name=mypg1 \

--admin-user=pgspotops --admin-user-password=topsecret123 **--connstr-output-only**)"

...

INFO Current Spot discount rate in AZ eu-north-1a: -75.5% (spot $126.6 vs on-demand $516.2)

...

psql (16.4 (Ubuntu 16.4-1.pgdg24.04+2))

SSL connection (protocol: TLSv1.3, cipher: TLS_AES_256_GCM_SHA384, compression: off)

Type "help" for help.

pgspotops@postgres=>

~6x savings compared to similar RDS!

** Assumes local AWS CLI setup*

“UI” - CLI / Docker params or a “manifest”

```
pg_spot_operator \
  --instance-name pg1 \
  --region eu-north-1 \
  --cpu-min 2 \
  --storage-min 100 \
  --storage-type network
```

```
docker run -e PGSO_REGION=eu-west-1 \
  -e PGSO_RAM_MIN=128 \
  -e PGSO_STORAGE_TYPE=local \
  -e PGSO_CHECK_PRICE=y \
  -v ~/.aws:/root/.aws:ro
  pgspotops/pg-spot-operator:latest
```

```yaml
api_version: v1
kind: pg_spot_operator_instance
cloud: aws
region: eu-south-2
instance_name: hello-aws
expiration_date: "2024-12-22 00:00+03"
vm:
  cpu_min: 4
  ram_min: 16
  storage_min: 500
  volume_iops: 10000
os:
  extra_packages: [ pgbadger, postgresql-16-cron ]
  ssh_pub_key_paths: [ ~/.ssh/my_key.pub ]
user_tags:
  app: backend
postgresql:
  admin_is_superuser: false
  app_db_name: app
  admin_user: admin
  tuning_profile: oltp  # none | default | oltp | analytics | web
  admin_user_password: !vault |
    $ANSIBLE_VAULT;1.1;AES256
    306433643563343037396265346239376137333865353466613631663231386
```

Sustainable Open Source

- We believe a great addition to the Postgres ecosystem
  - Making Postgres even more accessible and fun to play with

- Looking for a pre-seed round to bring out an Enterprise Edition
  - HA, Azure + Google cloud, etc
  - Hybrid provisioning - Spot only if DT budget allows
  - Extra developer hands needed to hack away the nitty-gritty

github.com/pg-spot-ops/pg-spot-operator

Licence: Functional Source License, Version 1.1, Apache 2.0 Future License

(Meaning: use however you want for any non-compete purpose, totally free after 2 years)

info@pgspotops.com